Chapter 5

Data and Results

Knowing the longitudinal phase space of the electron beams in the E164 experiment is important for understanding a variety of effects that occur in the plasma, primarily, the acceleration of the beam itself. To have confidence in the technique, we first seek to find verification that matching with LiTrack provides accurate data. There are several ways to test this method.

5.1 Verifying the Technique

As our technique is indirect, we have sought a variety of sources of confirmation that it is accurate. When we calculate the beam profiles, we can then correlate the properties of each bunch with several other measurements. If these all track one another in the expected ways, that greatly increases confidence in the method.

To evaluate the accuracy of our technique, we use data from several runs taken in 2004. As illustration, we focus on several 100 shot sequences of data taken at 1 Hz on 13 July of 2004, known in our database name convention as runs "07131cw" and "07131dc." In all of the runs from this day, the accelerator was set to create the shortest bunches possible to investigate acceleration at the very highest plasma densities achievable with our oven: 3.5×10^{17} cm^{-3}. The natural variations of the accelerator's various parameters mean that we would expect a variety of beam profiles at the position of the plasma.

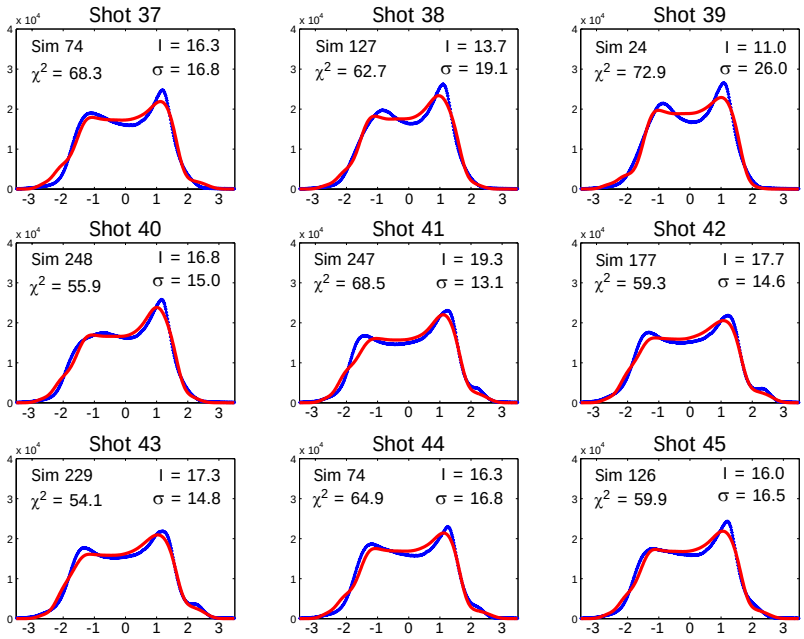

Figure 5.1: Energy spectra for nine sequential shots in Run 07131dc, each matched to one of the 255 simulations with slightly varying parameters. Horizontal axis is in units of δ. Blue is data and red is the best matching simulation spectrum. The peak current in kA and bunch width in μm for the corresponding simulation are indicated.

With all parameters close to optimal, we produce beams as short as 12 μm with corresponding peak currents of about 20 kA. When any of the various linac parameters discussed in Chapter 3 jitter away from optimum, the accelerator produces substantially longer bunches with less intense peak currents. This natural variation allows us to probe several phenomena of interest and is what allows us to see if calculated bunch properties do, in fact, track other measured quantities.

Figure 5.1 shows the matching of our simulations to nine shots from Run 07131dc. Each shot has a slightly different energy spectrum, but still matches one of the simulations, which are different from one another in just a few of the possible variables. Specifically, the compressor cavity voltage is not well known, so is allowed to take the values 42.3, 42.5 and 42.7 MV. Similarly, the phase of our bunch coming from the damping ring is allowed to vary somewhat in the range from 0.7 to about 1.4 mm. Lastly, the overall phase of the linac was scanned over a range of about one and a half degrees, with most shots determined to lie within a range of 0.4°.

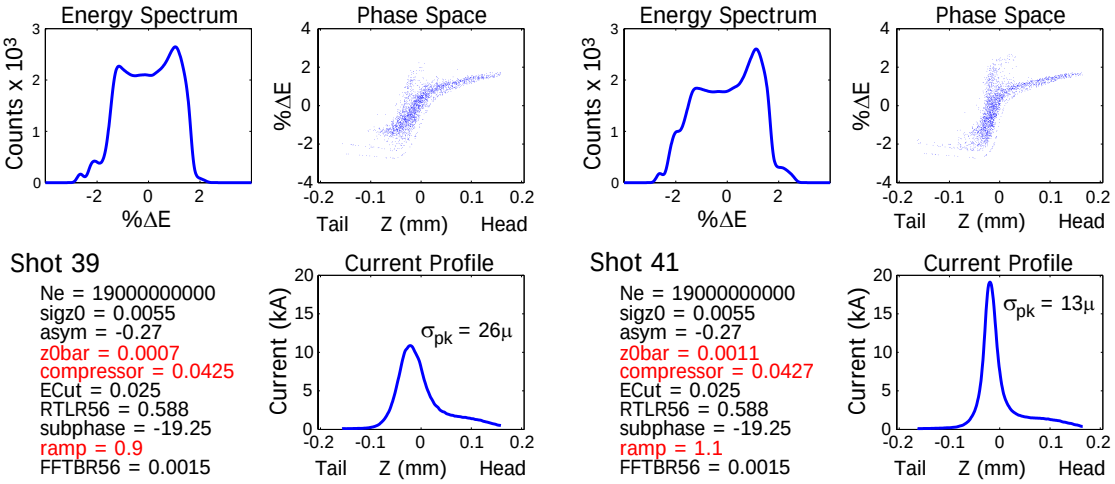

Figure 5.2: Simulation phase spaces, the left shows the long current pulse associated with Shot 39 of Figure 5.1 and the right is the short pulse matching Shot 41.

The overall phase of the accelerator has been measured to vary on short (less than 30 second) timescales, so we always recreate this variation in the suite of simulations for matching. The compressor cavity temperature is also known to vary, sometimes rapidly, when PEP requires changes in the main accelerator. Similar changes are observed throughout the linac, so there are both long term diurnal drifts and much more dramatic jumps of as much as a full degree in just one or two seconds. As discussed in Chapter 4, changes of even 0.1 or 0.2 degree can make appreciable changes to the energy spectrum. This parameter has the largest variability compared to the scale at which changes in it make observable differences in the energy spectrum. Overall linac phase must always be scanned in doing a set of simulations.

For reference, Figure 5.2 shows the phase spaces calculated for shots 39 and 41. They represent occasions where small changes in the linac led to peak currents which were a factor of two different, even though the initial conditions were similar on the scale of our ability to measure the three important quantities highlighted in red.

Determination of the current profile allows us to make a variety of predictions about the behavior of the electron bunches in the plasma and at other diagnostics. We discuss four different ways in which we connect the determined profiles with other measurable quantities, as a means of verifying the technique.

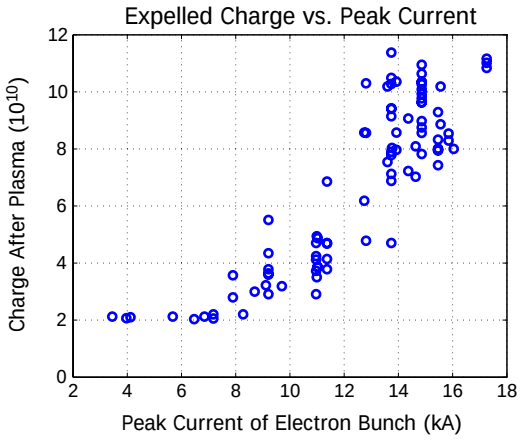

Figure 5.3: The total charge measured downstream of the plasma is related to the peak intensity of the incoming electron beam. Data is from Run 07131cw.

5.1.1 Trapped Charge Measurements

The mechanism is not yet understood, but we have observed that immediately downstream of the plasma, more charge often comes out than was sent in, by a factor of up to five or even more. This is probably some sort of trapping of plasma electrons by the strong wake, analogous to that seen in laser wakefield acceleration experiments.

Trapped particles must start from rest, rather than ~ 28.5 GeV, ending with only a few GeV of energy. They will be lost in the sequence of strong quadrupoles downstream of the plasma and are not seen at the diagnostic Cherenkov screen.

The downstream charge measuring toroid is close to the plasma, so could be influenced by the significant radiation there. We also cannot yet measure the energy of these particles, which could allow us to understand sources for what we see.

Nevertheless, for the above data run, if we plot the peak current of each individual machine shot against the measured charge flowing down the pipe immediately after the plasma, we clearly see that the more intense incoming bunches associate with much larger amounts of charge downstream. See Figure 5.3.

The exact mechanism of trapping is not understood, but it seems reasonable that a more intense electron bunch will drive the plasma wake harder and be more likely to trap particles. As we have substantial variation in the length of the 100 bunches in this run, we have a natural way to see the effect of changing bunch length on

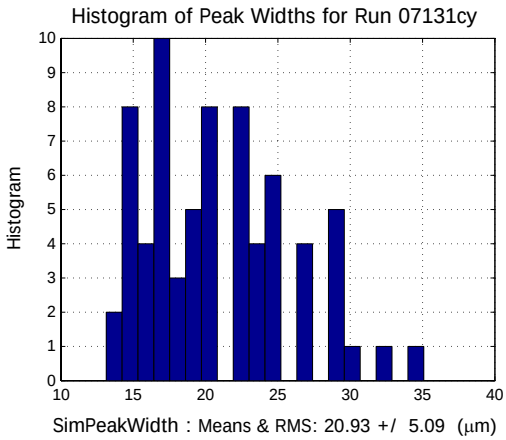

Figure 5.4: Histogram of the calculated bunch lengths in microns for Run 07131cy.

trapping. This provides an interesting result as well as evidence that the simulations do give an understanding of the beam's incoming phase space.

5.1.2 Autocorrelation Measurements

The autocorrelation discussed in § 4.2 is a multi-shot measurement that gives some idea of the bunch length for a given set of accelerator conditions. The measured bunch length is of the order 18 μm long, in reasonable agreement with the shortest bunches that simulation tells us that we can achieve.

As we have seen, the accelerator varies in its input conditions over time scales shorter than the multiple minutes required to build up the autocorrelation trace. Whenever the accelerator strays from ideal conditions, the bunches will be longer.

For example, if we look at a histogram of the peak widths, as determined by comparison with simulations, from Run 07131cy, we obtain Figure 5.4. This shows that we get substantial variation of bunch lengths from roughly our expected minimum of 12 μm to 30 and 40 μm. The most common length is of order 20 μm.

We recall that in choosing events to include in the autocorrelation trace, we had used the most probable total CTR power value as a cut. If total CTR power is inversely related to bunch length, we expect that we had selected the most common bunch length to measure. The autocorrelation value of 18 μm for that bunch length

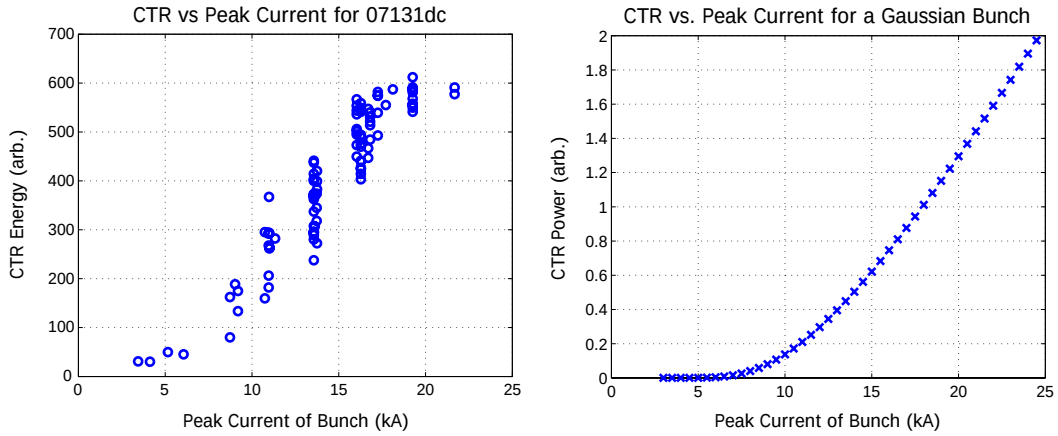

Figure 5.5: Plot of total Coherent Transition Radiation power as a function of the peak intensity of the driving electron bunch. Data is from Run 07131cw.

therefore agrees quite well with the bunch length derived from simulation. This gives confidence in both methods, as they are totally independent of one another.

5.1.3 Pyro Peak Measurements

The autocorrelation measurements give confidence that our short bunches have an overall absolute length close to that estimated from the simulation comparisons. A further piece of useful information comes from the simpler measurement that measures only the total CTR power coming from the bunch. Clearly, we cannot perform the bunch length scan for each shot of the accelerator, but we can measure the total broadband power emitted by each individual bunch on a shot by shot basis.

To find the expected total power for bunches of various lengths, we integrate Equation (4.5) numerically. For each bunch length, the minimum wavelength is taken to be 0.6 times σ_z. As discussed in § 4.2, wavelengths beyond about 100 μm are progressively attenuated by the apparatus itself, so we choose to stop the integration at 150 μm.

We can integrate over all angles, and the effect of the acceptance of our system is taken care of by the choice of limit on wavelength above.

The left hand side of Figure 5.5 shows our measured CTR power plotted against the peak current as inferred from simulations for Run 07131dc. Clearly, short high

current bunches produce more CTR than the longer low current bunches. The right hand side of the same figure shows the theoretical curve we expect.

That the total CTR power we detect as a function of the bunch's peak power follows the general trend that we expect from theory gives confidence that we do understand the longitudinal profile of our bunches.

5.1.4 Comparison with Ionization Measurements

A less direct verification of our understanding of the beam properties comes from understanding the process of Lithium ionization. Because we seek to use our intense beams to ionize the Lithium, it is important to verify the threshold at which ionization occurs. One of the ways to investigate this is to change the electric field of the bunch in a controlled way and see at what point ionization is initiated.

Caolionn O'Connell, in her thesis on the E164 experiment [41], investigated the onset of Lithium ionization for different beams. Understanding the bunch longitudinal profile was particularly useful for the case where the beam was changed in length to see when it would finally have strong enough fields to dissociate Lithium.

With the high charge bunches that we normally use: 1.8×10^{10} particles, the ionization of Lithium has been observed to happen readily, so in this investigation of the threshold for ionization, the bunch charge was dropped to about 0.9×10^{10}.

The spot size at the entrance to the plasma was held at approximately $15\,\mu$m. Thus the peak fields associated with the different bunches were determined only by their instantaneous currents.

