Online Emittance Monitoring for the New Ion Source and RFQ at Fermilab

Purnima P. Balakrishnan
Fermi National Accelerator Laboratory, Batavia, IL 60510
(Dated: August 17, 2012)

The old ion source and Cockroft-Waltons at the Fermilab Linac are being replaced with a new ion source and RFQ. The plan is to install them this fall during the shutdown, but in order to do so, the beam outputted by the RFQ must have the correct energy, current, and emittance to be accepted by the Linac. Currently, emittance data is gathered using emittance probes controlled through ACNET, but calculations are done offline, after the collection of data. This paper describes a new ACNET console program which gathers emittance data and performs all data analysis and calculations in realtime. This allows emittance measurements to be done more rapidly and easily, while still giving the option of double checking or doing more detailed analysis offline. This program was used to measure the horizontal and vertical emittances downstream of the RFQ, at two different power levels. The normalized horizontal and vertical emittances were measured as 0.6π and 0.4π mm·mrad, respectively. While these emittance values are acceptable for injection into the Linac, more work may need to be done to optimize the efficiency and power usage of the RFQ before its installation.

I. INTRODUCTION

Due to accelerator design constraints, high energy beams must be created through multiple stages. The first stage for the accelerators at Fermilab, called the Linac injector or pre-accelerator, is the Cockroft-Walton, which uses a large DC voltage to accelerate H^- ions to 750 keV. The ions are injected into the Linac, which accelerates them through a series of cavities containing oscillating electromagnetic waves to 400 MeV. The ions are then stripped of their electrons, converted into bare protons, and accelerated by a series of synchrotrons operating at different energies: the Booster (8 GeV), Main Injector (150 GeV), and finally the Tevatron (980 GeV) [1].

Because each accelerator is larger than the previous one, each successive stage takes longer to fill with particles and reach the final energy. For the Tevatron to fully power up, the Booster has to cycle almost 800 times [1], and the Linac even more. With plans of even higher intensity beams, the early stage accelerators must be increasingly robust. The current Cockroft-Waltons started running in 1968, and although they have been fairly reliable, have recently been having problems creating a consistent beam [2]. As they become older, they will become even harder to maintain; for this reason, the Accelerator Division plans to replace the Cockroft-Waltons with newer technology which is more reliable, more efficient, smaller, and easier and faster to fix in the event that something does break.

A. The New Linac Injector

The new injector for the Linac consists of two swappable 35 keV H^- sources, quadrupoles, solenoids, and corrector magnets for steering, a chopper, and a radio-frequency quadrupole (RFQ) to create bunches and accelerate the beam to 750 keV. The RFQ consists of a cavity containing four copper rods—in a quadrupole configuration—on support stems and tuning plates, which can be adjusted to tune the frequency response of the cavity. Inputted RF power produces an oscillating electromagnetic field which accelerates and focuses the beam. In addition, the rods have sinusoidal modulations in them, with a wavelength that grows with the beam energy; these modulations modify the field inside the cavity to separate the particles into bunches. The entire injector design is based off of the injector at Brookhaven National Laboratory in the 1990s, which has shown to be very reliable and easy to maintain [2].

Fermilab is currently in a shutdown, and because of time constraints, the RFQ will be installed in the fall and winter of this year. Although the Tevatron and many of the prior accelerators are shutdown, the Linac is still partially running, creating neutrons for cancer treatment at the Neutron Therapy Facility. The patients have strict treatment schedules, so the RFQ needs to work properly when installed to minimize any downtime. Therefore, the entire pre-accelerator system must be tested prior to installation to see if it matches designed parameters. There must be ways of measuring the outputted beam energy, RF frequency, transmission efficiency, and emittance, to make sure that they are acceptable for injection into the Linac. In addition, once the beam is running, it is important to have a way to easily monitor these parameters in order to keep the system running properly.

B. Emittance

The transverse emittance of a beam is a combined measure of how large and how divergent the beam is. The size of a beam can be characterized by the position, x, of particles relative to the beam center, while divergence can be described either by their transverse momentum p_x, or by their trajectory angle from the beam, x', which are linearly related for fast particles[1]. Alone, these numbers do not mean much—a large beam may be focusing together or defocusing apart, which are very different scenarios—but together can say a lot about the beam.

