

Detection and clearing of trapped ions in the high current Cornell photoinjector

S. Full,^{*} A. Bartnik, I. V. Bazarov, J. Dobbins, B. Dunham, and G. H. Hoffstaetter

CLASSE, Cornell University, Ithaca, New York 14853, USA

(Received 22 July 2015; published 3 March 2016)

We have recently performed experiments to test the effectiveness of three ion-clearing strategies in the Cornell high intensity photoinjector: DC clearing electrodes, bunch gaps, and beam shaking. The photoinjector reaches a new regime of linac beam parameters where high continuous wave beam currents lead to ion trapping. Therefore ion mitigation strategies must be evaluated for this machine and other similar future high current linacs. We have developed several techniques to directly measure the residual trapped ions. Our two primary indicators of successful clearing are the amount of ion current removed by a DC clearing electrode, and the absence of bremsstrahlung radiation generated by beam-ion interactions. Measurements were taken for an electron beam with an energy of 5 MeV and continuous wave beam currents in the range of 1–20 mA. Several theoretical models have been developed to explain our data. Using them, we are able to estimate the clearing electrode voltage required for maximum ion clearing, the creation and clearing rates of the ions while employing bunch gaps, and the sinusoidal shaking frequency necessary for clearing via beam shaking. In all cases, we achieve a maximum ion clearing of at least 70% or higher, and in some cases our data is consistent with full ion clearing.

DOI: 10.1103/PhysRevAccelBeams.19.034201

I. INTRODUCTION

In an accelerator's vacuum chamber, any residual gas is rapidly ionized by collisions with the electron beam. At high beam currents, the resulting positive ions become trapped inside of the negatively charged beam and can cause a variety of effects including charge neutralization, coherent and incoherent tune shifts, optical errors, beam halo, beam losses, 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 potential effects.

Ion clearing methods are understood well enough to mitigate ion effects in most storage rings and synchrotrons [2–4]. In low repetition rate linacs, ion trapping is typically not observed because the ions have time to drift out of the center of the beam pipe between bunches. However, future linacs, such as the Cornell Energy Recovery Linac (ERL), enter a new continuous wave (CW) high current regime where ion trapping is unavoidable [1]. Therefore it has become especially important to anticipate what clearing methods will be most effective in future high current linacs.

Up until this point, many simulations have been created to study ion trapping and clearing in ERLs [5,6]. In general,

these simulations are difficult to create and can have prohibitively long run times, particularly if one simulates space charge repulsion between a large number (upwards of millions) of ions. Additionally, these simulations have yet to be verified by experiment, because data for this range of beam parameters is very scarce.

The high current Cornell photoinjector is one of the few linacs in the world where different ion mitigation strategies can be experimentally tested for this parameter regime. 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.

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 [5]. Although clearing electrodes can achieve significant ion clearing, often reducing the beam neutralization fraction to just a few percent [3,4], it is important to explore other clearing methods as well. This is because the electrodes often benefit from the combined use of other techniques. In addition, electrodes are often 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 allow the ions to drift out of

^{*}sf345@cornell.edu

the center of the beam pipe while the beam is absent. This technique has been employed in many electron storage rings with great success [3,7]. 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, one of the side effects of this technique [1]. Secondary concerns involve beam instabilities such as the fast ion instability, which occur because ions still accumulate over the course of a single bunch train [8]. 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 Antiproton Accumulator Ring [9] and more recently in the Metrology Light Source (MLS) [10]. 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 [10]. The necessary frequency is typically close to the ion oscillation frequency, and is usually determined experimentally by trial and error [3,10].

All three methods can result in significant clearing, and can be even more effective when deployed in tandem [2]. 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.

