

Chapter 4

Apparatus and Techniques

To perform the measurements in this thesis, a variety of systems installed in the FFTB were important. We first discuss the hardware common to all aspects of E164: the plasma source, the diagnostics for understanding the beam transversely, and the Cherenkov based energy spectrometer which is the central diagnostic for acceleration.

In 2003, we added new hardware at the upstream end of the FFTB. This location was advantageous because there was space available for our 3 ton magnet and associated hardware, but also because there is substantial horizontal dispersion, as discussed in § 3.6.3. Using this magnet, as we discuss below, allows the *non-destructive* measurement of the electron beam's energy spread before it enters the plasma.

We compare these energy spectra, for every data shot, with simulations of the main accelerator. This enables us to determine the phase space for every beam as it enters the plasma and provides greater power in understanding our acceleration and other effects, as discussed in Chapter 5. The specific hardware and spectrum measuring techniques, as well as a discussion of the software we use are the core of this chapter, but first we discuss the existing hardware common to all parts of E164.

4.1 Plasma Source and Diagnostics

The acceleration experiments use a Lithium plasma. The intense electric fields of our electron beams are so strong that they will fully ionize Lithium vapor quickly




Figure 4.1: Schematic of the plasma cell with diagnostics for understanding the electron beam before and after the plasma. Light from the OTR foils is imaged to view the transverse size of the beam and a series of electron optics image the plasma exit to the Cherenkov radiator while dispersing the beam to measure the spectrum.

compared to the transit time of the beam, as discussed in Chapter 2. We provide a column of Lithium vapor with variable length and density in order to investigate acceleration under a variety of conditions.

After the beam's energy is changed by its interaction with the plasma, we measure the final energy spectrum as a way to understand the plasma's effect on the beam.

To aid in properly focusing the electron beam into the plasma, we also use several beam diagnostics immediately before and after the plasma. We discuss these first.

Figure 4.1 shows a schematic of the beamline at the end of the FFTB where our plasma experiments take place. The major features of relevance are the OTRs, the plasma source and the spectrometer.

4.1.1 OTR Diagnostics

To take an image of the electron beam's transverse components both entering and exiting the plasma, we insert thin metal foils in the beam path to create optical transition radiation (OTR).

A general discussion of transition radiation can be found in [60]. When the electron beam transits the foil, the collapsing dipoles created by the electrons rushing toward their image charges in the foil radiate light across a broad spectrum. As with synchrotron radiation, wavelengths short compared to the bunch length are radiated incoherently, and longer wavelengths will have a coherent nature.

We focus on the incoherent portion for this measurement, as it lies in the visible and is easily viewed with standard cameras. Also, incoherent radiation gives a direct image of the beam, because each electron has an equal probability of radiating, while coherent radiation processes, by definition, are affected by groups of electrons. We want to know where each electron is independent of the others in order to build up a picture of the beam profile.

Radially polarized radiation is emitted from the point where each electron enters the metal and then spreads out from that point. The distribution of the radiation into a given solid angle is only dependent on θ , the angle of the radiation away from the central axis. For ultrarelativistic beams impinging on an infinitely wide metal foil with a large (real portion of the) dielectric constant, the formula for the radiation intensity at a given photon energy and into a given solid angle is, in CGS:

$$\frac{d^2 I}{d\omega d\Omega} = \frac{e^2}{\pi^2 c} \left| \frac{\sqrt{\epsilon} - 1}{\sqrt{\epsilon} + 1} \right|^2 \frac{\theta^2}{(\theta^2 + 1/\gamma^2)^2} \quad (4.1)$$

The angle where the peak of the radiation happens is, perhaps unsurprisingly, given by $1/\gamma$. In order to have the radiation come out to the side where we can view it, the foil is mounted at 45° to the beam's direction of motion. The radiation thus comes out centered on the direction perpendicular to the electron beam path.

There are several concerns that have been raised about using OTR to produce an image of a beam. The first is that the radiation pattern has the above mentioned peak at $\theta = 1/\gamma$. If this value represents the angular width of the radiation distribution, the resolution limit that we could image would go as $\gamma\lambda/2$, and for visible light with our highly relativistic beam, we could not resolve points on the metal foil less than about 1 cm apart.

Fortunately, this conclusion is actually wrong by a large margin. As discussed in [61], there is still substantial power per solid angle radiated to angles much larger than $1/\gamma$, as the intensity only drops as $1/\theta^2$. Additionally, the total solid angle to be radiated into grows linearly with the angle, so the number of photons coming out at a given angle only drops as $1/\theta$. A more useful metric for the distribution of this light than the naive value of the peak location at $1/\gamma$ is therefore the RMS spread of

the photons in angle, $\sqrt{\langle\theta^2\rangle}$. We calculate this quantity for (4.1), and calculate out to some $\theta_{max} \gg 1/\gamma$, where θ_{max} is governed by the actual angle subtended by the optical element being used to capture the light. We find the surprising result that the effective width of the beam at the front of the lens depends primarily on the size of the lens we are using, even though $1/\gamma$ is very small, because there is still so much power out at large angles:

$$\sqrt{\langle\theta^2\rangle} = \frac{\theta_{max}}{\sqrt{2\ln(\gamma\theta_{max})}} \quad (4.2)$$

To capture OTR light, we have a camera lens of roughly 5 cm diameter situated about 25 cm away from the screen. The angle subtended by the camera is therefore $\theta_{max} \approx 0.2$, and for our case, the resolution is about *2500 times* better than expected from naively taking that the characteristic angle is governed by the angle at which the highest flux of photons is emitted, $\theta = 1/\gamma$. We thus expect that for 1:1 imaging in our optical system, we will have diffraction limited resolution of $5\text{ }\mu\text{m}$, which is smaller than the pixels on our camera, and we can see the beam's profile clearly. Images of the beam have shown features with roughly this size scale, so the resolution has been as good as expected.

It is worth noting that, unlike the naive case, where resolution gets worse linearly with increasing particle energy, this function depends only very weakly on the relativistic factor of the particles, going as the square root of a logarithm. Therefore, not only is OTR useful in the current SLAC linac, but it should be useful even as electron energies go much higher in the future.

A second concern about using OTR for beam imaging is that there is a distance from the source of radiation, known as the formation length, L_f , within which the Coulomb field of the bunch moving at some v and the radiation field moving at c are not well separated. This is equivalently the distance scale within which we are in the near field, and paraxial optics do not fully apply. This length grows rapidly with the energy of a relativistic particle and is given by $L_f \approx 2\lambda\gamma^2$ [62]. For optical radiation at 500 nm from our 28.5 GeV beam, this analysis predicts that $L_f = 3.1\text{ km}$! Our camera is only 25 cm away from the source of radiation, and is thus very much in the near field. There certainly remain theoretical questions about exactly what happens with

the radiation impinging on a camera in the near field. However, earlier experiments measured the beam's transverse size with a wire scanner while simultaneously looking at OTR images of the beam nearby. The sizes were consistent and imply that the possible resolution was not compromised significantly from being in the near field.

The beam has very high peak intensity, so an appropriate material for the foil is required. The main damage mechanism for substances placed in the beam comes from the heat load of impact ionization of the material. The rapid deposition of heat can cause small cracks between neighboring crystals from the differential expansion in a kind of shock. The total amount of heat can also simply melt a hole in the foil.

To minimize the ionization, we turn our gaze to metals with low atomic number. Naively obvious choices from the first few members of the Group 1 and 2 metals of the periodic table are ruled out because of toxicity (Beryllium), brittleness (e.g. Calcium) or extreme chemical reactivity (e.g. Lithium). Additionally, the material will heat up rapidly at the point where the beam transits, requiring a high melting point. Lastly, great stress is induced between locations of high and low heating, requiring high mechanical strength. Aluminum has a low melting point and is not particularly strong, so the first good choice is Titanium, which can be made into very thin foils.

Our foils were obtained from Arizona Carbon Foil Company and are grown by vapor deposition on a polished substrate to a thickness of $1\text{ }\mu\text{m}$. The substrate is then removed to leave only the foil. Each is tensioned slightly and vacuum epoxied to a stainless steel ring for support.

Our expectation that they would be suitable was borne out, as they have held up well when hit by the beam.

To view the OTR light, we use Photometrics SenSys 12 bit cameras with 768×512 pixels, where each pixel covers an area of 9×9 microns. To capture the light, we use a Nikon f/2.8 AF Micro-Nikkor 105 mm lens. At 25 cm, the lens captures enough light that we expect tens of billions of photons to be collected, for a strong signal. We set the lens to the 1:1 imaging mode such that the pixels on our CCD each see that same area on the Titanium foil. We have plenty of light, and in seeking to image with high resolution, we employ a blue filter to reduce chromatic aberrations. There is also a linear polarizer in front of the lens, as it is expected to improve the resolution

in the plane perpendicular to the direction of polarization somewhat [63]. As with other cameras in the experiment, a computer reads out the data at 1 Hz.

We can measure the beam's size both before and after the location of the plasma, where the beam is nominally focused. We have the ability to change the final focusing quadrupole strengths, and can observe the changes in spot size on the upstream and downstream OTRs as the focal point moves. This allows us to verify that we have placed the focal point at the entrance to the plasma as we intend.

Also, any transverse tails that the beam has coming from the main linac will be visible on the upstream OTR camera, and are an indication of input conditions where the beam has a tilt. In § 5.3.2, we see the utility of knowing the transverse properties of the beam before it enters the plasma.

The downstream OTR camera provided the first indication that we appear to be trapping electrons from the plasma itself and accelerating them. We discuss this effect further in § 5.1.1.

These two OTR based diagnostics are an important part of understanding the focusing and other transverse properties of our beams in the vicinity of the plasma oven, and even pointed to unexpected phenomena in the plasma.

4.1.2 Heat Pipe Oven

As we ionize a neutral gas with the beam itself, there are several possible ways to hold gas in a specific region for use in the plasma wakefield experiments. The most obvious simply uses thin metal foils at either end of a section of pipe, into which gas can be pumped in or out to change the density.

When we tried to use such a setup briefly in E164, we found that the downstream metal foils were rapidly punctured, allowing the gas to escape. Damage happened sometimes as fast as after a few minutes.

We therefore used a heat pipe oven, which uses temperature differences and a buffer gas to constrain Lithium vapor to a specific region. The buffer gas is kept from the main vacuum systems by thin Beryllium foils, which are situated a meter or two away from the regions with plasma, and are therefore not damaged.




Figure 4.2: Diagram showing the essential elements of our heat pipe oven.

The theory behind the heat pipe oven is described by its original designer C. R. Vidal in [40]. In our implementation, we place a few grams of Lithium inside a pipe which is about 3 cm in diameter and about 50 cm long. We maintain the entire inside of the pipe under an inert atmosphere of Helium gas. The Helium prevents any chemical reactions with the Lithium and is also central to the operation of the heat pipe oven.

With Lithium in place, we heat the central region of the pipe. This heating causes the Lithium to melt and evaporate at a rate dependent on the temperature. Outside this heated region, there are cooling collars wrapped around the pipe at either end to maintain the Helium gas on either side at something close to room temperature.

In this arrangement, a schematic of which can be seen in Figure 4.2, the Lithium evaporates and flows either direction away from the center of the heated region. This flow is at something like the speed of sound, and the rapid Lithium flow simply sweeps all Helium atoms from the central region.

Eventually, the Lithium encounters the cold Helium gas at either end and cools off rapidly. Having cooled, the Lithium condenses on the sides of the pipe where a mesh is installed. The molten Lithium is then drawn back toward the heated central region by the mesh's wicking effect where it again evaporates.

There is a substantial flow of heat from the central regions toward the Helium, but no net transfer of material. We have a central region with essentially a pure Lithium vapor, then wings on either side in which the Lithium density progressively drops as the Helium fraction rises, then the outer regions of pure, cold Helium.

The two parameters of interest concerning the oven are the density of the Lithium vapor inside and the length of the region of vapor. The density is simply controlled by the pressure of the Helium buffer gas. If we evaporate enough Lithium to expel the Helium, the so-called “oven condition,” then the Lithium pressure must equal that of the room temperature Helium gas outside.