At any point in time, we can plot the angle versus position of all particles in a beam, creating a distribution in phase space. For most beams, this distribution is elliptical. The emittance is defined as the half-axis product (the area divided by π) of this ellipse in phase space. This quantity is conserved as the beam travels through various steering devices, so knowledge of beam distribution at one point can reveal a lot about the distribution at another point along the beam path.

A single particle passing through dipoles and quadrupoles acts like a harmonic oscillator. In this case, the emittance is given as a function of the Twiss parameters α, β, and γ, which describe the size and change in size of the oscillation envelope:

$$\varepsilon = \gamma x^2 + 2\alpha x x' + \beta x'^2 \tag{1}$$

However, for entire beams, the actual size of the beam is difficult to characterize, as there will always be some number of particles everywhere in the beam pipe. It is better to use the RMS emittance, which depends on the RMS size of the beam [3]. This can be calculated as

$$\varepsilon_{rms} = \sqrt{\sigma_x^2 \sigma_{x'}^2 - \mathrm{Cov}_{xx'}^2} \tag{2}$$

where σ_x^2 and $\sigma_{x'}^2$ are the variances in position and angle, respectively, and $\mathrm{Cov}_{xx'}$ is their covariance. Larger variances in position and angle will create thicker or taller, and therefore bigger ellipses, but a larger covariance will cause the ellipse to be thinner, but angled, decreasing the emittance. The Twiss parameters for the beam as a whole can also be calculated as

$$\alpha = \frac{-\langle x x' \rangle}{\varepsilon_{rms}} \qquad \beta = \frac{-\langle x^2 \rangle}{\varepsilon_{rms}} \qquad \gamma = \frac{-\langle x'^2 \rangle}{\varepsilon_{rms}} \tag{3}$$

where $\langle x \rangle$ signifies the average of x over the particles in the beam.

As the particles are accelerated, their longitudinal velocity increases, while their transverse velocities remain the same. This gives a false sense that the beam is squeezing together as it is accelerated. In order to preserve the emittance as a constant parameter of the beam, we can normalize it relativistically:

$$\varepsilon_{\text{norm}} = \gamma \beta \varepsilon \tag{4}$$

where $\beta = v_z/c$ and $\gamma = 1/\sqrt{1-\beta^2}$ depend on the velocity of particles along the beam. This normalized emittance does remain constant even as the beam is accelerated, so is a good way to compare emittances between different stages of acceleration.

[1] $x' = \tan^{-1}(\frac{p_x}{p_z}) \approx \frac{p_x}{p_z}$ for small p_x or large p_z by the small angle approximation

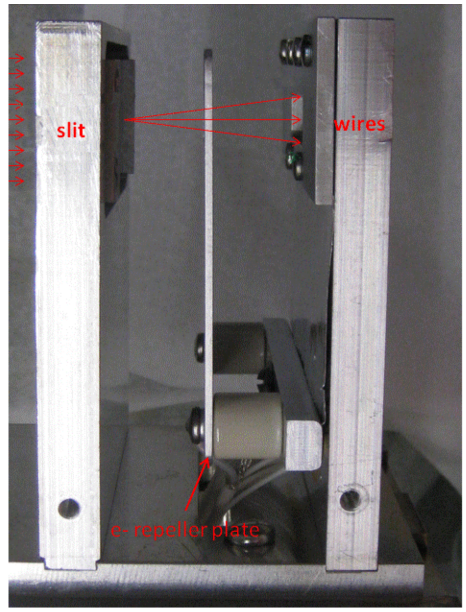

FIG. 1: A photo of the inside of an emittance probe. The beam, in red, passes through a slit and then hits the wires with some distribution based on the angle of incidence. This causes a voltage along the wires which can be read out. There is a plate in the center, kept at 70 V to repel electrons back towards the wires.

The rest of this paper describes a new program to measure and calculate both the RMS and normalized RMS emittances for the new pre-accelerator.

