Dark Current Simulations for the Cornell ERL

Christie Chiu*
Massachusetts Institute of Technology
Cambridge MA, 02135

August 12, 2011

The proposed Cornell ERL contains two linacs that we would like to assess for dark current. Our dark current simulation takes an integer number of particles per cavity and phases per particle from the user, then creates these particles with a weighted charge according to the Fowler-Nordheim field emission model. It tracks these particles through the accelerator lattice and returns the particles' final positions, energies, and angles of impact. We implement several features in our simulation, including asequential element tracking, time-based Runge-Kutta tracking, global and element coordinate systems, querying EM field grid data, RF cavity autophasing, rootfinding, and element walls. After running the simulation code with a cyromodule lattice, we find that the majority of charge hits the cavity irises, but the particles colliding with the adjacent pipes have the higher energies. Our main limitations at this point are accounting for Coulomb forces between dark current particles, adding asymmetric geometries and curved elements, and simulating secondary electrons from dark current.

1 Introduction

Field emission is the release of charged particles from a surface due to local electromagnetic fields. Dark current is the flow of these particles through the medium into which they were released. They accelerate according to the fields in the accelerator, so they could end up virtually anywhere on the accelerator walls with a large range of energies. Levels of dark current are difficult to measure; because these destinations are not well known, there are no obvious locations to place Faraday cups. Other measurements involve measuring the effects of dark current, such as radiation, and therefore lack in precision [1]. Therefore, a dark current simulation is quite appealing as a method to assess dark current levels.

The proposed Cornell ERL contains two linacs which accelerate and decelerate the particle beam but are also likely field emitters due to their geometries and high time-dependent electromagnetic fields. Each linear accelerator is composed of cryomodules, which in turn are made up of six RF cavities. The RF cavities contain electromagnetic fields that oscillate with a frequency of 1.3 GHz such that they provide maximum acceleration to the particle beam. We have developed a general dark current simulation code for an arbitrary lattice; a lattice is the arrangement of components (such as pipes, quadrupoles, and cavities) that make up the accelerator or part of the accelerator. We then ran it with the linac lattices to assess dark current in the linacs and the extent of damage that it may cause. Our input variables are the particle phase space coordinates and relative weight of the particle, which is represented by an effective charge. The particle's phase is defined according to the RF cavity oscillation, and will be determined so that

the phases of all particles will be evenly spaced over one RF period. Our outputs, then, are the final phase space coordinates, time of impact, and angle of impact of each particle.

The dark current simulation code is written in Fortran 90 and uses the `Bmad` subroutine library. `Bmad` was developed at Cornell beginning in 1996 by David Sagan and contains the tools necessary to simulate particles in accelerators and storage rings [2]. We also utilize Mathematica, a computational program, to generate an input file and process the output data from the simulation. The simulation features several new methods such as time-based tracking, querying electromagnetic fields, and element walls; it also uses existing Bmad routines for particle beam simulation which include RF cavity autophasing, global and element coordinate systems, and Runge-Kutta numerical integration.

Our simulation of a cryromodule suggests several key points. First, particles are emitted with nonnegligible charge only in a narrow band of phases. Second, most of the dark current gets deposited at the cavity irises, but with energy very close to the electron rest energy. The particles that do attain higher energies tend to collide with the nearby pipes and other lattice elements. Finally, a very few number of particles leave the cryomodule entrance or exit with a significant amount of energy, but their effective charge is so small that this may be negligible.

2 Dark Current

In accelerators, we use electromagnetic fields to focus, direct, and accelerate charged particle bunches. Here, our field emission surfaces are the interior walls of the accelerator, the dark current particles are electrons, and the medium is the beamline vaccuum. The emission and movement of these electrons are very often completely unintended and separate from the particle beam. In fact, the particle beam does not need to be on for dark current to exist; we require only that the fields be present. Now we will look at the dark current origins, differences from the particle beam, and possible harm from dark current to develop a better understanding of this phenomenon.

2.1 Origins of Dark Current

The Fowler-Nordheim model for electron field emission applies to bulk metals, and so is a good approximation for our system. It considers electron tunneling through a triangular potential barrier to give emission current as a function of the local electric field as well as two parameters A_{FN} and β_{FN}. The parameter A_{FN} is the effective emitter area, and usually takes on values between 10^{-7} and 10^{-19} m^2 [3]. The model also uses a field enhancement factor β_{FN} that usually falls within the range 50 to 300 [3]. Although the values of A_{FN} and β_{FN} will be higher for surfaces that emit more electrons at any given field strength, they do not have strict physical interpretations which makes it very difficult to rigorously determine their values. For our simulations, we let A_{FN} equal 10^{-17} m^2 and β_{FN} equal 300. These estimates border the upper bound of the ranges typically seen, which should give us an upper bound on the severity of dark current in the accelerator.

If we include the values for our parameters and the work function for our Niobium cavities, we find that the Fowler-Nordheim equation is:

$$I_{\mathrm{FN}} = 3.85 \cdot 10^{-24} \cdot \mathrm{A} \cdot \left(300 \cdot E \cdot \frac{\mathrm{m}}{\mathrm{V}}\right)^2 \cdot \exp\left(\frac{-5.464 \cdot 10^{10}}{300 \cdot E \cdot \frac{\mathrm{m}}{\mathrm{V}}}\right)$$

From this equation, we see that surfaces with a high local electric field will have a higher dark current. Two main factors exist which determine this electric field strength. First, metal surfaces with a small radius of curvature have higher field strengths. Second, the dirtiness or roughness of surfaces contributes to local field strength for similar reasons, because we can think of the unwanted particles and obtrusions as having very small radii of curvature. We account for the first factor when we calculate and define our electromagnetic fields in the accelerator. The latter manifests itself through

the values of A_{FN} and β_{FN}, with the assumption that all surfaces have equal cleanliness and smoothness.