As we have seen in § 2.3, Lithium ionizes rapidly once the space charge fields of the beam rise above roughly 5 to 6 GV/m. Given the instantaneous current of the highly relativistic bunch, then the radial electric field at all points in the associated infinitessimally thin pancake at that z location is given by:

$$E_r(r,z) \;=\; \frac{I(z)}{2\pi\epsilon_0 c\,\sigma_r^2\,r}\int_0^r e^{-\frac{r^2}{2\sigma_r^2}}\,r\,dr \;=\; \frac{I(z)}{2\pi\epsilon_0 c}\;\frac{1}{r}\left(1-e^{-\frac{r^2}{2\sigma_r^2}}\right) \tag{5.1}$$

We recall that the pancake's peak field occurs at $r = 1.6\,\sigma_r$. With a transverse

size of 15 μm, this peak electric field of our bunch is given by the engineering formula:

$$\widehat{E}_r(z)\ [GV/m] = 1.81\, I(z)\ [kA] \tag{5.2}$$

Peak fields of 5 to 6 GV/m thus will correspond to about 3 kA of beam current for our beam radius. Away from $r = 1.6\,\sigma_r$, the fields are lower, so full or nearly full ionization of the ion column is not expected to happen until the beam current rises to 4 or 5 kA.

In one run of 200 images from 16 July 2004, the bunch length (and peak current) was varied by changing the overall phase of the accelerator. Comparison of the X-Ray based spectrometer images with simulation shows that the bunch lengths ranged from about about 130 μm to about 20 μm during this sequence of runs. After selecting only those events within a tight range of charge at the plasma and after removing the few events which failed to match well to any of the set of simulations, 115 shots remain.

It is useful to create a condensed view of all 115 Cherenkov spectrometer images. We take each image, an example of which is the right hand one in Figure 4.19, and then sum it sideways to create a single pixel wide lineout. We can then place these lineouts immediately next to one another and order them by the calculated peak current of the incoming electron bunch to get a picture of what is happening to the Cherenkov energy spectrum as the bunch length changes. Figure 5.6 shows such an image, with a plot of the increasing peak current of the bunches below. The overall trend of increasing energy spread at the Cherenkov screen with increasing peak current is very clear.

We see that for the first few images, there is a bit of noise, but there is no ionization. Once the peak current rises a little above 2 kA, the Cherenkov spectra begin to broaden, primarily with particles dropping in energy. The only mechanism available to decelerate these particles in our experiment is that they are driving a plasma wave, so they clearly must have started to ionize the Lithium. With increasing peak current, we see that the deceleration of the particles increases to some maximum, at which point it mostly levels off. There appear to be two interrelated effects causing this behavior. The first is that we go from having no ionization to having full

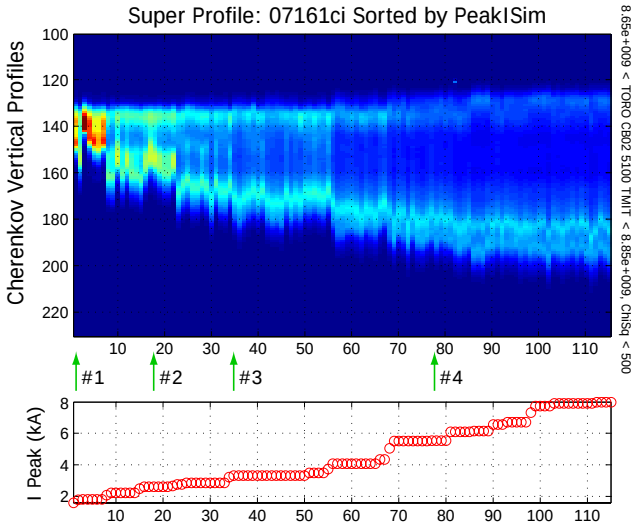

Figure 5.6: Cherenkov spectra summed to single pixel width and ordered by the peak current of the incoming beam as determined through simulations.

ionization, allowing for ever denser plasmas in which to drive wakes. Additionally, the more intense bunches ionize the beam ever earlier within the bunch, so that a larger fraction of the particles can participate in driving a wake. We quantify these statements below.

We know the current profile of the electron beam and its transverse size at the plasma. Thus, we also know the electric field at all points, and can use Eq. (2.47) to find the instantaneous ionization rate. The peak electric fields of our relativistic electron bunch lie at the $r = 1.6\,\sigma_r$ point, so we integrate the total expected ionization in the annular sheath at that optimal radius and plot it for four of the calculated profiles from Figure 5.6, giving Figure 5.7.

The first current profile corresponds to the very first shot. The second profile is that which matched shots 16 through 21, and had a peak current 2.6 kA. The third profile is for shot 34, with a peak current of 3.2 kA, and the last profile is for shots 77 to 80, with peak current of 5.5 kA. For all of the bunches, the ionization at other radii than $1.6\,\sigma_r$ will be lower, which is why, for example, the third profile shows that we expect full ionization, but the corresponding Cherenkov lineout does not yet have the maximal amount of deceleration there. Only part of the plasma sheath

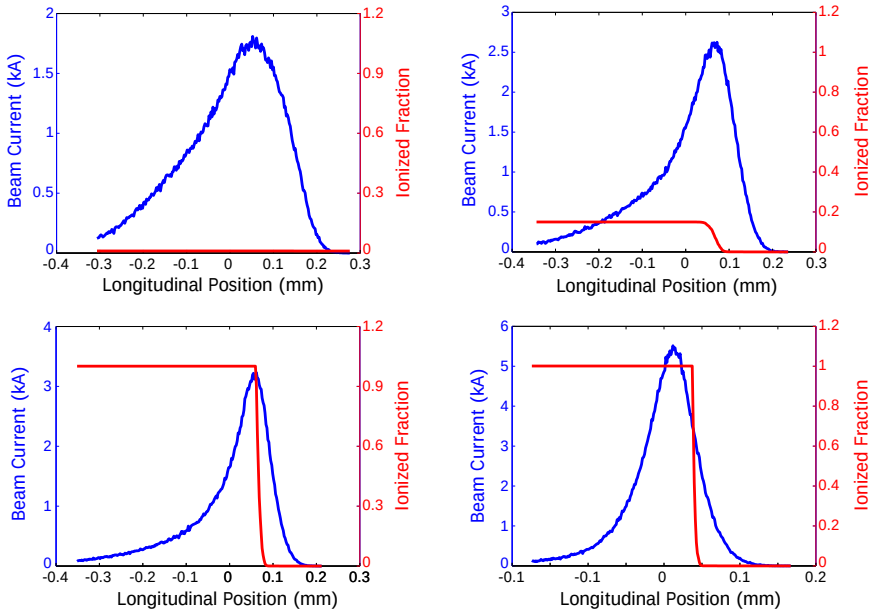

Figure 5.7: Current profiles for the four bunches indicated in Figure 5.6. Ionization is rapid compared to the bunch length above about 3.5 kA.

experiences full ionization and therefore can produce the largest fields available. Once the ionization is complete at some early point in the bunch, the wakes we drive should not change significantly, which is what we see for the final third of the shots.

In looking at these various plots in conjunction with Figure 5.6, it is clear that ionization at $r = 1.6\sigma_r$ turns on rapidly compared to the beam transit time when the current rises above 3 or 3.5 kA, corresponding (for a beam radius of 15 μm) to peak fields of about 5 GV/m, just as predicted by theory. Additionally, the *entire* ion column seems to be ionized for peak currents above 5 or 6 kA. Our understanding of the bunch profiles and of ionization reinforce each other very well, and these numbers have immediate consequences in the following section.

5.1.5 Post-Plasma Energy Spectrum Features

During the summer 2004 data run, we observed an unexpected feature on the Cherenkov based electron beam spectrometer. Specifically, there was a large amount of charge at or near the highest incoming energies which appeared to be unaffected by passage

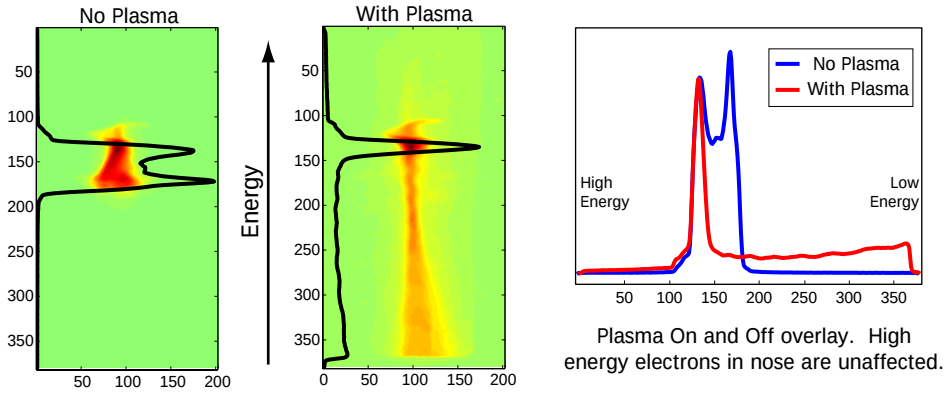

Figure 5.8: Two Cherenkov spectrometer images for beams with nearly identical incoming energy spectra. On left is a shot without, right is with plasma. High energy particles for both shots make it through with no effect as shown on right hand overlay. The axis corresponding to energy has units of pixels on our camera.

through the plasma. With a normal Gaussian bunch that ionizes the Lithium early on, so that almost all but the tailmost electrons are decelerated, this is puzzling. The feature was better explained after the experimental runs, when we simulated the accelerator in detail to match the observed energy spectra.

For data runs from 13 July 2005, we have two shots with nearly identical incoming energy spectra. The first Cherenkov spectrum was with no plasma in the beam path, and the second is of the beam after going through the plasma, see Figure 5.8. With these two shots, we can compare the Cherenkov energy spectra to get a good idea of what changes downstream as we insert the plasma cell into the beam. For reference, Figure 5.9 shows the phase space we have determined for this pair of shots.

If we look closely at the current profile of the bunch in Figure 5.9, we note that there is a long "nose" at the front where the current of the bunch remains low for about 100 μm. As we saw previously, until the current rises above 3 to 5 kA, the Lithium is not ionized. With no plasma for the initial portion of the bunch, nothing should happen to the energies in this population of high energy incoming electrons.

The "Nose" of our bunch, before the current ramps up to over 5 kA, contains roughly a third of the total charge in this instance. We have also observed that in many cases, about one third of the total charge goes through the plasma unaffected.

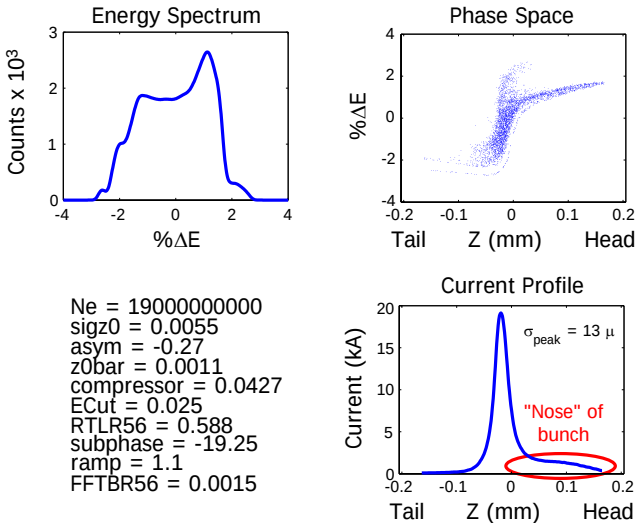

Figure 5.9: Plot of the simulated phase space for the two accelerator shots discussed in this section. One shot had plasma off, and one had the plasma on.

Our understanding of the bunch current profile and of Lithium ionization allow us to explain the initially puzzling collection of charge at high energy that we observed on the Cherenkov spectrometer even with the plasma cell in.

5.1.6 Uses for Knowledge of the Phase Space

Having verified through these various techniques that simulations do, indeed, tell us the phase space, we can move on to understanding various aspects of the beam's interaction with the plasma. We use knowledge of the beam's phase space to inform investigations of hosing instabilities, and of greatest importance, the plasma acceleration effect itself.

5.2 Understanding Acceleration

The main motivation for seeking to know the longitudinal phase space of the beams is naturally so we can understand acceleration more fully. We can use the knowledge of the phase space in several ways.

5.2.1 Direct Acceleration Determination

An important application of the phase space reconstructions is to use them to determine the magnitude of acceleration that we actually have achieved in our accelerator. The technique discussed here is similar to that presented in [79], and the basic idea is that identifying the precise incoming energy of particles in the tail allows a more accurate determination of their overall energy gain.

In E164, we have a short plasma and as discussed previously, our incoming bunch has a substantial energy spread in order to make it short enough for our experiment. As we see above in Figure 5.9, the intrinsic energy spread of the beam coming into the plasma has a full width of 4%, or nearly 1.2 GeV.

With peak gradients expected to be less than 50 GeV/m over a plasma of only 10 cm, the best possible energy gain does not dwarf the beam's own energy spread. As a result, there is potentially substantial uncertainty in determining the actual energy gain by particles when we look at the images of the Cherenkov spectrometer downstream of the plasma.

If we can identify the particles which are being accelerated and know their incoming energy, then we can improve the accuracy of our gradient measurement.

Identifying Particles

We have seen in a variety of simulations of our plasma wakefield accelerator that the wake forms a surprisingly long distance behind the bunch. In linear wakefield theory, we would expect, as shown in Chapter 2, cf. (2.41), that the wake will be at a maximum when $k_p\xi = \frac{\pi}{2} \rightarrow \xi = \frac{\lambda_p}{4}$ behind the bunch. This is because the electrons are assumed to have been blown out gently and immediately recollapse after the peak current portion of the beam passes. From fully blown out to recollapsed back to the axis is one quarter of a plasma wavelength.

Related to this, in linear theory, we see that there is an optimal bunch length for any given plasma density, where we match using the condition that $k_p\sigma_z = \sqrt{2}$ for Gaussian bunches. For too long a bunch, the wake is not strongly driven, and for too short a bunch, there are no particles available in the tail to be accelerated.

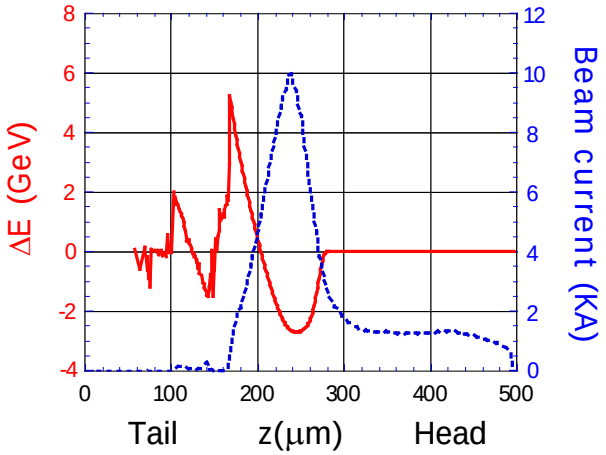

Figure 5.10: Plot of wake developing behind a bunch simulated in QUICKPIC. The plasma density is 2.8×10^{17} cm^{-3} with an associated plasma wavelength of $\lambda_p = 63\,\mu$m. The maximum accelerating field can be seen to lie almost a full λ_p behind the peak current point of the electron bunch. Adapted from [80].