II. MEASUREMENT HARDWARE

A. Emittance Probes

The emittance measurements for the RFQ, as well as for the Cockroft-Walton and Linac, are made using a multi-wire emittance probe. This apparatus is a metal box with a slit in the front and a set of wires parallel to the slit along the back wall, as shown in figure 1. The box is attached to the end of a probe, which can be moved in and out of the beam path, in a direction perpendicular to the slit.

As the probe moves into the beam, some of the H^- particles pass through the slit. Some of these will then hit a wire, depending on the angle of entrance through the slit, and cause a large voltage across the wire. Occasionally, the ion will be energetic enough to knock electrons off of the wire, thereby decreasing the read voltage. In order to counteract this, a bias voltage of 70 V is applied between the wires and a central plate; this attracts the electrons back. The resulting voltage measurement across each wire is linearly proportional to the number of incident ions.

The voltages on each wire, which select for angle, are recorded for each position of the slit. This allows us to measure approximately how many ions are at each point in phase space. There are two probes at each location in the beam. One sweeps horizontally, and one vertically. This allows us to calculate the emittance in each direction, since the accelerator components may not be completely symmetric.

While this apparatus is very useful, there are two major problems. First, the measurement is destructive, meaning that the probe blocks the entire beam. Second, the wires in the probe can be damaged from constant exposure to the energetic particles in the beam. In order to minimize beam downtime and damage to the equipment, the measurement has to be done quickly and efficiently.

B. Test Stand

Since the RFQ is still under development, it is constantly being modified and worked on, so it is not always running, nor is everything always hooked up. The test stand is a convenient place to test various components, such as sources, probes, or in this case, programs. It currently contains a source similar to the one in the pre-accelerator, a 35 keV extractor, an electric lens to focus the beam, two emittance probes, and a lead-glass shield to stop the beam, as shown in figure 2. This provides beam which can be used for emittance measurements. The only major difference in measurements comes from the energy difference between the test stand beam and one that has passed through the RFQ. A more energetic beam spreads out much less in the same distance, as mentioned earlier, so the test beam will be larger. Despite this, the test stand is a perfect place to test the emittance program, because it produces a beam, has working probes, and there aren't too many other components that will affect the beam.

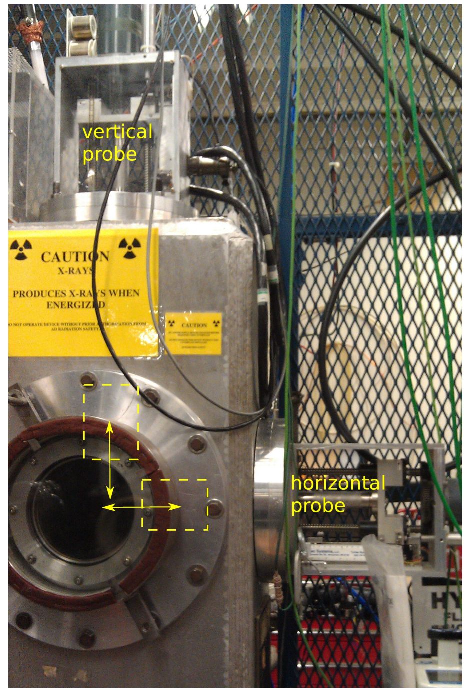

FIG. 2: A front view of the test stand, with the beam coming out of the page. The source, extractor, and Einzel lens are on the far side of the high-voltage fence. The motors for the probes protrude from the stand.

III. SOFTWARE FRAMEWORK

A. ACNET Controls System

The Fermilab controls system, ACNET, encompasses controls for all of Fermilab. It provides a framework to make and run console programs, including packages to write user interfaces, communicate with devices (such as motors), access a central database, and produce graphics. Programs for ACNET can be written in C++, Java, or Fortran. Once compiled, the program is placed on a "page", which can be accessed through the console. Opening the page runs the program in the current window.

ACNET makes it very simple to decipher user actions, called "interrupts". An enter or left click will trigger a keyboard interrupt, which contains the location of the click. Periodically, at 15 Hz, a periodic interrupt is sent to the program. A click on the leftmost corner of the window will send a termination interrupt, which will cause the program to exit.