Thus, given any point on an accelerator element wall and a time, we can query the electric field strength and direction to determine the level of dark current emitted from that point. Notice that the Fowler-Nordheim equation only holds if the electric field is pointing into the wall; that is, the dot product of the field vector and the normal to the wall is negative. No electrons can be emitted if the electric field would push nearby electrons into the wall.

2.2 Differences Between Dark Current and the Particle Beam

The behavior of the particle beam in an accelerator is well-known because all accelator elements are designed to shape or direct the beam in a certain way. Because the nature of dark current is so different, however, we cannot apply any of this knowledge to dark current.

First, the particle beam is created in the injector with a specific phase relative to the accelerating RF cavities, specific bunch length, and nonzero momentum. It also has a narrow emittance, that is, it occupies a small horizontal and vertical phase space. On the other hand, dark current comes from field emission of any element surface, as long as the electric field is strong enough. The electrons originate from the walls, their initial momentum is typically very small, and they may or may not be emitted in bunches.

Second, the particle beam has a specific path, namely, along the accelerator beamline. The time-dependent electromagnetic fields that the beam encounters are typically phased to give the beam particles maximum acceleration. Magnet strengths are tailored according to the beam energy to keep the beam away from the walls as well as shape the beam itself. Field emitted electrons, however, encounter widely varying directions and magnitudes of electromagnetic fields on account of their origins and phases. Consequently, their paths are much more erratic- the particles may decelerate at times, and even reverse direction. Multiple particles starting at the same point on the accelerator wall but emitted at different phases could theoretically collide at completely different destinations with completely different energies. Therefore, we must consider time of emittance relative to the RF cavity phasing as well as original location.

The momentum spectra for the particle beam and dark current differ greatly as well, mainly due to their differences in intial momenta and path. The particle beam attains the maximum energy possible, simply because that is the accelerator's main goal, but dark current could have final energies anywhere up to that maximum value.

2.3 Effects of Dark Current

The major effects of dark current include heating of accelerator elements, harm to nearby electrical components, activation of nearby materials, and demagnetization of permanent magnets in the accelerator. Field emitted electrons that collide with higher energies are more likely to create more detriment, so we want to know the final energy spectrum of dark current in any given accelerator.

2.3.1 Heating of Accelerator Elements

When stray electrons hit the accelerator walls, some of their energy is deposited as heat. This is harmful because many elements are sensitive to temperature, such as the niobium RF cavities that must be held below 2 K to be superconducting. Past studies of dark current have also revealed that some locations heat up above the melting temperature of the metal, which makes the surface uneven and may even cause splashes of liquid metal [4]. This raises the possibility of a situation where dark current causes increased roughness of the wall surface, which then increases the likelihood of field emission [3, 4, 5].

2.3.2 Harm to Electrical Components

There are electrical components near the accelerator walls to monitor key signals and take data. If the dark current collides with an area too close to these instruments, their deposited energy can cause both short term and long

No.	Element	Length (m)	Description
1	Gatevalve	0.075	
2	Taper	0.05	
3	HOM Absorber	0.06	
4	Copper Pipe	0.0476	Drift element
5	Niobium Pipe	0.113275	Drift element
6	Niobium RF Cavity	1.0	Contains oscillating EM fields to accelerate particle beam
7	Niobium Pipe	0.113275	Drift element
8	Copper Pipe	0.0476	Drift element
⋮	⋮	⋮	⋮
39	HOM Absorber	0.06	
40	Taper	0.05	
41	Corrector	0.15	
42	BPM	0.075	
43	Quadrupole	0.45	Four permanent magnets used to focus the particle beam
44	Gatevalve	0.075	
45	Intermoduleflex	0.488	

Table 1: Listing of cryomodule elements in order. Elements 3 through 8 repeat six times.

term damage. We usually refer to the short term effects as single event upsets (SEUs), and they include a memory bit flip or current surge [1]. On the other hand, a single hard error (SHE) describes long term damage such as permanently setting a bit, and are much more destructive [1]. It may be difficult to detect when these errors occur, but if we know the distribution of dark current final energies and destinations we can take the appropriate measures to prevent them from happening.

2.3.3 Activation

Electrons emit bremsstrahlung when they interact with particles in the accelerator walls and decelerate. Given a sufficient energy, this radiation can knock out neutrons from atoms in the wall. We then have neutron radiation, which can spread throughout the wall in a series of chain reactions and make it radioactive. Of course, this radiation is very harmful to people working on or near the accelerator. This makes it one of the key motivations to better model and understand dark current; by knowing where the locations most prone to radioactivation are, we can install radiation shielding to ameliorate the problem.

2.3.4 Demagnetization of Permanent Magnets

Finally, field emitted electrons can also cause demagnetization of permanent magnets in the accelerator [1]. This happens through radiation, and the level of demagnetization has several factors including the direction of magnetization and temperature of demagnetization [6]. Accelerators have permanent magnets for several purposes, for example: dipoles bend the beam, quadrupoles help focus it, and undulators make it emit X-rays for experiments requiring such radiation. Once these magnets start losing their strength, the quality of the beamline and radiation it produces will also decline.

3 ERL Linac

Given the nature of dark current and its multiple harmful effects on the accelator, simulating dark current would be quite helpful. In particular, we want to simulate the linacs because their time-dependent RF cavity fields complicate predicting field emission and stray electron paths. In addition, RF cavity fields tend to reach much higher magnitudes than fields in other elements because their goal is to maximize beam bunch acceleration. That, combined with the fact that cavities are composed of many elliptical curves, some of which have relatively small radii of curvature, make the cavities much more likely field emitters.