For parameters as in E164, plasma electrons are expelled ballistically by the short electron beam passing through. After the point of peak current, the plasma electrons are still moving *away* from the beam axis before they are ultimately pulled back in by the ion column. Being relativistic, the electrons cannot accelerate arbitrarily as they would in simple harmonic motion. Thus, the plasma electrons oscillate with a phase delay and a substantially longer period than predicted by the naive linear theory, where we assume motion as in a classical harmonic oscillator.

With these two effects, the wake forms not at a distance behind the main current pulse of $\xi = \frac{1}{4}\lambda_p$, but at $\xi \lesssim \lambda_p$. The simulation output shown in Figure 5.10 shows this long distance. The simulation was performed for a 10 kA peak current beam in a plasma of density 2.8×10^{17} cm^{-3} with associated plasma wavelength of $63\,\mu$m. The wake is strongest about $65\,\mu$m behind the peak current point of the bunch, just in front of the point where the blown out plasma electrons come back to the axis.

With this, we know that the particles which will be accelerated are relatively far into the tails of our bunch. With even our shortest available plasma wavelength of 56 microns, the distance behind our bunch for largest acceleration is more than one or two σ_z behind the peak.

The wake evolves from decelerating to accelerating to decelerating again as we move progressively backward behind the bunch. This means that electrons in the tail will see more or less acceleration as a function of their position.

In an ideal plasma accelerator, we would have a very short bunch of electrons trailing the main drive bunch in a position such that all of the witness bunch is accelerated strongly and equally. With the single bunch setup of E164, we have particles at all of the phases of the evolving plasma wake, so some are accelerated more strongly than others. This manifests itself as the jet of accelerated particles that we see on our Cherenkov detector. The particles with highest energy are whichever lie at exactly the best longitudinal position behind the bunch.

Finding the Acceleration Strength

An example data event from 13 July 2004 where we seek to know the gradient showed many particles being accelerated. On the right side of Figure 5.11 we see the energy spectrum as measured on our Cherenkov diagnostic after the plasma with a strong jet of particles above the incoming maximum energy. We can see the substantial number of particles which have gained a continuum of energies at least up to the top of the screen, about 2 GeV above the highest incoming energies.

Improvements to the E164 Cherenkov based spectrometer have given a larger energy acceptance so that all particles will be viewable, but here we are constrained by the available hardware, and some of the particles may have been accelerated by more than we could see.

Using the measured energy spectrum from before the plasma, we have found the phase space for the above event, as shown in the left portion of Figure 5.11. The intrinsic energy spread is quite large.

At a density of $3.5 \times 10^{17}\text{cm}^{-3}$, with even our shortest plasma length of about 10 cm (see Chapter 4), there is enormous acceleration. Recall that when we claim the plasma is 10 cm long, that refers to its FWHM, and the actual flat top portion which is at the full nominal plasma density is only a few cm long.

The maximally accelerated particle must have seen an accelerating gradient of at least 2 GeV in 10 cm, or 20 GeV/m. However, the peak gradient achieved in the

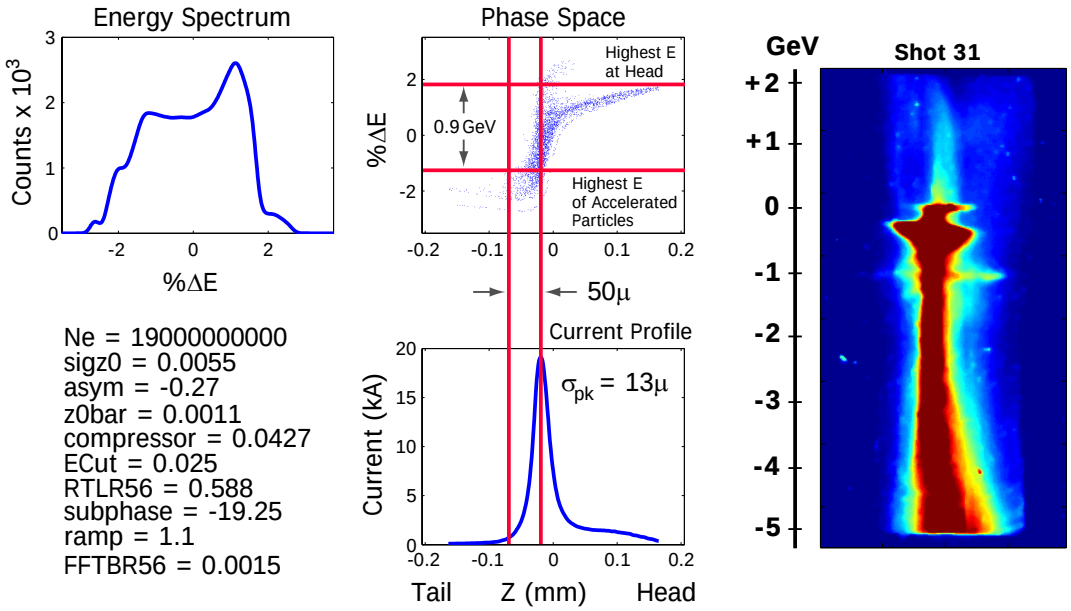

Figure 5.11: Incoming phase space for Shot 31 of Run 07131co and post-plasma spectrometer image showing substantial acceleration.

full density flat top must be higher than this. Further, we can show that the actual energy gain to these particles was even greater than the observed 2 GeV.

As discussed in the previous section, best acceleration happens nearly a full plasma wavelength behind the point of maximum current in our bunch, and at a density of 3.5×10^{17} cm^{-3}, the point of maximum acceleration thus lies about 50 μm behind the position of the bunch's peak current. In Figure 5.11, we draw a line on the simulation phase space corresponding to that position.

As we can see, the particles in a position to be accelerated all lie below about -1.5% in energy relative to the mean particle in the bunch. The head of the bunch, which we have seen will be unaffected by the plasma, has its highest energy particles at nearly +2%. Just for the accelerated particles to become visible, they must have gained more than 3%, or over 0.9 GeV of energy. Thus, when we see 2 GeV of gain, the actual acceleration was greater than we can directly observe, and our peak gradient at 3.5×10^{17} cm^{-3} is now calculated to be nearly 30 GeV/m. Further discussion of this technique for direct gradient determination can be found in [79].

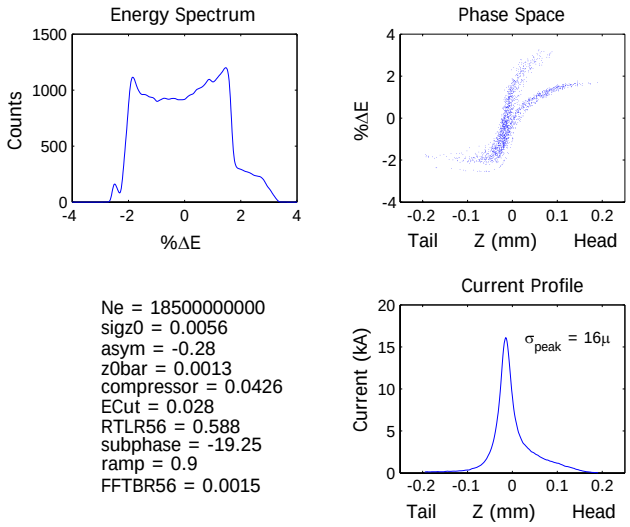

Figure 5.12: The phase space to which we match the three shots shown in this section.

5.2.2 Acceleration Properties Under Varying Conditions

Naturally, we seek to understand the acceleration for more than just one plasma density. Toward this end, we performed a series of data runs in July of 2004 with substantially different plasma densities and therefore, different plasma wavelengths. Although linear theory does not apply, there is still an optimal bunch length for each plasma density. Long bunches drive gentle wakes, but have many particles in the tail to be accelerated. Short bunches create very strong wakes, but have few particles far enough behind the main portion of the beam to see any acceleration.

During one week, we used plasma densities from 1.5×10^{17} to 3.5×10^{17} per cubic centimeter, so that the associated plasma wavelength shrank from 85 to 56 microns. Thus, we would expect that if we sent the same bunch into these different plasma densities, we would see differing amounts of acceleration.

We found the phase space for each event in the various data runs, and so could say what the approximate bunch length was for every shot. To do a direct comparison of one plasma density versus the others, we choose only a subset of the data events from various days which all matched to the same phase space from simulation.

A bunch profile that matched to data from each plasma density has a σ_z for the main peak of about 16 μm. Figure 5.12 shows the output from LiTrack. Linear theory

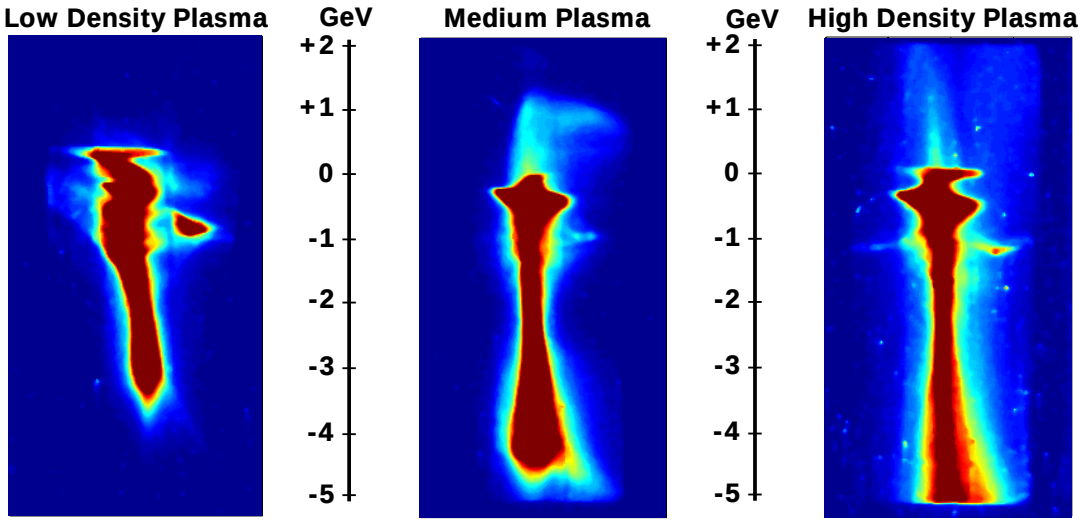

Figure 5.13: We compare the effect of different plasma densities on similar incoming beams. At left is 1.5×10^{17} cm^{-3} with no clearly accelerated particles. In the center, the plasma density is 2.5×10^{17} cm^{-3} and we clearly see an accelerated tail. On the right is the densest plasma: 3.5×10^{17} cm^{-3}, with many strongly accelerated particles. The approximate energy scale shown is relative to the highest incoming energies.

would predict the plasma density with maximum energy gain to be 2.2×10^{17} cm^{-3}, but we have non-Gaussian bunches with uneven wings. Also, the previously described phase delays in the formation of the wake make it less easy to predict the ideal plasma wavelength. We expect that we need a higher density for maximum acceleration than predicted by linear theory.

Having chosen events with the same incoming conditions, we then quantify the acceleration as a function of plasma density.

Figure 5.13 shows representative shots from the three densities of plasma. (Note that the first image had a slightly higher centroid incoming energy. The accelerator varies somewhat from day to day.) The first image represents the effect of the 1.5×10^{17} cm^{-3} plasma, where there are no accelerated particles visible. The second image shows the acceleration after the higher density plasma of 2.5×10^{17} cm^{-3}, where we now clearly see a wisp of electrons which have been accelerated beyond the highest incoming energies. The third image, with plasma density of 3.5×10^{17} cm^{-3}, shows particles with energies all the way to the top of the image. Beyond that point, we

cannot see the electrons accelerated by more than about 2 GeV, though it is clear that there must be some such electrons.

Using the same technique as in the previous section, we can calculate the gradient available for acceleration in each of these cases.

The distance behind the peak current of the bunch that the wake lies changes from 75 μm to 50 μm between these cases, but for each of them, the particles that experience that maximum wake start out lower than the head particles by almost 4% = 1.1 GeV. For the case with 1.5×10^{17} per cc, we do not see any accelerated electrons, so the gradient is ≤ 11 GeV/m.

With the higher plasma density of 2.5×10^{17} cm^{-3}, we see acceleration of about 1.2 GeV above the highest incoming energy for a maximum gradient experienced by a particle in this bunch of the order of 23 GeV/m.

At the highest density of 3.5×10^{17} cm^{-3}, we now see acceleration of a full 2 GeV, so as with the case in the previous section, we see a gradient ≥ 31 GeV, limited by our spectrometer range. It appears that the maximum acceleration for our 16 μm bunch would actually be achieved in a yet higher density plasma, but 3.5×10^{17} cm^{-3} is the maximum density in our oven.

This analysis is direct and gives information about the absolute maximum acceleration, but it suffers from the significant problem that we do not know the exact plasma length very well, and thus cannot quote a gradient to very high accuracy. With multi cm "wings" on either side, it is difficult to separate out the effect just from the highest density region in the center. The next section discusses a more involved technique which allows better investigation of the gradient as well as information about transverse deflections such as hosing.

5.3 Acceleration Analysis - Statistical

A more accurate way to understand acceleration is to change the plasma length and see how the acceleration is affected. Plotting these versus each other gives another way to determine the gradient. Because the wings where the Lithium density drops to zero have the same length whatever the central plasma length, any questions about

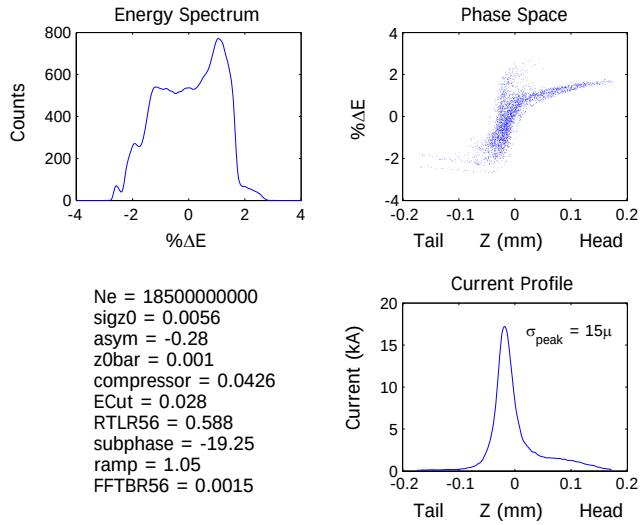

Figure 5.14: The phase space for the beams analyzed in this section.

the effect of the wings can be removed as common mode noise, and we only see the differential effect of adding or subtracting plasma length at the full desired density.

As discussed in Chapter 4, changing the oven heater power changes the plasma's length, but not its density. When reducing the power, changes in length require about 15 minutes to stabilize. As discussed in Chapter 4, for different plasma densities, the changes in length as a function of power are similar.

