manuscript accepted 2 June 2019 to *JGR: Space Physics* doi:10.1029/2018JA026262

Statistical study of electron bunching in auroral Langmuir waves

Micah P. Dombrowski[1], Jim LaBelle[1], Craig A. Kletzing[2], Scott R. Bounds[2], Iver H. Cairns[3], Stephen R. Kaeppler[4]

[1]Department of Physics and Astronomy, Dartmouth College, Hanover, NH, USA.
[2]Department of Physics and Astronomy, University of Iowa, Iowa City, IA, USA.
[3]School of Physics, University of Sydney, Sydney, Australia.
[4]Department of Physics and Astronomy, Clemson University, Clemson, SC, USA.

Key Points:

- The CHARM-II sounding rocket observed 57 correlations between electrons and Langmuir wave phases.
- Correlations reveal a specific relationship between electron beam flux gradient and energy flow.
- Simulation shows time-of-flight effects on the particle distribution are a plausible explanation.

Corresponding author: Micah P. Dombrowski, `micah.p.dombrowski@dartmouth.edu`

Abstract

Langmuir waves are a fundamental consequence of electron beam-plasma interactions in space physics. Wave particle correlator experiments have proven to be an effective way to probe the physics of Langmuir wave growth, damping, and particle trapping, but previous experiments were limited mainly to special cases. The CHARM-II rocket flown into active nighttime aurora included a correlator and encountered Langmuir waves along much of its trajectory. Careful consideration of the data yields 57 events in which significant correlations were detected between electrons and Langmuir wave phases. The majority of the selected correlations indicated energy exchange slightly dominant over particle trapping. Examination of the component of the correlation resulting from energy exchange reveals 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 resistive component of the electron bunching implied energy transfer from the beam particles to the wave field, and when the electron beam flux was decreasing, the reverse occurred. Two possible explanations for this effect are the time variation in the interaction of a given energy with the waves in a beam which is time-varying due to dispersion effects, and short-term enhancements in wave growth or damping on a time-varying beam due to time-of-flight effects on the particle distribution. The latter mechanism is investigated and found plausible using a numerical test particle simulation.

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 are ubiquitous in space plasma, including, as examples: the solar wind, where they generate radio bursts (Lin et al., 1981); planetary foreshocks (Filbert & Kellogg, 1979; Gurnett et al., 1981); and the auroral ionosphere (Kintner et al., 1995; M. H. 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 intensive study, particularly with regard to their eigenmode structures (Malaspina et al., 2012; Graham & Cairns, 2013), nonlinear processes (Graham & Cairns, 2014), three-dimensional effects (Malaspina & Ergun, 2008; Dombrowski et al., 2012), and wave-particle correlations (Ergun, Carlson, McFadden, TonThat, & Clemmons, 1991; Ergun, Carlson, McFadden, & Clemmons, 1991; Muschietti et al., 1994; Kletzing et al., 2005; Kletzing & Muschietti, 2006).

Dedicated rocket- or satellite-borne wave-particle correlator experiments have proven to be an effective way to probe the interactions between Langmuir waves and electrons in space plasmas. A detailed summary theory of expected theoretical results from such instruments 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. The width of the resonant part of the electron energy spectrum is inversely related to the length of the wave packet, and the resonant component is unipolar, whereas the reactive component is bi-polar; that is, the reactive component is negative over part of the energy range and positive over the other part. This signature implies that if the energy resolution is too coarse, the reactive component will not be detected at all because its positive and negative parts will cancel out.

An early version of wave-particle correlator flown on a sounding rocket in auroral plasma detected a strong correlation between beam electrons and Langmuir/upper hybrid wave electric fields over an interval of several hundred seconds (Gough et al., 1990). Ergun et al. (1998) flew a wave-particle correlator on an auroral rocket which worked by sorting detected electrons according to the quadrant of the Langmuir wave phase. This experiment revealed evidence of wave-particle interactions but did not address resistive versus reactive components. More recently, Kletzing et al. (2005) reported results from a wave-particle correlator with very good energy resolution (10%) and phase resolution (22.5 degrees). Launched into nighttime aurora, the experiment measured correlations associated with a relatively small number of the most intense Langmuir waves encountered, and the results gave a strong indication of wave trapping of the bunched electrons in those examples. A wave particle correlation experiment was flown on the Freja spacecraft but was not used much in flight (M. Boehm et al., 1994).

The success of these previous wave-particle correlator experiments inspired subsequent rocket experiments using a correlator very similar to that of Kletzing et al. (2005). The CHARM-II sounding rocket (NASA mission 40.025) 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 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. 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 by Kletzing et al. (2016).

The CHARM-II correlator measured over fifty examples of significant Langmuir wave-electron correlations—over an order of magnitude more than all previous experiments combined. The large number of examples for the first time reveals a previously unobservable pattern relating the phase of the correlation to the temporal or spatial gradient in the causative beam electrons. Section 2 describes the instruments which make up the Correlator system and the form of the data. Section 3 describes the method used to identify significant correlation events, and then presents several case study events from the data set as well as statistical observations regarding the data set as whole. Section 4 summarizes these results and analyzes two models developed to explain the observations. A numerical test-particle simulation used to test the models is described in Appendix Appendix A.

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 covering the range of potentially resonant energies with high resolution, and the correlation hardware itself, which processes these data streams 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 a fixed, axial boom parallel to the payload's spin axis, which is maintained within ten degrees of the ambient magnetic field by an attitude control system (ACS). Therefore, this potential difference provides an estimate of the the parallel component of the electric-field, exactly the desired component for maximum sensitivity to Langmuir waves near the plasma frequency, which have electric fields predominantly parallel to the background magnetic field. Active preamplifiers inside each spherical probe assure that the antenna functions as a double-probe over the 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 optimize usage of the dynamic range. The regulated HF signal directly modulates a 5 MHz-bandwidth S-band transmitter, and the resulting waveform is continuously digitized in the ground telemetry station at 10 MHz, with 12-bit resolution. The AGC control signal is digitized onboard at 20 kHz and telemetered on a separate digital telemetry link. This instrument is the latest iteration of a design which has flown on numerous other rocket campaigns in both $E_{\parallel}$ and $E_{\perp}$ configurations, including HIBAR (Samara et al., 2004), PHAZE II (McAdams et al., 1998), SIERRA and RACE (Samara & LaBelle, 2006a, 2006b), and ACES (Kaeppler et al., 2013).

For detection of electrons associated with the Langmuir waves, CHARM-II used eight 'bagel' particle detectors provided by the University of Iowa. These detectors, named for the bakery items they resemble, have energy acceptance ranges of 10% and are characterized by a large geometric factor, as required to provide good counting statistics for correlation with high-frequency waves. 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 when the payload is aligned to the magnetic field. The bagel detectors are tuned to logarithmically spaced energy levels ranging from 200 to 1050 eV. The resulting approximately 10% energy resolution is significantly better than achieved in early correlator experiments, and has been shown in recent experiments to be adequate for detecting both the resistive and reactive components of wave-particle correlations in nighttime aurora (Kletzing et al., 2005).

The University of Iowa Wave-Particle Correlator combines the above described Langmuir wave and electron measurements to measure wave-particle correlations. The correlator was previously flown on two rocket missions, RACE and CHARM, and is described in detail by Kletzing et al. (2016). 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, locked to the frequency of the highest-amplitude component of the incoming wave. When Langmuir waves are unstable, they dominate the HFE wave signal, and under this condition the PLL produces a clean square-wave locked in phase to the Langmuir wave, but at sixteen times its frequency. This square wave can then be used to sort incoming electron counts from each of four detectors into phase bins which accumulate during an integration period—1 ms in the CHARM-II case, corresponding to hundreds of wave periods per timeslice. CHARM-II included two correlators, each analyzing electron counts from four of the eight 'bagel' detectors.

CHARM-II also included an Electrostatic Electron Pitch Angle Analyzer (EEPAA), a 'top-hat' style detector which counts electrons with a 50 ms integration time, sorted into 15°-wide pitch angle bins and 47 logarithmically spaced energy bins from 15 eV to 15.5 keV. Flux-gate magnetometers, Langmuir probes, and DC/ELF/VLF electric field and wave measurements rounded out the instrument complement.

3 Data Presentation

Figure 1 shows a survey of wave and electron spectra from the active period of the CHARM-II flight, 325-800 s after launch. The top panel is an energy spectrogram of 15° and 30° pitch angle EEPAA data, the middle panel is a 300-900 kHz spectrogram of HFE data, and the bottom panel is a plot of $\log_{10}$ of total counts detected by the 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. In the HFE data, the upper cutoff to noise which is near 500 kHz at 350 s but rises to 800 kHz after 700 s is interpreted as the plasma frequency ω_p, which acts as an upper bound to whistler modes in 'underdense' plasma ($\omega_p < \omega_c$). In the auroral ionosphere ω_c typically exceeds 1 MHz, and the plasma frequency is clearly much lower than the upper hybrid frequency ($\sim$ 1.4 MHz), and therefore easily identified despite slight ($<$ 10 degrees) variations in the rocket's alignment parallel to the magnetic field. For resonant electrons in the energy range of the 'bagel' detectors, it can also be deduced that the Langmuir wavelengths will be tens of meters. There are many instances visible in Figure 1 where increased particle counts are accompanied by wave activity near ω_p. The activity near 650 to 660s corresponded to a particularly powerful event in which the peak electric field, estimated to be 1-3 V/m, momentarily saturated the on-board DC electric field instrument.