The ERL will contain two linacs, Linac A and Linac B. Linac A consists of 35 identical cryomodules. Similarly, Linac B consists of 29 cryomodules, all identical to each other and to those in Linac A [7]. Table 1 lists the elements in each cryomodule. There are a total of six RF cavities, each flanked by both niobium and copper pipes and separated from the others by an HOM absorber. Each cavity consists of seven cells; the geometry of these cells is determined by a series of ellipses and tangents. A three dimensional image of a cavity is shown in Figure 1.

The cryomodules have time-dependent RF electromagnetic fields whose purpose is to accelerate the particle beam. They oscillate at 1.3 GHz and are timed such that at any given instance, any two adjacent cells will have electric fields of the same magnitude, but opposite direction, as seen in the z- and x-components of the electric field in Figure 2a and Figure 2b. This is so that as the beam bunch travels through, it attains the maximum acceleration possible. Looking at the magnitude of the electric field in Figure 2c, however, notice that the strength of the field around the irises is comparable to the strength in the center of the cells along the beamline. Such field magnitudes around the irises imply that these are key locations for field emission. Finally, Figure 2d shows the magnetic field magnitude in the y-direction, or the angular direction in this case. Adjacent cells also have the same magnitude of magnetic field but opposite direction, as with the electric fields, but the magnetic fields are very small near the beamline. Therefore, the particle beam does not feel a strong force from the magnetic field, whereas dark current might.

Quadrupole magnets are placed in the cryomodules to focus the particle beam. Successive cryomodules will have stronger quadrupoles because as the beam gains energy from the accelerating cavities, it requires stronger electromagnetic fields to focus the beam by a comparable amount. All other elements are optically equivalent to drifts, or pipes without fields, and can be simulated as such.

4 Methodology

Accelerator simulations at Cornell were all beamline simulations, so the program architectures available were optimized according to particle beam properties. Because dark current differs widely, we have to implement some new methods. The following subsections discuss preexisting and new simulation algorithms.

4.1 Particle Creation

There is a separate program to create emitters around each cavity and give them a weight according to how likely they are to emit dark current. We then start a particle from each emitter whose charge is not the charge of an electron, but instead an effective charge that represents the weighted current over a period of time. Given any number of particles per cavity and number of phases per particle, we place each particle so that they are evenly spaced along the edge and have a random angle ϕ rotated from the x-axis about the center axis. If these electrons all had a value ϕ equal to zero, and we created one hundred particles per cavity, the emitters would be located at the blue points in Figure 3. The charge equals the current calculated by the Fowler-Nordheim equation, multiplied by the time difference in adjacent phases, then multiplied by the particle's distance from the center axis. This is because areas of the wall further from the center will

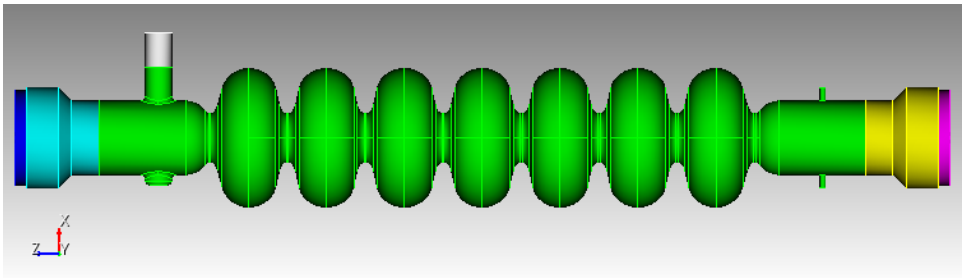

Figure 1: A cavity for the Cornell ERL. Image courtesy of LEPP.

have a larger area, and we need to normalize to account for this. Each electron begins with zero momentum.

Some particles will be created with very small charges (on the order of 10^{-50} C and smaller), which are negligible when we look at the statistics of these particles. To cut down on run time of our code and subsequent data analysis, we only create particles if they have a charge within a certain limit of the maximum. This limit can be set by the user.

4.2 Sequential Element Tracking

Particle beams typically travel from one element to the following element, simply because that is how the accelerator is designed. Dark current, on the other hand, presents a complication in that it may be directed backwards through the elements, forwards, or a combination of both. We account for this in our simulation by keeping track of which end each particle exits to determine which element to track through next.

4.3 S-Based and Time-Based Tracking

If we let s be the distance along the beamline, with $s = 0$ being the very beginning of the accelerator, then we know that the coordinates for the particle beam are well-defined for all positive values of s (up to the total length of the beamline). We then express all coordinates as a function of s rather than a time t [8]. This method is advantageous because we want to know a beam's position and momenta at a given location in the beamline rather than at a certain time. After all, for example, we have instruments at specific locations along the beamline so our measurements are location-based, not time-based. Additionally, with synchrotrons we have the beam passing through any given s position a multiple of times, which makes it more reasonable to refer to the beam according to location rather than time.

However, given any s coordinate, the coordinates for a dark current particle may not be well-defined. They could be multiply-valued if the particle passed by more than once, or undefined if the particle never reached that lattice section. It then becomes advantageous to track with all coordinates a function of time t.

4.4 Coordinate System Conversion

`Bmad` gives each accelerator lattice element its own coordinate system, and both s-based and t-based tracking are designed for the element coordinate system[1]. To track through multiple elements, then, we require a series of coordinate transformations. Given the architecture of the simulation driver program and tracking subroutines, the best solution is to transform from the lab reference frame to the element frame before tracking, and to perform the inverse after tracking.

Notice that in the linac, all elements have a straight beamline through their center so our s-axis is equivalent to the z-axis. We will refer to this dimension as s. The x-axis is defined according to Figure 3, and the y-axis follows from the right hand rule. We also keep track of the momenta in each direction as the particle travels, giving us a total of six dimensions for s-based tracking and seven in t-based tracking.

[1] See Chapter 9 of the `Bmad` manual for more information about the coordinate systems used [2].