II. EXPERIMENTS

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 time scales 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 [11], 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 beam

position monitors (BPMs) connected to a spectrum analyzer [10,12]. 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, noncirculating 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. 1. Before leaking gas, no such excess radiation was previously observed in the 10–20 mA range, indicating that this extra radiation (presumably bremsstrahlung) was caused entirely by beam-gas

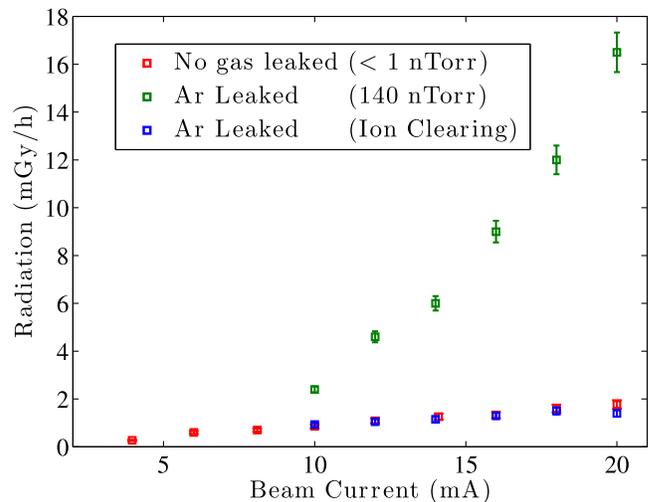


FIG. 1. After leaking gas into the beam pipe, background radiation levels rose dramatically due to bremsstrahlung generated by beam-ion interactions. Removing the trapped ions using clearing electrodes (shown above) or other clearing methods reduced this excess radiation to normal background levels.

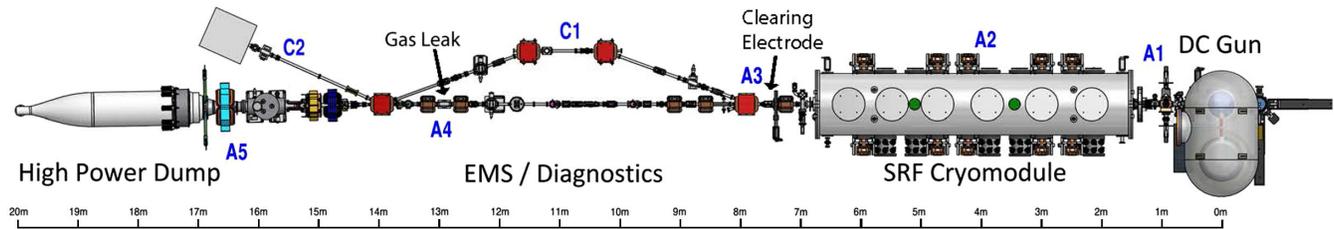


FIG. 2. A schematic of the photoinjector. The ion clearing electrode was installed just after the beam exits the SRF cavity at A3, and gas was leaked in at the end of A4. Radiation measurements were taken at several locations between sections A3 and A4 (next to the beam pipe).

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. 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 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.

A. 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 and its location in the beam line is shown in Fig. 2. The

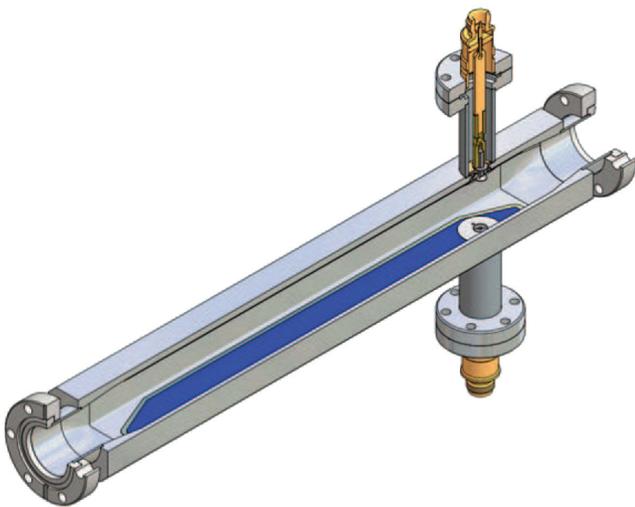


FIG. 3. The ion clearing electrode used during the experiments. The electrode coating is highlighted in blue, and is approximately 35 cm long and 3.5 cm wide.

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 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 toward 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 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 and 20 mA, is shown in Figs. 4 and 5. The beam current was varied by changing bunch charge (from 5 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. 5. This observation is typical throughout our experiments.