The correlator system returned 490 seconds of valid data, providing a matrix with accumulated counts $s(t, p, E)$, where t labels each of 490,000 1-ms timeslices, p (=0...15) labels each of the 16 phase bins, and E (=1...8) labels each of the 8 energy levels. The interpretation of

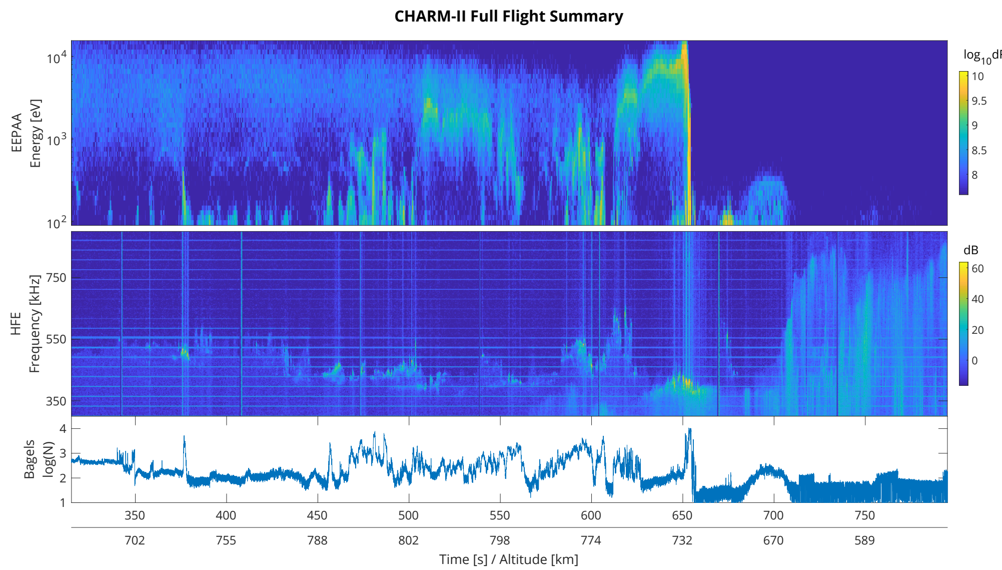

Figure 1. A summary plot of the active period of the CHARM-II flight. Both time after launch and altitude are shown on the x-axis. Top is a spectrogram of 15° and 30° pitch angle EEPAA particle data, with energy on the y-axis, and $\log_{10}$ of the differential energy flux [eV/(cm^2 * ster * s * eV)] 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. Inverted-V structures occur at 610-660 s and 500-560 s, the latter tenuous. The upper cutoff to noise in the HFE panel is interpreted as the plasma frequency. The electron density increases dramatically after 700 s, and broadband noise signals, presumably whistler-mode auroral hiss, occurs at frequencies below f_{pe}.

the phase bin, in terms of the phase between the electrons and the input waveform, varies with the frequency due to phase shifts in the electronics and other delay sources, all of which were carefully calibrated pre-flight. Therefore, a first step taken to aid comparability of timeslices is to shift all phase 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, a more effective quantitative analysis is achieved by examining the Poisson z-score for the data,

$$\sigma(t, p, E) = \frac{s(t, p, E) - \bar{s}(t, E)}{\sqrt{\bar{s}(t, E)}}, \tag{1}$$

where $\bar{s}$ is the mean particle count of the timeslice, $\bar{s} = \frac{1}{16}\Sigma_p s(t, p, E)$. This parameter is a measure of particle over- or under-density in each phase bin with respect to the expected value, assuming all phases are equally probable. The left panel of Figure 2 displays example results of this analysis, showing z-scores expressed as a color scale, with white being zero, red expressing negative values (deficit of counts relative to expected value) and blue expressing positive values (surfeit of counts relative to expected value), for 101 timeslices (on the x-axis), as a function of the electric field wave phase (on the y-axis). As mentioned above, calibration data have been applied to shift the phase bins from the raw values so that they reflect phase relative to the electric field wave.

Figure 2 shows a high-$|\sigma|$ event lasting for multiple timeslices in the middle of the plot (near $t = 50$ ms). However, for most of the times and phase bins the z-scores are insignificant, appearing as 'salt and pepper' in the plot. Cursory examination of the entire data set reveals that most timeslices can be discarded due to a lack of significance. Furthermore, many features that appear significant turn out to be results of natural and instrumental interference. For example, for each timeslice the correlator records whether the PLL is 'locked' or not, and to what frequency it is locked. 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, as this says nothing about the presence of interesting activity at that time, or whether the correlator is locked to the appropriate Langmuir frequency (which can be determined by comparing the 'lock' frequency to the HFE data). 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, based on theory (Kletzing & Muschietti, 2006), 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 automated reduction of the correlator data to identify discrete events—such as that shown in the middle ($t \sim 50$ ms) of Figure 2—is to re-bin the 16 phase bins into four reduced phase bins. Figure 3 shows the scheme: the 16 phase bins p (orange) are summed into four reduced phase bins p' (color coded by p'),

$$s'(t, p', E) = \sum_{n=0...3} s(t, 4p' + n, E), \tag{2}$$

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

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

A z-score,

$$\sigma'_q(t, p', E) = \frac{s'(t, p', E) - \bar{s}}{\sqrt{\bar{s}}}, \tag{4}$$

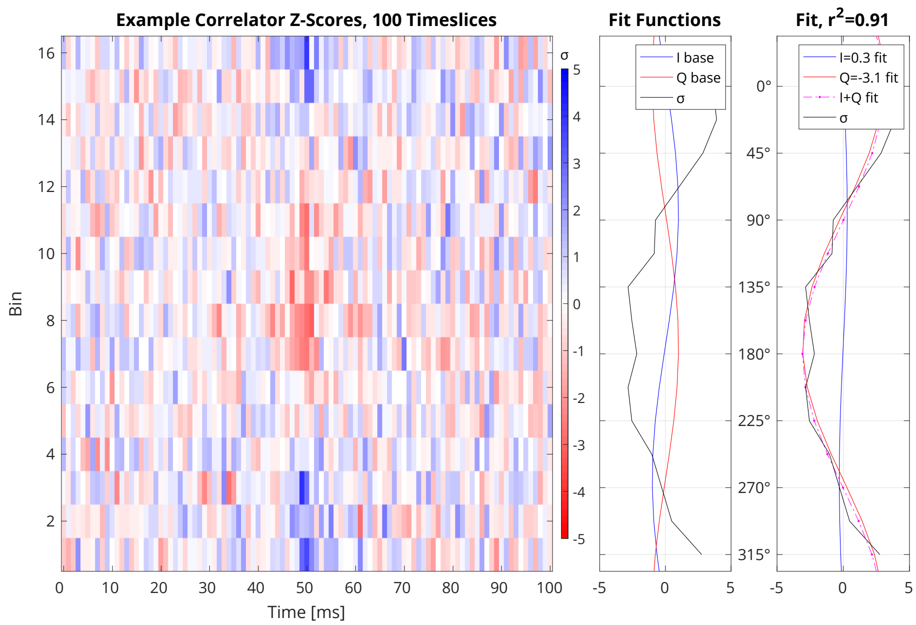

Figure 2. An example plot showing (left) the Poisson z-scores (σ) vs. phase and time for 101 ms of CHARM-II Correlator data. The z-score from -5 to 5 is shown with the blue-to-red color scale, electric field phase is on the y-axis, and time (in relative ms) on the x-axis. A high-$|\sigma|$ event lasting for multiple timeslices occurs in the middle of the plot, but for many of the timeslices, the z-scores 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 goodness-of-fit value. In this case the Q component dominates, implying electrons trapped in the wave. Only 15 of the 57 events were of this type.

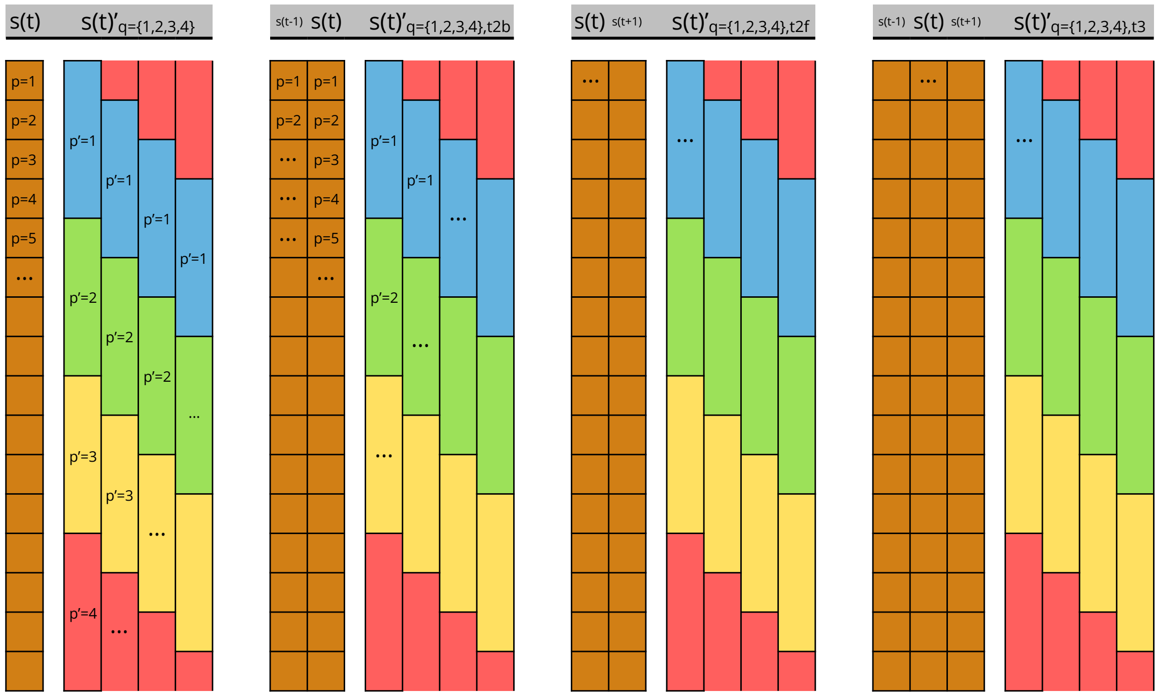

Figure 3. A map of the rebinning done at each timeslice to the raw correlator counts. Each large group is a time integration (singlet, doublets, triplet), and each of those contains four shifted rebinnings from the original 16 correlator bins of $s(t,p,E)$ (shown in dark orange), to the 4 bins of $s'_q(t,p',E)$ (shown in other colors, grouped by p').

was then calculated for each of these reduced bins.