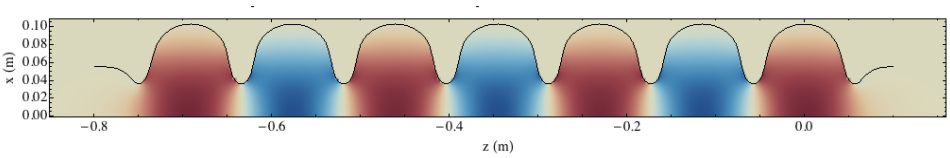

(a) z-component of electric field

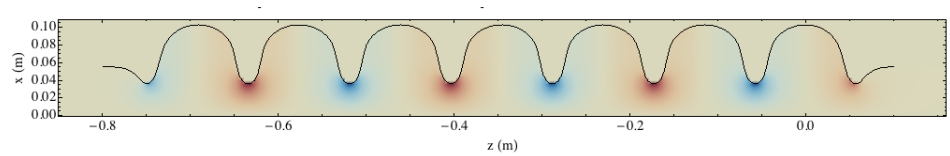

(b) x-component of electric field

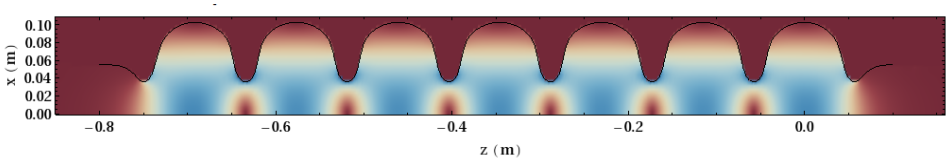

(c) Electric field magnitude

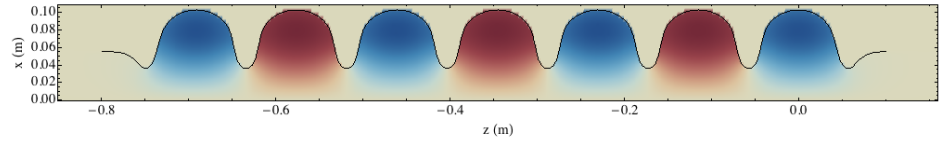

(d) y-component of magenetic field

Figure 2: Cross section and electromagnetic field maps of an RF cavity for the Cornell ERL at a given point in time. Images courtesy of Christopher Mayes.

4.5 Beamline Fields and Querying EM Fields

When simulating the particle beam, we can typically use a transfer matrix for the RF cavities- that is, there exists a matrix that, when multiplied by a particle's coordinates at the entrance of the cavity, gives us its coordinates at the exit. Dark current particles, however, can emerge from the cavity with a wide variety of changes in position and momentum. The transfer matrices are then invalid, and we must track each particle through the cavity by iterating over small time steps and querying for the electromagnetic fields at each step. We can then calculate where the particle will go according to the Lorentz force. Thus instead of a transfer matrix, we require the electromagnetic fields for the whole volume over a complete period.

Unfortunately, a data file containing such EM field data can be quite large and increase our simulation run time considerably. Therefore, we simplify our data under the assumption that the cavities, and therefore cavity fields, are rotationally symmetric about their center axis[2]. First, we generate a two-dimensional grid that can be superimposed on our cavity element along its xs-plane. At each grid point, we save the necessary EM field data, and create our data file. When we query for the field strengths and directions at any given point in space, we rotate the point around the center axis until it lies in the xs-plane. The three-dimensional volume then reduces to a two-dimensional plane, and we are left with a simple bilinear interpolation from the grid data to get our EM field at the desired point. Figure 4 shows the vector field data for the electric field at a given point in time for a portion of our RF cavity.

[2]This is not entirely the case; see §6.

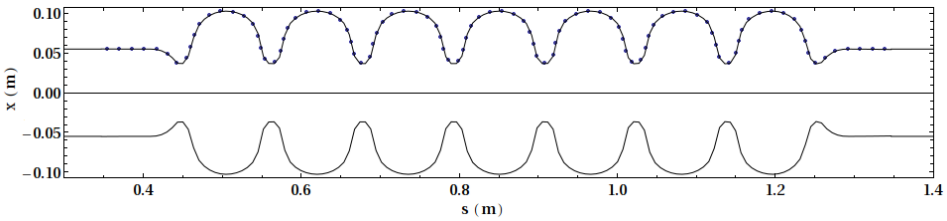

Figure 3: Emitter locations for dark current originating from 100 evenly-spaced locations along the cavity wall, in the xs-plane.

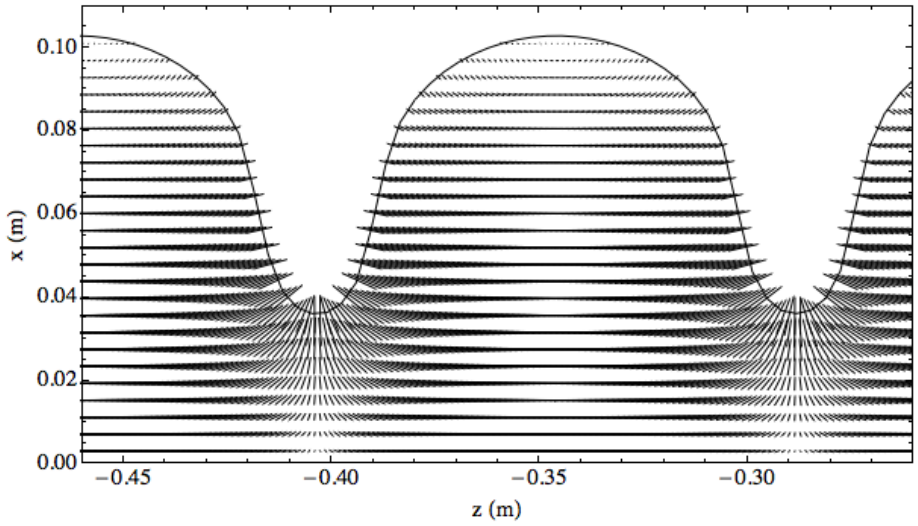

Figure 4: Electric field grid data used in our dark current simulation, shown in vector field form. Image courtesy of Christopher Mayes.

