ION TRAPPING IN HIGH INTENSITY LINACS

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Steven Full

May 2017

ION TRAPPING IN HIGH INTENSITY LINACS

Steven Full, Ph.D.

Cornell University 2017

A charged particle beam will rapidly ionize any residual gas in an accelerator's vacuum chamber, and, if that beam is negatively charged, the resulting positive ions can become trapped within the beam. Ion trapping has often been observed in circular accelerators, but has never before been seen in single-pass linear accelerators. However a new class of high intensity linacs will be the first such linear accelerators that experience ion trapping. In the Cornell photoinjector, we have recently observed this phenomenon for the first time, and we will share our experiences conducting experiments to study ion trapping, as well as theory and simulations modeling the phenomenon.

We start by outlining theories to determine whether or not ion trapping will occur in an accelerator. We describe in detail some of the effects that ions can have on a beam, including emittance growth, optical errors, beam losses and even beam instabilities. The severity of these effects varies widely depending on the accelerator in question, so we offer up several simulation techniques that can be used to predict their occurrence, as well as ion signatures that can be observed experimentally to confirm the presence of ions. We share results from experiments that tested three major ion clearing methods: ion clearing electrodes, bunch gaps, and beam shaking. Results obtained from these experiments are supported by various theories and simulation codes. Finally, because taking beam property measurements in the regime where ion trapping occurs can be difficult, we offer up a new design for a rotating wire scanner capable of obtaining beam profiles at high beam intensity.

Acknowledgments

Acknowledgments

Table of Contents

List of Figures

List of Tables

CHAPTER 1

Introduction

In an accelerator's vacuum chamber, any residual gas is rapidly ionized by collisions with the charged particle beam. At high beam currents of electrons or negatively charged particles, the resulting positive ions become trapped inside of the negatively charged beam and can cause a variety of effects, including emittance growth, optical errors, beam halo, beam losses, excessive radiation, or even beam instabilities [1]. Even with improvements in vacuum technology, ions can fully neutralize a beam within seconds for vacuum pressures as low as 1 nTorr. Therefore one must directly remove the trapped ions to avoid or mitigate these effects.

Until very recently, ion trapping has never been observed in a single pass linac, although it has been observed many times in circular accelerators such as the Advanced Light Source (ALS) [2], the CERN antiproton accumulator [3], and the Metrology Light Source (MLS) [4]. However a new class of high brightness linacs has emerged that is capable of reaching parameter regimes under which ion trapping is of great concern. These machines are capable of achieving beam currents in excess of 10 mA, the approximate threshold at which ion effects often become observable [4; 5]. The high current Cornell photoinjector is one of the few linacs in the world where ion trapping occurs. Therefore, we have developed theories, created simulations and conducted experiments to study ion effects in the photoinjector, so that ion problems in future linacs may be preempted and appropriate mitigation options may be implemented.

Our first goal in Chapter 1 is to offer methods to determine whether or not ion trapping will occur in an accelerator. First, we show ways to calculate the ionization rates for collision and field ionization processes, and offer a way to estimate the time necessary for the ions to fully neutralize a beam. We also establish two major trapping conditions to determine what criteria dictate whether or not ion trapping occurs. The first method models the beam as a continuous charge distribution that traps all ions within it, and shows that ion trapping occurs only when this approximation is valid. The second trapping condition takes into account the bunch structure of the beam. This trapping condition considers the kicks the ions receive as bunches pass by, in order to obtain a stability criterion for the trapped ions' motion. We also share evidence of the first ever observations of ion trapping in a single pass linac.

In Chapter 2 we outline ways in which ions cause negative effects on the beam. Although direct impacts of trapped ions on beam properties have yet to be observed in linacs, much work has been done to study and measure ion trapping in circular accelerators. These results can be used as a guide to predict what phenomena will occur and be observable in a linac. We provide general descriptions of each phenomenon, and, when possible, measurements of beam-ion effects taken in the Cornell photoinjector and at other accelerators such as the MLS in Berlin, Germany. We also present results from simulations designed to model the effects of non-linear beam focusing due to ions in the photoinjector. In general we show that the non-linear focusing has significant impacts on beam parameters that can be difficult to correct using linear optics systems, ultimately rendering ion clearing methods a necessity.

Chapter 3 contains two major sections. In the first, we share results from experiments testing three different methods of removing trapped ions. The first ion clearing strategy involves directly removing the trapped ions using an electric field generated by an ion

clearing electrode – an instrument similar to a large parallel plate capacitor. We present data and analytical formulas that allow one to predict the necessary voltage to achieve full ion clearing for any beam parameters. The second ion clearing method involves introducing periodic bunch gaps in order to give time for ions to drift out of the center of the beam pipe between bunches. We present data that suggests the amount of ion clearing does not depend strictly on the frequency or length of the bunch gaps, but depends primarily on the total time the beam was absent. Finally, the third method of ion clearing involves shaking the beam at a specific frequency to drive a beam-ion coupling resonance which knocks the ions out of the center of the beam. We derive an analytical formula which fits very well with our data for several gas species at beam currents between 10 - 20 mA.

In the second section of chapter 3, we show results from several simulations designed to further explain our experimental data. The first uses a 3D partial differential equation solving software package known as FENICs to model the geometry of complex clearing electrode designs. The goal is to find the minimum electric field required to overwhelm the electric potential of the beam, which is modeled as a long charge distribution. The second simulation tracks the motion of ions longitudinally through a DC beam and along the length of the accelerator. The simulation shows that the ions tend to accumulate at areas of beam size minima, making them the optimal locations to place clearing electrodes. We also demonstrate the ability of RF cavities to either trap ions or act as a barrier to prevent their longitudinal motion through the cavities. Finally, we developed a 2D Poisson solver using MATLAB to model the phenomenon of beam shaking. It was found that the simulation results qualitatively agree with our experimental findings, confirming that beam shaking can result in significant reductions in trapped ion density.

Finally, chapter 4 discusses some of the challenges of taking beam property

measurements at high beam intensity - the parameter regime at which ion trapping occurs. The biggest issue is that high beam power will melt almost all materials that come into contact with the beam within microseconds or less. This renders most traditional interceptive beam diagnostics obsolete. To measure beam profiles, we therefore designed a new rotating wire scanner which is capable of slicing the beam with a carbon filament at speeds in excess of 20 m/s – fast enough that the wire will not melt during a single scan. We present results from bench tests designed to calibrate the speed of the wire as well as measure the relative uncertainty of the wire speed due to wire vibrations. We also conducted tests in the beam, and showed that the instrument typically achieved measurement uncertainty between 5-10% at a scanning speed of 20 m/s for a large number of measurements.

1.1 Ionization rates

The time it takes to fully neutralize a beam can determine whether or not ions are a concern in an accelerator. Ion trapping occurs when the ion creation time is greater than the natural ion clearing time (primarily determined by the space charge repulsion of ions). Beam neutralization times can vary anywhere from hours at a vacuum chamber pressure of 10^{-12} torr to seconds at a pressure of 10^{-9} torr.

There are several ionization mechanisms in high brightness accelerators. The predominant method in the Cornell energy recovery linac (ERL) photoinjector is collision ionization [1; 6]. The rate of collision ionization can be determined by considering the density of the residual gas, as well as the collision ionization cross section of the gas species. More specifically, we can calculate the amount of time it takes to accumulate as many ions as beam electrons per unit length [1]

Table 1.1: Collision ionization cross sections for common residual gas species at a beam energy of 5 MeV, as well as the full beam neutralization time from collisions for a vacuum pressure of 10^{-9} torr.

Species	C_1	C_2	σ_{col}	τ_{col}
H_2	0.695	8.115	$2.0 \times 10^{-23}\ m^2$	6.7 s
H_20	3.24	32.26	$1.1 \times 10^{-22}\ m^2$	1.2 s
CO	3.70	35.14	$1.0 \times 10^{-22}\ m^2$	1.3 s
N_2	3.74	34.48	$1.2 \times 10^{-22}\ m^2$	1.1 s

$$\tau_{col} = \frac{1}{\sigma_{col}\rho_{gas}\beta_{rel}c} \tag{1.1}$$

where τ_{col} is the neutralization time, σ_{col} is the collision ionization cross section of the gas species, ρ_{gas} is the density of residual gas inside of the beam, and $\beta_{rel}c$ is the speed of the electrons. The density of the residual gas ρ_{gas} can be determined by using the ideal gas law, such that

$$\rho_{gas} = \frac{p_{gas}}{k_b T} \tag{1.2}$$

where p_{gas} is the residual gas pressure, k_b is Boltzmann's constant, and T is the temperature of the residual gas. The density ρ_{gas} describes the density of the gas that the beam will ionize, or the density of ionized gas that will become multiply ionized. In the following analysis higher ionization states are not considered.

The ionization cross sections for different gas species are energy dependent, but they vary little over several orders of magnitude of beam energy. For example, the collision ionization cross sections of H_2 for a beam energy of 10 MeV and 5 GeV differ only by roughly a factor of 1.5 [1]. Ionization cross sections for several gas species have

Table 1.2: The estimated residual gas composition for CBETA during initial and final stable operation modes.

Operation Mode	H_2	CO	H_2O
Initial	50%	30%	20%
Stable	78%	12%	10%

previously been measured experimentally for low energy beams (from 100 kev to 2.7 MeV) [7; 8]. The ionization cross sections are determined by the empirical formula

$$\sigma_{col} = 4\pi\left(\frac{\hbar}{\beta mc}\right)^2\left(C_1 + C_2\left(\ln[\beta^2/(1-\beta^2)] - 1\right)\right) \tag{1.3}$$

where m is the mass of the electron, $\frac{\hbar}{\beta mc} = 1.874 \times 10^{-24}\ m^2$, β is the relativistic beta of the electron beam, and C_1, C_2 are empirically determined constants for each gas species, shown in Table 1.1.

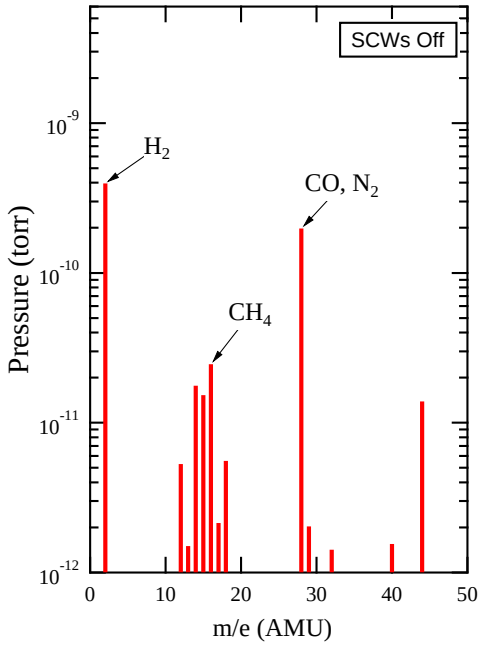

Figure 1.1: A typical gas composition measurement in the Cornell ERL photoinjector. Check CBETA DR for projected gas composition.

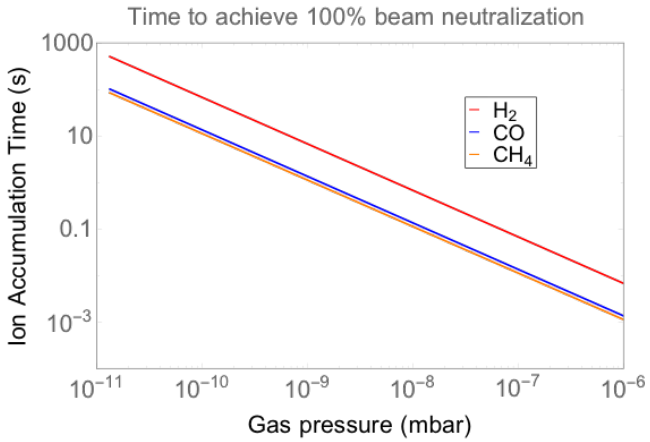

Figure 1.2: The time necessary to fully neutralize an electron beam for various gas species, as calculated using eqn. (1.1).

A typical gas composition for the Cornell ERL photoinjector is shown in Fig. 1.1. In Table 1.2 we also share the expected residual gas composition for the Cornell-Brookhaven Electron Test Accelerator (CBETA), a high current ERL currently being designed and constructed at Cornell. These gas distributions can be used as a guide to determine the dominant ion species trapped in a given accelerator.

Using these as a reference, we have calculated the collision ionization times for several relevant gas species in Table 1.1 and plotted ionization times as a function of residual gas pressure in Fig. 1.2, as calculated using eqn. (1.1). These calculations suggest that the beam will become fully neutralized within seconds, even for extremely low vacuum pressure. Thus, this suggests that ion accumulation is inevitable during regular operation if the ions accumulate around the beam.

The residual gas can also be ionized via tunneling ionization in the beam's electric field [8]. This phenomenon occurs when the bunch generates large enough electric fields to suppress the potential barrier confining an electron, thereby allowing the electrons to

escape via tunneling. The transition rate for ionization in the tunneling regime is given by [8]

$$W = 8\frac{\alpha^3 c E_{ion}}{\lambda_c^2 e E_{max}} \exp\left[-\frac{4\alpha E_{ion}}{3\lambda_c e E_{max}} \right] \tag{1.4}$$

where α is the fine structure constant, λ_c is the Compton wavelength of an electron, E_{max} is the peak electric field of the bunch, and E_{ion} is the ionization energy. Tunneling ionization requires extremely strong electric fields, and is very sensitive to the field strength due to the exponential in this formula. In general, the effects of tunneling ionization are not significant until the peak electric field of the beam is comparable to or exceeds 1 V per angstrom [8]. Calculations show that the rate of tunneling ionization is essentially zero for CO+ ions (with E_{ion} = 11.21 eV) inside a typical beam found in the photoinjector. These estimates were calculated for a beam current of 10 mA, a 2 mm round rms beam size, and an E_{max} found using eqn. (1.9) at one beam σ.

In storage rings, synchrotron radiation can also ionize the residual gas between the center of the beam and the vacuum chamber walls. Although the rate of synchrotron radiation ionization is not always negligible, the density of the generated ions is rather low, and they will form a halo around the beam which has little impact on beam dynamics [1; 8]. Therefore ions generated in this way can generally be neglected.

After they are created, the ions accumulate until their density approaches the density of the beam. At that point, the positive charges of the ions begin to screen the negative electric fields generated by the beam, and the ions clear themselves out by their own space-charge repulsion. This effect, which we call full beam neutralization, places an upper limit on the maximum ion density that can be achieved.

There are also several scenarios where ions will not accumulate at all. If gaps between

bunch trains are sufficiently large, the ions can simply drift out between the bunch trains. This is the case for most linear accelerators. When the gaps are not long enough for ions to drift out of the beam, as in most ring accelerators, the gaps may still be long enough for ions to be over-focused out of the beam center. The next section contains detailed calculations that predict several criteria used to determine whether or not ion accumulation can occur in any given accelerator.

1.2 Ion trapping conditions

1.2.1 Long bunch gap trapping condition

Ion trapping is not observed in linacs or other accelerators with low repetition rates because the ions will simply drift out of the center of the beam due to their initial thermal velocity. This threshold can be estimated by comparing the initial thermal velocity of an ion to the time between bunches. A typical ion will have a velocity somewhere between 10,000 cm/s and 200,000 cm/s at room temperature. For example, at 300 K a single proton will have a thermal velocity of roughly 150,000 cm/s, and ionized CO would have a velocity of roughly 30,000 cm/s. If a beam pipe has a radius of approximately 2 cm, then it will take an ion on the order of 100 μs to drift out of the beam. Therefore ion trapping will not be observed in a typical linac with a repetition rate on the order of Hz. If the drift time is not significantly faster than the bunch spacing, then other methods must be used to determine whether or not ion trapping will occur.

1.2.2 Ion focusing force and oscillation frequencies

A simple estimate that involves a comparison of the ion's transverse oscillation frequency to the bunch repetition rate can be used to determine whether or not trapping occurs. If the bunch repetition rate is sufficiently high, one can invoke a DC beam approximation (meaning that the beam is no longer modeled as bunches and gaps, but instead it is treated as one continuous charge distribution with no gaps). The ions will then oscillate in the beam's transverse potential at a certain frequency.

If that frequency is significantly smaller than the bunch repetition rate, then the ion will see many bunches pass by during a single transverse oscillation. This means that the DC approximation is valid, and the ions will continuously experience focusing forces by the beam and remain trapped within the center of the beam. If the calculated ion oscillation frequency is larger or comparable to the beam repetition rate, then the DC approximation is not valid, and other methods must be used to determine whether or not trapping will occur. These methods (described in the next section) must take into account bunch structure.

The transverse ion oscillation frequency can be derived by first imagining that the DC beam has a round transverse Gaussian distribution. Linearizing the electric force exerted on the ions by the Gaussian transverse charge distribution eventually yields the ion oscillation frequency [9]. The Coulomb force generated by an infinitely long, rotationally symmetric Gaussian beam can be derived using Gauss's law

$$\oint_S \vec{E} \cdot d\vec{S} = \frac{1}{\epsilon_0} \int_V \rho dV \tag{1.5}$$

where ρ is the beam's charge density, and S and V denote surface and volume integrals, respectively. The left hand side is given by

$$\oint_S \vec{E} \cdot d\vec{S} = E 2\pi r L \tag{1.6}$$

where r is the distance from the center of the beam to the location of the electric field E, and L is some arbitrary length dimension. The right hand side is given by

$$\frac{1}{\epsilon_0}\int_V \rho dV = \frac{2\pi L}{\epsilon_0}\int_0^r \frac{\lambda e}{2\pi\sigma_r'^2}\exp\left(-\frac{r'^2}{2\sigma_r'^2}\right)r'dr' \tag{1.7}$$

where r is the distance from the center of the beam, λ is the number of electrons per unit length, and σ_r' is the rms width of the electron beam. Thanks to the cylindrical symmetry of this problem, the integral of the Gaussian, which would normally be difficult, is now trivial. Its solution is given by

$$\int_0^r \exp\left(-\frac{r'^2}{\sigma_r'^2}\right)r'dr' = \sigma_r^2\left[1-\exp\left(-\frac{r^2}{2\sigma_r^2}\right)\right]. \tag{1.8}$$

Thus, by putting together eqns. (1.6) and (1.7), solving for the electric field E and multiplying it by the ion's charge e, we obtain the round beam force

$$F(r) = \frac{\lambda e^2}{2\pi\epsilon_0 r}\left[1-\exp\left(-\frac{r^2}{2\sigma_r^2}\right)\right]. \tag{1.9}$$

According to simulations of the photoinjector in General Particle Tracer (GPT) [10; 11], the beam in the photoinjector is very nearly round for the experimental parameters used in this thesis [9], making this an appropriate approximation for our case. By linearizing this force, we are able to treat the ion's motion inside the beam as a simple harmonic oscillator. The linearized force in this case is given by

$$F(r) \approx \frac{\lambda e^2}{4\pi\epsilon_0\sigma_r^2}r \tag{1.10}$$

and the corresponding equation of motion is now

$$\frac{d^2r}{dt^2} + \omega_i^2 r = 0 \tag{1.11}$$

where ω_i is the oscillation frequency of the ions. Using the linearized form of the force (1.10), it follows that this oscillation frequency is given by [12]

$$\omega_i = \sqrt{\frac{2r_p c}{e}\frac{I}{A\sigma_r^2}} \tag{1.12}$$

where I is the beam current, A is the atomic mass of the ion species and r_p is the classical proton radius. This formula has been experimentally verified using several gas species (N_2, Ar and Kr) at beam currents from 10 – 20 mA. More information on this derivation and the experiments can be found in section 3.1.3.

In the DC approximation ions are always trapped. If the ion oscillation time is slower than the bunch gap duration, then the DC approximation holds and the ions remain trapped. But if the oscillation time is faster than the bunch gap duration, the DC approximation no longer holds and the bunch structure must be taken into account to determine whether or not trapping occurs.

1.2.3 Bunch structure and ion trapping

The following analysis avoids the DC approximation and determines whether ions are trapped within the beam, or whether over-focusing removes them from the beam center. We first examine the case of regular bunch patterns. We will begin by using the linearized transverse force for an asymmetric Gaussian transverse charge distribution.

$$F(\vec{y}) = \frac{e^2}{4\pi\epsilon_0}\frac{\lambda}{\sigma_x+\sigma_y}\begin{pmatrix} \frac{x}{\sigma_x} \\ \frac{y}{\sigma_y} \end{pmatrix} \tag{1.13}$$

We assume that the electron beam has n electrons per bunch, and that the bunches are spaced a distance ΔL apart. When the ions partially neutralize the electron beam by some fraction f, the average charge density is given by $en/\Delta L(f-1)$. If we examine the electron beam's force in the vertical dimension only, then the ions experience a focusing force [1]

$$F_y(y) = -\frac{e^2}{4\pi\epsilon_0}\frac{n}{\sigma_s}\Big(1-f\frac{\sigma_s}{\Delta L}\Big)\frac{1}{\sigma_x+\sigma_y}\frac{y}{\sigma_y} \tag{1.14}$$

This electric force kicks the ions such that

$$\Delta\dot{y} \approx -\frac{1}{A_{ion}}\Big(1-f\frac{\sigma_s}{\Delta L}\Big)\frac{nr_pc}{\sigma_x+\sigma_y}\frac{y}{\sigma_y} \equiv \alpha y \tag{1.15}$$

where A_{ion} is the mass of the ion species and r_p is the classical proton radius. Inside of the electron beam, an ion both drifts according to its velocity such that $y = y_0+\dot{y}\Delta t = y_0+\dot{y}\sigma_s/c$, and is focused by the beam so that it's velocity receives a kick $\dot{y} = -\alpha+\dot{y}_0$. This allows us to construct a transport matrix similar to that of a quadrupole kick under a thin lens approximation

$$\begin{pmatrix} y \\ \dot{y} \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ -\alpha & 1 \end{pmatrix}\begin{pmatrix} 1 & \frac{\sigma_s}{c} \\ 0 & 1 \end{pmatrix}\begin{pmatrix} y_0 \\ \dot{y}_0 \end{pmatrix} \tag{1.16}$$

During the gap, the ions experience a force due to their natural space charge repulsion, given by

$$F_y(y) = \frac{e^2}{4\pi\epsilon_0}\frac{nf}{\Delta L}\frac{1}{\sigma_x+\sigma_y}\frac{y}{\sigma_y} \tag{1.17}$$

This corresponds to a velocity kick

$$\Delta\dot{y} \approx \frac{f}{A_{ion}} \frac{\Delta L - \sigma_s}{\Delta L} \frac{n r_p c}{\sigma_x + \sigma_y} \frac{y}{\sigma_y} \equiv \beta y \tag{1.18}$$

During the gap the ions both drift a distance $y = y_0 + \dot{y}\Delta t = y_0 + \dot{y}\sigma_s/c$ and receive a kick due to space charge repulsion $\dot{y} = \beta + \dot{y}_0$. This results in a defocusing transport matrix given by

$$\begin{pmatrix} y \\ \dot{y} \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ \beta & 1 \end{pmatrix} \begin{pmatrix} 1 & \frac{\Delta L}{c} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} y_0 \\ \dot{y}_0 \end{pmatrix} \tag{1.19}$$

Thus, an ion follows the trajectory given by multiplying eqns. (1.16) and (1.19)

$$\begin{pmatrix} y \\ \dot{y} \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ \beta & 1 \end{pmatrix} \begin{pmatrix} 1 & \frac{\Delta L}{c} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -\alpha & 1 \end{pmatrix} \begin{pmatrix} 1 & \frac{\sigma_s}{c} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} y_0 \\ \dot{y}_0 \end{pmatrix} \tag{1.20}$$

This motion is stable when the trace of the transport matrix has an absolute value between -2 and 2, such that

$$4 > \alpha \frac{\Delta L_g}{c} - \beta \frac{\sigma_s}{c} = \frac{n_e r_p \Delta L_g}{A_{ion}(\sigma_x + \sigma_y)\sigma_y} \left(1 - 2f \frac{\sigma_s}{\Delta L}\right) \tag{1.21}$$

where $\Delta L_g = \Delta L - \sigma_s$ is the bunch gap length. Thus the linear ion trapping condition is given by.

$$A_{ion} \geq \frac{n_e r_p}{4(\sigma_x + \sigma_y)\sigma_y} \Delta L_g \tag{1.22}$$

when $\sigma_s \ll \Delta L$ or $f \ll 1$. This formula also applies to the horizontal direction, by swapping σ_x for σ_y. Because the trapped ion mass is inversely proportional to the beam size, the smaller of the two beam sizes will lead to the most stringent trapping condition. It is worth pointing out that this formula does not take into account multiply ionized gas molecules - it is only for the single ionization case.

Equation (1.22) tells the mass of the ions that will remain trapped within the beam. For example, the CBETA project expects an ion trapping condition $A_{ion} = 0.01$ at 10 mA (assuming every RF bucket is filled), which means that all ion species with a mass number higher than 0.01 will remain trapped within the beam (i.e. all of them). So according to this theory, ion clearing mechanisms are an absolute necessity for CBETA.

One can modify our previous result to consider the effect of long intermittent gaps of a certain duration T_g and at a certain frequency T_p, in the same vein as the work of Sakanaka [Ref: Sakanaka]. During the gap, there are N_{gap} empty RF buckets, followed by $N_{bunches}$, such that

$$\begin{pmatrix} y \\ \dot{y} \end{pmatrix} = \left\{ \begin{pmatrix} 1 & 0 \\ \beta & 1 \end{pmatrix} \begin{pmatrix} 1 & \frac{\Delta L}{c} \\ 0 & 1 \end{pmatrix} \right\}^{N_{gap}} \left\{ \begin{pmatrix} 1 & 0 \\ \beta & 1 \end{pmatrix} \begin{pmatrix} 1 & \frac{\Delta L}{c} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -\alpha & 1 \end{pmatrix} \begin{pmatrix} 1 & \frac{\sigma_s}{c} \\ 0 & 1 \end{pmatrix} \right\}^{N_{bunch}} \begin{pmatrix} y_0 \\ \dot{y}_0 \end{pmatrix} \tag{1.23}$$

In a similar fashion as before, this matrix is stable when the absolute value of its trace has a value less than 2. Because of the nature of this matrix, it is better to perform the trace calculation numerically for a variety of different gas species.

This method must be used to analyze ion trapping conditions for any accelerators with long bunch gap patterns, as opposed to regular filling patterns. For example, it would have to be used to determine ion clearing gaps for CBETA [13], an analysis which will be conducted in the future.

1.3 Evidence of ion trapping in the Cornell photoinjector

We have recently found definitive evidence of ion trapping in the Cornell photoinjector - the first time ion trapping has ever been observed in a single pass linac [5]. There are two

primary pieces of evidence that support our claim. First, during 70 mA, 5 MeV operation, ion clearing electrodes reduced background radiation by over 50%. Beam-ion interactions generate bremsstrahlung radiation, as we have demonstrated in Fig. 1.3 by injecting Ar gas into the beam line at 10 – 20 mA. Using clearing electrodes reduced this excess radiation to background levels. Therefore, we believe that ions must have been trapped within the beam at 70 mA, because ion clearing electrodes also helped reduce radiation background radiation (even at normal vacuum chamber pressures),

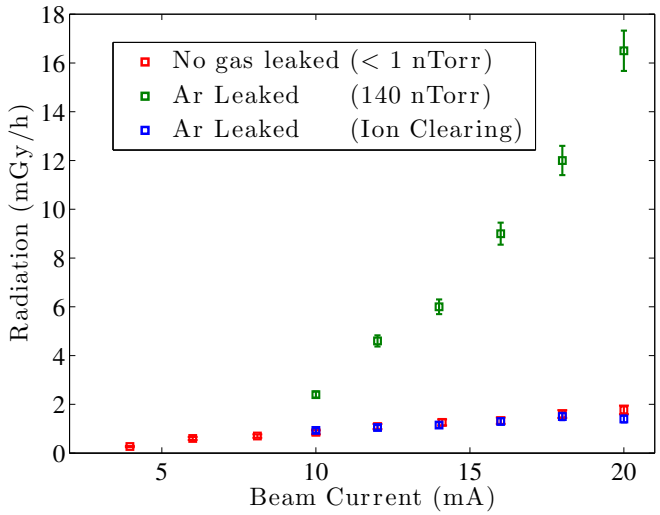

Figure 1.3: After injecting Ar gas into the beam pipe, background radiation levels rose dramatically due to bremsstrahlung generated by beam-ion interactions. Removing the trapped ions using clearing electrodes or other clearing methods reduced this excess radiation to normal background levels. Provided the beam current is high enough (anecdotally > 50 mA), this excess radiation is visible even without injecting gas, and is also significantly reduced by using ion clearing electrodes. This suggests that ion accumulation does indeed occur at high beam current.

Second, during reliability test runs at 20 mA and 350 keV, we observed beam trips that limited stable machine operation to approximately 10 – 15 minutes, as shown in Fig. 1.4. The beam trips were the direct result of the gun's high voltage power supply tripping off. Employing ion clearing techniques, primarily clearing electrodes and/or bunch gaps (described in chapter 3), allowed stable beam operation for at least 24 hours, leading us to

conclude that ions were the cause of the trips. Note that no testing was done for more than 24 hours, but it is expected that stable operation could safely be continued beyond this point if desired.

