

Simulations and Testing of Retarding Field Analyzers for Electron Cloud Monitoring

Lee McCuller

07/02/09

Abstract

My summer internship project was to develop a test stand and perform simulations to understand the response of retarding field analyzers (RFAs) used for studying the electron cloud, an intensity dependent cause of instability. My primary goal was to build a test stand to emulate measuring the E-cloud using an electron gun, except in a controlled environment. The purpose of these measurements is to understand the RFA response to incoming electrons of given energies, while accounting for tray fields and geometric effects similar to a beamline environment. Simulations were also enhanced to give deeper understand. The project culminated in results tying the experimental and simulation measurements.

1 Forward

This work is done as part of the Lee Teng internship at Fermilab with the help of my supervisor Robert Zwaska. I am also getting support for my project from Kevin Duel who helped us decide on and draft the test vessel design, and from Cheng-Yang Tan who helped with the ecloud simulations and RFA studies.

2 Introduction

My summer project was to assemble a test stand for new Retarding Field Analyzers (RFAs) to be used for electron cloud (ecloud) studies in the Main Injector. During machining of the stand, I worked on enhancing simulations of the RFA response using SIMION. This duty occupied most of my work time, but in the last week I managed to get test stand data to compare with simulations.

The ecloud is a phenomenon in positively charged beam-lines. It is a gas of electrons that is formed from a cascade of secondary emission electrons from the beam pipe. The process starts from the residual gas in the chamber; the beam passes through that and knocks off electrons from the gas. These electrons are accelerated by their attraction to the beam, and these accelerated electrons then impact the beam pipe with enough energy to eject more electrons through secondary emission, causing an exponential generation of electrons. This cascade reaches equilibrium when space charge from the gas of electrons cancels the attraction to the beam. Space charge forces act also on the beam, causing undesirable effects. The electrons push the beam around and lead to head-tail and multi-bunch instabilities. See figure 1 for a graphical description.

This phenomenon is affected by several beam parameters. The heating of the electrons is sensitive to the flight time of the electrons and to the bunch lengths and bunch spacings of the beam. The generation of electrons is sensitive to the energy gain and the surface properties of the the beam pipe.

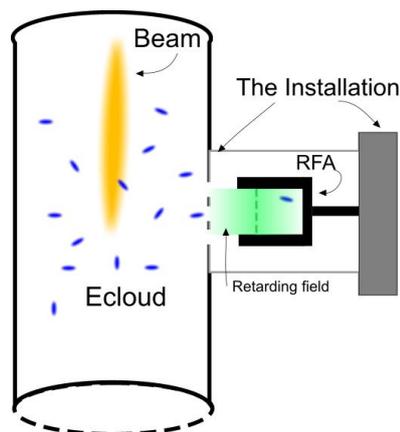


Figure 1: Depiction of electron cloud in beam pipe with RFA installation

And finally, the space charge limitations are affected by beam intensity. Since these factors affect an exponential process, small changes can have effects with orders of magnitude differences. The intensity dependence is particularly worrisome in this process, since there is a risk that the ecloud might limit the maximum intensity in the project X upgrades.

The ecloud has been seen at many of the high energy accelerators and it has been studied in several machines to help characterize it. Unfortunately since it is such a nonlinear process that depends on the accelerator geometry, as well as beam parameters, it behaves uniquely in each machine. Simulations have been done for the main injector we need real data to ensure our understanding of the phenomenon.

2.1 RFA Assembly and Installation

The retarding field analyzers are sensors that can detect current from a beam of charged particles. They consist of a collecting plate behind a wire screen. The wire screen is set at a retarding voltage to act as a high pass filter in the energy of passing electrons. Our installation consists of 4 RFA's in the Main Injector. 1 is an old design from the Advanced Photon Source that has already seen use, and the remaining 3 are a new design that I studied for my project. They will be placed to give data in two pipe sections: one section with a graphite coating to reduce secondary emission, and the other as a control section without coating. The newly made RFA's will also have the coating applied to them since their reading is dependent on SE characteristics of their electrodes.

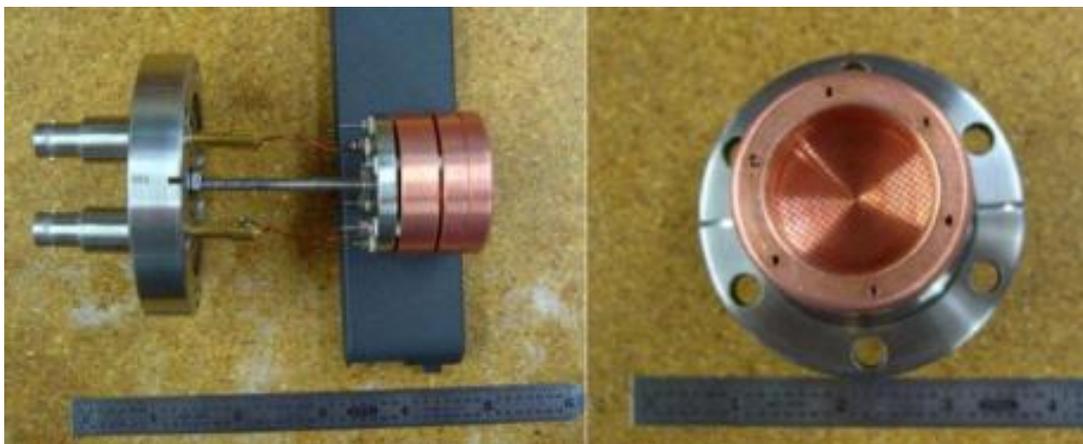


Figure 2: Side and top views of the new RFA design. This picture comes from a test RFA without the graphite coating

The installation of the Argonne RFA taught us several features necessary for good ecloud data from the Main Injector. Whereas the Argonne design uses two grids for field shaping, the new analyzers use a cupped shape around a single wire grid. This gives a natural defocussing-before, focusing-after effect that should reduce the effects of secondary emission on the collection current and reduces losses due to impacts with grid wires. The original installation also had a substantial amount of noise in the output signal, which we hope to prevent by adding a low-pass filter and amplification applied to the signal in the tunnel. The filter will reduce noise from the beam RF and the amplification will reduce line noise in transmission.

3 RFA Simulation

In order to develop a new RFA with characteristics suited for this application, the program SIMION was used to model the RFA geometry and associated fields and to develop a phenomenological model for the system. This program lacks features such as secondary emission electrons and certain data taking features that would be useful for this and future projects. I have added some features to more accurately model the RFAs using this program.