In order to evaporate enough Lithium for the oven condition, sufficient power must be supplied to the heater. Then by adding more power, the Lithium will flow further before finally condensing, leading to a longer oven. It is worth saying explicitly that once the oven condition is reached, heater power has no effect on the Lithium vapor density, only oven length.

4.1.3 Determining the Oven Length for E164

We seek to know the length of our oven as a function of the vapor density and the heater power that we are using. There are three regions, as mentioned above, all with different properties. In the low temperature regions, there is no Lithium vapor, in the center of the oven, there is only Lithium, and there are the two regions with varying Lithium density, often referred to as the “wings.”

When measuring the changing Lithium density profile along the length of the oven, it is not feasible to measure the changing partial pressure of the Lithium vapor directly. However, we can measure the temperature at each point with thermocouples; there is a well defined relationship between the vapor density of Lithium and the temperature. This relationship has been found empirically for a variety of metals in the first column of the periodic table [64]. The equation for Lithium is given by:

$$P = \exp \left(7.530 \times 10^{-4} T - 2.053 \ln T - \frac{1.943 \times 10^4}{T} + 32.60 \right) \quad (4.3)$$

where the pressure P is in Torr and the temperature T is in Kelvin. Figure 4.3 plots




Figure 4.3: Plot of Lithium vapor pressure versus absolute temperature in the region of interest for E164.

this equation for the temperature range of interest in E164. At these temperatures, a vapor density of $3.5 \times 10^{17} \text{cm}^{-3}$ corresponds to a pressure of 46 torr. Even our highest vapor pressure is less than one tenth atmospheric, and the ideal gas law applies very well in this regime. Knowing the temperature and the pressure, we can calculate the density from the standard relation $PV = nRT$.

The temperatures of interest are very high, and limit the ability to increase the Lithium density. We chose to make the pipe from A.I.S.I Type 310 stainless steel, the properties of which are discussed in [65]. This Chromium-Nickel steel has the best tolerance for elevated temperatures of the standard stainlesses. When used at a constant high temperature, it can be used to 1150°C , but when cycled repeatedly from room temperature, the maximum temperature that is recommended is 1040°C . The damage mechanism for repeated heating seems to be that oxide layers form on the surface at high temperatures and then fall off due to differential thermal expansion whenever the material is cooled.

Scaling can eventually weaken the pipe, and constrains our ability to produce high density plasmas. An adequate margin of safety to prevent rupture of our pipe – with molten Lithium inside – meant that we limited the highest temperatures achieved to about 1020°C , giving a vapor pressure of 46 Torr and density of $3.5 \times 10^{17} \text{cm}^{-3}$. We have observed no scaling of the pipe so far from operation at these temperatures.




Figure 4.4: Oven profile plot at $2.4 \times 10^{17} \text{ cm}^{-3}$. The FWHM is 13.2 cm.

To measure the oven's vapor density profile, we mount the oven on a stand where we can insert a thermocouple into the middle of the oven and pull it through to record the temperature profile along the longitudinal axis. We can use the above formula to convert to partial pressure of Lithium and determine the oven length from the FWHM of the density profile.

There are wings on either side where the Lithium density drops to zero. Interestingly, the width of these wings appears to be independent of density over the range investigated for E164. The distance from the half density point to where full density is achieved is about 5 cm for all cases. Figure 4.4 shows a sample plot for a peak density of $2.4 \times 10^{17} \text{ cm}^{-3}$. There is a central flat topped region and then we can clearly see the wings on either side. A useful metric is the FWHM, which is 13.2 cm.

There is a slight asymmetry to the curves, and we attribute that to the fact that the thermocouple was at the end of a long metal rod along the central axis of the oven. That rod can conduct heat, causing the thermocouple to read slightly higher temperatures than it should when the rod is extended through the central region. Similarly, the thermocouple reads slightly lower temperatures than expected when the portion of the rod nearby is in cooler regions. This effect appears to be small, but could give a systematic error on the measurement of the FWHM.

For the several densities of particular interest to the investigations of this thesis, we measured the oven's length at several values of the heater power.

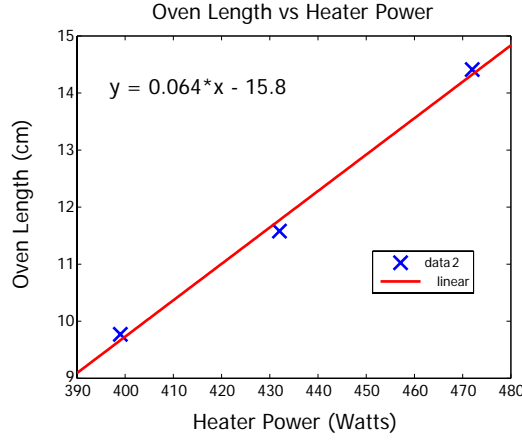


Figure 4.5: Plot of changing oven FWHM with increasing heater power for an oven with plasma density $1.5 \times 10^{17} \text{ cm}^{-3}$.

Figure 4.5 shows the change in oven FWHM as a function of the heater power. The measurements for this were taken at UCLA, where the oven was originally built. The laboratory there is maintained at or below about 20°C . In the FFTB tunnel, the temperature is often closer to 40°C , so we speculate that the same heater power will lead to a somewhat greater oven length than plotted here. This is because the oven is surrounded by a thick layer of insulating bricks, but substantial heat nonetheless flows out through that layer of insulation. In a colder room, it takes somewhat more power to maintain the interior of the oven at a given temperature.

The oven is at 1000°C , so this difference in external temperature should have only a small effect. We will ultimately be primarily concerned with *changes* in the oven length (see § 5.3), and these changes should be similar in the FFTB and at UCLA.

We have such curves for a variety of plasma densities, and Table 4.1 shows the slope of FWHM versus heater power over a ten-fold range in density. Knowing the oven length is certainly necessary to measure the gradient in our plasma.

Given the substantial wings on either side of the oven, if we seek to have a region at the nominal full vapor density, the shortest oven possible is about 10 cm, where the full vapor density is achieved only at the exact center of the oven. By adding more power, we can lengthen the oven, making the peak density region grow, while retaining wings of the same size.

Oven Density (cm^{-3})	Slope (mm/W)
0.3×10^{17}	0.68 ± 0.05
0.6×10^{17}	0.70 ± 0.04
1.5×10^{17}	0.64 ± 0.04
2.0×10^{17}	0.58 ± 0.03
2.5×10^{17}	0.56 ± 0.03
3.5×10^{17}	0.58 ± 0.04

Table 4.1: Oven length versus heater power slope for different oven densities. Each value for the slope of length versus heater power is calculated from several measurements of the FWHM at different heater powers. Errors are estimated from the spread in data points. The presence of the thermocouple may also add a systematic error.

Unfortunately, the heater used in E164 is damaged if too much power is put into it. A practical limit was about 500 Watts. At that power level, the oven cannot be longer than about 15 cm, so the range of oven lengths available was limited. Nonetheless, by changing the oven length, the region of full density can be changed in length from about 0 to about 4 or 5 cm long, allowing measurements of the change in acceleration with different oven lengths to be performed, see 5.3.

The other half of the calculation of gradient requires knowledge of the change in particle energies, and that is measured on the Cherenkov based energy spectrometer.

4.1.4 Cherenkov Spectrometer

The electron optics between the plasma cell and the spectrometer involve three quadrupoles and four long permanent magnet dipoles. We can image the exit of the plasma to our spectrometer location and also have substantial vertical dispersion.

At the location some 25 meters downstream from the plasma, with the needed small β_y and large η_y , we send the beam through a piece of aerogel 1 mm thick. When the electrons go through this material, they radiate many optical photons from the Cherenkov effect, as their speed exceeds that of light in the aerogel. The light from each electron comes off in a forward cone with an opening angle of several degrees, and we place a pick-off mirror to the side of the electron beam path to capture a portion of the light from each electron's cone, while allowing the beam to pass by.

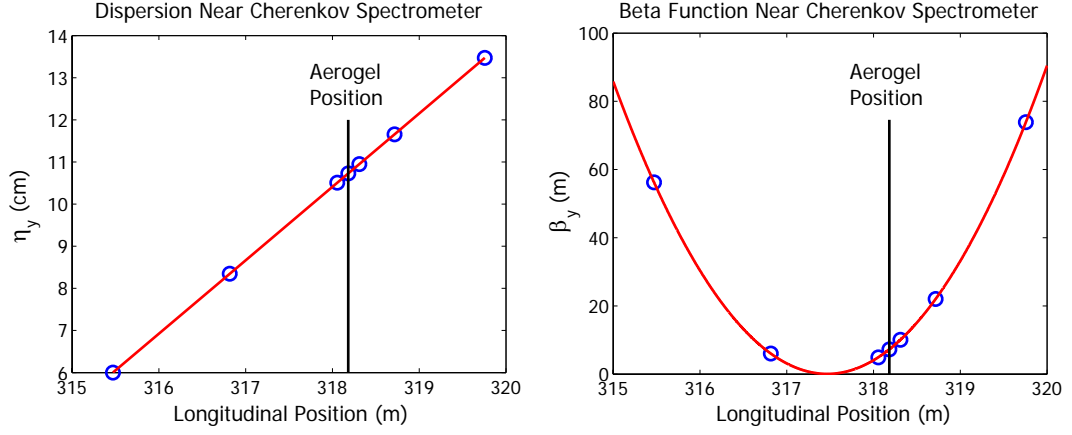


Figure 4.6: Output from ELEGANT showing the vertical dispersion and beta function for the electron beam in the region of our Cherenkov radiation based spectrometer.

This technique of creating Cherenkov radiation to observe the electron positions was originally used in the first plasma wakefield experiments involving long bunches with σ_z of about $650 \mu\text{m}$. For such long bunches, a streak camera can be used to measure the energy spectrum of different longitudinal slices of the beam, allowing a direct measure of the acceleration in the tail. Clearly, this requires a prompt method of radiation production, so the Cherenkov screen was used. With the ~ 40 fsec bunches in E164, we can no longer use the streak camera, so have no specific requirement that the radiation be prompt, but the system had worked well and did not need to be changed.

After picking off some of the Cherenkov light, we then use several mirrors to transport the signal to our 16 bit Princeton Instruments VersArray CCD camera located near the beamline. The camera lies in a region of substantial background radiation, so a lead hutch was constructed to shield the delicate CCD. A Nikon AF Micro-Nikkor 200 mm f/4 lens imaged the Cherenkov light to the camera.

The dispersion at the aerogel location can be measured through surveying, and calculation of its value using beam optics simulations has also been performed in ELEGANT [66]. Both methods give that the dispersion is 10.5 cm with an estimated error of a few millimeters. Figure 4.6 shows the output from ELEGANT with the vertical dispersion and beta function plotted in the region of our spectrometer.

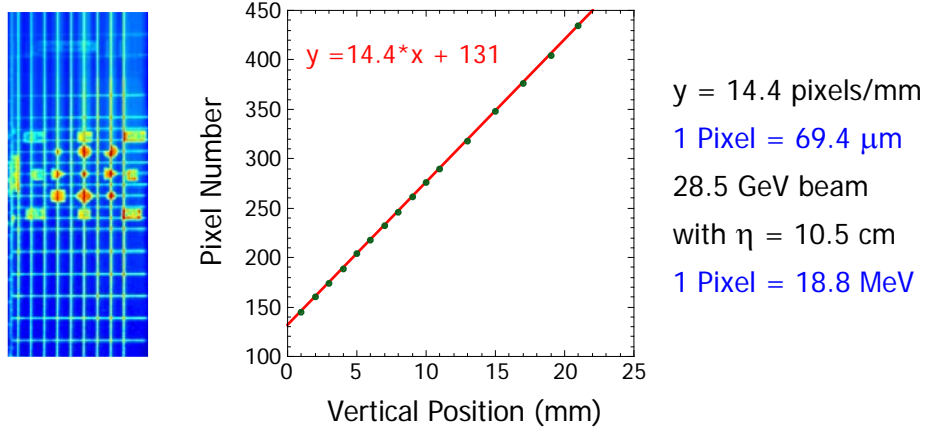


Figure 4.7: Sample image of the Cherenkov Aerogel with the calibration grid visible behind. A plot of the number of pixels per millimeter, combined with the dispersion, allows calibration of the energy per pixel.