B. Existing Emittance Programs

The current program for running emittance scans was written by W. Marsh in 2009 [4]. The page is shown in figure 3. It scans an emittance probe through a user-specified range of positions, stopping at each position to record the voltages across each wire, and the actual position of the probe. These values are stored in an array, and written to an Excel file at the end of the run, to be analyzed at a later time. The probes can then be moved out of the beam.

The program is structured as a finite state machine. It is at any one time in a single state, such as idle, going to the next position, or reading data. Every time the program receives a periodic interrupt, it checks the state of the system and then runs methods accordingly. For example, if the state of the program is moving to the next position, it will set the motor position, monitor the progress, and when the motor has stopped turning, change the state to reading data. This structure allows the user to easily abort the scan at any time, simply by changing the state to idle.

At each step in the scan, the program requests values for a set of devices, including motor position, beam cur-

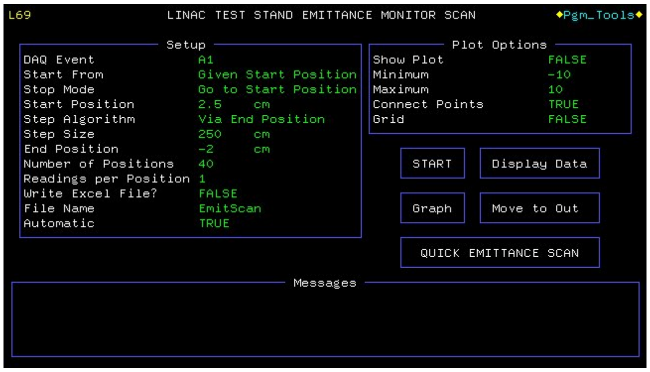

FIG. 3: The original emittance scan program. The user can specify which positions to scan through, and the program outputs wire voltages to an Excel file. The graph option outputs one pixel per (x, x') point, black if the voltage is over 10% of the maximum, representing beam, and white if below, creating a very coarse image of the beam.

FIG. 4: The main screen of the modified emittance program. The setup menu includes many more user inputs, including probe position, probe angular resolution, beam energy, and emittance cutoff.

rent as measured by a toroid, and wire voltages. These are stored locally in an array. At the end of the run, if the user has specified that a file should be output, the data is written to a new window, copied to an Excel file, and placed online through a built-in ACNET method. If the user forgets to specify this option before the scan, or if the scan is aborted, this data can be displayed but not easily written to a file.

The Excel file can be retrieved from http://www-bd.fnal.gov/excel and imported into a Mathematica notebook which calculates emittance according to equations 2 and 4. Because the voltage readings are proportional to the number of ions at a specific position and angle, they are used as a distribution function. The average position $\langle x \rangle$ and angle $\langle x' \rangle$, and higher moments, are calculated by weighting each angle and position by this distribution function and integrating, assuming that all positions and angles are equally spaced. These values can then be used to calculate the emittance and the Twiss parameters (α, β, γ). A user inputted energy is used to calculate the normalized emittance, and finally, Mathematica can plot density plots in phase space, overlayed with the theoretical elliptical distribution.

There are many shortcomings with this current procedure of measuring emittance. Firstly, a failure to click the "write to file" option will result in the loss of the data for that run. Likewise, a failure to change the filename before conducting another run will result in the overwriting of the previous run's data. Most importantly, the extra step of importing the data into Mathematica is inefficient and inconvenient. Most console computers don't have this software installed, so data analysis must be done on a separate computer after all of the data has been gathered. In addition, the notebook which does the calculation has a few different versions, and not everybody has access to the file. Ideally, the calculations would be performed online, just after data collection, so that any operator or person with access to ACNET can easily take an emittance scan and get analyzed results.