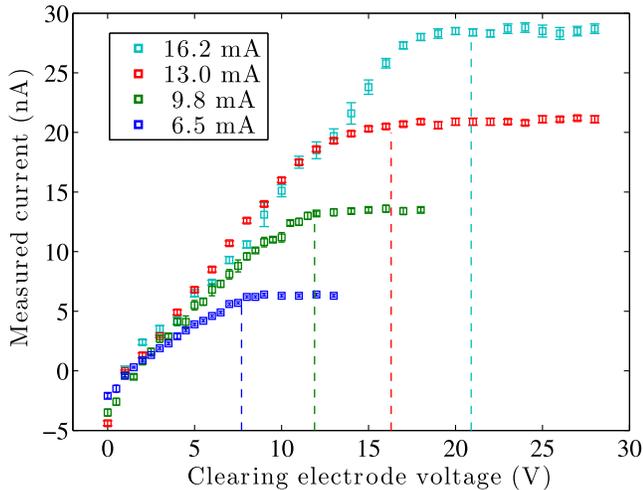


FIG. 4. A picoammeter was used to measure the ion current striking the clearing electrode for different applied voltages. The vertical dotted lines mark the minimum voltage required for full ion clearing, as predicted using equation (2).

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 [4]

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

where λ is the number of electrons per unit length, e is the elementary charge, and σ_b is the rms transverse beam size. Full clearing occurs when the clearing electrode's field exceeds the beam's peak electric field (at $r = \sigma_b$). For electrodes separated by a distance d , the field of the

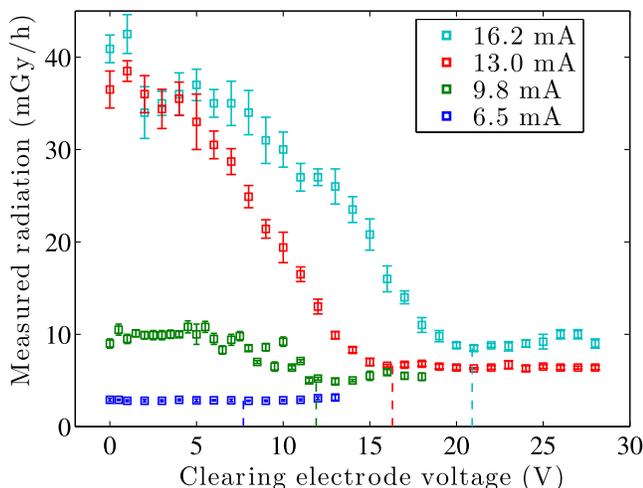


FIG. 5. Background radiation levels also decreased while employing the clearing electrodes due to the absence of bremsstrahlung caused by beam-ion interactions. The vertical lines again indicate the minimum voltage needed for complete ion clearing, as computed by Eq. (2).

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 $E_{\text{electrode}} = V_{\text{electrode}}/d \geq E_{\text{beam}}$. This yields the minimum voltage required

$$V_{\text{electrode}} \geq \frac{\lambda e}{2\pi\epsilon_0} \frac{d}{\sigma_b}. \quad (2)$$

Although we cannot take direct beam profile measurements in order to determine σ_b , we can obtain estimates using general particle tracer (GPT) [13], 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 [13], 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 the beam size significantly. Using the vertical beam size from GPT and a clearing electrode separation of $d = 40.6$ mm, Eq. (2) predicts clearing voltages that agree with measurements to within a few percent, as shown in Figs. 4 and 5. Our calculated values are shown in Table I. Vertical dotted lines were drawn in Figs. 4 and 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]

$$\tau_{\text{create}} = \frac{1}{\sigma_{\text{col}}\rho_{\text{gas}}c} \quad (3)$$

where σ_{col} is the collision ionization cross section of the gas species (these values are readily available [14]), 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 $L/c\tau_{\text{create}}$ ions. The beam current therefore produces $L/c\tau_{\text{create}} \cdot I/e$ ions per second in this section. Thus the maximum ion current is

TABLE I. The minimum clearing electrode voltage necessary for full ion clearing, calculated using Eq. (2). The rms transverse beam sizes σ_x and σ_y were obtained using GPT simulations [13].