To account for events which span multiple timeslices, two 'doublet' sets and one 'triplet' set were constructed at each timeslice, acting as additional arrays of timeslices (the large divisions in Figure 3). For these, the means and the counts in each reduced phase bin were integrated in time over the two or three raw timeslices, with the doublet sets defined as

$$s'_{q,t2b}(t,p',E) = \sum_{\tau=t-1,t} s'_q(t,p',E) \text{ and } s'_{q,t2f}(t,p',E) = \sum_{\tau=t,t+1} s'_q(t,p',E), \tag{5}$$

and the triplet

$$s'_{q,t3}(t,p',E) = \sum_{\tau=t-1,t,t+1} s'_q(t,p',E), \tag{6}$$

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 the minimum and maximum, Δ, provides a scalar measure of how well a given timeslice matches the expected signature of a wave-particle correlation event.

Initially, a simple global Δ threshold value was used to identify events, resulting in hundreds of candidate events; however, thorough investigation of these candidates revealed many false positives, including many events which were disqualified after examination of diagnostic and contextual data. Figure 4 shows the two major factors which lead to disqualifying candidate events. Panels (a) and (b) contain an energy spectrogram derived from the bagel detectors, and a line plot of the logarithm of counts from the selected bagel, with an 'X' at bottom indicating when the correlator is locked. The spectrogram in panel (c) shows HFE data concentrated around the plasma frequency—about 530 kHz in this case—with a white line indicating the frequency to which the correlator was locked. The overlaid dashed black line indicates the hand-estimated Langmuir frequency at the center time. Note that the signal at this estimated frequency is relatively weak, short-lived, and surrounded by noise. The final two panels show raw HFE waveforms on two different timescales, with the shorter timescale in panel (e) revealing a powerful interfering spike-like signal at a cadence of about 15 μs. This signal, originating on the rocket payload from an unknown source, is frequently seen in the HFE data, and in some cases is the highest-amplitude component of the waveform.

Either factor—weak signal at the Langmuir cutoff or obvious payload interference—would reduce confidence in the correlator phase binning. While in principle the correlator hardware filters out low-frequency interference, here it combines with a poorly defined Langmuir cutoff to yield a highly unstable correlator lock frequency, fluctuating between unrealistically high and low values. In panel (b) we see that the hardware's lock status is intermittent around the center time, definitively indicating a lack of significant coherence in the data. Together, these indicators motivate the discarding of this event and many others like it.

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 considerations such as those illustrated in Figure 4. Figure 5 shows an overview of the final set of 57 thresholded, hand-screened events, as a scatter plot versus bagel number and time, with a histogram versus time overlaid. The final set of events is tabulated and plotted in the Supporting Information. From the histogram it is clear that a majority of the selected correlator events occur during times of high flux at lower energies coincident with the boundaries of inverted-V structures at higher energies, particularly the tight cluster near 652 seconds coincident with the end of the second inverted-V structure. 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 5, showing that the majority of the events were at 260 and 630 eV, with 12 and 23 events, respectively. The 820 and 1050 eV bagels had no qualifying timeslices.

The large number of correlation events inspires a search for statistical patterns in the data which would not have been discernible in previous experiments due to their small number of events. The first step is to decompose the events into resistive and reactive components. The resistive component consists of particles which oscillate in phase or 180 degrees out of phase with the electric field of the Langmuir wave, and the reactive component particles oscillate 90 or 270 degrees out of phase (Kletzing & 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 has a significant value, and it is this form that the event-identification method focused on.

To examine the correlator data in light of the theory, correlator z-scores at each timeslice were fit, using a nonlinear least squares analysis, to a quadrature function vs. bin number p,

$$-I \sin\left((p - p_0) * \frac{\pi}{8}\right) - Q \cos\left((p - p_0) * \frac{\pi}{8}\right), \tag{7}$$

adjusting I and Q to provide the best fit to the observed z-scores, starting from $I = Q = 0.5$. The negative signs and p_0 are determined by calibration data relating the electric-field phase to the bins. The coefficients I and Q in principle represent the strengths of the In-phase and Quadrature signals, i.e. the resistive and reactive linear components. An example set of fits

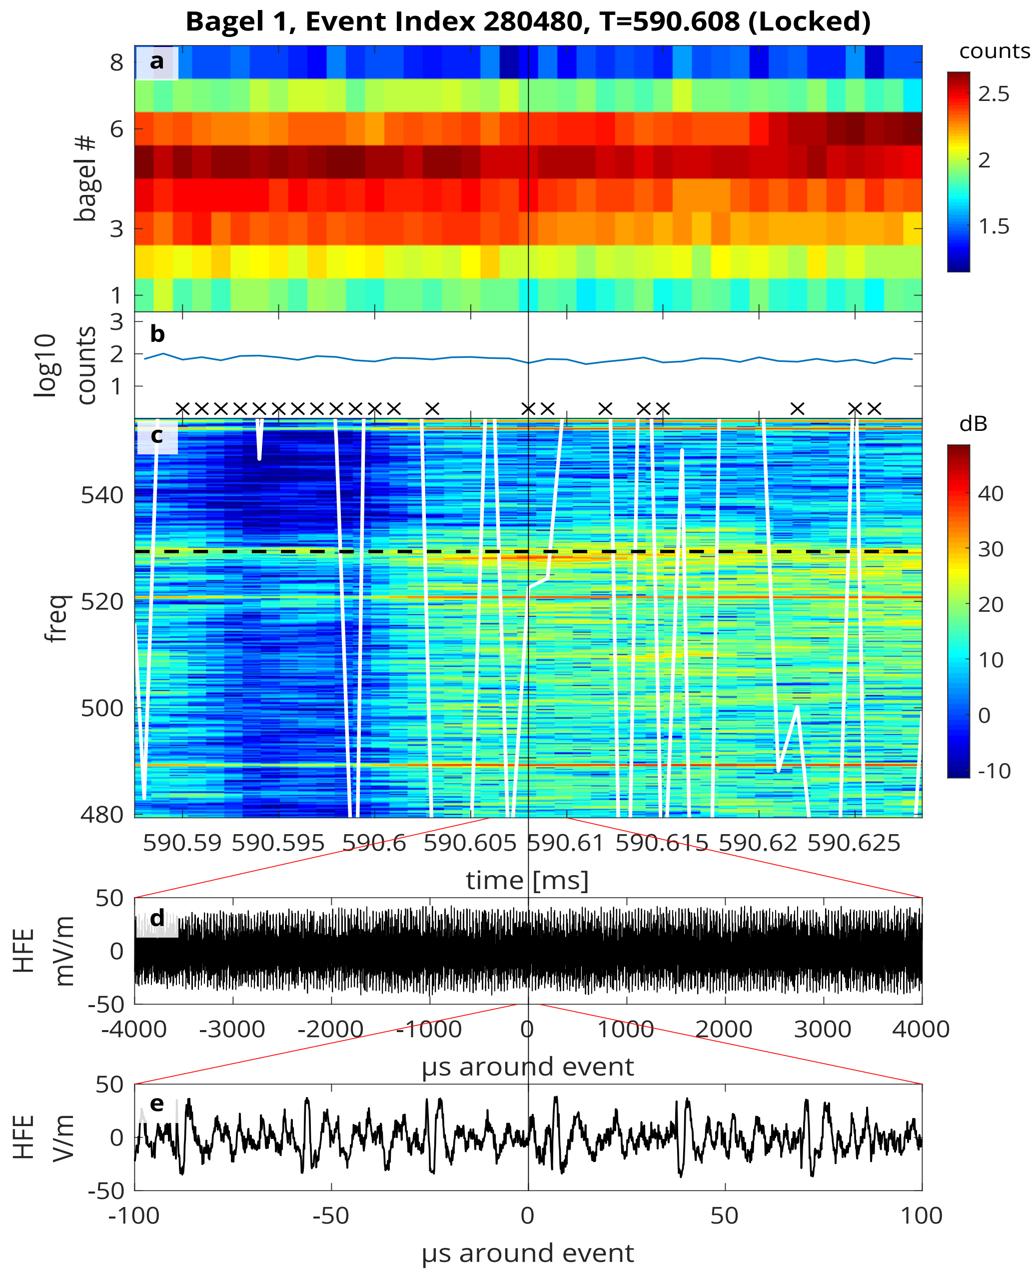

Figure 4. An example of data used to hand-screen events. (a) A spectrogram of 'bagel' detector data; (b) the logarithm of raw total counts for the selected bagel detector, with the correlator lock/no-lock status displayed at the bottom as an X for *lock*; (c) an HFE spectrogram covering the relevant frequency range for Langmuir waves, with overlays for the correlator lock frequency (solid white trace) and plasma frequency identified by inspection (dashed black trace); (d) HFE electric field waveform for a much smaller time interval: ± 4000 μs around the center time; (e) HFE waveform for an even more expanded time interval: ± 100 μs. This event was discarded because it shows a weak and noisy Langmuir cutoff, strong periodic interference at ~ 15 μs cadence, and a fluctuating lock frequency with sporadic *lock* status.