At each density, 1.5, 2.5 and $3.5\times10^{17}\,\mathrm{cm}^{-3}$, we changed the oven length by several centimeters in the course of about two hours. The most rapid changes come when cooling down, so we started with the highest power. In reducing the oven length by about 4 cm, we typically took 8 to 12 evenly spaced data runs of 100 shots each.

To make direct comparisons of the acceleration, we need to have similar incoming electron beams from each data run and for each of the plasma densities explored. We found that electron beams with the phase space as shown in Figure 5.14 were present in all of our data runs to be discussed below, and therefore provided a good basis for statistical analysis. Nonetheless, due to the natural variations of the main accelerator, some runs have only 1 or 2 matching shots, some have more than 40.

With the various data events all having the same incoming conditions, we seek some metric for the amount of acceleration experienced by the beam in each event.

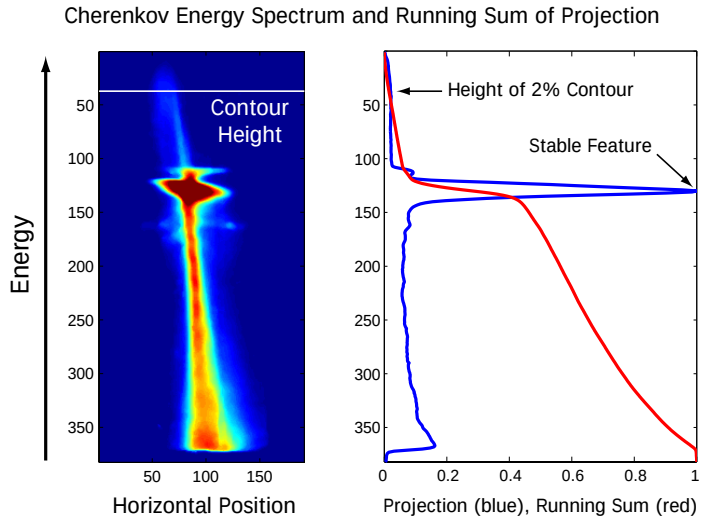

Figure 5.15: On left is a Cherenkov spectrometer image from 13 July 2004. On right is the projected energy spectrum in blue scaled so that the maximum is unity for plotting. In red is the running sum of the projection which is used to calculate the height of features such as the 2% contour labelled at top.

One way to quantify the acceleration that we see is simply to add up the total charge which has higher energy after the plasma than any particle had coming in. A more useful measure is to calculate contours of the beam in energy.

When dispersed at the Cherenkov screen, the vertical position of the electrons is dominated by their energy. If we calculate a running sum of the beam starting at the highest energy, we can calculate contours above which 5% of the beam's charge lies, above which 10% or 50% lies, and so on. Because a relatively small fraction of the beam is accelerated under even the best conditions, we focus on the contours in the range of about a half to several percent. Figure 5.15 shows a sample image from the Cherenkov screen and the calculation of the running sum which gives these contours.

As the incoming beam energy fluctuates slightly from shot to shot and between days, the relevant metric for the height of our 2% contour is not its absolute height, but its height above some stable feature in the beam itself. Thanks to the "nose" on our various beams before ionization takes place, we have such an unchanging feature. We calculate the height of the contour relative to this high energy spike, which is unaffected by the plasma, such as we have previously seen in Figure 5.8.

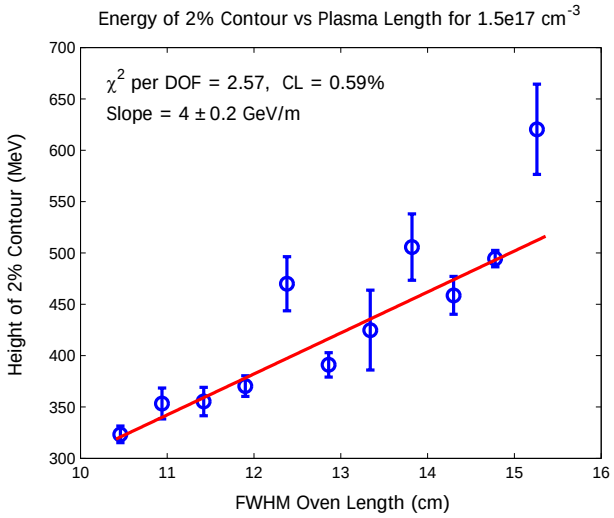

Figure 5.16: Plot of change in acceleration for the 2% contour vs oven length for plasma density 1.5×10^{17} cm^{-3} with linear fit.

5.3.1 Lowest Plasma Density

We use the above technique to analyze data taken on 11 July 2004 with the plasma set to a density of 1.5×10^{17} cm^{-3} to see what gradient had been achieved. By varying the oven power between the 390 and 460 Watts, we were able to change the oven length by about 5 cm in length. We took one set of 100 shots at each setting of the power, where the settings were 7.5 Watts different from one run to the next.

To have a consistent basis for comparing various events, we use only the shots from each run that match to the phase space shown in Figure 5.14. The number of successful matches varies, with between 2 and 17 matches for each of the runs here.

Having selected the subset of shots from each run, we calculate the height of various contours. As relatively few particles are accelerated, we calculate the gradient for only the 0.5% through 4% contours. For each shot, we calculate the height of the various contours above the stable feature afforded by the beam's "nose." We then plot the height of each contour versus the length of the plasma to see the gradient. An example of such a plot is shown in Figure 5.16 for the 2% contour.

The linear fit does not match the data all that well; the confidence level is quite low indicating that the linear fit is not capturing all of the physics or that the error

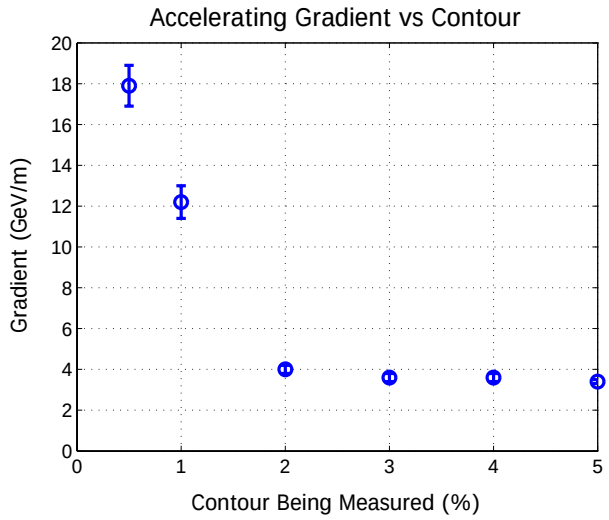

Figure 5.17: Plot of our accelerating gradient vs. the contour at which we choose to measure it for plasma density of 1.5×10^{17} cm^{-3}. Maximum achieved gradient is extrapolated to be about 20 GeV/m.

bars are incorrectly small. A discussion of the statistical techniques used to provide confidence levels can be found in Appendix A.

Nonetheless, by eye one can see that the gradient is something of the order of 4 GeV/m as given by the linear fit. We can create such plots for the various contours and it is then useful to plot the measured gradient versus the percent contour being studied. That way, we can extrapolate toward 0 to find the maximum gradient produced in our plasma. Such a plot is shown in Figure 5.17.

% Contour	**Gradient [GeV/m]**	**C.L.**
0.5	17.9 ± 1	0%
1	12.2 ± 0.8	0%
2	4.0 ± 0.2	0.6%
3	3.6 ± 0.2	0%
4	3.6 ± 0.2	0%
5	3.4 ± 0.1	0%

Table 5.1: Various parameters of the gradient for the lowest density plasma of 1.5×10^{17} cm^{-3}

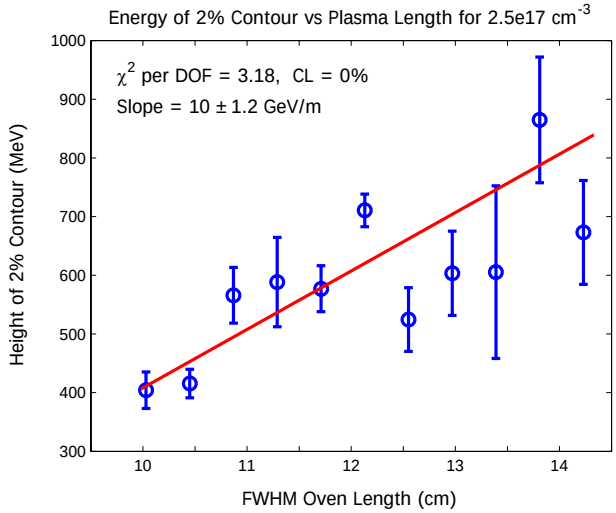

Figure 5.18: Plot of change in acceleration of the 2% contour vs oven length for the intermediate plasma density 2.5×10^{17} cm^{-3} with fit to a line.

We only include plasma lengths greater than 9 or 10 cm, where there are always visibly accelerated particles. Because the gradient is calculated for *changes* in plasma length, we remove questions about the size of the "wings" in the Lithium distribution, and no longer care exactly what the tail particle incoming energies had been.

We reiterate that the errors quoted on the gradients are unrealistically small, as the confidence levels for each linear fit are very low. Still, we can see that the gradient grows as we look at smaller and smaller groups of particles near the peak of the observed acceleration. Extrapolating to 0, we can estimate that the peak achieved gradient for our plasma of 1.5×10^{17} cm^{-3} was of the order of 20 GeV/m.

5.3.2 Intermediate Plasma Density

We perform a similar analysis with the 2.5×10^{17} cm^{-3} plasma and get interesting results. If we measure the acceleration of our 2% contour with beams that have the same longitudinal profile as mentioned above, we see the results in Figure 5.18.

The gradient determined by the fit is stronger than for the lower density plasma, as we would expect. Again, the line fit is definitely not capturing all of the physics. Unlike the previous case, however, the positions of the points are somewhat suggestive.

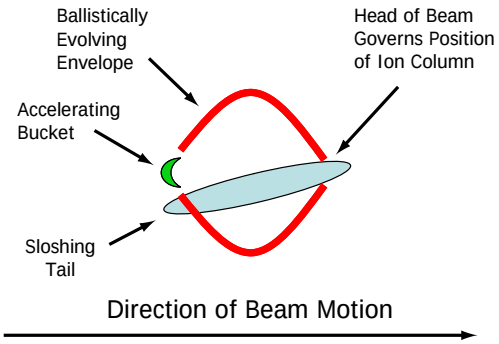

Figure 5.19: Diagram of the transverse size of the accelerating bucket. Particles that start off to one side will oscillate transversely at the betatron frequency, passing through the region of strong acceleration *twice* per oscillation.

We recall that we have seen transverse asymmetries in the beams coming into our plasma. If those offsets are in the tailmost particles, they will oscillate back and forth in the focusing fields of the ion column created by the beam blow-out. The strongest accelerating fields are located right on axis behind the head of the bunch. Although offset tail particles lie at the right longitudinal position, they will be in the right transverse position to experience the peak acceleration only part of the time, seeing an acceleration that oscillates in strength. The schematic of this effect is presented in Figure 5.19.

In each betatron oscillation, they will enter and then leave the accelerating region twice. Thus, not only will the average acceleration be lower than it would otherwise be, but we will see modulation of the effect at twice the betatron frequency, providing a useful hallmark to distinguish this particular effect from other possible explanations.

We can see on the upstream OTR screen that our beams do, in fact, come into the plasma with transverse asymmetries. Although we cannot distinguish the tail from the head in these images, offset tails are a reasonable explanation. We consistently notice that there are asymmetries to the same side. Figure 5.20 shows an example of the beam with a clear tail to one side.

Interestingly, the transverse profile of the bunch clearly does not match to a Gaussian. Rather, to quantify the asymmetry, we fit the bunch to an asymmetric Lorentzian, which is a very good match. Lorentzians are normally associated with

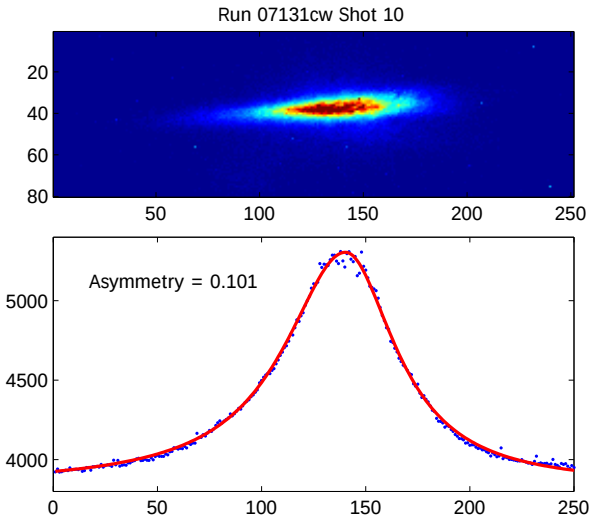

Figure 5.20: Sample image from the upstream OTR foil showing the small tail off to the left. The projection of the image onto the horizontal axis gives a curve which is well fit by an asymmetric Lorentzian.

resonance phenomena, so perhaps the periodic transverse focusing in the main accelerator optics participates in creating this shape. Our equation is:

$$L = C + \frac{W}{1 + \frac{2(1+sgn(x)A)^2}{x^2\sigma_L^2}} \tag{5.3}$$

where C is the pedestal on which our Lorentzian sits, W is the amplitude, A is the asymmetry factor and σ_L is the width parameter of the curve. This is by direct analogy with our definition of the asymmetric Gaussian (3.1).

With a transversely asymmetric beam coming in, it is natural to investigate what happens transversely to the tails of the beam *after* the plasma. As the Cherenkov spectrometer is imaging in x, we can see horizontal deflections of the beam, just as we saw them in the upstream OTR images.

The technique is to take a full Cherenkov image and find the height of whichever contour is of interest, as usual. We then take a horizontal lineout of the Cherenkov image at that height and fit the curve to a Gaussian. The transverse position that we quote is simply the mean value from the Gaussian fit. This is illustrated in Figure

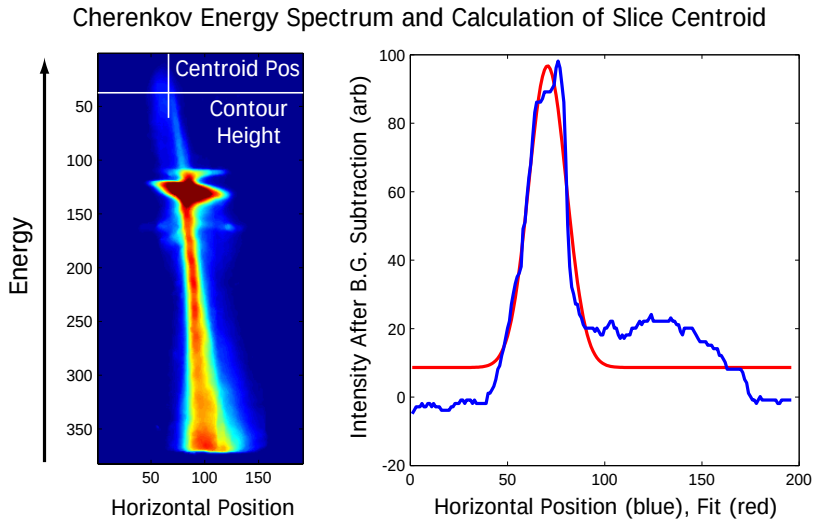

Figure 5.21: By taking a slice from the full image, we can then fit it to a Gaussian to get an estimate of where the centroid lies. In this case, the mean position in the Gaussian fit is not in the same place as the raw centroid would be, but this technique is more robust against noise such as X-ray hits.