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

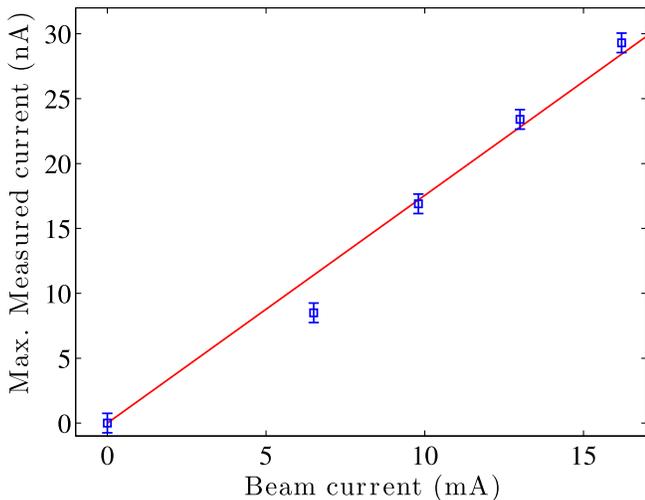


FIG. 6. The longitudinal range of the clearing electrode can be estimated by using the maximum ion current measured by the clearing electrode. The range can be extracted from a best fit line (in red) to our data.

$$I_{\max} = \frac{I_{\text{beam}}}{c} \left(\frac{L_{\text{create}}}{\tau_{\text{create}}} \right). \quad (4)$$

The ratio $L_{\text{create}}/\tau_{\text{create}}$ can be found by treating it as a fit parameter to our data, shown in Fig. 6. Note that in Fig. 6 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 300 K, this corresponds to a beam neutralization time τ_{create} of roughly 5.2 ms. Therefore, using the fit parameter found in Fig. 6, 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.

B. 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 [13]) 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 superconducting radio frequency (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 predefined limits and tripped off the machine at around 8 mA. We have a preexisting feedforward system [15], 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 ions travel longitudinally down the beam pipe toward 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 Figs. 7 and 8. The radiation data in Fig. 8 shows that the trapped ions are removed even without the clearing electrode turned on, confirming that the bunch gaps are the dominant

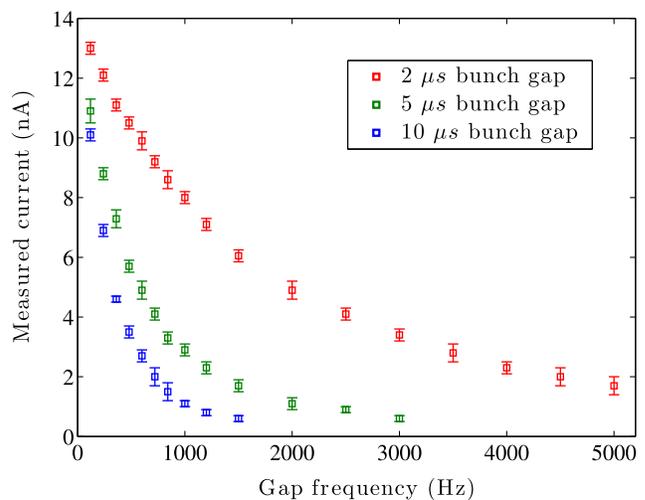


FIG. 7. Increasing the frequency and duration of bunch gaps reduces the number of trapped ions that reach the clearing electrode. For each data point, the beam current was held fixed at 10 mA.

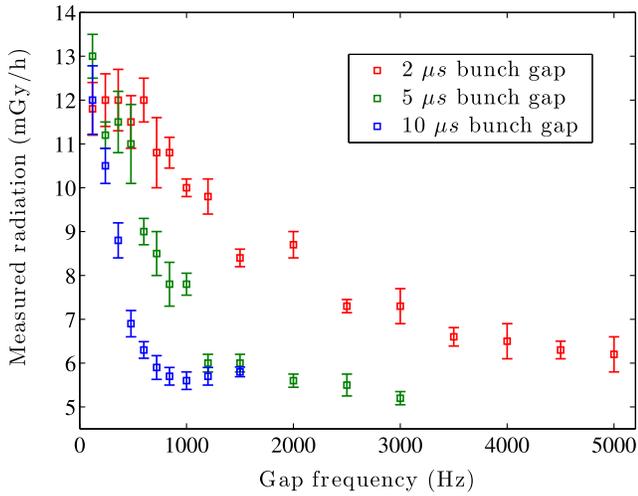


FIG. 8. The radiation caused by beam-ion interactions is reduced by increasing the frequency and duration of bunch gaps. This data was obtained while the clearing electrode was turned off. For each data point, the beam current was held fixed at 10 mA.