SIMION works in a very modular system. The entire simulation is encapsulated in a workbook, which holds the main coordinate system. Inserted into the workbook are Potential Array (PA) files which are each a grid of cells containing either electrodes or free space. The electrodes are set at a voltage and SIMION 'refines' them to generate the potential at each point in space using an iterative solution to the Laplace equation. These PA files define the electrode geometry of the simulation and can have reflection or cylindrical symmetry applied to them. The workbook represents potential arrays in space with their symmetry applied as well as with arbitrary scaling and orientation transformations. Because solutions to the Laplace equation can be superimposed, SIMION allows connected electrode regions to be adjusted in real time which it calls "fast adjusting". This method is useful with user control.

The RFA simulation uses a workbook in this manner, but requires some caveats in order to simulate a realistic grid. This is necessary since the grid requires an extremely dense array to describe its geometry due to its thin wires. With the wires only 5 cells thick, and the grid only occupying 1/16th of the true grid area, the PA file is 100mb. The trick used is to transport electrons from an ideal grid region in the RFA PA to a smaller non-ideal grid in a virtual mosaic. This mosaic transport remembers the electron's offsets so that they may be returned to the RFA. Toggling this transport can establish results with ideal and non-ideal grids.

My work was to enhance the user program to manage each particle and to add secondary emission capabilities. The modification design is highly coupled to how SIMION fly's particles through a workbook. The SIMION user-program model allows the program to be as passive as possible. The user program is a script written in LUA that registers a series of callbacks that SIMION calls at various stages of the electrons' lifetimes. Difficulties arise because SIMION treats particles in batches, even when it is actually simulating particles one by one, so advanced user programs need to have additional structure to account for this batch mode of operation. User program callbacks are called on a per particle basis and each of SIMION's callbacks have linked variables registered that correspond to particle parameters such as location or velocity.

SIMION's order of operations is to first read a .fly2 file which describes the batch of particles to fly including its size, starting location and velocity distributions. Then SIMION calls an initialize function on each of the particles. This function can modify energy or location and essentially override the characteristics that the .fly2 file had determined. Here one can check for the first particle and call a beginning of fly function or "global initialization" type function. After particle batch initialization, The user program is then allowed to do a fast-adjust on the potentials before the fly (they can also be done during the fly but this simulation uses static fields). After thorough initialization, the ions go through a flying loop. SIMION internally adjusts their position and velocity using Runge-Kutta integration over the EM field forces. After that, it checks for collisions and calls the "other_actions" function. Here the user program applies the transport trick and checks for electrode impacts to generate secondary emission electrons. SIMION then allows the user program modify the electric and magnetic fields experienced by ions.

After all of the particles have impacted a surface, SIMION calls "terminate" on each particle. Our program does not need this for each individual so it waits for the last particle and calls a global "fly completed" type function. After this is called SIMION checks if you want to begin another fly and repeats this process (so it uses the same .fly2 and the user program is responsible for changing particle parameters between flies).

3.1 Details of Added User Program Enhancements

The developed user program has added two abstract components to manage particles. It has a particleRun manager which stores all the different parameters that you want a data-point to apply. This includes the magnetic field magnitude and direction, the electron energy, parameters related to secondary emission, and the number of particles to run at these parameters. This manager is generated and called directly by the SIMION callback functions to make adjustments. The particleRun class also stores a queue of future particles to generate from secondary emission. This queue is necessary since the batch mode operation does not allow one to generate SE electrons as they are created. The queue is filled and SE electrons are created at initialization in the next fly. For individual particle data, the particleRun class includes a list of all currently flying particles. This avoids the anonymity that the batch mode running enforces by allowing individual tagging of particles. Particle tagging gives records of where particles impact, among other statistics. It is also ultimately used for storing transport parameters and

SE data.

When the current particleRun has taken enough data, control moves to the other abstracted component, the runManager. This class has the duty of recording the data stored in a particleRun to a file, and it has the duty to start new runs with incremented parameters (such as particle kinetic energy). In implementation, the runManager class uses a LUA type iterator to feed it new particleRun instances.

The choice to use iterators to generate new runs allows one to control the parameters without needing to modify the particleRun or runManager classes themselves. The user-program is given a script file to execute, which then gives it the iterator to use for the runManager, allowing more control of automated data taking.

3.2 Details of Grid Transport Trick

As described above, the user-program employs a trick in order to reduce the memory requirements of a full size realistic grid. To implement this trick, the workbook employs a series of "dummy" PA's. As a particle is flying, it knows which PA it is currently being influenced by. The user-program checks this variable to see if the particle has entered the kick out PA. If so then it employs some modular arithmetic to determine where in the grid it should be transported to. This math effectively tiles the grid across the RFA. The offset is recorded to the particle, once moved, and the particle flies in the grid PA as normal. the grid PA is enclosed in a return PA, which the transporter checks for. Once the particle enters the this PA it is retranslated by its offset.

In the flying view of SIMION, one can see that this does indeed preserve flight continuity. There are some catches to this trick though. Since the grid is of limited size. The return PA section must also check to make sure that the particle has not simply escaped the transverse boundaries of the grid PA. if it has, then it is transported around the grid and its offsets are adjusted (This oversight led to a reduction in the perceived efficiency of the RFA under a non-ideal grid in Tan's original implementation). The other catch is that the transporter must also check that it is not returning the particle to a spot inside the RFA wall, since the walls are not present in the grid PA.

3.3 Details of Secondary Emission Implementation

All of the abstractions added above were implemented so that secondary emission could be modeled with SIMION. SE is difficult to model in SIMION due to several program limitations. SIMION's representation of the electrode surfaces is the minimum necessary to do electric field calculations. It does not store metal surface properties and, more importantly, it has no knowledge of the surface normal. When a particle impacts a surface, our user-program detects this, and knows the location and velocity of the impact. A phenomenological SE yield model is employed to determine the number of expected particles generated.