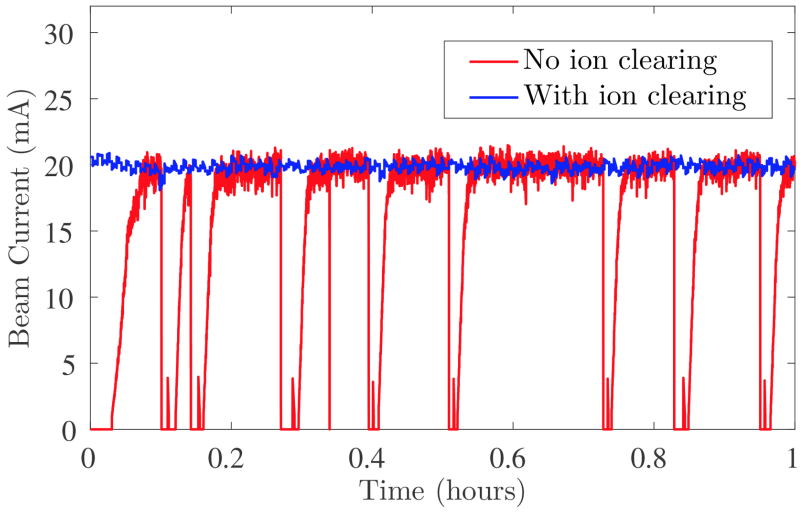

Figure 1.4: A plot of beam current as a function of time in the Cornell photoinjector. During reliability test runs at 20 mA, intermittent machine trips occurred every 10-15 minutes. Using ion clearing electrodes completely remedied this issue, leading to the conclusion that trapped ions were the primary culprit of the machine instability.

Although we have yet to determine the exact mechanism of these trips, we currently have two theories. The first theory is that trapped ions drift backwards and strike the cathode, ejecting particles that then cause arcing. This would ultimately trip off the high voltage power supply. This process of the ions striking and destroying the center of the cathode, known as ion back bombardment, is a problem unique to DC guns, and is slightly different from pure ion trapping as described in this thesis. In a DC gun, the electric field used to accelerate the electron beam to 350 keV will instead push ions back towards the center of the cathode. These ions do not necessarily need to be trapped within the beam in order for this phenomenon to occur, so that it could still occur even at repetition rates too low to observe ion trapping. In the past, ion back bombardment is expected and successfully handled during normal operation by activating the cathode

slightly off center. However, ion back bombardment in the photoinjector has not always been linked to these type of trips, suggesting that our atypical machine setup during the reliability tests played a role in our observations.

The second theory is that dust particles become trapped inside of the beam and drift longitudinally towards the DC gun, where they eventually cause arcing that trips off the power supply. This scenario differs slightly from the ion back bombardment theory in that the ions themselves cause the trips, as opposed to particles ejected from the cathode. In order to be true, this theory would require ions to be generated over a large area, remain trapped in the beam and drift backwards towards the cathode - something that we believe is likely to be occurring.

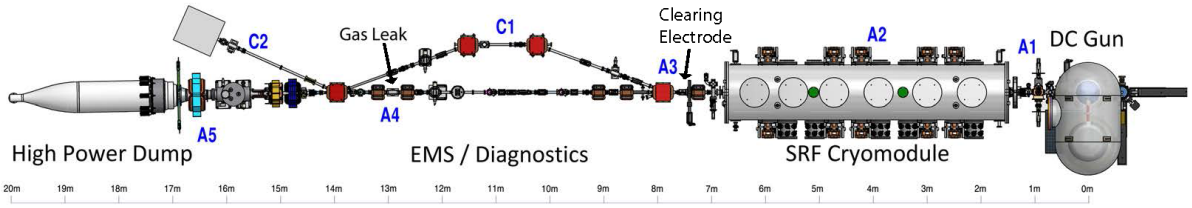

Figure 1.5: A schematic of the Cornell photoinjector. The superconducting RF cryomodule (which creates a potential barrier that prevents ion motion through the cavities) was removed during the reliability test runs, allowing the ions along the entire beam line to drift backwards towards the cathode.

In both cases, the absence of a superconducting RF (SRF) cryomodule (see Fig. 1.5) during these 350 keV test runs plays a major role in our observations, and explains why the machine trips due to ions were observed only during 350 keV operation, and not at normal 5 MeV operation. This is because the SRF cavities normally act like a barrier to the ions, impeding their longitudinal motion, so that the removal of the cavities allows ions from the entire beam line to drift longitudinally back towards the cathode.

The SRF cryomodule acts as a barrier to ions because particles in a rapidly oscillating electric field (rapid relative to the motion of the ion) experience a ponderomotive electric

force, which has a corresponding potential that is proportional to the square of the peak electric field of the RF waveform. This ponderomotive potential is strong enough to prevent almost all ion motion through the cavity, and as a result will reflect ion motion back towards its source. This dynamic is later demonstrated via simulations in section 3.2.2. Therefore, while the SRF cavity is in place during normal high current 5 MeV operation, any ions created in sections A3, A4 and A5 of the beam line (again in Fig. 1.5) will remain contained in those regions and not reach the photocathode. Removing the SRF cavity during 350 keV operation allowed all of these normally isolated ions to enter the DC gun, and significantly contribute to the arcing at the cathode and subsequently tripping the beam.

Ultimately, these two observations demonstrate that ion trapping can, and indeed does, occur in high current linacs. While we are quite convinced that the ions are trapped in the beam, at the present time the effects of trapped ions on beam dynamics and quality in a linac cannot be definitively measured due to the great difficulty of taking measurements at high beam current in a linac. The development of a new fast wire scanner for beam profile measurements, discussed in detail in chapter 4, will allow for such studies in the photoinjector in the future. In the meantime, we can begin to use previous observations from circular accelerators as a guide to understanding and predicting other negative effects of ion accumulation in linacs.

CHAPTER 2

Trapped ion effects on beams

2.1 Ion impacts on beam properties

Once trapped, the ions cause a variety of effects on the beam which can be challenging to identify, measure and quantify theoretically. These range from relatively minor problems such as focusing and incoherent tune shifts, to more severe problems such as beam instabilities. Although ion effects are mostly harmful, they can also be beneficial in rare instances - for example by screening beam space charge in low energy accelerators.

In the past, ion trapping has been extensively studied in circular accelerators via theory [12; 14; 15], simulation [16], and experiments [3; 6; 17; 18]. However, simulation work has been extended to the realm of linacs only recently [1; 19], and experimental observations of ion induced beam effects in linacs remain very rare except for the data presented in this thesis. Fortunately, the vast majority of effects predicted for and observed in circular accelerators are directly applicable to linacs, so future linac designers can rely on them as a general guide. In this section we will describe the possible impacts that ions can have on beam quality and beam operation. Various signatures of ion trapping will be pointed out, so that ion problems can be properly identified and diagnosed when encountered during ordinary machine operation. We will also, when possible, provide ways to predict the effects that will occur using various theories and

simulation codes. These results should allow one to determine if ion trapping in future accelerators is severe enough to warrant the planning of mitigation strategies.

2.1.1 Non-linear focusing

The predominant effect of trapped ions is the highly non-linear beam focusing they cause throughout the accelerator lattice. This non-linear focusing can cause severe changes to beam optics which cannot be easily corrected using linear optics. As a result, non-linear focusing is one of, if not the biggest, concern that one must study and simulate when determining the importance of ion effects in future accelerators.

The simplest way to model this focusing phenomenon involves invoking a DC approximation. The trapped ions can be thought of as a solid charge distribution spread throughout the accelerator. As the ions are created via collision ionization, they drift longitudinally towards beam potential minima, as will be illustrated later in section 3.2.2. This motion causes the density of the ion charge distribution to vary longitudinally along the length of the accelerator. In the presence of a clearing mechanism (including their own natural space charge repulsion), the density of ions eventually reaches a steady state equilibrium. One can thus treat the ions as a type of lens spread throughout the entire accelerator with a strength that varies according to the equilibrium ion density.

The ions possess a very sharply peaked transverse distribution [16; 19; 20], which will be shown later in detail in section 2.2.1. This distribution is ultimately responsible for the highly non-linear focusing force caused by the ions. Some examples of beam changes due to non-linear focusing, as well as the methods we use to model these focusing effects, will be shown in greater detail in section 2.2. It will be shown that, in general, this DC ion

model is sufficient for a first attempt at determining trapped ion effects on a beam in future accelerators.

2.1.2 Measured tune spreads

In circular accelerators, ions can lead to very pronounced tune spreads. In fact, an unexpectedly large tune spread is likely the first symptom of ion problems that one will observe during normal beam operation. This effect is most readily observed by analyzing the tune spectrum of the beam, which can be obtained by attaching a spectrum analyzer to a beam position monitor (BPM) to measure the beam's transverse oscillations.

An example of such a tune spread is shown in Fig. 2.2. This particular measurement was taken at the Metrology Light Source (MLS), a low energy (105-630 MeV) electron storage ring user facility located at the Helmholtz-Zentrum Berlin (HZB), Germany. The MLS is a heavily ion-dominated machine. As soon as the beam current exceeds a few mA at 105 MeV, the beam strongly blows up in all 3 spatial dimensions and begins to oscillate at several different frequencies [4]. In order to overcome these ion effects, they employ clearing electrodes as well as shake the beam using white noise to clear out all of the trapped ions. These methods prove very successful at clearing ions [4], and we have also verified their effectiveness experimentally in sections 3.1.1 and 3.1.3.

Unfortunately, a comprehensive theory that completely predicts and quantifies this tune spread does not yet exist, as the width or location of peaks does not scale in any clear way with several parameters, including beam current and vacuum pressure. This lack of scaling laws has been observed in measurements taken both at the MLS, where beam size was varied [4], and at the Cornell Electron Storage Ring (CESR), where

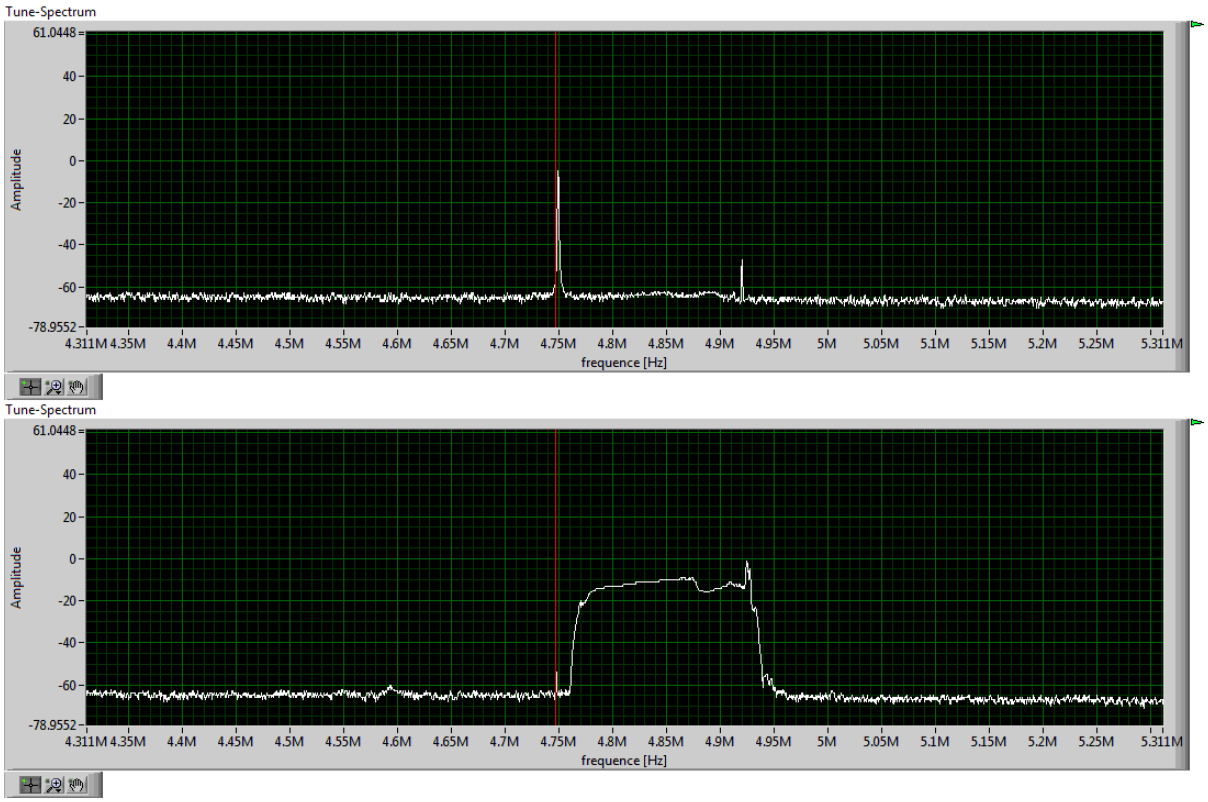

Figure 2.1: The tune at the MLS under normal operating conditions (top), as compared to incoherent tune spreads caused by trapped ions (bottom).

vacuum pressure was changed [18]. Both the magnitude of the peaks and the frequencies which the peaks occur at shift unpredictably as these parameters are changed linearly.

Nevertheless, the tune spreads appear to be rather benign and harmless, and machine stability can be achieved even while they are present. Whether or not they are tolerable during regular operation at a facility would largely be up to the discretion of the operators. Still, the tune spreads are very important as a clear and easy method to confirm the existence of ion trapping in a beam.

2.1.3 Sidebands around the tune

Another common signature of ion trapping in rings are sidebands around the tune. Each ion species possesses a characteristic ion oscillation frequency, which depends on the mass of the ion (see section 3.1.3). Over many turns in a circular accelerator, the ions tend to couple with the beam, causing it to oscillate at a similar frequency. This ultimately results in several sidebands around the tune, each corresponding to a different ion species.

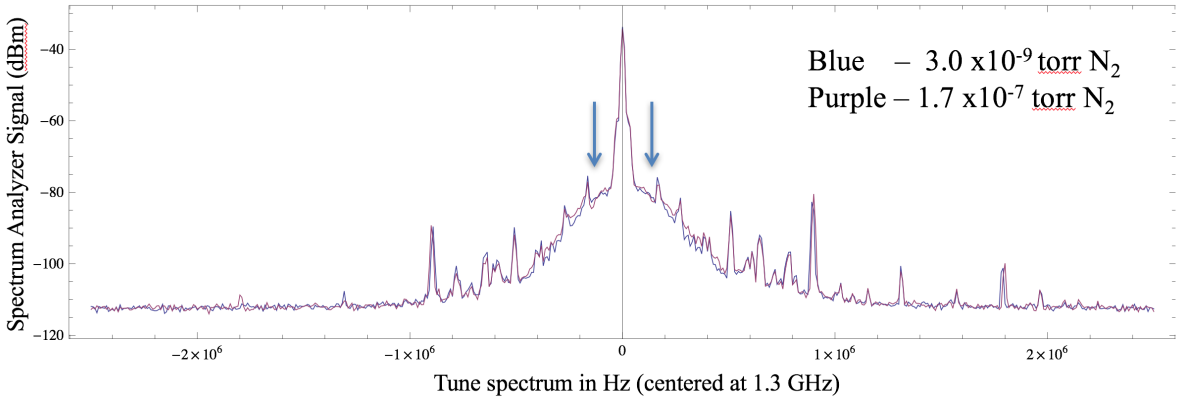

Figure 2.2: A beam spectrum taken while N_2 gas was leaked into the vacuum chamber of the Cornell ERL photoinjector at two different pressures. Equation (3.15) predicts that, for a 10 mA, 2 mm beam, beam-ion coupling should produce peaks in the 10-100 kHz range. For N_2, the location of the expected peaks is indicated with blue arrows. However, these peaks are unobservable due to large amounts of background noise induced by the photoinjector's RF systems at the time of the experiment. Increasing the vacuum pressure from 3.0×10^{-9} torr to 1.7×10^{-7} torr had very little effect on these observations.

These sidebands have been observed in many circular accelerators, including SPEAR-8 [21], the MLS [4], and CESR [18]. In theory, they should be observable in linear accelerators as well. The two measurements differ in one minor way, however: When measuring transverse beam oscillations in a linac using a BPM, one will notice a sharp peak located at the repetition rate of the machine, whereas in a circular accelerator that central peak would be the tune. But in either case, the sidebands due to beam-ion coupling should be located around the central peak.

We made several attempts to observe these tune sidebands by injecting N_2 gas into the Cornell photoinjector, however we were unable to find them because of excess background noise (which was most likely caused by our RF system, and not ions). Regardless of this, it is questionable as to whether or not beam-ion coupling would have been observed even without the noise. The Cornell photoinjector is rather short – about 14 m in total length – and for our experimental setup (shown later in section 3.1) the ion interaction region was only approximately 2 m due to the location of our gas injection system. This means that the beam-ion coupling would need to produce an observable effect after only roughly 7 ns of interaction time (i.e. 2 m / βc), which is likely not enough time for any reasonably large beam oscillations to be excited. Thus, ultimately, our lack of measurements prevents us from drawing any definite conclusions about sidebands in linacs, but this should not prevent one from taking similar measurements in the future.

Although sidebands are a clear signature of ions, they tend to be a rather harmless effect, in the same vein as tune spreads mentioned above. This is of course provided that the oscillations do not drive a resonance or induce an instability, such as the Fast Ion Instability. In that case, an entirely different set of analysis must be performed.

2.1.4 Fast Ion Instability

One of the most insidious effects caused by ions is the fast ion instability. This instability arises when the beam and ion column form a sort of coupled oscillator. The ion column and beam enter a positive feedback loop, which can eventually result in total beam loss. The exact mechanism is as follows: A bunch passes through the beam pipe and ionizes the residual gas. Because the resulting ions are generated within the bunch, it provides the ions with their initial transverse coordinates. Any bunch to bunch jitter will cause the

subsequent bunch to pass by at a slight transverse offset from the previous bunch. As new ions are created, the old ions will drift towards the passing bunch. This process continues for several bunches, until the ions begin oscillating back and forth collectively (as opposed to individual ion oscillations). This collective ion oscillation can begin to pull individual bunches back and forth, resulting in a positive feedback loop. The amplitude of these beam-ion oscillations can increase very rapidly over a short period of time. In the best case scenario, this leads to increases in vertical beam size until the beam size eventually reaches a stable equilibrium. In the worst case scenario, it leads to total beam loss.

As its name implies, the fast ion instability can occur very quickly, typically in milliseconds or less. In general, this means that the ions must be cleared faster than the rise time of the instability in order to achieve stable beam operation. Therefore one's clearing methods should be chosen to fit this demand. In storage rings and linacs where bunch gaps are the predominant clearing method, this effect can be substantial even within a single bunch train. This means that either bunch gaps must occur more frequently, or other clearing mechanisms must also be implemented.

The rise time is a good indicator as to whether or not this instability will be a concern in any given accelerator. The rise time of the fast ion instability is given by [1; 15]

$$\tau_{fast} = \frac{1}{n^2} \frac{\gamma \sqrt{3\sigma_y(\sigma_x + \sigma_y)^3}\sqrt{A_{ion}}}{4\rho_{gas}\sigma_{col}r_e c\beta_y \sqrt{n_e^3 r_p \Delta L}} \tag{2.1}$$

where n is the nth bunch in the train, β_y is the vertical beta value, $\gamma = E_e/(m_e c^2)$ is the relativistic gamma factor of the electron bunch, and r_e is the classical electron radius. As an example, in the case of the Cornell ERL, the rise time occurs on the order of

milliseconds [1]. Clearing times typically occur on the order of μs, therefore the Cornell ERL should not suffer from the fast ion instability.

The fast ion instability was first observed at the Advanced Light Source (ALS) []. During their initial experiments, they observed large vertical beam size growth, likely due to beam-ion coupling. After some time, this beam size growth reached equilibrium, suggesting that the instability remained under control. While taking this data, they found that their feed-forward system also operated at a frequency close to that predicted by fast ion instability theory, suggesting that beam blow up caused by the fast ion instability was being suppressed. This is perhaps one of the most clear indicators of this instability to date.

Many labs throughout the world also attempted to verify, with varying degrees of success, the fast ion instability by injecting gas into their beam lines in order to artificially induce the effect. These include CESR [] and the Pohang Light Source (PLS) []. In most cases, their experimental findings did not result in total beam loss. Instead, they observed a rapid increase of vertical beam size along the length of the bunch train, which eventually reached an equilibrium. However these results still indicate that the fast ion instability is still a very real phenomenon.

Ultimately, the fast ion instability remains one of the most important ion effects to consider. The fact that the instability can occur during a single bunch train can render bunch gaps – arguably the most effective clearing method – completely moot. Thus it is important to determine whether or not this effect will be significant for an accelerator. The most effective way to do this is to consider the rise time of the beam.

2.1.5 Halo, particle losses, and Bremsstrahlung radiation

Halo generation and particle losses are perhaps the most problematic effect for high brightness electron beam sources. Due to very high beam power, even a small fraction of beam losses (roughly $\geq 10^{-4}$) or beam halo can result in significant radiation spikes as the lost electrons strike the beam pipe [22]. This can place a limit on maximum achievable current due to machine protection system limits. In addition, beam losses in general can result in melted beam components due to high beam power. The only way to ameliorate these problems is to directly remove the trapped ions from the center of the beam.

Beam-ion collisions also generate bremsstrahlung radiation. Depending on an accelerator's safety protocols, this can result in trips due to machine protection system shutdowns. Thus, depending the accelerator may require significant additional shielding to operate at high current. In the photoinjector, it was been found that utilizing ion clearing electrodes significantly reduces this radiation, confirming that it is due primarily to beam-ion bremsstrahlung and not beam-gas scattering (see 3.1.1). Thus, as long as radiation shielding is sufficient and ion clearing methods are implemented, this should not be a huge concern.

2.1.6 Ion back bombardment

Ion back bombardment is a problem unique to DC photoinjectors. As the ions are created inside of the DC gun vacuum chamber, they are pushed back towards the photocathode by the DC voltage and strike its center, resulting in the gradual destruction of the center of the cathode [22–24]. Because of the large DC electric field in the gun, this can occur regardless of whether or not ion trapping conditions are met.

Figure 2.3: A photocathode used in the Cornell photoinjector. The cathode (oxidized, dark gray circle) is placed off center in order to avoid being destroyed by ion back bombardment (center light gray circle) [25].

Ion back bombardment has been observed many times and must be compensated for in daily high current operation. An example of a photocathode whose center has been destroyed by ion back bombardment in shown in Fig. 2.3. One of the best solutions to this problem is to activate the cathode slightly off center [22]. While the ions will track backwards towards the cathode and destroy its center over time, the edges remain intact. Therefore the cathode must have a large active area or be grown off center in order to allow the operator to circumvent problems with ion back bombardment. Although it sounds as if using the cathode off center would impact beam quality and operation, it has consistently been shown over the course of several years that it does not [22]. As a result, if this method is implemented then ion back bombardment can be safely ignored.

2.2 Modeling ion focusing

The simplest way to model the effects of trapped ions on a beam is to treat the ion density as a very long charge distribution with a fixed transverse shape. One can think of it simply as a type of lens spread throughout the accelerator. The width and maximum strength of this ion lens vary along the length of the accelerator in proportion to the transverse beam size as well as the local ion density.

This ion lens then produces a non-linear focusing force on the beam. In order to calculate this force, one must assume that the ions possess an appropriate transverse distribution and calculate that distribution's electric field. Once one obtains the necessary electric fields, they can be converted to velocity kicks as needed and inserted into a desired simulation program.

2.2.1 Simulated transverse ion density

The shape of the trapped ion density is well known, and has been predicted theoretically [20] and confirmed in simulations with and without space charge [16; 19]. It takes on a sharply peaked form, which can be seen in Fig. 2.4. This shape can be expressed analytically as [20]

$$\rho(x) = \frac{1}{\pi\sqrt{2\pi}\sigma_e}\exp\left(-\frac{x^2}{4\sigma_e^2}\right)K_0\left(\frac{x^2}{4\sigma_e^2}\right) \tag{2.2}$$

where K_0 is the modified Bessel function of the second kind, and σ_e is the rms beam size of the electron beam in which the ions have been created. This takes on an asymptotic form near the center of the ion distribution given by

$$\rho(x) \approx -\frac{1}{\pi\sqrt{2\pi}\sigma_e}\exp\left(-\frac{x^2}{4\sigma_e^2}\right)\left[\ln\left(\frac{x^2}{8\sigma_e^2}\right)+\gamma_c\right] \tag{2.3}$$

where the constant $\gamma_c = \int_0^\infty e^{-}x\ \ln x\, dx = 0.5772156649$.

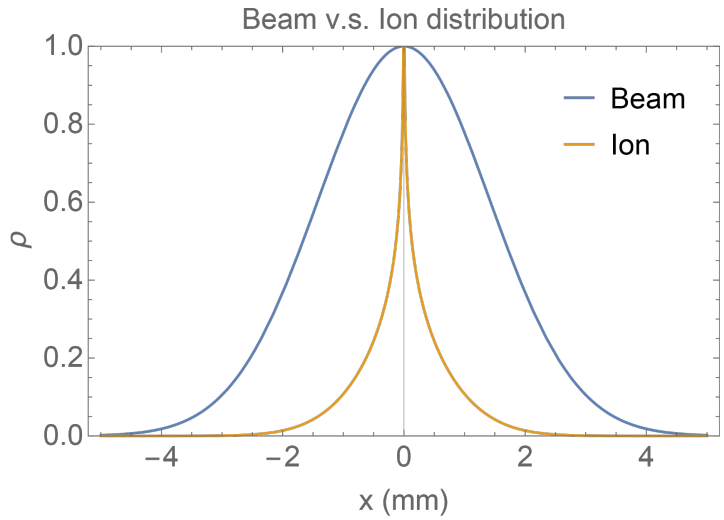

Figure 2.4: The transverse charge density of ions trapped in a round Gaussian beam, as predicted via theory [20].

The ion distribution is slightly more complicated in the intermediate regime when the gaps are long enough to clear only some, but not all, of the ion species. Ions on the edge of the beam get over-focused less than those in the center of the beam, because the focusing force created by the beam decreases in strength outside one beam σ. Ions towards the center of the beam experience a stronger force due to space charge repulsion. Therefore, ions in the center of the beam will be cleared first, resulting in an "ion halo" that forms around the beam. The amount and thickness of this halo will depend on the length of the bunch gap, because that dictates what fraction of ions will be over-focused out of the center of the beam. However, in general we do not consider this intermediate regime when performing our simulations.

Because of its unusual shape, the sharply peaked distribution's electric field is clearly nonlinear, but it is difficult to express it analytically in a closed form. As a result, for the sake of simplicity we chose to model the transverse ion distribution using both constant and Gaussian distributions, and adjusted the width and peak value of these electric fields in order to closely match those calculated numerically from the peaked distribution above. Another scientist, L. Wang, has shown that assuming an ion distribution transverse rms size $\sigma_i = \sigma_e/\sqrt{2}$ causes a Gaussian distribution to create an electric field that very closely matches the electric field created by the actual sharply peaked distribution [20]. As a result, a Gaussian beam with this rms size assumption is our best approximation for the real ion density. We will compare the effects of a constant charge density ion distribution, which generates a purely linear force in the center of the beam, to that of the Gaussian distribution, which generates a non-linear force.

2.2.2 Constant ion density model

For simplicity, we assume that the ion distribution has the same rms size as the beam along the accelerator. We also assume that the number of ions per unit length is proportional to the beam current. This is because more electrons per bunch allow more ions to be created before the ions fully screen the space charge forces of the beam and escape. One could modify this assumption by taking into account the fact that the ion density varies along the accelerator due to the ion's longitudinal motion. However, due to the upper limit on the ion neutralization fraction, this modification would likely have the effect of reducing the total amount of non-linear focusing caused by the ions. Therefore our first assumption describes the worst case scenario of this phenomenon.

Assuming a round transverse ion distribution of constant charge density, the beam's electric field is given by [17]

$$E(r) = \begin{cases} \frac{\lambda e f}{2\pi\epsilon_0}\frac{r}{\sigma_r^2}, & \text{if } r \leq \sigma_r \\ \frac{\lambda e f}{2\pi\epsilon_0}\frac{1}{r}, & \text{if } r \geq \sigma_r \end{cases} \tag{2.4}$$

where λ is the number of electrons per unit length, e is the elementary charge, σ_r is the rms transverse beam size, and f is the ionization fraction (or number of trapped ions per unit length divided by the number of beam electrons per unit length).