5.21. In order to convert the observed transverse positions to something meaningful, we use the magnification of our electron optics to calculate the actual size of the transverse deflections at the plasma exit.

Using this technique, we can find the average horizontal position of the slices at the height of various contours for each plasma length. If tails are oscillating in the focusing fields of the plasma, we should see sinusoidal oscillations at the betatron frequency. We can then plot our tail transverse positions vs. plasma length, as in Figure 5.22.

The amplitude of these oscillations is of the order of the transverse size of the beam, σ_r. The wavelength for the oscillation returned by the fitting routine is a bit larger than the expected 2.23 cm, but we already know that our understanding of plasma length as a function of heater power is only accurate to about 10%, which could account for some of the discrepancy between these two values. Combined with the uncertainties in the fit parameters, our measured value is close enough to the expected value to believe that this really does represent betatron oscillations of an offset tail.

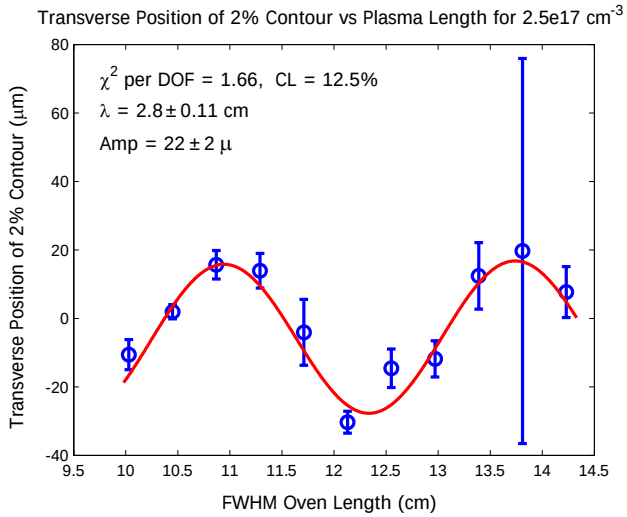

Figure 5.22: Horizontal position of the 2% contour versus plasma length and find that its oscillations could be consistent with simple betatron focusing of an initial offset in the accelerated tail.

And the observed 1.26 ± 0.05 cm period of oscillation in the energy gain is roughly half this value, as we would expect.

With all of this information, we fit the acceleration data not to a line, but we superimpose a sinusoid on top of the expected linear increase in energy with plasma length. If we fit the points again, but to a curve which is the sum of a line and a sinusoid, we get better agreement with the data. This is plotted in Figure 5.23.

There are now more parameters to the fit, and we show the most important ones on each graph, namely the wavelength of the oscillation and its amplitude. The wavelength we see for this particular fit is 1.26 ± 0.05 cm. This number is close to half of the betatron wavelength for our 28.5 GeV beam in a plasma of this density: 1.12 cm, and is close to half the value obtained by fitting to the transverse oscillations. This is exactly the relationship we expect if the transverse oscillations affect the acceleration.

Furthermore, if the oscillations in acceleration that we see are caused by this transverse motion of the tail particles, then we expect that the oscillation amplitude will be reduced in tandem with the decreasing overall acceleration as we investigate ever more inclusive contour lines. This is exactly what we see at right in Figure 5.24.

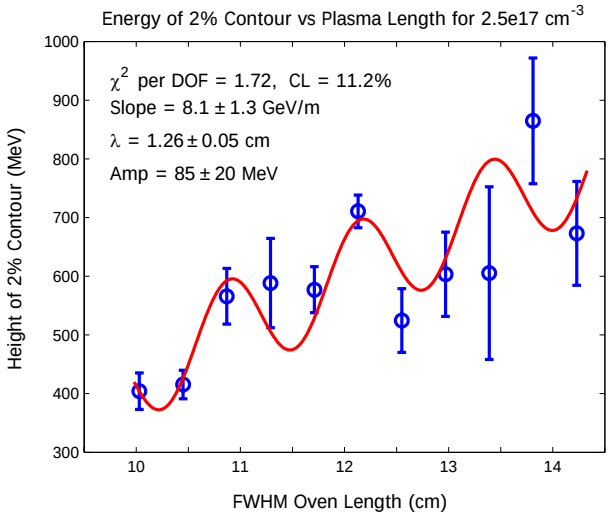

Figure 5.23: Plot of change in acceleration of the 2% contour vs oven length for plasma density 2.5×10^{17} cm^{-3} with fit to a line plus sinusoid.

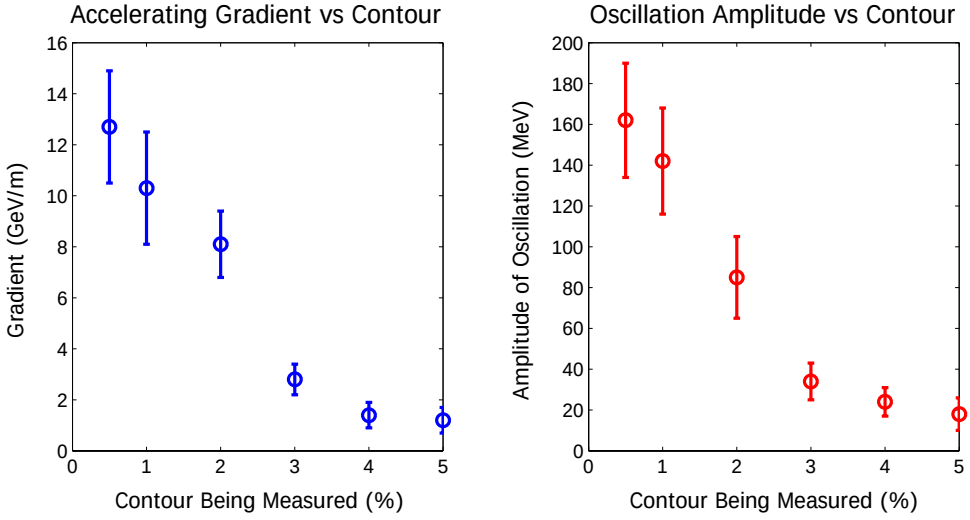

Figure 5.24: Plot of our accelerating gradient vs. the contour at which we choose to measure it for plasma density of 2.5×10^{17} cm^{-3}. Maximum achieved gradient is extrapolated to be about 14 ± 3 GeV/m.

% Contour	Gradient [GeV/m]	λ [cm]	Amplitude [MeV]	C.L.
0.5	12.7 ± 2.2	1.67 ± 0.08	162 ± 28	1%
1	10.3 ± 2.2	1.27 ± 0.05	142 ± 26	1%
2	8.1 ± 1.3	1.26 ± 0.05	85 ± 20	11%
3	2.8 ± 0.6	1.27 ± 0.06	34 ± 9	24%
4	1.4 ± 0.5	1.37 ± 0.11	24 ± 7	22%
5	1.2 ± 0.5	1.63 ± 0.14	18 ± 8	66%

Table 5.2: Various parameters of the gradient and observed oscillations in that gradient for the medium density plasma of 2.5×10^{17} cm^{-3}

If we look at the wavelength of the oscillation in acceleration amplitude from Table 5.2, we see that there is a group of similar values for the contours 1% through 4%, and the average value is about 1.3 cm. This is close to the expected period of oscillation for transverse effects, and gives further evidence that these oscillations cause the uneven acceleration as a function of oven length.

We see also that the oscillation amplitude changes in much the same way as the overall accelerating gradient, which again implies that this is caused by the above mechanism. We note that the overall gradient is actually lower than for the case of the lower density above. This is to be expected if the tail spends a significant amount of time outside of the main accelerating bucket. So these three lines of inference give confidence that we understand the effects on acceleration. As a side effect of this investigation, we can say that we do not see evidence of hosing growth, but transverse deflections can still be a problem for the overall achievable gradient.

5.3.3 Highest Density Plasma

Lastly, we investigate the highest plasma density achieved, 3.5×10^{17} cm^{-3}. As observed in § 5.2.2, we should again have the highest gradient.

If we again plot the transverse deflections for the 2% contour as a function of plasma density, we obtain Figure 5.25. The transverse oscillations are somewhat greater in amplitude than those observed with the intermediate plasma density. The machine was likely in a slightly different state the day this data was taken.

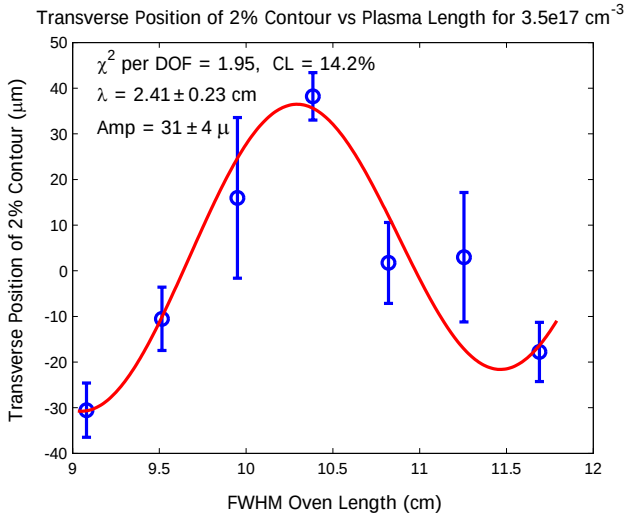

Figure 5.25: We plot the horizontal position of the 2% contour versus plasma length and find that its oscillations could be consistent with simple betatron focusing of an initial offset in the accelerated tail.

We expect that there will be oscillations in the energy, and that the amplitude of those oscillations will be greater than for the medium plasma density. We therefore use the line plus sinusoid fitting routine again. Figure 5.26 shows the result of finding the gradient for the 2% contour of shots into this high density plasma.

We see that the wavelength of oscillations in the transverse dimension is very close to twice the wavelength of the energy modulations that we observe, as expected with our understanding of what is happening inside the plasma. At 1.22 cm and 2.41 cm, both wavelengths are greater than expected theoretically, at 0.94 cm and 1.88 cm, respectively. This is probably because our understanding of the plasma length versus heater power is not perfect. If we refer to Table 4.1, we see that the value of the slope for 3.5×10^{17} cm^{-3} looks anomalously large. Incorrectly large values cause the wavelength to appear longer than it is (and reduce the measured gradient).

We plot the gradient and oscillation amplitude achieved vs. each contour investigated. This is presented for our high density plasma in Figure 5.27, where the estimated maximum gradient for the 0% contour would be at about 22 ± 5 GeV/m. If we assume that the differences we observe between the measured betatron wavelength and the theoretical value are because we have a poor calibration of the oven

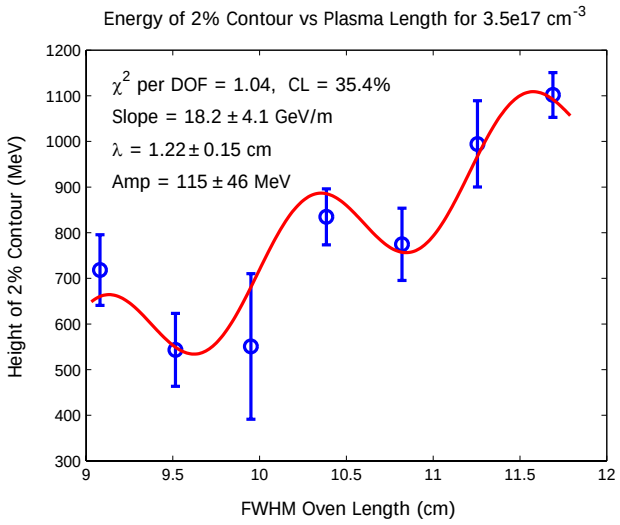

Figure 5.26: Plot of change in acceleration of the 2% contour vs oven length for plasma density 3.5×10^{17} cm^{-3} with fit to a line plus sinusoid.

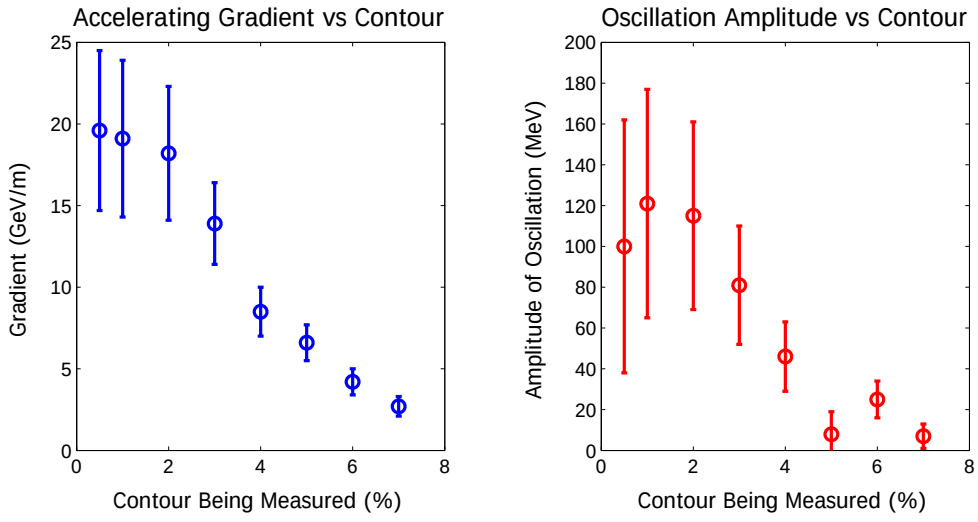

Figure 5.27: Plot of our accelerating gradient vs. the contour at which we choose to measure it for plasma density of 3.5×10^{17} cm^{-3}. Maximum achieved gradient is extrapolated to be about 22 ± 5 GeV/m.

% Contour	Gradient [GeV/m]	λ [cm]	Amplitude [MeV]	C.L.
0.5	19.6 ± 4.9	1.55 ± 0.22	100 ± 62	16%
1	19.1 ± 4.8	1.19 ± 0.14	121 ± 56	28%
2	18.2 ± 4.1	1.22 ± 0.15	115 ± 46	35%
3	13.9 ± 2.5	1.22 ± 0.13	81 ± 29	18%
4	8.5 ± 1.5	1.22 ± 0.13	46 ± 17	9%
5	6.6 ± 1.1	1.17 ± 0.47	8 ± 11	3%
6	4.2 ± 0.8	1.14 ± 0.07	25 ± 9	11%
7	2.7 ± 0.6	1.15 ± 0.24	7 ± 6	1.5%

Table 5.3: Various parameters of the gradient and observed oscillations in that gradient for the highest density plasma of 3.5×10^{17} cm^{-3}

length versus heater power, we can scale the gradient by the ratio of the observed betatron wavelength to that which we expect at a density of 3.5×10^{17} cm^{-3}. Doing so, we obtain a gradient of 28.5 ± 6.5 GeV/m, which is in good agreement with the more direct analysis presented in § 5.2.2.