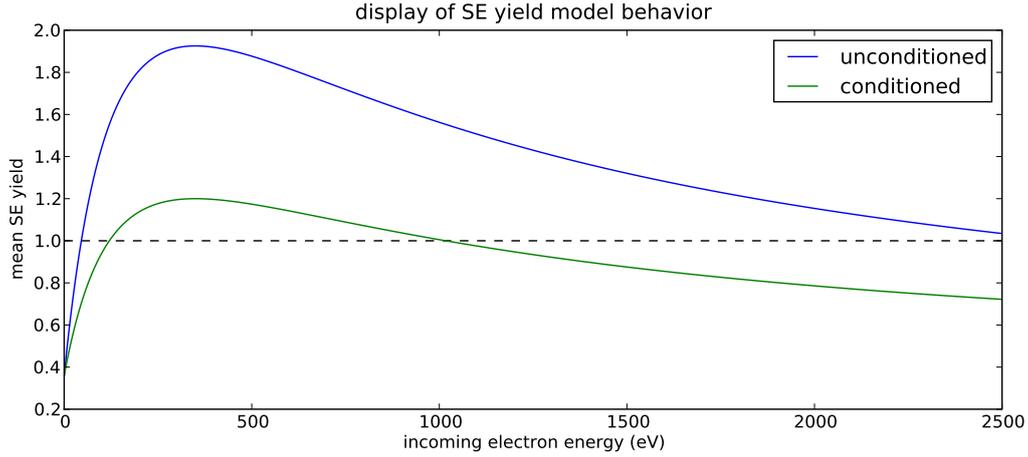
The SE yield model is based on a function fit from spectroscopy data taken by Bob Zwaska and others. The formula for this fit is

$$Yield = Y_0 + (Y_{max} - Y_0) * \frac{4 \frac{E_{In}}{E_{max}}}{\left(1 + \frac{E_{In}}{E_{max}}\right)^2}$$

where Y_0 is the yeild at 0 incoming energy, Y_{max} is the maximum yield, and E_{max} is the energy at maximum yield.

The SEY model fits the spectroscopy data extremely well and we have fit parameters for treated and untreated stainless steel. For simulation purposes, this model produces a random number of SE electrons with a Poisson distribution and mean equal to the yield number. I used an algorithm by Knuth to produce numbers with a Poisson distribution of a given mean. Figure 3 shows the response curve of this model and gives the parameters used in the simulation. For all of the simulation data included in this paper, the unconditioned steel parameters were used.

Each SE electron created is colored differently for inspection purposes, and for each SE electron to generate, we need to know its outgoing energy, and its outgoing direction. The outgoing direction has a uniform distribution over the 2π arc-radian hemisphere away from the surface. this distribution is generated in two parts, determining the surface normal, and applying that distribution to the normal.



	Unconditioned Stainless Steel	Conditioned Stainless
Y_0	.36	.36
Y_{max}	1.9255	1.2
E_{max}	350eV	350eV

Figure 3: Here we have a display to understand the response of the SE yield model. We see that there is a large region in which the expected yield is greater than one. This will show to be important since SE contributions to signal current can override the beam absorption contribution.

Since SIMION represents its surfaces in a grid, generating a realistic surface normal is a nontrivial process. From boundary conditions, The electric field off of the surface is always normal to the surface. This fact is used to approximated the normal in this finite element model. SIMION's scripting language LUA exposes a lot of particle data but does not, by default, have an interface for finding the electric field in PA's. Our program uses a current SIMION beta release and activates "early access mode" to gain this interface in our LUA user-program.

Each PA has a unique grid size, which sets a scale across which the electric field should be averaged to get the normal direction. Our computation works in PA local coordinates to naturally work across this scale. It takes the electric field in the 27 cell, 3 by 3 by 3 cube of PA grid cells centered at the impact point. It ignores those cells that are electrodes and the remaining values are averaged. This vector is then set to unit length to give the normal.

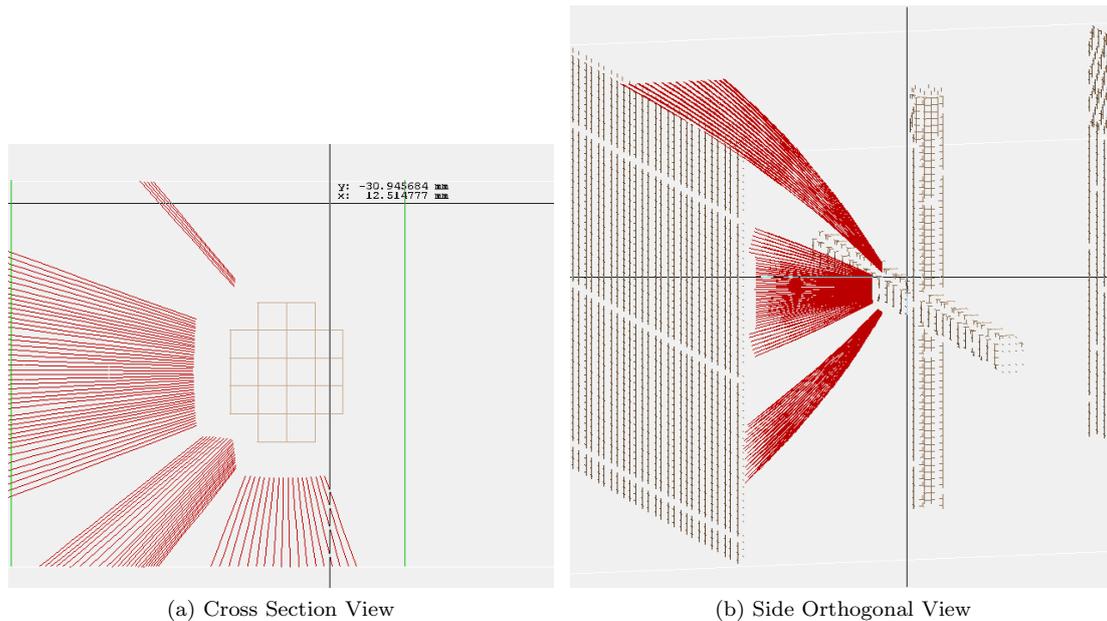


Figure 4: In each of these figures rays are shot from the normals computed using the described method. Ideally, the rays would appear in a perfect circular fan; However, grid computations put errors at the surface. Figure (a) shows where SE starting positions will be projected out from the electrode initially in their flight. The value is such to prevent immediate re-impacts with the surface.

Sign checks are made since the electric field lines can be pointing inside or outside the surface. The Runge-Kutta flying algorithm often has particles that impact just below the electrode surface, in which case the SE particle can't escape. Checking for this is most convenient in these methods while PA coordinates are available, so that the surface normal computation can also flag to toss out SE for the impact.

Once the normal, N is found, the distribution around it must be implemented. This is done by creating a uniform distribution of vectors in the positive X-axis hemisphere, and then rotating these so that the X-axis points in the direction of the normal.