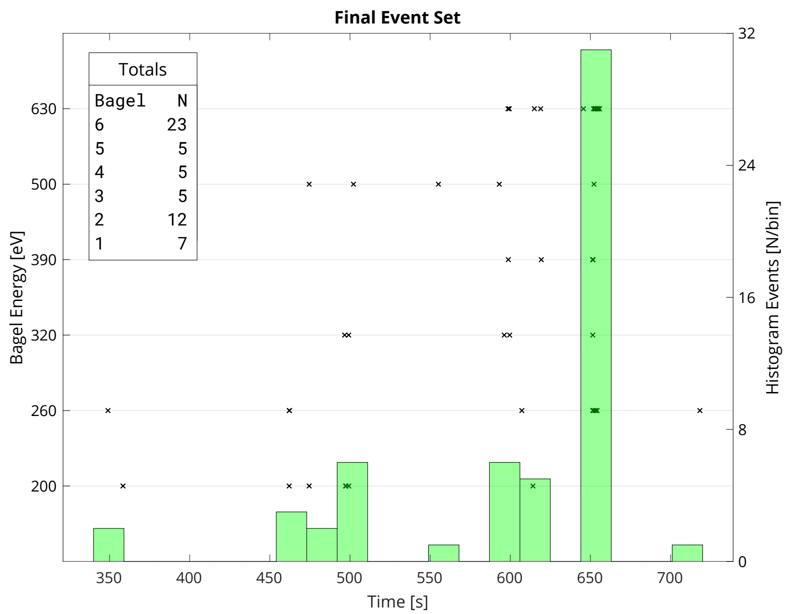

Figure 5. An overview of the final event set $\mathbb{S}$ vs. time for the whole flight, with the per-bagel totals in the table inset top left. Individual events are displayed as x-marks on lines corresponding to bagels on the left vertical axis. Overlaid on this is a histogram of events vs. time, corresponding to event/bin counts on the right vertical axis.

are shown in the line plots on the right of Figure 2, in which the fit has been applied to the middle ($t = 50$ ms) timeslice, with bins on the y-axis (aligned to the left σ plot)—in both line plots, the black line shows the σ values for the center timeslice. In the left plot, the solid blue line shows a normalized (unity amplitude) wave in-phase with the calibrated electric field, corresponding to I, the resistive component, and the dashed red line is the reactive Q component. The right plot shows fitted forms, with blue and red the fitted, separate I and Q contributions, 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. While many of the selected events show r^2 values below 0.5 (r below 0.707), examination of the fitted data shows that they are successful at reproducing the sinusoidal characteristics with correct phases in the correlator data, even when r^2 is as low as 0.2 (r as low as 0.45), which covers the vast majority of the selected events (51 of the 57). Plots of the event fits are included in the Supporting Information.

Figure 6 shows results of fitting I and Q functions for the 57 selected events, expressed as the log of the magnitude of the resistive-to-reactive coefficient ratio, $\log_{10}|I/Q|$, with vertical lines drawn to divide the events by bagel. Some basic statistics of I/Q and $|I/Q|$ are inset on the lower right, and the events are color-coded by their r^2 goodness-of-fit value. Individual plots of the fits are provided in the Supporting Information. There is significant scatter, but the majority of events (42 out of 57), have $\log_{10}|I/Q| > 0$, implying they tend to be dominantly resistive, although most events have both a resistive and reactive component. This result could arise partly from experimental bias: the reactive component is bipolar, implying that if the energy resolution is insufficient the correlation is more easily washed out than that of the resistive component, which is unipolar in energy. However, this bias only applies if the energy resolution is insufficient. The observation of both reactive and resistive components in multiple events suggests that the energy resolution is adequate for many cases, as is also

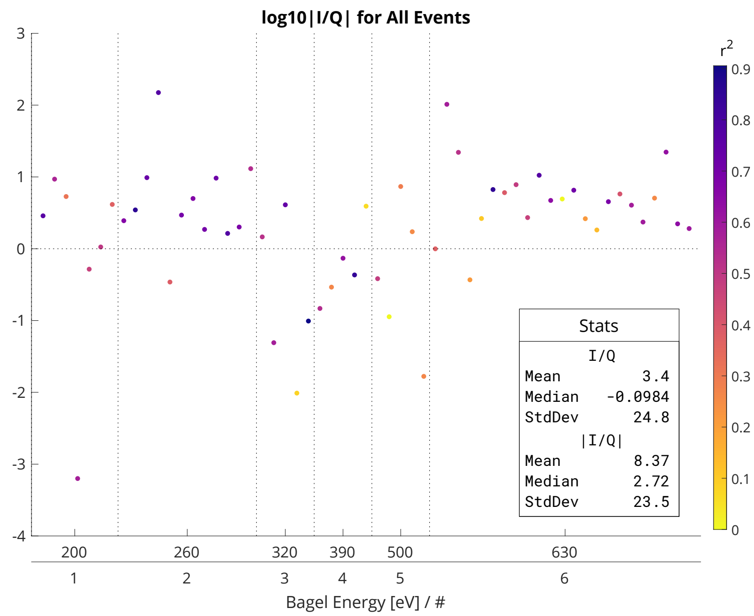

Figure 6. The ratio of the resistive to the reactive component ($\log_{10}|I/Q|$) of each of 57 selected wave-particle correlation events, with vertical lines dividing the events by bagel. Statistics of I/Q and $|I/Q|$ are inset in the lower right, and events are color-coded by their r^2 goodness-of-fit value. Note that the majority of events show more resistive power in the fit than reactive. Individual plots of the fits are provided in the Supporting Information.

known from previous correlator experiments (Kletzing et al., 2005) and suggested by theory (Kletzing & Muschietti, 2006). The tendency of the resistive component to be dominant may appear somewhat counter-intuitive, since one might expect the wave growth/damping stage to be of short duration compared to the static case in which electrons are trapped in the wave but not significantly exchanging energy. However, as shown and discussed below, the majority of selected cases correspond to time variations of the causative electron beam, conditions which might be associated with wave growth/damping and hence energy exchange between electrons and the waves.

Figure 7 shows four example events. In each case, for context, the top panel shows an electron spectrogram from the EEPAA, with high energy resolution and range but poor time resolution, covering one second centered on the event. The lower three panels show 41 ms of data around the event: a spectrogram of the total counts of each bagel vs. bagel energy, the values of I and Q, and a color plot of σ versus wave phase and time in the same format as the left panel of Figure 2.

The top left panels of Figure 7 show data from an event at 607.203 s after launch, involving 260 eV electrons. There is evidence of higher-energy (>500 eV) beams in the EEPAA data leading up to the event, but no significant fluxes at higher energy during the event. The bagel spectrogram reveals sporadic, short-duration (2-5 ms) 260 eV beams. The selected event occurs in the middle of the time frame, at 607.203 s, identifiable as a column of high σ values with strong negative (red) values at phases <180 degrees, and strong positive (blue) values at phases >180 degrees. This phase pattern corresponds to a fit in which $I < 0$, called a 'negative-I' event. The negative-I event occurs on the leading edge of the short-duration beam. A bit further on in the record, associated with the short-duration beam at 607.219 s, there occurs a negative-I event on the leading edge of the beam, and a positive-I event on the trailing edge.