clearing mechanism. This data also agrees with previous experiments at SPring-8 [7] 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 9 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

$$f_{\text{avg}}(R_g) = \frac{1}{1 + \left(\frac{\tau_1}{\tau_2}\right) \left(\frac{R_g \Delta L_g}{1 - R_g \Delta L_g}\right)} \quad (5)$$

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

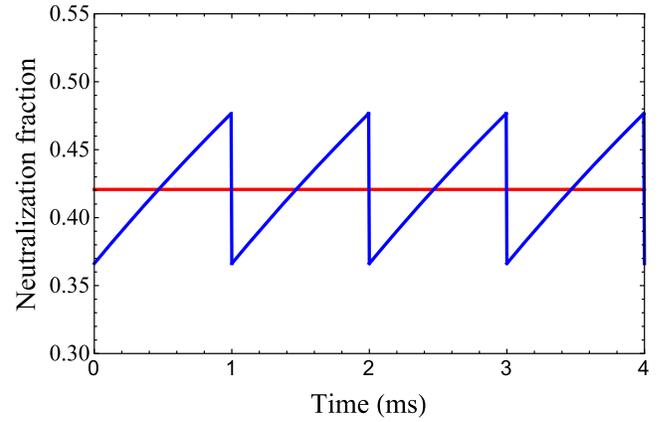


FIG. 9. Ions are created via collision ionization while the beam is on and decay exponentially during the bunch gaps. The equilibrium neutralization fraction, indicated by the red line, was found using Eq. (5) for a bunch gap duration of $5 \mu\text{s}$, a gap frequency of 1 kHz, and the ion creation to clearing ratio given in Table II.

fit curves are compared with our data in Fig. 10, and the best fit ratio of τ_1/τ_2 for each curve is shown in Table II. 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. 7 because it has less experimental uncertainty.

Assuming that τ_1 is roughly the time it takes to achieve full beam neutralization (i.e., $\tau_1 = \tau_{\text{create}} = 5.2 \text{ ms}$), then the ratio τ_1/τ_2 can be used to predict a clearing time of $21 \mu\text{s}$. 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

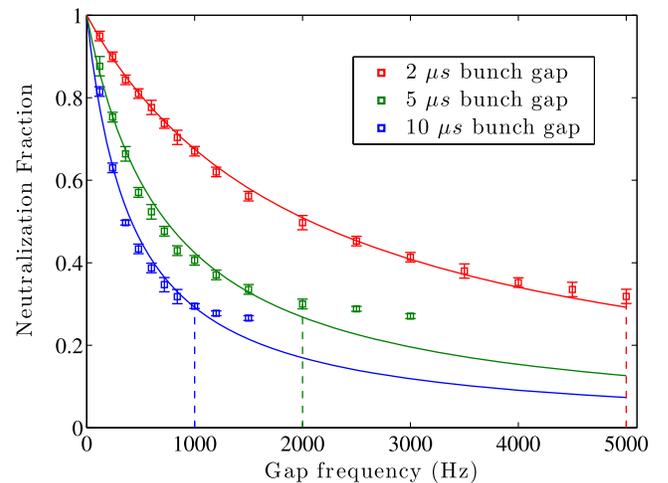


FIG. 10. Increasing the frequency and duration of bunch gaps reduces the trapped ion density as shown by the residual ion current hitting a clearing electrode. The curves are best fits obtained using Eq. (5). The vertical dotted lines mark the locations where a 1% reduction in beam current results in an approximately 70% reduction of trapped ion density.

TABLE II. 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

17.5 μs , which is of the same order of magnitude as this estimate for τ_2 .

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. 10 shows that an approximately 70% reduction in ion density can be achieved while retaining nearly 99% 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% reduction in ion density would require over a 30% reduction in maximum 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.