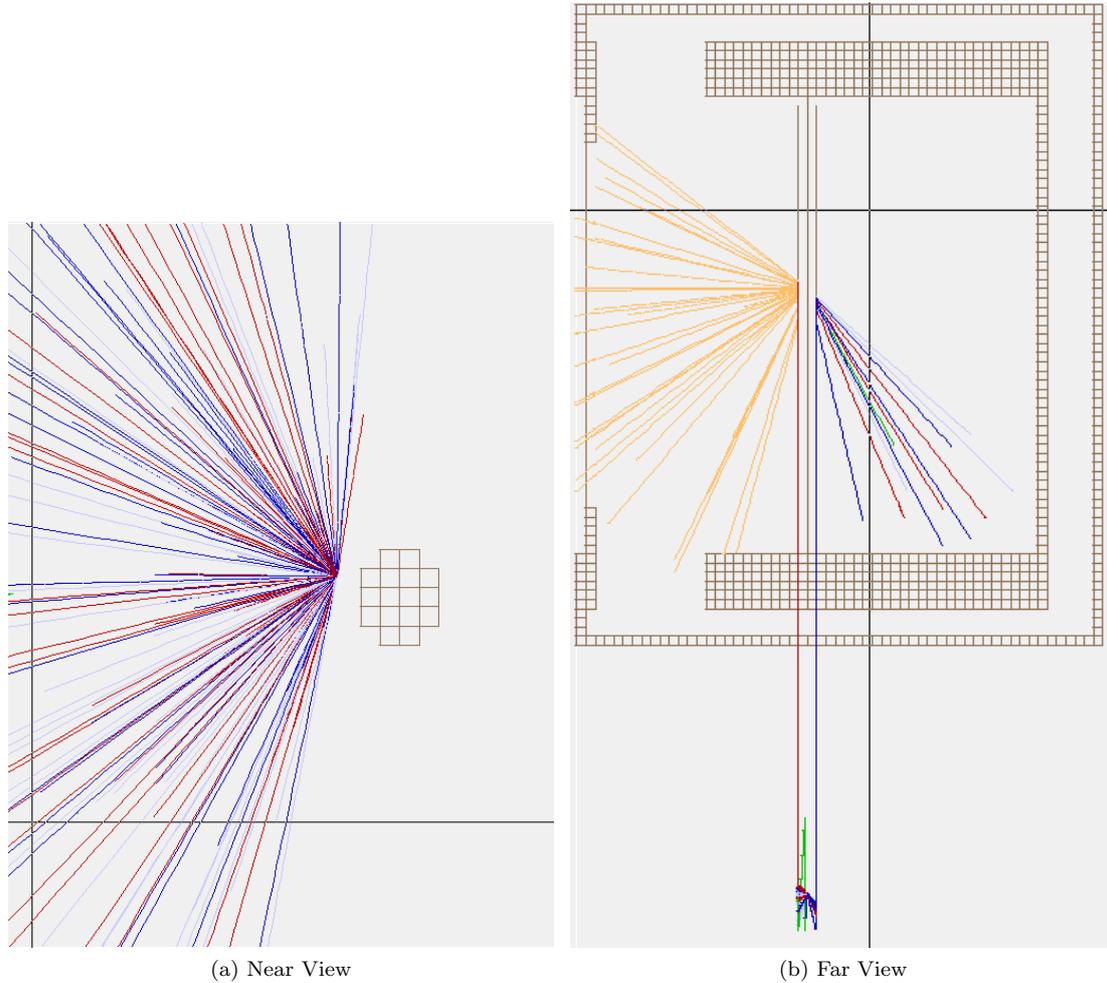


Figure 5: Shown here are rays shot in a uniform hemispherical distribution to validate the analysis over the distribution functions.

To generate the hemispherical distribution, I used spherical coordinates and the inverse PDF method for generating θ and ϕ . The differential area element for integrating a spherical surface in these coordinates is $r^2 \sin(\phi)$, with $r = 1$ for a unit sphere. With the inverse PDF method, uniformly distributed numbers from 0 to 1 are passed through $\arccos()$ to get the proper distribution for ϕ . θ can be generated from a uniform distribution from 0 to 2π .

$$V = \begin{bmatrix} \cos(\phi) \\ \cos(\theta)\sin(\phi) \\ \sin(\theta)\sin(\phi) \end{bmatrix}$$

These parameters are plugged into this vector formula to get their linear coordinates. I use a rotation matrix to point the x-axis through the direction of the normal. To generate this rotation matrix I arbitrarily choose a vector orthonormal to N (by subtracting 1 to N_z , projecting out N_z and renormalizing). I then use a cross product to get a final orthonormal basis-set, with which to construct an orthogonal matrix. In reality, such a matrix has two vector degrees of freedom, a "looking" direction which we give, and an orientation direction. Here the orientation direction is ignored since the components that it affects already have a random orientation.