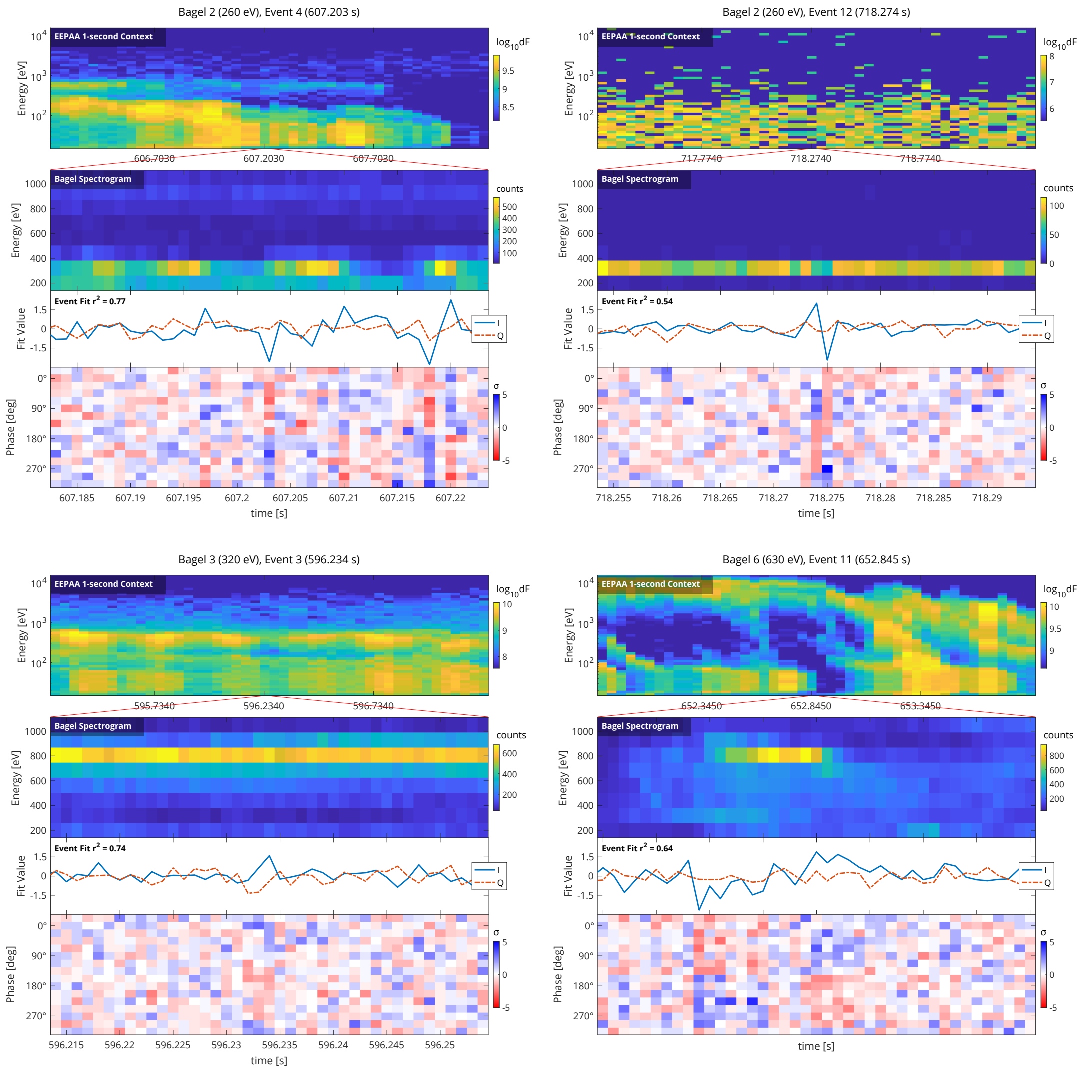

Figure 7. Example wave-particle correlation events from among 57 identified by a combination of automated and manual screening. Each set of four panels includes: (a) one second spectrogram of EEPAA data showing context for the event; (b) 40-ms spectrogram of bagel data; (c) line plots of reactive/resistive fit values; and (d) σ versus wave phase versus time.

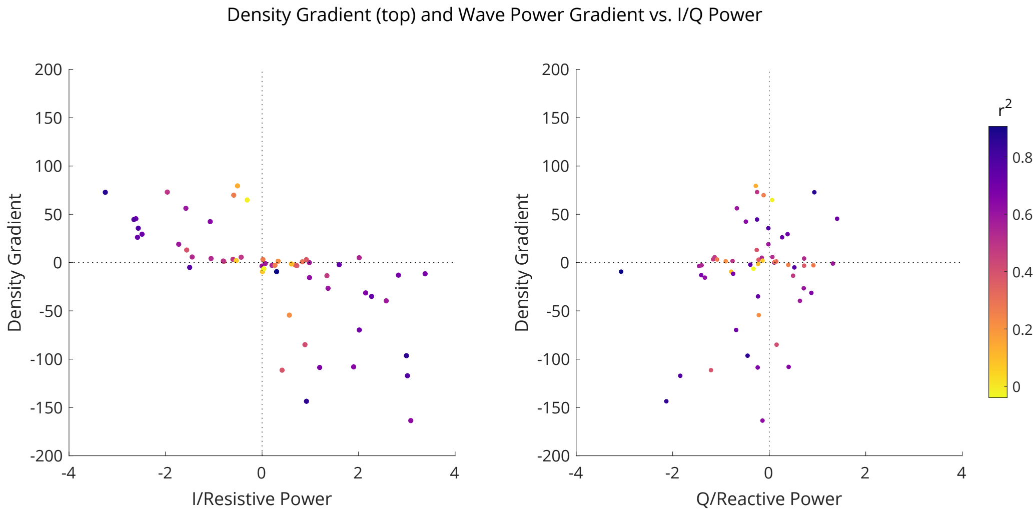

Figure 8. Scatter plots of bagel count (i.e. resonant electron density) gradient vs. I (reactive, left column) and Q (resistive) fit values, with color as the r^2 goodness-of-fit value. Note the clear relation in the left plot, and the lack of such on the right.

This pattern persists in the other three examples of Figure 7. The top right panels show two closely spaced examples, at 718.274 s and also involving 260 eV electrons, occurring on the trailing edge of a beamlet and the leading edge of a subsequent one. Again the example on the trailing edge is positive-I, and that on the leading edge is negative-I type. The examples in the bottom panels differ in that they span multiple timeslices rather than a single timeslice as in the top panels. The bottom left panel, at 596.234 s and involving 320 eV electrons, shows a positive-I event spanning a couple of timeslices that occurs in a region in which the 320 eV and 820 eV beam densities are dropping off. The bottom right panel, at 652.845 s and involving 630 eV electrons, shows two events, each spanning a range of timeslices at the front and back edge of a short-duration beam which is most prominent at 820 eV, but shows some structure at lower energies. Following the pattern, the event on the leading edge where the beam density increases is negative-I type, and that on the trailing edge where the beam density decreases is positive-I type.

To explore this relationship further, Figure 8 shows scatter plots of the temporal gradient of the electron beam flux ($\nabla_t n_B$) inferred from the bagel detectors versus the I and Q fit values. The points are colored according to their r^2 goodness-of-fit value from the reactive/resistive fitting. The left panel demonstrates from this statistical approach the correlation between the I value and the beam flux gradient which was illustrated by multiple examples in Figure 7. 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}, implying that the null hypothesis—that there is no relation between I and $\nabla_t n_B$—should be rejected. The right panel of Figure 8 shows that there is no pattern evident between the Q fit value and $\nabla_t n_B$. Linear regression of the right panel has a small t-statistic and p-value ≥ 0.15, suggesting that the relationship between I and $\nabla_t n_B$ is significant. Table 1 summarizes additional statistical tests performed on the data, showing that a two-sample Kolmogorov-Smirnov test finds a significant difference between the $\nabla_t n_B < 0$ and $\nabla_t n_B > 0$ distributions of I, and that in linear cross-correlations I and $\nabla_t n_B$ are significantly correlated.

Correlation might be expected between the I and Q fit values and the gradient in wave power for cases in which the resistive component dominates, because such correlations imply exchange of energy between electrons and waves. Such a correlation should be observable if the wave power at a given frequency is dominated by the wave mode that is interacting with the energy of electrons found to be correlated; however, the expected correlation would be on timescales less than the 1 ms interval over which the correlation is detected. The wave data appear inadequate to address this issue unless the gradient is larger than tens of dB over 0.1 ms, and gradients of this amplitude were not observed during any of the events.

Table 1. Statistical tests[a] on the Figure 8 scatter plots.

	I-$\nabla_t n_B$		Q-$\nabla_t n_B$	
	s	p[b]	s	p
Linear Regression	-6.97	4.0×10^{-9}	2.26	0.027
Kolmogorov-Smirnov Test	1[c]	2.8×10^{-8}	0	0.23
Cross-Correlation	-0.68	7.8×10^{-9}	0.24	0.062

[a] These statistics evaluate relations from I and Q to $\nabla_t n_B$. The 's' columns contain the significant output of the given test: from top to bottom, the t-statistic, null-hypothesis rejection, and correlation coefficient.

[b] In all tests, note the extremely low p-value of the I-$\nabla_t n_B$ relation compared to the Q-$\nabla_t n_B$ p-value, implying that this column's comparisons yield strong statistical results.

[c] The Kolmogorov-Smirnov test's result shows only that the null hypothesis (that the distributions are the same) is accepted or rejected. Here, note that it is rejected for the I-$\nabla_t n_B$ relation.

4 Discussion

During the CHARM-II flight, the correlator detected 57 wave-particle correlation events deemed significant after application of the automatic algorithm followed by manual inspection and screening. An analysis of the reactive and resistive components, I and Q, of these events reveals a correlation between a positive electron beam flux temporal gradient $\nabla_t n_B$ at a given bagel energy level, as in the case of an electron beam appearing at that energy level, and a negative value of I for the coincident wave-particle correlation event, and a similar relation between a negative $\nabla_t n_B$ and positive I. Careful consideration of the correlator calibration confirms that 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 in a case study of data from the RACE sounding rocket. It is possible that the extreme wave power during the majority of the 57 selected events may mask such an effect, particularly given that the amplitude modulation typical of the bursty Langmuir waves observed in the RACE experiment is not prevalent in the CHARM-II HFE data.

It is plausible that time-of-flight effects on a time-varying electron beam may lead to a situation in which Langmuir waves experience spikes in growth or damping like those seen in Figure 7. Given a beam population 'turned on' promptly at some altitude above the rocket, and a 'warm' background population of degraded secondary electrons associated with the beam, as depicted in Figure 9, the highest-energy particles arriving earliest at the rocket location may create a positive-slope region in the reduced parallel distribution function. This will produce wave growth, characterized by transfer of energy from electrons to waves and detected by the correlator as a negative-I type wave-particle correlation. In contrast, when the beam 'turns off,' it is also possible, depending on the relative energies and densities of the two electron populations, that an enhanced negative slope will appear in the reduced distribution function as the higher-energy particles disappear first, leading to enhanced damping of the waves due

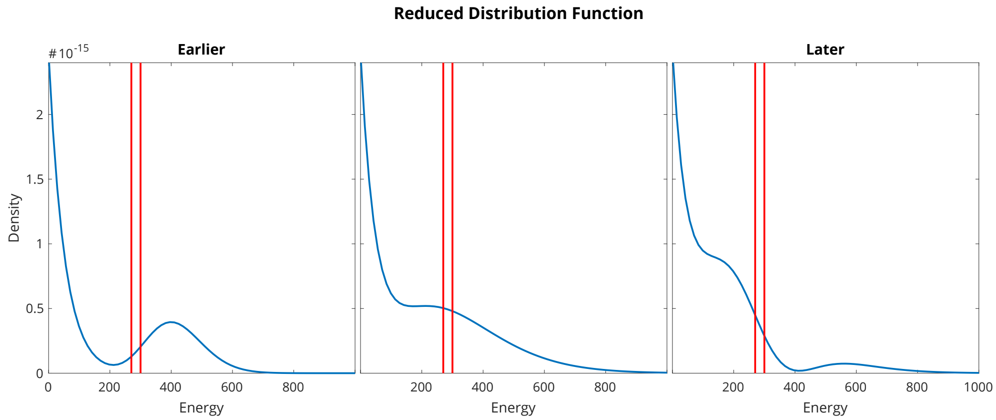

Figure 9. A cartoon showing an explanation for the relation seen between I and $\nabla_t n_B$. As a beam appears, the high-energy particles are the earliest to arrive, leading to an exaggerated positive slope and wave growth (left). The remainder of the particles and distribution relaxation then yield 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).

to interactions with electrons in the narrow energy region, and resulting in a wave-particle correlation event of the positive-I type. The presence of short-lived beam features in the top-left and bottom-right plots of Figure 7, and their temporal relation to the nearby correlator events, provide evidence suggesting this mechanism.