IV. INCLUSION OF ONLINE CALCULATIONS

I made a large number of modifications to the existing emittance scan program, as evident from figure 4, the foremost being the addition of online calculations. After the scan has been run, the program calculates emittance and the Twiss parameters, and generates plots, based on the functionality of the Mathematica worksheet. First, any voltages below 10% of the maximum voltage are cut out of the data set, and set to zero. This threshold cuts out background noise along the wires, with 90% beam inclusion being a standard [3]. This modified distribution function is used to calculate various moments, such as the average position of ions $\langle x \rangle$ and mean square angle $\langle x'^2 \rangle$. This is done using the numerical integration technique of rectangular summation, and assuming that the step sizes Δx are variable. These moments are then used to calculate the emittance and Twiss parameters, using equations 2 and 3.

Beam energy is not a simple measurement, and is usually done using separate equipment. The program allows for user input of the beam energy and uses this to calculate the normalized emittance, as in equation 4. The normalized and unnormalized emittances and the Twiss parameters are displayed for the user to see.

The program also has an option to produce graphs. When the user presses the "Graph" button, a new window opens containing a contour plot in green of the wire voltages in phase space. The Twiss parameters are used

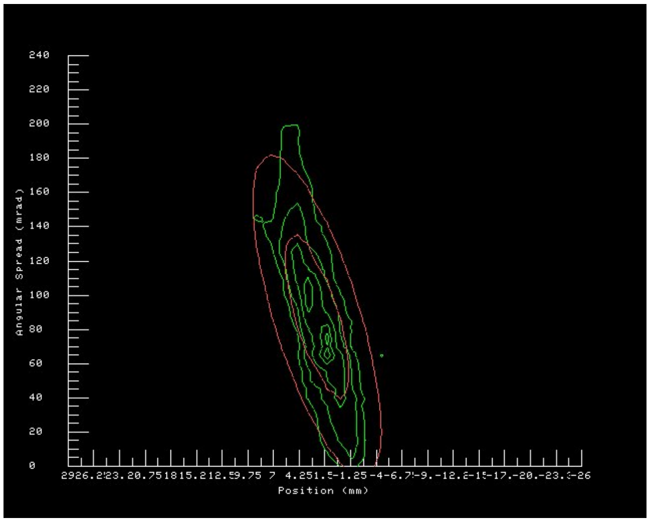

FIG. 5: A contour plot created by the program of a vertical scan of beam in the test stand. The green contour lines enclose regions of 90, 75, 50, 25, and 10 percent of the highest beam density. The red ellipses represent an ideal elliptical beam with the same emittance, the inner one containing one sigma of the beam, the outer one containing the entire beam, up to a cutoff.

to calculate the theoretical emittance as a function of x and x', as in equation 1, and contours of where these theoretical emittances match the RMS calculated emittance are overlayed in red, producing ellipses. An example of a graph made by the program is shown in figure 5.

A. Dealing with Bad Wires

Often, an emittance probe will have a few wires which are broken. This may be due to a bad connection, damage from the beam, a short with another wire, or another reason. The presence of a single bad wire does not warrant replacement of an entire probe, but it can affect calculations. The program allows the user to identify bad wires that they may notice after a scan, and redo all calculations ignoring these wires.

When the user presses the "Identify Bad Wires" button, a small window pops up with the numbers of the different wires in the probe, as shown in figure 6. Clicking on a number will toggle the wire's status as good (green) or bad (red). The next time data is analyzed, either by running a new scan or by clicking "Recalculate Emittance", this status is taken into account. For any wire marked as bad, the recorded voltages are thrown out. The voltage is instead given by a linear interpolation between the neighboring good wires. This new beam distribution function is used to recalculate all relevant beam parameters.

FIG. 6: The bad wires identification window allows the user to mark wires to ignore in data analysis by clicking on the wire number.

