Chapter 3

Wave-Particle Correlation in the Auroral Ionosphere: CHARM-II

3.1 Introduction

Langmuir waves, also known as electron plasma waves, are one of the most fundamental properties of a plasma, having been first observed in discharge plasmas in the early days of plasma physics (*Langmuir* 1928). They result from the interaction of electron beams with plasmas and hence are ubiquitous in space plasmas, including, for example, the solar wind, where they generate radio bursts (*Lin et al.* 1981), and planetary foreshocks (*Gurnett et al.* 1981; *Filbert and Kellogg* 1979) and auroral ionospheres (*Kintner et al.* 1995; *Boehm* 1987; *McAdams* 1999; *Samara* 2005), where they mediate energy transfer between the beam and thermal plasmas. Langmuir waves can generate nonlinear structures of fundamental interest to plasma physics, as well as linear eigenmode effects in inhomogeneous plasmas (*McAdams et al.* 2000; *Ergun et al.* 2008). Due to their significance and abundance in the space environment, Langmuir waves are a subject of current study; particularly with regard to their eigenmode structures (*Malaspina et al.* 2012), three-dimensional effects (*Malaspina and Ergun* 2008; *Dombrowski et al.* 2012), and wave-particle correlations (*Ergun et al.* 1991b,a; *Muschietti et al.* 1994; *Kletzing et al.* 2005).

Particle correlation experiments have proven to be an effective way to probe wave-particle interaction physics in space plasmas. A detailed theory of expected results from such instruments regarding Langmuir waves is given by *Kletzing and Muschietti* (2006). The phase bunching of the electrons in the field of the wave can be considered as a superposition of two components, a 'resistive' component which is in phase with the wave electric field and represents energy transfer either from wave to particles or vice versa, and a 'reactive' component which is in quadrature phase with the wave field and is a signature of electrons trapped in the wave. An early version flown on a sounding rocket in auroral plasma determined a strong correlation between beam electrons and Langmuir/upper hybrid wave electric fields over a several hundred second interval (*Gough et al.* 1990). A wave particle correlation experiment was flown on the Freja spacecraft (*Boehm et al.* 1994). *Ergun et al.* (1998) flew

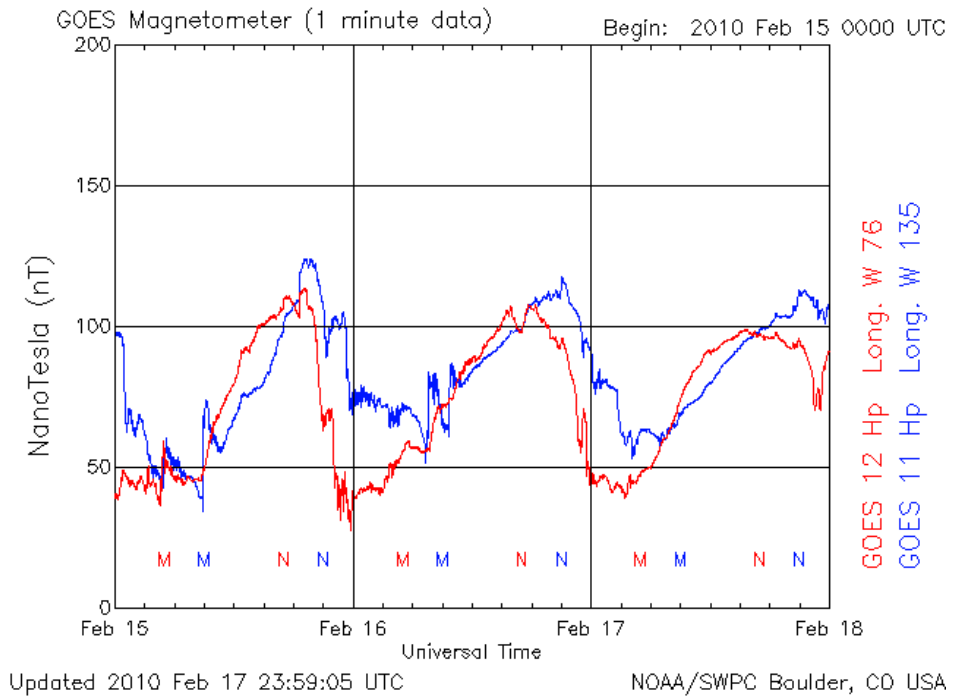

Figure 3.1: GOES magnetometer data for 15-18 February 2010. The CHARM-II launch occurred at 9:49 UT on the 16th, motivated by the preceding steep ~ 20 nT bay in the field.

a wave-particle correlator on an auroral rocket. The design was very similar to the instrument described below, except that the wave period was divided into only four phase bins rather than sixteen; the experiment revealed evidence of wave-particle interactions but could not resolve resistive versus reactive components. *Kletzing et al.* (2005) reports results from an instrument nearly identical to the one described below, launched into nighttime aurora. For various reasons, the experiment measured correlations associated with a relatively small number of the most intense Langmuir waves encountered, but the results gave a strong indication of wave trapping of the bunched electrons in those examples.

The success of previous wave-particle correlator experiments inspired a series of rocket experiments denoted the Correlation of High-Frequency and Auroral Roar Measurements (CHARM). The first CHARM launch experienced a payload system failure that precluded any correlator data-taking. A re-flight, CHARM-II, was launched from the Poker Flat Research Range near Fairbanks, AK, at 9:49 UT/22:46 MLT on 16 February 2010, reaching an apogee of 802 km. The launch, shown in Figure 3.2, was into an active substorm expansion phase, characterized by a 20 nT bay in the H-component of the magnetic field observed by GOES 11, as seen in Figure 3.1. The payload carried a Dartmouth High-Frequency Experiment and University of Iowa Correlator, as well as a number of other primary and contextual instruments. One particularly intense event encountered was reported on by *Kletzing et al.* (2011).

This work presents a comprehensive investigation of the entire wave-particle correlator data set from the CHARM-II mission. Section 3.2 covers the instruments which make up the

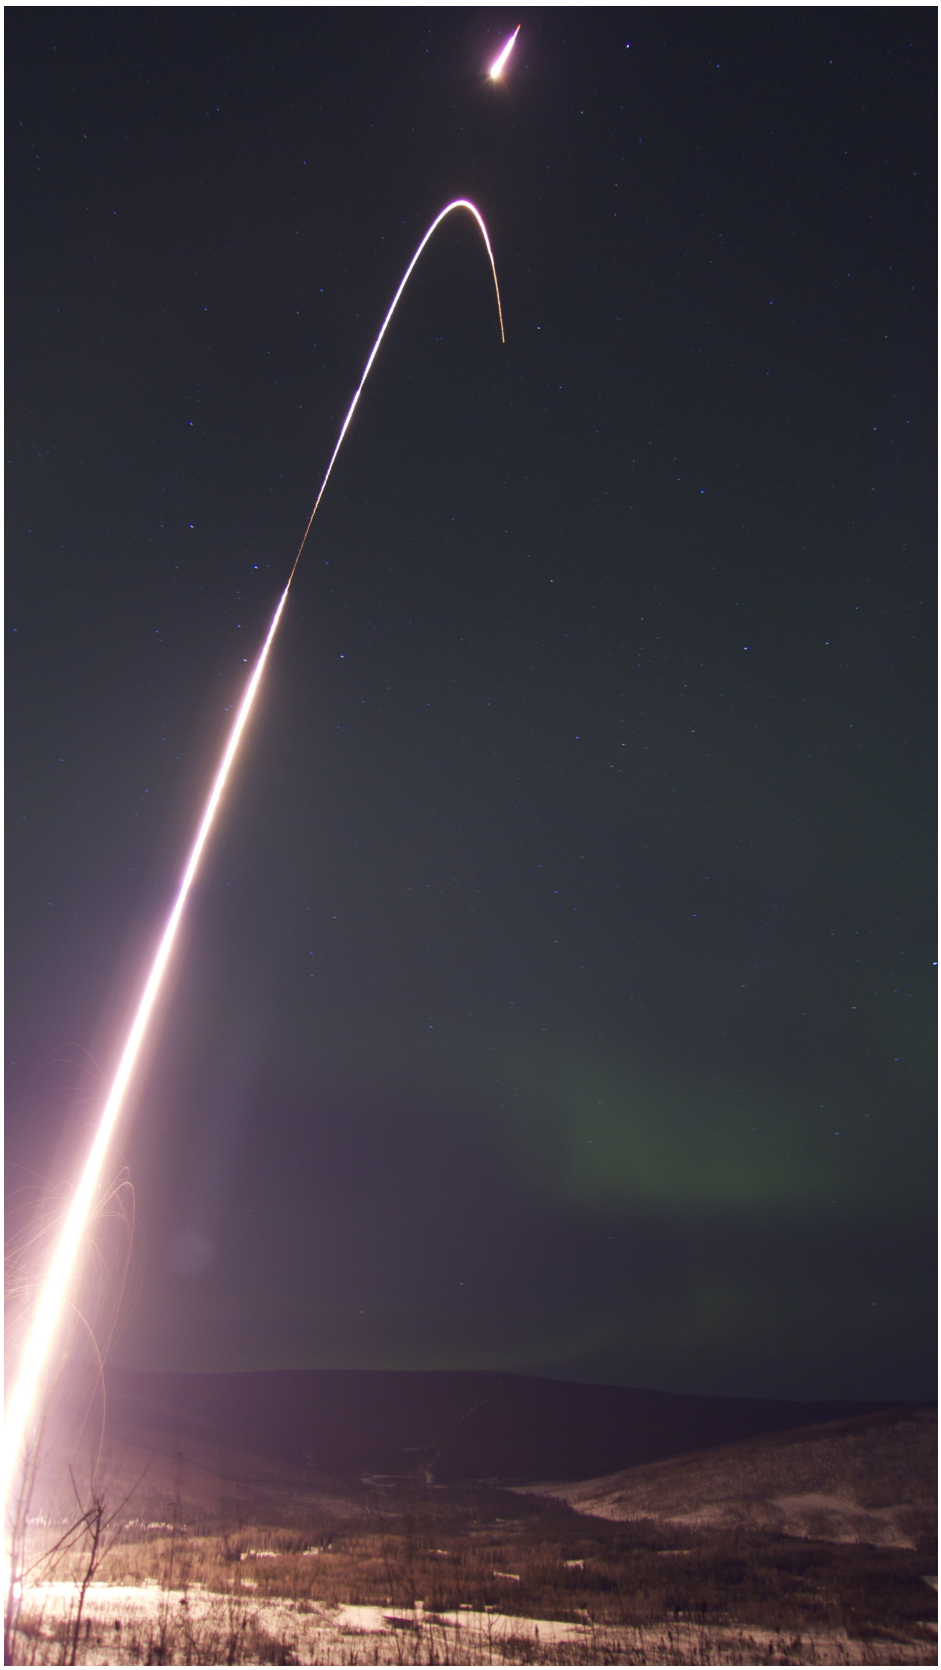

Figure 3.2: A photo of the CHARM-II launch.

Correlator system, and the form of the returned raw data. Section 3.3 presents the methods used to identify significant correlation events, a measure which characterizes individual events, several contrasting events from the final set, and observations regarding the set as whole. Section 3.4 summarizes these results and offers two theories developed to explain the observations. Finally, Section 3.5 develops a numerical test-particle simulation to test the plausibility of these theories.

3.2 Instrumentation

Accurate, in-situ correlation of Langmuir waves and electrons requires three primary pieces: a wave instrument covering the range of frequencies in which Langmuir waves are expected, high-speed particle detectors at a range of potentially resonant energies, and the correlation hardware itself, which processes these data streams, and reduces and returns the desired statistics.

The Dartmouth High-Frequency Experiment (HFE) detects the potential difference between two 2.5 cm spherical probes, separated by 30 cm along the payload's spin axis. This ΔV signal provides an estimate of the the axial component of the electric-field, which is mainly parallel to the ambient magnetic field, since the payload is kept field-aligned to within $\sim 10^\circ$ by an attitude control (ACS) system. Active preamplifiers inside each spherical probe assure that the antenna functions as a double-probe over the entire 0-5 MHz frequency range. The signal is band-pass filtered to a 100 kHz to 5 MHz band, and regulated by an Automatic Gain Control (AGC) system to enhance the dynamic range. The AGC control signal is sampled onboard at 20 kHz and telemetered with other digital data. The regulated HF signal directly modulates a 5 MHz-bandwidth S-band transmitter, and the resulting waveform is continuously digitized at the ground telemetry station at 10 MHz, with 12-bit resolution. This instrument is the latest iteration of a design which has flown on numerous other rocket campaigns in both $\boldsymbol{E}_\parallel$ and $\boldsymbol{E}_\perp$ configurations, including HIBAR (*Samara et al.* 2004), PHAZE II (*McAdams et al.* 1998), SIERRA, RACE (*Samara and LaBelle* 2006), and ACES (*Kaeppler et al.* 2011).

The University of Iowa Wave-Particle Correlator similarly has heritage on numerous sounding rocket missions, including RACE and CHARM, and is described in detail by *Kletzing et al.* (2005). The Correlator takes an input waveform from the HFE, and uses it to control a phase-locked loop (PLL) circuit running at 16 times the frequency. The PLL phase-locks onto the frequency of the highest-amplitude component of the incoming wave, and restores to baseline, maintaining a 50% duty cycle. In the case of the HFE signal, the waveform is strongly dominated by the component at the Langmuir frequency when plasma waves are unstable. Under this condition, the PLL produces a clean, square-wave version of the Langmuir wave. This wave is then divided into 16 bins along its phase, and incoming counts from each of four detectors are sorted into these bins during an integration period—1 ms in the CHARM 2 case—corresponding to hundreds of wave periods per timeslice.

For the CHARM-II mission, two correlators were flown, each receiving particle data from four 'bagel' particle detectors. These detectors, named for the baked goods they resemble,

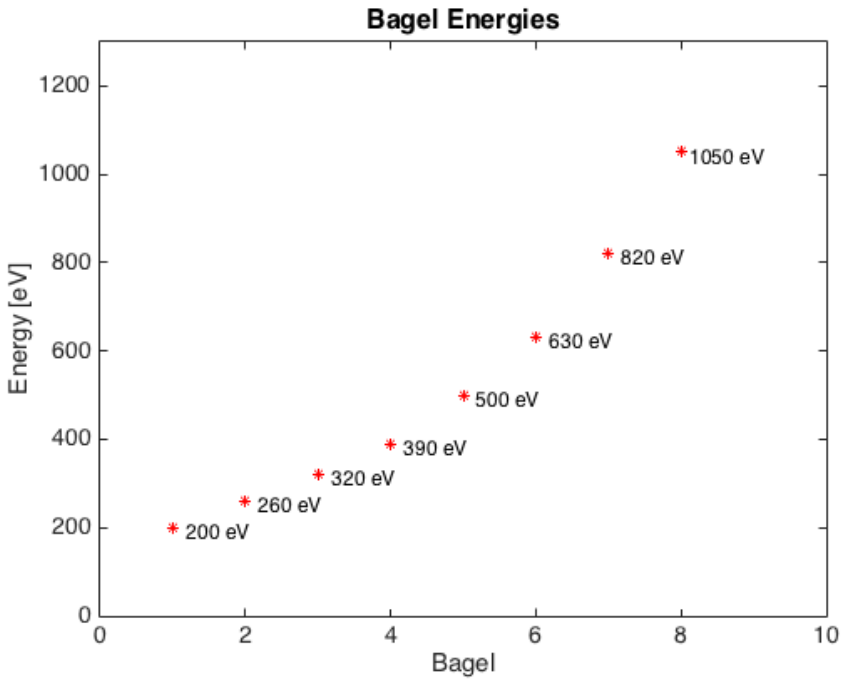

Figure 3.3: The energies of the eight 'bagel' particle detectors.

have energy acceptance ranges of 10%, and are characterized by a large geometric factor, as required for correlation with high-frequency waves (*Kletzing and Muschietti* 2006). The detectors are aligned with the rocket's spin axis, with a 10°-wide field of view, and thus are always observing the field-aligned component of incoming particles. The bagel detectors were tuned to logarithmically spaced energy levels ranging from 200 to 1050 eV (see Figure 3.3).