With a number of quadrupoles between the plasma and the Cherenkov screen, there is a potential to introduce dispersion from misalignments in the quadrupole positions. To obviate this problem, all quadrupoles were aligned at the beginning of the experimental run to have no effect on the beam orbit with changing magnet strength, so do not contribute to the dispersion.

To understand the energy resolution of our spectrometer, we must know the dispersion and the vertical span viewed by each pixel of our camera. If Δd is the pixel width, or equivalently, the pixel separation, then the corresponding energy range per pixel is given by:

$$\Delta E = \frac{\Delta d}{\eta} E_0 \quad (4.4)$$

To measure the pixel's viewing area, a thin piece of paper is placed behind the aerogel with lines spaced 1 mm apart. We then image the grid of lines to our CCD to determine how many pixels span each millimeter. Figure 4.7 shows that there are 14.4 pixels spanning each millimeter, so the pixels each view an area of $69.4 \mu\text{m}$. With a dispersion of 10.5 cm at the aerogel, the energy dispersion is therefore 18.8 MeV/pixel for a beam with central energy of 28.5 GeV.

As this spectrometer lies far from the exit of the plasma, we must verify that it is imaging in both x and y . We verify directly that we are imaging in the horizontal plane

by changing the location of beam focus in the vicinity of the plasma and observing that the minimum horizontal spot size at the aerogel occurs when the upstream waist is placed at the location of the plasma exit. The horizontal spot size that we see on the Cherenkov screen is typically of order $400\text{ }\mu\text{m}$ and the incoming beam optics are set up to create a $15\text{ }\mu\text{m}$ spot at the location of the plasma, so the magnification ratio is about 25 times.

With y , we have no way to verify the betatron imaging directly, but because we are able to observe relatively sharp features of the energy spectrum, we estimate that any blurring from the finite β_y amounts to less than one or two pixels on the camera.

This Cherenkov based spectrometer is the central diagnostic for the entire E164 experiment and shows the energy changes resulting from the plasma wake. As there is minimal dispersion in x , and we use optics to focus the beam in that plane, the Cherenkov screen can also be used to investigate deflections in the horizontal plane that may happen to the electron beam in its passage through the plasma.

4.2 Coherent Transition Radiation Measurement

A central concern of the E164 program is to know the length of the electron bunches coming from the linac. The use of Coherent Transition Radiation (CTR) can give information on the bunch length directly, and several different measurements based on CTR are of interest.

The setup for the CTR measurement is shown in the schematic of Figure 4.8. We send the electron beam through a thin titanium foil some 20 meters before it enters the plasma cell. The CTR radiating from the foil is collected by an off-axis paraboloid to be sent through our interferometer. As we are dealing with wavelengths of light in the roughly 5 to 100 micron regime, we use a vacuum window made from high-density polyethylene (HDPE) to provide reasonable transmission. This window also has the advantage of being opaque to the incoherent visible light portion, which could confuse the signal. Greater detail on the interferometer and the technique can be found in [67], and we summarize the main results useful for understanding the bunch structure.




Figure 4.8: Scanning interferometer for determining the electron bunch length.

The beam as a whole emits transition radiation as its electric field encounters the discontinuity of the conducting surface. With a Gaussian bunch, for all wavelengths where $\lambda \gtrsim 0.6 \sigma_z$, the radiation will be coherent [68]. For E164, this broadband CTR from the foil is emitted at a range of wavelengths. As we expect to produce bunches with a longitudinal RMS as short as 12 microns, the shortest wavelengths that can be produced coherently by this shortest bunch will be $\lambda_{min} \approx 7 \mu\text{m}$. This corresponds to frequencies of over 40 THz.

One way to use CTR to understand bunch length would be to measure the overall spectrum and inverse Fourier transform the result. This requires precise detectors that work over a very broad range of wavelengths. The method chosen for E164 was to split the radiation and send it through an interferometer to obtain the direct autocorrelation of the bunch's electric field.

In both cases, extracting the exact bunch length requires an assumption about the functional form of the bunch's profile, such as Gaussian or hyperbolic secant (as with many fast lasers). Nonetheless, choosing different functional forms has only minor effect on the calculated width. For example, the RMS width of the autocorrelation of a Gaussian pulse is $\sqrt{2}$ longer than the original pulse, whereas the hyperbolic secant pulse has autocorrelator RMS width about 1.5 times longer than that of the original

pulse. In general, we can estimate the actual pulse length to be shorter by roughly 1.4 to 1.5 times compared with the measured autocorrelation trace's RMS.

As with OTR, we are interested in the radiation pattern that we expect to observe, because knowing the spectrum helps to interpret results, see the end of this section. Theory again predicts a long formation length. Assuming that we are in the far field, the total power emitted as CTR per unit solid angle and per frequency is given by [69]:

$$\frac{dW}{d\Omega d\nu} = \frac{N_b^2 e^2}{4\pi^3 \epsilon_0 c} \frac{\beta^2 \sin^2 \theta}{(1 - \beta^2 \cos^2 \theta)^2} \exp \left[-\frac{4\pi^2 \nu^2}{c^2 \beta^2} (\sigma_z^2 + \beta^2 \sigma_r^2 \sin^2 \theta) \right] \quad (4.5)$$

where N_b is the number of particles, β is the particle speed normalized to c , and θ is the angle at which we look at the radiation.

We note that this function must be convolved with a function that takes into account the diffraction of the CTR from the finite source size in order to give the actual distribution of radiation in the far field [69]. We will integrate this numerically in Chapter 5, when we compare the total power observed against theoretical expectations. Clearly, shorter bunches will produce more total CTR power because the first term in the exponential becomes less negative with shrinking σ_z .

The above equation, and following estimates, all calculate based on assuming far-field optics. The formation length for CTR is longer than the size of our experimental apparatus, as it is for OTR, so these calculations must be considered approximate.

At the location where we measure the CTR, the electron beam is only about 60 by 170 microns in transverse RMS size. Thus, for even the highest energy photons we expect to produce, the source plane is not very many wavelengths across.

If we use estimates based on far field optics, we see that the radiation will diffract away from its point of emission rapidly, as the divergence angle of a beam of radiation is of order $\lambda/d \approx \lambda/2\sigma_r$ [70]. Thus, in the vertical direction, the $7\mu\text{m}$ shortest wavelength produced has a characteristic full angle of about $60\text{ mrad} \approx 3.3^\circ$; all longer wavelengths diffract away more rapidly. In the horizontal direction, diffraction is less severe, with about 20 mrad divergence for the shortest wavelength.

An off-axis paraboloid mirror intercepts the broad cone of radiation to collimate it so that it can be sent through the arms of a Michelson interferometer. The paraboloid

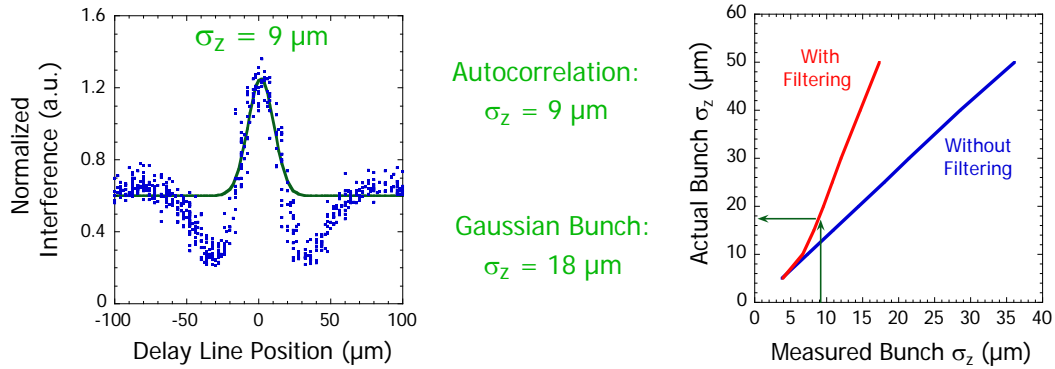


Figure 4.9: Autocorrelation Results. Courtesy Mark Hogan.

is a standard 2 inch optic held at 45° , so the effective width is about 1.4 inches. This optic is situated about 6 inches away from the CTR foil, so any radiation coming from the foil with a characteristic angle of divergence greater than about a third of a radian in the vertical plane or a quarter of a radian in the horizontal will not be collected fully. Given our source size, wavelengths beyond $40 \mu\text{m}$ are not fully captured vertically, and wavelengths beyond $90 \mu\text{m}$ are not fully captured horizontally. Again, as we are in the near field, these estimates should be considered approximate.

The detectors that we use for the radiation are pyroelectric elements (pyros) which absorb across a wide spectrum and produce a voltage in proportion to changes in their temperature. The detector elements are quite small, so have a bandwidth well above the 1 Hz repetition rate of our electron beam and measure each shot independently.

For autocorrelation traces, the reference detector is used to select only shots that have total CTR power within some narrow range. We know that the length of the bunches coming from the linac is not stable shot to shot, so it is convenient to select shots having total CTR power that is at the most probable value. We will see in Chapter 5 that that corresponds to bunches of order $20 \mu\text{m}$ long.

Scanning the interferometer arm distance over many shots, we can build up the bunch's autocorrelation function, as shown in Figure 4.9.

There are several effects which make interpreting the data not straightforward. Any head to tail asymmetries of the electron bunch are washed out in the autocorrelation, which is inherently an even function. Also, the broadband Terahertz radiation

being emitted by our source gets absorbed at various wavelengths by the high-density polyethylene (HDPE) vacuum window and the air. Moreover, we have seen that longer wavelengths of the radiation are not fully collected. Lastly, there are etalon effects in the beamsplitters. These elements have multiple resonances right in the brightest portions of our emitted radiation. This causes the bunch to appear artificially short and also causes the two dips visible on either side of the peak.

Taking these effects into account, we have calculated what we expect to see in the data for various bunch lengths. With the anticipated correction ratio as indicated in the right graph of Figure 4.9, we expect that the observed data actually correspond to a bunch length of $\leq 18\mu\text{m}$ in the measurement run presented. We will see in Chapter 5, that this is close to the answer provided by simulations. The correlation of these answers provides confidence that each is correct.

4.3 Non-Invasive Energy Spectrum Measurement

Many diagnostics which seek to measure the energy spectrum of an electron beam are destructive in that they present a physical barrier to the beam. This increases the transverse emittances and robs the beam of energy, also causing an increase in energy spread.

We are careful to create tightly focused beams which are short in time, so a physical barrier will clearly not be advantageous. Recalling a technique used in the SLC to measure the energy spread [71], we decided to use a large magnet as the basis for our spectrometer.

By deflecting the electron beam slightly in the vertical direction, each electron makes a vertical stripe of synchrotron radiation. As the electron beam is already dispersed according to energy in the horizontal plane, the sum of all such stripes is imprinted with the energy spectrum of the electrons and can be thought of as a sort of analog “bar code.”

We intercept these X-ray photon stripes on a scintillator to convert the signal to visible light. A scientific-grade CCD camera provides a digitized image of our bar code. Summing this image in the vertical direction, we get a curve which reflects




Figure 4.10: Schematic of the synchrotron radiation producing chicane.

the energy spectrum of our electrons. A schematic of the process of translating the electron positions to a visible signal is shown in Figure 4.10.

We need to separate the electron beam and the X-rays, so the magnet was configured as a chicane. The electron orbit returns to what it had been before the magnet, while the synchrotron radiation comes off in an upward fan such that it can be intercepted by a scintillator which is held just above the electron beam path.

4.3.1 Potentially Destructive Effects on the Beam

Although we do not present a physical barrier to the beam, there are still possible deleterious effects from its passage through the magnetic fields, so we calculate them.