Despite the increase in transverse oscillation amplitude compared to that for the intermediate plasma density, we see an increase in accelerating gradient compared to the case of 2.5×10^{17} cm^{-3} above.

5.3.4 Overall Results and Conclusions

We have demonstrated that energy spectrum measurements, coupled with detailed simulations of beams in the SLAC main accelerator, do, in fact, give an understanding of the longitudinal profile. The matching worked well, but should be even more reliable if the sources of blurring in the spectrum, primarily coming from the electron beam's significant β_x, were reduced. Nonetheless, this understanding then allows several investigations of the beam-plasma interaction.

Knowing the incoming energy of the particles which are accelerated gives a better understanding of their overall energy gain, and, therefore, the gradient. We can also use the energy spectra more simply to select a subset of nearly identical incoming bunches in order to compare the effect of different plasma densities and lengths on those bunches.

The most important benefit of this second technique is that we can measure the achieved gradient in a manner which does not depend on the regions in which the plasma density is changing. We have seen that the strongest acceleration is achieved when we use the highest density plasma available, although simple linear theory would predict that best acceleration for our bunches of about $\sigma_z \approx 18\,\mu$m should be with a plasma density of 2.5×10^{17} cm^{-3}. This indicates that continued progress in making denser plasmas could be more beneficial than initially thought.

In the course of these investigations, it became apparent that transverse oscillations of the tailmost particles are responsible for a decrease in accelerating gradient as well as oscillating energy gain to the particles. Further investigation is warranted to understand more fully these effects.

Due to the nature of the data available as of July 2004, where we could change the oven length only modestly, more accurate conclusions about gradients and hosing must await the case when the plasma length can be changed by larger amounts, and those experiments are underway as of the second half of 2005.

Chapter 6

Conclusions

With current technologies for particle acceleration nearing their theoretical maximum gradients, new accelerating techniques are necessary. Plasma acceleration with lasers has already demonstrated more than a thousandfold increase in gradient, but has long suffered from difficulties in propagating the laser over a significant distance.

Beam based plasma wakefield accelerators get around such difficulties, but can only achieve the same high gradients with very short electron bunches. Having created the necessary bunches of RMS length 40 fsec, we lose the ability to understand their longitudinal phase space directly, as no currently available technique can resolve such bunches in time. To solve this problem, we used an indirect method to understand the longitudinal phase space of our electron bunches. This indirect technique has proven crucial to understanding a variety of effects that happen with short electron beams driving a plasma wakefield.

As the technique is new, we have verified its accuracy by comparing with a variety of other ways of understanding the electron beams at SLAC.

We have seen several different ways to apply the phase space information to understanding acceleration, where we have applied the technique to measure incredibly high gradients of nearly 30 GeV/m over macroscopic distances. We found that the gradient does, in fact, increase with greater plasma density, though the gradient is weakened when transverse deflections become significant compared to the transverse size of the beam.

Understanding the phase space of short beams is not only useful for E164, but has already been applied in designs for the Short Pulse Picosecond Source (SPPS) experiment at SLAC, where short bunches are used to create intense X-Ray radiation in a wiggler, as a test for the Linac Coherent Light Source (LCLS). Researchers at DESY have also inquired about the technique for measuring the energy spread of beams in a non-destructive way, and it may well find broad applicability in future machines.

Continued progress in accelerating gradient requires ever more sophisticated methods for understanding the short electron bunches that we use, and we have demonstrated one such in E164.

Appendix A

Linear and Nonlinear Fitting to Data

We discuss the useful and generally applicable case of linear least squares fitting both for background and to motivate the discussion of the main topic of interest: nonlinear least squares fitting.

The linear case is nice because it admits of an exact analytic solution which can be used to check numerical methods we seek to employ for the nonlinear case.

A.1 Linear Least Squares

The term "linear least squares fitting" does not require that the function we are using for the fit be linear in the independent variable, e.g. x. Rather, the fit function must be linear in our *fit parameters*, which we will denote as α_i. To perform linear least squares fitting, we require that our function $F(x)$ have the form:

$$F(x) = \alpha_1 f_1(x) + \alpha_2 f_2(x) + \ldots + \alpha_m f_m(x) \tag{A.1}$$

In this case, the $f_i(x)$ can be anything we want, x^2, $\cos x$, $J_\nu(x)$, or whatever else, as long as each α_i multiplies a separate function of x.

A.1.1 Linear Fitting Example

Without proving any of the formulae, we outline the analytic method for finding the linear least squares fit to data with error bars by following the procedures and notation presented in Orear's monograph "Notes on Statistics for Physicists, Revised" [81].

The problem we often have is to fit experimental data with error bars to a theoretical curve. For the reader's reference, when that curve is just a line, the procedure of least squares fitting is sometimes referred to as "linear regression," a term coming from statistics in the social sciences and business world.

The best way to demonstrate the technique for finding the best fit to our data and for finding the errors in the fit parameters is through a concrete case. Orear's Example 6 poses the problem that we have four experimental data points which we wish to fit to a parabola. (This can also be referred to as "quadratic regression.") Namely:

$$\vec{x} = (-0.6, -0.2,\ 0.2,\ 0.6)$$
$$\vec{y} = (5,\ 3,\ 5,\ 8)$$
$$\vec{\sigma} = (2,\ 1,\ 1,\ 2)$$

where $\vec{x}$ is the vector of horizontal positions at which we took our data points, and is assumed to represent perfectly known values. $\vec{y}$ gives the experimentally measured values at the various x positions, and $\vec{\sigma}$ represents the quoted errors in the y values.

We seek to use a quadratic fit function $F(x)$, so we write it as

$$F(x) = \alpha_1 + \alpha_2 x + \alpha_3 x^2 \quad \longrightarrow \quad f_1 = 1 \quad f_2 = x \quad f_3 = x^2 \tag{A.2}$$

With the three functions known, there is a straightforward method for calculating the fit parameters and their errors. The first step is to calculate what is generally known as the Hessian Matrix, **H**. Luckily, this is easy to do for linear fitting:

$$H_{ij} = \sum_a \frac{f_i(x_a) \cdot f_j(x_a)}{\sigma_a^2} \tag{A.3}$$

where a is an index referring to the various points. In this example, a runs from 1 to 4 such that we evaluate the functions f_i at each x. As i and j are interchangeable, $\mathbf{H}$ is clearly symmetric. For this example, as with all cases where $F(x)$ is composed of three subfunctions, we need only calculate six terms:

$$H_{11} = \sum_a \frac{1}{\sigma_a^2} \qquad H_{22} = \sum_a \frac{x_a^2}{\sigma_a^2} \qquad H_{33} = \sum_a \frac{x_a^4}{\sigma_a^2}$$
$$H_{12} = \sum_a \frac{x_a}{\sigma_a^2} \qquad H_{13} = \sum_a \frac{x_a^2}{\sigma_a^2} \qquad H_{23} = \sum_a \frac{x_a^3}{\sigma_a^2}$$

It is straightforward to verify that performing these sums will give:

$$\mathbf{H} = \begin{pmatrix} 2.5 & 0 & 0.26 \\ 0 & 0.26 & 0 \\ 0.26 & 0 & 0.068 \end{pmatrix} \tag{A.4}$$

The Hessian matrix is not directly useful until we take its inverse, $\mathbf{V}$. Doing this by hand is the worst combination of unpleasant and prone to calculational errors for all but the smallest of matrices. Thus, we normally let MATLAB perform the inversion with the delightfully simple command: `V = inv(H)`. We obtain:

$$\mathbf{V} = \begin{pmatrix} 0.664 & 0 & -2.54 \\ 0 & 3.847 & 0 \\ -2.54 & 0 & 24.418 \end{pmatrix} \tag{A.5}$$

Before we can calculate our fit parameters and their errors, we need to construct one final vector, $\vec{u}$. We have made no reference yet to the actual y values that we have measured, so it is reasonable that they enter into the calculation at this point:

$$u_i = \sum_a \frac{y_a \cdot f_i(x_a)}{\sigma_a^2} \quad \longrightarrow \quad \vec{u} = (11.25,\ 0.85,\ 1.49) \tag{A.6}$$

Armed with these various vectors and matrices, we can directly find the best fit

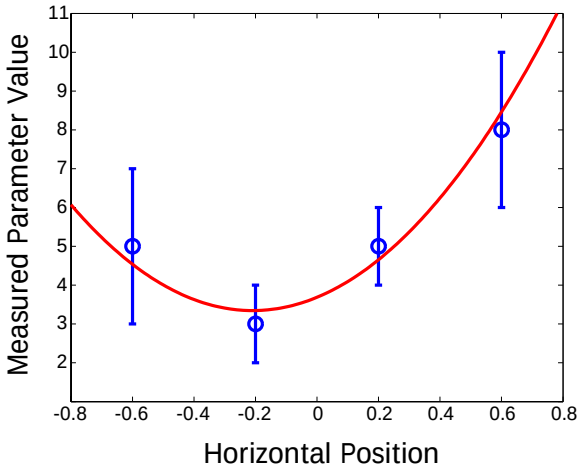

Figure A.1: Best fit plot. The blue points with error bars reflect the measured values and the red curve is the best fit function superimposed.

parameter values, often denoted by α_i^*:

$$\vec{\alpha^*} = \vec{u} \cdot \mathbf{V} \quad \longrightarrow \quad \alpha_1^* = 3.685 \quad \alpha_2^* = 3.270 \quad \alpha_3^* = 7.808 \tag{A.7}$$

To find the errors in the fit parameters, we just take the square root of the diagonal elements of our error matrix, $\mathbf{V}$:

$$\Delta\alpha_1 = 0.815 \quad \Delta\alpha_2 = 1.96 \quad \Delta\alpha_3 = 4.94 \tag{A.8}$$

If we are interested in the cross correlations between the fit parameters, they are simply the corresponding matrix elements of $\mathbf{V}$:

$$\overline{\Delta\alpha_1\Delta\alpha_2} = 0 \qquad \overline{\Delta\alpha_1\Delta\alpha_3} = -2.54 \qquad \overline{\Delta\alpha_2\Delta\alpha_3} = 0 \tag{A.9}$$

In summary, we have the best fit function which is shown in Figure A.1. We can quote only one or two significant digits, so write:

$$F(x) = (3.7 \pm 0.8) + (3.3 \pm 2.0)x + (7.8 \pm 4.9)x^2 \tag{A.10}$$

A.2 Interlude - Goodness of Fit

We have so far finessed the question of how good a job we have done in fitting our data to the assumed functional form. The traditional metric for how well the fit curve represents the data is called the χ^2, or "chi-square." This has the straightforward and intuitive definition that the differences between each data point and the fit function at that point are scaled by the known errors, then squared and summed:

$$\chi^2 = \sum_a \left[\frac{y_a - F(x_a)}{\sigma_a} \right]^2 \tag{A.11}$$

Generally more useful than the raw χ^2 is the normalized value, χ^2 per degree of freedom. As the reader recalls, the number of degrees of freedom, ν, is just the number of data points minus the number of free parameters in the fit. Thus χ^2/ν naturally builds in the concept that there must be more data points than degrees of freedom in the fit for us to be able to quote a meaningful result.

Depending on the number of degrees of freedom, we can determine how likely it is that we have chosen an appropriate function to fit our data. The intervals in Figure A.2 tell us this information. What the contours show is the percent chance that if we perform a subsequent identical experiment, we will get a χ^2/ν value greater than the one we have measured. The ideal case is where we are on the 50% contour, meaning that we have found the most likely function.

These "confidence level" curves can be calculated by using a generalization of the error function known as the Incomplete Gamma Function, $P(a, x)$ [82]. The following discussion is based on the treatment in Chapter 6, *Special Functions*. The Incomplete Gamma Function is calculated in much the same way as erf(x) except that the normalization is given by the Gamma Function. We recall that $\Gamma(a)$ is itself the generalization of the factorial function to non-integral a and is defined:

$$\Gamma(a) = \int_0^\infty e^{-t}\, t^{a-1}\, dt \tag{A.12}$$

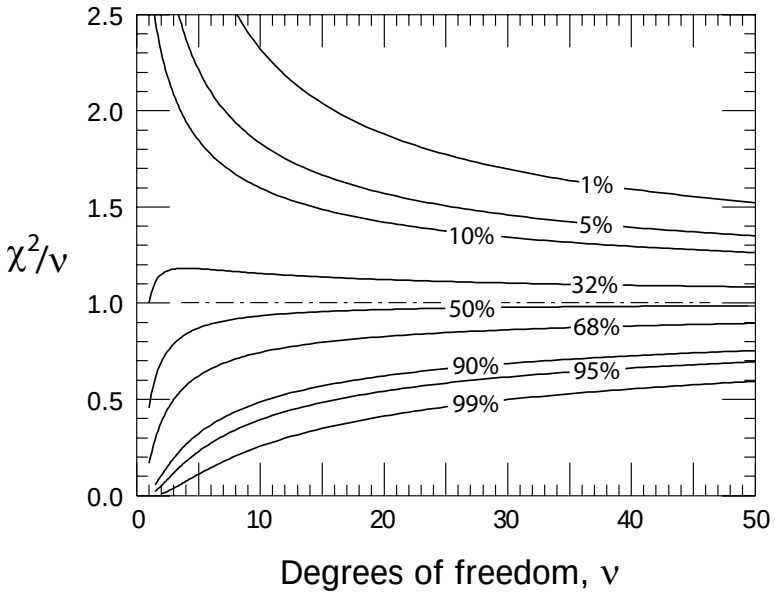

Figure A.2: Confidence interval curves for various χ^2/ν as a function of ν. In the example from § A.1.1, ν is just one, and the most trusted value for χ^2/ν would be about 0.5. Reproduced from [83].

With this normalization, we can define $P(a, x)$ to be:

$$P(a,x) = \frac{1}{\Gamma(a)} \int_0^x e^{-t}\, t^{a-1} dt \tag{A.13}$$

We use this to find our Confidence Level, CL, by setting a to be $\nu/2$ and x to $\chi^2/2$ and by taking the complement of $P(a, x)$. So for a given raw χ^2 and number of degrees of freedom ν, our confidence level is given by:

$$\mathrm{CL} = 1 - P\left(\frac{\nu}{2}, \frac{\chi^2}{2}\right) \tag{A.14}$$

In MATLAB, we calculate $P(a, x)$ with `gammainc`, but must be careful because a and x are *reversed* in order relative to the standard mathematical notation. Thus, we might implement the above as `CL = 1 - gammainc(ChiSq/2,DOF/2)`.