4.6 Runge-Kutta Tracking

Numerical integration methods are used to track particles through EM fields because all tracking is based on discrete distance or time steps. Indeed, it would be too computationally intensive to try and generate continuous functions representing the paths that particles take as a result of the fields they encounter, even for the beam bunch particles. `Bmad` offers several integration methods for s-based tracking, some of which have adaptive stepsizes for optimization and some with fixed stepsize. For t-based tracking, we use the fourth-order integration method Runge-Kutta with a constant time step for simplicity.

4.7 Autophasing

When we send a particle beam through a cavity, the relative phasing between the beam and cavity fields is timed such that the beam gains the maximum acceleration possible. Similarly, `Bmad` autophases the cavities so that a particle sent into the beginning of the lattice at time 0 will leave the lattice with the greatest energy attainable. We keep autophasing on while tracking field emission particles because the timing is crucial to their paths and final coordinates.

4.8 Lattice Element Walls

The particle beam should never hit the accelerator walls. Because of this, previous simulations only roughly specified wall geometries for the accelerator. We implement wall structures for dark current so that we can simulate the creation of particles on the walls and determine any subsequent collisions with the walls.

`Bmad` wall structures are defined for an element through a series of cross sections per-

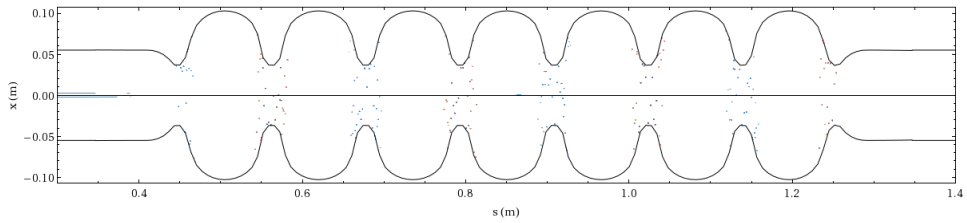

(a) First cavity

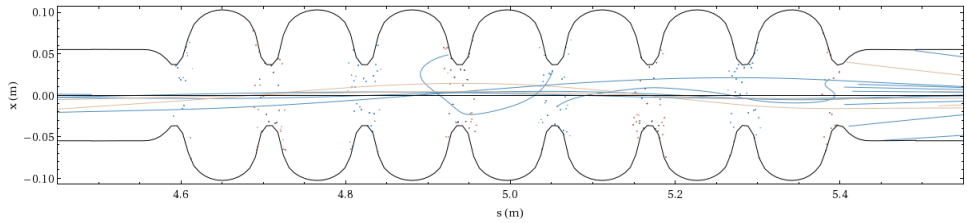

(b) Fourth cavity

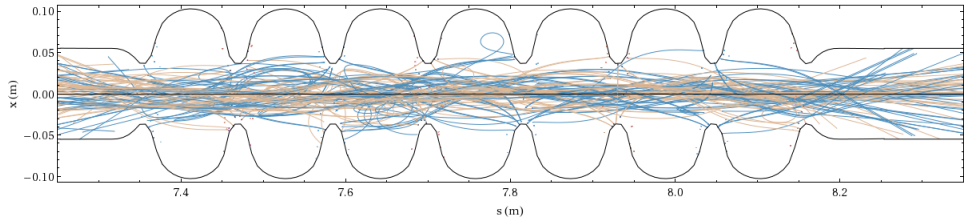

(c) Sixth cavity

Figure 5: Particle tracks in three of the six cavities of the cryomodule, color coded by the particle's phase.

pendicular to the s-axis [2]. All of our cross sections are circles, so they are specified by its center (which is always the origin[3]) and radius.

The walls between cross sections are determined through linear interpolation, unless additional parameters are given for cubic interpolation [2]. For our dark current simulations, the former method is a better choice for several reasons. First, we model any non-cavity walls as hollow cylinders, which naturally suits linear interpolation. Second, it is much less convoluted given how we generate our cavity walls; because the cavity must fit the EM field grid data perfectly, we obtain our wall data from the same source as the grid data. It comes as a series of data points, so we decimate the data then create a cross section at each point. On the other hand, cubic interpolation would require us to pick out key points from the data set so that the walls would still be represented accurately, which is nearly impossible to automate. We would also need to calculate the cubic coefficients in the form that `Bmad` requires. Third, to create evenly-spaced field emitters we measure the cavity wall length. Of course, linearly-defined walls make this much easier. Finally, if straight lines connect the cross sections rather than cubic functions, we can place our field emitters comfortably near the wall and between cross sections because that section of wall will always be convex.

[3]We assume our cavities to be rotationally symmetric, but this is not completely accurate; see §6

4.9 Rootfinding

Regardless of whether it is from the beam or field emission, whenever a particle exits an element we need rootfinding to determine its last coordinate. This is to pass on to the next element, after the approprate coordinate changes, to continue tracking. However, both s- and time-based trackers take discrete steps, so the last tracking coordinate will often go past the end of the element. To remedy this, we use a rootfinder that takes smaller stepsizes and