A more realistic scenario is to look at an asymmetric beam, such that $\sigma_x \neq \sigma_y$. For an elliptical beam, we define several factors for the sake of clarity

$$\begin{aligned} A &= \frac{\lambda e f}{2\pi\epsilon_0} \\ B &= x^2 + y^2 - \sigma_x^2 - \sigma_y^2 \\ C &= x^2\sigma_y^2 + y^2\sigma_x^2 - \sigma_x^2\sigma_y^2 \\ t &= \sqrt{\frac{B^2}{4} + C} + \frac{B}{2}. \end{aligned}$$

where σ_x and σ_y are the transverse beam sizes. When outside of the ion column, the transverse electric fields are then given by [26]

$$\begin{aligned} E_x &= \frac{2Ax}{\sqrt{\sigma_x^2 + t}(\sqrt{\sigma_x^2 + t} + \sqrt{\sigma_y^2 + t})} \\ E_y &= \frac{2Ay}{\sqrt{\sigma_y^2 + t}(\sqrt{\sigma_x^2 + t} + \sqrt{\sigma_y^2 + t})}. \end{aligned} \tag{2.5}$$

And inside of the ion column, the transverse electric fields are given by

$$E_x = \frac{Ax}{\sigma_x(\sigma_x + \sigma_y)}$$
$$E_y = \frac{Ay}{\sigma_y(\sigma_x + \sigma_y)}. \tag{2.6}$$

Both equations reduce to the round beam case in the limit that $\sigma_x = \sigma_y$.

2.2.3 Gaussian ion density model

The sharp kink in the electric field for a constant charge distribution will result in an unnatural "tearing" in phase space. To avoid this, we also examine the fields generated by a transverse Gaussian distribution. This is a somewhat better choice because of it's smoothness near $r = \sigma_r$.

The electric field generated by an infinitely long, rotationally symmetric Gaussian beam can be derived using Gauss's law (as was done in section 1.2.2), and is given by

$$E(r) = \frac{\lambda e f}{2\pi\epsilon_0 r}\left[1 - \exp\left(-\frac{r^2}{2\sigma_r^2}\right)\right] \tag{2.7}$$

where

The Bassetti-Erskine electric fields for a Gaussian distribution that is uniform in the z direction are given by [27]

$$E_x - iE_y = \frac{-i\lambda e f}{2\epsilon_0 \sqrt{2\pi(\sigma_x^2 - \sigma_y^2)}}\left[w\left(\frac{x+iy}{\sqrt{2(\sigma_x^2-\sigma_y^2)}}\right) - \exp\left(-\frac{x^2}{2\sigma_x^2} - \frac{y^2}{2\sigma_y^2}\right) w\left(\frac{\frac{x\sigma_y}{\sigma_x} + \frac{iy\sigma_x}{\sigma_y}}{\sqrt{2(\sigma_x^2-\sigma_y^2)}}\right)\right] \tag{2.8}$$

where $w(z)$ is the complex error function, also known as the Faddeeva function.

2.3 Tracking through the Cornell ERL photoinjector

Our next step was to track a beam through the ion column by inserting these electric field models into a simulation code. We chose to use GPT [10], the same simulation code shown to be in good agreement with experimental results in the Cornell photoinjector [11]. This simulation models the electron beam from its creation at the photocathode, through several SRF cavities and beam optics devices, until it reaches a beam dump approximately 15 m away. For reference, a schematic of the Cornell ERL photoinjector is included later in chapter 3, Fig. 3.2.

The simulation tracks individual particle trajectories in the beam using a very fast multigrid Poisson solver, which is capable of accurately simulating the space charge forces between electrons in the beam - a necessary feature for a low energy linac (4-12 MeV). It treats each beam line element as a set of electric and magnetic fields in all 3 spatial dimensions. Therefore the ion column could be inserted into GPT solely by using the electric fields outlined in the previous section - no transport matrices are required.

Rather than treating the ion column as an external beam line element with a fixed width, however, the width of the ion column was allowed to vary along the length of the accelerator according to the rms beam size. Because the ion column had a profound effect

on rms beam size as a beam was tracked through it, this ion beam line element could not be pre-computed, unless an iterative solution was implemented (i.e. the beam size obtained from one simulation was used to modify the beam line element, and this process was repeated until a solution was converged on). This iterative approach would be prohibitively time consuming. Instead, the electric field for the ion distribution was inserted directly into the space charge algorithms for GPT, so that the values of the rms beam width instantaneously could be taken into account. This method was used to obtain all of the simulation results in the following section.

2.3.1 Comparison of ion distribution models

Our first goal is to compare the shapes of the electric field generated by a round constant and a round Gaussian ion charge distribution. The electric field for each distribution is shown in Fig. 2.6. In both cases, the electric field is approximately linear in the center of the ion distribution, and falls off roughly as $1/r^2$ very far outside of the distribution. The primary difference between the two models lies when $r = \sigma_{ions}$. In this region, the Gaussian distribution is rather smooth, but the constant charge distribution has a very sharp "kink" which will cause tearing and discontinuities in several examined beam parameters.

One can examine many transverse beam parameters after the beam has traversed an ion column at full beam neutralization that is several meters long. Changes in beta functions and similar rms beam parameters due to the different models are very similar, and it is difficult to distinguish between the effects of different models. However, when one examines beam profiles, as is done in Fig. 2.6, the effects of the "kink" in the constant charge distribution become very obvious. For these simulations we have used standard

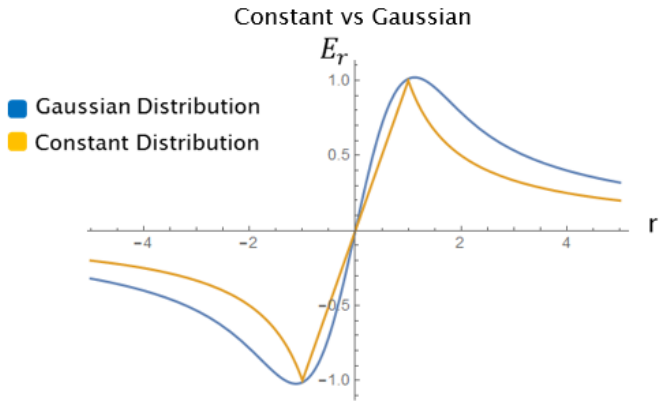

Figure 2.5: Comparison of transverse electric fields for a constant charge distribution model and a Gaussian charge distribution model.

daily optics settings in the Cornell ERL photoinjector, and assumed a 40 mA beam. In both instances, the constant charge distributions create a beam profile which is very sharply peaked in the center of the beam. Note that the shape of the beam profile is not itself triangular - but rather a smooth rounded distribution with a "bump" on top of it. This unnatural discontinuity highlights the flaws of the constant charge distribution model. The Gaussian charge distributions, however, result in beam profiles with a much smoother shape, suggesting that they are a superior model to use. Because of this, we prefer to use this model when simulating ion effects.

In these simulations we have assumed that the rms beam size of the Gaussian distribution is $\sigma_i = \sigma_e/\sqrt{2}$. We do so in order to approximate the shape of the electric field generated by the sharply peaked ion distribution shown in Fig. 2.4. The effect of this assumption can be tested by varying the rms radius of the ion distribution and observing any changes in beam profiles, as shown in Fig. 2.7. Here it is shown that changing σ_i between several values close to the rms beam size ($\sigma_i = \sigma_e$ and $\sigma_i = \sigma_e/sqrt2$) produces very little observable change in beam profiles. This suggests that changing the ion distribution radius may not be so important, provided the changes in radius are not too

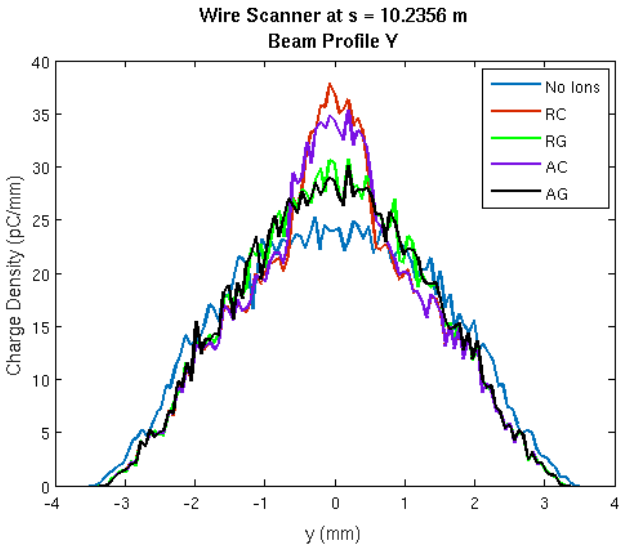

Figure 2.6: Comparison of simulated beam profiles obtained by tracking through different ion distributions in the Cornell ERL photoinjector.

drastic. Problems arise only when the ion distribution radius is chosen to to be very small compared with the rms beam size ($\sigma_i = \sigma_e/10$). In this case the beam profile again obtains a very sharp, discontinuous peak in its center similar to the constant charge distribution model, suggesting that this is an inappropriate model to use. Again, this result supports the assertion that it is necessary to choose values close to σ_e. As long as that is done, then the results obtained from the simulation should not change significantly.

However it is worth pointing out that a very small ion distribution radius may actually be observed in reality, provided that the beam is not fully neutralized. If the beam is fully neutralized then the ion distribution will spread out as far as possible, filling up the beam. But during partial neutralization it is theoretically possible that ions on the edge of the beam will escape first, leaving a very densely packed and sharply peaked core in the center of the beam. Thus if sharply peaked beam profiles such as those shown in Fig 2.7 are ever obtained via measurement, this may explain why.

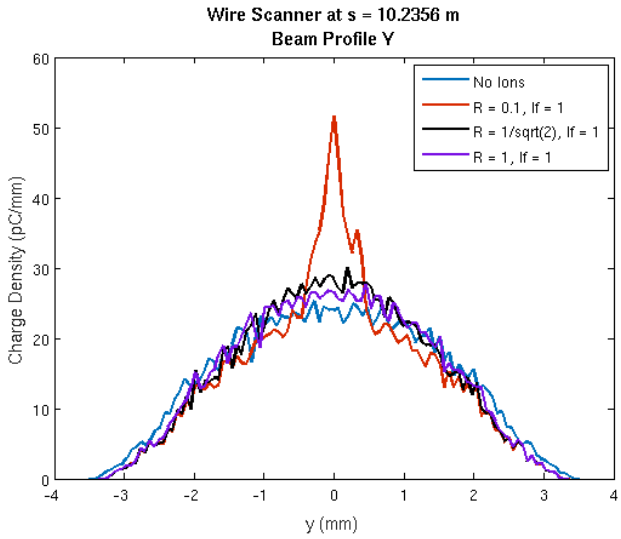

Figure 2.7: Comparison of beam profiles for different ion distribution radii, assuming a round Gaussian charge distribution.

2.3.2 Effects of beam neutralization fraction

One of the most important purposes of a simulation code tracking a beam through an ion column is to determine the amount of ion clearing necessary to minimize the impact of ions on beam dynamics. The way to determine this requirement is to vary the neutralization fraction f of the ion distributions and observe the effect on beam parameters. We have done this for two separate cases, shown in Fig. 2.8 and Fig. 2.9. In the first case, we tested the effects of neutralization fraction on the beam parameters mentioned above for the Cornell photoinjector (i.e. a 40 mA beam, 30.8 pC at 1.3 GHz with standard optics settings). In the second case, we have used the optics settings necessary to achieve the ultimate beam parameters for the CBETA project [13; 28] (also 40 mA, but with different optics settings and 125 pC at 325 MHz).

The simulations show that reducing the neutralization fraction to approximately 20% can negate any visible changes in beam profile in our first simulation code in Fig. 2.8. Even neutralization fractions approaching 100% product very little visible

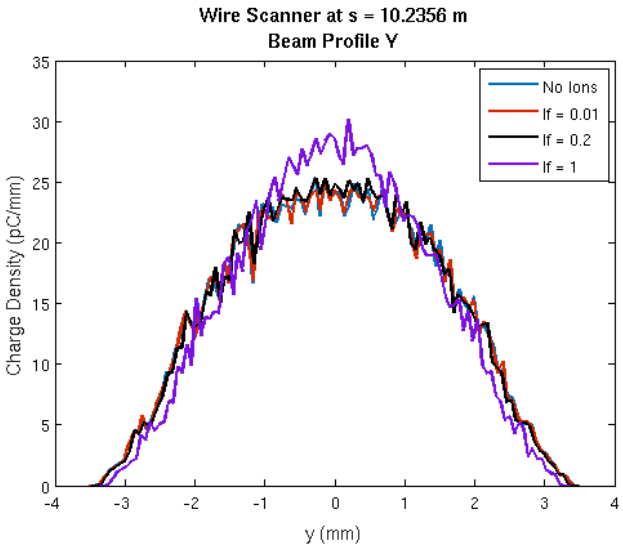

Figure 2.8: Comparison of vertical beam profiles resulting from different ionization fractions in the Cornell ERL photoinjector. The effects of ions on beam profiles can only be distinguished at high ionization fractions.

change in beam profile, suggesting that it may be difficult to conduct experiments in the photoinjector to observe ion effects using only beam profile measurements. The right half of Fig. 2.9 shows that a high neutralization fraction causes the beam to initially focus very hard and cause significant emittance growth. Smaller neutralization fractions reduce this amount of focusing, resulting in even less emittance growth, with a 1% neutralization fraction nearly returning beam parameters to normal ion-free optics settings.

2.3.3 Effect of ions on beam parameters

The non-linear ion focusing has a considering impact on rms beam parameters, including beam size and emittance, provided the strength of the ion column is sufficiently high. Space charge suppression due to excess ions, shown in Figs. 2.10 and 2.11, actually causes the beam to focus more easily. This effect is not necessarily bad, but in fact could be beneficial if taken advantage of. However, because of the non-linearity of the ion

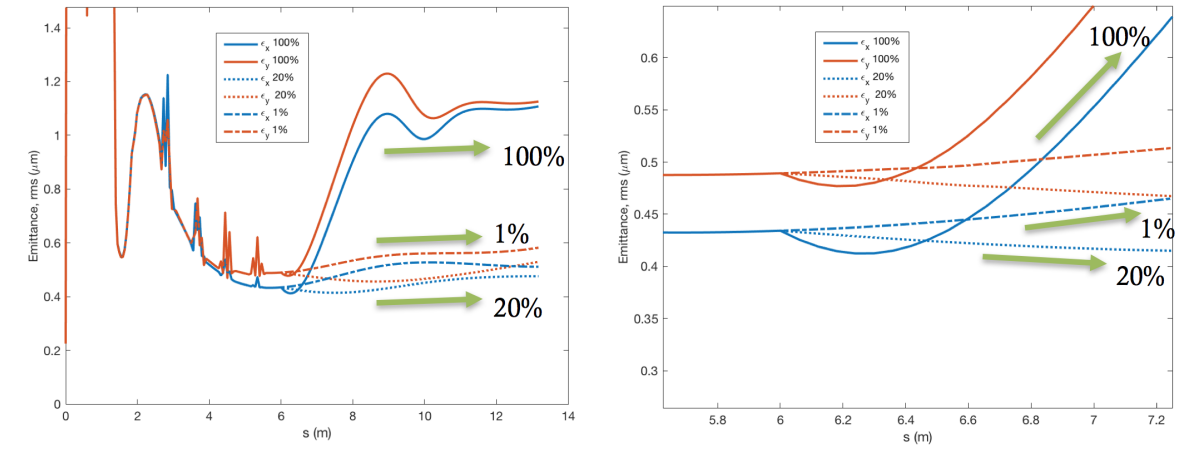

Figure 2.9: Simulations showing beam size for different degrees of ion neutralization. The plot on the right magnifies the region around the point where the ion column begins in order to highlight the strength of the ion focusing.

focusing, the ions can also result in significant emittance growth, as shown in Fig. 2.12. This emittance growth would be difficult to correct using linear optics, and in general suggests that the equilibrium ionization fraction must be reduced to avoid this effect.

It is interesting to note that the curves in Figs. 2.9 and 2.10 appear almost identical. This is no accident, and highlights an important feature of ion simulations. The strength of the ion lens is determined by the number of ions per unit length. Therefore a 40 mA, 125 pC beam at 20% neutralization (effectively a fully neutralized 8 mA, 125 pC beam) will create an ion column that is the same strength as one created by a fully neutralized 6 mA, 125 pC beam. Because the bunch charge in the 40 mA and 6 mA simulations are held fixed at 125 pC, (since the beam current is changed by adjusting the beam repetition rate), this means that the two 125 pC bunches will follow nearly identical trajectories through the same ion column. Therefore, in an actual accelerator it is possible for beams of identical bunch charges but different repetition rates to trace through the same trajectories only if the neutralization fractions are different. If the beams are both fully neutralized, the higher repetition rate beam will more strongly experience the effects

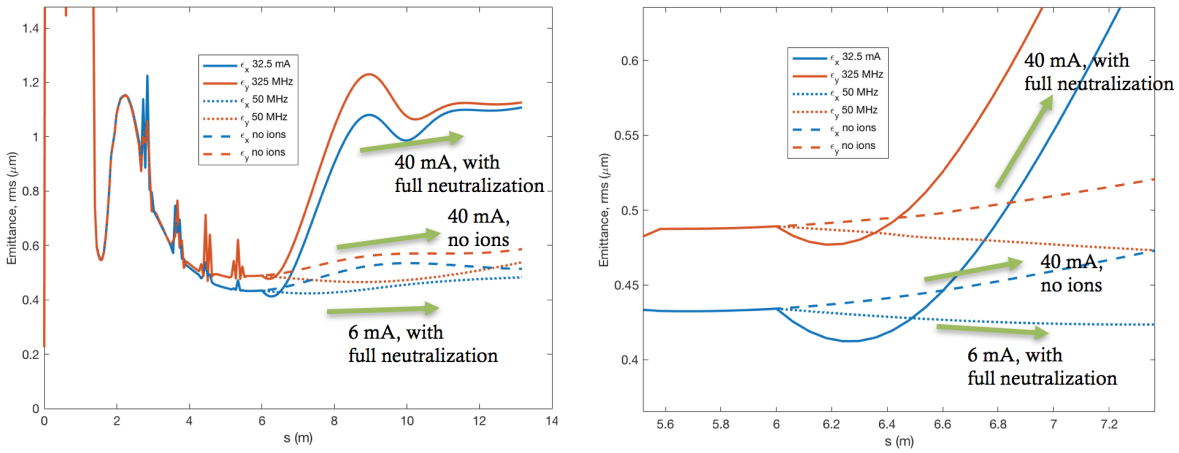

Figure 2.10: Simulations showing the effect of ions as beam current is increased. The plot on the right magnifies the region around the point where the ion column begins. The simulations assume that the ions accumulate until they achieve full neutralization, meaning that the strength of the ion kicks will increase with beam current.

of ions. Therefore one could conduct an experiment by holding bunch charge fixed and slowly increasing the beam repetition rate in order to look for changes in beam trajectories due to ions. If this is done, any changes in beam optics at higher repetition rates would be a very clear indication of ion trapping.

2.3.4 General conclusions

As a general rule of thumb, for both simulation cases tested above the ion effects only become observable above approximately 1 mA. This is because the strength of the ion "lens" depends directly on the maximum beam current. A higher beam current allows more ions to accumulate, which increases the strength of the ion kicks. This current threshold is consistent with experimental observations in the photoinjector [9] and at other accelerators [4].

However, after examining many simulation results, the only real conclusion one can

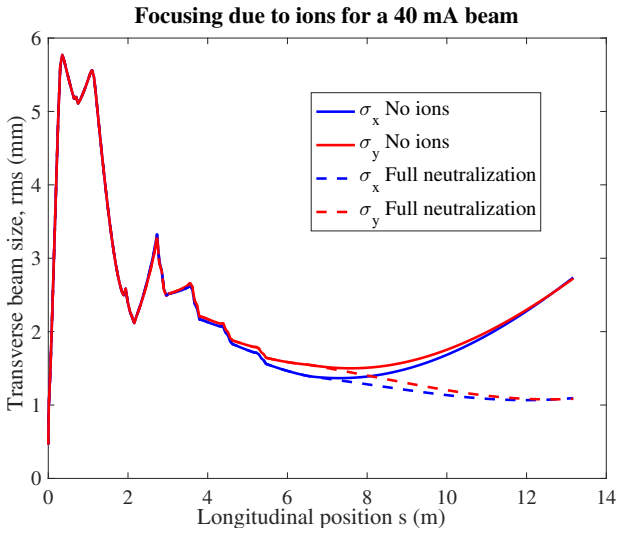

Figure 2.11: Simulations showing the focusing effects of an elliptical Gaussian ion distribution on a 40 mA beam. The ion column begins at s = 6.5 m, immediately after the beam has exited the last SRF cavity of the injector linac.

draw is that the trapped ions significantly change beam optics for sufficiently high neutralization fractions. Because the ion density increases with beam current, the amount of ion clearing necessary to minimize impacts on beam quality will vary largely with beam current and between different accelerators. In the case of the Cornell photoinjector and planned CBETA beam optics, a neutralization fraction of approximately 10% results in a significant decrease in ion-induced beam changes. The optics changes themselves are also not necessarily always harmful, and can be in fact beneficial, for example in the case of screening space charge repulsion. As a result, the only option to determine the severity of ion effects in future accelerators, and judge whether or not mitigation techniques are required, is to simulate their effects in detail.

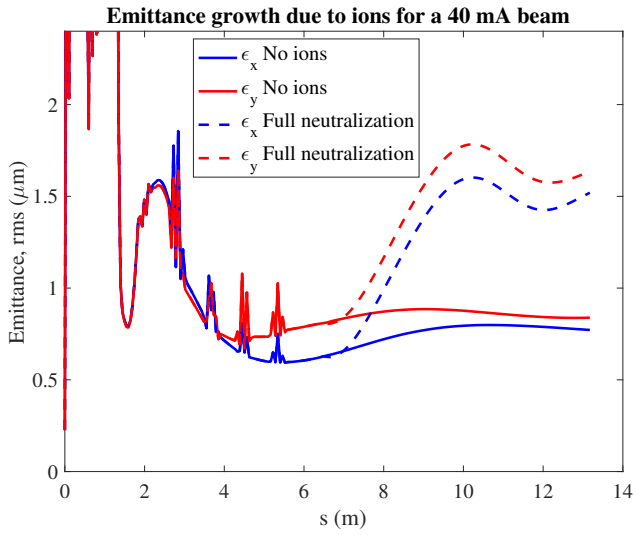

Figure 2.12: Simulations showing the effects of an elliptical Gaussian ion distribution on a 40 mA beam. The ion column begins at s = 6.5 m, immediately after the beam has exited the last SRF cavity of the injector linac.

CHAPTER 3

Ion Clearing Methods: Simulations and testing

The high current Cornell photoinjector is one of the few linacs in the world where different ion mitigation strategies can be experimentally tested for a parameter regime in which ion trapping occurs. We have carried out a series of experiments to test the effectiveness of three different clearing methods: DC clearing electrodes, ion clearing bunch gaps, and resonant beam shaking. In addition, we have developed two separate simulations to aid in the design and placement of clearing electrodes. The process of beam shaking was also confirmed by creating a space-charge simulation code to model beam-ion interactions.

Clearing electrodes are essentially a parallel plate capacitor with an applied DC voltage in the range of 1 V to a few kV. The electrodes are designed to overwhelm the electron beam's attractive potential and allow the ions to escape from the center of the beam pipe. They are best employed in areas of high ion concentrations, which tend to be near beam size minima where ions accumulate due to longitudinal motion [16]. Although clearing electrodes can achieve significant ion clearing, often reducing the beam neutralization fraction to just a few percent [6; 17], it is important to explore other clearing methods as well, because electrodes can be expensive to deploy around the recirculating path [1]. This is especially true in large machines with relatively low beta functions (on the order of meters), as electrodes would need to be installed at most beam size minima.

In the second clearing method, short gaps between bunches are introduced to

either allow the ions to directly drift out of the center of the beam, or to produce an over-focusing that allows ions to reach larger and larger amplitudes until they leave the center of the beam. This technique has been employed in many electron storage rings with great success [6; 29]. In storage rings, this method is typically implemented by leaving a fraction of the ring empty at any given time. In linacs and photoinjectors, this is done by turning the beam off for a given duration and at a certain frequency. Unfortunately, several problems emerge when applying this method to ERLs and CW linacs. The primary concern is that ERLs are particularly susceptible to RF beam loading, because of the large loaded quality factors in their RF cavities [1]. Secondary concerns involve beam instabilities such as the fast ion instability, which occur because ions can accumulate over the course of a single bunch train [14]. One must choose the proper gap duration and frequency to achieve clearing while avoiding these other detrimental effects.

The third method, beam shaking, is a technique that was successfully employed in past accelerators such as the Anti-proton Accumulator Ring [30] and more recently in the Metrology Light Source (MLS) [4]. This method involves applying a time varying voltage to a kicker or other electrode to transversely shake the beam and resonantly clear any trapped ions. Even shaking amplitudes much smaller than the transverse beam size can result in significant reductions in trapped ion density. Typically a single sinusoidal frequency is used, although broadband white noise has also shown to be effective [4]. The necessary frequency is typically close to the ion oscillation frequency, and is usually determined experimentally by trial and error [4; 6].

All three methods can result in significant clearing, and can be even more effective when deployed in tandem [3]. In our experiments, we have examined each method independently to compare their effectiveness at clearing trapped ions, and have also developed several empirical models to explain and analyze our data. The following

sections outlining these experimental results were taken directly from a recent publication in Physical Review Accelerators and Beams [9], written by the author of this thesis.

3.1 Experimental tests of clearing methods

Instead of measuring the effects of ions on the beam, we instead directly studied the trapped ions. We chose to do this because the Cornell photoinjector is a relatively short accelerator, so any changes in beam dynamics due to ions may be difficult to observe directly. Another contributing factor is that most traditional beam diagnostics are not viable in the photoinjector's parameter regime. Due to the beam's high power at full current operation, any traditional interceptive beam diagnostics such as viewscreens, slits or wire scanners will quickly melt (with timescales typically on the order ms or lower). Additionally, because the photoinjector is a low energy linac, we are unable to use synchrotron or diffraction radiation to take measurements. Our best option, a fast beam profile monitor recently developed at Cornell for use in high intensity accelerators [31], was unfortunately not available for use at the time of these experiments.

Larger machines, such as synchrotrons or storage rings, may observe the tune spectrum of a beam using BPMs connected to a spectrum analyzer [4; 21]. Ion-beam interactions lead to incoherent tune spreads and sidebands around the tune, so this is probably the easiest way to observe ion effects. This measurement technique was attempted in the photoinjector, but no ion signatures were observed, even after leaking gas to increase the residual vacuum chamber pressure by a factor of 100. This is likely because the small scale, non-circulating nature of the photoinjector means that beam-ion coupling must be visible on the spectrum analyzer after an interaction region of only about 6 m, which is simply too short.

Instead we used two primary indicators of accumulated ions. The first was a direct measurement of the trapped ion density using our clearing electrode. By applying a DC voltage to the clearing electrode, the ions are drawn out of the center of the beam pipe, strike the clearing electrode and are measured by a picoammeter connected in series with the electrode. The total ion current reaching the clearing electrode depends on the applied voltage, as will be shown below. A sufficiently high voltage (in our case, only 28 V) will draw out all trapped ions in the vicinity of the clearing electrode.

We also used our radiation monitors as a secondary, indirect way of observing the trapped ion density. The high power of the ERL photoinjector's beam generates large amounts of radiation, primarily created by beam losses and beam halo striking the beam pipe. When the beam current was increased above 10 mA after gas injection, measured radiation levels rose sharply above normal background levels, as shown in Fig. 3.1. Before leaking gas, no such excess radition was previously observed in the 10–20 mA range, indicating that this extra radiation (presumably bremsstrahlung) was caused entirely by beam-gas interactions. All clearing methods significantly reduce this radiation, usually returning it to background levels.

The experiments were performed in an approximately 8 m long straight section immediately after the beam exited the final accelerating cavity, as shown in Fig. 3.2. Either N_2, Ar or Kr gas was leaked into the beam pipe so that the dominant ion species was known during the experiment. The pressure in the beam pipe was increased to approximately 100–150 nTorr, as compared to typical values of 1–2 nTorr or less measured during normal operation.

The photoinjector is designed to operate with a beam energy of 5–15 MeV and beam currents up to 100 mA, corresponding to a bunch charge of 77 pC at a repetition rate of

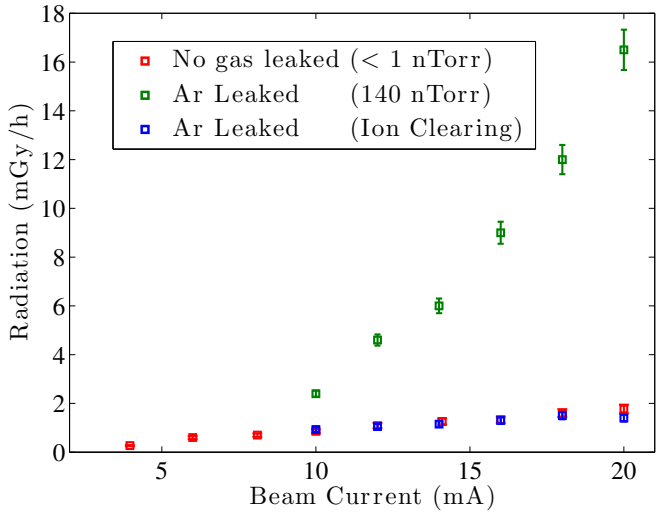

Figure 3.1: After leaking gas into the beam pipe, background radiation levels rose dramatically due to bremsstrahlung generated by beam-ion interactions. Removing the trapped ions using clearing electrodes (shown above) or other clearing methods reduced this excess radiation to normal background levels. (Note that this figure was previously used in section 1.3 as Fig. 1.3 to illustrate a separate point).