An additional contextual instrument of interest is the Electrostatic Electron Pitch Angle Analyzer (EEPAA), a 'top-hat' style detector which counts electrons, sorted into 15°-wide pitch angle bins and 47 logarithmically spaced energy bins from 15 eV to 15.5 keV, with a 50 ms integration time.

3.3 Data Presentation

Figure 3.4 shows a summary plot of the active period of the CHARM-II flight, with both time after launch and altitude on the x-axis. On top is a spectrogram of EEPAA data, with energy on the y-axis, and the log of differential flux as color intensity. Middle is an HFE spectrogram, with frequency on the y-axis, and color following wave power in decibels. Finally, the bottom is a plot of $\log_{10}$ of total counts among all eight Bagels. In the EEPAA data, an inverted-V structure is clear from approximately 610 to 660 s, with a more tenuous one from 500 to 560 s. The upper cutoff to noise which is near 500 kHz on the left of the HFE panel is interpreted as the Langmuir frequency ω_p, which acts as an upper bound to whistler modes in 'underdense' plasmas, where ω_p is less than the cyclotron frequency ω_c. From this,

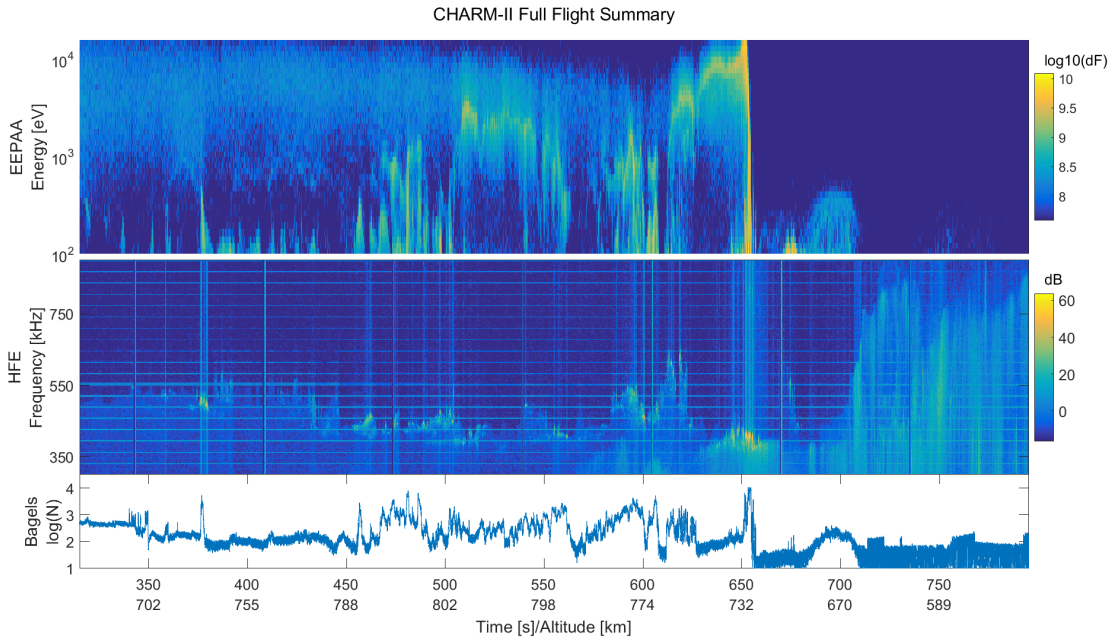

Figure 3.4: A summary plot of the active period of the CHARM-II flight. Both time after launch and altitude are shown on the x-axis. Top is a spectrogram of EEPAA particle data, with energy on the y-axis, and $\log_{10}$ of the differential flux [N * eV/(cm^2 * ster * s * eV)] as color intensity. Middle is an HFE spectrogram, with frequency on the y-axis, and color following wave power in decibels. Bottom is a plot of $\log_{10}$ of the total counts among all eight Bagels detectors. Note a clear inverted-V structure at 610 to 660 s, and a more tenuous one from 500 to 560 s. The upper cutoff to noise near 500 kHz on the left of the HFE panel is interpreted as the Langmuir frequency ω_p.

it is clear that the Langmuir frequency is much lower than the Upper-Hybrid Frequency (~ 1.4 MHz), and so easily selected for despite variance in the rocket's alignment parallel to the magnetic field. Given the bagel energy range it can also be deduced that the wavelengths in question will range from approximately 10 to 60 meters. Finally, it is clear that there are many instances where increased particle counts are accompanied by wave activity near ω_p. The activity near 650 to 660s was a particularly powerful event which saturated all onboard electric-field instruments at its peak, though the HFE data (and thus the Correlator) was only unusable for a few milliseconds. The peak Langmuir-wave electric-field intensity was estimated between 1 to 3 V/m, with the Langmuir frequency near 350 kHz.

The Correlator system returned approximately 489 seconds of valid raw data, providing a matrix with counts $s(t,p,E)$ at each of 488,869 timeslices (t), 16 phase bins ($p = 0...15$), and 8 energy levels ($E = 1...8$). While the phase between particles and the input waveform varies based on frequency, and due to daisy chaining of the HFE signal between the two Correlators, a first step taken to aid comparability of timeslices is to shift all bins to the same baseline, based on the recorded Correlator frequency. While the most direct way of looking at this data might be to display raw counts vs. phase and time, it is generally more edifying to

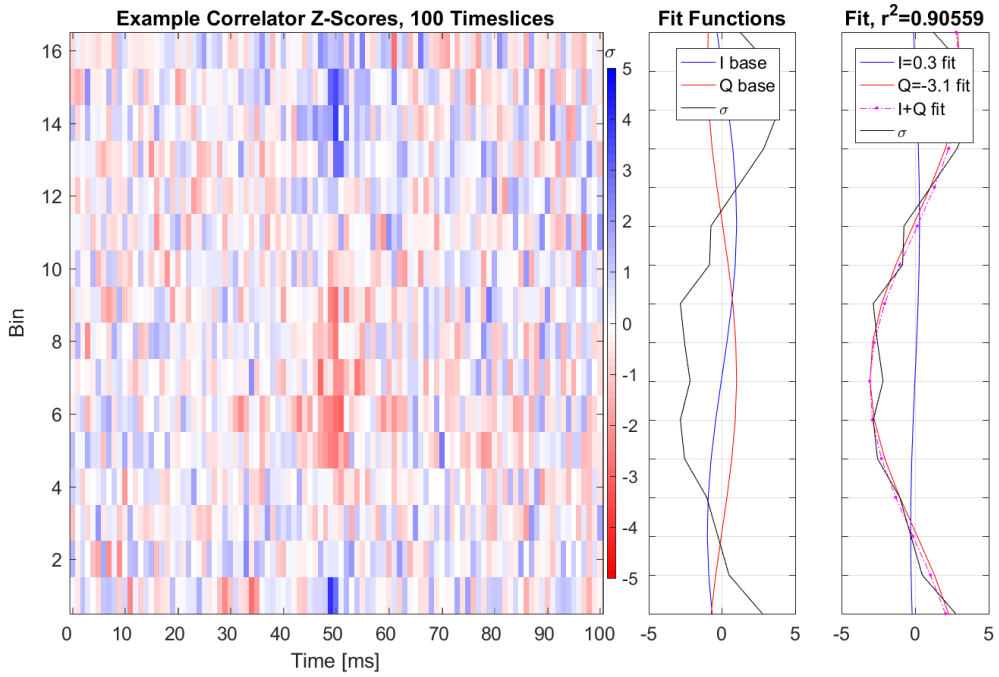

Figure 3.5: An example plot showing (left) the Poisson z-scores (σ) vs. phase and time for 101 ms of CHARM-II Correlator data. Z-score from -5 to 5 is shown as blue-to-red color scale, electric field phase is on the y-axis, and time (in relative ms) on the x-axis. Note that while a high-$|\sigma|$, multi-timeslice event is clearly visible in the middle timeslices, many more timeslices are insignificant. In the left line plot, the σ values for the central timeslice are shown, along with the base (unity amplitude) functions for an I/Q Resistive/Reactive fit, while in the right line plot they are shown with the fit amplitudes and their sum. The fit is reasonable, as shown by the r^2 value.

examine the Poisson z-score for the data,

$$\sigma(t,p,E) = \frac{s(t,p,E) - \bar{s}(t,E)}{\sqrt{\bar{s}(t,E)}},$$

where $\bar{s}$ is the mean particle count of the timeslice, $\bar{s} = \frac{1}{16}\Sigma_p s(t,p,E)$. This is a measure of particle over- or under-density in a given bin, with respect to the mean for that timeslice. This analysis yields plots as in Figure 3.5 (left), showing z-scores for 101 timeslices, with time on the x-axis, electric field wave phase on the y-axis (with each timeslice shifted such that the zero phases are aligned), and z-score shown as color.

While Figure 3.5 shows significant timeslices, it also shows large regions of low significance and noise. Likewise, perusal of the complete set of timeslices, $\mathbb{F}$, makes clear that many timeslices can be discarded due to a lack of significance and/or natural and instrumental interference, and indeed each timeslice can be classified by Correlator telemetry as 'locked' or not, and less than 15% of the data set has both Correlators locked. In addition, a timeslice cannot be considered reliable merely from the presence of a lock state at that timeslice, and

this says nothing about the presence of interesting activity at that time. With so much data, manual inspection was not a practical or desirable method to identify reliable lock or significance, so an automated algorithm for event identification was developed.

The reduction analysis is motivated by an expectation of how significant wave-particle correlations will manifest themselves in the data: as a sine wave in the phase bins, with a quarter of the bins having a statistical excess of counts, and a quarter having a deficit. The pattern arises because of the bunching of the particles in the electric field of the Langmuir waves to which the PLL is locked.

The first step towards reduction of the Correlator data to identify discrete 'events' is, as shown in Figure 3.6, re-binning the 16 raw data phase bins p (orange) into four reduced phase bins p' (color coded by p'),

$$s'(t,p',E) = \sum_{n=0...3} s(t,4p'+n,E),$$

in order to emphasize the expected pattern of a quarter of the bins having over and under-dense counts. The re-binning was done four times ($q = 1...4$, depicted as individual columns of p' bins), shifted by one raw bin for each, to cover all possible patterns that might result from a wave-particle correlation event,

$$s'_q(t,p',E) = \sum_{n=0...3} s(t,4p'+n+q,E),$$

and the z-score,

$$\sigma'_q(t,p',E) = \frac{s'(t,p',E) - \bar{s}}{\sqrt{\bar{s}}},$$

was then calculated for each of these reduced timeslices.

To account for events which span multiple timeslices, two 'doublet' and one 'triplet' sets were constructed at each timeslice, acting as additional arrays of timeslices (the large divisions in Figure 3.6). For these, the means and the counts in each reduced phrase bin were integrated in time over the two or three raw timeslices, with the doublet sets defined as

$$s'_{q,t2b}(t,p',E) = \sum_{\tau=t-1,t} s'_q(t,p',E) \text{ and } s'_{q,t2f}(t,p',E) = \sum_{\tau=t,t+1} s'_q(t,p',E),$$

and the triplet

$$s'_{q,t3}(t,p',E) = \sum_{\tau=t-1,t,t+1} s'_q(t,p',E),$$

i.e. the doublets integrate either the timeslice prior or after, and the triplet both. All three sets then have their associated $\sigma'_{q,t2b}$, $\sigma'_{q,t2f}$, and $\sigma'_{q,t3}$. Thus, the process yields four total σ' arrays over the four q values, times four timeslice arrays (singlet, doublets, triplet), or sixteen total arrays.

Finally, as a criterion to identify timeslices with interesting events, we find the global minimum and maximum σ' over the sixteen arrays at each timeslice. The difference between

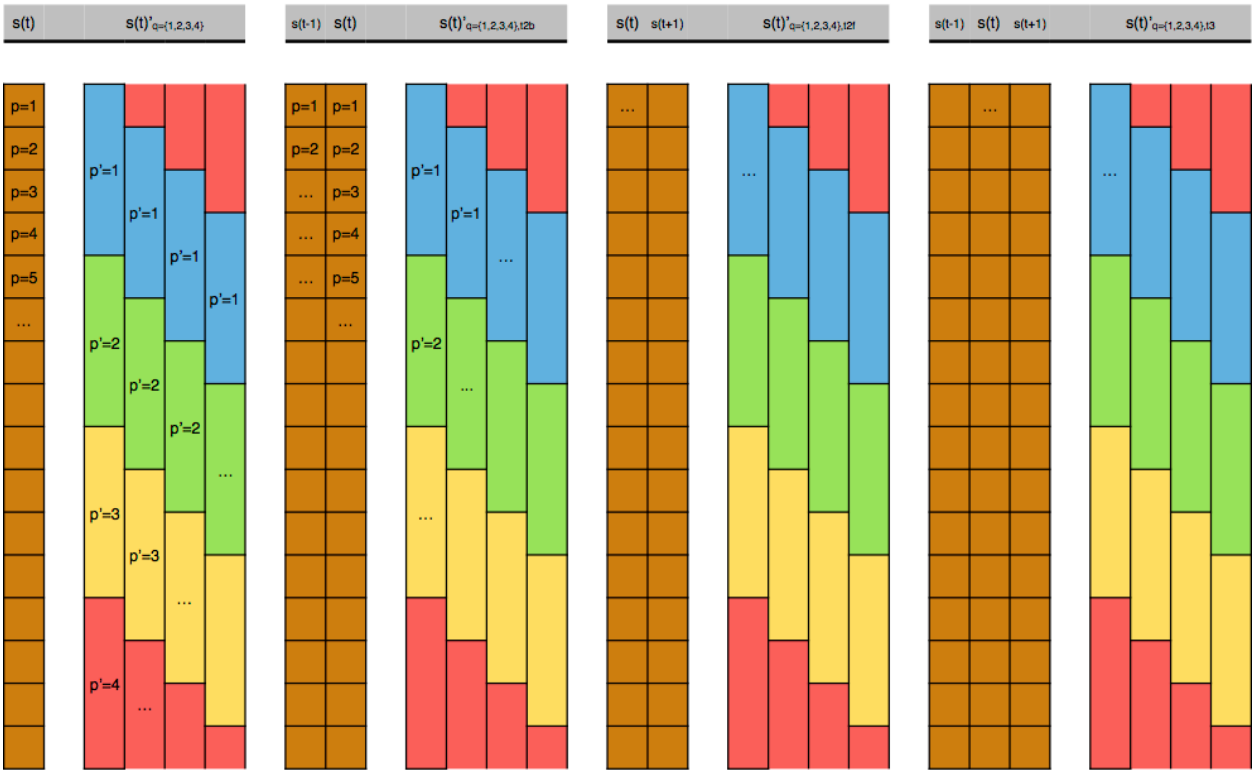

Figure 3.6: A map of the rebinning done at each timeslice to the raw correlator counts. Each large group is a time integration (singlet, doublets, triplet), and each of those contains four shifted rebinnings from the original 16 correlator bins of $s(t,p,E)$ (shown in orange), to the 4 bins of $s'_q(t,p',E)$ (shown in other colors, grouped by p').

that min and max, Δ, is then a scalar measure of how well a given timeslice matches the expected signature of a wave-particle correlation event.