An extreme case of such time-of-flight effects which is easy to perceive physically is the dispersive electron beam which has been observed frequently, e.g. by Kletzing et al. (2016). This type of dispersive beam, in which higher energies arrive before lower energies, 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 (Kletzing & Hu, 2001; Chen et al., 2005). The beam energy decreases with time from left to right, and 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 yields Langmuir wave growth, and the latter case corresponds to $df/dv_{\|} < 0$ which damps Langmuir waves. 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.

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. The easiest way to try to visualize and confirm the effect is by using a numerical simulation.

To this end, numerical calculations have been performed using a a test-particle simulation tool described in Appendix Appendix A. The basis for this calculation is Liouville's theorem which allows individual test particles to stand for portions of the electron phase space distribution. A library of particle trajectories has been created by launching test particles into a model converging magnetic field, and allowing 5000 km of simulated flight. For computational efficiency the calculation proceeds backwards, with test particles representing the full distribution function at the low-altitude end traced backwards in time and space to find out which region of phase and physical space they connect to at the top end. It is then possible to model arbitrary input electron distributions at the top end, and time variations in those, and determine the electron distribution as a function of time at the bottom end. As a final step in the numerical calculation, the reduced distribution is calculated from the full distribution function, and the Langmuir wave growth or damping rate is calculated from the reduced distribution function. Details of this multi-step numerical tool are given in Appendix Appendix A.

For this paper, the simulation was run with ionospheric background parameters $f_{pe} = 400$ kHz and $T_{iono} = 2000$ K, secondary background $T_{bg} = 2 \times 10^5$ K with $n_{bg} \simeq 1.07 \times 10^6$ m^{-3}, and beam parameters $T_{beam} = T_{bg}/5$ and $n_{beam} = n_{bg}/50$ with a velocity shift $\delta(t) = 400$ m/s.

Figure 10 shows the final results of the simulation of Langmuir wave growth/damping, resulting from a beam 5000 km above the observation point that was turned on promptly at $t = 5$ s and off promptly at $t = 10$ s. The top panel shows the ratio of the beam density to the background density at the top end, as a fraction of time as imposed in the simulation. The middle panels show the maximum Langmuir wave growth rate γ, in units of inverse seconds, as a function of time—and as a color scale (red=growth, blue=damping)—for two selected intervals, on the left starting two seconds after the beam was turned on, and the right starting two seconds after it was turned off. The lower panels show growth/damping rate on a symmetric-logarithmic color scale as a function of wavenumber k and time for the same two selected time intervals. All these calculations correspond to a selected frequency range just above the plasma frequency, with parameter ω_t (described in Appendix Appendix A) ranging from approximately ω_p to $1.02\omega_p$, corresponding to $k = 0.01$ to 1. The lower limit of k is determined by the resonance condition $kv_\parallel = \omega_p$. For the theory to be viable $v_\parallel$ should be kept far enough below c that any relativistic corrections remain small—in our simulation, where $\omega_p = 400$ kHz, this sets a lower bound on k of $10^{-2}m^{-1}$.

Significantly, Figure 10 shows both a growth rate spike during the beam arrival and a damping enhancement during beam departure as predicted by the qualitative argument above, thus supporting the mechanism advanced for the relationship between the direction of energy flow in the wave-particle correlation and the gradient in the beam density. 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.

The bottom panel of Figure 10 shows that the strongest growth and damping are associated with the long-wavelength modes; this 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 time frame 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 are suggestions of short-timescale structures in the simulation events, for which the growth and damping rates exhibit multiple shorter-timescale peaks. The timing of these fine features, however, has been found to shift significantly based on parameters of the simulation, such as shifting the beam up and down in energy, or removing parts of the beam distribution function, although the overall result showing enhanced growth and damping is robust (Dombrowski, 2016). These tests suggest that the fine structure seen in Figure 10 may depend on aspects of the simulation system.

Another variable is the fact that the instrumentation is in motion. Given that the payload's total velocity ranges from 1300 to 2600 m/s during the flight period of interest and the

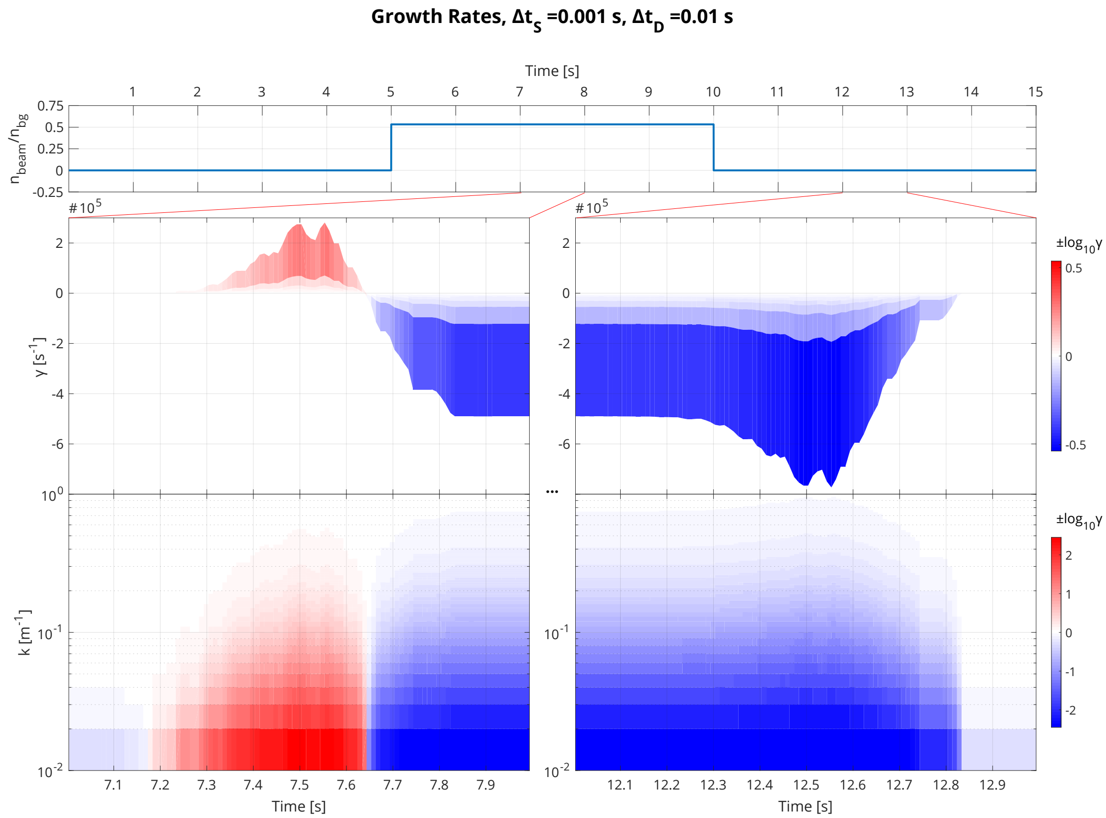

Figure 10. Results from the simulation: Langmuir wave growth rate γ, versus k and time (on the horizontal axis), calculated for ω_t from approximately ω_p to $1.02\omega_p$. The top panel shows the n_{beam}/n_{bg} at the top, while the two columns are zoomed into the times at which the bulk of the particles arrive during beam turn-on (left) or depart with beam shutoff (right), after an approximate 2-second travel time—for the fastest particles—across the length of the simulation. 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. The lower limit of the plot is determined by the resonance condition $kv_{\parallel} = \omega_p$, with $v_{\parallel}$ kept to where any relativistic correction is small.

dynamic, non-uniform nature of the aurora, it is plausible that the short duration of detected events is due to payload transitioning in and out of plasma structures with limited spatial extent. To deconvolve the effects of payload motion on correlation measurements, synchronized, multi-point observations would be required.

Further investigation using more sophisticated tools is required to show whether the mechanism described here can produce temporospatial structuring in wave-particle correlations on the observed few-ms timescales. The numerical simulation described here does, however, prove the overall plausibility of the mechanism.

5 Conclusions

The CHARM-II sounding rocket carried a Langmuir wave receiver, eight large-geometric factor upward-looking electron detectors tuned to tightly spaced energies, and a wave particle correlator to an apogee altitude of 802 km in substorm aurora, returning data above 600 km from approximately 500 s after launch. Applying both automated and manual event selection, 57 timeslices containing wave-particle correlation events were identified as statistically significant 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 resistive component of the electron bunching implied energy transfer from the beam particles to the wave field, 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. (2016).

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.

Appendix A Simulation

The aim is to simulate a minimal-complexity environment that is sufficient to probe whether an electron beam with reasonable characteristics, and which shows significant Langmuir wave growth upon its appearance, also shows enhanced wave damping as it disappears. This goal requires calculating the short-timescale behavior of the wave growth rate as the electron distribution function evolves in time due to the temporal turning on and off of the beam, combined with travel-time effects in the converging magnetic field geometry. These factors compel the use of a complex, flexible—and computationally intensive—test particle simulation system. Its application to this case shall use simple, gyrotropic magneto-kinetic parameters, with no inter-particle interaction or wave-particle scattering.