C. 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 [16]

$$F(r) = \frac{\lambda e^2}{2\pi\epsilon_0 r} \left[1 - \exp\left(-\frac{r^2}{2\sigma_b^2}\right) \right] \quad (6)$$

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 [13], 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

$$\frac{d^2 r}{dt^2} + \omega_i^2 r = 0 \quad (7)$$

where ω_i is the oscillation frequency of the ions. Using the linearized form of (6), it follows that this oscillation frequency is given by [17]

$$\omega_i = \sqrt{\frac{2r_p c}{e} \frac{I}{A\sigma_b^2}} \quad (8)$$

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. 11. Measurements were taken for several gas species, including N_2 , Ar, and Kr, in order to confirm that

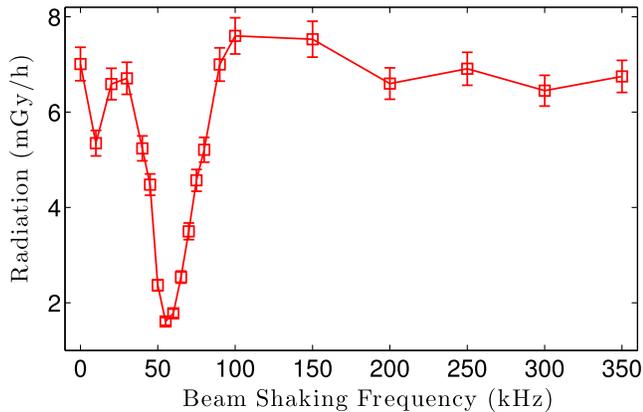


FIG. 11. Shaking the beam at frequencies near the ion oscillation frequency eliminates the excess radiation caused by beam-ion interactions.

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 [10] and not the Cornell photoinjector.

In the course of our experiments we attempted to confirm the three scaling laws predicted by Eq. (8): 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 12 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

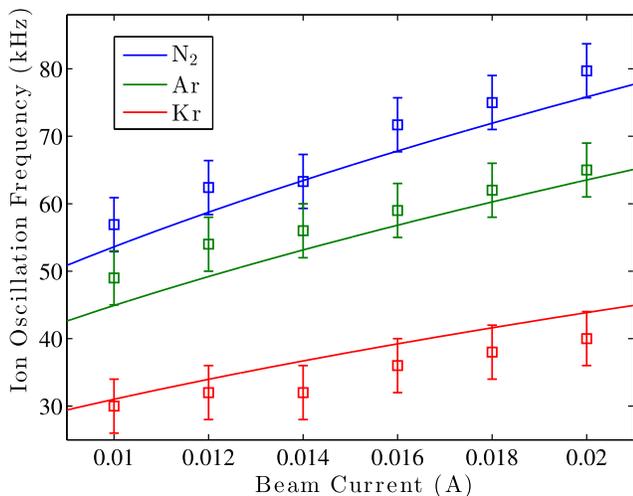


FIG. 12. Resonance frequencies for various beam currents and ion species. The circles represent data points, while the lines indicate best fits in the form of Eq. (8), where the beam size is used as a fit parameter.

± 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 $\sigma_b = 4.2$ mm was used to obtain the fit curves for the data in Fig. 12. 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 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. 13. The simulated beam sizes for the three different optics settings used during the experiment are shown in Fig. 14.

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

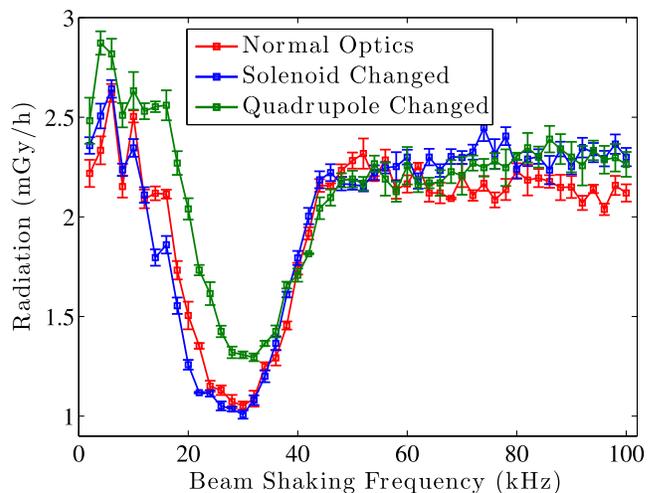


FIG. 13. Radiation levels for various beam shaking frequencies. Changing the beam size by over a factor of 2, as shown in Fig. 14, does not result in a significant shift of the resonance frequency, in disagreement with theory.