In general, if we have a confidence level of only 0.1%, then our fit is not very

meaningful. Either we are trying to fit to the wrong function, or we have underestimated the errors of each individual measurement. Conversely, if we are on the 99.9% contour, then the fit is "too good," and we have overestimated our errors or otherwise made a mistake. It is worth noting that when data is fudged to make it "better," a hallmark can be that the χ^2/ν is on a contour closely approaching the 100% confidence interval.

In the example from § A.1.1, the χ^2/ν is 0.35, placing us near the 60% contour, and indicating that the data are very much consistent with our assumed parabolic function. We note that being near the 50% contour does *not* in and of itself prove that we can exclude all other functional forms for the fit function.

We can only exclude other possible fit functions one by one. If the χ^2/ν associated with fitting to a specific function leads to a CL of 0.1%, then we say that that particular function does not match the data well. In this example, if we fit to a simple line, we obtain that χ^2/ν is 1.42, placing us at perhaps the 20% confidence interval. That is not a terrible result, so we cannot dismiss out of hand a linear fit to the data without real theoretical reasons to eschew this possibility.

It is nonetheless correct to say that the parabolic fit is a better match to the data. This difficulty in confidently choosing which is the correct fit function in this example is a direct result of the low number of degrees of freedom (1 or 2!), and adding even one more data point should enable a more confident discrimination between these two possible fit functions.

As we normally have a theory and therefore a hypothesis as to what functional form we would expect, the most robust result we can claim is simply whether or not we are confident that we have observed the functional form predicted in our hypothesis.

A.3 Theory Behind the Hessian Matrix

The most important part of determining the best fit and the errors on the fit parameters is the Hessian matrix, **H**. We have seen *how* to construct this matrix, but we need to understand *why* we did what we did when we move to the case of nonlinear least squares fitting.

General background leading to the results I quote subsequently can be found in [82], specifically Chapter 14 on modelling of data. The following discussion is adapted primarily from §14.4, *Nonlinear Models*.

The general numerical approach to finding fit parameters is to give approximate values and then vary them repeatedly, calculating χ^2 for each combination of parameters. Minimizing χ^2 gives the best fit parameters as long as the search starts close enough to the global minimum. The Hessian matrix at any point can be thought of as the curvature matrix of the χ^2 merit function at our set of fit parameters, $\vec{\alpha}$. The curvature is simply one half times the second derivative matrix:

$$H_{ij} = \frac{1}{2}\frac{\partial^2\chi^2}{\partial\alpha_i\partial\alpha_j} \tag{A.15}$$

At the risk of being repetitive, we again present the definition of χ^2 to show how its derivatives are calculated, and this time make explicit the dependence of F on the data and the fit parameters:

$$\chi^2(\vec{\alpha}) = \sum_a \left[\frac{y_a - F(x_a;\vec{\alpha})}{\sigma_a}\right]^2 \tag{A.16}$$

We take the gradient of χ^2 with respect to our fit parameters:

$$\frac{\partial\chi^2}{\partial\alpha_i} = -2\sum_a \frac{y_a - F(x_a;\vec{\alpha})}{\sigma_a^2}\,\frac{\partial F(x_a;\vec{\alpha})}{\partial\alpha_i} \qquad i = 1, 2, \ldots, m \tag{A.17}$$

Taking the second partial, we obtain:

$$\frac{\partial^2\chi^2}{\partial\alpha_i\partial\alpha_j} = 2\sum_a \frac{1}{\sigma_a^2}\left[\frac{\partial F(x_a;\vec{\alpha})}{\partial\alpha_i}\frac{\partial F(x_a;\vec{\alpha})}{\partial\alpha_j} - \left[y_a - F(x_a;\vec{\alpha})\right]\frac{\partial^2 F(x_a;\vec{\alpha})}{\partial\alpha_i\partial\alpha_j}\right] \tag{A.18}$$

As we can see, the components of **H** depend on both first and second partial derivatives, but the second derivatives are almost always ignored. Loosely speaking, this is for several reasons. First, the second derivatives are multiplied by the difference between the data and the fit function. An even modestly good fit makes this number

relatively small. Second, the sign of these difference terms will randomly be positive or negative, causing them to tend to cancel one another. Third, including the second derivatives can make the fit even less reliable in the circumstance where the model is not an ideal fit or when there are random outlying points.

Recalling that our factor of one half in the definition of the curvature matrix cancels the 2 in front of our sum, we ignore the second derivative terms and give the approximated definition for the elements of **H**:

$$H_{ij} = \sum_a \frac{1}{\sigma_a^2} \left[\frac{\partial F(x_a; \vec{\alpha})}{\partial \alpha_i} \frac{\partial F(x_a; \vec{\alpha})}{\partial \alpha_j} \right] \tag{A.19}$$

It is immediately apparent that this is the general formula which gave us the method of constructing the Hessian matrix in our earlier example. We had required that F be a linear combination, $F = \sum \alpha_i f_i(x)$. Thus, the derivatives with respect to each α_i trivially return just the $f_i(x)$ such that the bracketed term in the above sum will always be $[f_i(x_a) \cdot f_j(x_a)]$ for linear least squares fitting.

Although nominally a second derivative matrix, our approximated Hessian is really constructed from combinations of first partial derivatives. Those partial derivatives are themselves the elements of a matrix known as the Jacobian, **J**. (Confusingly, the determinant of **J** is sometimes also called the Jacobian.) In terms related to the above, and recalling that a is the index of each data point, the definition of the elements of the Jacobian matrix is:

$$J_{ai} = \frac{\partial F(x_a; \vec{\alpha})}{\partial \alpha_i} \tag{A.20}$$

where i spans the number of fit parameters we are using and the derivative is evaluated at the data points x_a. Thus, in this formulation, the matrix must always be taller than it is wide, corresponding to a nonzero number of degrees of freedom.

Explicitly, we can now define the Hessian as:

$$H_{ij} = \sum_a \frac{1}{\sigma_a^2} \left[J_{ai} \cdot J_{aj} \right] \tag{A.21}$$

To be more concrete, the Jacobian for our earlier example uses the partials with respect to the various α_i:

$$\frac{\partial F(x;\vec{\alpha})}{\partial \alpha_1} = 1 \qquad \frac{\partial F(x;\vec{\alpha})}{\partial \alpha_2} = x \qquad \frac{\partial F(x;\vec{\alpha})}{\partial \alpha_3} = x^2 \tag{A.22}$$

Evaluating these partials our 4 data points, we obtain the simple matrix:

$$\mathbf{J} = \begin{pmatrix} 1 & -0.6 & 0.36 \\ 1 & -0.2 & 0.04 \\ 1 & 0.2 & 0.04 \\ 1 & 0.6 & 0.36 \end{pmatrix} \tag{A.23}$$

The first element of the Hessian matrix is the sum of the first column after we square each element and divide by σ_a^2. Similarly, H_{23} is given by the sum of the term by term products of column 2 and column 3, again divided by the corresponding σ_a^2.

In the linear case, this really just gives theoretical background for the method already discussed at the beginning of this appendix. Association with the concept of the Jacobian is useful because that is what MATLAB actually produces in doing numerical fits.

A.4 Nonlinear Least Squares Fitting

We can still sometimes calculate the Hessian even if $F(x_a;\vec{\alpha})$ is not linear with respect to the various α_i, and therefore the partial derivatives are difficult to find analytically. Better yet, we can let computers give a numerical approximation to the derivatives and never do any algebra at all.

As discussed in Chapter 5, our acceleration data from the plasma appears to have the functional form of a line plus a sinusoid. Fitting to this takes us away from linear theory, because several fit parameters appear inside the sine function, and we no longer have a linear combination. For readability, instead of using the notation α_i

for the fit parameters, we just use the obvious symbols in the fit function:

$$F(x) = (Ax + B) + C \sin\left(\frac{2\pi}{\lambda}x + \phi\right) \tag{A.24}$$

There is no straightforward way to solve this analytically, though one could perform a Taylor expansion on the functions with redefined fit parameters. This is, in general, a nightmare, so we let computers do what they do best: perform lots of calculations to give a numerical answer which will be close to the correct one.

MATLAB is widely used, and has the additional advantage that its syntax is generally straightforward and resembles pseudocode. So for this example, I will quote actual MATLAB code to provide illustration of the technique for nonlinear fitting including errors.

A good way to find the best fit numerically is to minimize the χ^2 directly. That is implicitly what is done with more user friendly routines in MATLAB, such as `lsqcurvefit`. Here, we need to use a less convenient, but more powerful minimization routine. In order to get all of the parameters needed for us to be able to quote errors in the fit parameters, we do some of the work explicitly ourselves.

To find the best fit parameters, we use the function `lsqnonlin`. Because we are doing a very specific task, it is worth noting that we cannot (as of MATLAB 7.0) use the closely related `lsqcurvefit`, because it passes data in a slightly different way.

The routine `lsqnonlin` has a variety of bells and whistles. Its basic purpose, however, is to minimize some function that the user provides with respect to various parameters whose seed values are also provided by the user.

We have to define the χ^2 function in a specific way so that MATLAB will calculate it correctly. Because `lsqnonlin` squares and sums our function for us, we calculate the individual terms at each data point which combine to make χ^2, calling them the scaled differences $\vec{D}$:

$$D_a = \frac{y_a - F(x_a; \vec{\alpha})}{\sigma_a} \tag{A.25}$$

If we compare the D_a with the $F(x_a; \vec{\alpha})$, we see that the only differences are that we have *already* divided by σ_a and that we will have an overall minus sign in the

derivatives with respect to the α_i, as F is subtracted from y by convention. Thus, the terms of the Jacobian of $\vec{D}$ will be scaled by the $1/\sigma_a$ and will be the negative of the Jacobian of $F(x_a; \vec{\alpha})$. In constructing the Hessian, we must multiply a pair of Jacobian terms together, so the two minus signs will always cancel.

Our earlier definition of the Hessian of F (A.19) had a separate factor of $1/\sigma_a^2$ for each term, but that is now included in the Jacobian of $\vec{D}$. Having redefined the function we wish to minimize thus makes the Hessian matrix elements trivial to compute:

$$H_{ij} = \sum_a J_{ai} \cdot J_{aj} \tag{A.26}$$

In the previously mentioned case where we seek to match to a line plus sinusoid, we first have to write the function which will take the difference between each data point and the fit curve and then scale these differences by the errors. We require the vector of parameter starting conditions `StartParam` and the structure array `Data` which has the vectors $\vec{x}$, $\vec{y}$ and $\vec{\sigma}$ as its three fields. The function is:

```
\% Fit to a line plus a sinusoid with five parameters
function Differences =  LinePlusSinusoid(StartParam,Data)

\% Extract the data to be fitted
Horizontal = Data.Horizontal;
Vertical   = Data.Vertical;
Error      = Data.Error;

\% The initial conditions for the fit
Slope      = StartParam(1);
Intercept  = StartParam(2);
Amplitude  = StartParam(3);
Wavelength = StartParam(4);
Phase      = StartParam(5);
```

```
\% Calculate the function at the data points
Curve = Slope*Horizontal + Intercept + ...
   Amplitude*sin(2*pi*Horizontal/Wavelength + Phase);

\% Take the differences scaled to the errors, lsqnonlin  squares and
\% sums automatically in creating the Chi Squared.
Differences = (Vertical - Curve)./Error;
```

We now write a short script to call `LinePlusSinusoid` inside `lsqnonlin`. We assume that we have defined the vector `StartParam` and the structure array `Data`. We could have created a simple array with x, y, and σ, but this is more mnemonic.

In the input list of parameters, we just give empty sets for several of the parameters such as upper bounds that we do not wish to use as constraints.

To get parameters such as χ^2, we must give a sequence of variable names inside the brackets on the left of the equals sign into which `lsqnonlin` can write the results. These parameters must be given in order, so we still include variables for items we do not care about. For ease of comparison with MATLAB help, in the following code, we use the same names for output parameters as appear there and indicate which ones we care about by capitalizing them. `Resnorm` gives the raw χ^2 and `Jacobian` is the numerically calculated $\mathbf{J}$. It is now easy to calculate $\mathbf{H}$, invert it to make $\mathbf{V}$, and take the square root of the diagonal elements to find the errors on the previously returned fit parameters:

```
\%******* Fit to Line and Sinusoid with lsqnonlin *******
StartParam = [Slope0 Intercept0 Amplitude0 Wavelength0  Phase0];
[FitParams,Resnorm,residual,exitflag,output,lambda,Jacobian] = ...
    lsqnonlin(@LinePlusSinusoid, StartParam,[],[],[],DataToFit);

DOF = length(DataToFit.Horizontal) - length(StartParam)
ChiSqPerDOF = Resnorm/DOF
CL = 1 - gammainc(Resnorm/2 , DOF/2)
```

```
Slope          = FitParams(1);
Intercept      = FitParams(2);
Amplitude      = FitParams(3);
Wavelength     = FitParams(4);
Phase          = FitParams(5);

\%--------- Error Analysis on Fit Parameters -------------
for i = 1 : length(FitParams)
    for j = 1 : length(FitParams)
        H(i,j) = sum(Jacobian(:,i).*Jacobian(:,j));
    end
end
V = inv(H);
for k = 1 : length(FitParams)
    Deltas(k) = sqrt(V(k,k));
end
\%*******************************************************
```

With `FitParams` and `Deltas` in hand, we have the fit with its errors and are done with the statistics. It is now up to us to interpret the results.

To give an example of applying this code, we apply it to data discussed in Chapter 5 for our highest density plasma with 3.5×10^{17}. This code and a similar variant that fits simply to a line allows us to discriminate whether the line plus sinusoid or the simple line is superior. Our data vectors corresponding to `DataToFit.Horizontal`, `DataToFit.Vertical` and `DataToFit.Error` are:

$$\vec{x} = (0.37, 0.80, 1.22, 1.65, 2.08, 2.51, 2.93, 3.36)$$
$$\vec{y} = (780, 700, 1320, 1160, 1380, 1720, 1660, 1700)$$
$$\vec{\sigma} = (91, 64, 153, 124, 228, 67, 316, 60)$$

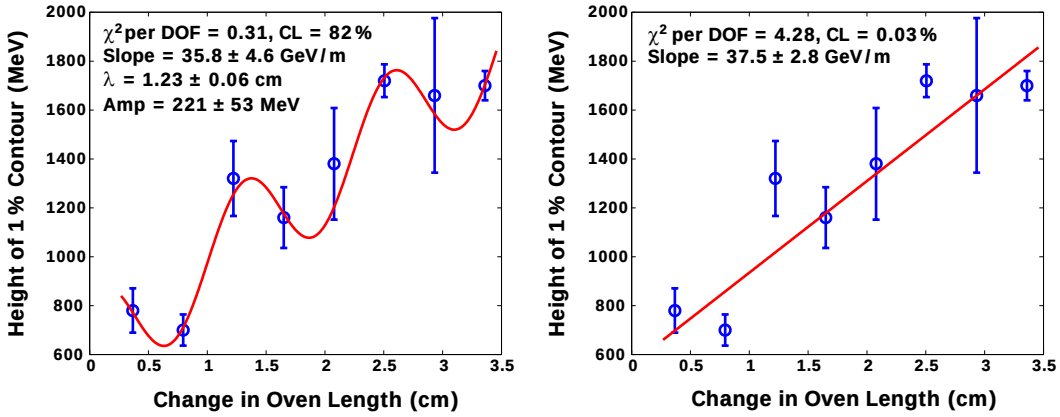

Figure A.3: Energy gain as we lengthen the plasma cell. Data fit much better to a line plus a sinusoid than a simple line.