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

$$f(\bar{x}, \bar{v}, t) = f(\bar{x}', \bar{v}', t'), \tag{A1}$$

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 a positive-upward, field-aligned coordinate, with $z = 0$ corresponding to the beam generation altitude. To simulate only those particles which will arrive at the 'detection point' at $(x = y = 0, z = -5000$ km) requires using a deterministic (i.e. time-reversible) simulation method: we watch for when test particles originating at the detection point with an upward velocity 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 & 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 period, 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.

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

$$\bar{B} = -\frac{zr}{L^2}\hat{r} + \frac{1+z^2}{L^2}\hat{z} = -\frac{xz}{L^2}\hat{x} - \frac{yz}{L^2}\hat{y} + \frac{1+z^2}{L^2}\hat{z}, \tag{A2}$$

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 wave-particle correlator, particles were launched at 31 energies linearly spaced in velocity, with energies ranging from 25 to 1225 eV. At each energy, particles were launched at each 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. To ensure that our hemispherical segments sweep out constant solid angles $\Delta\Omega$—herein set to 0.001 steradians—requires that the azimuthal launch angles ϕ be related to by their separation $\Delta\phi$ to the pitch angle separation $\Delta\alpha$ such that $\Delta\phi = \frac{\Delta\Omega}{\sin(\alpha)\Delta\alpha}$. We determine and use the integer number of angles ϕ which most closely achieve this relation.

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 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, $t_0 \approx 714$ ns, $r_0 \approx 0.337$ m, and $v_0 = 0.00989$ c, corresponding to 25 eV.

Figure A1 shows some basic diagnostics of the output of this simulation, with expected trends compared to total energy and launch pitch angle. Once the table $\mathbb{T}(\mathsf{E}, \alpha)$ has been created, it becomes possible to reverse the velocities and launch distributions downward, towards the simulated 'detector'.

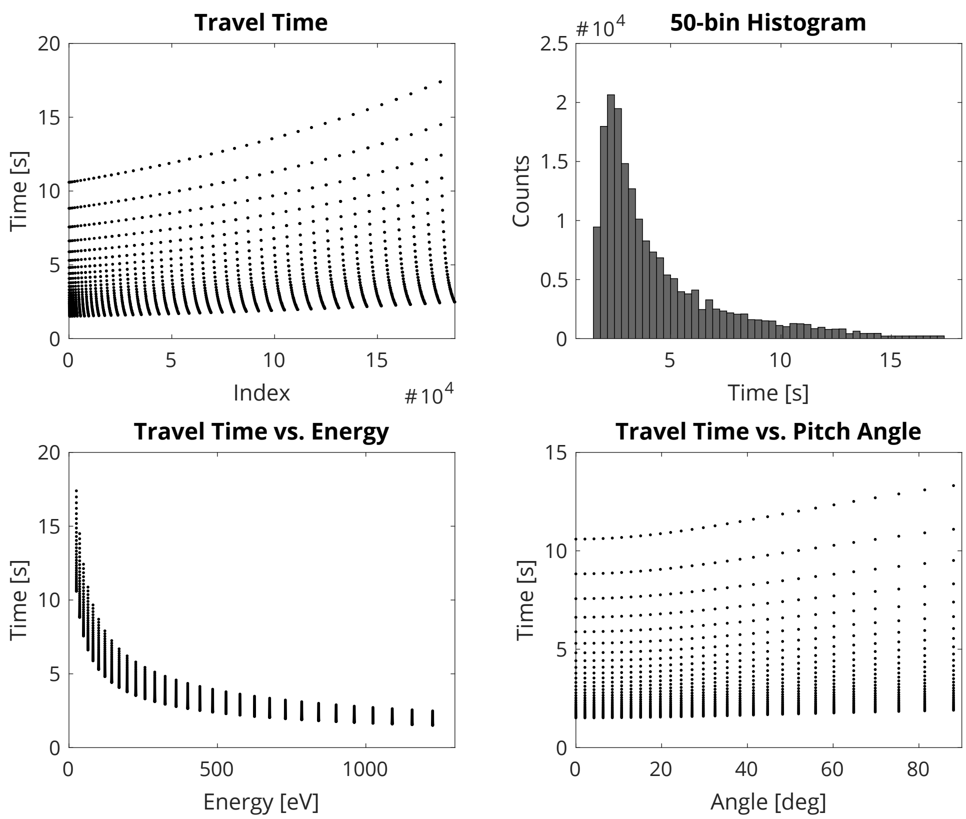

Figure A1. Various diagnostics of the input characteristics and resultant travel times in the test-particle simulation. All time axes are travel time along the length of the simulation. Energy and pitch angle in the lower row are the launch values at the detection point (high magnetic field).

The beam 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 beam is 'turned on' (or off) instantly at a specified time by adding and removing an additional component to the particle distribution sum. The environment 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, the x and y positions of particles in the simulation are neglected.

For purposes of calculating growth rate, the distribution at the rocket altitude is presumed composed of three components. First, there is a cold background, f_{iono}, with a 2000 K temperature which is taken to be uniform and constant, unaffected by the particle tracing, and mainly plays the role of contributing the largest part of the plasma density, hence essentially determining the plasma frequency. A secondary background distribution serves as the main background to the 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. It is defined as an isotropic Maxwellian,

$$f(|v|) = n_e\,(2\pi)^{3/2}\;v_{th}^3\;\mathsf{e}^{-\left(\frac{v}{2v_{th}}\right)^2} \Rightarrow f_{bg} = n_{bg}\left(\frac{m_e}{2\pi k T_{bg}}\right)^{3/2}\mathsf{e}^{-\frac{m_e v^2}{2kT_{bg}}}, \tag{A3}$$

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 third population is the beam, which is also originally Maxwellian, isotropic, and homogeneous in velocity space at the top of the simulation region, except that its parameters can be considered time-varying $T_{beam}(t)$, $n_{beam}(t)$, and a velocity shift $\delta(t)$:

$$f_{beam}(|v|, t) = n_{beam}(t)\left(\frac{m_e}{2\pi k T_{beam}(t)}\right)^{3/2}\mathsf{e}^{-\frac{m_e(v-\delta(t))^2}{2kT_{beam}(t)}}. \tag{A4}$$

In practice, $T_{beam}(t)$ and $n_{beam}(t)$ are set as fractions of the secondary background values. The final distribution is the sum of these,

$$f(v,t) = f(v_\parallel, v_\perp, \phi, t) = f_{iono} + f_{bg} + f_{beam}(t). \tag{A5}$$

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

$$f(v_\parallel, v_\perp, t) = \sum_{i\forall\phi} f(v_\parallel, v_\perp, \phi_i, t)\Delta\phi_i, \tag{A6}$$

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

Time variations have been neglected until now because the hemispheric interpolation introduces no time dependence, and so the azimuthal sum has none either. Simulating the distribution function from the electrons being emitted effectively continuously from a source region requires setting 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 $10\Delta t_S$. The next step is interpolating and reducing away a dimension from the test-particle simulation, while taking travel times into consideration. 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 $t_0 + t_T \le \tau$ and $> \tau - t_D$. There is also an implicit sum here as the detector integration is accounted for, which requires its own weighting value for f, simply the ratio of the launch time and the integration time, i.e.

$$f(v_\parallel, v_\perp, \tau) = \frac{\Delta t_S}{\Delta t_D} f(\mathbb{J}_\tau). \tag{A7}$$

Therefore, the second step in obtaining a one-dimensional reduced distribution function, after summing over the azimuthal angles, is to perform a time-limited, time-weighted sum over the perpendicular velocities, following the standard Landau theory for parallel propagation (Ergun et al., 1993, e.g.). This is slightly complicated by the large number of unique $v_\parallel$ values. Defining '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,

$$f(v_{\parallel,i}, \tau) = \sum_{j \forall v_\perp ; \mu_\parallel} (v_{\perp,j+1} - v_{\perp,j}) \frac{f(\mu_\parallel, v_{\perp,j}, \tau) + f(\mu_\parallel, v_{\perp,j+1}, \tau)}{2} v_{\perp,j}, \tag{A8}$$

where the factor $v_{\perp,j}$ is the phase-space cell weighting. Figure A2 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.

Finally, Langmuir wave growth rate can be calculated from the time-integrated one-dimensional reduced distribution function in $v_\parallel$. For a given cold ionospheric background plasma frequency ω_p, wave vector $\mathbf{k} = k_\parallel \mathbf{B}/B$, and test wave frequency ω_t given by the dispersion relation, the growth rate is

$$\gamma(f(v_\parallel), k_\parallel, \omega_p, \omega_t, \tau) = \left(\frac{d\epsilon}{d\omega} \right)^{-1} \mathrm{Sign}[k] \ \frac{\pi \omega_p^2}{k^2 n_e} \ \left[\frac{\partial f(v_\parallel, \tau)}{\partial v_\parallel} \right]_{k v_\parallel = \omega_t}, \tag{A9}$$

where ϵ is the dielectric function, approximated as $1 - \frac{\omega_p^2}{\omega_t^2}$ for cold plasma. The derivative $\partial f(v_\parallel, \tau)/\partial v_\parallel$ is calculated at a test velocity related to the beam parameters, specifically the closest $v_\parallel$ value to the resonance condition, ω_t/k.