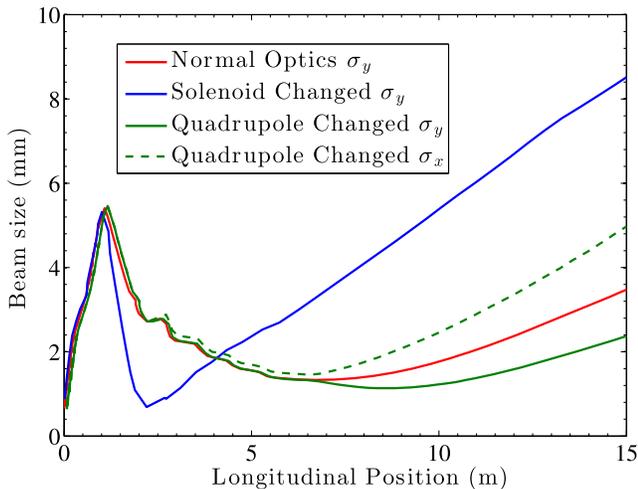


FIG. 14. The three optics settings used while obtaining the data in Fig. 13. The beam was round for both the normal and solenoid optics settings. The clearing electrode and gas leak were located at 7.5 and 13 m, respectively, and the beam exits the SRF cavity at 6.7 m.

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. 14. 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.

III. CONCLUSIONS

Experiments were performed to test the effectiveness of three different ion clearing techniques in the Cornell ERL photoinjector. During the first experiment, we installed a DC clearing electrode and measured the amount of ions it removed as a function of clearing voltage. It was found that the voltage necessary for maximum clearing can be predicted by considering the suppression of the transverse potential of a constant charge distribution beam. The total measured ion current can also be used to estimate the effective longitudinal range of the clearing electrode. In the second experiment, we introduced bunch gaps and determined the amount of clearing for different combinations of bunch gap duration and frequency. The amount of clearing depended only on the total time the beam was turned off, and was independent of the bunch gap duration and frequency. Finally, we used an electrode to shake the beam sinusoidally and resonantly clear out any trapped ions. In this case, the shaking frequency needed to induce a resonance was the trapped ion oscillation frequency.

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 [5]. 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 1 kV [2,10]. 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% 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 [11]. 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.

ACKNOWLEDGMENTS

The authors would like to thank Atoosa Meseck for helpful discussions concerning the planning of this experiment. This work was supported by financial assistance from the U.S. Department of Energy (Grant No. DE-SC0012493) and the National Science Foundation (Award No. NSF-DMR 0807731).

APPENDIX A: 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

$$f_1(t) = 1 - (1 - f_0)e^{-t/\tau_1} \quad (\text{A1})$$

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. 15.

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

$$f_2(t) = f_m e^{-(t-T_1)/\tau_2} \quad (\text{A2})$$

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. 15). 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 ,

$$f_0 = \frac{1 - e^{-\frac{T_1}{\tau_1}}}{1 - e^{-\frac{T_1+T_2}{\tau_1}}}, \quad (\text{A3})$$

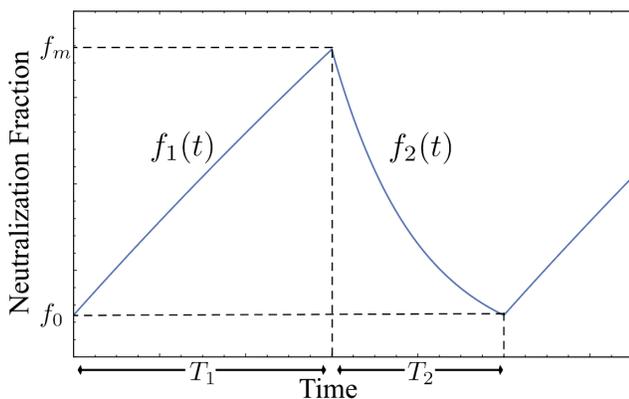


FIG. 15. A sketch of the neutralization fraction while employing bunch gaps. The neutralization fraction increases while the beam is on until some time T_1 , then decays exponentially during the bunch gaps with a duration of T_2 .