Coherent Synchrotron Radiation Longitudinal Effects

The bunch is already very short at the point of our magnetic chicane, so we must verify that the effects of Coherent Synchrotron Radiation (CSR) will not be of concern. This radiation is at much longer wavelengths than the incoherent synchrotron radiation we seek to use. We discuss the properties of the incoherent portion in § 4.3.3.

The wavelength cutoff for CSR is of the order of the bunch length: $\lambda \gtrsim 0.6 \sigma_z$ (the same as for Coherent Transition Radiation, c.f. Chapter 5). CSR normally affects

the energy spread of an electron beam as well as its transverse emittances.

The Linac Coherent Light Source (LCLS) program at SLAC is very concerned with such effects in their proposed bunch compressor chicanes, and several useful formulae for estimating the energy spread due to CSR can be found in the LCLS Design Report [72]. In particular, we seek to know the energy spread induced by CSR, and it is given by:

$$\sigma_\delta \approx 0.22 \frac{Nr_e L_B}{\gamma \rho^{2/3} \sigma_z^{4/3}} \quad (4.6)$$

where N is the number of particles in our bunch, r_e is the classical electron radius, L_B is the length of the bends in our magnet, γ is the relativistic factor, ρ is the bend radius and σ_z is the bunch length.

Our magnet (see next section and Figure 4.12) has four bends of ≤ 50 cm length each. The peak magnetic field available is 1.3 T, with corresponding radius of curvature of roughly 75 m. An upper bound for the CSR effect is given by assuming that the entire bend length is at this full field. We obtain from (4.6) that σ_δ is no greater than 6.5×10^{-5} for the effect from all four magnets.

To find the final energy spectrum of the beam, we simply convolve a Gaussian of width σ_δ with the incoming spectrum, whose energy spread is already 1.1% in RMS. We can immediately see that this convolution will make no discernible change, because the added energy spread changes are negligible compared to the 1.1% RMS.

Coherent Synchrotron Radiation Transverse Effects

For changes in the emittance, we use a complicated approximate formula for the increase due to CSR found on page 7-38 of the LCLS design report:

$$\frac{\epsilon}{\epsilon_0} \approx \sqrt{1 + \left(\frac{0.11}{3}\right)^2 \frac{r_e^2 N^2}{\gamma \epsilon_N \beta} \left(\frac{L_B^6}{\rho^5 \sigma_z^4}\right)^{\frac{2}{3}} (L_B^2(1 + \alpha^2) + 9\beta^2 + 6\alpha\beta L_B)} \quad (4.7)$$

where ϵ_N is the normalized emittance before the bend, and α and β are the Twiss parameters for the beam at the location of the bend.

Given these values at the magnetic chicane, the increases in emittance for x and y are about 0.02 and 0.2 percent, respectively. This will have no effect on our beam.

Wake Fields from the Scintillator Crystal

In passing through an aperture, part of the energy stored in the beam's electric field is scraped off, causing the particle to lose energy. Our crystal sits on one side of the beam path about 2 mm away, but a useful limit on the wakefields comes from the more easily calculated geometry of a circular aperture of radius 2 mm; we follow the treatment in [73]. Any wake loss from the crystal will be something less than half the value calculated. For a bunch of charge q and length ℓ , where $\gamma\ell$ is much greater than the radius of the aperture, r , we find that the energy lost by each particle is given by:

$$W \approx \frac{N., e^2}{2\pi\epsilon_0\ell} \ln\left(\frac{\gamma\ell}{2r}\right) \quad (4.8)$$

For our bunches with 1.8×10^{10} particles, a FWHM length of as little as $25 \mu\text{m}$, and a distance to the crystal of 2 mm, we find that the maximum energy lost per particle is about 12 MeV. This is an overestimate not only because the crystal covers less area than an equivalent radius circular aperture, but because the fields are not completely reflected by the crystal, which is very much not a perfect conductor. Based on geometric arguments, every particle loses a bit of energy, with the largest expected energy loss something less than 5 or 6 MeV. The whole beam simply loses a tiny amount of energy, and there is no discernible effect on its structure.

Bunch Compression in the Chicane

Lastly, we can say that there is no measurable bunch compression occurring in our chicane. We know that the bend angles are only 2.5 mrad, and for a chicane with such small angles, we have an approximate formula for the R_{56} , where ΔL is the distance between the first and second bends, which must be the same as between the third and fourth bends [74]:

$$R_{56} \approx -2 \frac{L_B^2}{\rho^2} \left(\Delta L + \frac{2}{3} L_B \right) \quad (4.9)$$

Our chicane has $R_{56} \approx -17 \mu\text{m}$ and even extreme opposite energy particles will move by less than half a micron in z relative to one another. This tiny change compared to the bunch length is dwarfed by the effects of the FFTB dogleg's R_{56} .



Figure 4.11: Photograph of our magnet rewired as a chicane. The orange hoses are for water cooling. On the front face is one of the two added flux return plates.

Overall, these effects of the chicane on the electron beam are essentially zero, and we treat the beam as returning to its nominal orbit unchanged. Thanks to the bends in the chicane, however, synchrotron photons come off at a variety of small upward angles to be intercepted by a phosphor about 2 m downstream of the magnet.

4.3.2 Magnet Hardware and Details

Due to space constraints, we needed to deflect the electron beam by several milliradians in each bend of the chicane in order to have the X-Rays be at least a few mm above the electron beam by the time they were intercepted by our phosphor about 2 meters downstream.

To bend the 28.5 GeV beam sufficiently, we needed a magnet with high fields sustained over distances of more than a few centimeters. To bend our beam by 2.5 mrad requires that $\int \vec{B} \cdot d\vec{l} \approx 0.25 \text{ T}\cdot\text{m}$.

The available hardware was a large electromagnet 1.8 meters long and weighing some 3 tons, see Figure 4.11. This had originally been designed ca. 1980 for use as a four period wiggler with the beams in SPEAR of up to approximately 3 GeV. The




Figure 4.12: On left is a schematic of the rewired magnet. At right is a plot of the magnetic field strength versus longitudinal position going through the first bend.

magnet has nine separate pole pieces, where the central seven are the same and at either end are two poles with half the integrated $\vec{B} \cdot d\vec{l}$ of these central poles.

With over ten times the beam energy of SPEAR, the deflections from single poles in the wiggler would not have been large enough, so we rewired the magnet as a chicane by combining neighboring poles. The first 1 1/2 poles together constitute the first bend, the central 3 poles make bends two and three and the final 1 1/2 poles are the final bend. We can see a schematic of this change on the left in Figure 4.12.

The magnet was originally an air core magnet where the separation between the two halves could be adjusted to adjust the field. The magnetic fields available were substantially weaker because the magnet had to drive field across not only the air gap between the poles but also across the gap between the two halves of the yoke.

Requiring the strongest fields possible, we closed the yoke gap and held the magnet with two large steel plates. One of them is visible behind the forest of orange cooling hoses in Figure 4.11. These act as flux return for the magnet, increasing the field strength between the poles. With the plates, the magnet achieves about 1.3 Tesla at its maximum rated current of 800 Amps. The right side of Figure 4.12 is a plot of the magnetic field as we progress through the first bend. These measurements were taken with a Hall probe pulled through the magnet at 1 cm intervals. Included on the plot is the total cumulative bend angle for a beam at 28.5 GeV. We see that each of the four bends in our chicane will deflect the beam by 2.5 mrad, as required.

4.3.3 Synchrotron Radiation Properties

The electrons see a varying magnetic field in the course of traversing each bend, and therefore at each of these points, they have a different synchrotron radiation spectrum. We can calculate this instantaneous spectrum by following the treatment in [46]. A useful quantity to know is the total synchrotron power over all wavelengths, P_γ , which is given (in MKS) by:

$$P_\gamma = \frac{e^2 c}{6\pi\epsilon_0} \frac{\gamma^4}{\rho^2} \quad (4.10)$$

where, again, γ is the relativistic factor of the electron and ρ is the bend radius of the beam in the magnet.

To calculate the spectrum, rather than just the average power, we first need the critical frequency, which is the frequency above and below which half the radiated power lies. This ω_c is given by:

$$\omega_c = \frac{3c\gamma^3}{2\rho} \quad (4.11)$$

When we have calculated ω_c , it is natural to normalize all frequencies to it, using $x = \omega/\omega_c$. In these normalized units, all synchrotron radiation spectra look the same. To derive the actual spectrum, we just scale by ω_c and find the appropriate constants to give it unit area. Then we will just need to multiply by the total power, P_γ , to give the actual power at each frequency. Note that this is similar to what can be done with the blackbody radiation from an object at finite temperature, where ω_c is the characteristic parameter of the spectrum, as T is for thermal radiation.

It turns out that the spectrum's shape is given by an integral of a Modified Bessel Function of $5/3$ order. With P_γ and ω_c , we can determine the power spectrum for a single electron. The equation is:

$$P(\omega) = \frac{9\sqrt{3}}{8\pi} \frac{P_\gamma}{\omega_c} x \int_x^\infty K_{5/3}(x) dx \quad (4.12)$$

As our magnetic field varies, so will the spectrum's peak value and total power. To find the total spectrum in our bend magnet, we find the instantaneous spectrum associated with the magnetic field strength at each point through the magnet. We

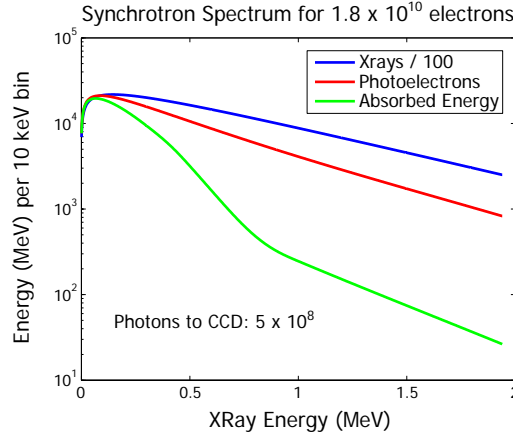


Figure 4.13: Synchrotron radiation spectrum for 28.5 GeV beam with 1.8×10^{10} electrons in the first chicane bend. Blue is the X-ray spectrum divided by 100 for plotting. Red shows the photoelectron spectrum from the (small) fraction of X-rays intercepted by the crystal. Green shows the energy ultimately deposited in the crystal.

use the data as plotted on the right hand side of Figure 4.12 as the basis for the calculation. Adding up the effects from each part of the bend, we get a numerical approximation to the integrated final spectrum of radiation that we expect for each electron and multiply it by the 1.8×10^{10} particles in each bunch. Although the calculation refers to ω , we simply multiply by \hbar to give the more conventional units of photon energy on the horizontal axis.

The first bend's total spectrum is shown in the blue curve of Figure 4.13. There is substantial power in the spectrum, and we divide by 100 for plotting purposes so the curve will lie close to the other curves discussed in the following section.

Note that the total energy lost by each electron in going through the chicane is about 800 keV, composed of photons primarily in the range around 100 keV. The loss of 800 keV is insignificant compared to the beam's average incoming energy. Changes to the energy *spread* are also small, as the average electron emits something like eight photons. Therefore, variation in energy loss between electrons is expected to be of order $\sqrt{8} \cdot 100 \text{ keV} \approx 300 \text{ keV}$, some 10 ppm of the beam energy. The incoherent synchrotron radiation has even less effect on the beam's final spectrum than the insignificant effect of the coherent radiation.




Figure 4.14: Scintillation crystal in its Aluminum holder.

4.3.4 The Scintillator Crystal

To detect the X-rays, we use a $10 \times 10 \times 0.08$ mm crystal made of Cerium doped Yttrium Aluminum Garnet scintillator (YAG:Ce). Figure 4.14 shows our yellowish crystal in its mount which holds it several millimeters above the electron beam path.