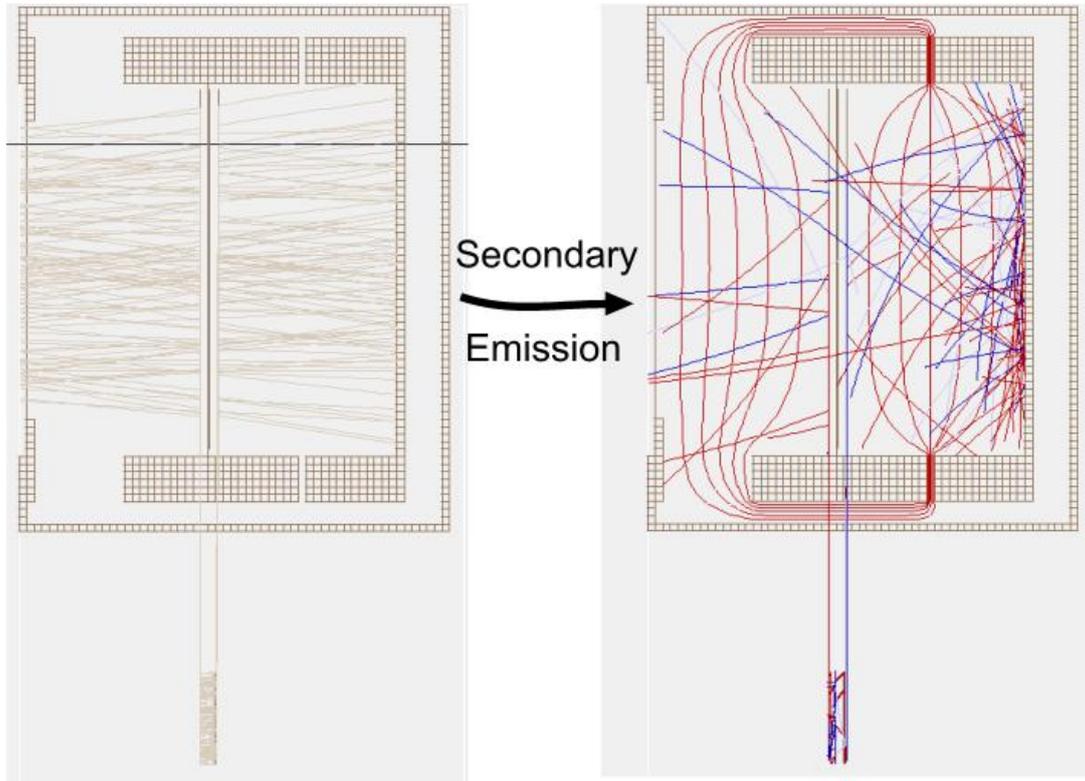


Figure 6: Here I show two batches, the first are 100eV electrons passing through a -10V grid potential with ease. These form SE electrons which run in the second batch as shown.

Future enhancements to the SE system could include recording all impact points for a set of flies. This would allow us to run simulations and to immediately see the effect of differing SE model parameters, without the overhead of flying beam particles. As it stands we have the multi-batch system shown in figure 3.3

4 Simulation Data

Running the simulation was a good first step before taking test stand data. It gave a qualitative understanding of the nature of the results to be expected. The ability to visualize the electron behavior in the RFA is extremely useful for understanding real signal output.

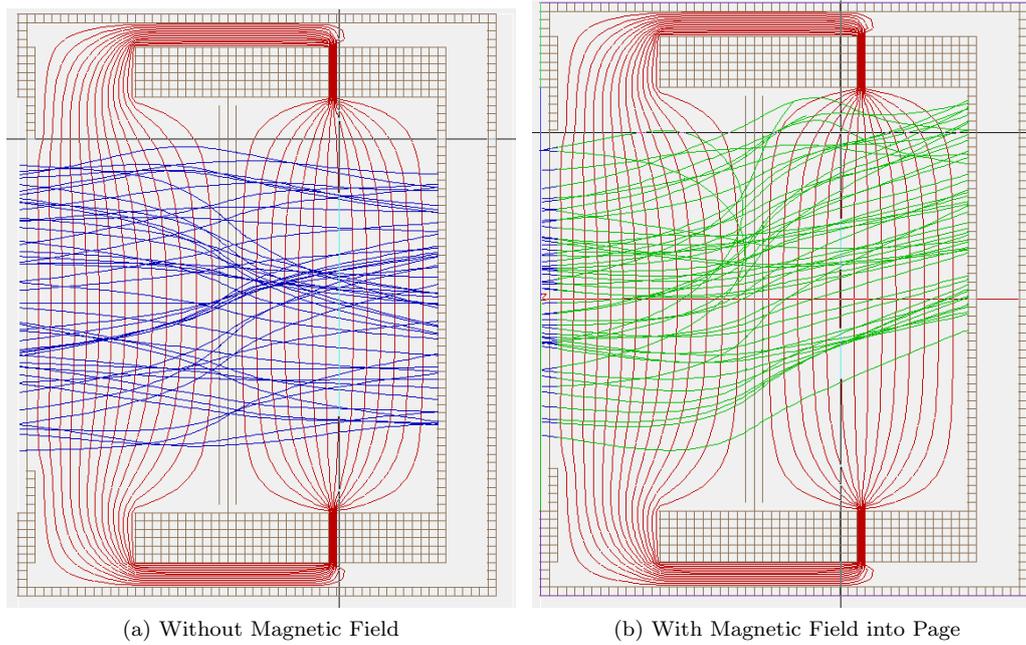


Figure 7: Here is a cross section view of SIMION flights with and without magnetic field applied. Grid transport has been disabled here for ease of view. The electrons here are slightly higher energy than the grid potential, showing the focusing effects during flight.

The simulation data also gives a profile of the capture efficiency of incoming electrons. Tan used this data to optimize the analyzers for high signal, but my use has been to produce efficiency curves across electron energy. These profiles are necessary since an RFA is not a perfect high pass analyzer. Furthermore, External magnetic fields can greatly affect the response of the RFA by changing the path of the electrons, as demonstrated in Figure 7.

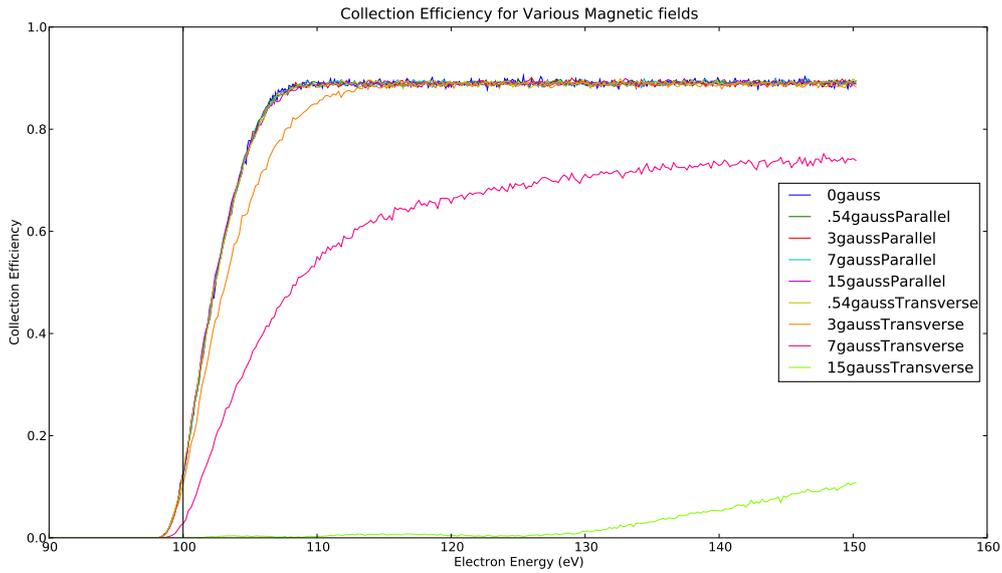


Figure 8: Efficiency profiles across grid voltage for 100eV incoming electrons. Notice that only transverse (into page) magnetic fields seem to have any effect on the signal and that even small fields can substantially disturb the profile.