1.3 GHz. During these experiments we used a 5 MeV beam and varied the beam current from 1–20 mA by changing bunch charge.

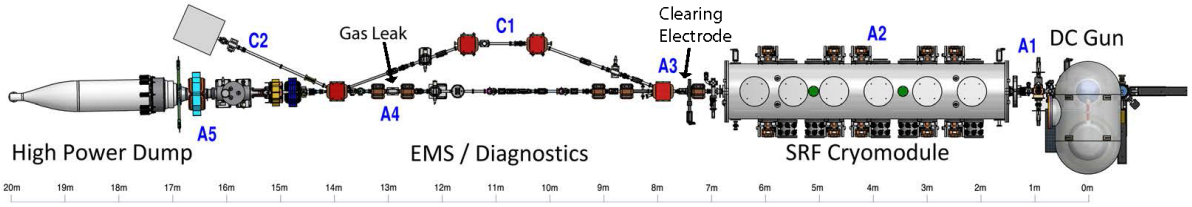

Figure 3.2: A schematic of the photoinjector. The ion clearing electrode was installed just after the beam exists the SRF cavity at A3, and gas was leaked in at the end of A4. Radiation measurements were taken at several locations between sections A3 and A4 (next to the beam pipe). (Note that this figure was previously used in section 1.3 Fig. 1.3 to illustrate a separate point).

3.1.1 Ion clearing electrodes

Although it is possible to use button or stripline BPMs to clear ions, the photoinjector uses a specially created ion clearing electrode. The device schematic is shown in Fig. 3.3 and its location in the beam line is shown in Fig. 3.2. The electrode was oriented vertically during all experiments. The clearing electrode surface is approximately 35 cm long and 3.5 cm wide, and it consists of two layers. The bottom layer is a 0.30 mm thick alumina dielectric coating (in contact with the beam pipe), and the top is a 0.20 mm thick tungsten electrode coating. Alumina was chosen because it is a dielectric material commonly used in vacuum, and tungsten was chosen because it is a well developed thermal-spray material that has a linear thermal expansion coefficient comparable to alumina. The top electrode was attached to a voltage supply, while the other was attached to ground. The electrode's geometry is tapered to reduce wake fields, and it has been designed to allow for a maximum voltage of approximately 4 kV. A picoammeter was attached in series with the voltage supply in order to measure the trapped ion current that was removed by the electrode.

During this experiment we leaked N_2 gas into the beam vacuum chamber to raise the background pressure from a nominal value of less than 1 nTorr to 117 nTorr. This ensured that we knew the dominant ion species present during the experiments. After they are created via collision ionization, the ions drift longitudinally towards beam size minima. This was taken into consideration when choosing beam optics for the experiment.

We varied the applied voltage on the clearing electrode from between 0 V and 28 V to test its effectiveness at clearing ions. We looked at two signatures: the ion current striking the clearing electrode, and the background radiation observed by nearby radiation monitors. Our data, taken for various beam currents between 5 mA and 20 mA, is shown

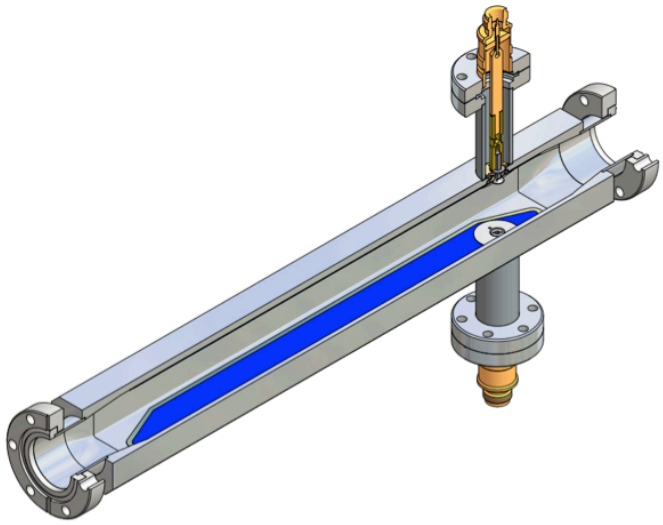

Figure 3.3: The ion clearing electrode used during the experiments. The electrode coating is highlighted in blue, and is approximately 35 cm long and 3.5 cm wide.

in Fig. 3.4 and Fig. 3.5. The beam current was varied by changing bunch charge (from 5 pC to 12.5 pC) at a constant repetition rate of 1.3 GHz.

Even when it was turned off, the clearing electrode measured a small background current. The measured current was typically -4 nA (using the convention of a positive ion current). This was true for both the clearing electrode and bunch gap experiments. Also note that the measured radiation only exceeds background levels for beam currents greater than or equal to 10 mA, as can be seen in Fig. 3.5. This observation is typical throughout our experiments.

The required voltage for maximum ion clearing can be predicted as follows: Assuming a round beam of constant charge density, the beam's electric field is given by [17]

$$E_{beam}(r) = \begin{cases} \frac{\lambda e}{2\pi\epsilon_0}\frac{r}{\sigma_b^2}, & \text{if } r \leq \sigma_b \\ \frac{\lambda e}{2\pi\epsilon_0}\frac{1}{r}, & \text{if } r \geq \sigma_b \end{cases} \tag{3.1}$$

where λ is the number of electrons per unit length, e is the elementary charge, and σ_b is

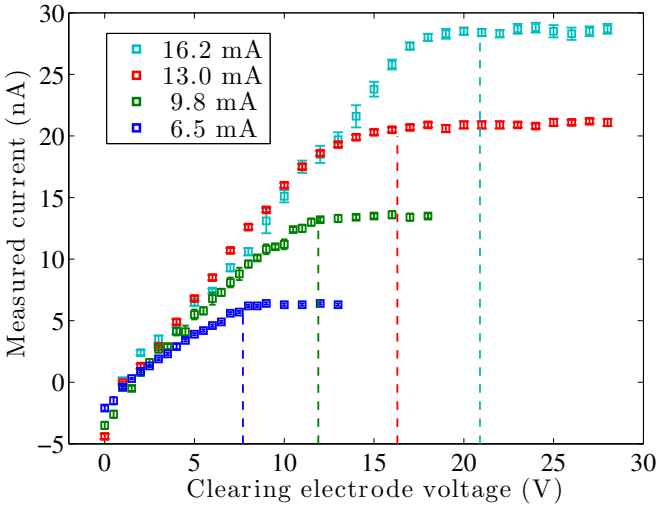

Figure 3.4: A picoammeter was used to measure the ion current striking the clearing electrode for different applied voltages. The vertical dotted lines mark the minimum voltage required for full ion clearing, as predicted using eqn. (3.2).

the rms transverse beam size. Full clearing occurs when the clearing electrode's field exceeds the beam's peak electric field (at $r = \sigma_b$). When this is true, the transverse beam potential is fully suppressed, as illustrated in Fig.3.6, such that no ions may remain trapped within the beam. For electrodes separated by a distance d, the field of the electrode is given by $E = \frac{V}{d}$, as the clearing electrode is very nearly a parallel plate capacitor. Therefore we examine the case where $E_{electrode} = V_{electrode}/d \geq E_{beam}$. This yields the minimum voltage required

$$V_{electrode} \geq \frac{\lambda e}{2\pi\epsilon_0}\frac{d}{\sigma_b} \tag{3.2}$$

Although we cannot take direct beam profile measurements in order to determine σ_b, we can obtain estimates using GPT [10; 11], a 3D space charge simulation code that models the photoinjector. This simulation is found to be in good agreement with results at low beam currents [11], although it has yet to be experimentally verified for beam currents above 1 mA. For now we assume that the ions or other high current effects do not change

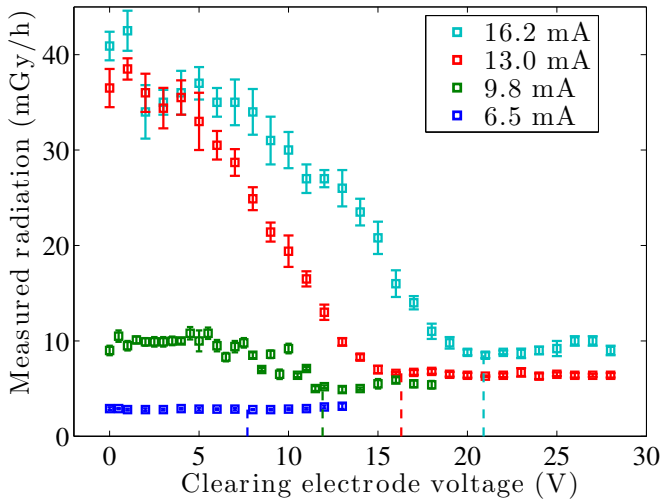

Figure 3.5: Background radiation levels also decreased while employing the clearing electrodes due to the absence of bremmstrahlung caused by beam-ion interactions. The vertical lines again indicate the minimum voltage needed for complete ion clearing, as computed by eqn. (3.2).

the beam size significantly. Using the vertical beam size from GPT and a clearing electrode separation of $d = 40.6$ mm, eqn. (3.2) predicts clearing voltages that agree with measurements to within a few percent, as shown in Fig. 3.4 and Fig. 3.5. Our calculated values are shown in Table 3.1. Vertical dotted lines were drawn in Fig. 3.4 and Fig. 3.5 to guide the eye and make it easier to compare these calculated values with our data.

The maximum measured ion current can be used to obtain an estimate of the longitudinal range of the clearing electrode. In the absence of clearing, ions will accumulate (via collision ionization) until the total number of ions per unit length equals the total number of beam electrons per unit length. The time it takes to accumulate as many ions as electrons per unit length is given by [1]

$$\tau_{create} = \frac{1}{\sigma_{col}\rho_{gas}c} \tag{3.3}$$

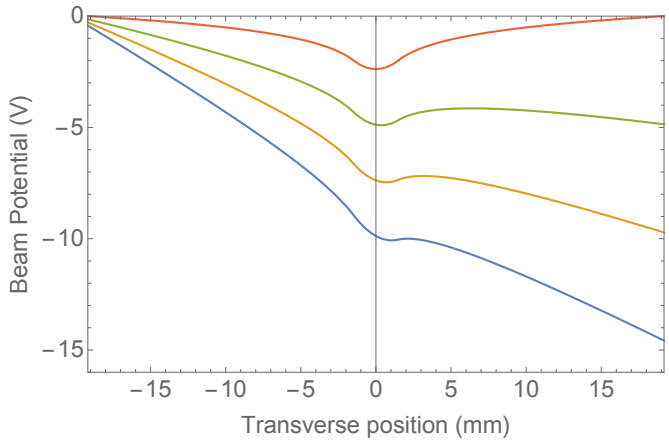

Figure 3.6: As the clearing electrode voltage is increased, the transverse beam potential is suppressed. At the minimum voltage required for full clearing, no ions can remain trapped within the beam.

Table 3.1: The minimum clearing electrode voltage necessary for full ion clearing, calculated using eqn. (3.2). The rms transverse beam sizes σ_x and σ_y were obtained using GPT simulations [11].

Current (mA)	σ_x(mm)	σ_y(mm)	Voltage (V)
6.5	2.11	2.05	7.7
9.8	2.07	2.00	11.9
13.0	2.03	1.94	16.3
16.2	1.98	1.89	20.9

where σ_{col} is the collision ionization cross section of the gas species (these values are readily available [7]), ρ_{gas} is the residual gas pressure, and c is the speed of light.

Over a longitudinal region L, there are λL beam electrons. We define a region L_{create} such that all of the ions over this region are removed by the clearing electrode. On average, an electron needs the time τ_{create} to produce one ion. In the length L, each electron produces $L/c\tau_{create}$ ions. The beam current therefore produces $L/c\tau_{create} \cdot I/e$ ions per second in this section. Thus the maximum ion current is

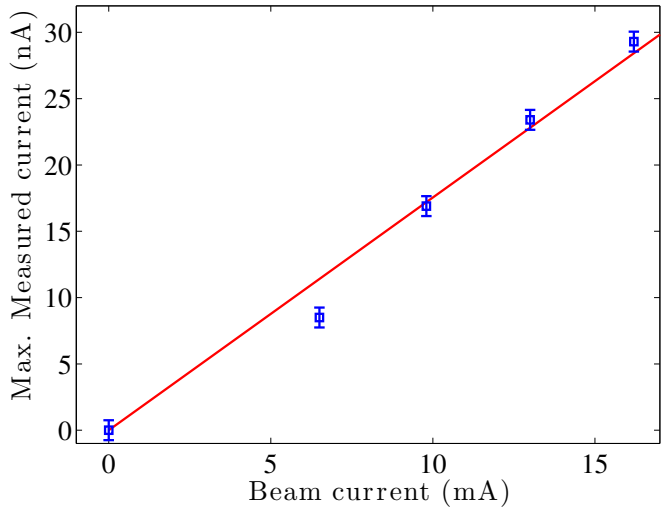

Figure 3.7: The longitudinal range of the clearing electrode can be estimated by using the maximum ion current measured by the clearing electrode. The range can be extracted from a best fit line (in red) to our data.

$$I_{max} = \frac{I_{beam}}{c}\left(\frac{L_{create}}{\tau_{create}}\right) \tag{3.4}$$

The ratio L_{create}/τ_{create} can be found by treating it as a fit parameter to our data, shown in Fig. 3.7. Note that in Fig. 3.7 we have adjusted our data to take into account background levels.

If the value for τ_{create} is known, then we can obtain an estimate for the longitudinal range L_{create}. For a N_2 gas pressure of 117 nTorr and an assumed temperature of 300K, this corresponds to a beam neutralization time τ_{create} of roughly 5.2 ms. Therefore, using the fit parameter found in Fig. 3.7, the creation length L_{create} should be roughly 2.7 m. This estimate seems reasonable, considering that the distance between the gas leak and the clearing electrode is roughly 5 m. In reality, it is possible that the longitudinal range will increase with the applied clearing voltage. In the future, taking data points for more beam currents should allow us to find a more accurate estimate.

3.1.2 Bunch gaps

While storage rings can create gaps simply by leaving a fraction of the ring empty at any given time, CW linacs require the introduction of a short bunch gap every few milliseconds. In our experiments this was achieved by using a Pockels cell (normally used to select pulses for our low repetition rate emittance measurements [11]) to reject laser pulses with a given duration and at a certain frequency. Due to the hardware limitations of our Pockels cell, we were unable to create a bunch gap larger than 10 μs. This experiment was performed immediately following the ion clearing electrode experiments, so the gas pressure remained at 117 nTorr for N_2.

During these gaps, the regulation of the fields in the SRF and buncher cavities struggled to handle the sudden change in beam loading. As the beam current was increased, the error in the field amplitude and phase increased until they reached their pre-defined limits and tripped off the machine at around 8 mA. We have a pre-existing feedforward system [32], originally designed to handle the analogous situation when there are short bunch trains for emittance measurements. This system was able to completely remove the amplitude and phase errors in the SRF cavities without any modifications. However, for the buncher, the feedforward became unstable above a certain amount of gain. At the time of the experiment, we decided to just limit the gain rather than investigate the cause of the instability. As a result, we were limited in beam current to around 20 mA by the remaining phase error in the buncher. Ultimately, we believe that this is not a fundamental limitation to this bunch gap clearing method. With more work, we believe we could solve this problem in the future.

When employing bunch gaps, a fraction of the trapped ions drift transversely out of the beam during the gaps and into the vacuum chamber walls. The remaining trapped

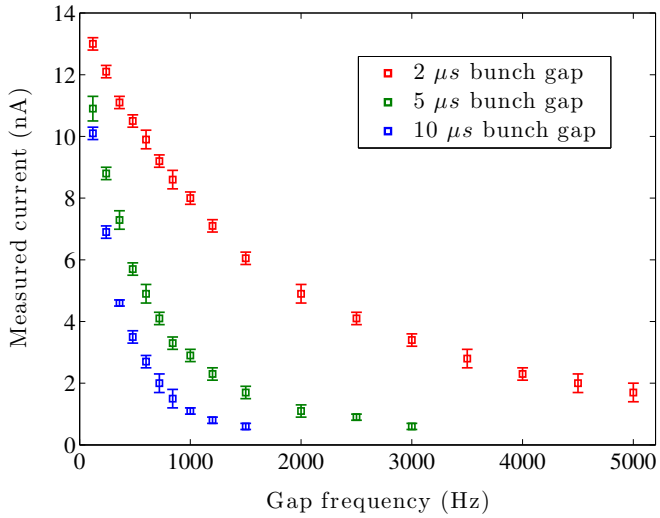

Figure 3.8: Increasing the frequency and duration of bunch gaps reduces the number of trapped ions that reach the clearing electrode. For each data point, the beam current was held fixed at 10 mA.

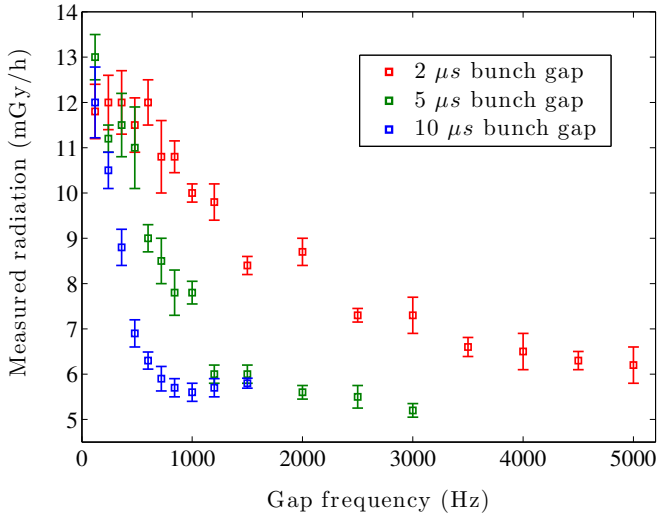

Figure 3.9: The radiation caused by beam-ion interactions is reduced by increasing the frequency and duration of bunch gaps. This data was obtained while the clearing electrode was turned off. For each data point, the beam current was held fixed at 10 mA.

ions travel longitudinally down the beam pipe towards our clearing electrode and are measured by the picoammeter. We applied a large enough voltage (28 V) to the clearing electrode to ensure maximum ion clearing. Thus we are measuring the amount of ions that remain trapped in the beam after clearing via bunch gaps. Data for an average beam current of 10 mA was taken for various bunch gap lengths and frequencies, and is shown in Fig. 3.8 – 3.9. The radiation data in Fig. 3.9 shows that the trapped ions are removed even without the clearing electrode turned on, confirming that the bunch gaps are the dominant clearing mechanism. This data also agrees with previous experiments at SPring-8 [29] that observed a significant decrease in bremsstrahlung radiation when employing bunch gaps.

For now we have devised an empirical model to explain our data. In our simple model we assume that while the beam is on, ions are created via collision ionization. The neutralization fraction should never exceed 1, because when the beam is fully neutralized, its potential well will be suppressed and the ions will begin to drift out of the center of the beam. We model the process of increasing neutralization fraction during an electron bunch train using $f(t) = 1 - (1 - f_0)\exp(-t/\tau_1)$, where τ_1 is a time constant that defines the ion creation rate and f_0 is the initial neutralization fraction (which is not necessarily 0 in the steady state). While the beam is off, the trapped ion density decays exponentially with a time constant τ_2. Figure 3.10 illustrates this creation and clearing process.

The average value of the neutralization fraction (i.e. steady state solution) determines the amount of clearing we have observed experimentally. From our model, the average ionization fraction is given by

$$f_{avg}(R_g) = \frac{1}{1 + \left(\frac{\tau_1}{\tau_2}\right)\left(\frac{R_g \Delta L_g}{1 - R_g \Delta L_g}\right)} \tag{3.5}$$

where τ_1 is the characteristic creation time, τ_2 is the characteristic clearing time, R_g is the bunch gap frequency, and ΔL_g is the bunch gap duration. Note that this is an approximate form, valid only for $(1/R_g - \Delta L_g)/\tau_1 \lesssim 0.5$ and $\Delta L_g/\tau_2 \lesssim 0.5$. The full expression is derived in Appendix A. The parameters R_g and ΔL_g are well defined in the experiment, but the ratio of the two time constants τ_1 and τ_2 must be determined empirically from our data. Our fit curves are compared with our data in Fig. 3.11, and the best fit ratio of τ_1/τ_2 for each curve is shown in Table 3.2. Even though the clearing electrode may have had a small impact on the clearing rate τ_2, we chose to compare our model to the data in Fig. 3.8 because it has less experimental uncertainty.

Assuming that τ_1 is roughly the time it takes to achieve full beam neutralization (i.e. $\tau_1 = \tau_{create} = 5.2$ ms), then the ratio τ_1/τ_2 can be used to predict a clearing time of 21 μs. This number is consistent with clearing rates estimated using the ion oscillation frequency (to be further explained in the next section). For example, according to our data at 10 mA, a N_2 ion has an oscillation period of 17.5 μs, which is of the same order of magnitude as this estimate for τ_2.

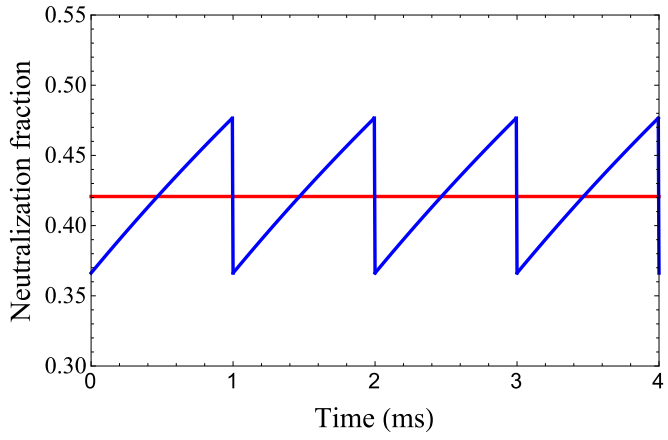

Figure 3.10: Ions are created via collision ionization while the beam is on and decay exponentially during the bunch gaps. The equilibrium neutralization fraction, indicated by the red line, was found using eqn. (3.5) for a bunch gap duration of 5 μs, a gap frequency of 1 kHz, and the ion creation to clearing ratio given in Table 3.2.

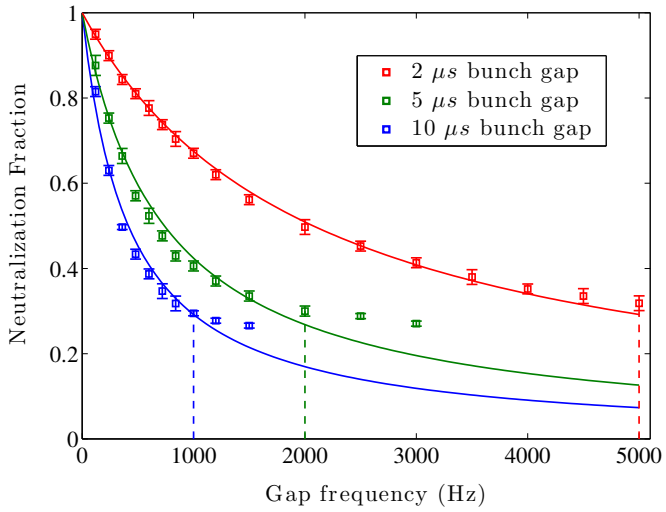

Figure 3.11: Increasing the frequency and duration of bunch gaps reduces the trapped ion density as shown by the residual ion current hitting a clearing electrode. The curves are best fits obtained using eqn. (3.5). The vertical dotted lines mark the locations where a 1 percent reduction in beam current results in an approximately 70 percent reduction of trapped ion density.

Of particular note is that this simple model predicts the average amount of clearing depends only on the total time the beam is turned off ($R_g \Delta L_g$). This observation agrees with our data. For example, a 2 μs gap at 5 kHz achieves the same amount of clearing as a 10 μs gap at 1 kHz. In other words, the two data sets overlap when the horizontal axis is adjusted so that it becomes the total time the beam is turned off. This suggests that the bunch gap method offers some flexibility, and may potentially allow the user to avoid RF beam loading problems by choosing the correct combination of gap duration and frequency.

Our data in Fig. 3.11 shows that an approximately 70 percent reduction in ion density can be achieved while retaining nearly 99 percent of the maximum beam current. To further reduce the number of trapped ions, one must increase the bunch gap frequency or duration, and introduce even more beam downtime. According to our model, a 99 percent reduction in ion density would require over a 30 percent reduction in maximum

Table 3.2: The ratio of creation to clearing times, found empirically from fits to our data.

Gap duration (μs)	τ_1/τ_2
2	2.4×10^2
5	2.7×10^2
10	2.4×10^2

beam current. This is unacceptable for most ERL applications. However, this is a large extrapolation of our model, and more data must be taken to determine the true limits of this clearing method. The shortest possible gap that can still achieve clearing also has yet to be determined. This is something we would like to examine in future experiments.

3.1.2.1 Derivation of average neutralization fraction

To explain our bunch gap clearing data, we seek to obtain an expression for the average beam neutralization fraction as a function of bunch gap duration and frequency. We begin by considering the creation rate of ions while the beam is turned on. Ions are generated via collision ionization inside of the beam at a constant rate, therefore the ion density increases linearly with time. However, the neutralization fraction can never exceed 1, because the positive ions will eventually screen the electron beam's negative potential and the ions will begin to drift out of the center of the beam pipe. We model this behavior by assuming the neutralization fraction has a functional form of

$$f_1(t) = 1 - (1 - f_0)e^{-t/\tau_1} \tag{3.6}$$

where τ_1 is a time constant that defines the creation rate and f_0 is the initial neutralization fraction (which is not necessarily 0 in the steady state). This expression applies only while the beam is on, up until some time T_1, as shown in Fig. 3.12.

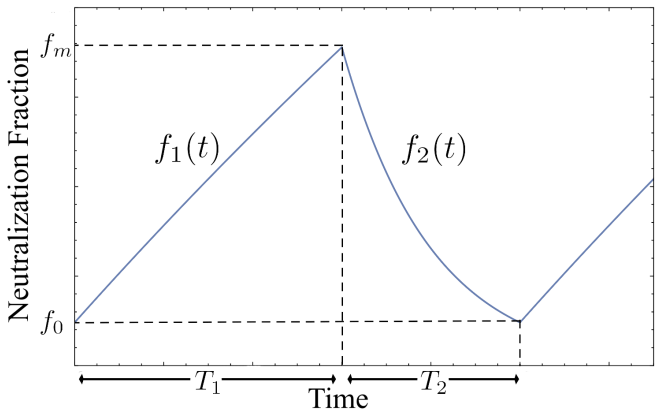

Figure 3.12: A sketch of the neutralization fraction while employing bunch gaps. The neutralization fraction increases while the beam is on until some time T_1, then decays exponentially during the bunch gaps with a duration of T_2.

During the bunch gaps, the neutralization fraction decays exponentially as given by

$$f_2(t) = f_m e^{-(t-T_1)/\tau_2} \tag{3.7}$$

where τ_2 is the characteristic clearing time, and the constant f_m is the maximum neutralization fraction reached in the steady state (because $f_2(T_1) = f_m$ in Fig. 3.12). Using this, along with the fact that the total function must be continuous at both $f_1(T_1 + T_2) = f_2(T_1 + T_2) = f_0$ and $f_1(T_1) = f_2(T_1) = f_m$, one can obtain expressions for both f_0 and f_m,

$$f_0 = \frac{1 - e^{\frac{T_1}{\tau_1}}}{1 - e^{\frac{T_1}{\tau_1} + \frac{T_2}{\tau_2}}}, \tag{3.8}$$

$$f_m = 1 - (1 - f_0)e^{-T_1/\tau_1}. \tag{3.9}$$

Now that $f_1(t)$ and $f_2(t)$ are well defined, the average ionization fraction can be found via integration over the full time interval $T_1 + T_2$. This is given by

$$f_{avg} = \frac{1}{T_1}\int_0^{T_1} f_1(t) \;\; + \;\; \frac{1}{T_2}\int_{T_1}^{T_1+T_2} f_2(t). \tag{3.10}$$

Performing this integral yields the formula for the average neutralization fraction

$$f_{avg} = \frac{T_1}{T_1+T_2} - 2\frac{\tau_1-\tau_2}{T_1+T_2}\frac{\sinh\left(\frac{T_1}{2\tau_1}\right)\sinh\left(\frac{T_2}{2\tau_2}\right)}{\sinh\left(\frac{T_1}{2\tau_1}+\frac{T_2}{2\tau_2}\right)}. \tag{3.11}$$

This expression can be simplified for small T_1/τ_1 and T_2/τ_2 by using the approximation $\sinh(x) \approx x$. This ultimately yields an approximate form of the average neutralization fraction given by

$$f_{avg} = \frac{1}{1+\left(\frac{\tau_1}{\tau_2}\right)\left(\frac{T_2}{T_1}\right)}. \tag{3.12}$$

This approximation works well for $T_1/\tau_1 \lesssim 0.5$ and $T_2/\tau_2 \lesssim 0.5$. By substituting $T_2 = \Delta L_g$ for the bunch gap duration and $T_1 = 1/R_g - T_2$ for the time the beam is on, one can obtain the expression found in eqn. (3.5).