To seed the fits, we just use `polyfit` to give us the linear portion's two parts. For the line plus sinusoid, we then use trial and error to find reasonable starting values for the amplitude, wavelength, and phase. When we get a reasonably close starting value, the fits always return to those shown in Figure A.3 for even substantial changes in any one of the fit parameters. For reference, we give the seed values for the line plus sinusoid fits. When fitting only to a line, our function is simpler than `LinePlusSinusoid` and we only seed it with the first two parameters:

$$(A_0, B_0, C_0, \lambda_0, \phi_0) = (353, 645, 300, 1.3, 5.0) \tag{A.27}$$

We plot the two ways of fitting this data side by side to compare. Clearly, the addition of the sinusoid makes for a vastly superior fit. In that case, there are 3 degrees of freedom, and the χ^2/ν of 0.31 gives us a degree of confidence of 82%. Perhaps we have slightly overestimated the errors, but this is a very believable value.

For the linear fit, with χ^2/ν of 4.28, the degree of confidence is 0.026%, and a line is extremely unlikely to represent the data fully.

Bibliography

[1] W. K. H. Panofsky and M. Breidenbach. Accelerators and detectors. *Review of Modern Physics*, 71:S121–S132, 1999.

[2] R. M. Barnett et al. Review of particle physics. particle data group. *Physical Review*, D54:128–131, 1996.

[3] J. W. Wang et al. Progress toward NLC/GLC prototype accelerator structures. In *International Linear Accelerator Conference (LINAC 2004)*, 2004.

[4] I. Wilson. The compact linear collider (CLIC). *Phys. Rept.*, 403-404:365–378, 2004.

[5] D. P. Pritzkau. *RF Pulsed Heating*. PhD thesis, Stanford University, 2001.

[6] D. P. Pritzkau and R. H. Siemann. Experimental study of rf pulsed heating on oxygen free electronic copper. *Phys. Rev. ST Accel. Beams*, 5:112002, 2002.

[7] S. S. Manson. *Heat Transfer Symposium*. University of Michigan Engineering Research Institute, 1953.

[8] R. H. Siemann. *Symposium on Electron Linear Accelerators: In Honor of Richard B. Neal's* 80^{th} *Birthday*. Stanford Linear Accelerator Center, 1997.

[9] A. E. Vlieks et al. *Breakdown Phenomena in High Power Klystrons*. SLAC-PUB-4546, Stanford Linear Accelerator Center, 1988.

[10] P. B. Wilson. Scaling linear colliders to 5 TeV and above. In *ITP Conference on Future High Energy Colliders*, 1996.

[11] G. Mourou et al. Laser-induced breakdown by impact ionization in SiO_2 with pulse widths from 7 ns to 150 fs. *Applied Physics Letters*, 64:3071–3073, 1994.

[12] B. C. Stuart et al. Laser-induced damage in dielectrics with nanosecond to subpicosecond pulses. *Physical Review Letters*, 74:2248–2251, 1995.

[13] R. L. Byer et al. Progress of the laser electron accelerator project at stanford university. Presented at IEEE Particle Accelerator Conference (PAC 2001), Chicago, Illinois, 18-22 Jun 2001. Pub. in 'Chicago 2001, Particle accelerator' 108-110.

[14] A. Modena et al. Electron acceleration from the breaking of relativistic plasma waves. *Nature*, 377:606–608, 1995.

[15] Y.B. Fainberg et al. Wakefield excitation in plasma by a train of relativistic electron bunches. *Plasma Physics Reports*, 20:606–612, 1994. and ref. therein.

[16] T. Tajima and J. M. Dawson. Laser electron accelerator. *Physical Review Letters*, 43:267–270, 1979.

[17] K. Nakajima et al. Proof-of-principle experiments of laser wakefield acceleration using a 1 ps 10 TW Nd:glass laser. In *Proc. AIP Conf. Advanced Accelerator Concepts (AAC 1995)*, pages 145–155, 1994.

[18] S. P. D. Mangles et al. Monoenergetic beams of relativistic electrons from intense laser-plasma interactions. *Nature*, 431:535–538, 2004.

[19] C. G. R. Geddes et al. High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding. *Nature*, 431:538–541, 2004.

[20] J. Faure et al. A laser-plasma accelerator producing monoenergetic electron beams. *Nature*, 431:541–544, 2004.

[21] L'OASIS Group, Lawrence Berkeley National Laboratory. Private Communication, 2005.

[22] K. G. Steffen. *High Energy Beam Optics*. John Wiley, New York, 1965.

[23] S. Wang et al. X-ray emission from betatron motion in a plasma wiggler. *Physical Review Letters*, 88(14):135004, 2002.

[24] T. Katsouleas. Plasma wakefields - linear theory. Lecture given at Joint US CERN Japan Russia Accelerator School (JAS 2002). Long Beach, CA, 2002.

[25] J. C. Stewart and K. D. Pyatt. Lowering of ionization potentials in plasmas. *The Astrophysical Journal*, 144:1203–1211, 1966.

[26] P. Debye and E. Hückel. *Z. f. Phys.*, 24:185, 1923.

[27] J. M. Dawson. Nonlinear electron oscillations in a cold plasma. *Physical Review*, 113:383–387, 1959.

[28] E. Esarey et al. Overview of plasma-based accelerator concepts. *IEEE Transactions on Plasma Science*, 24:252–288.

[29] W. K. H. Panofsky and W. Wenzel. Transverse deflection of charged particles in radiofrequency fields. *Review of Scientific Instruments*, 27:967, 1956.

[30] R. Keinigs and M. E. Jones. Two-dimensional dynamics of the plasma wakefield accelerator. *Physics of Fluids*, 30:252–263, 1987.

[31] W. Lu et al., 2004. Private Communication.

[32] T. Katsouleas. Physical mechanisms in the plasma wake-field accelerator. *Physical Review A*, 33:2056–2064, 1986.

[33] T. Katsouleas et al. Beam loading in plasma accelerators. *Part. Accel.*, 22:81–99, 1987.

[34] C. Joshi et al. High-energy density plasma science with an ultrarelativistic electron beam. *Physics of Plasmas*, 9:1845–1855, 2002.

[35] P. Muggli et al. Meter-scale plasma-wakefield accelerator driven by a matched electron beam. *Physical Review Letters*, 93:014802, 2004.

[36] P. Chen et al. Energy transfer in the plasma wakefield accelerator. *Physical Review Letters*, 56:1252–1255, 1986.

[37] R. D. Ruth et al. A plasma wakefield accelerator. *Part. Accel.*, 17:171–189, 1985.

[38] J. Krall and G. Joyce. Transverse equilibrium and stability of the primary beam in the plasma wakefield accelerator. *Physics of Plasmas*, 2:1326–1331, 1995.

[39] D. Whittum et al. Electron-hose instability in the ion-focused regime. *Physical Review Letters*, 67:991–994, 1991.

[40] C. R. Vidal. Heat-pipe oven: A new, well-defined metal vapor device for spectroscopic measurements. *Journal of Applied Physics*, 40(8):3370–3374, 1969.

[41] C. L. O'Connell. *Plasma Production via Field Ionization*. PhD thesis, Stanford University, 2005.

[42] N. B. Delone M. V. Ammosov and V. P. Krainov. Tunnel ionization of complex atoms and of atomic ions in an alternating electromagnetic field. *Sov. Phys. JETP*, 64:1191–1194, 1986.

[43] D. L. Bruhwiler et al. Particle-in-cell simulations of tunneling ionization effects in plasma-based accelerators. *Physics of Plasmas*, 1:2022, 2003.

[44] B. V. Podobedov. *Saw-Tooth Instability Studies at the Stanford Linear Collider Damping Rings*. PhD thesis, Stanford University, 1999.

[45] R. L. Holtzapple. *Longitudinal Dynamics at the Stanford Linear Collider*. PhD thesis, Stanford University, 1996.

[46] M. Sands. The physics of electron storage rings: An introduction. SLAC-0121, Nov. 1970. 172pp.

[47] K. L. F. Bane et al. High-intensity single bunch instability behavior in the new SLC damping ring vacuum chamber. Presented at 16th IEEE Particle Accelerator Conference (PAC 95) and International Conference on High Energy Accelerators, Dallas, Texas, 1-5 May 1995, p. 3109-3111.

[48] J. Haissinski. Exact longitudinal equilibrium distribution of stored electrons in the presence of self-fields. *Il Nuovo Cimento*, 18B(1):72, 1973.

[49] R. Erickson (Editor). SLC Design Handbook. Stanford, USA: SLAC (1984) Nonconsec. Pag.

[50] K. L. Brown. A second order magnetic optical achromat. SLAC-PUB-2257.

[51] David C. Carey. Why a second order magnetic optical achromat works. *Nucl. Instr. Meth.*, 189:365, 1981.

[52] K. Bane and P. B. Wilson. Longitudinal and transverse wake potentials in SLAC. (talk). In High Energy Accelerators Proceedings, Geneva 1980, 592-596.

[53] V. P. Smirnov V. E. Balakin, A. V. Novokhatsky. VLEPP: Transverse beam dynamics. Proceedings of the 12$^{\text{th}}$ International Conference on High Energy Accelerators. Fermilab: 119, 1983.

[54] J. T. Seeman et al. Measured optimum BNS damping configuration of the SLC linac. Contributed to 1993 Particle Accelerator Conference (PAC 93), Washington, DC, 17-20 May 1993. Published in IEEE PAC 1993:3234-3236.

[55] K. L. F. Bane, A. Mosnier, A. Novokhatsky, and K. Yokoya. Calculations of the short-range longitudinal wakefields in the NLC linac. *eConf*, C980914:137–139, 1998.

[56] E. Keil. *Nucl. Instr. Meth.*, 100:419, 1972.

[57] K. L. F. Bane and P. Emma, 2005. Private Communication.

[58] P. Emma, L. Bentson, and P. Krejcik. A new bunch compressor chicane for the SLAC linac to produce 30-fsec, 30-kA, 30-GeV electron bunches. Contributed to the 8th European Particle Accelerator Conference (EPAC 2002), Paris, France, 3-7 Jun 2002. Published in 'Paris 2002, EPAC 02' 683-685.

[59] P. Emma et al. Measurements of transverse emittance growth due to coherent synchrotron radiation in the SLAC SPPS bunch compressor chicane. Prepared for Particle Accelerator Conference (PAC 03), Portland, Oregon, 12-16 May 2003. Published in IEEE PAC 2003:3129-3131.

[60] V. L. Ginzburg. Transition radiation and transition scattering. Proceedings, 16th International Cosmic Ray Conference, Vol. 14, Kyoto, 1979. Pages 42-50.

[61] X. Artru et al. Experimental investigations on geometrical resolution of electron beam profiles given by OTR in the GeV energy range. Laboratoire de l'Accelerateur Lineaire Report No. 97-30, 1997.

[62] V. L. Ginzburg and V. N. Tsytovich. Several problems of the theory of transition radiation and transition scattering. *Physics Reports*, 49(1):1–89, 1979.

[63] M. Castellano and V. Verzilov. Spatial resolution in optical transition radiation beam diagnostics. *Phys. Rev. Special Topics - Accelerators and Beams*, 1.

[64] A. G. Mozgovoi. The saturated vapour pressure of lithium, sodium, postassium rubidium, and cesium. *High Temperatures - High Pressure*, 19:425–430, 1987.

[65] Allegheny Ludlum Steel Corporation. *Stainless Steel Fabrication.* Allegheny Ludlum Steel Corporation, Pittsburgh 22, PA, 1959.

[66] M. Borland. ELEGANT: A flexible SDDS-compliant code for accelerator simulation. Advanced Photon Source LS-287, September 2000.

[67] P. Muggli et al. Coherent transition radiation to measure the SLAC electron bunch length. In *Proceedings of the 2005 Particle Accelerator Conference (PAC) Knoxville, TN, 16-20 May 2005.*

[68] M. L. Ter-Mikaelian. *High-energy Electromagnetic Processes in Condensed Media.* Wiley-Interscience, New York, 1972.

[69] M. A. Geitz. *Investigation of the Transverse and Longitudinal Beam Parameters at the TESLA Test Facility Linac.* PhD thesis, Universitat Hamburg, 1999.

[70] A. E. Siegman. *Lasers*. University Science Books, Sausalito, 1986.

[71] J. Seeman et al. SLC energy spectrum monitor using synchrotron radiation. Contributed to Linear Accelerator Conf., Stanford, CA, Jun 2-6, 1986. Published in Stanford LINAC86:441.

[72] J. Arthur et al. Linac Coherent Light Source (LCLS) conceptual design report. pages 7–36 to 7–41. SLAC-R-593.

[73] P. Sprangle et al. Laser driven acceleration in vacuum and gases. *AIP Conf. Proc.*, 398:96–105, 1997.

[74] E. R. Colby. *Design, Construction, and Testing of a Radiofrequency Electron Photoinjector for the Next Generation Linear Collider*. PhD thesis, The University of California, Los Angeles, 1997.

[75] R. H. Evans. *The Atomic Nucleus*. McGraw-Hill, New York, 1955.

[76] A. W. Chao (ed.) and M. Tigner (ed.). Handbook of accelerator physics and engineering. pages 215–217. World Scientific, Singapore (1999).

[77] J. D. Jackson. *Classical Electrodynamics, Third Edition*. John Wiley and Sons, New York, 1998.

[78] K. Bane P. Emma. LiTrack: A fast longitudinal phase space tracking code with graphical user interface. In *Proceedings of the 2005 Particle Accelerator Conference (PAC) Knoxville, TN, 16-20 May 2005*.

[79] M. J. Hogan et al. Multi-GeV energy gain in a plasma-wakefield accelerator. *Phys. Rev. Lett.*, 95:054802, 2005.

[80] M. M. Zhou et al. Modeling self-ionized plasma wakefield acceleration for afterburner parameters using QUICKPIC. Submitted to Particle Accelerator Conference (PAC 05), Knoxville, Tennessee, 16-20 May 2005.

[81] J. Orear. Notes on statistics for physiciscts, revised. *Cornell Laboratory for Nuclear Studies*, 82/511, 1982.

[82] W. H. Press et al. *Numerical Recipes: The Art of Scientific Computing (FORTRAN Version)*. Cambridge University Press, Cambridge, 1990.

[83] L. Montanet et al. Review of particle properties. particle data group. *Phys. Rev.*, D50:1173–1823, 1994.