Another concern to address is how having a potential on the collector changes the response curves. Ideally, it should have no effect, but magnetic fields and secondary emission make could make it useful.

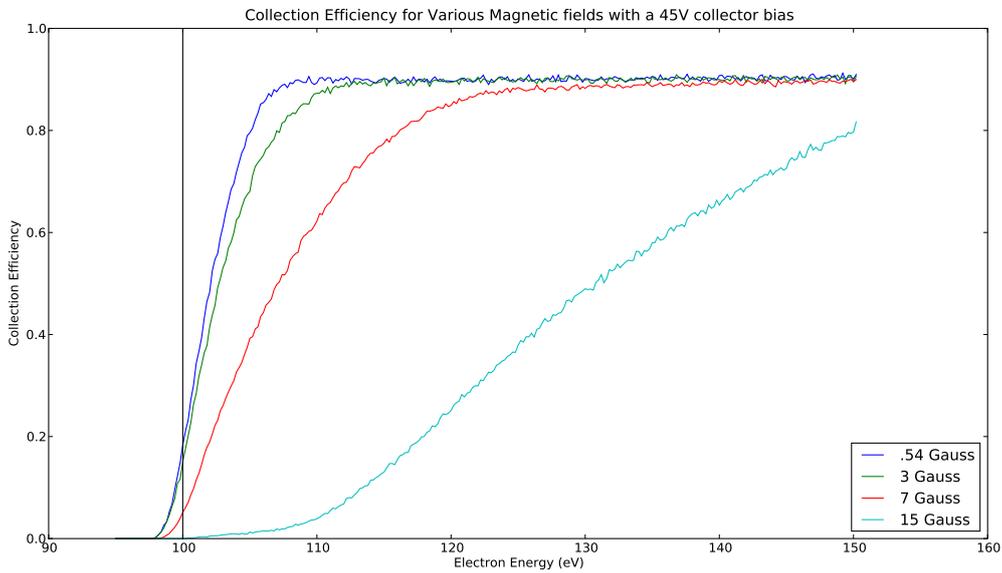


Figure 9: Efficiency profiles across grid voltage for 100eV incoming electrons and an added 45V bias on the collector. Transverse magnetic fields still affect the signal, but the bias greatly reduces the effects of high fields.

Scans that include secondary emission are included later to be compared against real data. More

of these profile curves should be done to characterize the effects of magnetic fields. Secondary emission was not included in these results, but would be useful to have to see how much these curves change, especially at low beam energy.

5 Test Stand Setup

The primary focus of this project is to develop and use a test stand for the RFA. We are using an electron gun to provide a beam of electrons at a known energy and spread to compare gun emission current to RFA collector current. The primary difficulty in this test stand will be preparing a vessel to house the gun and RFA in vacuum.

The vessel is designed to emulate the beam pipe effects on the ecloud, as well as to facilitate measurements. It consists of a central pipe with end flanges, one for connecting the gun and the other for the RFA. In order, there are several components along the pipe. First is an extrusion to attach vacuum pump, second is an angled pipe with flange for a viewing window. This angled pipe meets at the vessel along with another pipe with a flange to attach a mount for a phosphor screen, to be viewed in the window. immediately after these connections is a grill visible in Figure 10. This grill will also be present on the RFA mount to the beam pipe to prevent impedance. For the vessel, the screen has the advantage that the electron gun should be calibrated using the light emitted from electrons off the screen, so the phosphor screen should not be a requirement (though we now know that this was wrong).

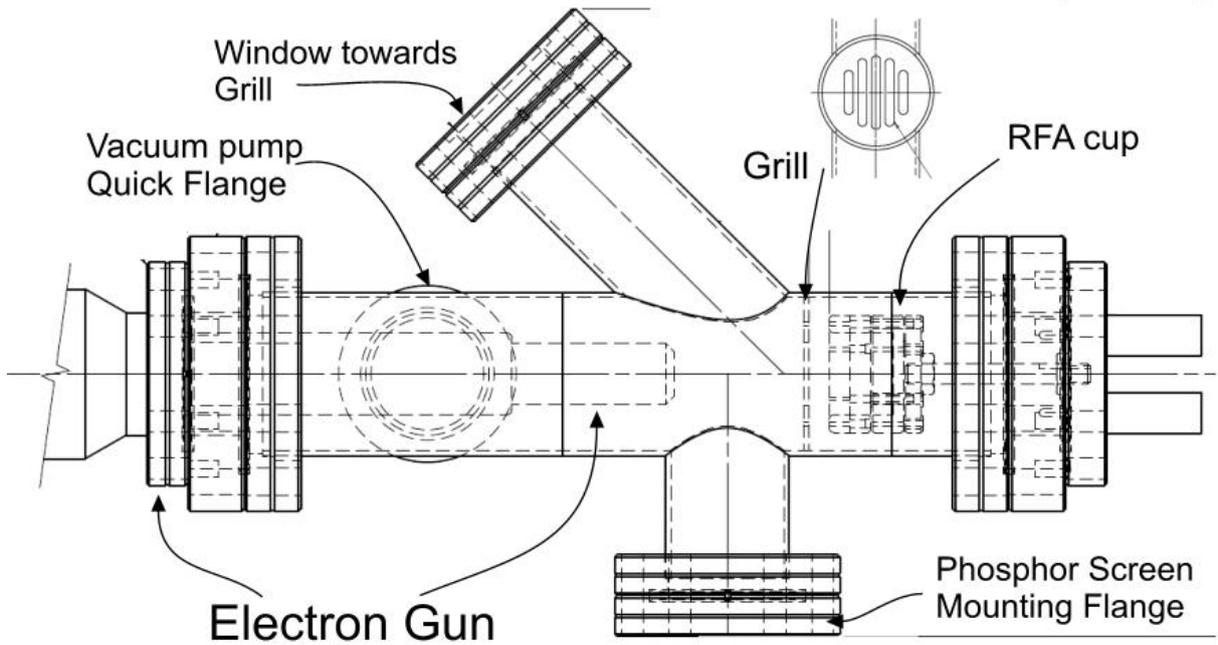


Figure 10: Annotated view of vessel design.

6 Procedure

The measurements taken by me for the coated RFA installed in the test stand are a first run set of data. As such they tell us more about the calibration stand than they do about the RFA. We found that our plan to use photon emission from the grill for gun calibration will not work. It turns out that

the hot emission cathode produces enough light projected on the grill to completely mask the glow from electrons.