We chose this particular crystal because it is mechanically very strong, and has good light output centered in the green portion of the spectrum, where our camera is most sensitive. Due to its strength, we could use a thin crystal without fear of cracking it in the process of mounting to the aluminum holder. The thin crystal reduces blurring from multiple scattering of photoelectrons as they transit the crystal. We discuss this effect in § 4.3.7.

The thickness of crystal presented to the X-ray beam is greater by a factor of $\sqrt{2}$ because the crystal is at a 45° angle to the photon beam. This allows our camera to view it head-on with all parts in focus. (The camera is very close to the crystal with a zoom lens such that the depth of focus is only about one hundred microns.)

We address the issue of spatial resolution below, but first we must know how much energy actually gets deposited in the crystal from our synchrotron X-rays. We needed a crystal thick enough to give an appreciable signal.

X-Ray Absorption

The Atomic Nucleus by Evans [75] gives a variety of useful curves and formulae for the interaction of photons and energetic particles with matter.

We use these curves first to find the expected absorption of X-rays by our thin crystal. Specifically, a graph from p. 686 of Evans for Aluminum is most useful. Because of its greater density, YAG:Ce will stop X-rays 1.74 times as much as Aluminum, so we scale the absorption values accordingly.

The red curve in Figure 4.13 shows the significantly smaller total energy that is actually absorbed by electrons in the crystal, primarily through the photoelectric effect and some Compton scattering. With an effective thickness of $80 \times \sqrt{2} = 113 \mu\text{m}$ of material, we expect the absorption to vary from 1.1% at 10 keV to about 0.3% at 2 MeV. At 1 MeV, a photon coming from the magnet has only a 0.45% chance of being absorbed. We note that essentially all of the energy of the photon is transferred to the liberated electron.

Electron Range and Energy Deposition in YAG:Ce

The scattered electrons stop rapidly at low energies, but the faster electrons will fly out the back of the crystal, depositing only a fraction of their energy into it before escaping. On p. 625 of Evans is a useful empirical formula for the range of energetic electrons through matter when the electrons have energies between 10 keV and 3 MeV, which covers our spectrum well. When E is in MeV, and ρ is in g/cc, we find that the range, given in microns, of electrons in material is:

$$R = \frac{4120 E^n}{\rho} \quad n = 1.265 - 0.0954 \ln E \quad (4.13)$$

For YAG:Ce, with a density of 4.57 g/cc, we find that 10 keV electrons go less than a micron. At 100 keV, electrons have a range of about $30 \mu\text{m}$ and 1 MeV electrons can go about $900 \mu\text{m}$. See the left side of Figure 4.15.

With both the photoelectric effect and, for higher energy X-Rays, Compton scattering, the electrons are not ejected parallel to the direction of the incident photon. There is a spread in angles of emission, but for the lowest energy photons which can

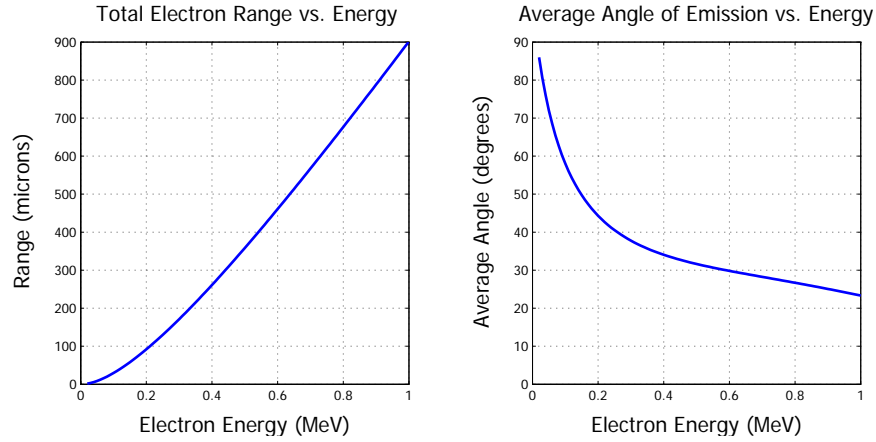


Figure 4.15: Total range and average angle of emission for electrons of varying energies transiting the YAG:Ce crystal.

still liberate electrons from the ground state, the angle of emission tends to be in the direction of the electric field, namely at 90° relative to the photon incident direction.

As the photon energies become higher, the momentum of the photon causes the electrons to be ejected more and more in the direction of the incident X-Ray. At about 0.2 MeV, electrons are ejected, on average, at 45° and by 1 MeV, the average angle is just under 25° , as shown on the right hand plot of Figure 4.15.

The electron loses energy through many small collisions, so its trajectory tends to continue along a fairly straight line. Given this, we can calculate the longitudinal and transverse components of the range for the average emitted electron. We do so in Figure 4.16, and note that these curves will be relevant later for calculations of resolution. For the moment, though, we concern ourselves with the longitudinal range, as it allows estimates of how much light the scintillator will produce.

The electrons liberated by the X-Ray beam are born evenly throughout the thickness of the crystal, so the average of how much energy is deposited by them must take that into account; with a tiny range, any electron born at 10 keV anywhere but the exact back surface of the crystal will deposit all of its energy into the crystal.

After electron energies rise much above about 0.25 MeV, their longitudinal range becomes roughly the same as the crystal thickness. If the electron is born at the exact front of the crystal, it will deposit all of its energy, but if it is born at the

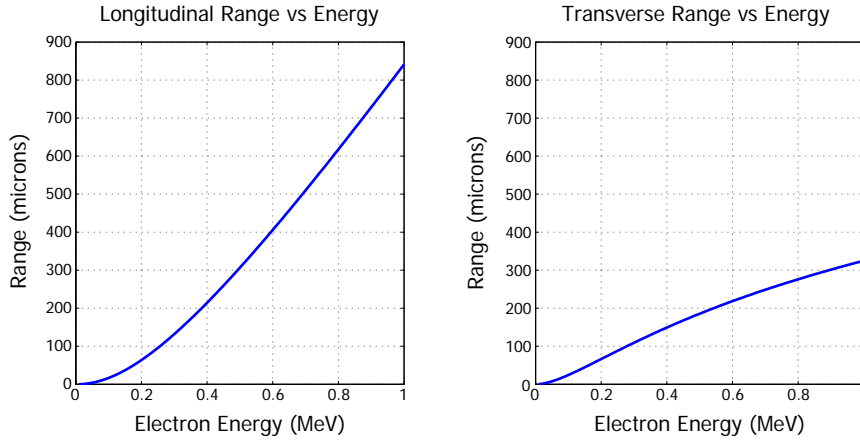


Figure 4.16: Average longitudinal and transverse ranges for electrons of varying energies transiting the YAG:Ce crystal.

back surface, it will deposit none. Averaging over the possible places where the particle could be born, we see that, statistically speaking, 0.25 MeV photons which get absorbed ultimately deposit about half of their energy in the crystal.

At 1 MeV, the longitudinal range is about $850\ \mu\text{m}$, much longer than the crystal thickness. According to [76], the loss of energy by an electron in matter per unit distance is nearly constant with electron energy, varying only logarithmically. We can thus say that the fraction of the electron's energy which is deposited in our crystal for particles born at the front surface should be approximately $113/850 = 13\%$. For particles at the back surface, again no energy is deposited. Averaging, we see that 1 MeV electrons deposit about 6.5% of their energy before exiting the crystal.

Combining the low chance of X-ray absorption and the average energy loss of the resulting electrons, we can calculate the average energy deposited by each photon. For example, the average 1 MeV photon, which gives birth to a ~ 1 MeV electron, ultimately deposits just $0.45\% \times 6.5\% = 2.9 \times 10^{-4}$ of its energy in the crystal.

The green curve in Figure 4.13 reflects the overall energy that is actually deposited in the crystal as a function of the incoming photon energy. Above about 200 keV, where the electron longitudinal ranges start to become significant relative to the crystal thickness, an ever larger fraction of the electrons exit the back of the crystal, causing a rapid drop in deposited power. Once the electron range is significantly

larger than the crystal thickness (above 0.6 to 0.8 MeV), the nearly constant dE/dx of high energy particles in the crystal means that we enter the more shallowly sloped region where the fraction of energy deposited goes roughly as $1/E$.

Photon Production in YAG:Ce and Capture in Our Camera

Knowing how much energy is actually deposited in the crystal, we can calculate how much optical light will come out for us to observe. The manufacturer of our YAG:Ce scintillator claims that per MeV of deposited energy in the crystal, about 8,000 photons with a spectrum centered at a wavelength of 550 nm will be created. Knowing the solid angle subtended by our camera lens, we can then calculate how much light the camera will capture from the isotropic emission. In this case, we expect about half a billion photons to reach the CCD. As discussed in § 4.3.6, this is a good signal. A major concern had been whether this thin crystal would have enough signal for the camera, and it worked well when installed.

Reflections back and forth between the two polished faces of our crystal are a possibility which could reduce the resolution. Additionally, if photoelectrons are continually ejected from the crystal, it will be charged to a very high voltage, potentially leading to damage. To ameliorate both effects, we purchased the crystal with a protective coating. The crystal was coated on both sides with a 75 nm Indium Tin Oxide (ITO) transparent conductive layer. The ability to conduct allows the crystal to avoid possible damage from electrostatic fields. The specific ITO thickness was chosen to make a quarter wavelength antireflection coating for the scintillation central wavelength of 550 nm, where ITO has an index of refraction of about 1.6.

Light leaking out from the edges of the crystal is visible, but in software we choose a portion of our camera images which excludes all crystal edges to remove this as a source of possible error in our measurements.

Having the appropriate crystal, we need a way to hold it above the electron beam and actuate it up and down so that it is close to, but not intersecting the beam path. Given our maximum 2.5 mrad divergence between the X-Rays and the electrons, and about 2 meters of propagation distance, the x-rays which are *furthest* from the electron beam at the crystal are only 5 mm above, so adjustability at the sub-millimeter scale




Figure 4.17: Schematic of the stainless steel vacuum chamber for the scintillator crystal. It provides ports for alignment and viewing by two cameras simultaneously.

is needed. The actual mount for the crystal is pictured in Figure 4.14. To move the assembly up and down above the beam, we use a stepper motor driving a linear actuator that feeds through into the vacuum chamber.

4.3.5 The Vacuum Chamber

In order to hold the YAG:Ce crystal above the electron beam at a variable height and also to be able to view the crystal from various angles, we constructed a special purpose vacuum chamber with a variety of ports for alignment and viewing of the crystal. Figure 4.17 shows a schematic of the arrangement used.

Although only one camera is necessary to view the crystal and provide the high quality images useful for later data analysis, it was useful to add the ability to have a second analog camera viewing the back side of the YAG:Ce crystal. The crystal mount was purposely made in the shape of an upside down “U” (see Figure 4.14). Thus the crystal is visible from both sides, making this possible.

The operators of the main accelerator sometimes used the spectrum information coming from this second camera to tune the beam. As the spectra for different running conditions are distinct, this was a useful diagnostic for them, and could be used at all times, unlike the scientific CCD camera for E164.

4.3.6 The Camera

For imaging the light coming from our scintillator crystal, we used a SenSys camera made by Photometrics. This line of air-cooled cameras uses scientific grade CCD imagers which are sensitive to relatively low light conditions.

We imaged onto the CCD with a standard 105 mm AF Micro-Nikkor lens made by Nikon. This has up to an $f/2.8$ aperture for good light gathering. Due to space constraints, the lens was situated about 20 cm from the YAG:Ce crystal, giving a demagnification of 3 times and a depth of focus of about $100\text{ }\mu\text{m}$ as mentioned above.

Our model, the KAF 1400, has a CCD with 1317×1035 pixels, each 6.8×6.8 microns in size. The quantum efficiency in the region of 550 nm is 42%, and the ADC saturates at between 10,000 and 80,000 electrons, depending on the sensitivity setting. The corresponding noise levels per pixel range between 12 and 20 electrons. We used the High Dynamic Range gain setting with full well capacity of 40,000 electrons and expected noise of 17 electrons. The digitization is with 12 bits, corresponding to about 10 electrons per count. Readout of the camera was performed at up to 1 Hz by a computer located in the FFTB tunnel.