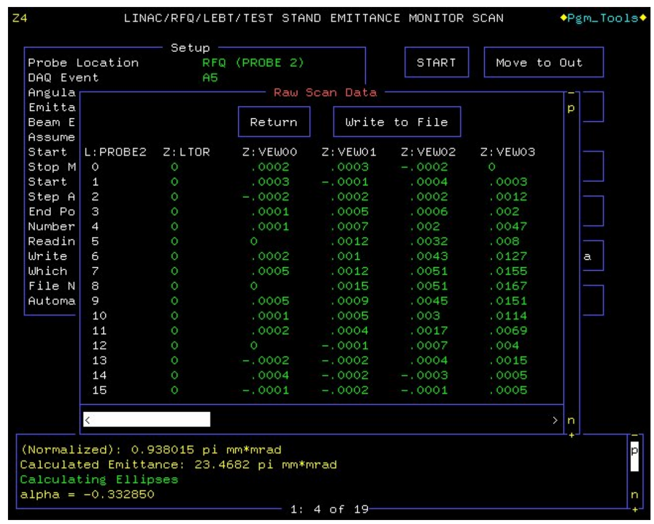

FIG. 7: The data table window showing simulated data. The user can scroll through data and view all of it, for either raw data, averaged data (if multiple readings are taken at each probe position), or analyzed data (bad wires interpolated, noise cut out).

B. Other Features

The new program has quite a few other features designed to make scanning emittance easier. The first is the ability to scan for the position of the beam, if the range of positions to scan is unknown. This can either be done by pressing "Quick Emittance Scan" or setting the start mode to "Scan for Position". The program runs a very coarse scan, with low position resolution. Any resulting voltages less than 5% of the maximum (a lower threshold than that used for the emittance calculations)

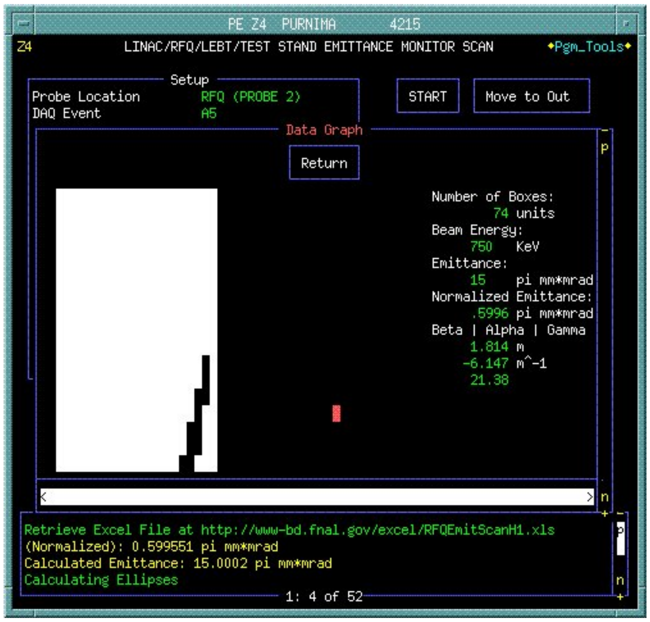

FIG. 8: Displayed outputs for the horizontal distribution (RMS emittance, Twiss parameters)

are cut out, and the first and last positions containing any wires with non-zero voltage are identified. The program then runs a finer scan between these positions. This process may not work properly if bad wires are not properly labelled.

Additionally, the user can inspect the data manually by clicking "Display Data". This brings up all collected data in a table format, as shown in figure 7. The user can look at the raw data or analyzed data, which is after interpolation of voltages for bad wires and after cutting out the background 10%. Any of these can be written to an Excel file by clicking "Write to File".

Last but not least, the program can control one of multiple probes. The user has a choice between the test stand probes, Probe 2, which is installed during testing of the RFQ, and Probes 3 and 4, which are currently installed in the Linac beyond the Cockroft-Waltons. Because the angular resolution of probes is adjustable, this variable is also user-modifiable.

V. RESULTS

The first step in testing this program was to test the algorithms used to calculate emittance. This was done by calculating voltages (up to a scaling factor) that would be caused along the wires by a beam with a two-dimensional Gaussian distribution in phase space. These voltages were then treated like real data, and the theoretical and calculated emittances were compared.

The majority of testing of the algorithms was done using the test stand. The stand was used to test the motion of the probes, the quick scan function, and bad wire identification. Values and plots outputted by the program were compared with output from the original Mathematica notebook.