Figure 11: Here is my lab with the setup used. Notice the pump in the lower left, feeding to a valve with pressure gauge and vessel. The gun electrons and RFA power supply are on their respective sides, along with an oscilloscope measuring collector signal. Foil was needed to shield splits in the wire due to the noise picked up from the TEL magnets.

When we first took data we did the following procedure: We set up the electron gun for various energy configurations. The gun manual has two useful items in it or this. There is a section for normal operating procedure, which we followed in applying the settings to the gun to ensure safety and to model the distributor's usage. The manual also includes a set of graphs of source voltage to emission current in the back. These were valuable since they give the settings for the grid and anode voltages for a given electron energy. On the first run we needed these predetermined settings in lieu of measuring the beam profile with the grill.

With the gun electronics set, we connected the RFA through to an oscilloscope to measure the signal. This signal measurement was compared against the emission current to gain an understanding of the efficiency. The initial signal readout gave that the grid and collector current were similar. With 200eV electrons, waving a 5 gauss magnet near the vessel produced a largely anti-correlated signal from the electrodes. At 2keV the magnet gave highly correlated changes in the signal. We know from simulation data and future measurements that with no potentials applied to the collector or grid that the signal is strongly affected by secondary emission off of the grid and collector. This anti-correlation could be explained by low energy SE electrons being diverted back to the collector by the magnet rather than impacting the grid.

After these initial explorations, we used two nine volt batteries to apply voltages to the grid or collector. In both cases the incoming electron energy was too high to be greatly affected by an 18 Volt potential, so this exercise was to characterize how the signal would change with a potential pushing or pulling at secondary emission electrons.

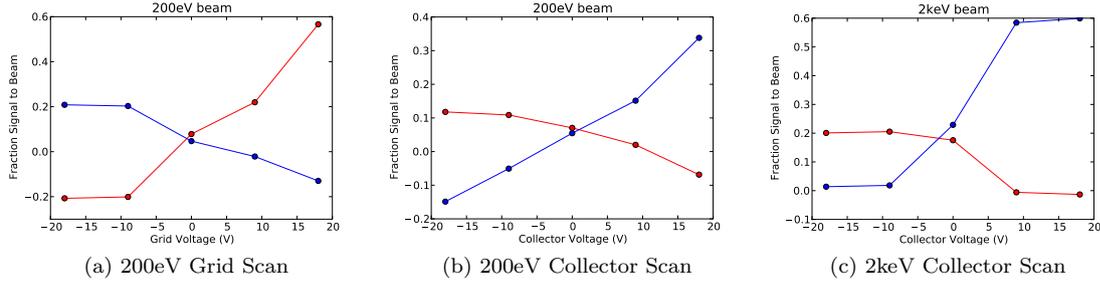


Figure 12: Here we see the results of the scans with the battery. For each graph, Blue indicates signal from the collector whereas red is signal from the grid. In part (a) The battery voltage was applied to the Grid, but in parts (b) and (c) the voltage was applied to the Collector. The negative values shown demonstrate the significance of secondary emission in the signal, since it can be greater in magnitude than the original signal. Other conclusions are difficult to draw, but from the change in the Grid current we get an idea of how much grid SE contributes to collector current with a potential applied.

After these preliminary tests, we had the confidence (and urgency) to go ahead and do a profile of the signal current across the retarding voltage. At this point we had installed a high voltage power supply that could provide a potential to the grid. Along with this I decided to take more methodical calibration data to ensure that our measurements would be good. For these measurements I decided to use the Emission Current Control feature of the gun electronics box. This feature allows one to set the emission current and the box will use feedback circuits to ensure the source voltage maintains the current. The ECC allows me to use the signal current for direct comparison without having to first correct for changing emission current. For these measurements I decided to use a 100eV beam, which does not have a preconfigured parameter set from the booklet.

To set up a good 100eV beam, I set my emission current and followed the book procedure for the other parameters. The ideal grid voltage is one that minimizes the source current to emission current ratio. I discovered that having a grid voltage of zero did this, despite the high variability the electronics allows for this parameter. I also found that the signal seemed independent of its value. Next, I optimized the first anode voltage. This anode also helps reduce source current, but it also affects the beam characteristics. The gun reference warns that it must be high enough to give a well behaved spot size. From stepping through values I gathered that 140V gave the strongest signal from the RFA and I didn't worry about the source current changes (they stayed roughly constant past 100V).

The final diagnostic scans to be done were over the focus and over the deflections. I did a quick optimization of the focus first, then the deflection scans, so that a final focus scan would be centered in the RFA collector (which we assume is the point in the deflections of maximum signal). For these diagnostics, I took the signal data with -90V on the grid, which should be safely on the high signal plateau.

7 Final Measurements

The diagnostics gave me confidence that I could produce meaningful data. The first dataset comes from a diagnostic tool, the deflection scan over the X direction. Figure 7 shows the results of the scan. The results are very pleasing to see since we can not directly visualize the beam. The multiple peaks tell us that our beam width is relatively small. Since there are fewer peaks than there are grill lines, the beam is likely larger than the grill line width.