3.1.3 Beam shaking

In addition to their longitudinal drifting, the ions oscillate transversely in the beam's potential well. One can imagine that the ion cloud and electron beam form a coupled oscillator. By driving the beam at the trapped ions' oscillation frequency, a resonance is induced that kicks the ions out of the center of the beam. This characteristic frequency should depend on the beam size and beam current.

In order to determine the frequency of trapped ion oscillations inside an electron

beam, we must first calculate the force acting on the ions. The Coulomb force generated by an infinitely long, rotationally symmetric Gaussian beam can be derived using Gauss's law, and is given by

$$F(r) = \frac{\lambda e^2}{2\pi\epsilon_0 r}\left[1 - \exp\left(-\frac{r^2}{2\sigma_b^2}\right)\right] \tag{3.13}$$

where r is the distance from the center of the beam, λ is the number of electrons per unit length, and σ_b is the rms width of the electron beam. According to our simulations [11], the beam in the photoinjector is very nearly round for our experimental parameters, making this an appropriate approximation for our case. By linearizing this force, we are able to treat the ion's motion inside the beam as a simple harmonic oscillator. The equation of motion in this case is then

$$\frac{d^2r}{dt^2} + \omega_i^2 r = 0 \tag{3.14}$$

where ω_i is the oscillation frequency of the ions. Using the linearized form of (3.13), it follows that this oscillation frequency is given by [12]

$$\omega_i = \sqrt{\frac{2r_p c}{e}\frac{I}{A\sigma_b^2}} \tag{3.15}$$

where I is the beam current, A is the atomic mass of the ion species, and r_p is the classical proton radius. This formula can be used to estimate the frequency needed to clear out the ions. Over the course of this experiment, we attempted to test the validity of this theory as well as the scaling laws it predicts.

During this experiment the clearing electrode was used to shake the beam vertically. It was placed approximately 1 m from the exit of the accelerating cavity. A sinusoidally time

varying voltage was applied to the electrode using a function generator and high voltage amplifier. The voltage was applied to only one side of the clearing electrode, while the other side was attached to ground. Oscillation frequencies were predicted to be in the 10-100 kHz range, so this is the primary range over which the experiment was performed.

Because our clearing electrode was being used to shake the beam, we could not measure the residual ion density using the picoammeter and electrode. We were instead forced to rely solely on our indirect radiation measurements. When the ions are cleared from the center of the beam pipe at resonance, the excess radiation caused by beam-ion collisions should vanish. Thus, by measuring this radiation as a function of beam shaking frequency and noting the frequencies where the radiation vanishes, we are able to determine the oscillation frequencies of the ions. The maximum voltage applied to the clearing electrode was adjusted as needed to completely clear the radiation at resonance, but the shaking amplitude never exceeded 0.5 mm for beam sizes of approximately 2–4 mm. Generally results were visible for a shaking amplitude of roughly 0.1 mm. An example of a typical measurement is shown in Fig. 3.13. Measurements were taken for several gas species, including N_2, Ar and Kr, in order to confirm that the measured resonance frequency scaled correctly with ion mass.

An attempt to shake the beam using broadband white noise was made, but this method did not result in any observable reduction in radiation. At the present time, we do not understand why this method works in the MLS [4] and not the Cornell photoinjector.

In the course of our experiments we attempted to confirm the three scaling laws predicted by eqn. (3.15): resonance frequency as a function of beam current, ion mass, and beam size. Because the resonance peaks were quite broad, a fitting algorithm was used to fit the data, and the maximum value was taken as the resonance frequency.

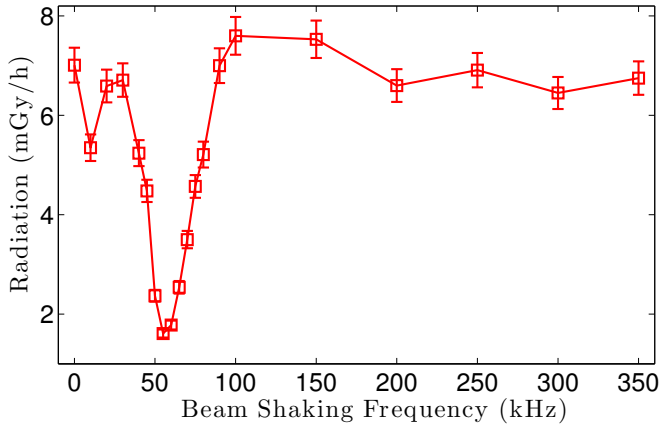

Figure 3.13: Shaking the beam at frequencies near the ion oscillation frequency eliminates the excess radiation caused by beam-ion interactions.

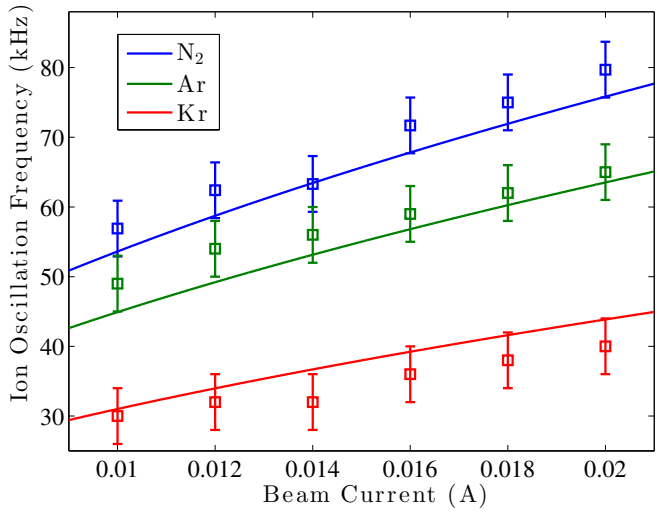

Figure 3.14: Resonance frequencies for various beam currents and ion species. The circles represent data points, while the lines indicate best fits in the form of eqn. (3.15), where the beam size is used as a fit parameter.

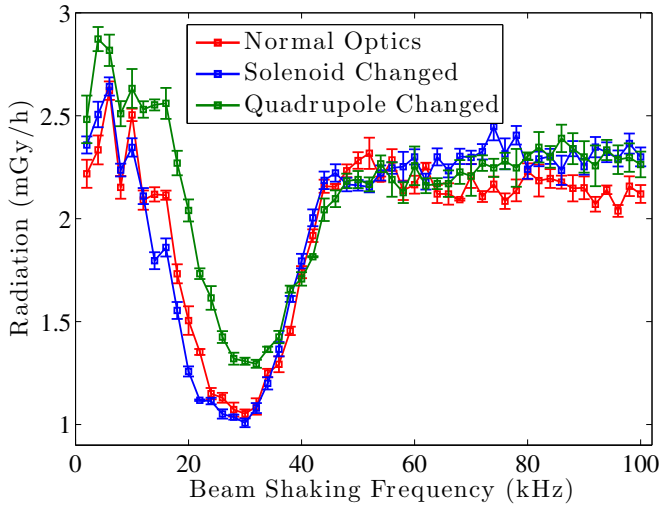

Figure 3.15: Radiation levels for various beam shaking frequencies. Changing the beam size by over a factor of 2, as shown in Fig. 3.16, does not result in a significant shift of the resonance frequency, in disagreement with theory.

Figure 3.14 shows the measured resonance frequencies for beam currents over the range 10-20 mA, and for three different gas species. Error bars for the data points are typically ± 3 kHz, and are due to systematic shifts in resonance frequency due to changing the electrode voltage, as well as statistical fluctuations.

Here it is shown that the resonance frequency scales as predicted with beam current and ion mass. This suggests that the resonance frequency required to clear the ions is indeed the ion oscillation frequency. Given our lack of actual beam size measurements during this experiment, and the fact that GPT has not been experimentally verified in this beam parameter range, the beam size was treated as a fit parameter for our data. A value of σ_b = 4.2 mm was used to obtain the fit curves for the data in Fig. 3.14. GPT predicts a beam size of approximately 2 – 3 mm between the clearing electrode and the gas leak, which is reasonably close to this fit value.

However, the resonance frequency did not scale with beam size, as predicted by our theory. Changing the beam size by almost a factor of 3 (using GPT simulations as a guide

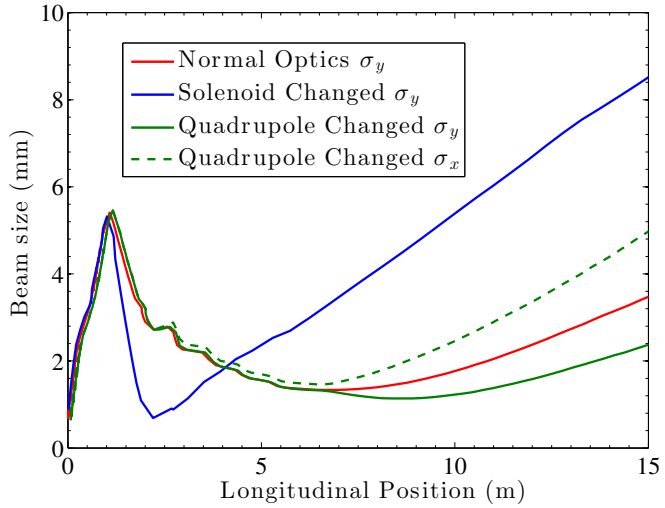

Figure 3.16: The three optics settings used while obtaining the data in Fig. 3.15. The beam was round for both the normal and solenoid optics settings. The clearing electrode and gas leak were located at 7.5 m and 13 m, respectively, and the beam exits the SRF cavity at 6.7 m.

during operation) at the gas leak using a solenoid or a quadrupole magnet lead to a negligible change in resonance frequency, as shown in Fig. 3.15. The simulated beam sizes for the three different optics settings used during the experiment are shown in Fig. 3.16.

There are a few possible explanations as to why no change was observed. This factor of 3 change in beam size was predicted using GPT, and it is possible that the beam size was not actually changing during the experiment. At the time of these experiments, the current version of GPT did not include ion effects on the beam, and they might be somehow relevant to keeping the beam size constant. Another possibility is that changing the optics settings only longitudinally moved the beam waist location while keeping the beam size constant at that waist. Because the ions accumulate at beam size minima, the longitudinal location of the minima does not matter as much as the transverse beam size at that location. However GPT does not predict this sort of behavior, as illustrated by comparing beam sizes in Fig. 3.16. In the future, repeating this

experiment with beam profile diagnostics is necessary to determine why we were unable to observe a dependence on transverse beam size.

3.1.4 Experiment summary

From an ion mitigation standpoint, clearing electrodes appear to remain the most straightforward option. A single electrode seems to clear most of the trapped ions in the photoinjector, especially because the region of interest is rather short (only about 5 m). A larger accelerator would require the deployment of clearing electrodes near most beam size minima and other pockets of high ion concentration. This may become difficult or expensive to implement in machines with low beta functions (on the order of m). In this case, simulations must be done to determine the optimal placement of clearing electrodes [16]. The voltage required for full ion clearing can be predicted using a simple formula, and this can be used to design a proper clearing electrode. For the photoinjector, the required voltage was rather small (28 V) compared to much higher energy accelerators which may require upwards of 1kV [3; 4]. This is due to a large difference in transverse beam size, which is typically mm in the photoinjector, as compared to beam sizes on the order of 10 μm in storage rings or synchrotrons.

In larger accelerators, beam shaking may be a more cost-effective option, because it only requires installing one or two electrodes to shake the beam. The question remains whether or not the shaking amplitude is tolerable, as transversely shaking the beam can lead to undesirable effects such as emittance dilution. However, in practice, shaking appears to work for amplitudes that are much smaller than the transverse beam size, which may lessen these drawbacks. Also, employing other clearing methods (such as

bunch gaps) in tandem with beam shaking may allow one to further avoid the drawbacks of this method by reducing the necessary shaking amplitude.

The most promising results from these experiments are the bunch gap measurements. Previously it was thought that this method was impossible in ERLs due to problems with RF beam loading [1]. However, the size and frequency of the bunch gaps appears to be rather flexible, as the amount of clearing depends only on the total time the beam is turned off. In addition, significant clearing can be achieved while retaining nearly 99 percent of the nominal beam current. With this new information in mind, an analysis of beam loading and bunch gap mitigation in ERLs merits further study.

In the future, we would like to continue these experiments with a new beam profile monitor capable of operating at high beam current [31]. This diagnostic is currently undergoing bench testing and should be available soon. This will allow us to determine transverse beam sizes and supplement our results. In addition, these measurements will be our first glimpse of any beam changes due to ions at high current in the photoinjector. An attempt will also be made to take data for even higher beam currents above 20 mA.

3.2 Simulations of clearing methods

We developed several simulation codes to support our observations above. First, we used a 3D Poisson solver to determine the maximum clearing voltage required to achieve full ion clearing for electrodes of unique geometries. Second, because the optimal placement of clearing electrodes along the accelerator is also important, we developed a simulation code to track the longitudinal motion of ions through the beam in order to confirm areas

of high ion concentration. Finally, we created a code to model the process of beam shaking and determine its effectiveness as an ion clearing mechanism.

3.2.1 Clearing electrode design

When designing a simple parallel plate style ion clearing electrode, eqn. (3.2) produces a good estimate of the required clearing voltage. Provided this voltage is sufficiently small (hundreds of V or less), biasing button BPMs with a DC voltage is generally sufficient to eliminate any trapped ions. Thus, an accelerator designer might be able to forgo designing specialty clearing electrodes and plan to react to ion problems on the fly with preexisting accelerator components. This is clearly the most cost effective option, provided the button BPMs can easily support biased operation.

However, if the required voltage exceeds roughly 1 kV, then button BPMs are no longer an option for several reasons. First, Type-N connectors (as opposed to common SMA connectors) cannot support a voltage above 1.5 kV [33]. Second, the button BPM structure itself may or may not be designed to sustain such a high voltage. And third, if one wants to both bias the button BPMs and still use them to detect beam positions, one would need to supply this voltage in a way that doesn't interfere with the high frequency performance of the signal path. This can be very challenging to achieve [33]. Thus, in general it is better to design dedicated clearing electrodes for a large required voltage.

Equation (3.2) predicts that successful ion clearing for a single pass in CBETA requires a voltage of ± 1 kV for a minimum beam size of 30 μm, creating such a need for a special clearing electrode. This abnormally high voltage is a direct result of extremely small rms beam sizes in the straight sections. Before fully committing to building many specialized electrodes (which can cost thousands of dollars per electrode) based on this formula

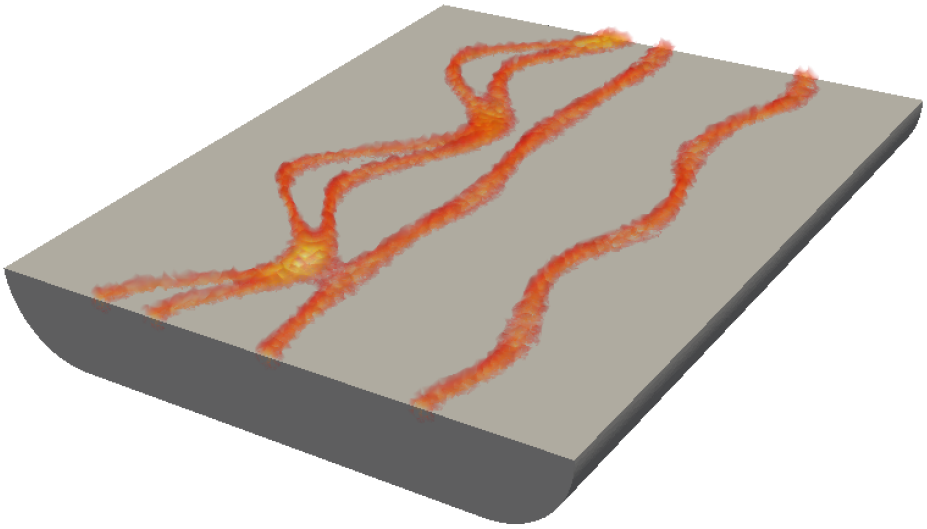

Figure 3.17: Four different DC beam charge distributions, corresponding to four simultaneous beam trajectories in the preliminary CBETA beam pipe. The beam size has been magnified by a factor of 50 from CBETA's smallest design value (roughly 30 μm) to improve visibility.

alone, we decided to develop a code to solve Poisson's equation and provide more accurate estimates of required clearing electrode voltages.

CBETA is also unique because it is an ERL which stores 8 beams at 4 different energies (42, 78, 114 and 150 MeV) [13; 28]. Each individual beam follows a different trajectory according to its energy, meaning that multiple beam potentials may overlap in certain places and ions will accumulate in deeper pockets according to how these beam trajectories overlap. Thus, the simple analytical formula may fail to take this into account also, giving us further motivation to develop a simulation code.

Outside of writing one's own simulation code, an option is to use a one of many 2D Poisson solver codes commonly used in accelerator physics, such as Poisson Superfish [34], and "stitch together" multiple 2D slices to form a full 3D solution. The other alternative is to use a fully 3D Poisson solver. We chose to use a relatively new (as of the

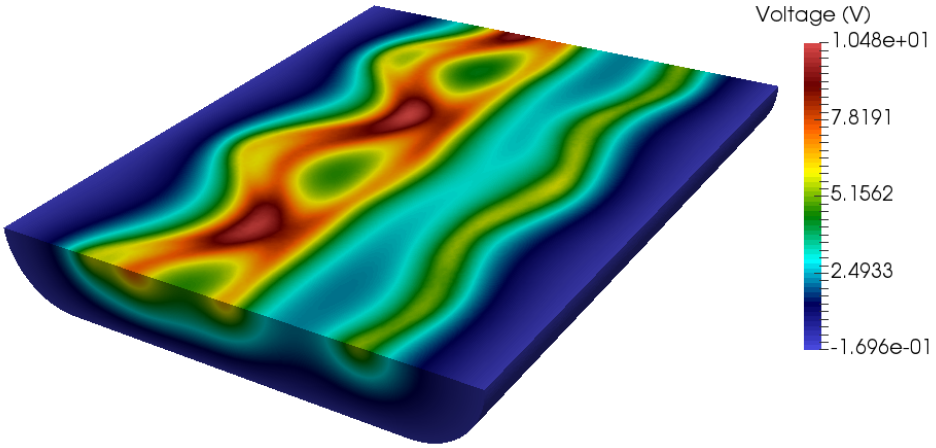

Figure 3.18: The four overlapping beam potentials of the beams shown in Fig. 3.17, obtained using FEniCS. These plots highlight areas of strong electric beam potential, and therefore indicate possible areas of high ion accumulation.

time of this writing) code package called FEniCS. The FEniCS Project is a collaborative project whose primary goal is to solve differential equations numerically using the finite element method [35; 36]. It is not built specifically to solve just Poisson's equation, but rather capable of dealing with any partial differential equation (PDE). FEniCS itself is actually a collection of many components (DOLFIN, FFC, FIAT, Instant, UFC, UFL, and others) all sewn together with a Python interface. Because the software package used the finite element method, it requires a mesh to model the relevant geometry, which we generated using Gmsh [37]. Paraview [38] was used for analyzing and viewing 3D renderings of solutions obtained from FEniCS.

As a first test, we inserted several DC beams into the initial concept CBETA beam pipe chamber design, as shown in Fig. 3.17. The DC beams have a Gaussian transverse charge density, and undergo betatron oscillations that match closely with CBETA's proposed design in the arc sections. Then we used FEniCS' algorithm to solve Poisson's equation, assuming that the beam pipe is grounded, and that the clearing electrodes were held fixed at a constant voltage. An example solution for a three button BPM configuration is

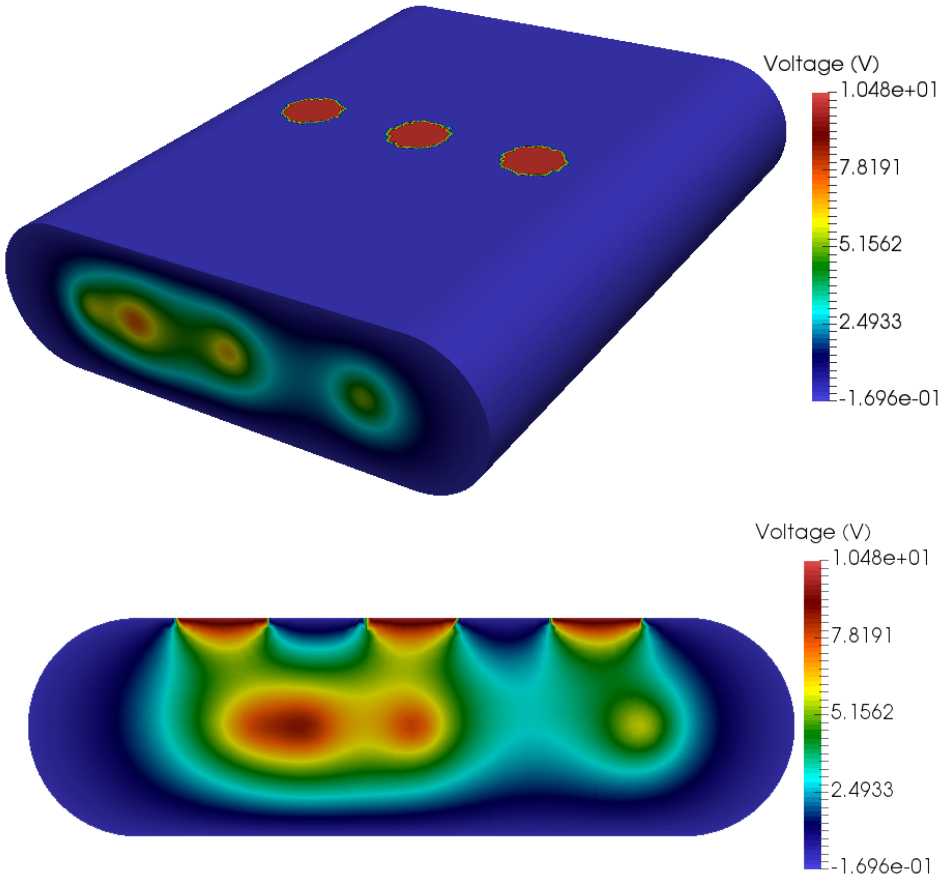

Figure 3.19: A three button BPM arrangement using the preliminary beam pipe dimensions for the CBETA project. A 2D cross section of the beam potential (bottom) allows one to observe if the beam potential is sufficiently suppressed to allow ions to drift out of the center of the beam pipe.

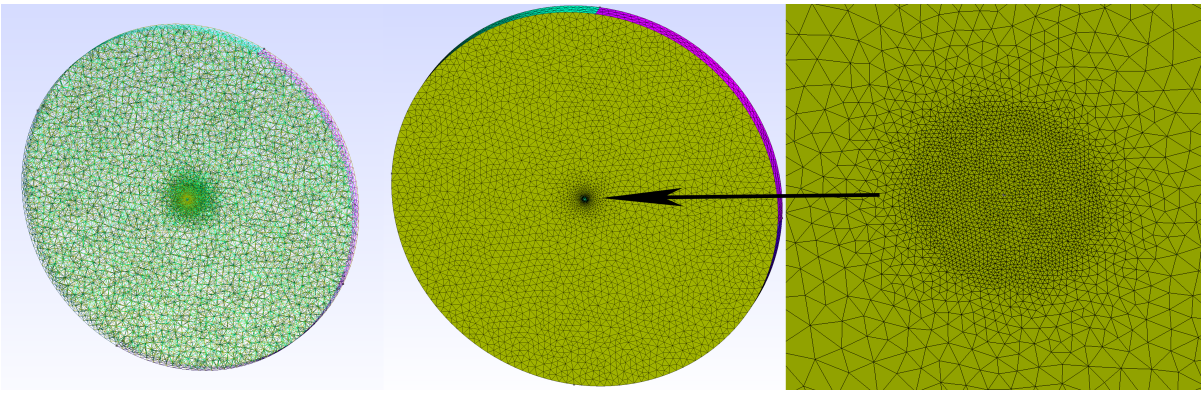

Figure 3.20: Very small beams sizes (30 μm) relative to the beam pipe radius (2.5 cm above) necessitated the use of a variable mesh size (2.5 μm in the center of the beam pipe above).

shown in Figs. 3.18 and 3.19. When the clearing electrode voltage is strong enough, the transverse beam potential will be suppressed, allowing the ions to escape. The suppression of the transverse beam potential can be determined by examining 2D cross sections of the 3D solution, as shown in the bottom half of Fig. 3.19.

The size of the beam used in the above example was exaggerated by a factor of 50 for the purposes of illustration. In reality, CBETA's beams are very small (as small as 30 μm) compared to a beam pipe size of several cm. This large length scale disparity leads to several problems during simulation. The biggest issue is that the mesh size must be very fine in order to resolve the potential near the center of the beam. In theory this can be solved by creating a very fine mesh size, however creating a μm sized mesh over several cm would create prohibitively large computation times and mesh file sizes. Thus, a variable mesh was used with a very fine spacing near the center of the beam and a much larger spacing outside of it, as is illustrated in Fig. 3.20.

The accuracy of this method was tested by comparing a numerical solution to the analytical formula for a round Gaussian beam inside of a round beam pipe. For several test cases, the numerical solution agreed with the analytical result to within roughly 20%

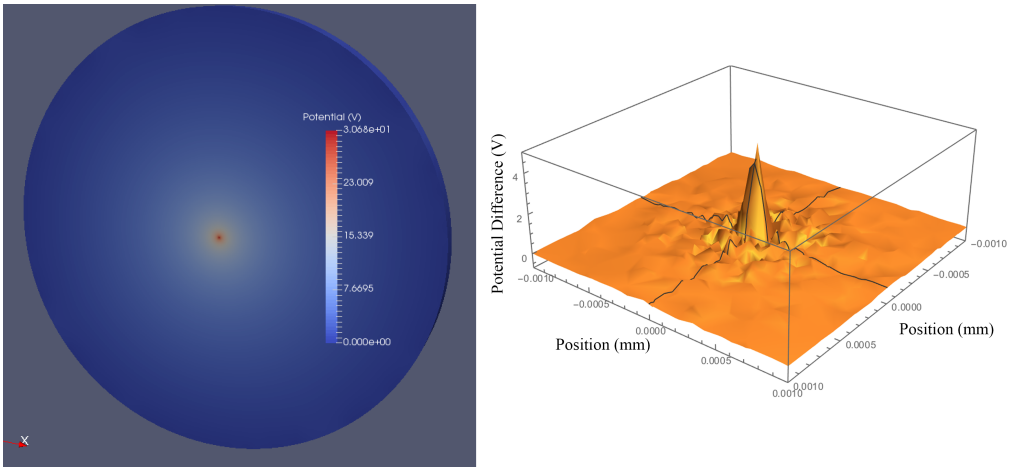

Figure 3.21: On the left, the numerically computed beam potential of a 50 μm beam inside a 2.4 cm beam pipe with a minimum mesh size of 2.5 μm. On the right, the difference between the potential obtained numerically and that found analytically using eqn. (**??**). Several test cases show that the simulation disagrees with the theory by roughly 20 %.

or better. An example of this agreement is shown in Fig. 3.21. More work must be done to determine if this is a fundamental limit of this technique (highly unlikely) or if these results can be improved with minor changes in algorithms and mesh designs. In the meantime, while this code does not achieve perfect agreement, it should generally be sufficient for making clearing electrode design choices in the future.

3.2.2 Longitudinal drift of ions

Once a clearing electrode design has been chosen, one must decide where they should be deployed along the beam line. Electrodes simply cannot be deployed everywhere, due to their negative effects on beam properties, as well as cost and physical space limitations. Ideally, for the sake of efficiency, one should place as few electrodes as possible in the places of highest ion concentration to achieve the maximum amount of clearing. Therefore we have developed a simulation to determine areas of high ion concentration.

We know that the ions oscillate transversely in the potential well of the beam. The depth of this beam potential also changes with beam size variations along the length of the accelerator. The strongest (i.e. deepest) parts of the potential wells occur at beam size minima, and this is where the ions will accumulate. Therefore we believe this is where the ion clearing electrodes should be deployed.

To test this hypothesis, we created a simulation code to track the longitudinal motion of ions through a DC beam. The simulation code GPT [10], which has been used in the past to model the Cornell ERL photoinjector [11], contains a very fast multi-grid Poisson solver ($O(N^{1.1})$) that is excellent for calculating space charge effects between particles. This helps make GPT a perfect tool for tracking individual ions inside of a DC beam.