After testing its accuracy, we used the program to measure both horizontal and vertical normalized emittances just downstream of the RFQ for various power settings. When 208 kW were put into the RFQ, this produced a beam of 38 mA, with a horizontal emittance of 0.60π and a vertical emittance of 0.44π mm·mrad. At a slightly lower power of 178 kW into the RFQ, the beam had a horizontal emittance of 0.60π and a vertical emittance of 0.45π mm·mrad, all normalized. The outputs for these scans are shown in figures 8, 9, and 10.

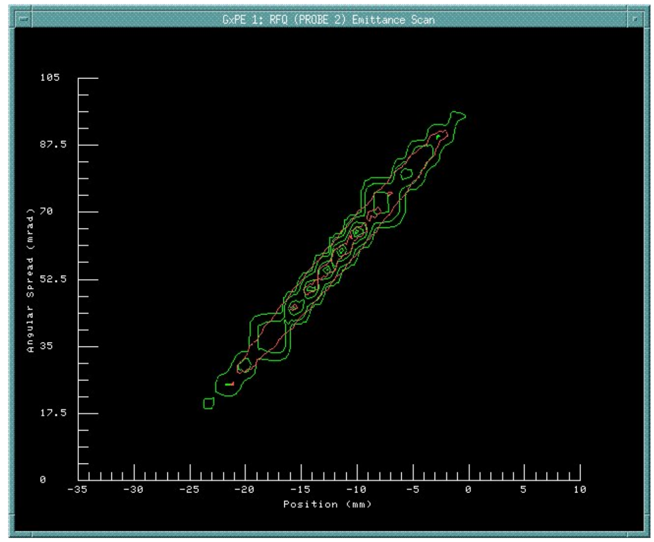

(a) 208 kW through RFQ, with a normalized emittance of 0.60π mm·mrad

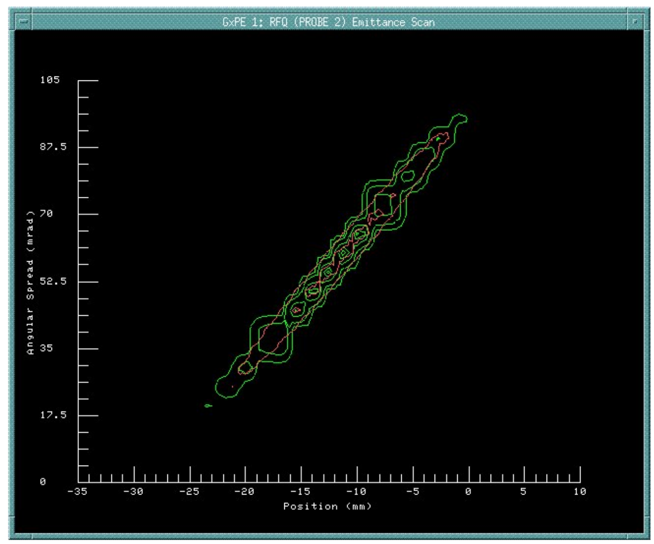

(b) 178 kW through RFQ, with a normalized emittance of 0.60π mm·mrad

FIG. 9: Program output for horizontal scans with various power inputted to the RFQ.