Initially, a simple global threshold was used to find potential events, resulting in hundreds of identified timeslices; however, thorough investigation of these revealed many false positives among these sets, including many events which were disqualified after examination of diagnostic and contextual data. Figure 3.7 shows the two major factors which lead to disqualifying timeslices. In the second plot from the bottom, a powerful interfering signal can be seen in the raw HFE waveform data, at a cadence of about 15 μs. This signal, originating on the rocket payload, is frequently seen in the HFE data throughout the flight, and in some cases is the highest-amplitude component of the waveform. While the Correlator filters out such a low frequency, its presence is often correlated with the second disqualifying factor: a dubious recorded PLL lock frequency. Visible as the solid white line overlaying the HFE spectrogram in the bottom plot, the lock frequency can either be unstable (i.e. not properly locked), or set to unrealistically high or low values compared to by-eye evaluation of the Langmuir cutoff (dashed black overlaying line). In either case, this drastically reduces confidence in the correlator phase binning correctness, and motivates the discarding of such events.

The large number of false identifications in the initial run also revealed that the Δ-threshold needs to be different for each bagel detector. With these issues in mind, the event identification algorithm was altered to iteratively optimize the Δ threshold for each bagel. Subsequently, manual screening was applied based on the considerations above. The 820 and 1050 eV bagels had no qualifying timeslices. Figure 3.8 shows an overview of the final set $\mathbb{S}$ of thresholded, hand-screened events, as x-marks at a given bagel and time, overlaid on a 20-bin histogram. From the histogram it is clear that majority of the events are in a tight cluster near 652 seconds, coincident with the strong event at the end of the major inverted-V structure seen in Figure 3.4. Two longer clusters of events are centered near 490 and 610 seconds. The per-bagel set sizes are shown in the table inset in the upper left of Figure 3.8, showing that the majority of the events were at 260 and 630 eV, with 12 and 23 events, respectively.

Linear analysis of the interactions of Langmuir wave packets with electrons has shown that the perturbation of the electron distribution function can be broken into two components: the 'resistive' component of trapped particles which oscillate in-phase with the electric field of the Langmuir wave, or 180 degrees out of phase, and the 'reactive' component which oscillate 90 or 270 degrees out of phase (*Kletzing and Muschietti* 2006). A strong resistive component is an indicator of wave-particle energy exchange, leading to wave growth or damping, while the reactive phase is associated with particle trapping. The summation of these two components will tend to have a sinusoidal form when either component shows significant activity, and it is this form that the event-identification method focused on.

To look at the Correlator data in a comparable manner, we fit the correlator timeslices to a quadrature function vs. bin number p,

$$-I\sin\left((p-p_0)*\frac{\pi}{8}\right)-Q\cos\left((p-p_0)*\frac{\pi}{8}\right)$$

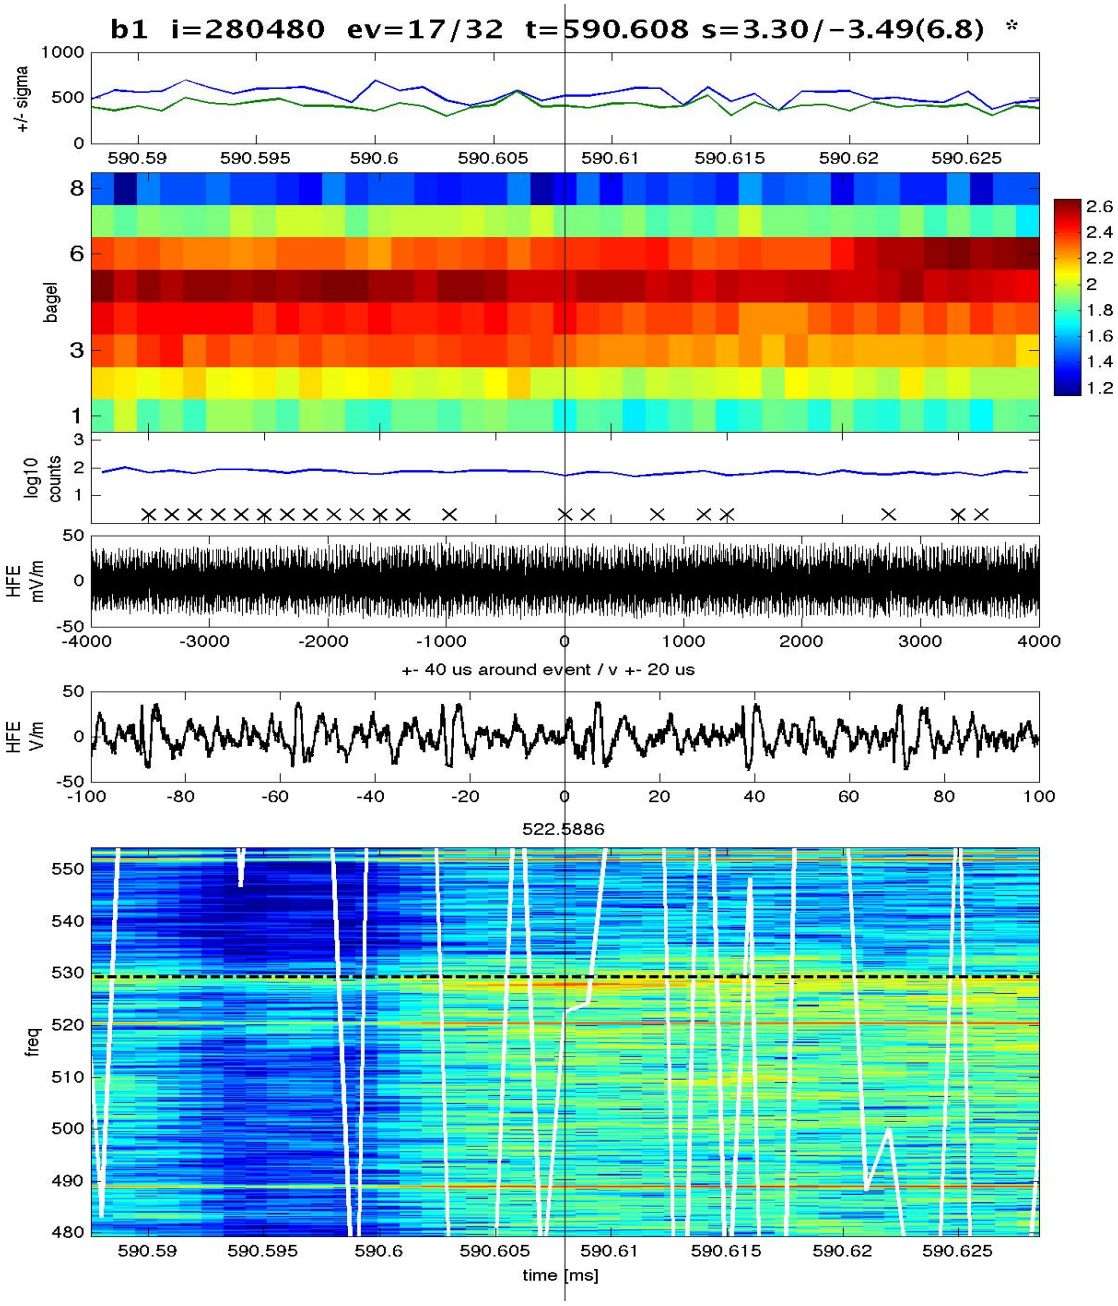

Figure 3.7: An example of plots used to hand-screen events. From the top, the plots show max(σ) and min(σ), a counts vs bagel & time spectrogram, the logarithm of raw counts for a single bagel (with PLL lock/no-lock status displayed at the bottom, as an X for 'lock'), two timescales of raw HFE waveforms, and an HFE spectrogram with PLL lock frequency (white, solid) and hand-picked Langmuir frequency (black, dashed) overlaid. This event was discarded because it shows both a wildly fluctuating lock frequency with only sporadic 'lock', and strong periodic interference of unknown origin at approximately 15 μs cadence.

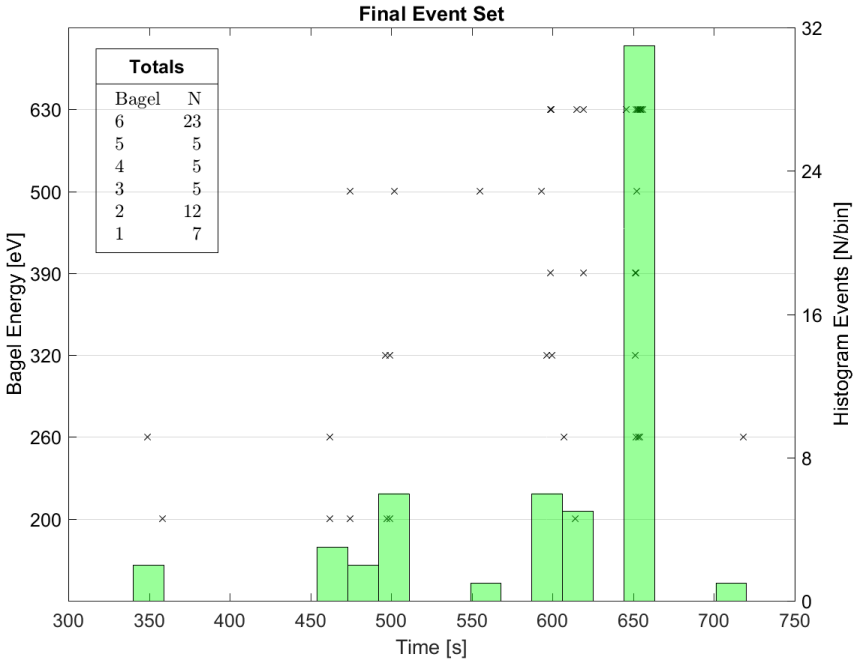

Figure 3.8: An overview of the final event set $\mathbb{S}$ vs. time for the whole flight, with the per-bagel totals in the table inset top left. Individual events are displayed as x-marks on lines corresponding to bagels on the left vertical axis. Overlaid on this is a histogram of events vs. time, corresponding to event/bin counts on the right vertical axis.

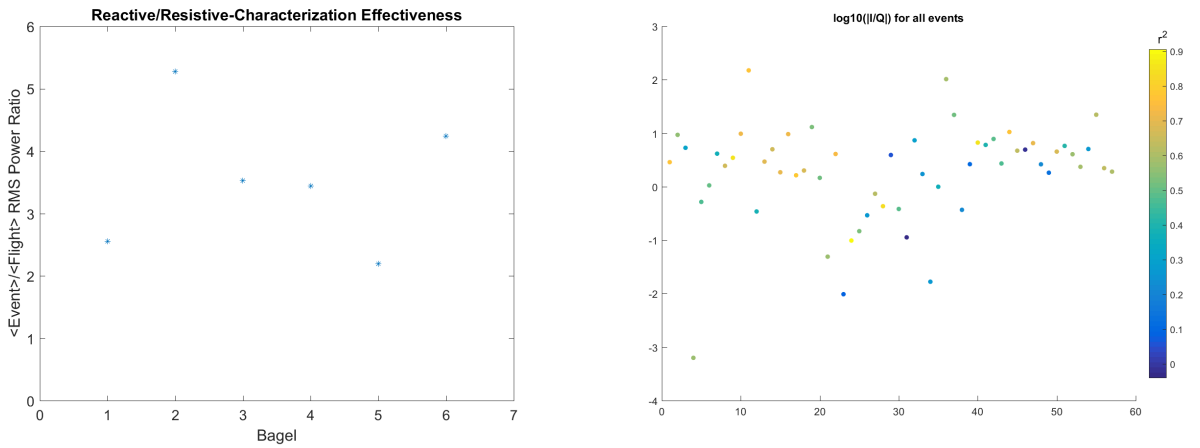

Figure 3.9: On the left, for each bagel, the ratio of $\langle\sqrt{I^2+Q^2}\rangle_{\mathbb{S}}/\langle\sqrt{I^2+Q^2}\rangle_{\mathbb{F}}$, showing that the events in $\mathbb{S}$ show significantly more reactive/resistive activity than seen in the general flight. On the right, a plot showing the resistive/reactive power ratios, as $\log_{10}|I/Q|$, for all of $\mathbb{S}$, showing some variance, and resistive activity in a majority of events.

optimizing I and Q to best fit the z-scores. The negative signs and p_0 are determined by calibration data relating the electric-field phase to the bins. After performing this type of fit on all timeslices, the coefficients I and Q are then magnitudes of the In-phase and Quadrature signals, i.e. the resistive and reactive linear components. An example set of fits are shown in the line plots on the right of Figure 3.5, with bins on the y-axis (aligned to the left $\boldsymbol{\sigma}$ plot)—in both line plots, the black line is the $\boldsymbol{\sigma}$ values for the center timeslice. In the left plot, the solid blue line is the calibrated electric field, corresponding to I, the resistive component, and the dashed red line is reactive Q component. The right plot shows fitted forms, with blue and red the fitted, separate I and Q, and the dot-dashed magenta line their sum (i.e. the actual fit function). For this timeslice, the r^2 goodness-of-fit is ~ 0.906, showing a reasonable fit, and with $I \sim 0.3$ and $Q \sim -3.1$ this event appears to be dominated by the reactive component.

Figure 3.9 shows diagnostics of this fitting. The plot on the left has the per-bagel ratios of the mean magnitudes I and Q when taken over only $\mathbb{S}$, compared to the mean magnitude for all of $\mathbb{F}$. From this it's clear that this form of analysis does a reasonable job of characterizing events, though it is worth noting that it yields enough false positives and negatives to make it less useful than the above-outlined method for event identification. The right plot shows the log of the magnitude of the resistive-to-reactive power ratio, $\log_{10}|I/Q|$, for $\mathbb{S}$, revealing significant variance, and that a majority of events are more resistive.

Figure 3.10 displays several events of interest from $\mathbb{S}$. Each shows, from the top, 1 second of contextual data from the EEPAA, a spectrogram of the total counts of each bagel vs. bagel energy, the values of I, Q, and $\sqrt{I^2+Q^2}$, and $\boldsymbol{\sigma}$ vs. wave phase.