We first assumed that the DC beam had a round Gaussian transverse charge distribution. The ions are first given an initial set of coordinates (x_0, y_0, z_0) chosen to be inside of the beam and initial velocities (v_{x0}, v_{y0}, v_{z0}) chosen to be the thermal velocities for the sake of simplicity. However, beam size variations along the length of the accelerator ($\sigma_r(z)$) create a longitudinal electric field component E_z that pushes the ions towards beam potential minima. We can derive an expression for E_z by beginning with the Maxwell-Faraday equation

$$\nabla \times \vec{E} = -\frac{\partial \vec{B}}{\partial t} \tag{3.16}$$

For the sake of simplicity, we study the case of a cylindrically symmetric beam. We assume that the electric field has no ϕ dependence, contains no E_ϕ component, has a non-zero longitudinal component E_z, and a non-zero radial electric field component given by

$$E_r(r,z) = \frac{\lambda e}{2\pi\epsilon_0 r}\left[1 - \exp\left(-\frac{r^2}{2\sigma_r(z)^2}\right)\right] \tag{3.17}$$

which follows directly from eqn. (3.13). We can rewrite eqn. (3.16) using the curl operator in cylindrical coordinates

$$\nabla \times \vec{E} = \left(\frac{1}{r}\frac{\partial E_z}{\partial \phi} - \frac{\partial E_\phi}{\partial z}\right)\hat{r} + \left(\frac{\partial E_r}{\partial z} - \frac{\partial E_z}{\partial r}\right)\hat{\phi} + \frac{1}{r}\left(\frac{\partial (rE_\phi)}{\partial r} - \frac{\partial E_r}{\partial \phi}\right)\hat{z} = -\frac{\partial \vec{B}}{\partial t}. \tag{3.18}$$

Because of our assumptions above, the only relevant component of this equation is given by

$$\frac{\partial E_r}{\partial z} - \frac{\partial E_z}{\partial r} = -\frac{\partial B_\phi}{\partial t} = 0 \tag{3.19}$$

If the region inside the beam contains no time varying magnetic flux as is the case in the DC beam approximation, we can neglect $\partial B_\phi / \partial t$. This allows us to obtain E_z from our transverse expression E_r. Taking the derivative of E_r and then integrating it yields the final result

$$E_z(r,z) = \frac{\lambda e}{4\pi\epsilon_0}\frac{\sigma_r\prime(z)}{\sigma_r(z)}\exp\left(-\frac{r^2}{2\sigma_r(z)^2}\right). \tag{3.20}$$

This resulting formula and eqn. (3.17) were used to generate an electric field beam line element, which was then inserted into GPT. Individual H_2 ions were created in regions outside of the SRF cavities and then allowed to drift through the DC beam potential. As predicted, simulations show (in Fig. 3.22) that the ions drift towards and accumulate at beam size minima, making them the best locations for placing clearing electrodes.

In our simulations we ignored ions created inside of the SRF cavities for several reasons. First, the vacuum pressure is so low in SRF cavities (as low as $10^{13} - 10^{14} torr$) that the ion accumulation rate is many orders of magnitude slower than the accumulation

rate in the regular beam pipe. This large disparity would lead to prohibitively long simulation times if one generates ions slowly over time (as opposed to simulating only a set number of ions). But a more important reason for ignoring ions created in the SRF cavities concerns the dynamics of the ion motion inside of the cavities: oscillating electric fields in the RF cavity create a ponderomotive force which pushes the ions towards regions of weak electric fields. This ponderomotive force can be expressed by [39]

$$\vec{F}_p = -\frac{e^2}{4m\omega^2}\nabla\left(E^2\right) \tag{3.21}$$

where e is the electric charge of the trapped particle, m is the trapped particle's mass, ω is the RF frequency and E is the magnitude of the oscillating electric field. This ponderomotive force creates an electric potential well that is proportional to the electric field squared, which will trap any ions that are created inside of the cavities. This effect was simulated using the same GPT code, and is illustrated in Fig. 3.23. It can be seen that ions created inside of the SRF cavities (red) remain trapped within the center of the cavity thanks to this ponderomotive potential (blue line). The ions created outside of the cavities (blue) are free to drift towards beam potential minima.

This GPT simulation is very fast and accurate for small accelerators on the order of meters, but may have prohibitively long run times for larger accelerators. An alternative to this method is to essentially ignore the transverse oscillations of the ions by utilizing the principle of adiabatic invariance [16]. The concept here is to turn a 3D problem (x, v_x, y, v_y, s, v_s) into a quasi-1.5D problem (J, s, v_s) by expressing the transverse ion motion using its action integral J. A particle exhibiting simple harmonic motion moves in a closed circular trajectory in (x, v_x) phase space. The variable J is essentially the area enclosed by an ion's phase space trajectory. The principle of adiabatic invariance states that J remains constant under a slowly varying force âĂŞ in our case, the longitudinal

electric force of the beam that changes slowly with beam size. So a particle's initial (x, v_x) determine J for all subsequent time. For any given J, one can calculate the average longitudinal kick a particle receives during a single transverse oscillation. Thus, a particle's motion becomes only a function of some initial J and the transverse beam size at any given longitudinal coordinate. This eliminates the need to simulate any transverse motion at all when simulating longitudinal ion movement.

We have constructed such a simulation, and it obtains results that are qualitatively similar to our GPT simulation. However, in our case the small size of the injector that is simulated (roughly 14 m) negates the speed advantage this method would otherwise offer for a much larger accelerator. This fact, coupled with the code's lack of support for simulating non-round beams and ion space charge, makes GPT the superior simulation method for our purposes.

3.2.3 Beam shaking

Shortly after performing the experiment outlined in section 3.1.3, we developed a simple simulation code to prove that beam shaking can successfully clear ions out of the center of the beam pipe. This code, developed solely in MATLAB, solved Poisson's equation in 2D using the finite element method. The distmesh package was used for mesh generation [40]. In the simulation, ions were created slowly over time according to the neutralization rate obtained using eqn. (1.1). They were given random initial positions inside of a Gaussian beam, and random initial thermal velocities.

We again invoked a DC Gaussian beam approximation, so that eqn. (3.17) could be used to calculate the electric field generated by the beam. The shaking motion was taken into account by moving the beam centroid sinusoidally back and forth. Maintaining a

boundary condition of zero along the walls of the beam pipe was achieved by using the method of images.

When the beam remained stationary, the ions slowly filled up the beam until they fully screened the beam's electric fields and begin to drift out of the center of the beam. The Poisson solver took into account the space charge repulsion between ions, which accelerated this natural clearing process and allowed ions to escape from the center of the beam before full neutralization occurred.

When the beam was shaking at a random frequency, this same phenomenon occurred. However, while shaking the beam at a particular resonant frequency, the collective ion motion coupled with the motion of the beam, and the ions oscillated back and forth with increasing amplitude until they left the center of the beam (Fig. 3.24. Because the ions undergo non-linear oscillations, eventually the oscillation amplitude became so large that the ion oscillation frequency shifts and the ion motion begins to decouple with the beam. At this point, the initially tight ball of ions blows out and begins to form a large cloud around the beam.

Although by eye it appears that many ions remain trapped within the center of the beam, shaking in fact results in a significant reduction in trapped ion density. Figure 3.25 shows that shaking the beam at resonance can reduce the equilibrium neutralization fraction from roughly 70% to 15%. This reduction agrees well with our experimental findings in section 3.1.3, and confirms that the theory of beam shaking explains our experimental results.

To better understand why ions escape from the center of the beam during shaking, one can examine the individual phase space trajectories of many different ions. Figure 3.26 shows that individual ions start at a small initial amplitude slowly acquire more and

more velocity. The ions follow a swirling pattern in phase space until they eventually escape from the center of the beam and strike the walls of the beam pipe.

Although this simulation agrees qualitatively with our experimental results, we are unable to obtain precise numerical agreement due to a lack of knowledge of our beam sizes during the experiment. In particular, it is not immediately clear from 3.16 what beam size to use, as it is not obvious where ions will primarily accumulate for some of the optics settings. In addition, the changing beam size along the accelerator would mean that the ion resonance frequency shifts as the ions drift longitudinally. This spread of resonant frequencies may help explain the width of the peaks shown in Fig. 3.15. Although this subject may merit further study, this simulation sufficiently achieved our initial goal of proving the effectiveness of beam shaking as an ion clearing mechanism.

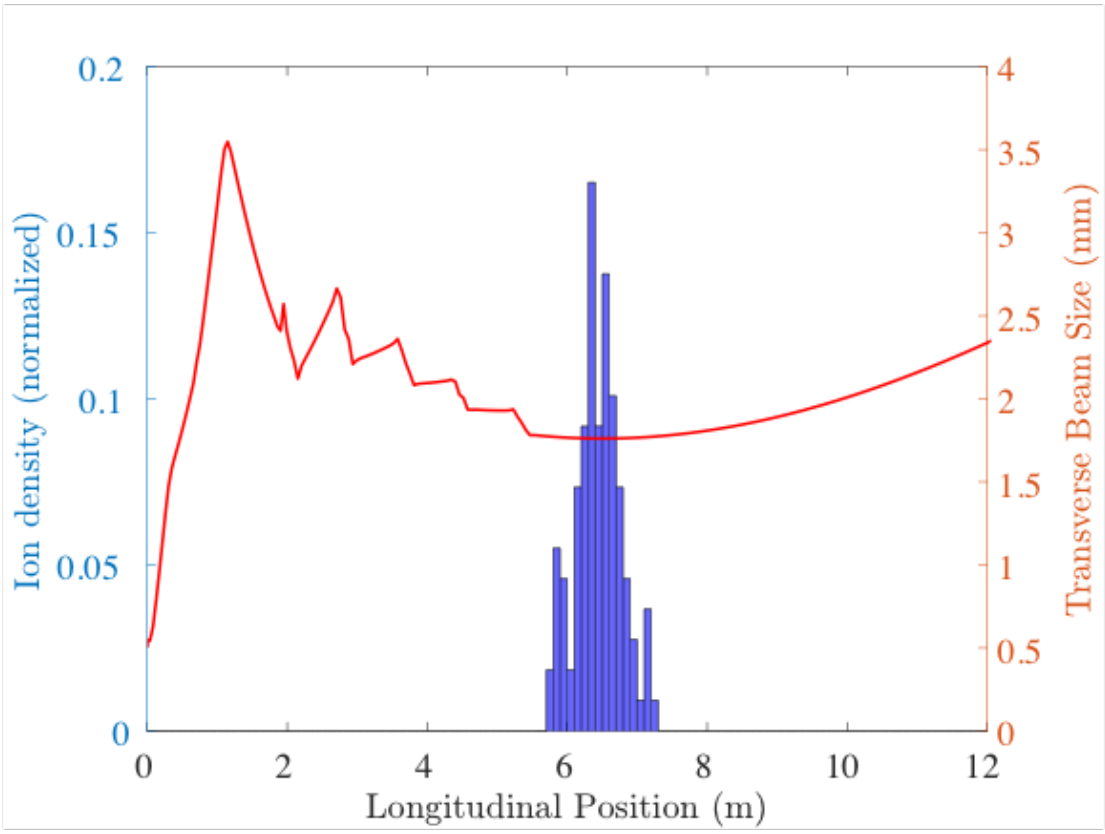

Figure 3.22: Simulated ion density along the Cornell ERL photoinjector. The ions, created between 5.5 m (just after the SRF cavities) and 12 m, drift towards the beam potential, which is directly proportional to beam size.

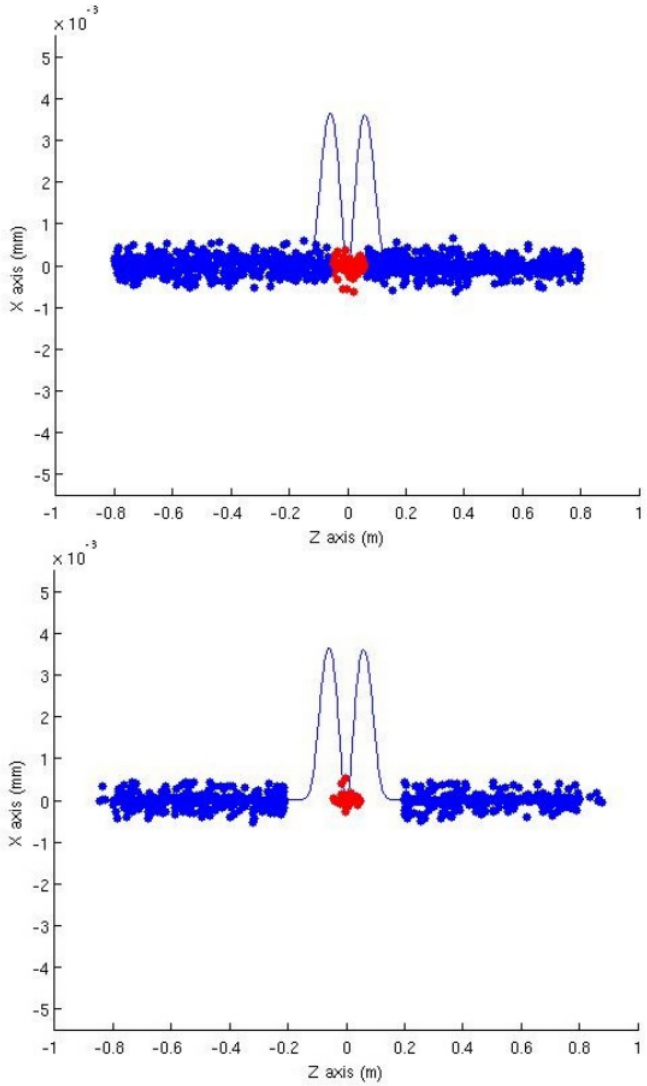

Figure 3.23: Two snapshots of a simulation tracking ions in an SRF cavity. The oscillating electric fields inside of an SRF cavity generate a pondermotive potential (blue line) proportional to the square of the electric field. Ions created in the cavities (red) remain trapped within this potential well, and unlike other ions (blue) do not drift into the remainder of the beam line.

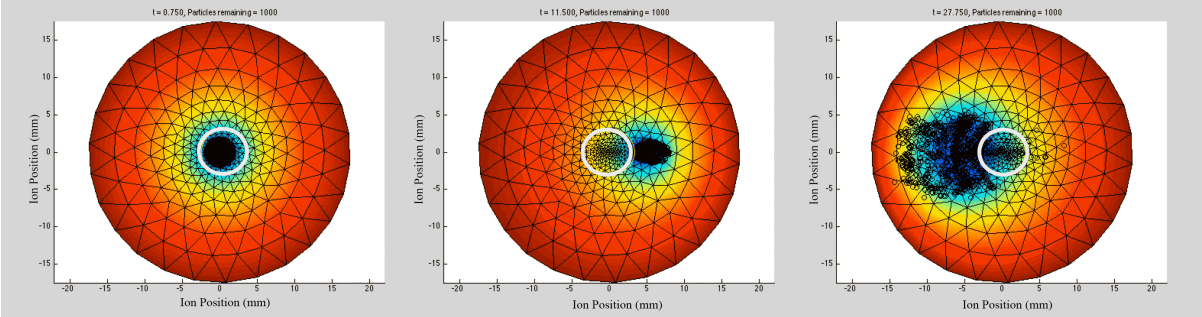

Figure 3.24: Shaking the beam at the ion oscillation frequency induces a resonance that drives the ions out of the center of the beam, as demonstrated using a 2D finite element method Poisson solver. The white circle represents the beam, while the blue dots represent individual ions.

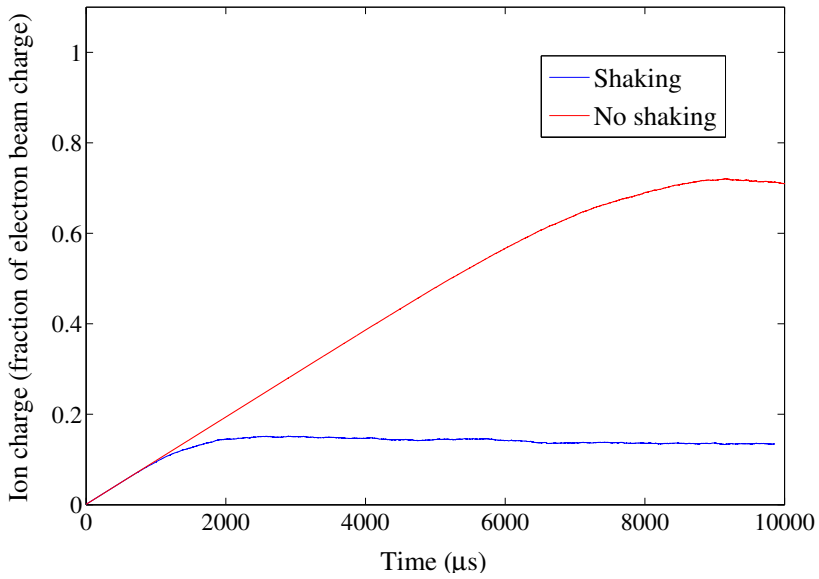

Figure 3.25: The equilibrium ion density achieved under normal conditions and after shaking the beam for 10 ms. In this scenario, shaking the beam reduced the neutralization fraction from roughly 70 % to 15 %.

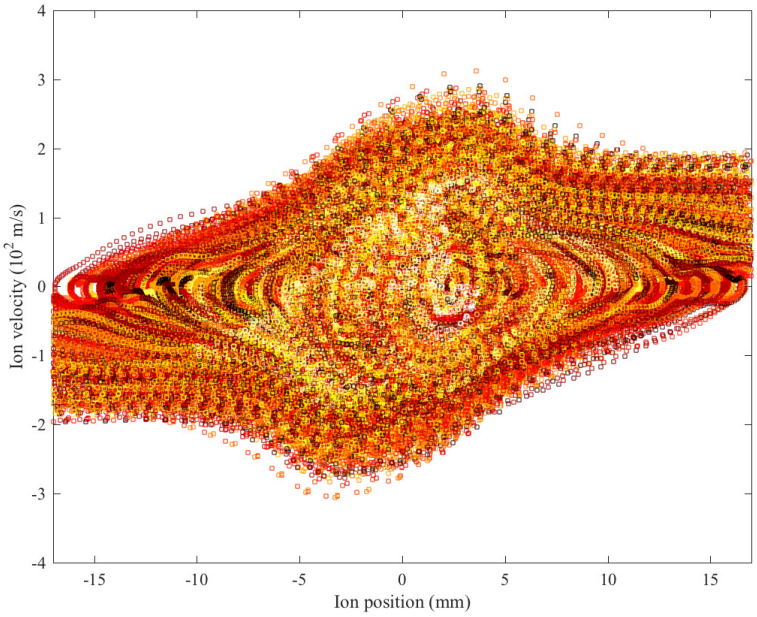

Figure 3.26: Phase space trajectories of many N_2 ions (here created with no thermal velocity) while shaking the beam at the ion resonance frequency.

CHAPTER 4

Beam diagnostics at high intensity

4.1 Diagnostics at high beam power

While we have several straightforward ways of determining whether or not ions are trapped within a beam, it is much harder to measure the beam properties in the parameter regime under which ion trapping occurs. Very few beam diagnostics suitable for linacs can operate at high average beam current (> mA) for several reasons. First, the photoinjector, with beam energies in the range of 5 - 15 MeV operating at 1.3 GHz, creates a very large beam power (on the order of 100 kW) concentrated into a small beam area (rms widths of mm or less). This means that the beam power is high enough to melt any material intercepting the beam within 10 microseconds or less. Second, the beam energy is low enough that synchrotron or diffraction radiation, which is often used to measure beam profiles at higher energies in circular accelerators, is generally not available.

Fortunately there are still several options for taking beam profiles at high current. The first and perhaps most desirable option is to construct a laser wire scanner – a device in which one shines a laser through the beam and looks for Compton scattering radiation to obtain beam profiles [**?**]. These measurements are non-interceptive and do not disrupt beam operation, which is a huge advantage. However their signal to noise ratio can be very poor without a proper setup, and they are difficult to construct and calibrate. Also,

they can have a large footprint, and therefore cannot be deployed in great numbers along the length of an accelerator. A second option is to use a beam-gas luminescence monitor [**?** **?** **?**], which looks at radiation generated by beam-gas collisions in a small pocket of gas in order to obtain beam profiles. However, this device would not allow us to do a controlled experiment studying ion effects because this apparatus would directly interfere with our goal of comparing measurements at high and low vacuum pressures.

A third option is to use a traditional wire scanner [41] that can slice through the beam at such a high speed (> 20 m/s) that a high specific heat capacity wire does not have time to melt during a single measurement. This is the option that we chose to pursue at Cornell. The device uses carbon wires, a material capable of withstanding temperature rises up to 3600 K that is still durable enough to not break during a scan [**?**]. One can determine the minimum required scanning speed to avoid melting by considering the transverse beam size and beam energy, keeping in mind that the diameter of the wire does not determine its temperature rise to first order [**?**]. Experiments found that it is necessary to deposit 10^{10} electrons onto a $3\mu m^2$ area in order for these wires to break [**?**]. Using these numbers and assuming that a carbon wire with a diameter of 34 μm passes through a 100 mA beam, one finds that the wire must travel through the beam at 20 m/s in order to survive a scan [31] – the same speed required by wire scanners at the Large Hadron Collider's (LHC's) interaction point [**?**].

In this chapter, we describe a design for a fast rotating wire scanner that is capable of meeting these speed requirements, while also being cheap enough for mass production. In this design we fix a carbon wire to a blade at one end and rapidly spin it over a small distance using a two planetary gear setup. The smaller gear is attached to a much larger gear, and the smaller gear makes several rotations for every single rotation of the large gear. By accelerating the wire gradually over a much larger path length, this two gear

setup allows the carbon wire to achieve a large speed without breaking. In addition, the two gear design allows the wire scanner to have a low footprint – it fits in a flange with a diameter of approximately 35 cm. A single gear design of comparable size would result in a much more rapid acceleration over a shorter distance, which would likely break the wire. The wire scanner is also made from readily available components (with only two custom parts), which significantly drives down the cost of construction. In the following sections, we present results from tests both inside and outside of the Cornell photoinjector. These tests sought out to ensure that the carbon wire survives scans, and, because it is fixed only at one end, that its vibrations do not significantly impact measurement precision.

4.2 Wire scanner design

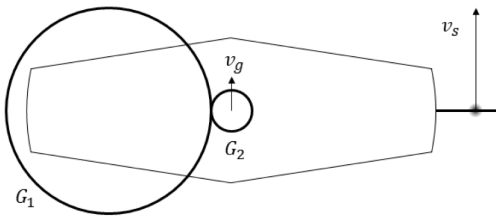

Figure 4.1: A diagram showing the two gear setup of the wire scanner. A smaller gear G_2, traveling at a linear speed v_g, rotates around a much larger stationary gear G_1. This setup gradually accelerates the carbon wire over a large path length, eventually allowing it to slice through the beam with a large velocity v_s [31].

Figure 4.1 shows the two gear design of the wire scanner. In the design, a smaller gear G_2, which holds the blade with the carbon wire attached, rotates around a much larger stationary gear G_1. One can calculate the scanning speed of the wire in terms of the linear speed of the smaller gear v_g [31]

$$v_s = v_g\left(\frac{D}{R_2} + 1\right) \tag{4.1}$$

where D is the distance from the center of the smaller gear G_2 to the beam and R_2 is the smaller gear's radius. However in practice the distance between the smaller gear and the beam is not necessarily the same between successive measurements due to small variations in beam trajectories. This results in some uncertainty in wire scanning speed, which places a restriction on the ultimate precision of the device. One can estimate this uncertainty as

$$\frac{\delta v_s}{v_s} = \frac{\frac{\delta D}{R_2} + 1}{\frac{D}{R_2} + 1} \approx \frac{\delta D}{D} \tag{4.2}$$

where we make the assumption that $\delta D \ll R$ and $\delta D \ll R_2$. Assuming that the beam position is only known to within $\delta D = \pm 5$ mm, and using the distance from the center of the gear to the beam of the present design, $D = 82.4$ mm, we obtain a relative scanning speed variation of $\delta v_s / v_s = 0.061$.

Figure 4.2 shows a photo of the outside of the wire scanner, and Fig. 4.3 shows a rendering of the inside gear assembly. All of the parts for the wire scanner are standard off-the-shelf components, except for the aluminum gear box housing the larger gear and the blade that holds the carbon wire. Two pairs of vacuum flanges enclose the system. The first pair has a 337 mm outer diameter with a 152 mm diameter inner hole, while the second pair has a diameter of 203.2 mm and encloses this inner hole. One of the smaller flanges has three windows which aid in wire alignment and allow one to check the carbon wires for damage between scans. A ferrofluidic rotary feedthrough (Thermionics FRMRE-275-38CL) connects a stepper motor outside of the vacuum chamber to the larger gear while maintaining regular vacuum pressure levels.

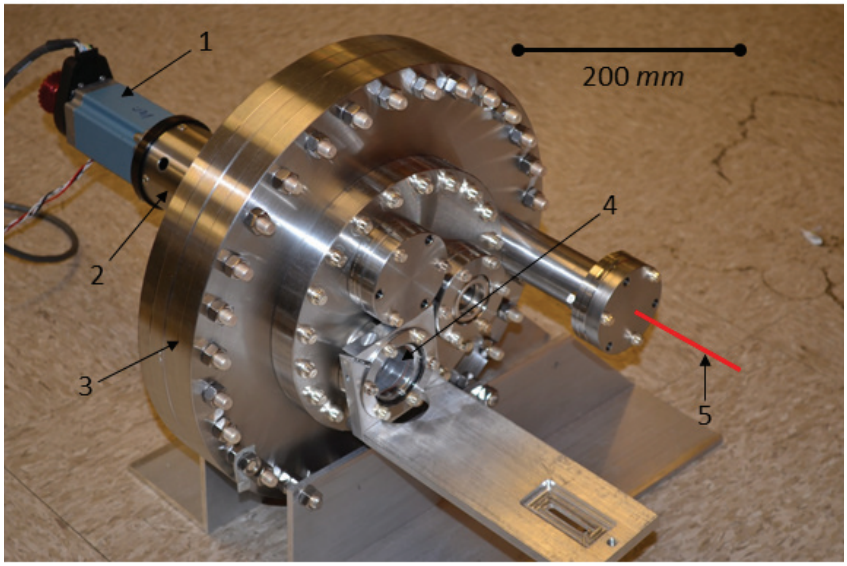

Figure 4.2: A photograph of the outside of the wire scanner. 1: a stepper motor, 2: the ferrofluidic rotary feedthrough for connecting the stepper motor to the gear system, 3: vacuum flanges for housing the entire unit, 4: a vacuum flange with windows for viewing the wire between scans, 5: the path the beam follows through the instrument [31].

Two blades materials were tested: one made from aluminum, and one made from a silica wafer with a metal absorbing layer. The silica blade was designed to avoid the wire absorbing, and possibly being destroyed by, RF fields excited in the wire scanner housing by the electron beam [**?**]. The broadband RF fields could be absorbed by matching the square resistance of the metal absorbing layer to the vacuum impedance [31]. This would prevent the appearance of standing waves in the housing. However, in practice we found that the aluminum blade was generally sufficient for use in the photoinjector, as none of these RF absorption effects were actually detected for beam currents as high as 35 mA when using the aluminum blade.

Three major factors place a limit on the maximum scanning speed of the device: the maximum torque rating of the rotary feedthrough, the friction between the gears, and the moment of inertia of the rotary feed through, gears, wire holding blade, and motor rotor.

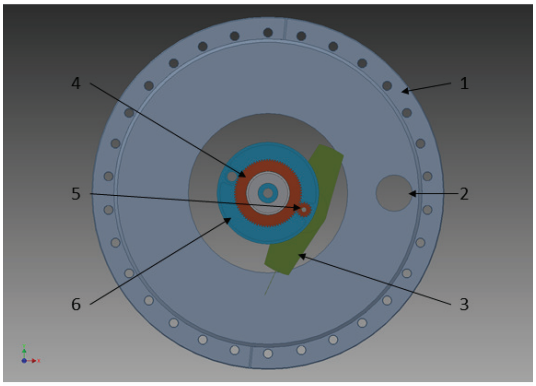

Figure 4.3: A 3D rendering showing the planetary gear setup. 1: vacuum flange enclosing the system, 2: a cut out for the beam pipe, 3: the carbon wire attached to the blade, 4: the larger stationary gear G_1, 5: the smaller rotating gear G_2, 6: the rotating custom-made gear box [31].

Dicronite – a lubricant designed for use in high vacuum environments – was used to lubricate the gear system in order to reduce friction as much as possible. Although the effects of friction were generally small, the gears sometimes locked up and caused the wires to jerk suddenly, which at times resulted in uncertainty in the wire position. Therefore finding methods to further reduce this friction would significantly improve the usability of the device. Holes were also cut into the blade to reduce its moment of inertia, although this may not have been necessary, because the moment of inertia of the system was not a limiting factor in wire scanner operation. With the device's current moment of inertia and a rotary feedthrough with a maximum torque limit of 1.06 N-m, we were able to reliably achieve scanning speeds as high as 30 m/s, which is sufficient for our purposes.