The predicted half billion photons from our scintillator were imaged onto a region of about 300×200 pixels. This gave an expected value of some 3000 photons converted to detectable electrons in each pixel and typical ADC counts of 300. Actual images show a lower average count per pixel of more like 200, with uniform and easily subtracted background counts per pixel of about 60. The actual light seen was roughly half what had been expected. Given the various uncertainties in the calculation of the estimated light yield, this was still reasonably close to expectations, and in practice we had a good signal with no difficulties seeing the spectrum.

4.3.7 Spectrometer Resolution

Computer simulations of the beam properties in the FFTB indicated that the available locations for our spectrometer would provide good spectra. Figure 4.18 shows the Twiss parameters in the available region, and plots the actual location that was finally selected due to space constraints.

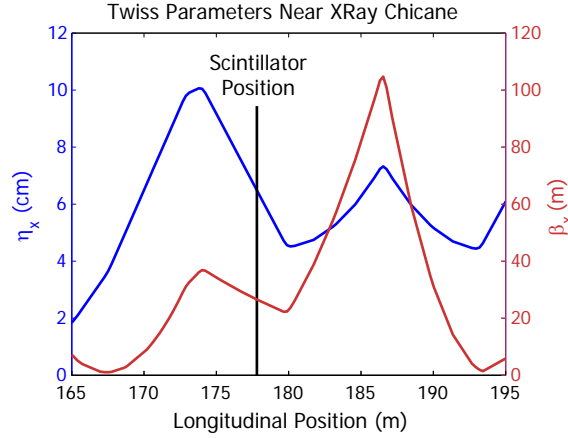


Figure 4.18: Plot of horizontal Twiss parameters at the start of the FFTB.

The fundamental ability of a spectrometer to provide good spectra at any location has to do with the balance between the dispersed and betatron spot sizes there. Our non-zero beta function acts to blur our ability to view the spectrum.

There are also a series of other resolution compromising effects that are unique to our indirect method of observing the electron positions. We discuss each of these sources of blurring in order, starting with the beam properties.

Electron Beam Focusing

The first resolution limiting effect comes about from the imperfect focusing of the electron beam. We recall the usual formula for the spot size in a region of dispersion:

$$\sigma = \sqrt{\beta\epsilon + (\eta\delta)^2} \quad (4.14)$$

where δ is the fractional energy spread as discussed in Chapter 3.

In the circumstance of an ideal cold beam with very small emittance ϵ or a tightly focused beam with small β , the (horizontal, say) position of a particle would be determined solely by its energy, and we map energy to position with perfect fidelity.

For real beams, the $\beta\epsilon$ term will always play a role, and its contribution is to cause uncertainty in particle positions because even particles of the exact same energy will be spread out over the betatron spot size. Mathematically, equation 4.14 can be

thought of as the result of a convolution of the energy spectrum with a Gaussian of width $\sigma_\beta = \sqrt{\beta\epsilon}$. This convolution has a smoothing effect which reduces the information content of the spectrum.

Referring to Figure 4.18, we can see that the dispersion is about 6.4 cm. The RMS (correlated) energy spread is 1.1%, giving a spot size σ_δ from horizontal dispersion that has an RMS of about 1 mm. The spot size from the betatron focusing of the beam is not so much smaller than this value. With typical normalized emittance of about $\epsilon_x = 50 \mu\text{m}$ and $\beta = 26 \text{ m}$, we find $\sigma_\beta = \sqrt{\beta\epsilon_x/\gamma} = 153 \mu\text{m}$. This gives the uncertainty in position for an electron of a given energy, and turns out to be the dominant source of blurring in our spectrum.

Three quadrupoles in the vicinity of the X-Ray spectrometer can be used to alter the R_{56} of the FFTB for varying bunch compression at the plasma. For all data discussed in this thesis, these magnets were unchanged, leaving R_{56} at 1.5 mm, but when R_{56} is set to 1 or 2 mm, the dispersion changes to 5.6 cm and 7.2 cm, respectively.

Electron and Photon Trajectories

The spectrometer magnet is a large object almost 2 meters in length, and the vacuum chamber with our scintillator lies almost a meter downstream. This means that the position where the photons are radiated is in a region of significantly larger dispersion than the location where they are intercepted.

In the horizontal plane, each electron radiates an X-ray which propagates in a very narrow cone centered in the same direction as the electron is moving. This will lead to slight blurring, as discussed next, but even though the photons are born some meters upstream, they are still, on average, in the exact same horizontal place as their originating electrons when both reach the scintillator. This is why we must calculate the dispersion and beta function for the electrons at the location of the detector.

X-ray Production

When the electrons radiate X-rays in our chicane bends, the X-rays do not come off in a perfectly collinear fashion. From each electron, the radiation is expected to

come off in a cone with a spread of angles in a way analogous to the case of Optical Transition Radiation.

The photon cone is centered on the direction of motion of the source charge, and the distribution of angles about that direction is governed by the frequency of the emitted radiation. The RMS spread in angles σ_θ is given by [77]:

$$\sigma_\theta \approx \frac{1}{\gamma} \left(\frac{\omega_c}{3\omega} \right)^{\frac{1}{2}} \quad (4.15)$$

where γ is the relativistic factor of the emitting electron while ω_c is the critical frequency for the synchrotron radiation and ω is the frequency for which we seek to know the opening angle of our cone.

Even for photons with only four percent the critical energy (of which we do not have many), the opening angle of the cone of radiation is only about three times as large as the value $1/\gamma$, and all photons with higher than the critical energy are even more collimated. As our relativistic factor is about 56,000, even the lowest energy photons come out in a cone narrower than $60 \mu\text{rad}$ and the vast majority of the power comes out within an angle more like $20 \mu\text{rad}$.

With some 2 meters of propagation to our scintillator screen, the RMS spread of the X-rays is of order $40 \mu\text{m}$ and this blurring effect is much smaller than that from the imperfect focusing of the electron beam.

X-ray Absorption

We discussed the thickness of our scintillator crystal above, and one of the reasons we chose to use a thin crystal came from concerns about resolution. We refer again to Figure 4.16 in order to estimate the blurring arising from the fact that electron trajectories are not collinear with the incident X-Rays.

As noted in § 4.3.4, above about one quarter of an MeV, the average electron has a longitudinal range greater than the thickness of our crystal as presented to the photon beam. Understanding the resolution limits for electrons at this threshold in energy is useful, as it represents nearly the *maximum* possible blurring of our signal.

Having a longitudinal range slightly exceeding the $113 \mu\text{m}$ thickness of the crystal

as presented to the photon beam, every electron above 270 keV of energy will eventually exit the crystal. If that electron is born at the front surface of the crystal, its average transverse deflection is about $100\text{ }\mu\text{m}$ before leaving the crystal. That electron can be born with essentially equal probability anywhere through the thickness of the crystal, and particles originating at the back surface of the crystal clearly experience no transverse deflection before exiting. The average transverse deflection within the crystal experienced by each 270 keV electron is thus $50\text{ }\mu\text{m}$.

Electrons with less energy simply have a total transverse range which is smaller. Competing with that reduction in total range is the fact that a higher fraction of these lower energy electrons will traverse their entire range within the crystal, once their longitudinal range shrinks below the crystal's thickness. Thus, they will also go fully as far to the side as possible. It turns out that the increasing fraction of particles going to their full transverse range is overwhelmed by the rapid diminution of that range. For example, at 100 keV, most of the electrons have a longitudinal range short compared to the crystal's thickness and will be stopped in the crystal. However, those particles have a total transverse range of under $30\text{ }\mu\text{m}$ and will therefore cause less blurring than the average electron at 270 keV.

For electrons with greater than the threshold of 270 keV, we know that they are emitted with angles ever more closely parallel to the direction of the original incoming photon. Although their overall ranges are longer and longer, the transverse deflections achievable in our thin crystal will progressively be reduced because the typical angles shrink and the thickness of our crystal is fixed. For example, at 500 keV, the typical angle of emission is only 32° and an electron born at the front surface only goes $60\text{ }\mu\text{m}$ to the side before shooting out the back of the crystal. Again averaging over the thickness of the crystal, electrons at half an MeV only blur by $30\text{ }\mu\text{m}$.

The maximum blurring from the electron trajectories is about $50\text{ }\mu\text{m}$ and the average effect will be smaller, perhaps half as much. This is a minor source of blurring.

This estimation was borne out in the course of E164. In 2003, we had used a thicker crystal of $200\text{ }\mu\text{m}$. Changing in 2004 to the $80\text{ }\mu\text{m}$ crystal made no observable change in the resolution, implying that even the thicker crystal had blurring which was small compared to that from the imperfect electron beam focusing.

Blurring From the Camera

We took several images of the YAG:Ce scintillator and its holder to focus the images and check resolution. We can see the 1 by 1 cm YAG crystal in full frame views from the camera. This crystal spans about 880 pixels as viewed by the camera, so the distance between points on the crystal viewed by neighboring pixels is $11.4\ \mu\text{m}$.

When we look at sharp edges in images taken of the entire aluminum holder and crystal, the transition is clearly visible within two pixels, implying that blurring from the lens is at or below that scale, ~ 25 microns. This is also minor compared to the major resolution-limiting effect in the electron beam itself.

Total Blurring Effects

Nominally, we can resolve light coming from our scintillator that originates from points only $11.4\ \mu\text{m}$ apart, but the actual resolution is poorer due to blurring.

As mentioned above, we add the various blurring sources in quadrature. When we combine the four blurring effects, of size 153, 40, 25 and 25 microns, respectively, we come out with a total resolution wash-out with an RMS of $160\ \mu\text{m}$ corresponding to 14 pixels on the camera.

Energy Per Pixel in the Camera

At the scintillator, models of the electron optics indicate that the dispersion for our 28.5 GeV beam is near 6.4 cm for the common case of $R_{56} = 1.5\text{mm}$, so we would expect that each pixel will span about 5 MeV.

The dispersion at the location of our X-Ray based spectrometer is not well measured, so we do not have a good estimate of the accuracy of the simulations. In practice, a good way to calibrate this resolution directly is to compare it with the results from the Cherenkov spectrometer downstream from our plasma where the optics are well understood and have been measured many times. We have discussed the Cherenkov spectrometer and its calibration of 18.8 MeV/pixel in § 4.1.4. Because the X-ray based spectrometer is non-invasive, we can compare the energy spectra at both screens for the same shot to determine the relative calibration.

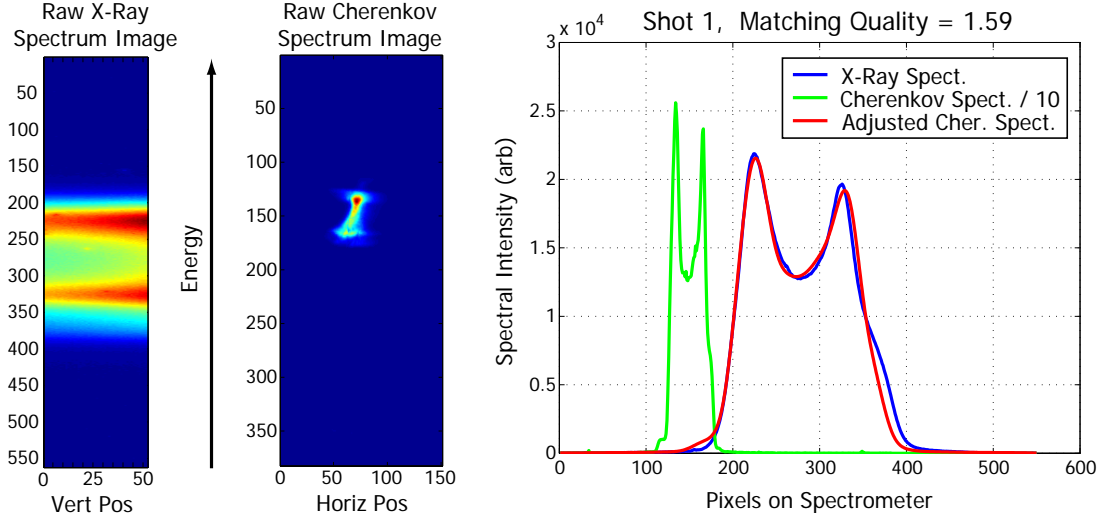


Figure 4.19: Comparison of X-Ray (left) and Cherenkov (center) based electron beam energy spectrum measurements. At right is a graph showing the two raw spectrum curves (blue and green) with the stretched and recentered Cherenkov spectrum (red) matching well. To match the X-Ray’s lower resolution, the green curve is convolved with a Gaussian of 10 pixels width after it is stretched horizontally.