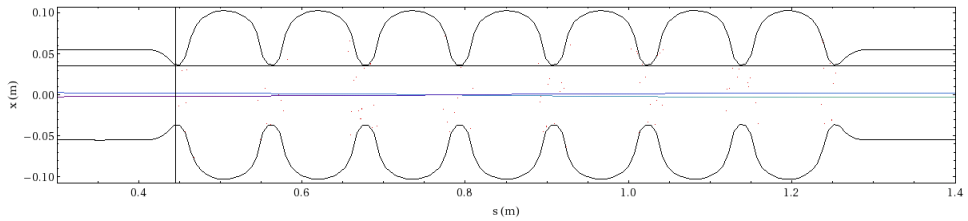

(a) First cavity

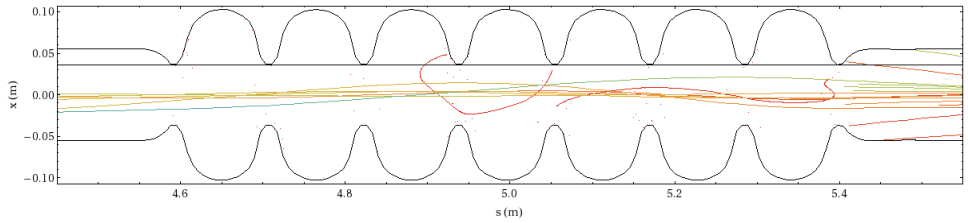

(b) Fourth cavity

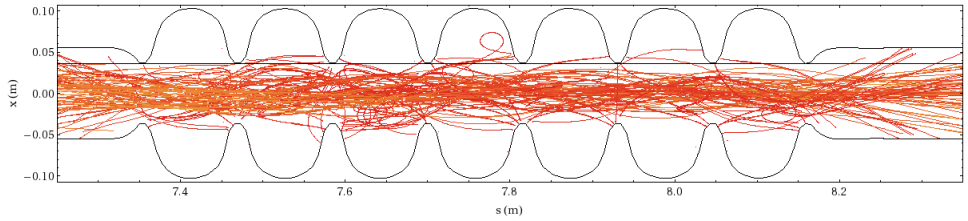

(c) Sixth cavity

Figure 6: Particle tracks in three of the six cavities of the cryomodule, color coded by energy with red representing the lowest energy.

finds the exact point of intersection[4]. Similarly, whenever a dark current particle collides with a wall, the rootfinder determines the exact location of impact. It also adjusts the particle's momentum accordingly, so that its final energy is accurate as well.

5 Preliminary Results

In this section, we will look at two data sets.

The first contains particle tracks, or phase space coordinates at all time steps, in one cryomodule. For each cavity, we look at 500 emitter locations evenly spaced along the walls with a random angle about the s-axis and 100 phases per location, evenly spaced over the RF period. Particles originating at these phases and locations are given a weighted charge according to the Fowler-Nordheim model and their distance from the s-axis. We then filter for particles with a charge within 10^6 of the maximum charge, giving us a total of 46,330 dark current particles to track.

The second data set contains only final phase space coordinates in one cryomodule. We create our particles in the same way and with 500 emitter locations per cavity, but with 200 phases. This gives us a total of 92,792 particles.

5.1 Dark Current Tracks

In Figure 5, particle tracks color coded by the particle's phase are shown for three cavities in the cryomodule. The particles in the cavities at the end of the cryomodule tend to have much longer tracks than the particles at the beginning, which may be an artifact of autophasing. Indeed, because the RF cavities are timed so that a beam bunch attains maximum acceleration after traveling through the cryomodule, there is a phase shift between any

[4]There is an allowed range within the exact point of intersection in which our last point can fall, which can be manually adjusted.

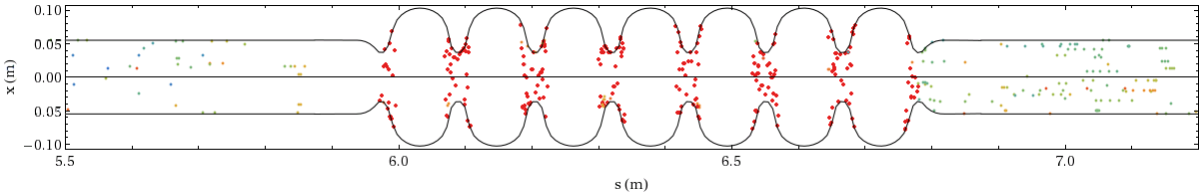

Figure 7: Particle destinations in the fourth cavity of the cryomodule and surrounding pipes, color coded by energy with red representing the lowest energy.

two cavities. As a result, two particles starting at the same time in two different cavities have differing relative phases, and so take different paths.

In the first cavity, Figure 5a, the tracks are very short which indicates that dark current does not travel far. Furthermore, the particles begin only around the cavity irises, areas of high curvature and maximum electric field. From this we see that particles are emitted with sufficient weighted charge only near the irises, which agrees with our previous understanding of field emission. The colors alternate between a tan and a blue from one iris to the next, demonstrating that particles are also emitted within a specific range of phases depending on the oscillating fields. From Figure 5b and Figure 5c, we see that most field emitted electrons travel near the center beamline rather than near the cell walls.

Figure 6 shows some particle tracks in the same select cavities of the cryomodule. This time, they are color coded by energy, which ranges from the rest energy of an electron (red) to 46.5 MeV (violet). So, we see that the electrons that have short or very curved tracks collide with low energy, while electrons with longer tracks straight through the cavities collide with much higher energies. The high-energy particles with the blue and purple tracks in the first cavity (Figure 6a), then, most likely came from a cavity towards the end and traveled backwards through the cryomodule.

5.2 Dark Current Destinations

The final phase space coordinates and angle of impact for dark current hold greater importance than the tracks because they will help us with radiation shielding calculations to minimize the harm done to the accelerator. For this data set of 92,792 particles, three exited the cryomodule's beginning with $1.44 \cdot 10^{-11}\%$ of the total weighted charge of all particles and a total energy of 126.57 MeV. Similarly, fourteen left through the exit with $9.56 \cdot 10^{-6}\%$ percent of the total weighted charge and an aggregate energy of 105.67 MeV. Although the total energy leaving the cryomodule is large, the weighted charges are so small that this is negligible.