The position and velocity profiles were chosen in order to maximize the path length over which the wire accelerates and make the acceleration experienced by the wire as gradual as possible. The wire started just outside of the beam pipe at the beginning of the scan and made a full rotation around the larger gear in order to reach maximum speed before reaching the beam. After slicing through the beam at maximum speed, the wire

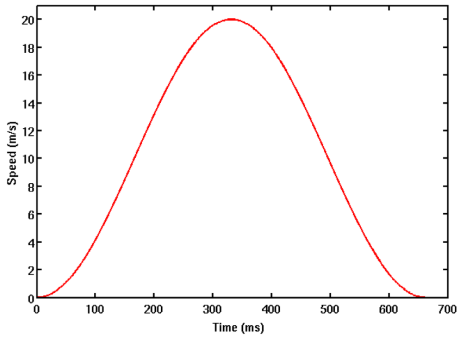

Figure 4.4: The stepper motor was programmed to follow a smooth velocity profile during a scan to avoid any abrupt accelerations which may break the carbon wire [31].

gradually slowed down over one full revolution of the larger gear until it came to a complete stop just on the other edge of the beam pipe. After this measurement cycle was completed, the wire was rotated backwards into its original starting position before another scan was initiated. One could theoretically begin another scan in the opposite direction by starting at the end position, however it was found that it was better to begin each measurement from the same location in order to achieve consistent wire speeds in the center of the beam pipe and thereby maximize measurement consistency.

Figure 4.4 shows the smooth sinusoidal velocity profile programmed into the stepper motor. The wire was made to reach its maximum velocity in the center of the beam. The wire completes a scan in less than 700 ms, and it takes only 2 ms for the wire to travel through the beam pipe. Even though the wire follows a sinusoidal curve, the velocity changes only by 0.0005 m/s over this time interval from one end of the beam pipe to the other, making it unnecessary to program a flat segment into the center of the velocity profile.

4.3 Bench tests

Wire scanner designs typically hold a wire by both ends in order to reduce wire vibrations [41**?**], and scanners designed in this way rarely exceed vibration amplitudes of more than 10 μm [**?**]. However, our design makes it necessary to fix the carbon wire to the blade by only one end. This can cause the wire to vibrate with amplitudes of several mm, which can have a significant impact on measurement precision by virtue of changing the maximum wire speed as it passes through the beam. To first order, we can model the vibrations of the wire as simple harmonic motion, with amplitudes given by $A(x) = A_0(x)\sin(\omega t + \phi)$ and speeds $v(x) = \omega A_0(x)\cos(\omega t + \phi)$. We were able to obtain an estimate of the wire vibration amplitude and frequency by observing the vibrations of the wire with a CCD camera. We measured the amplitude simply by inspection. To obtain the frequency, we shined a pulsed laser onto the vibrating wire and tuned the laser pulse frequency until the image of the wire became stationary. With a wire oscillation frequency of $\omega = 75 \pm 2$ Hz and maximum vibration amplitudes of $A_0 = 5$ mm, we calculated that wire vibrations can cause the wire scanning speed to vary by as much as $A_0\omega = 0.4$ m/s during a scan.

These vibrations place a limit on the maximum precision achievable by the device and, as a result, it is very important to determine whether or not this estimate is accurate. Therefore we conducted several bench tests in order to measure the actual vibration amplitudes of the wire. Our goal was to obtain images of the wire passing through the center of the beam pipe at high speeds. In theory this could be easily achieved by using a high speed camera, however none were available at the time of this experiment.

Instead we performed the experiment shown in Fig. 4.5, which entailed using a pulsed laser and a series of Fourier transform lenses to obtain multiple images of the wire

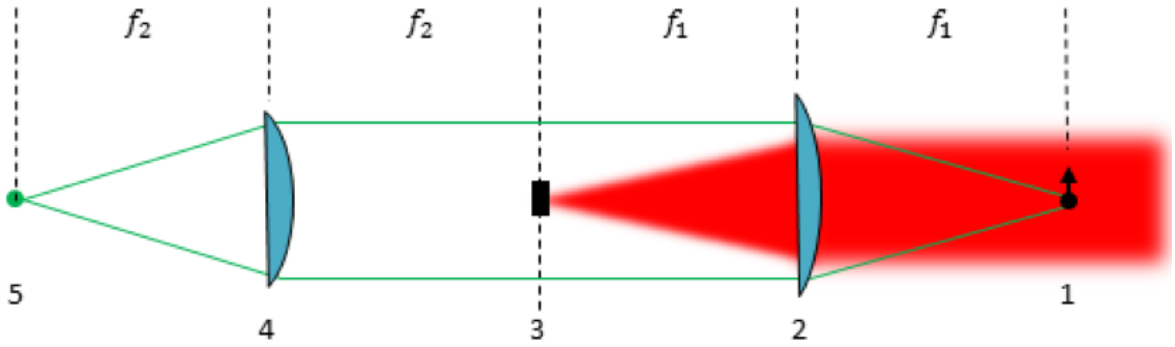

Figure 4.5: A schematic showing the setup used for the wire scanner bench tests. The goal was to take images of the wires in motion by using a modulating laser to obtain multiple wire images on a single CCD frame. 1: the moving wire is illuminated by a pulsed diode laser, 2: light from the laser passes through the first Fourier transform lens which collimates light scattered by the wire and focuses the remaining laser light onto an obstruction, 3: a spatial filter removes the zeroth order laser light focused by the first lens while allowing the scattered light to pass through uninhibited, 4: the light passes through a second Fourier transform lens which focuses light scattered by the wire onto a CCD camera, 5: a CCD camera is used to obtain an image of the moving wire [31].

on a single CCD frame with a long exposure time. The wire scanner was held fixed at a vacuum pressure of 10^{-7} torr in order to replicate actual operating conditions in the photoinjector. The pulsed laser was shined through a window onto the wire as it passed through the center of the beam pipe. Lens 2 focused the majority of the zeroth order laser light onto a strip of black electrical tape in order to block it out. The laser light that scattered off the wire was collimated by lens 2 and therefore not blocked by this tape. Lens 4 collected this collimated light and focused it onto a CCD camera, which was synchronized to take a long exposure snapshot just as the beam passed through the center of the pipe.

Figure 4.6 shows an image of a stationary wire taken with this setup, and Fig. 4.7 shows multiple images of a moving wire taken on a single CCD frame using a laser modulation frequency of 8 kHz. Background subtraction was used to remove any light scattered by the flange windows. In the moving image, the wire vibrations slightly

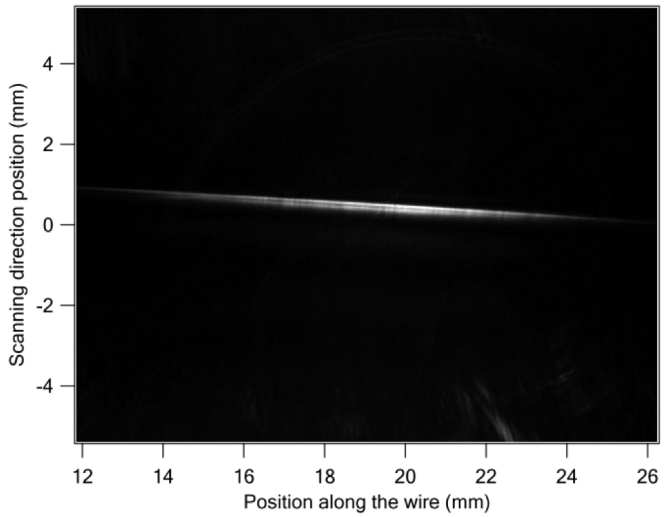

Figure 4.6: An image of a stationary wire attached to the blade of the wire scanner. The wire bends downward slightly due to the force of gravity. The image was taken using the bench test setup shown in Fig. 4.5 [31].

broaden the image, and do not appear as pronounced sinusoidal vibrations as one might expect.

We carefully analyzed this image by taking several vertical image profiles along the length of the wire, obtained by averaging over 50 pixels (0.37 mm). This analysis is shown in Fig. 4.8. Gaussian profiles were fit to the data in order to determine the exact locations of the peaks. We used the distance between adjacent peaks, the known magnification of our optical system, and knowledge of the laser pulse frequency to calculate the wire velocity at different points along the wire. We found that the wire speed varied slightly along the length of the wire, suggesting the presence of wire vibrations. The average speed along the length of the wire for both measurements were 20.2 ± 0.2 m/s between peaks 1 and 2, and 20.0 ± 0.3 m/s between peaks 2 and 3. The uncertainty in these measurements is consistent with our estimates (± 0.4 m/s), suggesting that our estimates were not far off from reality.

Increasing the modulation rate of the laser to 25 kHz reduced the broadening of the

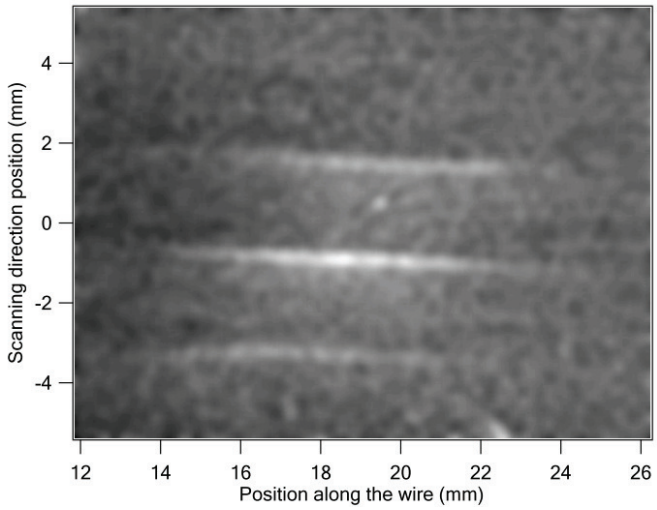

Figure 4.7: Several images of a single wire moving upward at 20 m/s, obtained by using a laser pulsed at 8 KHz to capture multiple images of the wire on a single long exposure CCD frame [31].

wire images and allowed us to obtain up to 5 beam profiles in a single image. The experiment was repeated using a scanning speed of 18 m/s, and our data was analyzed in the same way as Fig. 4.9, except this time we only analyzed the speed at the location where the wire intersects the center of the beam pipe. Results for this test are shown in Fig. 4.10. The average speed was found to be 18.1 ± 0.3 m/s, which is consistent with both our estimates and the results obtained using the 8 kHz modulating laser. Therefore from all of our data we can conclude that the uncertainty in the wire speed due to vibrations is approximately 0.3 m/s / 20 m/s = 0.015, or 1.5%. As will be shown later, this is likely a lower bound, and tests in the actual photoinjector show that the measurement uncertainty is closer to 5-10%.

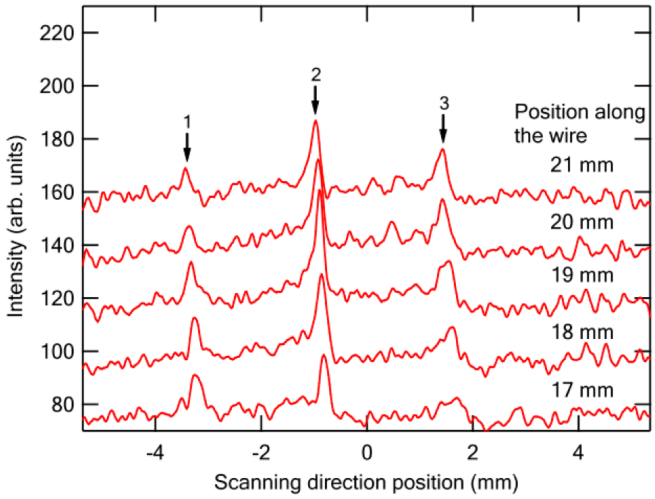

Figure 4.8: Vertical slices of Fig. 4.7 taken at different locations along the length of the wire, obtained by averaging over 50 pixels (0.37 mm). The profiles share the same level of background intensity, and are shifted vertically in order to place them on the same plot for the sake of convenience. Peaks 1, 2 and 3 were used to determine the speed of the wire by calculating the distance between adjacent wire images [31].

4.4 Tests at high beam current

After performing the bench tests, we tested the wire scanner in a chicane in the Cornell ERL photoinjector [**?**]. Beam profiles were obtained by measuring X-rays emitted when the wire sliced through the beam. The X-rays were detected using a scintillator crystal and silicon photomultiplier sensor (MicroSM-60035-X13 by SensL) combination. A preamplifier (Micro-EVB) and standard power supply from the same manufacturer were used for signal conditioning and to power the silicon photomultiplier sensor [31]. Because very high beam currents in the photoinjector produced large amounts of X-rays, the signal measured by the photomultiplier frequently became saturated. In order to combat this saturation we installed a neutral density filter (with attenuation ≈ 30) between the scintillator crystal and photomultiplier detector to reduce the number of photons (produced by the scintillator) that reached the detector. This allowed us to measure beam profiles for beam currents between 1 - 35 mA with very good signal-to-noise ratio.

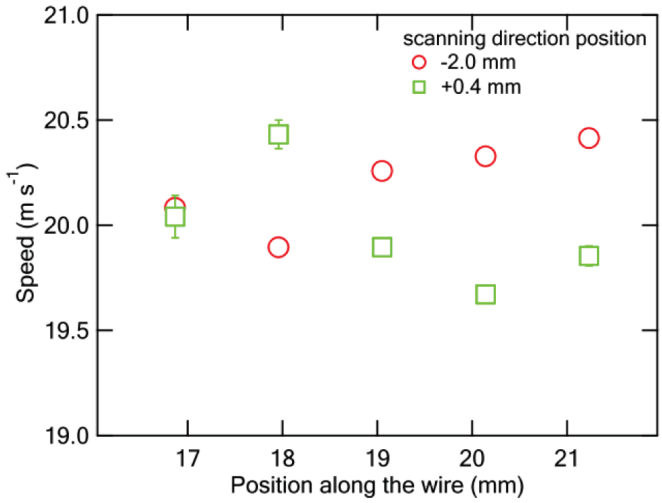

Figure 4.9: Wire speeds variations along the length of the wire, calculated using Fig. 4.8. The circles are values calculated using the distance between peaks 1 and 2, while the squares used peaks 2 and 3. Many of the data points have error bars smaller than the symbol size [31].

A digital acquisition device (Agilent U2500A USB DAQ) sampled the signal from the photomultiplier sensor at a rate of 2 $\times 10^6$ samples per second. Assuming a scanning speed of 20 m/s, this means that data points were sampled every 10 μm, which is sufficient for obtaining beam profiles with a wire of diameter 34 μm. The stepper motor controller (DMC-2183 by Galil Motion Control) was capable of starting data acquisition right before the wire crossed the center of the beam. Data could also be obtained by connecting the photomultiplier sensor directly to a common oscilloscope, and triggering the start of the data acquisition using the signal itself. This produced results that were consistent with those obtained using the DAQ, suggesting that this simpler setup was appropriate to use in a pinch.

We observed that between measurements the wire violently vibrated, because vibrations were excited when the wire abruptly stopped after it was moved to its starting position. The vibrations decayed slowly over time, but it could take several minutes to do so. By observing the wire using a video camera, we estimated that the wire vibrations

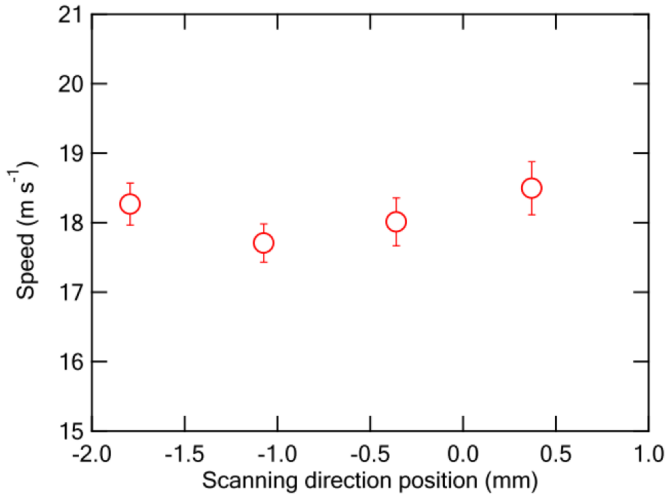

Figure 4.10: Wire speeds variations along the length of the wire when using a 25 kHz laser, calculated in the same way as Fig. 4.9. The higher modulation frequency allowed us to obtain 5 profiles in a single CCD image, as opposed to the 3 obtained when using a modulation frequency of 8 kHz [31].

had a damping time constant of $\tau = 60 \pm 5$ s. It was feared that these vibrations would have a significant impact on the precision of our measurements, which would ultimately force us to wait for the wire vibrations to completely stop between measurements. This would significantly increase the time between measurements – something we would like to avoid if possible. So our first goal was to test this behavior and see whether or not it was truly necessary to wait for these vibrations to stop between measurements.

We took two sets of data at 4 MeV with a 5 mA beam at a fixed repetition rate of 1.3 GHz. During the first set of measurements, we waited between scans approximately 180 s for the vibrations to decrease to a sufficiently small amplitude. We then used a smoothing spline fit to determine the full width at half maximum (FWHM) for each measured beam profile. In Fig. 4.11 we compared this to the case where we only waited for 20 s between scans (the fastest scan time possible), so that the wire did not have time to significantly stop vibrating. Only a small difference in beam profile reproducibility was found in both cases. The average FWHM for both data sets was found to be 2.79 ± 0.16

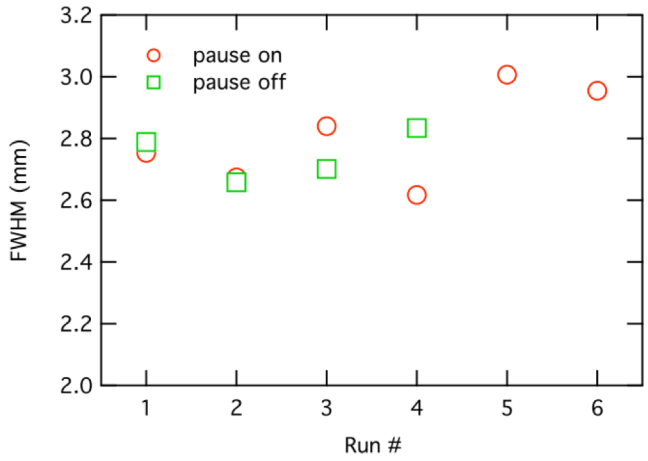

Figure 4.11: A comparison of the full width at half maximum of many beam profiles obtained with and without waiting times between measurements. Circles: Data obtained with a pause of 180 s between consecutive scans – a sufficient amount of time to let wire vibrations decrease to negligible amplitudes. Squares: Data taken with only a 20 s pause, so that the wire did not have time to stop vibrating between scans. Error bars are smaller than the symbol size for all data points [31].

mm, with a relative precision of 5.7%, which is slightly larger than the wire speed variations caused by vibrations during the bench tests.

There is one probable reason for this discrepancy. During the tests in the photoinjector it was feared that the wire might accidentally slip into the beam at a low velocity between scans, so the beam was turned off between measurements in order to avoid this possibility. The beam current in the photoinjector is increased by changing bunch charge at a fixed repetition rate, and may not return to the same exact value each time it is turned on. This could directly cause the slight variations in the width of beam profiles that were measured with the wire scanner. Because this is our only beam diagnostic capable of taking beam profile measurements at high beam current, we are unable to definitively determine if this is the case.

We also set out to test the dynamic range of the photomultiplier sensor for beam currents between 250 μA and 35 mA. Figure 4.12 shows this data, with several beam

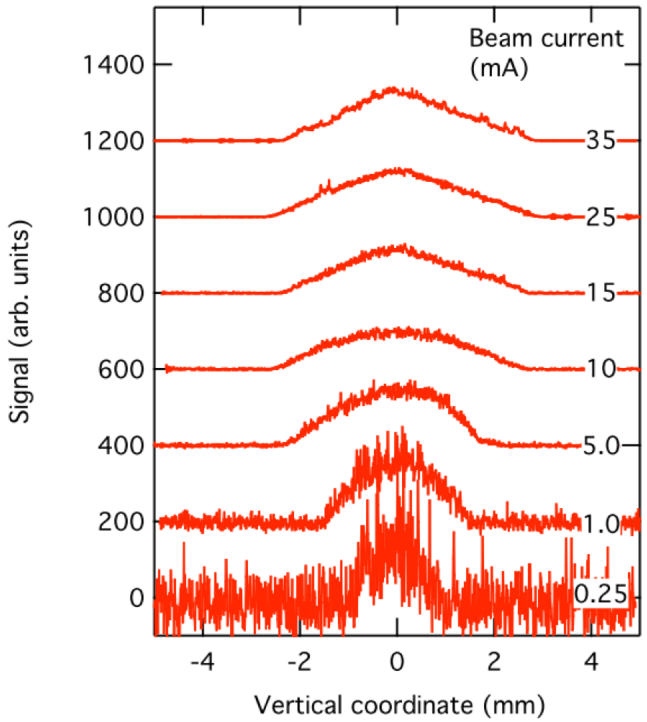

Figure 4.12: Several vertical beam profiles taken at beam currents between 0.25 and 35 mA. The profiles are stacked on top of one another on the same plot for the sake of comparison, and the amplitudes of the profiles are normalized to the same peak height [31].

profiles stacked on top of one another for the sake of clarity. The beam profiles were normalized by dividing the amplitude of the signal by the value of the beam current. The linearity of the system is quite good, as evidenced by the area under the curve remaining constant between measurements. As beam current is increased by increasing bunch charge, the profiles begin to take on a distinct triangular shape. This is believed to be a direct result of space charge repulsion – one of the dominant effects in the low energy photoinjector.

Beam profiles were also obtained by performing scans in opposite directions to confirm that the asymmetry measured in the beam was actually a property of the beam at 35 mA, and not an unwelcome artifact of the measurement process. Figure 4.13 confirms

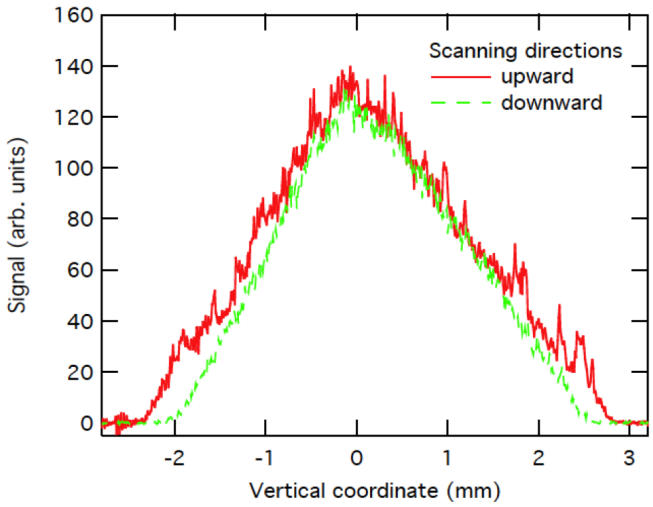

Figure 4.13: Two beam profiles obtained by scanning the wire in opposite directions in order to check that the asymmetry of the profiles was a property of the beam itself and not due to the measurement technique. Solid: Upward scan, Dashed: Downward scan [31].

that the profile obtained in both scanning directions is nearly identical, and the beam is actually asymmetric.

Finally, we can compare beam profiles obtained with the wire scanner to those obtained by a view screen at low beam current. Fig. 4.14 shows such a comparison for a beam current of 20 μA (0.015 pC at 1.3 GHz). Because the beam current and therefore X-ray signal was so low, a regular photomultiplier tube with a high dynamic range was used to detect the beam profile. The data for both curves agrees very closely, confirming that the wire scanner obtains accurate beam profiles at low current.

4.5 Tests with more durable carbon wires

The aforementioned tests were performed with very weak, brittle carbon wires originally baked at Cornell. These wires were so brittle that they would frequently break while

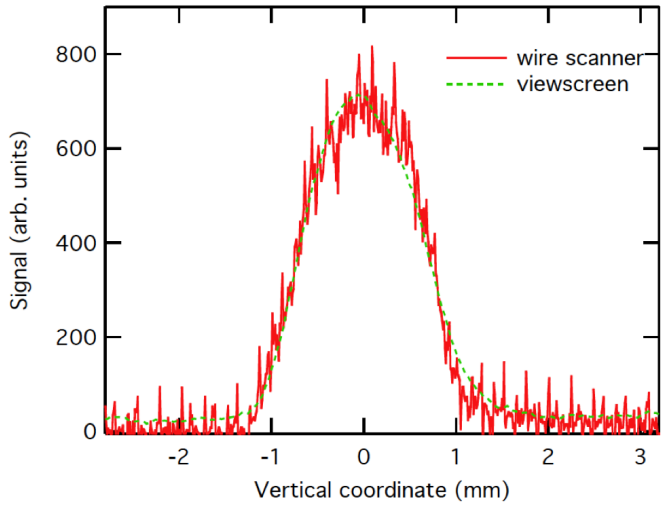

Figure 4.14: Comparison of the beam profile obtained with the wire scanner (solid line) to that obtained by a luminescent view screen (dashed line) for an average beam current of 20 μA (0.015 pC bunch charge) [31].

attaching them to the wire scanner blade using a pair of tweezers. Whether or not a wire would survive multiple spins was more or less determined by chance, forcing us to frequently open up the beam vacuum chamber to replace missing wires as necessary. This resulted in significant machine downtime. As a result, we set out to find stronger carbon wires to improve the durability of the device and hopefully allow for continuous operation for several weeks at a time.

We eventually chose to use 30 μm carbon wires purchased from [**?**]. These wires were very robust - so robust that they could be bent greatly without breaking, and could be handled rather violently without snapping into pieces. However these new wires were significantly less rigid than the old ones, and it was feared that new excessive vibrations would decrease the precision of the beam profile measurements. As a result, we performed beam tests in the photoinjector again, but this time with a stronger emphasis on measurement precision and reliability over many scans.

For this experiment we attached three wires to a single blade, allowing us to obtain

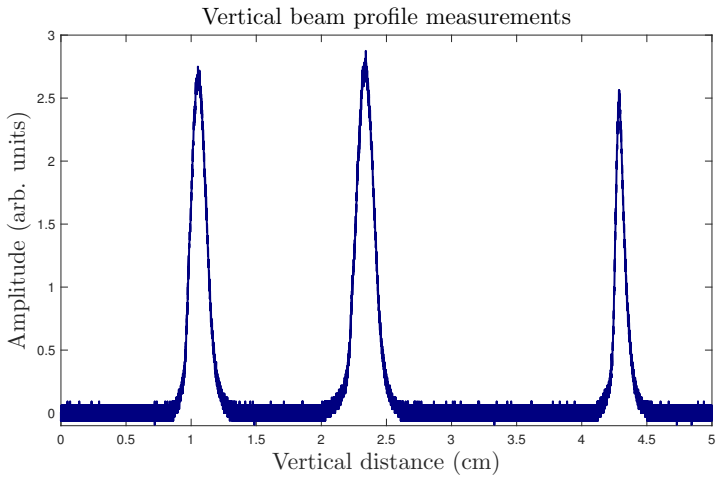

Figure 4.15: A vertical beam profile measurement taken with a scanning speed of 10 m/s. Three wires were attached to the wire scanner blade in order to obtain three beam profiles during a single slice through the beam.

three beam profile measurements during a single scan. An example measurement is shown in Fig. 4.15. This setup also allowed us to determine the amplitude of wire vibrations by measuring the distance between adjacent beam profiles.

Unfortunately, variations in wire speed between scans lead to inconsistent profile measurements, as shown in the top graphs in Figs. 4.16 and 4.17. These variations in scanning speed could be adjusted for by properly scaling the widths and amplitudes of the beam profiles. Specifically, the data was transformed by using $A(y) \rightarrow A(cy)/c$, where c is a constant denoting the mean of the rms widths divided by the rms width for that particular beam profile, such that $c = \langle \sigma_y \rangle / \sigma_y$. Making this adjustment allowed almost all of the beam profiles to overlap very nicely, as shown in the bottom graphs in Figs. 4.16 and 4.17. The higher scanning speed resulted in more consistent measurements that required less adjustments, suggesting that higher speeds are preferable for taking measurements.