Considering the top left stack of Figure 3.10, which shows data from an event identified in the 260 eV bagel's data, we see some evidence of higher-energy beams in the EEPAA data leading up to the event, but nothing significant during it. The bagel spectrogram shows

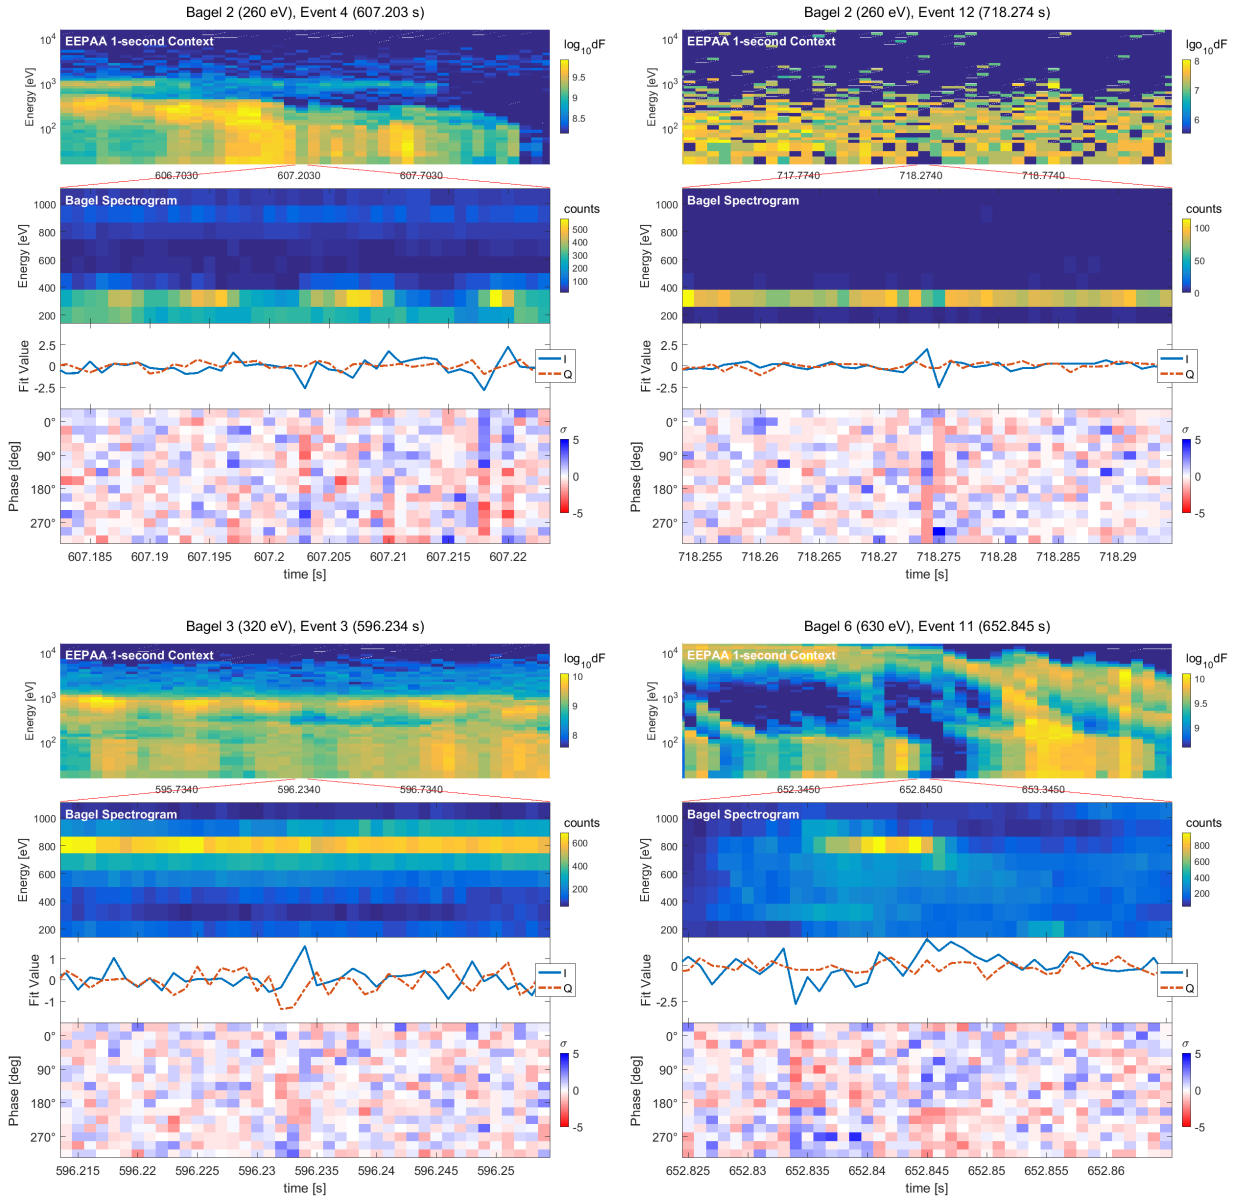

Figure 3.10: Four example events from the final, Δ-thresholded, hand-screened set $\mathbb{S}$, with 1 second of EEPAA/top-hat particle detector data on top, for context, followed by a spectrogram of bagel total counts vs. energy, a line plot of reactive/resistive fit values, and a $\boldsymbol{\sigma}$ vs. wave phase plot. Note the presence of multi-timeslice events in the $\boldsymbol{\sigma}$ data, as well as short-lived narrow-band beam features in the bagel spectrogram, the presence of different 'red over blue' and 'blue over red' regions in close temporal proximity, and the behavior of I and Q during these events.

sporadic, short-timescale (2-5 ms) beams in this bagel alone. The central timeslice in the σ plot is the event which passed the thresholding and hand-screening, identifiable as a vertical bar with strong blue bins in the top half, and strong red bins in the bottom half, showing the characteristics which were selected for by the thresholding process. There are several other timeslices around this event which show similar signals, or signals with red bins over blue bins, and it is clear that the resistive component of our I/Q fit is differentiating between the red-over-blue and blue-over-red cases. It also appears that the stronger beams show a relation to I, with negative-I events preceding beams, and positive-I following.

In the other plots of Figure 3.10, we see that while there are several cases of timewise-longer events, the majority are single-timeslice (1 ms) events as in the top left, a trend that holds throughout $\mathbb{S}$. The top-right stack shows extremely low counts on the EEPAA for unknown reasons. While only the lower-right stack shows a clearly distinct beam, the pattern does appear to hold that a negative I comes before or during a density increase, while a positive I correlates with a density decrease.

Figure 3.11 highlights the event previously presented by *Kletzing et al.* (2011), which covers a larger time and energy range. This figure shows the same parameters as those in Figure 3.10, but I/Q and σ plots are shown for the 200 to 630 eV bagels. The EEPAA data shows clear evidence of a dispersive beam appearing at or slightly before this event, lasting for approximately 200 ms, and the bagel spectrogram and σ plots also appear to show this at finer timescales. For this event, both I and Q show strong responses, and I appears to show the same association with the density gradient at a given bagel.

Finally, Figure 3.12 shows a set of events which take place over a similar time scale to those in Figure 3.11. Evidence for dispersion is less clear here, especially given the large gap between the 260 eV and 630 eV bagels in which events are visible. However, in each of these examples I displays the same relation to the beam as observed in Figures 3.10 and 3.11; that is, that a negative I is associated with a density increase, and vice versa.

Figure 3.13 shows scatter plots of I and Q values versus two selected parameters, the temporal gradient of the electron beam flux ∇n_B, inferred from the bagel detectors (upper row), and the temporal gradient in the RMS HFE waveform $\nabla < w >$ (lower row). The points are colored according to their r^2 goodness-of-fit value from the reactive/resistive fitting. The upper left panel demonstrates from this statistical approach the correlation between the I value and the beam gradient which has been illustrated by multiple examples in Figures 3.10 and 3.12. A clear trend is evident whereby negative I values correspond to positive beam flux gradients, and vice versa. A linear regression to these points returns a t-statistic value of -6.97 with a p-value of 4×10^{-9}. The other panels of Figure 3.13 show that there is no pattern evident between the Q value and ∇n_B, or between the I or Q values and $\nabla < w >$. Linear regressions of these sets all have small t-statistics, and p-values ≥ 0.15, strongly suggesting that the I-∇n_B relationship is the only significant one of those examined. Table 3.1 summarizes additional statistical tests performed on the data, showing that a Kolmogorov-Smirnov test finds a significant difference between the $\nabla n_B < 0$ and $\nabla n_B > 0$ distributions of I, and that I and ∇n_B are the only significantly correlated measures.

Additional by-eye comparisons were made between the I/Q responses and high-cadence

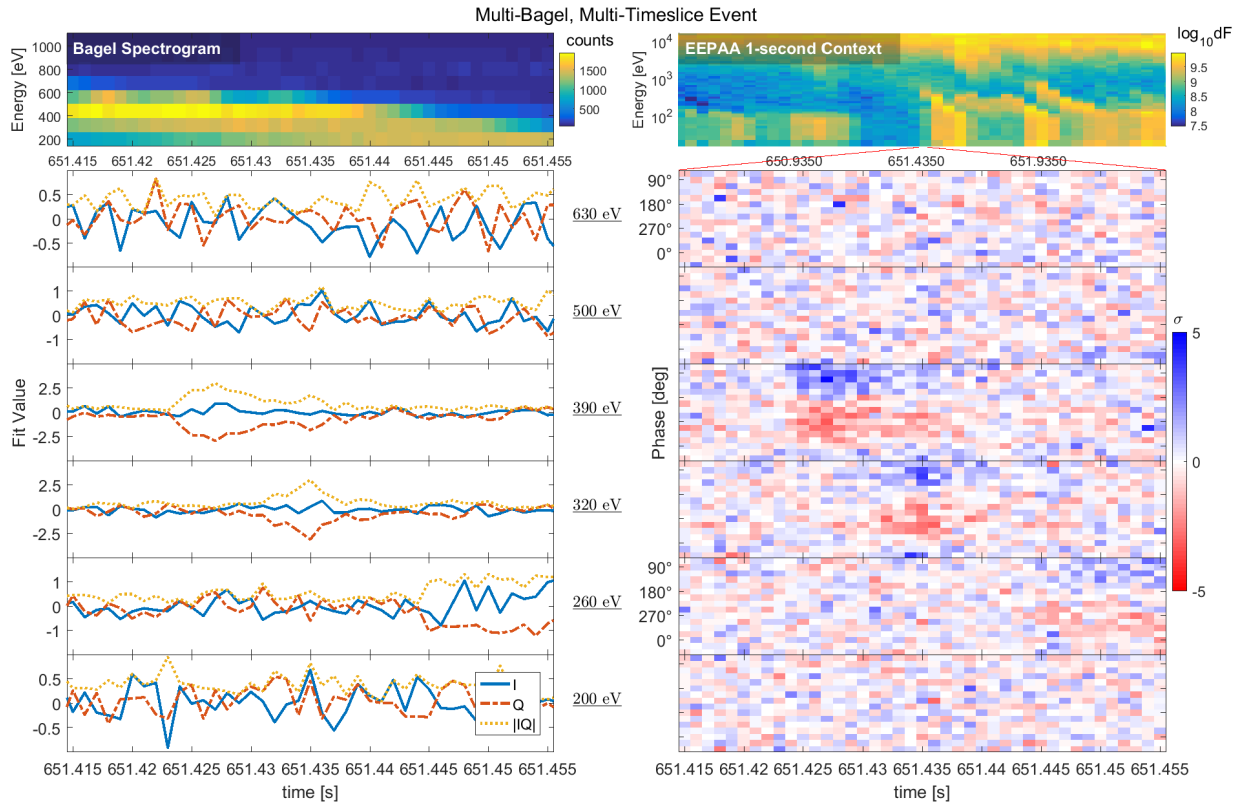

Figure 3.11: A single-event summary, showing bagel spectrogram (top left), EEPAA/top hat particle detector context (top right), and I/Q values (left column) and z-scores (right column) for each bagel. Note clear evidence of a dispersive beam passing through the 390 to 260 eV energy range, and the strong response in I and Q.

Table 3.1

STATISTICAL TESTS ON THE FIGURE 3.13 SCATTER PLOTS.

	I-∇n_B		Q-∇n_B		I-∇<w>		Q-∇<w>	
	s	p	s	p	s	p	s	p
Lin Regression	-6.97	4.0×10^{-9}	2.26	0.027	-0.36	0.72	1.27	0.21
K-S Test	1	2.8×10^{-8}	0	0.23	0	0.54	0	0.36
X-Correlation	-0.68	7.8×10^{-9}	0.24	0.062	-0.05	0.69	0.19	0.15

These statistics all evaluate relations from I and Q to ∇n_B and $\nabla < w >$. The 's' heading is general, referring to the significant output of the given test: from top to bottom, the t-statistic, null-hypothesis rejection, and correlation coefficient.

In all tests, note the extremely low p-value of the I-∇n_B relation, compared to the others.

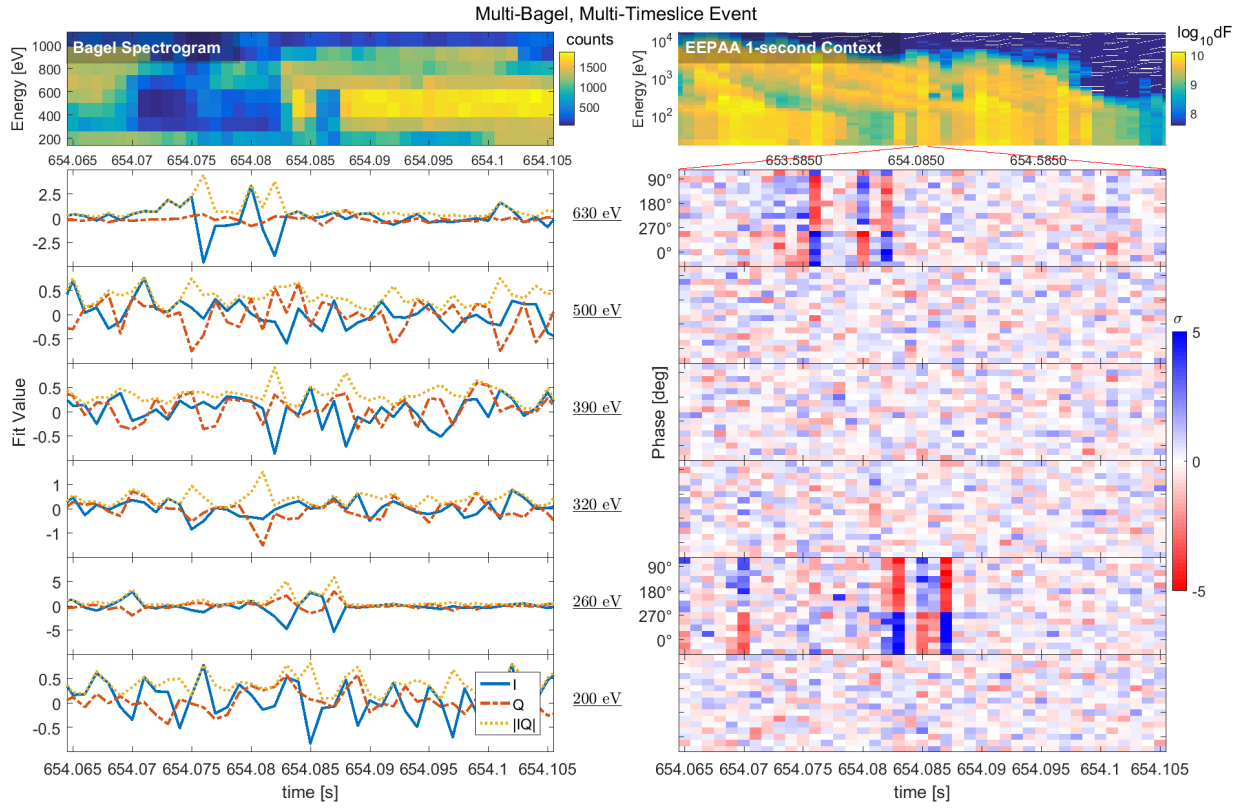

Figure 3.12: A single-event summary, bagel spectrogram (top left), EEPAA/top hat particle detector context (top right), and I/Q values (left column) and z-scores (right column) for each bagel. Note that while the EEPAA shows evidence of larger-timescale dispersive beams at higher energies, the same pattern does not seem to hold at the bagel energies.