Although we take full images of the beam, we are interested in just the energy spectrum curve. This is made by summing the images sideways to find the number of particles at a given energy. We can see example images for the same data event as viewed by the X-Ray and Cherenkov cameras in Figure 4.19. In both images, higher energy particles are at the top. We typically plot the energy spectra horizontally for ease of viewing and note that the raw data from our cameras has the high energy particles at *low* numbers. Both cameras have uniform background light and electronic noise, so we subtract the background from each image by taking the average values at the beginning and end of our curves.

When we plot the energy spectrum curve for these two diagnostics together, the curves can have an arbitrary horizontal offset in the central energy and the overall amplitude of the two spectra will be unrelated. On the right of Figure 4.19, we plot the spectrum from the non-invasive X-Ray spectrometer in blue and that from the Cherenkov camera (divided by 10 in y for ease of plotting) in green. The red curve is the result of our fitting routine as described below.




Figure 4.20: Matching Quality between the Cherenkov and X-Ray spectra as a function of the amount of stretching applied to the Cherenkov.

Clearly, the overall width of the Cherenkov curve is smaller than that of the X-Ray, so we are interested in the stretching ratio that will make them overlap most precisely, and try several values to see which gives best agreement.

To do this, we stretch the Cherenkov curve by a fixed ratio, which we scan between 2.8 and 4 times, by using a cubic spline interpolation and resampling. Then we blur the image by a Gaussian of a typical width of 10 pixels, to recreate the reduced resolution images from our X-Ray spectrometer.

Next, we use a fitting routine to align the two spectra horizontally and scale the Cherenkov image vertically to match the X-Ray. The parameter which is minimized by the fitting routine is the sum of the squared differences between the points on each curve, similar to a standard χ^2 metric.

For each of our stretching ratios, we go through a 100 shot run of data and align the Cherenkov and X-Ray images, calculating the “Matching Quality,” where smaller numbers reflect better matching. We then average for all 100 shots and plot it versus the stretch ratio that had been used. We see the results in Figure 4.20.

The best overlap happens when the Cherenkov images are stretched by 3.4 times, meaning that the X-Ray spectrum has a dispersion of 5.5 MeV per pixel. This is not

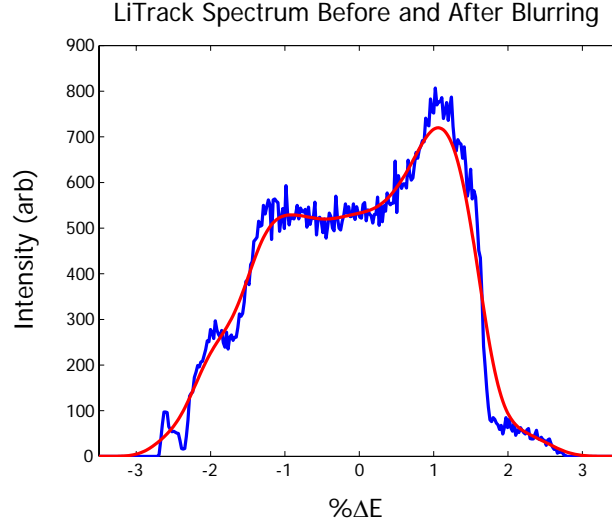


Figure 4.21: Example spectrum from LiTrack before (blue) and after mathematically blurring the image (red) to mimic the physical loss of resolution in the energy spectrometer. The blurring is performed by convolving a Gaussian of width 0.27% with the raw output from simulation, and several features are lost.

in complete agreement with the expectation from ELEGANT, and implies that the dispersion at our X-Ray about 10% smaller than expected. Because it is more closely based on a directly measurable quantity, we choose the value of 5.5 MeV/px, with an estimated error of 0.5 MeV/px.

Overall Energy Blurring Effects

Our estimated 14 pixels of uncertainty correspond to about 75 MeV, or 0.27%. This will cause washing out of some small features, but the major elements of our 4% full width energy spectra will still be visible. Better electron optics could reduce the dominant source of blurriness to the point where the other effects must be taken into account, but they are not of significant concern here.

Figure 4.21 illustrates the loss of resolution with a sample simulated spectrum from LiTrack before and after the expected blurring. The major features remain visible, but we lose information about the small peaks and valleys in the spectrum.

Knowing that there will be limited resolution, we have several options when we seek to compare data with the perfect fidelity spectra from simulation. As was used

in the relative calibration of the Cherenkov and X-Ray spectrometers, the technique employed in this thesis is simply to take the “ideal” spectra output by LiTrack and blur them out as shown in Figure 4.21. Then a comparison between simulation and experiment will be on an equal footing.

If the electron optics were better, the blurring could be reduced by several times, allowing finer details of the spectrum to be visible. That greater information would allow even better precision in the comparisons with simulation discussed next.

4.4 LiTrack

We now discuss some of the specifics of the simulation program LiTrack. This software suite for simulation is used to understand the beams from the SLAC main accelerator, and provides the energy spectra against which we can compare our data.

LiTrack was developed by K. Bane and P. Emma of SLAC [78]. Originally written for IBM VM machines, this code was rewritten more recently in MATLAB. LiTrack has been benchmarked against ELEGANT [66] for accuracy. Although it is a general code for linear accelerators, it was originally developed for modelling SLAC.

In many accelerators, longitudinal phase space can be investigated independent of the transverse phase space properties. LiTrack tracks particles only in the 2D longitudinal phase space given by z and $\delta \equiv \Delta E/E_0$, for a large increase in speed.

Acceleration is represented as a pure sinusoid without approximation, and compression from any optical element with bends is represented directly by the path length versus energy coefficients up to third order, e.g. R_{56} , T_{566} and U_{5666} . Wakefields are included as the point charge wake, which is then applied to the evolving time profile of the bunch. The effects of synchrotron radiation are also included.

Unlike slower particle tracking codes such as ELEGANT which explicitly define every optical element in the beamline, LiTrack combines many elements into single units. For example, the first third of the accelerator is rendered as a single object 809.5 meters long which has a final energy of 9 GeV and accelerates using an rf wavelength of 10.4969 cm. Longitudinal wakefields as discussed in Chapter 3 are calculated for the SLAC linac and their effect is included as a single element.

Similarly, the entire RTL is rendered as a single object with a defined R_{56} , T_{566} and central energy for the beam. The U_{5666} is small and ignored. For complex beam lines such as the RTL, these values are normally calculated using another code, such as ELEGANT. Lumping elements like this greatly speeds execution compared to alternative codes such that a typical simulation of the entire SLAC accelerator from the damping ring to the FFTB with 400,000 macroparticles takes about 25 seconds on a PC having a 2.4 GHz Pentium processor and 512 MB of memory.

The number of macroparticles needed for good accuracy depends on the complexity of the system being studied. For conditions similar to those that match our experimental data, experience shows that the results are fairly consistent beyond 100,000 particles and do not change once the number of macroparticles grows beyond about 300,000. To maximize accuracy without unnecessary computational time, we use 400,000 particles for the simulations discussed in this thesis.

With fast execution, it is feasible to simulate the accelerator many times with slight variations in the input parameters. However, there are usually ten parameters which can be adjusted. If each of them were scanned over only three values, we would have 3^{10} , or nearly 60,000 simulations to perform for all possible combinations.

In practice, we focus on several parameters which are likely to vary more than the others to reduce the number of simulations substantially. Even so, performing a set of simulations with several hundred members is commonplace, and would be impractical with a slower code than LiTrack.

Based on measurements in the accelerator, we have an approximate idea of the various input parameters for LiTrack, but the specific combination of parameters for any given data shot is subject to jitter in the machine, so cannot be measured with required accuracy. Thus, we need to simulate many possible states for the machine before trying to match to any given measured energy spectrum.

For a given set of data, we initially perform a rough scan, changing the accelerator parameters by comparatively large values. Once we find simulations that appear to match the data reasonably well, we redo the set of simulations with finer changes in the parameters of interest. This process is usually iterated several times to accurately find the phase spaces corresponding to a particular set of data.

Parameter	Typ. Value	Range	Units
Number of Particles	1.9	± 0.2	10^{10}
Bunch Length From Damping Ring	5.5	± 0.2	mm
Bunch Asymmetry From D.R.	-0.27	± 0.03	-
Phase From D.R.	1.0	± 0.5	mm
Compressor Cavity Voltage	42.5	± 0.5	MV
Energy Collimation in RTL	2.8	± 0.2	%
RTL R_{56}	-58.8	± 0.2	cm
Sub-Booster Phase	-19.25	± 0.25	degrees
Overall Linac RF Phase	1.0	± 2.0	degrees
FFTB R_{56}	-1.5	± 0.5	mm
Compressor Phase	90		degrees
RTL T_{566}	105.4		cm
Mean Energy at Sector 10	9		GeV
Sector 10 Chicane R_{56}	7.6		cm
Mean Energy at End of Linac	28.5		GeV
FFTB T_{566}	4		mm

Table 4.2: LiTrack parameters for simulations in this thesis. The top section shows variable values, and the bottom shows fixed parameters.

4.4.1 Parameters in the Linac

For the SLAC main accelerator simulation in LiTrack, ten parameters can be varied to find conditions matching those of the real machine. These parameters, their typical values and typical ranges are given in the first part of Table 4.2. These parameters could be substantially different for other experiments, but the listed values and ranges reflect conditions pertaining to E164 data runs as discussed in this thesis. The parameters which are expected to stay fixed are listed in the second part.

We discuss each accelerator parameter below, and for those which affect timing in the accelerator, we refer to Figure 4.22 to visualize their relationship.

The number of particles in the accelerator tells LiTrack how strong the wakefields will be. Although some investigations were performed with as few as 6×10^9 , normally we ran with as many electrons as allowed by instabilities in the damping ring, in the range of 1.9×10^{10} . This particular parameter is measured by various toroids, and is known for every data shot to within a few percent.




Figure 4.22: Timing schematic for LiTrack’s representation of the SLAC main accelerator in E164. Positions above the reference particle line indicate early arrival.

The bunch length and asymmetry of the beam coming from the damping ring are not independent parameters, but are determined by the total bunch charge that we have. We recall Figure 3.3, showing the measured values as a function of current. There were possible systematic errors, so the exact values are not known for a given bunch charge, but our best estimate is that for bunches of 1.9×10^{10} particles, we should use a bunch length of 5.5 mm and asymmetry of -0.27. Simulations with these values have agreed well with the data.

When we specify phases, they can be thought of as being taken relative to the timing of a speed of light particle traversing the main accelerator which had hit the compressor cavity exactly at zero crossing. This speed of light particle in the simulation is equivalent to the Master Oscillator in the real accelerator, which propagates timing information down the full length of the machine.

LiTrack treats the Compressor Cavity as always fixed in phase relative to the reference particle, but we may inject at various times into the cavity. The parameter known as “Phase From Damping Ring” is what we use to change injection timing, and has units of mm. This seemingly odd choice of units may be because LiTrack deals directly with longitudinal coordinates z and δ .

We normally inject the electron beam somewhat after the reference particle. This is the parameter known as the phase from the damping ring, and is a positive quantity

expressed in millimeters, with a value typically around 1 mm. Thus, there is a slight net deceleration of the beam in addition to the large energy chirp. This timing offset on injection into the accelerator propagates all the way through the main accelerator.