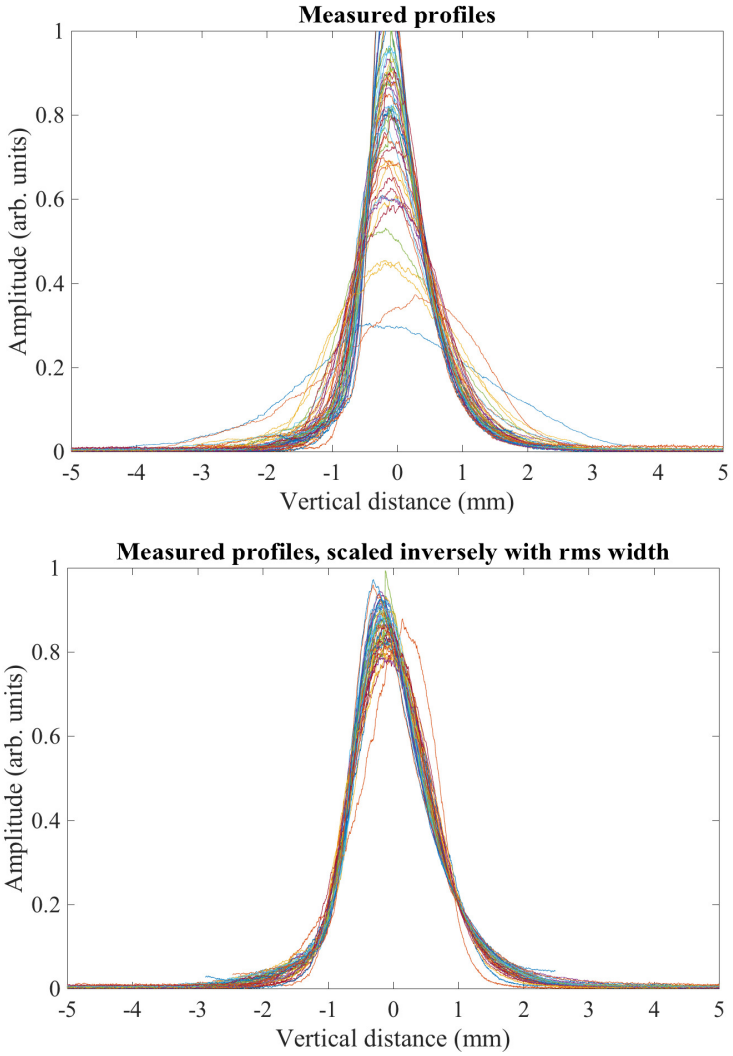

Figure 4.16: A collection of vertical beam profile measurements taken at a scanning speed of 10 m/s. After scaling the width and amplitudes of the raw beam profile measurements (top) to adjust for possible variations in wire speed, all of the beam profiles closely overlapped (bottom). Similar measurements in Fig. 4.17 show that higher scanning speeds lead to much better reproducibility.

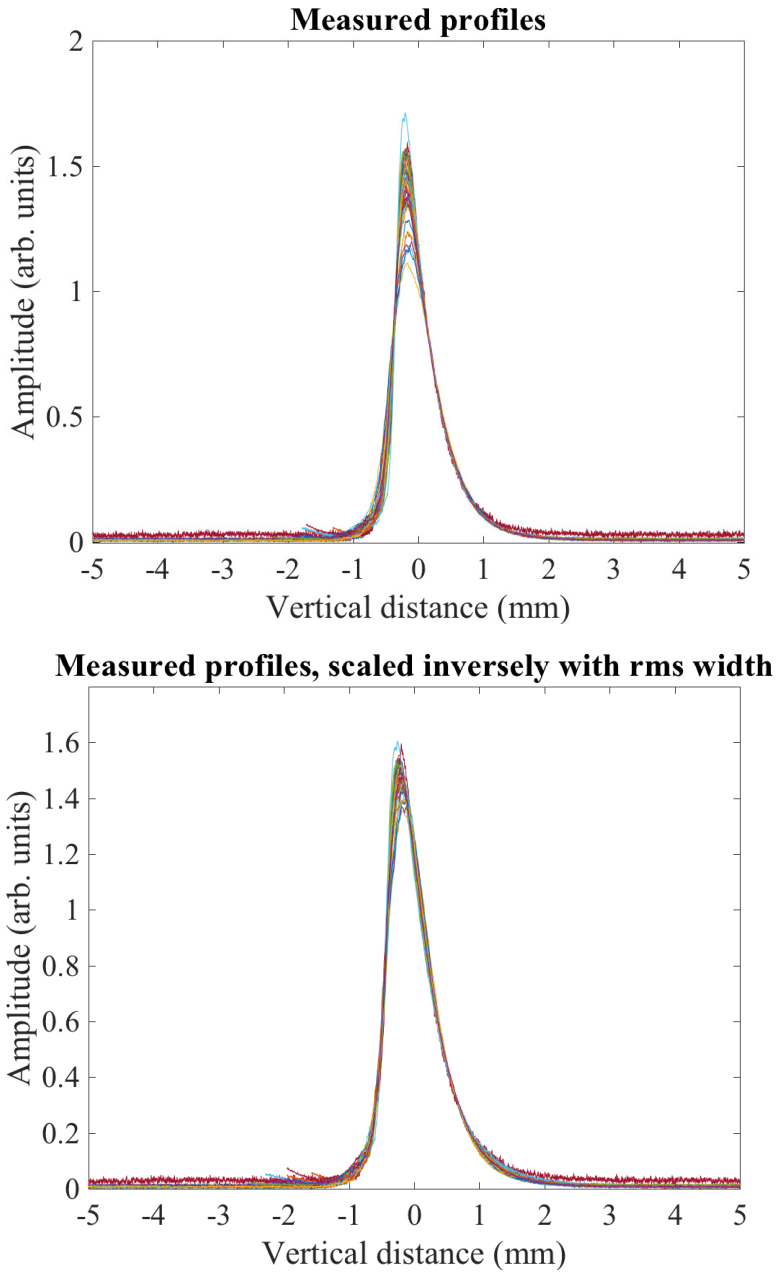

Figure 4.17: A collection of vertical beam profile measurements taken at a scanning speed of 20 m/s. After scaling the width and amplitudes of the raw beam profile measurements (top) to adjust for possible variations in wire speed, all of the beam profiles closely overlapped (bottom). Even without adjustments, these measurements where much more consistent than those obtained in Fig. 4.16 at a velocity of 10 m/s.

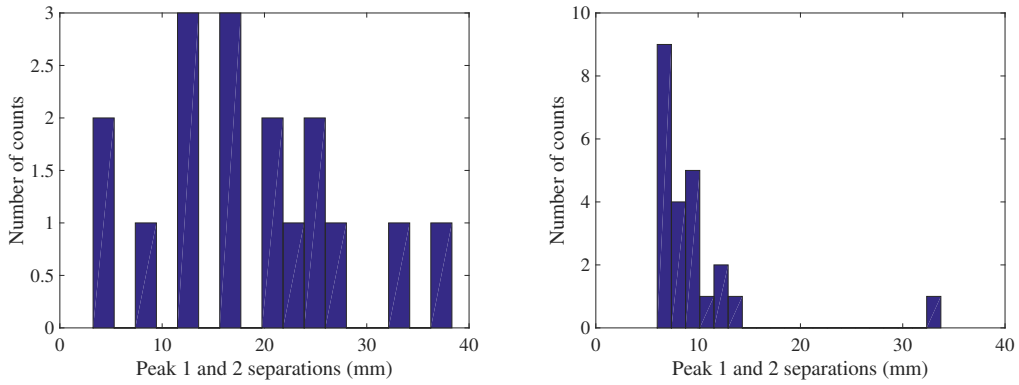

Figure 4.18: Histograms showing the separation between beam profiles obtained by adjacent wires for two different scanning speeds. Left: 10 m/s scanning speed, Right: 20 m/s scanning speed. High scanning speeds exhibit significantly smaller wire separations on average, almost by a factor of 4, suggesting that wire vibrations are less pronounced at higher velocities.

This measurement consistency was a direct result of smaller wire vibration amplitudes at higher scanning speeds. As can be seen in Fig. 4.18, at 10 m/s adjacent wires were separated by 10's of mm, whereas at 20 m/s they were separated only by 3-7 mm. These smaller amplitudes would cause less uncertainty in the speed of the wire as it passes through the center of the beam, leading to more precise measurements. Speed modulations at 10 m/s due to wire vibrations may also explain why some of the beam profiles in Fig. 4.16 are unusually broad.

We also compared the rms beam profile widths for two different scanning speeds, as shown in Fig. 4.19. The higher 20 m/s scanning speed resulted in significantly less measurement uncertainty - only 5.2%, compared to 16% at 10 m/s. This finding makes sense and is consistent with our previous results. Namely, smaller wire vibration amplitudes lead to more consistent beam profile measurements, which in turn leads to more consistent rms beam profile widths.

The general conclusion we can draw from this data is that one should operate at 20

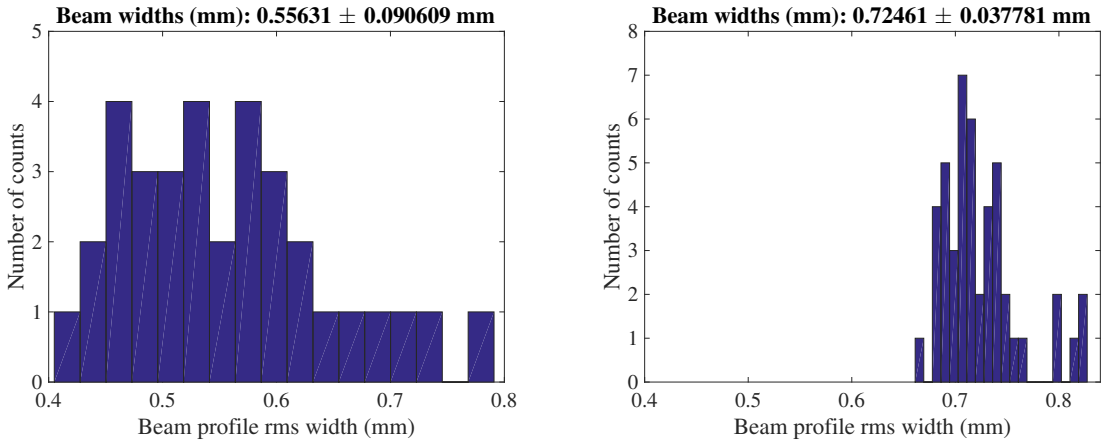

Figure 4.19: Histograms showing many rms beam profile widths obtained for two different scanning speeds. Left: 10 m/s scanning speed, Right: 20 m/s scanning speed. The 10 m/s scans had a relative uncertainty of 16%, while the 20 m/s runs had a relative uncertainty of 5.2%, suggesting that measurement precision improves at higher velocities. This is likely the result of decreased vibration amplitudes, as shown in Fig. 4.18.

m/s to minimize wire vibrations and decrease measurement uncertainties. However there are still some durability issues at high scanning speeds, and during the 20 m/s testing one of the three wires broke. It appears it broke because the wire was not rotating perfectly within the plane of rotation. Instead, it was slightly bent outwards out of the plane of rotation, which allowed it to get caught on the edge of the flange and break. Fortunately, this problem has a relatively straightforward solution. More carefully attaching the wire to the blade by taking special care to keep it completely flat within the plane of rotation would likely remedy this problem in the future.

4.6 Conclusions

Several improvements can be made to the wire scanner design in order to increase the precision and reliability of the device. First, as many moving parts as possible must be

removed from the gear and stepper motor setup in order to lessen the amount of locking up of the gears and slippage between measurements. The gear system must also be very well lubricated to minimize the effects of mechanical friction. Changing the carbon wire material or modifying the way it is attached to the blade can also prevent breakage in the future. This may lead to longer periods of continuous deployment without the need for constant wire maintenance – one of the biggest flaws of the device. If these changes are made, then perhaps the device will be ready for deployment on a massive scale in a large high current accelerator.

Ultimately, although this wire scanner design has several issues, it proves to be a useful device for taking beam profile measurements at high beam current, and is one of the few windows we currently have for studying ion effects at high beam current. In the future it is hoped that this device will be used to conduct experiments by leaking gas into the photoinjector and comparing beam profiles with and without excess gas. Comparing the beams with identical bunch charge at low and high repetition rates will also allow us to observe the effect of ions on the beam, as suggested in section 2.3.3. Then, finally, we can begin to put the ion issue to rest and determine whether or not the effects are severe enough to warrant the use of mitigation strategies.

CHAPTER 5

Conclusions and future work

Ion trapping will most certainly occur in future high current, high intensity linacs. In fact it has already been observed in the Cornell photoinjector, where trapped ions led directly to machine trips at beam currents exceeding 10 mA (section 1.3). Whether or not trapping occurs in future accelerators can be determined by examining the trapping conditions outlined in section 1.2, which depend largely on the maximum achievable beam current and the spacing between bunches. CBETA and the Berlin energy recover linac prototype (bERLinPro) [19; 42] are the next two linear accelerators most likely to observe ion trapping, although there are sure to be many more in the future.

It is clear that ions can have a significant and sometimes detrimental effect on the beam, but the precise impacts of ions must be determined via theory and simulations on a case by case basis for future accelerators. Several effects such as tune spreads or sidebands caused by beam-ion coupling may be less than ideal, but they can often be only a minor inconvenience and might have little impact on stable machine operation. Others, such as the fast ion instability, must be avoided to achieve any stable operation at all. The only definitive effect found in all ion dominated accelerators is that trapped ions create non-linear focusing forces leading to changes in beam optics, which can often cause emittance growth. This focusing cannot be easily fixed using linear optics systems, which necessitates the use of ion clearing strategies if the emittance growth is intolerable for that machine's design.

Fortunately there are many ion mitigation options which we have tested extensively in chapter 3. All three methods that we have tested – ion clearing electrodes, bunch gaps and beam shaking – resulted in significant reductions in the trapped ion density. A simple formula predicts clearing electrode voltages required for maximum clearing, which can be anywhere from 1-5000 V depending on the beam density, a fact that we confirmed experimentally. This formula can be used in the design of clearing electrodes. Low voltages are generally not a problem, as button BPMs can easily be biased for the task of clearing ions, but voltages exceeding 1 kV would require specialty clearing electrodes. The lone ion clearing electrode in the photoinjector was actually effective enough to maintain stable beam operation when ion problems were encountered during 350 keV runs, but much larger accelerators would require multiple clearing electrodes strewn along the length of the accelerator in order to achieve similar results. In that case, it becomes important to locate pockets of high ion concentration (generally beam size minima) via simulations to deploy clearing electrodes in the most effective manner.

Experiments suggested that the effectiveness of bunch gap clearing largely depends on the total time the beam is absent, and does not necessarily depend on the exact bunch gap durations and frequencies used. In the future more work must be done to understand why this result was found. Perhaps the data was taken in an intermediate regime between the two extremes where ion trapping always occurs (at small bunch gap lengths) and ion trapping never occurs (at large bunch gap lengths). Examining a wider range of bunch gap parameters experimentally would shine a light on this result, and would make for an interesting set of experiments in the future.

The results from the beam shaking experiments agreed excellently with scaling laws predicted by the ion oscillation frequency. However the exact beam size during the experiment still remains a mystery, and so assumptions using GPT simulations must be

made in order to achieve full agreement between the experimental results and theoretical predictions. Using the new wire scanner to measure rms beam size while repeating this experiment would eliminate these assumptions and resolve this issue definitively. Additionally, simulations and experimental testing of white noise beam shaking – a technique employed daily at the MLS – may be warranted in the photoinjector to see if this technique also results in significant reductions in trapped ion density. If white noise shaking works, this would allow one to clear a wide range of ion species at multiple locations along the accelerator, without the need for prioritizing certain ion species.

Overall, there are many directions that the study of ion trapping can move towards in the realm of theory, simulation and experiments. In the realm of theory, further explorations of the variable-length bunch gap theory at the end of section 1.2.3 would be very interesting for future ERLs with strict requirements on bunch filling patterns. This is something we have started to examine at Cornell, and explorations of this have also already begun at other labs [43]. Initial results at Cornell show that this theory creates a very complex, potentially non-linear parameter space with much room for interesting discoveries.

The "holy grail" method for determining the effectiveness of clearing methods would be to make a fully self-consistent simulation code, and track ions accumulation inside of an accelerator while tracking a beam through the ion density in the same code. This could not easily be achieved due to large variations in time scales and huge simulation regions, but an iterative approach could be used where the ion density is first calculated, then it is used to kick the beam, then the beam is used to shift the ion column, and this process is repeated until convergence is achieved. Any steps towards this goal would be significant, and in one fell swoop would lead to both a greater understanding of ion effects on beams and the effectiveness of clearing methods.

Finally, there are several experiments that can be performed to further study ion trapping. First and most importantly, one would like to measure beam properties at high current in a linac in order to determine the actual impacts of ions on beam properties. Although it has yet to be done, we are able to obtain beam profiles using the wire scanner design in chapter 4 to look for changes in rms beam sizes due to ions, and one can study ion dynamics as a function of beam size. Improvements in beam diagnostic technology, driven by diagnostic challenges in similar high current linacs, should also create new experimental opportunities for the study of trapped ions. With the proper diagnostics, one can compare changes in beam properties after deliberately introducing gas into the beam line to the case where no such gas is leaked. Introducing gas in this way would ensure that ion trapping was the dominant effect observed. Additionally, by using an accelerator that is capable of increasing beam current via changes in bunch repetition rate and not bunch charge (unlike the Cornell photoinjector), one can begin to determine the beam current threshold at which ion trapping occurs. This method could also potentially make ion effects observable without gas injection, a fact that would render it the definitive experiment on ion trapping.

Ultimately, the field of ion trapping is a very rich and robust field, and there remains much room for further study in the future. But fortunately, from a practical standpoint the problem of mitigating trapped ion effects is very nearly solved. If one is not so interested in understanding the trapped ions themselves and simply cares about achieving machine stability, then enough results have been obtained experimentally to aid in the development of appropriate ion clearing techniques. Hopefully this thesis serves as a useful guide for future linac designers concerned that ion trapping may impact their accelerator.

References

[1] G. H. Hoffstaetter and M. Liepe, "Ion clearing in an {ERL}," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **557** no. 1, (2006) 205 – 212. http://www.sciencedirect.com/science/article/pii/S0168900205019935. 1, 4, 5, 8, 13, 20, 26, 27, 45, 46, 53, 70

[2] J. Byrd, A. Chao, S. Heifets, M. Minty, T. O. Raubenheimer, J. Seeman, G. Stupakov, J. Thomson, and F. Zimmermann, "First observations of a ldquo;fast beam-ion instability rdquo; at the als," in *Proceedings of the 1997 Particle Accelerator Conference (Cat. No.97CH36167)*, vol. 2, pp. 1563–1565 vol.2. May, 1997. 1

[3] A. Poncet, "Ion trapping, clearing, beam-ion interactions," in *CAS-CERN Accelerator School: Vacuum technology, Snekersten, Denmark, 28 May-3 Jun 1999: Proceedings*. 1999. http://cds.cern.ch/record/455562/files/p165.pdf. 1, 20, 46, 69

[4] J. Feikes, M. Hartrott, G. WÃijstefeld, A. Hoehl, R. Klein, C. Koschitzki, and G. Ulm, "Observation of Ion Induced Effects and their Impact on the Performance of the MLS Electron Storage Ring," in *Particle accelerator. Proceedings, 23rd Conference, PAC'09, Vancouver, Canada, May 4-8, 2009*, p. TU5RFP004. 2010. http://accelconf.web.cern.ch/AccelConf/PAC2009/papers/tu5rfp004.pdf. 1, 22, 24, 42, 46, 47, 65, 69

[5] S. Full, A. Bartnik, I. V. Bazarov, J. Dobbins, B. Dunham, G. H. Hoffstaetter, and K. Smith, "Ion effects in high brightness electron beam linacs," in *Proceedings of the 28th Linear Accelerator Conference, LINAC 16, East Lansing, Michigan, USA, September 25-30, 2016* . 2017. 1, 15

[6] A. Poncet, "Ion trapping and clearing," in *Frontiers of Particle Beams: Factories with e+ e- Rings*, M. Dienes, M. Month, B. Strasser, and S. Turner, eds., vol. 425 of *Lecture Notes in Physics*, pp. 202–221. Springer Berlin Heidelberg, 1994. http://dx.doi.org/10.1007/3540565884_11. 4, 20, 45, 46

[7] F. F. Rieke and W. Prepejchal, "Ionization cross sections of gaseous atoms and molecules for high-energy electrons and positrons," *Phys. Rev. A* **6** (Oct, 1972) 1507–1519. http://link.aps.org/doi/10.1103/PhysRevA.6.1507. 6, 54

[8] T. O. Raubenheimer, "Ion effects in future circular and linear accelerators," *Conf. Proc.* **C950501** (1996) 2752–2756. 6, 7, 8

[9] S. Full, A. Bartnik, I. V. Bazarov, J. Dobbins, B. Dunham, and G. H. Hoffstaetter, "Detection and clearing of trapped ions in the high current cornell photoinjector," *Phys. Rev. Accel. Beams* **19** (Mar, 2016) 034201. http://link.aps.org/doi/10.1103/PhysRevAccelBeams.19.034201. 10, 11, 42, 47

[10] S. B. van der Geer and M. J. de Loos, "General particle tracer (gpt)," 2016. http://www.pulsar.nl/gpt/. 11, 35, 52, 77

[11] C. Gulliford, A. Bartnik, I. Bazarov, L. Cultrera, J. Dobbins, B. Dunham, F. Gonzalez, S. Karkare, H. Lee, H. Li, Y. Li, X. Liu, J. Maxson, C. Nguyen, K. Smolenski, and Z. Zhao, "Demonstration of low emittance in the cornell energy recovery linac injector prototype," *Phys. Rev. ST Accel. Beams* **16** (Jul, 2013) 073401. http://link.aps.org/doi/10.1103/PhysRevSTAB.16.073401. 11, 35, 52, 54, 56, 64, 77

[12] A. Chao, "Lecture notes on topics in accelerator physics," Tech. Rep. SLAC-PUB-9574, SLAC, 2002. 12, 20, 64

[13] G. Hoffstaetter, D. Trbojevic, *et al.*, "Cbeta design report: Cornell-brookhaven erl test accelerator,". 15, 39, 72

[14] T. O. Raubenheimer and F. Zimmermann, "Fast beam-ion instability. i. linear theory and simulations," *Phys. Rev. E* **52** (Nov, 1995) 5487–5498. http://link.aps.org/doi/10.1103/PhysRevE.52.5487. 20, 46

[15] F. Zimmermann, T. O. Raubenheimer, and G. Stupakov, "A fast beam ion instability," *Conf. Proc.* **C950501** (1996) 3102–3104. 20, 26

[16] G. H. Hoffstaetter and C. Spethmann, "Equilibrium ion distribution in the presence of clearing electrodes and its influence on electron dynamics," *Phys. Rev. ST Accel. Beams* **11** (Jan, 2008) 014001. http://link.aps.org/doi/10.1103/PhysRevSTAB.11.014001. 20, 21, 30, 45, 69, 79

[17] Y. Baconnier, A. Poncet, and P. Tavares, "Neutralization of accelerator beams by ionization of the residual gas," in *CERN Accelerator School: Course on General Accelerator Physics Jyvaskyla, Finland, September 7-18, 1992*. 1994. 20, 33, 45, 51

[18] A. Chatterjee, K. Blaser, W. Hartung, D. Rubin, and S. T. Wang, "Fast ion instability at the cornell electron storage ring test accelerator," *Phys. Rev. ST Accel. Beams* **18** (Jun, 2015) 064402. http://link.aps.org/doi/10.1103/PhysRevSTAB.18.064402. 20, 23, 24

[19] G. Pöplau, U. van Rienen, and A. Meseck, "Numerical studies of the behavior of ionized residual gas in an energy recovering linac," *Phys. Rev. ST Accel. Beams* **18** (Apr, 2015) 044401. http://link.aps.org/doi/10.1103/PhysRevSTAB.18.044401. 20, 21, 30, 111

[20] L. Wang, Y. Cai, T. O. Raubenheimer, and H. Fukuma, "Suppression of beam-ion instability in electron rings with multi-bunch train beam fillings," *Phys. Rev. ST Accel. Beams* **14** (2011) 084401. 21, 30, 31, 32

[21] L. Wang, J. Safranek, Y. Cai, J. Corbett, R. O. Hettel, T. O. Raubenheimer, J. Schmerge, and J. Sebek, "Beam ion instability: Measurement, analysis, and simulation," *Phys. Rev. ST Accel. Beams* **16** (Oct, 2013) 104402. http://link.aps.org/doi/10.1103/PhysRevSTAB.16.104402. 24, 47

[22] B. Dunham, J. Barley, A. Bartnik, I. Bazarov, L. Cultrera, J. Dobbins, G. Hoffstaetter, B. Johnson, R. Kaplan, S. Karkare, V. Kostroun, Y. Li, M. Liepe, X. Liu, F. Loehl,

J. Maxson, P. Quigley, J. Reilly, D. Rice, D. Sabol, E. Smith, K. Smolenski, M. Tigner, V. Vesherevich, D. Widger, and Z. Zhao, "Record high-average current from a high-brightness photoinjector," *Applied Physics Letters* **102** no. 3, (2013) 034105, http://dx.doi.org/10.1063/1.4789395. http://dx.doi.org/10.1063/1.4789395. 28, 29

[23] A. Bartnik, "Cathodes in cw operation at the cornell photoinjector," in *Photocathode Physics for Photoinjectors (P3): Jefferson Lab, VA, USA, Oct 17-19, 2016*. 2016. https://www.jlab.org/indico/event/124/session/7/contribution/53/material/slides/.

[24] A. Bartnik. Private communication, 2017. 28

[25] L. Cultrera. Private communication, 2017. 29

[26] G. Parzen, ""electric fields of a uniformly charged elliptical beam"," tech. rep., BNL/SNS Technical Note, 2001. 33

[27] E. Keil, "Beam-beam dynamics," Tech. Rep. CERN-1995-06, CERN, 1994. 34

[28] I. Bazarov *et al.*, "The Cornell-BNL FFAG-ERL Test Accelerator: White Paper," arXiv:1504.00588 [physics.acc-ph]. 39, 72

[29] M. Takao, H. Ego, Y. Kawashima, Y. Ohashi, T. Ohshima, and H. Saeki, "Observation of ion effects at the SPring-8 storage ring," in *Particle accelerator. Proceedings, 8th European Conference, EPAC 2002, Paris, France, June 3-7, 2002*, pp. 1562–1564. 2002. http://accelconf.web.cern.ch/AccelConf/e02/PAPERS/WEPRI037.pdf. 46, 58

[30] J. Marriner, D. Mohl, Y. Orlov, A. Poncet, and S. van der Meer, "Experiments and practice in beam shaking," *Part. Accel.* **30** (1990) 13–20. 46

[31] T. Moore, N. I. Agladze, I. V. Bazarov, A. Bartnik, J. Dobbins, B. Dunham, S. Full, Y. Li, X. Liu, J. Savino, and K. Smolenski, "Fast wire scanner for intense electron beams," *Phys. Rev. ST Accel. Beams* **17** (Feb, 2014) 022801. http://link.aps.org/doi/10.1103/PhysRevSTAB.17.022801. 47, 70, 88, 89, 91, 92, 93, 95, 96, 97, 98, 99, 100, 101, 102, 103, 104

[32] J. Dobbins, R. P. K. Kaplan, F. Loehl, and C. R. Strohman, "Control of RF Transients in Cavities Induced by Pulsed High Current Beams," *Conf. Proc.* **C1205201** (2012) 2891–2893. 56

[33] J. Dobbins. Private communication, 2016. 71

[34] R. F. Holsinger and K. Halbach, "Poisson superfish," 2016. http://laacg.lanl.gov/laacg/services/download_sf.phtml. 72

[35] M. S. Alnæs, J. Blechta, J. Hake, A. Johansson, B. Kehlet, A. Logg, C. Richardson, J. Ring, M. E. Rognes, and G. N. Wells, "The fenics project version 1.5," *Archive of Numerical Software* **3** no. 100, (2015) . 73

[36] A. Logg, K.-A. Mardal, G. N. Wells, *et al.*, *Automated Solution of Differential Equations by the Finite Element Method*. Springer, 2012. 73

[37] C. Geuzaine and J. F. Remacle, "Gmsh: a three-dimensional finite element mesh generator with built-in pre- and post-processing facilities," *International Journal for Numerical Methods in Engineering* (2009) . 73

[38] J. Ahrens, B. Geveci, and C. Law, *ParaView: An End-User Tool for Large Data Visualization*. Elsevier, 2005. 73

[39] D. R. Nicholson, *Introduction to Plasma Theory*. Wiley Publications, 2 ed., 1983. 79

[40] P.-O. Persson, "Distmesh," 2016. https://people.sc.fsu.edu/~jburkardt/m_src/distmesh/distmesh.html. 80

[41] P. Tenenbaum and T. Shintake, "Measurement of small electron-beam spots," *Annual Review of Nuclear and Particle Science* **49** no. 1, (1999) 125–162, http://dx.doi.org/10.1146/annurev.nucl.49.1.125. http://dx.doi.org/10.1146/annurev.nucl.49.1.125. 88, 94

[42] M. Abo-Bakr, W. Anders, R. Barday, A. Bondarenko, K. Bürkmann-Gehrlein, V. Dürr, S. Heßler, A. Jankowiak, T. Kamps, J. Knobloch, O. Kugeler, B. Kuske, P. Kuske, A. Matveenko, A. Meseck, G. Meyer, R. Müller, A. Neumann, K. Ott, Y. Petenev, D. Pflückhahn, T. Quast, J. Rahn, and S. Schubert, "Conceptual design report BERLinPro," tech. rep., Helmholtz Zentrum Berlin, 2012. 111

[43] Sakanaka, "Investigation of bunch-gap effects for curing ion trapping in energy recovery lin," *Journal* (Year) . 113