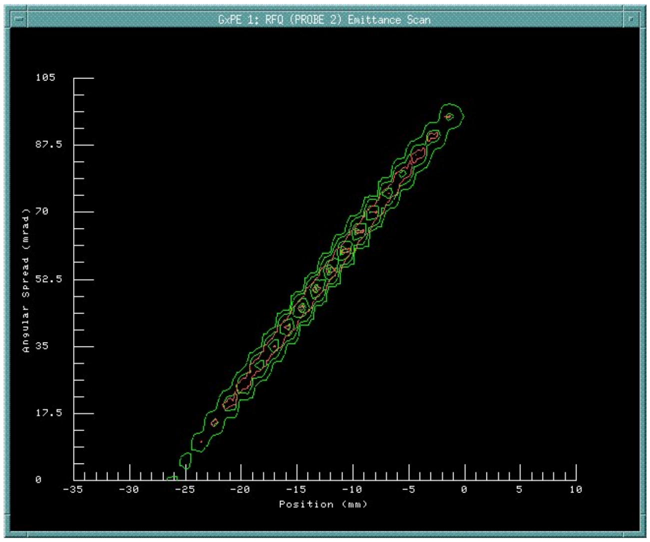

(a) 208 kW into the RFQ, with a normalized emittance of 0.44π mm·mrad

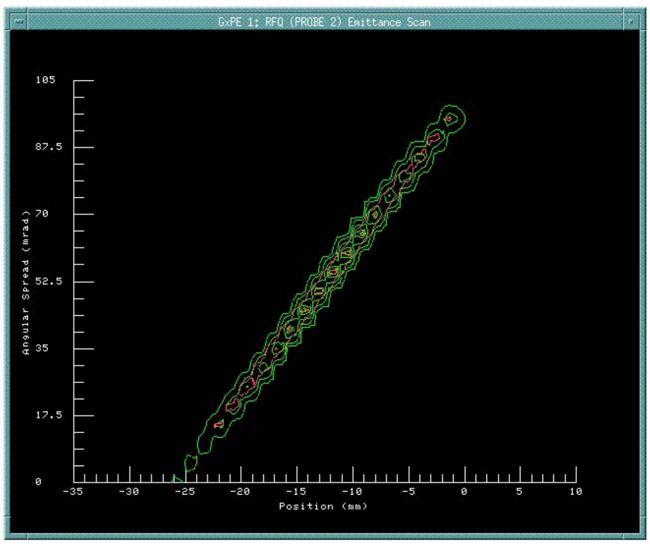

(b) 178 kW into the RFQ, with a normalized emittance of 0.45π mm·mrad

FIG. 10: Program output for vertical scans with various power inputted to the RFQ. The vertical beam distribution is slightly larger for lower power.

VI. CONCLUSIONS

In order for the RFQ to be installed, the emittance must be checked to make sure that the beam can be accepted by the Linac. The current method of emittance analysis is done offline, is not realtime, and can be very inconvenient. This method was improved by heavily modifying the current emittance program to include emittance calculations as well as plot generation. This greatly sped up emittance measurements of beam at the end of the RFQ.

Normalized emittances of the beam were found to be between 0.4π and 0.6π mm·mrad, meaning the beam is small enough in phase space to be accelerated by the Linac. More work still needs to be done on improving beam transmission efficiency and reduce the power needed to run the RFQ before it can be installed, but once the upgrade is complete, the main use of this program will be to perform quick scans during maintenance or tuning. The program can control multiple emittance probes, so there is no need to modify it for this purpose.

ACKNOWLEDGMENTS

This project was carried out during the summer of 2012 as part of the Lee Teng Internship at Fermilab. I am very grateful to my mentors, Dan Bollinger and Cheng-Yang Tan, for entrusting me with a part of their critical work upgrading the Linac injector. Their support and guidance throughout this summer has taught me a lot about this field, in addition to my specific project. I would also like to thank everyone else involved with the RFQ project, for patiently sharing their experience and knowledge with me. I have enjoyed working with everyone for the last ten weeks and wish them all the best.

I learnt a great deal about accelerator physics from the introductory class at the US Particle Accelerator School, and thank the instructors, Bill Barletta, Linda Spentzouris, and Elvin Harms.

Lastly, I appreciate the time and commitment of Eric Prebys and Carol Angarola put into organizing and smoothly running the internship program, and thank them for giving me the opportunity to work at Fermilab.

---

[1] D. A. Edwards and M. J. Syphers. *An Introduction to the Physics of High Energy Accelerators*. Wiley-VCH, 2004.

[2] C. Y. Tan et al. The 750keV Injector Upgrade Plan. Technical report, FNAL, November 2011.

[3] Martin P. Stockli. Measuring and Analyzing the Transverse Emittance of Charged Particle Beams. In *Beam Instrumentation Workshop 2006: Twelfth Beam Instrumentation Workshop*, volume 868, pages 25–62, 1-4 May 2006.

[4] W. Marsh. Linac Test Bench Emittance Monitor Scan. ACNET computer software, 2009.