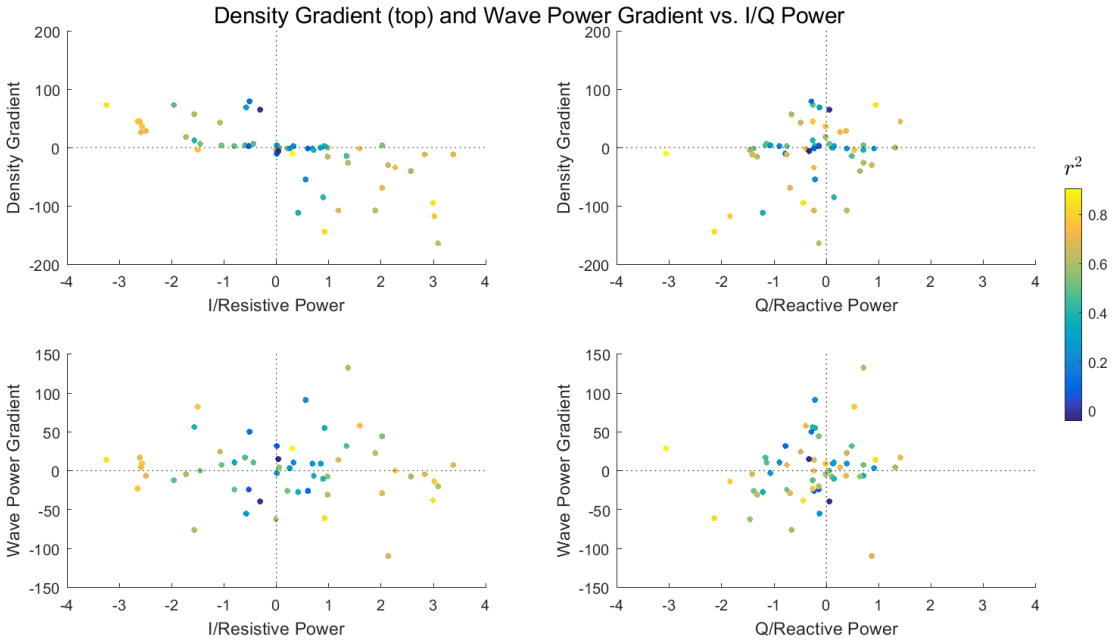

Figure 3.13: Scatter plots of bagel count gradients (top row) and HFE wave power gradients vs. I (reactive, left column) and Q (resistive) fit values, with color as the r^2 goodness-of-fit value. Note the clear relations in the I-to-count gradient (i.e. resonant electron density gradient) plot, and the lack of any relation in the wave power plots.

payload data streams. In particular, the Langmuir frequency—as judged by finding the whistler-mode cutoff on HFE spectrograms—was closely examined, but showed no obvious short-timeframe reactions in relation to I or Q changes. The I-∇n_B relation to-date remains the only relation seen.

3.4 Discussion

The above shows that the Correlator system has observed 57 potential wave-particle correlation events, after thresholding for significance and hand-screening to remove interference and potentially bad timeslices. An analysis of the reactive and resistive components of the Correlator event timeslices reveals a relation between a positive ∇n_B at a given bagel energy level, as in the case of an electron beam appearing at that energy level, and negative values of I for coincident Correlator events, and a similar relation between a negative ∇n_B and positive I. Given the calibration of the Correlators, the positive half of the electric-field waveform corresponds to a field pointing towards the bagels, and thus electrons being accelerated away from them. Thus, the observed relation is consistent with energy going from the beam to the wave field during a beam density increase, and the inverse for a density decrease. The lack of evidence for a relation between the HFE power and I and Q is curious, given prior observations by *Kletzing et al.* (2005) of such a relation on the RACE mission. It is possible

that the extreme wave power during the majority of the events in $\mathbb{S}$ may mask such an effect, particularly given that the amplitude modulation typical of bursty Langmuir waves is also not prevalent in the CHARM-II HFE data.

There are at least two relatively simple explanations for the pattern of correlations observed between in-phase wave particle correlation phase and the gradient in the beam flux. The first is motivated by the clear evidence for a dispersive beam in the event illustrated in Figure 3.11. In fact, this type of dispersive beam is the normal pattern for parallel electron beams in so-called Alfvénic aurora, in which the beams are accelerated by Alfvén waves at altitudes well above the rocket, and undergo dispersion as they propagate to lower altitudes, with fast electrons outrunning slower ones (*Kletzing and Hu* 2001; *Chen et al.* 2005). Figure 3.14 shows a schematic of the situation, as such a dispersed beam passes by a low-altitude rocket payload. The beam energy decreases with time from left to right, and as shown, the energy range of an appropriate fixed-energy particle detector will shift from lying below the peak energy of the beam to lying above the peak energy. In the former case the detected energy corresponds to the condition $df/dv_{\|} > 0$, which is destabilizing for Langmuir waves, and the latter case corresponds to $df/dv_{\|} < 0$ which is stabilizing. Under the former condition one expects waves resonant to the detector's energy to be growing, extracting energy from the beam, which would correspond to the negative values of the in-phase component of the electron-electric field correlation. Under the latter condition, the opposite energy flow would be expected, corresponding to wave damping at the detector energy. The expected signature in the phase of the electron bunching is exactly as observed.

The Figure 3.12 event is a case where the explanation given above falls short. While there is evidence in the EEPAA data of dispersion at higher energies, the alternating patterns in the σ data are difficult to explain, as is the large energy gap between the two bagels which show a signal. An additional theory does not exclude effects from the above behavior, but rather focuses on the beam appearance and disappearance at a small region of energy and pitch-angle space. The two populations involved in this theory are the 'warm' background electrons which are a degraded secondary population associated with a beam, and the beam population itself. Depicted in Figure 3.15, the bagels observe a slice of the incoming particle distribution along the parallel axis—as a beam appears, the highest-energy particles will

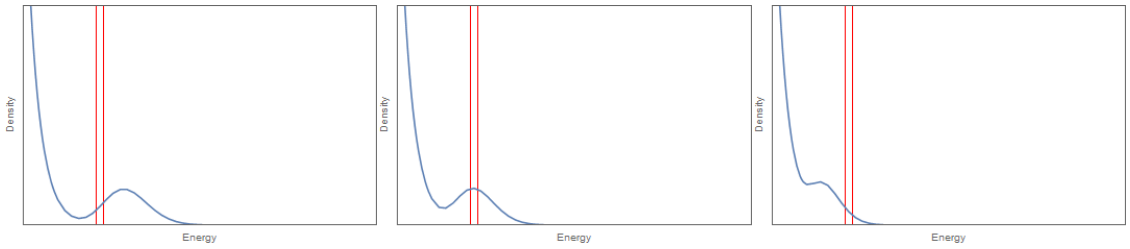

Figure 3.14: A cartoon showing the 'dispersive beam' explanation for the relation seen between I and Q and the particle count gradient. As time passes (left to right), the bagel detector's first sees the positive-slope region of the Maxwellian beam distribution, and then later a negative-slope region. Thus a beam passing downward through a bagel's energy range will see resonant wave growth, followed by enhanced damping.

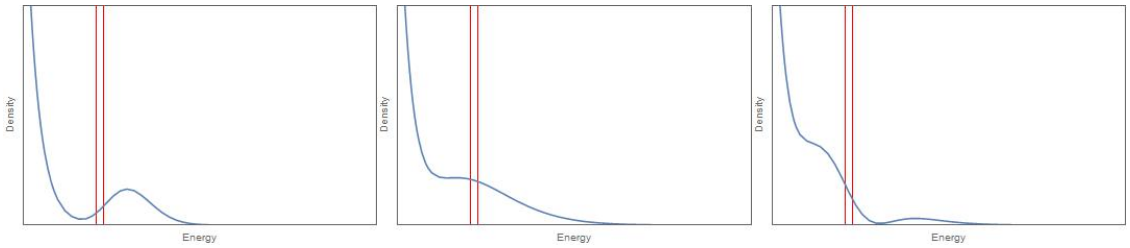

Figure 3.15: A cartoon showing a more in-depth explanation for the relation seen between I and ∇n_e. As a beam appears, the high-energy particles are the earliest to arrive, leading to an exaggerated positive slope and wave growth (left). The remainder of the particles and distribution relaxation then yield a plateau in middle times. Finally, when the beam turns off at the source, the high-energy particles are the first to disappear, and in the right configuration may yield an exaggerated negative slope, enhancing wave damping (right).

arrive first, and are likely to create a positive-slope region and wave growth. However, as the beam turns off, it is also possible, depending on the energies of the two populations, that an enhanced negative slope will appear as the higher-energy particles disappear first. This enhanced negative slope can then lead to enhanced damping in a narrow energy region. The presence of short-lived beam features in the top-left and bottom-right plots of Figure 3.10, and their temporal relation to the nearby correlator events, is compelling evidence supporting this theory.

Langmuir wave growth during an increase in the number of electrons at or near the resonant energy is generally expected because of the resultant instability, whether due to a beam moving into an energy range, or simply appearing at that energy. While subsequent damping is also expected, an impulsive enhancement of damping concurrent with the beam's disappearance, is, on the other hand, not an immediately obvious causal relationship.

Finally, in a high field, it is expected that unstable distributions will relax extremely swiftly, leaving mostly trapped, reactive particles to be observed; thus, the degree of resistive activity evident in the ratios on the right of Figure 3.9 is unexpected. These observations, and the fact that many of the events are narrow in time, many in a single timeslice, may suggest further structure at even shorter timescales.

A numerical analysis can confirm interpretations of the observed Langmuir wave growth and damping—and hence the phase of the in-phase wave-particle correlation for both positive and negative temporal gradients in the electron beam flux. It could also potentially allow probing of activity at shorter timescales. Towards this end, a method was developed to simulate the evolution of the electron distribution function, and thus the reduced distribution function and Langmuir wave growth rate, as an electron beam propagates through a vertical distance of approximately 5000 km in a converging magnetic field, with parameters close to those of the auroral ionosphere.

3.5 Simulation

The aim is to simulate a minimal-complexity environment that is sufficient to probe the questions at hand: does an electron beam with reasonable characteristics, and which shows significant Langmuir wave growth upon its appearance, also show enhanced wave damping as it fades? Can we say anything about the short-timescale behavior of the wave growth rate?

We shall define a 'reasonable beam' as originating with a Maxwellian distribution at a realistic altitude, and traveling through a magnetic field with an Earth-like mirror ratio. A similar analysis is performed by *Arnoldy et al.* (1999), and simpler methods such as one using a guiding-center approximation might suffice. However, interesting effects may be evident in this case and in broader applications, if detailed environmental attributes are taken into account, such as electric fields and agryotropic distributions. We therefore choose to develop a complex, flexible—and computationally intensive—test particle simulation system. Its development on and application to this case shall use simple, gryotropic magneto-kinetic parameters, with no inter-particle interaction or wave-particle scattering.

Following a numerical analog to the analytical method of *Cairns* (1987), we note Liouville's equation governs the evolution of a distribution function over time, and with no wave-particle or inter-particle scattering, we can simply write

$$f(\bar{x},\bar{v},t) = f(\bar{x}',\bar{v}',t'),$$

i.e. that the value of the distribution function at a source phase-space region $(\bar{x}',\bar{v}')$ at time t' is the same for the related region $(\bar{x},\bar{v})$ at time t. The test-particle simulation is used to relate the primed and unprimed regions, by creating a lookup table of particle travel times $\mathbb{T}(\mathsf{E},\alpha')$ for a range of source energies E and pitch angles α. These test particles are then treated as centers of regions in phase space, and are used to 'carry', in time $\mathbb{T}$, values of the source distribution function down to a corresponding region (E,α) at the observation point.

In this analysis, z is taken to be positive, downward, field-aligned coordinate, with $z = 0$ corresponding to the beam generation altitude. In order to only simulate particles which will arrive at our 'detection point' at $(x = y = 0, z = -5000$ km) we use a deterministic (i.e. time-reversible) simulation method, and originate our test particles at the detection point with an upward velocity, watching for them to cross a target plane at $z = 0$. The velocities can then be reversed for the later downgoing analysis. The 'Boris Method' is used—a standard, time-reversible particle pusher (*Boris* 1970; *Birdsall and Langdon* 2005). This method separates the effects of the electric and magnetic forces, dividing them into a half-impulse from any background electric field, followed by a rotation according to the magnetic field, and then another electric half-impulse.

Careful testing of energy conservation led to setting a unitless timestep of 0.01. The base of the time system is the electron cyclotron frequency, and so this is equivalent to each timestep moving each particle a hundredth of an orbit. For the input parameters used, this yielded a worst-case energy loss of 0.06% over the full length of the simulation.

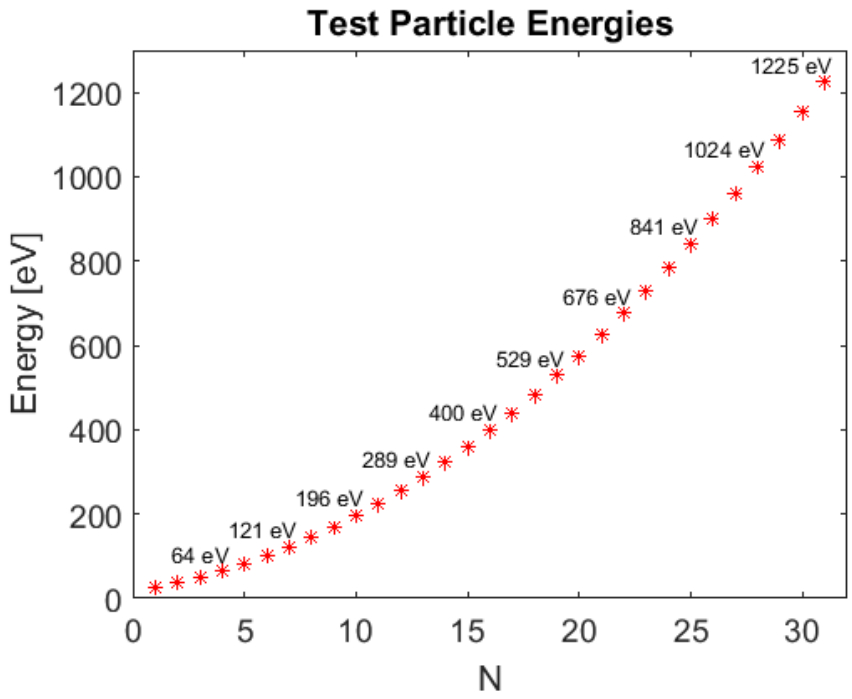

Figure 3.16: The 31 test-particle simulation launch-energy levels. Linearly spaced in velocity.