The voltage of the Compressor Cavity affects the first stage of compression. For different desired levels of compression, we can change this value, but it remained very close to 42.5 MV during the experimental run discussed in this thesis.

Although it was never changed during the experiments relevant to this thesis, the exact value of the R_{56} is not known exactly. It has been calculated to be in the region between -58.6 and -59.2 cm, and we have found best agreement with data using a value of -58.8 cm in LiTrack, so we use that value and do not scan this parameter in our series of simulations. The T_{566} value has been calculated to be -1.054 m and is also not scanned here.

The energy collimation in the RTL does not come from a specific object. Rather, it reflects the loss of some particles as they transit that complicated beamline with various places of reduced aperture. Because we do not have a specific aperture, the energy collimation that occurs throughout the RTL is also not known exactly. Good agreement was observed with an energy acceptance half width of 2.8%, and this parameter was not scanned.

The overall linac phase is phase shifted relative to the Master Oscillator in the real accelerator. In LiTrack, the speed of light particle traversing the compressor cavity at zero-crossing will go through the linac on crest if the overall linac phase is set to zero. When this quantity is positive, the reference particle will go through the linac behind the wave crest, because a positive value means that the phase has advanced ahead of the nominal position.

The first third of the linac is run, as discussed in Chapter 3, at -19.25° , placing the beam ahead of the rf crest in order to give the tail particles higher energy than the head for compression in the compressor chicane. This normally fixed value is known as the sub-booster phase, and is not absolute relative to the reference particle, but rather is added to the overall linac phase. Thus, if the overall phase were $+2^\circ$ and the sub-booster phase were set to -19.25° , then the actual rf crest in the first third of the linac would be only -17.25° behind the nominal reference particle. We recall that

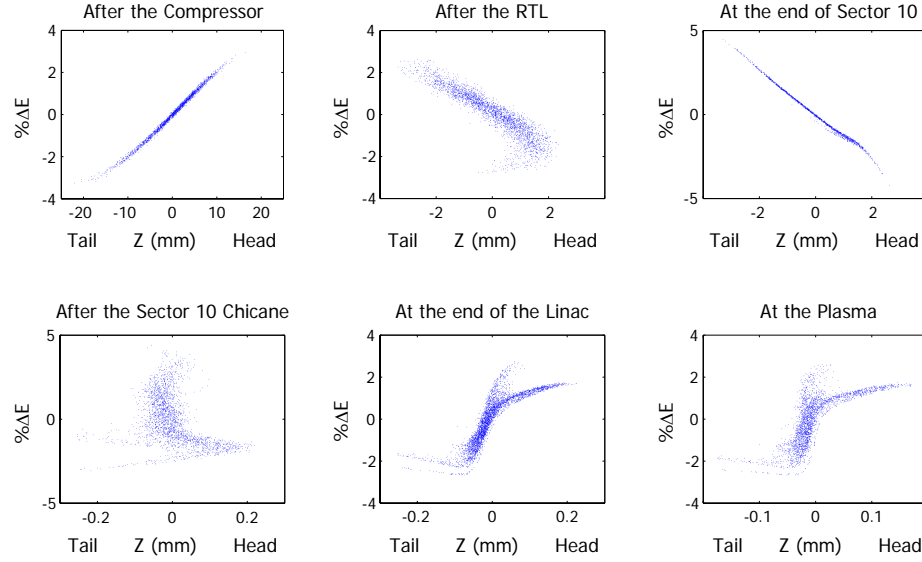


Figure 4.23: Progression of the beam phase space through the linac.

the bunch leaving the damping ring is usually slightly behind that reference particle, so the center of the bunch will be even *less* than 17.25° ahead of the wave crest.

The rf in the last two thirds of the linac is simply at the overall phase, and so bunch is only a few degrees behind the rf crest for near-maximum energy gain. The energy correlation we impart comes almost completely from the longitudinal wakes.

Sample phase spaces from the various elements of the accelerator have been shown in Chapter 3, and we include them again in Figure 4.23 to show the evolution of the phase space clearly. Note that energies are expressed as a fraction of the central energy, which changes by a factor of about 25 through the linac.

We have seen the evolution of a beam responding to the many elements in the main accelerator as illustration of the physics. The plots of the phase space in Chapter 3 come from particle dumps out of LiTrack at various points through the accelerator.

4.5 Finding the Phase Space for Data Events

When we seek to know the phase space of the electron beam for any given shot of the linac, we must compare its energy spectrum with phase spaces from simulations.

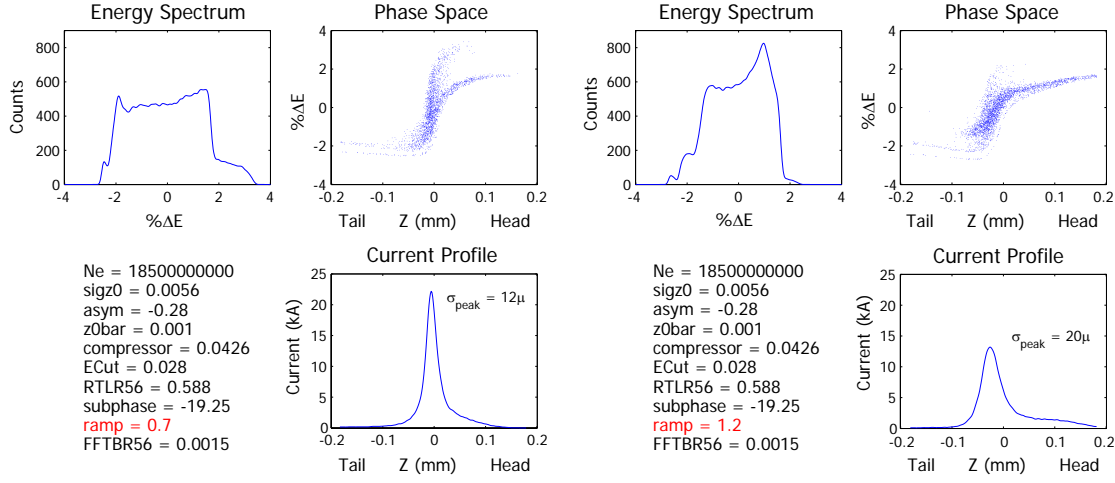


Figure 4.24: Plots of energy spectra and associated current profiles for two different values of overall linac phase. Less than half a degree leads to substantially different beams in the FFTB.

4.5.1 Simulation Parameters for this Thesis

As mentioned above, we do not know the exact state of the main accelerator at any given time. Luckily, we do know the approximate state, and therefore, we allow some of the parameters discussed in the previous section to vary. The overall linac phase, in particular, is known to shift in a range of a degree and sometimes more over time scales as short as a minute. Any given data run of, for example, 100 shots will take 100 seconds. During that span of time, there can be variations in phase of at least half a degree between the various shots. This half of a degree in overall phase can lead to a substantially different beam coming from the accelerator. Figure 4.24 shows the difference that 0.5° can make to the energy spectrum and current profile, even if all other parameters are held constant.

The other parameters above are less likely to vary on scales with such dramatic impact, but several have some effect. For the investigations discussed in Chapter 5, we performed a suite of simulations where three possibly varying parameters were allowed different values. In particular, we simulated conditions where the phase from the damping ring varied between 0.7 and 1.1 mm. The uncertainty in the compressor cavity voltage is larger than that required to change the final phase space, so we allow

it to have values of 42.3, 42.5 and 42.7 MV. Lastly, the overall linac phase varies over roughly a one and a half degree range between 0.3° and 1.9° .

In this instance, there are a total of 255 simulations we can compare against each data shot of interest. It is certainly possible to create a vast suite of simulations by varying all of the parameters in the simulation, but these are the three that experience has taught are most likely to vary and have an effect on the final phase space.

4.5.2 Comparison of Simulation and Data

After performing the series of simulations, we save a file which has information on the input conditions to LiTrack for each simulation, as well as each resulting energy spectrum and current profile. With this data file in hand, we can then compare the simulations to data in a separate program.

To perform the comparisons, we have a program written in MATLAB which first loads the incoming beam energy spectrum images which have been stored for later analysis such as this. It then sums the images sideways as discussed above to get spectrum curves that span several hundred bins. There is some background light in the images, so that is removed by taking the average of the first ten plus last ten bins and subtracting that value from each point.

Next, we load the suite of simulations previously performed. We bin them with our 5.5 MeV/pixel calibration to match the binning of the data.

Because we know that the data has some blurring in it, we then convolve each simulation with a unit area Gaussian whose RMS width is 10 pixels, or 0.21%.

At this point, there remain only two unknowns for the comparison, as with Cherenkov vs. X-Ray comparisons discussed in the section on calibration of the energy per pixel on our camera. There are energy feedbacks within the accelerator, and sometimes individual klystrons will go off line. As a result, the centroid energy of our electron beams varies by up to a few hundred MeV out of 28.5 GeV. This manifests itself as the spectrometer images moving up and down modestly. Summed sideways, and plotted, we see the spectrum curves moving sideways by a noticeable amount, though it is still smaller than the overall typical width of the curves.

Small changes in the total energy of the beam in the linac should have no significant effect on the various effects which allow us to produce short beams at the end of the FFTB. The simulations are always, by construction, perfectly centered on $\delta = 0$. Therefore, each simulation is allowed to move back and forth in the matching routine to get the best overlap with the data.

In the vertical direction, the scales are quite different. The spectra come from 12 bit digitized images. Each pixel therefore has a value that can range as high as 4095 and we sum a whole row of pixels, so the points on our spectrum can be over 100,000.

The simulations are typically run with 400,000 macroparticles. When binned across several hundred bins, the spectra have a maximum number of these particles per bin that is typically less than 1000. So in the fitting routine, the simulations are simply multiplied by a constant to account for the mismatch. The seed value for this constant can be calculated by the sum of the data spectrum curve divided by the sum of the simulation spectrum curve.

With approximate initial values, we can then perform the comparison. The heart of the program is a routine which takes the two curves, and calculates the χ^2 between them, seeking to minimize this value by varying the sideways shift of the two images and their overall vertical size. To calculate the χ^2 , we simply take the raw difference at each bin, square it and then sum over all bins. The only difference from a standard χ^2 is that we do not normalize by the errors, as they are not known here and are assumed to be constant. For shorthand, we will refer to this value as the χ^2 with the understanding of its modified meaning. The process is nearly identical to that described above for comparison of spectra from the Cherenkov screen and the X-Ray based spectrometer, and output from the fitting routine looks much like Figure 4.19.

To minimize χ^2 , we give the fit the ability to vary only the two parameters mentioned above. When the best overlap is found by the fitting routine, we record the χ^2 that was achieved between that particular simulation and the data spectrum. Clearly, if the simulation spectrum has a shape and overall width that is similar to that of the data, we get good overlap, and if the shape or overall width is wrong, then no amount of stretching or translating the curve sideways will allow the χ^2 to be small and the fitting routine will return a large value from the optimization.

Having found the best χ^2 for a given simulation and the data, we then compare the same data against the next simulated spectrum. We record which simulation yielded the smallest χ^2 and then repeat the entire process for the next data shot.

The number of times the fitting routine is called goes as the number of data shots times the number of simulations, so it is beneficial to have as small a number of simulations against which to compare the data of interest as possible, while still spanning the space of accelerator parameters that pertains to any given run. As mentioned in the previous section, 255 simulations was enough to provide a decent match to almost all of the data discussed in this thesis.

Having performed the comparisons, we then save a file with the information about the simulation which matched any given data shot, as well as the quality of that matching. Separately, we calculate properties of the images from the Cherenkov spectrometer and store them.

A suite of MATLAB analysis programs has been compiled by our group at SLAC, Accelerator Research Department B, and all data is formatted so that it can be read in and analyzed using this suite. We can thus easily plot calculated parameters of the bunch against measurements from Cherenkov images or from the accelerator, such as the charge or total amount of CTR power produced by any given bunch. This software is the source of many of the graphics in Chapter 5.

With all of these tools in hand, we can now investigate the longitudinal phase space and its effect on ionization, acceleration, and other interactions of the beam and the plasma.