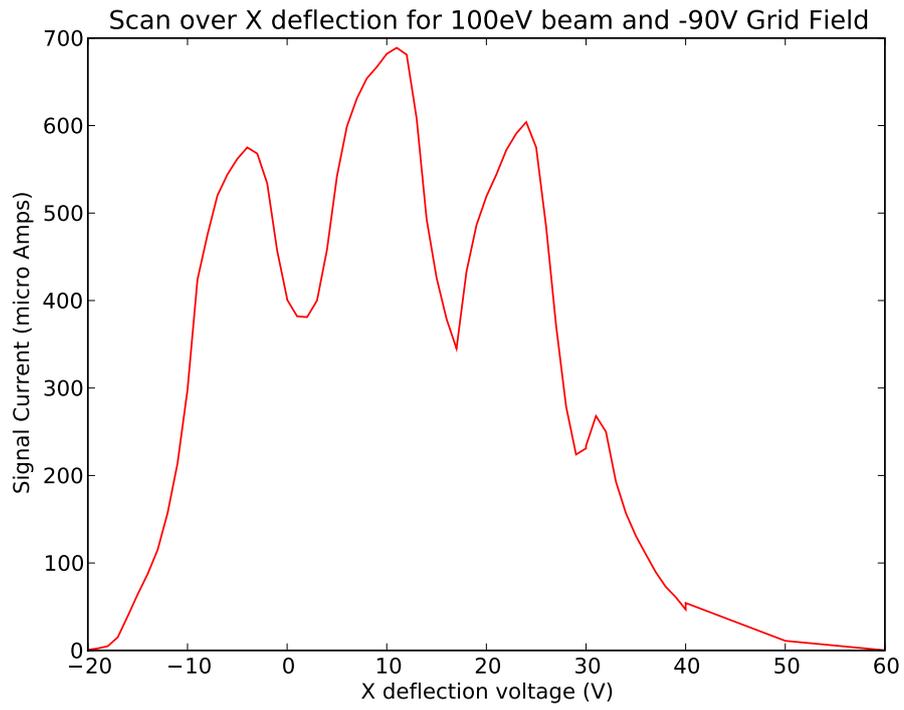


Figure 13: Scan over X deflection for 100eV beam. The multiple peaks are from the beam being occluded by the grill lines. These peaks are indicative of a beam width of one or two grill widths.

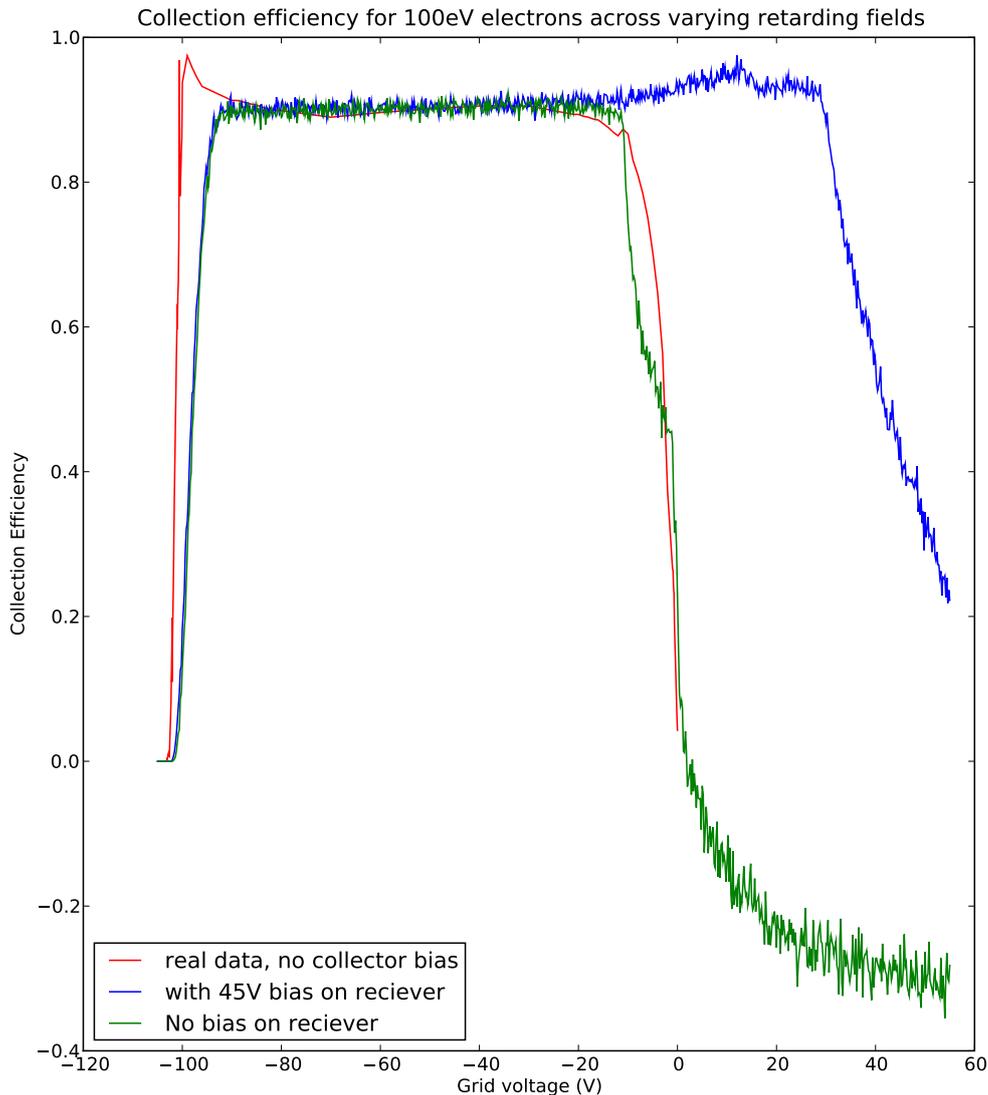


Figure 14: This is a scan across Grid voltage for 100eV electrons. The red is experimental data superimposed on the green and blue simulated data. The real data was originally a signal current at a set emission current. It was scaled arbitrarily to set the plateaus to the same value. The cusp in the real data is from error in the voltage reading in the power supply from switching to an over 100V value from a 0V-99.9V value.

The final dataset taken is in Figure 7. The overlay of real data with simulated data shows some interesting disparities. There are differences in the centers of the transitions, which could be due to the simulation models being poorly calibrated, or possibly to some systematic error in the power supply readout. At the -100V transition the slopes are substantially different. Also noticeable is the substantial maximum efficiency peak present in the real data but completely absent in the simulation. Is this a beam effect or is it due to unmodeled emission effects? One possibility could be due to grid focusing on a wrinkled, non-ideal wire grid. More simulation studies must be done to determine this.

At the 0V transition there are other interesting effects between the real data and the simulation

without a bias. The experimental curve seems caught between some sort of shoulder in the simulation. As stated in section 3.3, the model for the energy of SE electrons is incomplete. All of the energies are set at exactly $1/9$ of the impact electron energy. From that, the shoulder seems to be a transition from trapped to escaping collector SE electrons at that energy in the simulation. A more distributed energy spectrum should convolute that shoulder, bringing the transitions more into agreement.

The transition seen for the 45V collector potential simulation is interesting. Previous experiments show a much larger jump in signal at 0 than is seen. The disparity could be from our different geometry, but it is likely due to the same oversights that give different transitions at -100V. The sim does show a gradual increase, which is likely due to SE from the grid into the collector.

8 Conclusion

Altogether, this project has been an extremely interesting fusion of simulation and experimental setup. Seeing the test stand behave much like the simulations was extremely promising. With sufficient understanding of the retarding field analyzers, I believe that the data from the Main Injector installation should give a detailed picture of the electron cloud.