To allow a realistic amount of time/space for mixing of particles of different energies and pitch angles, a distance of 5000 km is used, corresponding to the distance from the bottom of the electron acceleration region to the ionospheric detection point. The background electric field is assumed zero, and the magnetic field is rotationally symmetric around $x = y = 0$, defined as

$$\bar{B} = -\frac{zr}{L^2}\hat{r} + \frac{1+z^2}{L^2}\hat{z} = -\frac{xz}{L^2}\hat{x} - \frac{yz}{L^2}\hat{y} + \frac{1+z^2}{L^2}\hat{z},$$

with L a scaling variable determined by our desired mirror ratio and target distance. For the following simulation, the mirror ratio was set to 5.

In order to fully cover the range of energies detected by the Correlator, particles were launched at the 31 energies plotted in Figure 3.16, linearly spaced in velocity, with energies ranging from 25 to 1225 eV. At each energy, particles were launched in a range of 41 pitch angles, with an angular spacing of $3\pi/256$. Because of the rotational symmetry of the simulation imposed by gyromotion, it is only necessary to launch particles at one azimuthal angle—the results can then be rotated to fill a velocity-space hemisphere at the detection point. As we would like our hemisphere segments to sweep out constant solid angles, we require that, for pitch angle α, the particle rotates to the integer number of azimuthal angles ϕ which most closely yield $\Delta\phi = \frac{\Delta\Omega}{\sin(\alpha)\Delta\alpha}$, where $\Delta\Omega$ is the solid angle, herein set to 0.001 steradians.

The simulation code was implemented in MATLAB, manually fragmented into 64 shards, and run in approximately five weeks on the Dartmouth Discovery cluster. Due to the finite

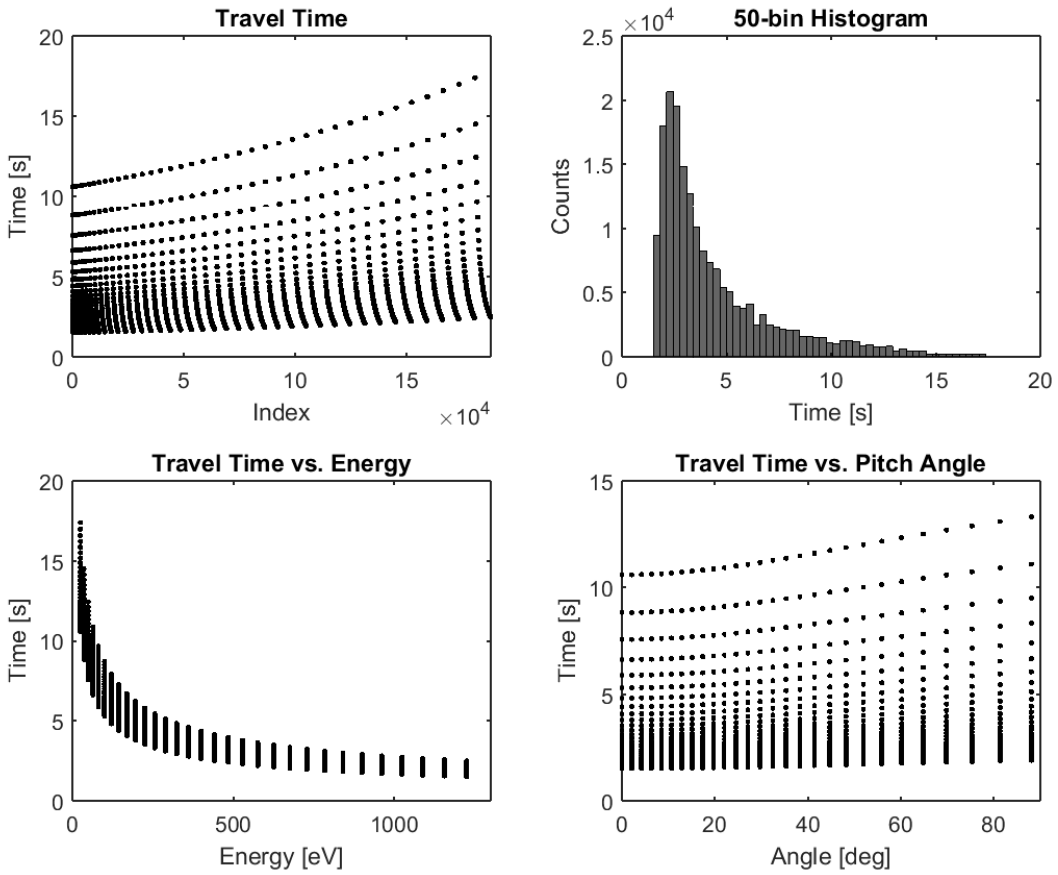

Figure 3.17: Various diagnostics of the resultant travel times for the simulated test-particles. Note that the energy and pitch angle in the lower row are the launch values at the detection point (high magnetic field).

timesteps, which are unlikely to land precisely on $z = 0$, interpolation was required to find exact crossing parameters. To enable this, the final 1000 timesteps for each particle were saved (with the very last step having $z > 0$), and gyro-orbit equations were fit to these, from which accurate final $z = 0$ position, velocity, and travel times come forthwith. All unitless values were then interpreted via inter-defined base values: $B_0 = 50$ microtesla, $t_0 = \sim 714$ ns, $r_0 = \sim 0.337$ m, and $v_0 = 0.00989$ c, corresponding to 25 eV.

Figure 3.17 shows some basic diagnostics of the output of this simulation, with expected trends compared to total energy and launch pitch angle. With a table $\mathbb{T}(\mathsf{E}, \alpha)$, we proceed to reverse the velocities and launch distributions downward, towards our 'detector'.

The distribution is imposed at the top (lowest B) on total velocity, $|v|$ (i.e. we assume $T_{\parallel} = T_{\perp}$ and a flat distribution across pitch-angle space), and sampled at the test velocities used in $\mathbb{S}$. The environment here is assumed to be both homogeneous and large enough that any generated region of velocity-space at the top will detected at the bottom. Thus, we can neglect the x and y positions of particles in the simulation.

We seek to simulate an instrument observing the electrons being emitted effectively continuously from a source region, which contains a stable background distribution, and has a strong beam appear, stabilize, and then disappear. Thus, we set a Δt_D period over which the detector bins incoming particles, and a Δt_S period between source distribution 'launches'. To achieve something approaching the appearance of a continuous source, Δt_D should be at least $10\Delta t_S$.

A secondary background distribution serves as the main background to our beam for growth-rate calculation purposes. It is a localized population, distinct from the much colder, higher-density ionospheric background, and is composed of a population of degraded beam particles with a higher density and temperature than the beam population. We define it as a constant Maxwellian,

$$f(|v|) = n_e\,(2\pi)^{3/2}\,v_{th}^3\;\mathrm{e}^{-\left(\frac{v}{2v_{th}}\right)^2} = n_{bg}\left(\frac{m_e}{2\pi k T_{bg}}\right)^{3/2}\mathrm{e}^{-\frac{m_e(v)^2}{2kT_{bg}}} = f_{bg},$$

where, for a given input temperature, the electron number density n_{bg} is interpolated from a table of values used by *Lotko and Maggs* (1981).

The beam is built with a similar form, except for time-varying parameters $T_{beam}(t)$, $n_{beam}(t)$, and a velocity shift $\delta(t)$:

$$f_{beam}(|v|,t) = n_{beam}(t)\left(\frac{m_e}{2\pi k T_{beam}(t)}\right)^{3/2}\mathrm{e}^{-\frac{m_e(v-\delta(t))^2}{2kT_{beam}(t)}},$$

where in practice $T_{beam}(t)$ and $n_{beam}(t)$ are set as fractions of the secondary background values.

The final distribution is then the sum of these, as in Figure 3.18,

$$f(v,t) = f(v_\parallel, v_\perp, \theta, t) = f_{iono} + f_{bg} + f_{n_{beam}}(t)$$

where the contribution from the ionospheric background f_{iono} with a 2000 K temperature is included for completeness, but its effect on the distribution function is vanishingly small at these high energies.

To dimensionally reduce these towards a parallel distribution function $f(v_\parallel)$, we first sum over the azimuthal angles. This is not a simple sum: as these are finite cells in velocity-space, we must weight each angular 'wedge' by its accompanying $\Delta\theta_i$, i.e.

$$f(v_\parallel, v_\perp, t) = \sum_{i\forall\theta} f(v_\parallel, v_\perp, \theta_i, t)\Delta\theta_i,$$

where $\Delta\theta_i$ is set by the pitch angle, as in the hemispheric interpolation.

We put off dealing with time until now because the hemispheric interpolation introduces no time dependence, and so the azimuthal sum has none either. The next step will be interpolating and reducing away a dimension from our test-particle simulation, so we must

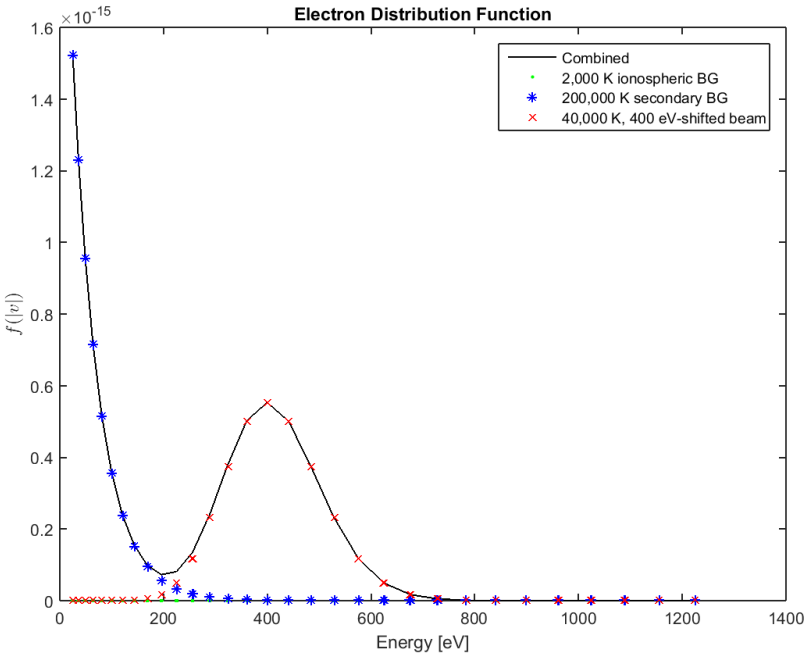

Figure 3.18: Imposed top distribution in $|v|$, showing ionospheric background distribution (green dots), secondary background (blue stars), a beam distribution (red Xs), and final combined sum (solid black line), with $T_{iono} = 2e3$ K, $n_e \sim 1.984e9$ m^{-3}, $T_{bg} = 2e5$ K, $n_{bg} = 1.066e6$ m^{-3}, $T_{beam} = 4e4$ K, $n_{beam} = 21.32e3$ m^{-3}, and $\delta = 400eV$.

take travel times into consideration beforehand. This is simply a set limitation at detector timeslice τ, such that the particles we consider are, henceforth, in the set $\mathbb{J}_\tau$ of particles whose launch time t_0 and travel time t_T fulfill $t_0 + t_T <= \tau$ and $> \tau - t_D$. There is also an implicit sum here as we are taking the detector integration into account, and this needs its own weighting value for f, simply the ratio of the launch time and the integration time, i.e.

$$f(v_\parallel, v_\perp, \tau) = \frac{\Delta t_S}{\Delta t_D} f(\mathbb{J}_\tau).$$

Next, we sum over perpendicular velocities to get a one-dimensional reduced distribution function, taken from the standard Landau theory for parallel propagation (e.g. *Ergun et al.* (1993)). This is slightly complicated by the large number of unique $v_\parallel$ values. We define a set of 'center points' in $v_\parallel$ as simply the points along the $v_\perp = 0$ axis. For each of these $v_{\parallel,i}$, all values of f with $v_\parallel$ in the range $\mu_\parallel = \left\{ \frac{v_{\parallel,i-1}+v_{\parallel,i}}{2}, \frac{v_{\parallel,i}+v_{\parallel,i+1}}{2} \right\}$ are summed over $v_\perp$ using a modification of the trapezoidal rule,

$$f(v_{\parallel,i}, \tau) = \sum_{j \forall v_\perp; \mu_\parallel} (v_{\perp,j+1} - v_{\perp,j}) \frac{f(\mu_\parallel, v_{\perp,j}, \tau) + f(\mu_\parallel, v_{\perp,j+1}, \tau)}{2} v_{\perp,j},$$

where the factor $v_{\perp,j}$ is the phase-space cell weighting. Figure 3.19 shows a color plot of the reduced distribution function vs. $v_\parallel$ and time, as well as several timeslices as the distribution evolves through the beam-arrival phase.

Now, with a time-integrated one-dimensional reduced distribution function in $v_\parallel$, we can calculate growth rates. For a given cold ionospheric background plasma frequency ω_p, wave vector $k = k_\parallel$, and test plasma frequency ω_t, the growth rate is

$$\gamma(f(v_\parallel), k, \omega_p, \omega_t, \tau) = \left(\frac{d\varepsilon}{d\omega} \right)^{-1} \text{Sign}[k] \, \frac{\pi \omega_p^2}{k^2 n_e} \left[\frac{\partial f(v_\parallel, \tau)}{\partial v_\parallel} \right]_{k v_\parallel = \omega_t},$$

where ε is the dielectric function, approximated as $1 - \frac{\omega_p^2}{\omega_t^2}$ for cold plasma. The derivative $\partial f_1 / \partial v_\parallel$ is calculated at a test velocity related to the beam parameters, specifically the closest $v_\parallel$ value to $\delta + T_{beam}/k_B$.

Ideally we would want to set Δt_D to match the Correlator's timeslices, 1 ms, and to allow the simulation to 'settle' for a long enough time between source changes that even the slowest particles reach the detector, approximately 14 s per change. However, given our above guideline that $\Delta t_S \leq \Delta t_D / 10$, this would require storage of prohibitive numbers of time-overlapping distributions. From Figure 3.17 we know that the majority of particles will have arrived within 5 seconds, so we use that as our settling time, and only use realistic Δt_D in special cases. Finally, we relate k values and ω_t values via an approximation of the warm plasma dispersion relation, $\omega = \omega_p + \frac{3}{2} k^2 v_{th}^2 / \omega_p$, where v_{th} is the background ionospheric thermal velocity $\sqrt{3 k T_{iono} / m_e}$.