Looking at the section of cyromodule in Figure 7, notice that many particles collide with low energy at the irises. The majority of these particles most likely originated at the same iris they collided with, as the particle tracks previously discussed suggest. However, the irises also have smaller apertures than other areas of the cryomodule; this makes them likely areas for dark current to hit. Particles of higher energies seem to be more prone to end up in the adjacent pipes, as seen by the yellow, green, and blue points.

To get a better idea of what energies are attained by field emitted electrons, we create a weighted histogram of energies, as shown in Figure 8. The majority of particles gain very little, if any, energy, but a good deal of charge lands with energies up to 10 MeV. The peaks in the weighted histogram for these energies may result from the discrete cavities transferring energy to each particle that passes through. There are also three single peaks around 12, 16.5, and 19 MeV that probably represent the particles that exit the cryomodule.

Similarly, we generate a weighted histogram of each particle's final s-coordinate to see where this energy is deposited; see Figure 9. Immediately we can distinguish to locations of

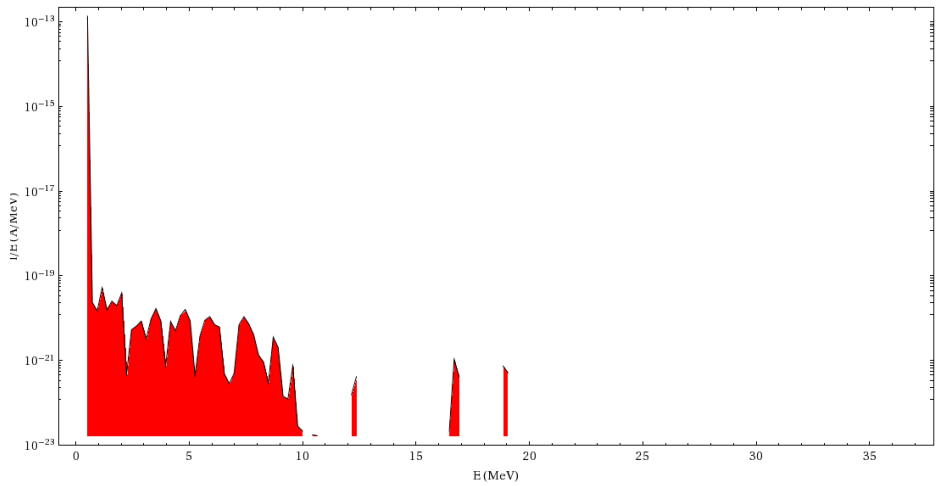

Figure 8: Histogram of final energies for dark current in one cryomodule, weighted by effective charge.

the six cavities and seven cells per cavity by the six clusters of seven or eight peaks each. Towards the end of the cryomodule, we see some more charge deposited in the pipes and other elements between cavities, albeit a very small amount compared to the total effective charge.

Overall, from this data it seems that most of the effective charge gets deposited at the irises, but the particles hitting the pipes have higher energies. The particles leaving the cryomodule have the highest energies, but there are so few of them in number that this is negligible.

6 Limitations

6.1 Coulomb Forces Between Dark Current Particles

The dark current simulator studies field emitted particles by creating them at a given point and letting them go one by one. This allows the tracker to iterate over all particles and process only one at a time. However, we know that in the actual linac, if two electrons were emitted at the same time very close to one another, they would experience a Coulomb repulsion and subsequently move away from one another unless the surrounding fields were strong enough to indicate otherwise. The tracker will not account for this. Frankly, it would be unreasonable to try to. To do so would require tracking all particles over each time step and updating of the EM fields between each time step. Given that we may easily take over hundreds of time steps, tracking multiple electrons at once is essentially impossible[5].

Field emission should not be a common occurance in the Cornell ERL because the niobium cavites undergo a rigorous cleaning process. Therefore, we can assume that the dark current emission locations are sparse and that the probability of dark current particles interacting with one another is incredibly low. Neglecting Coulomb forces between dark current particles by tracking them individually is a reasonable decision.

6.2 Asymmetric Geometries

To keep run time down for our dark current simulation, the RF cavity's EM field data is stored as a two-dimensional grid[6]. The grid is then rotated along the cavity center axis to create field data for the entire volume, which assumes symmetry. In reality, the cavities are not completely symmmetric. On one end, they contain a forward power coupler and sym-

[5] Beam simulations track all the particles in a bunch at once to study their collective properies such as bunch length or beta function. This is possible because they move at relativistic speeds, so the Coulomb Forces between them are negligible.

[6] For more information, see §4.5.

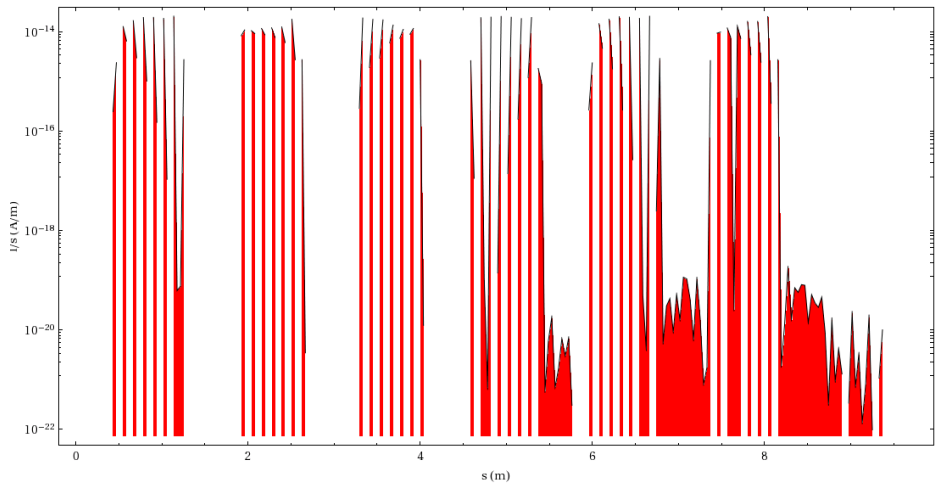

Figure 9: Histogram of final s-coordinate for dark current in one cryomodule, weighted by effective charge.

metrizing stub. On the other, two RF probe ports extend in the positive and negative x-directions. The cavity walls, too, assume rotational symmetry and therefore do not accurately represent the cryomodules.