$$f_m = 1 - (1 - f_0)e^{-T_1/\tau_1}. \quad (\text{A4})$$

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

$$f_{\text{avg}} = \frac{1}{T_1} \int_0^{T_1} f_1(t) dt + \frac{1}{T_2} \int_{T_1}^{T_1+T_2} f_2(t) dt. \quad (\text{A5})$$

Performing this integral yields the formula for the average neutralization fraction

$$f_{\text{avg}} = \frac{T_1}{T_1 + T_2} - 2 \frac{\tau_1 - \tau_2}{T_1 + T_2} \frac{\sinh(\frac{T_1}{2\tau_1}) \sinh(\frac{T_2}{2\tau_2})}{\sinh(\frac{T_1}{2\tau_1} + \frac{T_2}{2\tau_2})}. \quad (\text{A6})$$

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

$$f_{\text{avg}} = \frac{1}{1 + (\frac{\tau_1}{\tau_2})(\frac{T_2}{T_1})}. \quad (\text{A7})$$

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 Eq. (5).

-
- [1] G. H. Hoffstaetter and M. Liepe, *Nucl. Instrum. Methods Phys. Res., Sect. A* **557**, 205 (2006).
 - [2] A. Poncet, Ion trapping, clearing, beam-ion interactions, in *CAS-CERN Accelerator School: Vacuum technology, Snekersten, Denmark, 1999: Proceedings, 1999*.
 - [3] A. Poncet, Ion trapping and clearing, in *Frontiers of Particle Beams: Factories with e+ e Rings*, edited by M. Dienes, M. Month, B. Strasser, and S. Turner, Lecture Notes in Physics (Springer, Berlin Heidelberg, 1994), Vol. 425, p. 202.
 - [4] Y. Baconnier, A. Poncet, and P. Tavares, in *CERN Accelerator School: Course on General Accelerator Physics Jyväskylä, Finland, 1992, 1994*.
 - [5] G. H. Hoffstaetter and C. Spethmann, *Phys. Rev. ST Accel. Beams* **11**, 014001 (2008).
 - [6] G. Pöplau, U. van Rienen, and A. Meseck, *Phys. Rev. ST Accel. Beams* **18**, 044401 (2015).
 - [7] M. Takao *et al.*, in *Proceedings of the 8th European Particle Accelerator Conference, Paris, 2002* (EPS-IGA and CERN, Geneva, 2002), p. 1562.
 - [8] T. O. Raubenheimer and F. Zimmermann, *Phys. Rev. E* **52**, 5487 (1995).
 - [9] J. Marriner, D. Mohl, Y. Orlov, A. Poncet, and S. van der Meer, *Part. Accel.* **30**, 13 (1990).
 - [10] J. Feikes *et al.*, in *Proceedings of the 23rd Particle Accelerator Conference, Vancouver, Canada, 2009* (IEEE, Piscataway, NJ, 2009), p. TU5RFP004.

-
- [11] T. Moore *et al.*, *Phys. Rev. ST Accel. Beams* **17**, 022801 (2014).
- [12] L. Wang, J. Safranek, Y. Cai, J. Corbett, R. O. Hettel, T. O. Raubenheimer, J. Schmerge, and J. Sebek, *Phys. Rev. ST Accel. Beams* **16**, 104402 (2013).
- [13] C. Gulliford *et al.*, *Phys. Rev. ST Accel. Beams* **16**, 073401 (2013).
- [14] F. F. Rieke and W. Prepejchal, *Phys. Rev. A* **6**, 1507 (1972).
- [15] J. Dobbins, R. P. K. Kaplan, F. Loehl, and C. R. Strohman, *Conf. Proc. C1205201*, 2891 (2012).
- [16] E. Keil, CERN Report No. CERN-1995-06, 1994 (unpublished).
- [17] A. Chao, SLAC Report No. SLAC-PUB-9574, 2002 (unpublished).