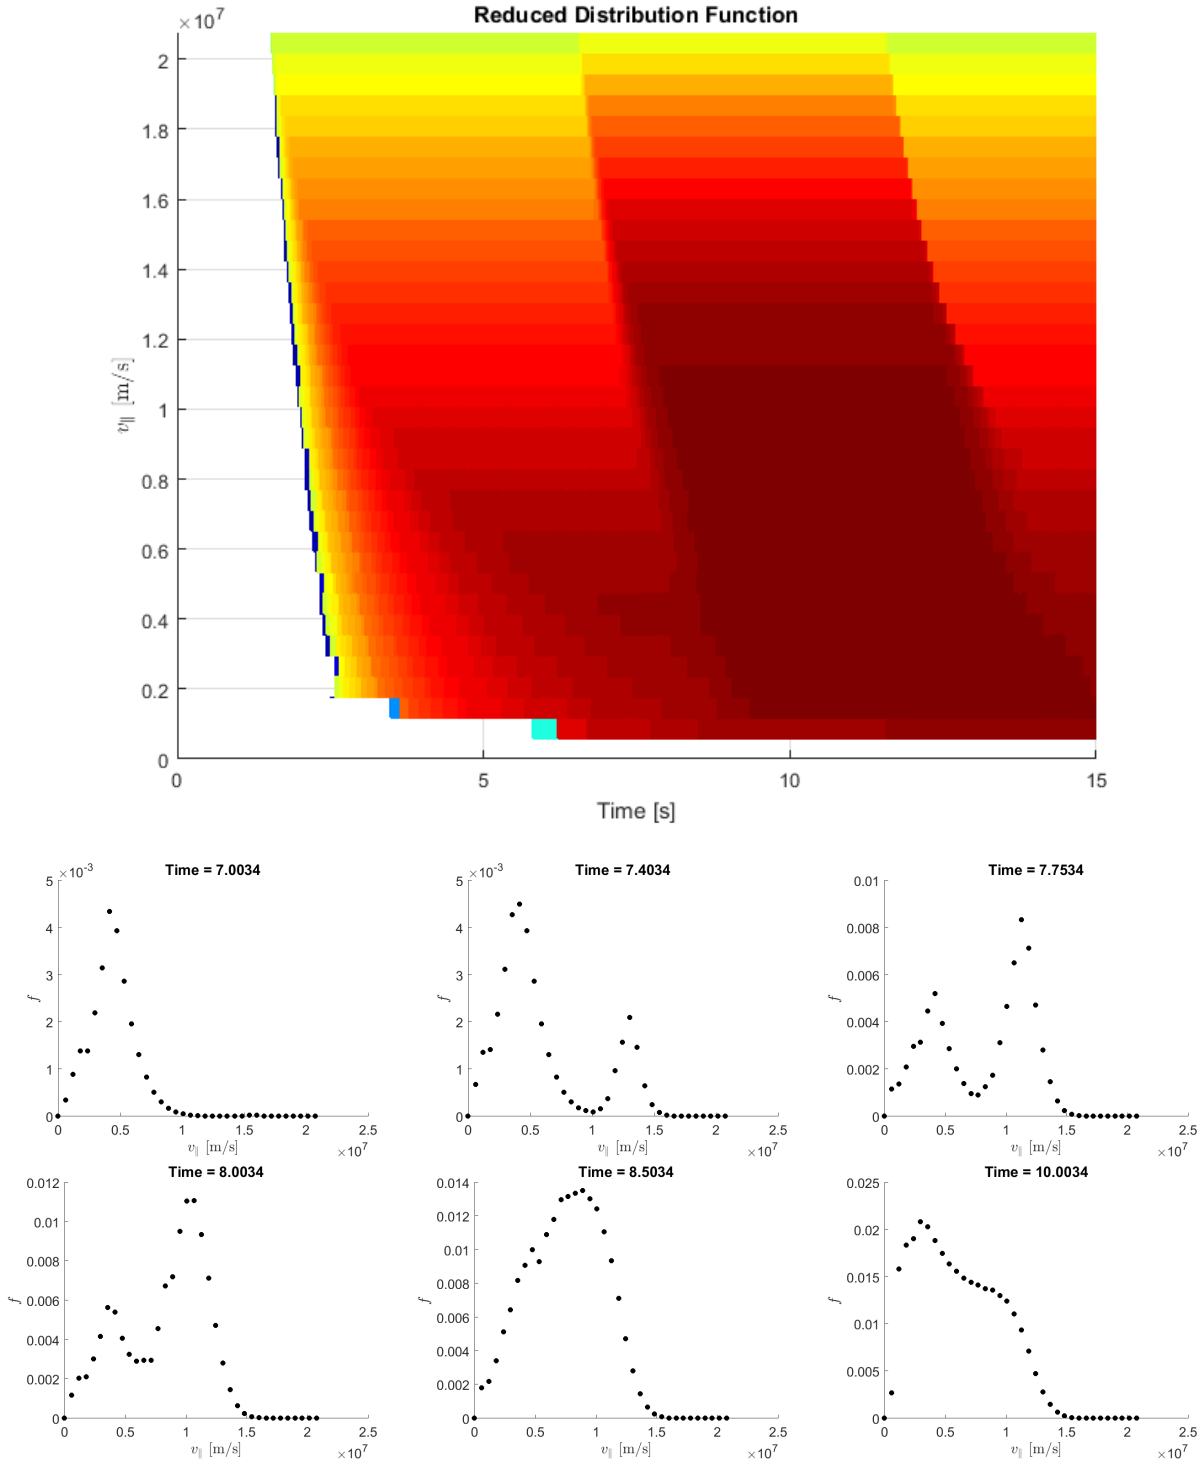

Figure 3.19: Reduced distribution function values. Top shows a color plot vs. $v_{\parallel}$ and time, for the entire time span of the test. Below, six timeslices from the beam arrival period, showing the formation and disappearance of a positive slope. Note that the continued appearance of low-energy particles near the bottom is the tail of the background distribution. It has not yet arrived before beam turn-on, but is considered too small to significantly affect the results.

Figure 3.20 shows the final results of the simulation: Langmuir wave growth rate γ, versus k and time (on the horizontal axis), calculated for ω_t from approximately $1.00022\omega_p$ to $1.0054\omega_p$. The top panel shows the n_{beam}/n_{bg} at the top, while the two columns are zoomed into the times at which the bulk of the particles arrive during beam turn-on (left), or depart with beam shutoff (right).

The top growth-rate panels show γ on the vertical axis, as well as in color scale (blue is negative, red positive). Both a growth rate spike during the beam arrival and a damping enhancement during beam departure are clearly visible. This result matches qualitatively the pattern observed in the phase of the in-phase, resistive component (I) of the wave-particle correlations during positive and negative gradients in beam flux.

In the lower panels, the color scale is still γ, and the vertical axis is the wavenumber k. The strongest growth and damping are associated with the long-wavelength modes at $k < 0.004$, which is generally expected as shorter-wavelength modes are more heavily damped. Growth at the long wavelengths is associated with the earliest-arriving, higher-energy particles, with the later lower-energy arrivals exciting some growth at shorter wavelengths.

The overall timeframe of the growth and damping peaks are of order 100 ms, which is significantly longer than most of the observed wave-particle correlation events; however, there is suggestion of shorter-scale time structure in the simulation events, particularly of a double-peak in the growth rate. To further probe the small-scale structure of these results, we can modify certain parameters of the distribution-building. Figure 3.21 shows the results of three such tests: moving the beam by half the inter-energy spacing, both up and down in energy, as well as removing half of the energies entirely. Motivated by the fact that realistic electron beams have lifetimes orders of magnitude smaller than the 5 s beams used in Figure 3.20, Figure 3.22 shows the result of a beam with identical parameters, but a lifetime of only 100 ms. The general impression of the above tests is that all of the small-scale structure seen in Figure 3.20 is heavily dependent on various aspects and limitations of the simulation system itself. Thus, these structures may bear little to no resemblance to any physical reality—even an approximation of such. Given this, as the simulation stands, no quantitative conclusions can follow regarding the small-timescale behavior of the growth rate.

3.6 Conclusions

The CHARM-II sounding rocket successfully carried a wave particle correlator to an apogee altitude of 802 km in substorm aurora. The correlator instrument locks onto the highest-amplitude waveform in the 100 kHz to 4 MHz range, ideally the wave at the Langmuir frequency, and bins incoming electrons at 8 energy levels into 16 phase bins. It returned data from > 400 seconds of flight time, and after both automated and manual event selection, 57 timeslices containing events of interest were selected and analyzed. Breakdown of the phase correlation data into resistive and reactive components revealed a striking relationship between electron beam dynamics and the nature of the wave-particle correlation: whenever the beam flux at the measured electron energy was increasing with time, the phase of the

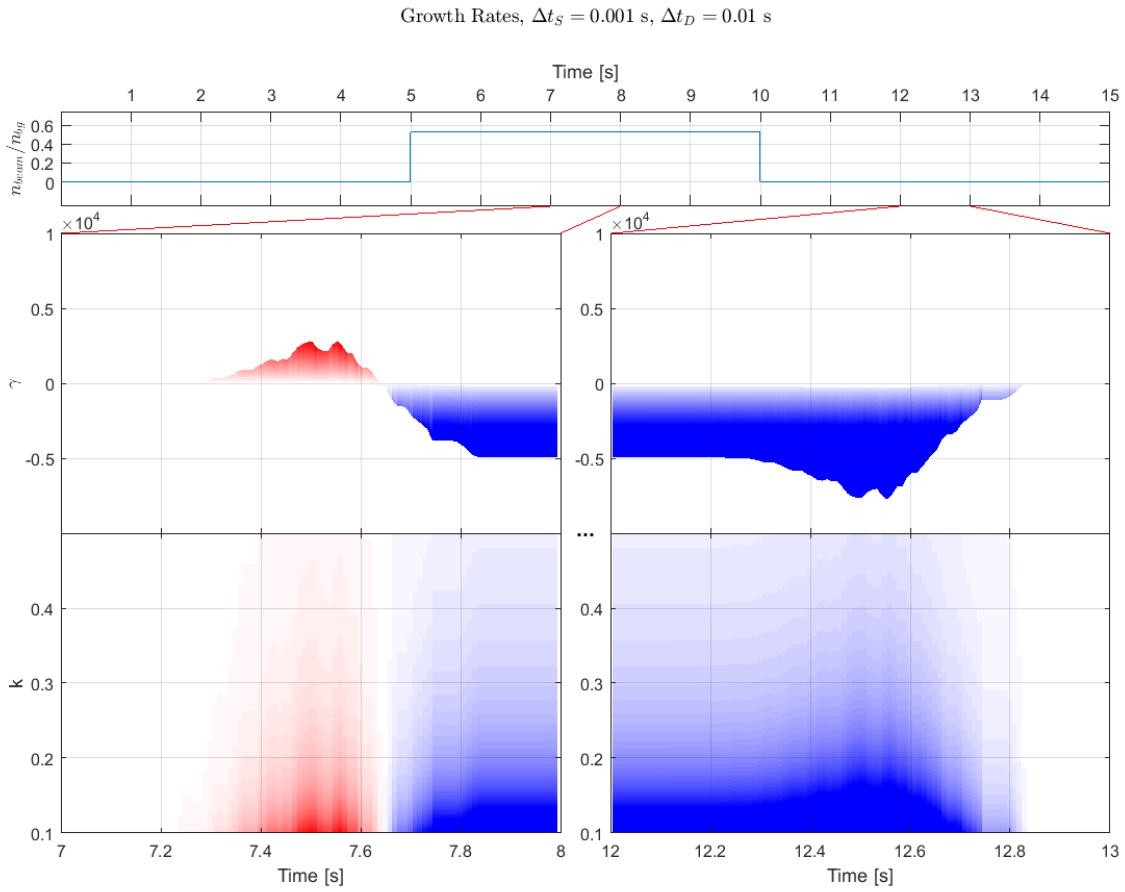

Figure 3.20: Results from the simulation: Langmuir wave growth rate γ, versus k and time (on the horizontal axis), calculated for ω_t from approximately $1.00022\omega_p$ to $1.0054\omega_p$. The top panel shows the n_{beam}/n_{bg} for at the top, while the two columns are zoomed into the times at which the bulk of the particles arrive during beam turn-on (left), or depart with beam shutoff (right). The top growth-rate panels show γ on the vertical axis, as well as in color scale (blue is negative, red positive), and both a growth rate spike during the beam arrival and a damping enhancement during beam departure are clearly visible, qualitatively matching the the pattern observed in the data. In the lower panels, the color scale is still γ, and the vertical axis is the wavenumber k.

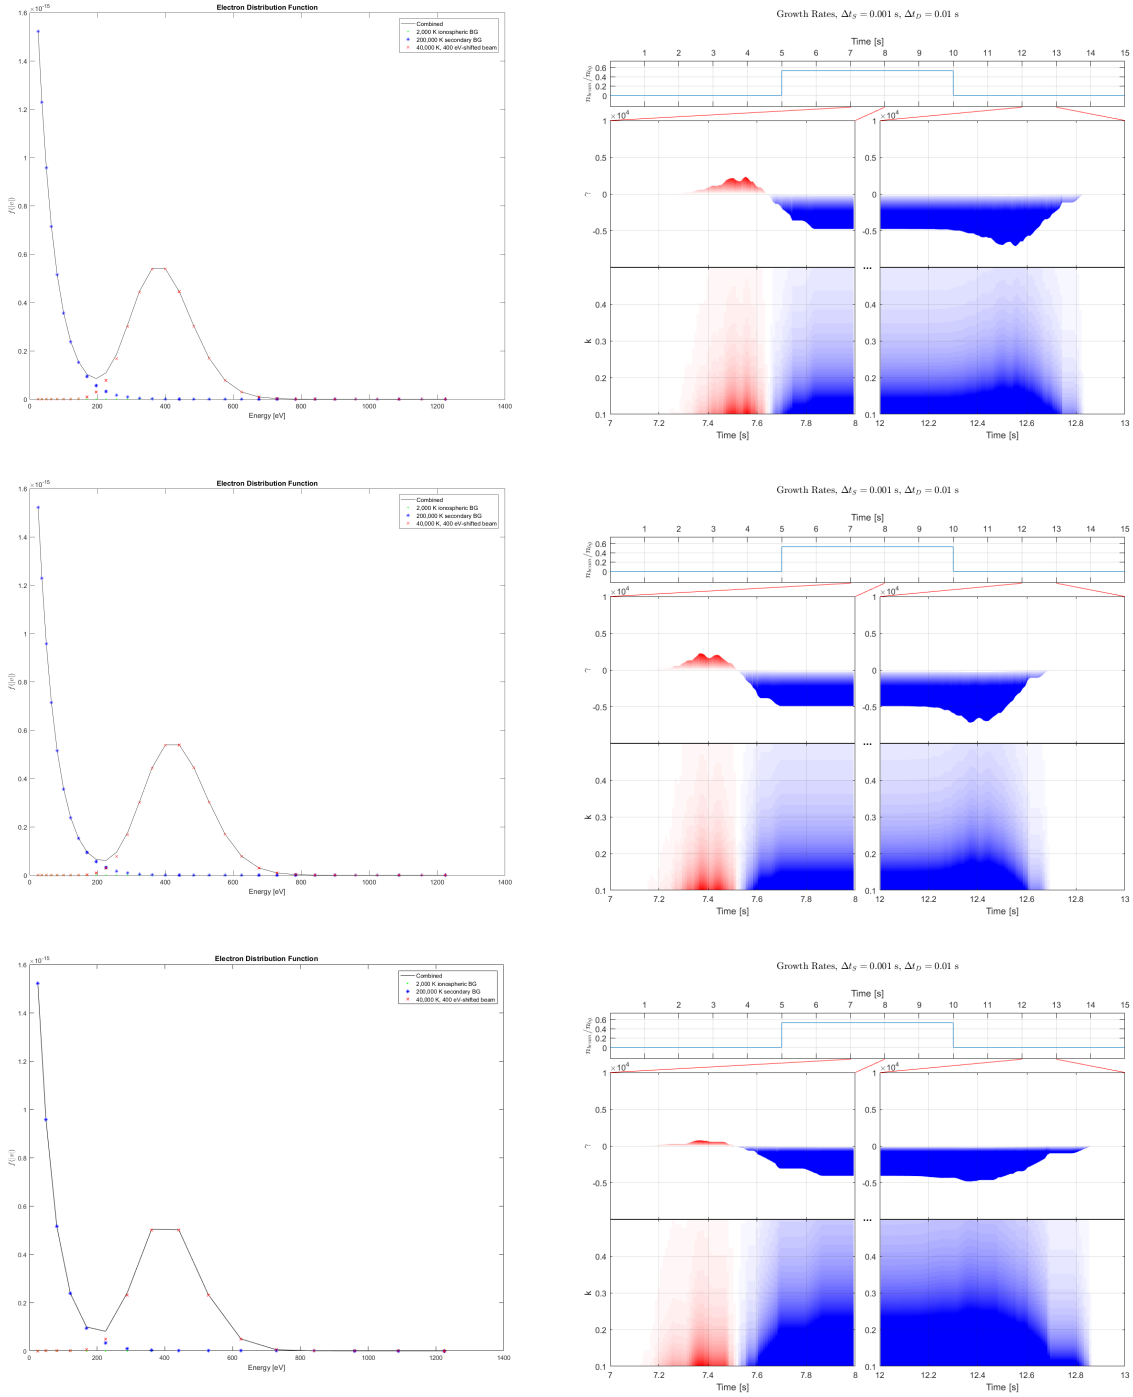

Figure 3.21: Test inputs and resultant growth rates, showing that much of the small-timescale structure seen in the growth rate is due to binning effects.