These asymmetries would not be very difficult to implement. We already have the complete asymmetric field data, and just need to replace the two-dimensional grid methodology with a three-dimensional grid and interpolation. Adjusting the wall structure may require some editing by hand, but would be time-consuming rather than difficult. Of course, if there is still a large amount of symmetry in the cavity, which there should be around the RF cells, we could place some logic in the code so that we could use the old method if the particle were in the symmetric portion of the cavity, and the new field data for the extra appendages that we've added on to the cavity. This would probably help decrease run time.

6.3 Curvilinear Coordinate System

The whole ERL contains both straight elements and curved ones such as dipoles. Our dark current simulation has coordinate changes in place that work with the former and should work with the latter. They have only been tested with the ERL linacs, which contain only straight elements. Additionally, `Bmad` allows lattice elements to have a tilt, pitch, or offset; they can be rotated about or translated along the x-, y-, or s-axes [2]. Our coordinate changes should account for these parameters, but this too has not been tested yet. Testing is not an issue once the ERL lattice has been created.

6.4 Secondary Electrons

When dark current particles collide with the wall, they may create secondary electrons. These particles are scattered or ejected from the point of impact into the beamline vacuum, where they follow the electromagnetic fields and travel as another form of dark current. Simulating secondary electrons is quite difficult, because their inital momenta depends on the angle of impact and energy of the incident charge. Furthermore, when these seconary electrons collide with the walls they may release even more charge into the vacuum, and we see that the rate at which particles are created exceeds the rate at which we track particles. In other words, secondary electron simulation is too computationally intensive to implement at this point.

7 Future Work

Other than accounting for the limitions described in §6, there are several features and

tasks to incorporate and complete that will improve the existing dark current simulation. The dark current simulation works without any obvious mistakes, but of course we need to check for any bugs or false assumptions. It would also be helpful to determine how accurately the simulation models dark current in the linac, although this would be difficult because so little is known about dark current.

7.1 Minimum Particle Count

Currently, the particle generating code will not necessarily meet a quota when creating particles; it only looks at a select number of locations and creates particles there if their weighted charge is sufficiently high[7]. To ensure that we create enough particles to get useful data from our simulation, we want to implement a minimum particle count. Then the code would check whether it had made enough particles, and if not, increase the space and/or phase resolution to create more particles until it met the requirement.

7.2 Simulation Optimization

The code needs to be profiled to ascertain which areas are taking the most time computationally and to see what can be done to make the code run faster. This would allow us to simulate more particles at once. Due to the structure of our code, we could also run our simulation in parallel; because we track each particle individually, they do not depend on each other. Finally, our Runge-Kutta tracking uses a fixed time step, and implementing an adaptive stepsize may make the simulation more efficient.

7.3 ERL Simulation

If our simulations show that there is a high level of dark current leaving the linacs, then it will be particularly important to include the rest of the ERL in the simulation. Most of the electrons leaving Linac A should not successfully travel through the bend to Linac B because of the strong dipole magnets located in the bend. However, the electrons leaving Linac B see a series of undulators and a larger radius of curvature, so they may travel farther and could cause some damage. Field emission is much more unlikely in other areas of the ERL given the geometries of other accelerator elements and the smaller electromagnetic field strengths. Regardless, for completeness the entire ERL should be simulated.

[7] See §4.1

7.4 Probability Distributions for A_{FN} and β_{FN}

By taking the values of A_{FN} and β_{FN} to be constant in our Fowler-Nordheim model, we assume that our cavities have uniformly cleanliness and smoothness. Of course, this is not the case, but these two parameters are empirically determined and not much is known about them. Assuming that we can find a realistic probability distribution for A_{FN} and β_{FN}, such a feature would make our simulation much more accurate.

7.5 Data Analysis

So far, we have conducted only a preliminary data analysis for our dark current simulation results. More thorough data processing, especially after we address the improvements discussed above, is necessary to develop a complete and accurate assessment of dark current in the Cornell ERL. More specifically, we want to know how much current leaves each linac and with what energy, as well as how much charge hits the wall at each position s along the beamline and with what energy and angle. This will in turn prepare us for radiation shielding calculations.

8 Acknowledgements

Special thanks to Christopher Mayes and Georg Hoffstaetter for their advice and guidance throughout this project. Many thanks to David Sagan as well for his assistance with Bmad. This project was supported by the NSF.

References

[1] Lars Fröhlich. PhD dissertation, Universität Hamburg, 2009.

[2] Dave Sagan. "The Bmad Reference Manual". Revision 15.2. 2011.

[3] Hasan Padamsee. *RF Superconductivity: Science, Technology, and Applications.* (Wiley-VCH, 2009).

[4] J. Norem. *et al.* Phys. Rev. Special Topics-AB **6**, 072001 (2003).

[5] Hasan Padamsee *et al. RF Superconductivity for Accelerators.* (Wiley-Interscience, 1998).

[6] Alexander Temnykh. Nuclear Instruments and Methods in Physics Research A. 587 (2008) 13-19.

[7] Christopher Mayes. PhD dissertation, Cornell University, 2009.

[8] Klaus Wille. *The Physics of Particle Accelerators: An Introduction.* (Trans. Jason McFall. Oxford University Press, 2001).