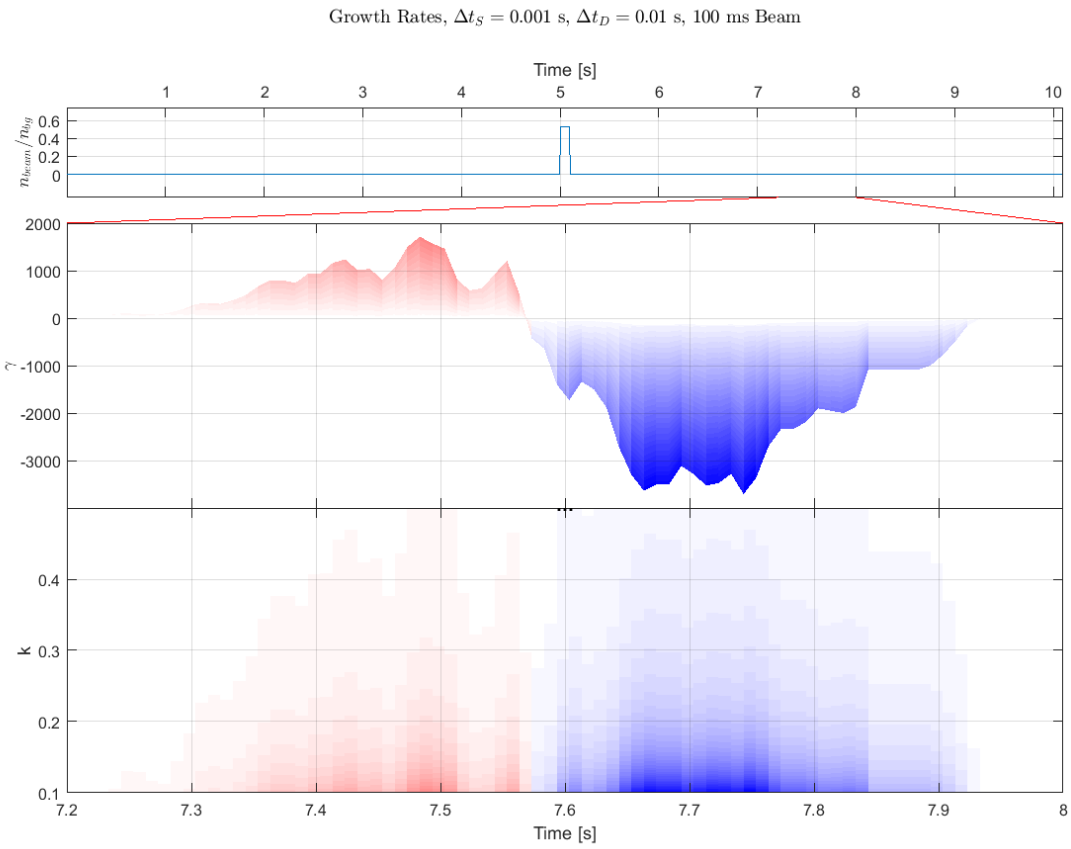

Figure 3.22: Growth rates resulting from a beam identical to that in Figure 3.20, but lasting only 1/50th of the time.

resistive component of the electron bunching implied energy transfer from the wave field to the particles, and when the electron beam flux was decreasing, the reverse occurred. This pattern was repeated for all events, and was particularly clear in several events, including the largest-amplitude event investigated by *Kletzing et al.* (2011).

Two related theories to explain this observation have been explored, one invoking the changing nature of the interactions of the electrons with a given Langmuir wave as the beam energy decreases, as typically occurs due to dispersion of an auroral electron beam accelerated several thousand kilometers above the interaction location; and the other invoking detailed features of the electron distribution function at ionospheric altitudes, arising when the electron beam is modulated at higher altitudes. A magneto-kinetic test-particle numerical simulation confirmed that for an electron beam which causes an impulsive increase in wave growth upon its appearance, its disappearance will be accompanied by an impulsive enhancement of wave damping within the same frequency range. The results therefore agree qualitatively with the experimental data from the CHARM-II rocket, though an exactly simulated quantitative representation has not yet been achieved.

The numerical simulation system developed consists of a flexible, cluster-enabled, node-independent suite of MATLAB scripts using the Boris Particle Pushing algorithm, modular background and beam distribution functions with arbitrary time-changing beam profiles, and adaptive distribution function dimensionality reducer algorithms. This system may be easily adapted to a multitude of tasks, including significant extensions of the work presented here, such as including non-uniform beam profiles and considering a broader class of Langmuir waves ($k_\perp \neq 0$) in the growth rate calculations. These and other extensions of the modeling may provide more quantitative tests of the time structure or magnitude of correlations observed in the CHARM-II wave particle correlator data.

Thanks to Dr. Scott Bounds and Dr. Stephen Kaeppler for assistance with the Correlator data and calibrations, and to Dr. Wayne Scales for providing the original magneto-kinetic test-particle code on which this simulation system was based.

Bibliography

Arnoldy, R. L., K. A. Lynch, and J. B. Austin, Energy and pitch angle-dispersed auroral electrons suggesting a time-variable, inverted-v potential structure, *J. Geophys. Res.*, *104*(A10), 22,613–22,621, 1999.

Birdsall, C. K., and A. B. Langdon, *Plasma Physics via Computer Simulation*, Institute of Physics Series in Plasma Physics, Taylor & Francis Group, New York, 2005.

Boehm, M., G. Paschmann, J. Clemmons, H. Höfner, R. Frenzel, M. Ertl, G. Haerendel, P. Hill, H. Lauche, L. Eliasson, and R. Lundin, The tesp electron spectrometer and correlator (f7) on freja, *Space Sci. Rev.*, *70*, 509–540, 1994.

Boehm, M. H., Waves and static electric fields in the auroral acceleration region, Ph.D. thesis, University of California, Berkeley, Berkeley, CA, 1987.

Boris, J. P., Relativistic plasma simulation-optimization of a hybrid code, in *Proceedings of the Conference on the Numerical Simulation of Plasmas (4th)*, edited by J. P. Boris and R. A. Shanny, 1, pp. 3–67, Plasma Physics Division, Naval Research Laboratory, National Technical Information Service, Washington, DC, 1970.

Cairns, I. H., The electron distribution function upstream of earth's bow shock, *J. Geophys. Res.*, *92*(A3), 2315–2327, 1987.

Chen, L. J., C. A. Kletzing, S. Hu, and S. R. Bounds, Auroral electron dispersion below inverted-v energies: Resonant deceleration and acceleration by alfvén waves, *J. Geophys. Res.*, *110*(A10), A10S13, doi:10.1029/2005JA011168, 2005.

Dombrowski, M. P., J. LaBelle, D. E. Rowland, R. F. Pfaff, and C. A. Kletzing, Interpretation of vector electric field measurements of bursty Langmuir waves in the cusp, *Journal of Geophysical Research (Space Physics)*, *117*, A09209, doi:10.1029/2012JA017741, 2012.

Ergun, R. E., C. W. Carlson, J. P. McFadden, and J. H. Clemmons, Langmuir wave growth and electron bunching: Results from a wave-particle correlator, *J. Geophys. Res.*, *96*(A1), 225–238, 1991a.

Ergun, R. E., C. W. Carlson, J. P. McFadden, D. M. TonThat, and J. H. Clemmons, Observation of electron bunching during landau growth and damping, *J. Geophys. Res.*, *96*(A7), 11,371–11,378, 1991b.

Ergun, R. E., G. T. Delory, E. Klementis, C. W. Carlson, J. P. McFadden, I. Roth, and

M. Temerin, Vlf wave growth from dispersive bursts of field-aligned electron fluxes, *J. Geophys. Res.*, *98*(A3), 3777–3787, 1993.

Ergun, R. E., J. P. McFadden, and C. W. Carlson, Wave-particle correlator instrument design, in *Measurement Techniques in Space Plasmas: Particles*, *Geophysical Monograph*, vol. Particles, pp. 325–331, American Geophysical Union, American Geophysical Union, 1998.

Ergun, R. E., D. M. Malaspina, I. H. Cairns, M. V. Goldman, D. L. Newman, P. A. Robinson, S. Eriksson, J. L. Bougeret, C. Briand, S. D. Bale, C. A. Cattell, P. J. Kellogg, and M. L. Kaiser, Eigeinmode structure in solar-wind langmuir waves, *Phys. Rev. Lett.*, *101*(051101), 2008.

Filbert, P. C., and P. J. Kellogg, Electrostatic noise at the plasma frequency, *J. Geophys. Res.*, *84*(A4), 1369–1381, 1979.

Gough, M. P., P. J. Christiansen, and K. Wilhelm, Auroral beam-plasma interactions: Particle correlator investigations, *J. Geophys. Res.*, *95*(A8), 12,287–12,294, 1990.

Gurnett, D. A., F. L. Scarf, W. S. Kurth, R. R. Shaw, and R. L. Poynter, Determination of jupiter's electron density profile from plasma wave observations, *J. Geophys. Res.*, *86*(A10), 8199–8212, 1981.

Kaeppler, S., C. A. Kletzing, S. R. Bounds, J. W. Gjerloev, B. J. Anderson, H. Korth, J. W. LaBelle, M. P. Dombrowski, M. Lessard, R. F. Pfaff, D. E. Rowland, S. Jones, and C. J. Heinselman, Current closure in the auroral ionosphere: Results from the auroral current and electrodynamics structure rocket mission, submitted Nov 2011, 2011.

Kintner, P. M., J. Bonnell, S. Powell, and J. E. Wahlund, First results from the freja hf snapshot receiver, *Geophys. Res. Lett.*, *22*, 287, 1995.

Kletzing, C. A., and S. Hu, Alfvén wave generated electron time dispersion, *Geophys. Res. Lett.*, *28*(4), 693–696, doi:10.1029/2000GL012179, 2001.

Kletzing, C. A., and L. Muschietti, *Phase Correlation of Electrons and Langmuir Waves*, *Lecture Notes in Physics*, vol. 687, chap. 13, pp. 313–337, Springer Berlin Heidelberg, 2006.

Kletzing, C. A., S. R. Bounds, J. LaBelle, and M. Samara, Observation of the reactive component of langmuir wave phase-bunched electrons, *Geophys. Res. Lett.*, *32*(L05106), doi:10.1029/2004GL021175, 2005.

Kletzing, C. A., S. R. Kaeppler, S. R. Bounds, J. LaBelle, and M. P. Dombrowski, Nonlinear wave evolution: Observation of electron phase bunching in auroral langmuir waves, Abstract NG23A-1478 presented at 2011 Fall Meeting, AGU, San Francisco, Calif., 5-9 Dec., 2011.

Langmuir, I., Oscillations in ionized gasses, *Proceedings of the National Academy of Sciences*, *14*(8), 627–637, 1928.

Lin, R. P., D. W. Potter, D. A. Gurnett, and F. L. Scarf, Energetic electrons and plasma waves associated with a solar type iii radio burst, *Astrophys. J.*, *251*(1), 364, 1981.

Lotko, W., and J. E. Maggs, Amplification of electrostatic noise in cyclotron resonance with an adiabatic auroral beam, *J. Geophys. Res.*, *86*(A5), 3449–3458, 1981.

Malaspina, D. M., and R. E. Ergun, Observations of three-dimensional langmuir wave structure, *J. Geophys. Res.*, *113*(A12108), 2008.

Malaspina, D. M., I. H. Cairns, and R. E. Ergun, Antenna radiation near the local plasma frequency by langmuir wave eigenmodes, *Astrophys. J.*, *755*(45), 2012.

McAdams, K. L., Sounding rocket based investigations of hf waves in the auroral ionosphere, Ph.D. thesis, Dartmouth College, 6127 Wilder Laboratory, Hanover, NH 03755, 1999.

McAdams, K. L., J. LaBelle, P. W. Schuck, and P. M. Kintner, Phaze ii observations of lower hybrid burst structures occurring on density gradients, *Geophys. Res. Lett.*, *25*(16), 3091–3094, 1998.

McAdams, K. L., R. E. Ergun, and J. LaBelle, Hf chirps: Eigenmode trapping in density depletions, *Geophys. Res. Lett.*, *27*(3), 321–324, 2000.

Muschietti, L., I. Roth, and R. E. Ergun, Interaction of langmuir wave packets with streaming electrons: Phase-correlation aspects, *Phys. Plasmas*, *1*, 1008, 1994.

Samara, M., Sounding rocket based investigations of whister, upper hybrid and langmuir waves in the auroral ionosphere, Ph.D. thesis, Dartmouth College, 6127 Wilder Laboratory, Hanover, NH 03755, 2005.

Samara, M., and J. LaBelle, Lf/mf whistler mode dispersive signals observed with rocket-borne instruments in the auroral downward current region, *J. Geophys. Res.*, *111*(A09305), 2006.

Samara, M., J. LaBelle, C. A. Kletzing, and S. R. Bounds, Rocket observations of structured upper hybrid waves at fuh = 2fce, *Geophys. Res. Lett.*, *31*(22), 2004.