Ideally, Δt_D should match the correlator's time resolution, 1 ms, and also 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 the above guideline that $\Delta t_S \leq \Delta t_D/10$, this would require storage of prohibitive numbers of time-overlapping distributions, so some sort of approximation is necessary. Figure A1 shows that the majority of particles arrive within 5 seconds, so taking that as the settling time is a reasonable approximation. The calculations shown in the paper use Δt_D of order 10-50 ms; short test calculations with $\Delta t_D \sim 1$ ms showed no obvious qualitative differences, but were impractical for the simulations in Figures 10 to A2. As an additional approximation, k and ω_t are determined via the warm plasma dispersion relation, $\omega_t = \omega_p + \frac{3}{2} k^2 v_{th}^2 / \omega_p$, where v_{th} is the background ionospheric thermal speed $\sqrt{3kT_{iono}/m_e}$. With these approximations Langmuir wave growth rates are calculated as a function of frequency, wavenumber and time at the rocket altitude, for postulated time varying beams injected at the top of the simulation region, revealing variations in growth or damping of the waves associated with time-of-flight effects on the particle distribution function at the rocket altitude.

Acknowledgments

The authors thank engineers David McGaw and Jeff Dolan for their work with the particle correlator and its components. The NASA/NSROC team led by Ted Gass executed a successful launch from Poker Flat Research Range. The authors thank Dr. Wayne Scales for providing the original magneto-kinetic test-particle code on which this simulation system was based, and Susan Schwartz and Bill Hamblen for assistance running the code on the Dartmouth Discovery cluster. CHARM-II sounding rocket data are archived at `https://dartmouth.app.box.com/v/CHARM2-rocket-40025-HF-data`, and simulation codes are available on GitHub: `https://github.com/dc-tpgr-sim`. Work at Dartmouth College was supported by NASA through grant NNX17AF92G; work at University of Iowa was supported by NASA through a sub-award of the same grant; work at the University of Sydney was supported by its Special Studies Program.

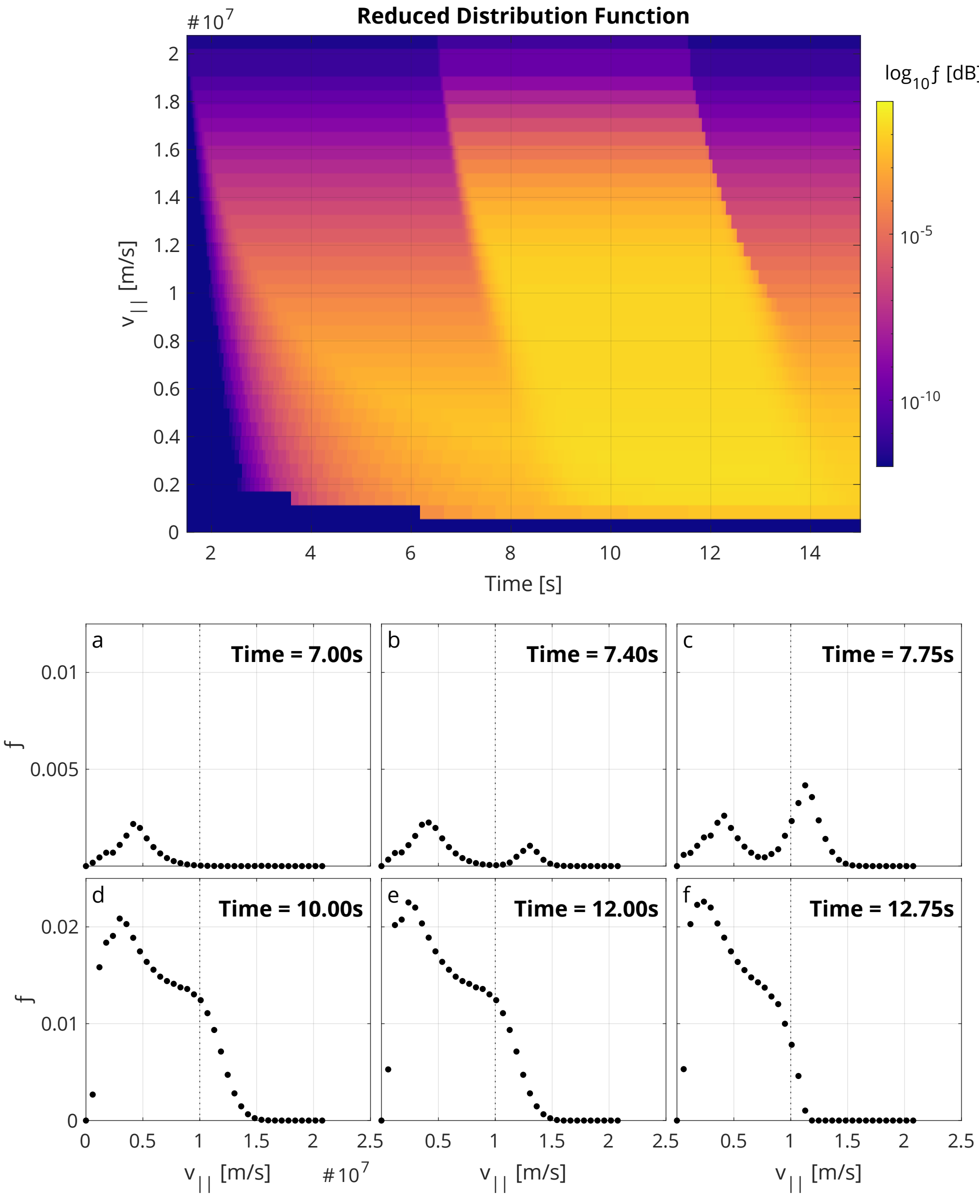

Figure A2. Reduced distribution function values. Top shows a color plot vs. $v_{\|}$ and time, for the entire time span of the test. Below, six timeslices from the beam arrival period, showing the formation and disappearance of a positive slope (compare panels c, d, and f to Figure 9). Note that the component of low-energy electrons centered near 0.5×10^7 ms^{-1} is the tail of the background (bg) distribution. It has not yet arrived before beam turn-on, but is considered too small to significantly affect the results.

References

Birdsall, C. K., & Langdon, A. B. (2005). *Plasma physics via computer simulation* (S. Cowley, P. Stott, & H. Wilhelmsson, Eds.). New York: Taylor & Francis Group.

Boehm, M., Paschmann, G., Clemmons, J., Höfner, H., Frenzel, R., Ertl, M., … Lundin, R. (1994, Apr). The tesp electron spectrometer and correlator (f7) on freja. *Space Sci. Rev.*, *70*, 509-540.

Boehm, M. H. (1987). *Waves and static electric fields in the auroral acceleration region* (Doctoral thesis). University of California, Berkeley, Berkeley, CA.

Boris, J. P. (1970, Nov). Relativistic plasma simulation-optimization of a hybrid code. In J. P. Boris & R. A. Shanny (Eds.), *Proceedings of the conference on the numerical simulation of plasmas (4th)* (p. 3-67). Washington, DC: National Technical Information Service.

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

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

Dombrowski, M. P. (2016). *Sounding-rocket studies of langmuir wave microphysics in the auroral ionosphere* (Doctoral thesis, Dartmouth College). doi: 10.1349/ddlp.1956

Dombrowski, M. P., LaBelle, J., Rowland, D. E., Pfaff, R. F., & Kletzing, C. A. (2012, Sep). Interpretation of vector electric field measurements of bursty Langmuir waves in the cusp. *Journal of Geophysical Research (Space Physics)*, *117*, 9209. doi: 10.1029/2012JA017741

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

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

Ergun, R. E., Delory, G. T., Klementis, E., Carlson, C. W., McFadden, J. P., Roth, I., & Temerin, M. (1993, Mar). Vlf wave growth from dispersive bursts of field-aligned electron fluxes. *J. Geophys. Res.*, *98*(A3), 3777-3787.

Ergun, R. E., Malaspina, D. M., Cairns, I. H., Goldman, M. V., Newman, D. L., Robinson, P. A., … Kaiser, M. L. (2008, Aug). Eigeinmode structure in solar-wind langmuir waves. *Phys. Rev. Lett.*, *101*(051101).

Ergun, R. E., McFadden, J. P., & Carlson, C. W. (1998). Wave-particle correlator instrument design. In *Measurement techniques in space plasmas: Particles* (Vol. Particles, p. 325-331). American Geophysical Union.

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

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

Graham, D. B., & Cairns, I. H. (2013, Sep). Constraints on the formation and structure of langmuir eigenmodes in the solar wind. *Phys. Rev. Lett.*, *111*, 121101. Retrieved from `http://link.aps.org/doi/10.1103/PhysRevLett.111.121101` doi: 10.1103/PhysRevLett.111.121101

Graham, D. B., & Cairns, I. H. (2014, Apr). Dynamical evidence for nonlinear langmuir wave processes in type iii solar radio bursts. *J. Geophys. Res.*, *119*(4), 2430-2457. doi: 10.1002/2013JA019425

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

Kaeppler, S. R., Kletzing, C. A., Bounds, S. R., Gjerloev, J. W., Anderson, B. J., Korth, H., … Heinselman, C. J. (2013). Current closure in the auroral ionosphere: Results from the auroral current and electrodynamics structure rocket mission. In A. Keiling,

E. Donovan, F. Bagenal, & T. Karlsson (Eds.), *Auroral phenomenology and magnetospheric processes: Earth and other planets* (pp. 183–192). American Geophysical Union. Retrieved from https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2011GM001177 doi: 10.1029/2011GM001177

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

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

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

Kletzing, C. A., LaBelle, J., Bounds, S. R., Dolan, J., Kaeppler, S. R., & Dombrowski, M. P. (2016, Feb). Phase sorting wave-particle correlator. *J. Geophys. Res. Space Physics*, *122*, 20969-2078.

Kletzing, C. A., & Muschietti, L. (2006). Phase correlation of electrons and langmuir waves. In J. LaBelle & R. A. Treumann (Eds.), (Vol. 687, p. 313-337). Springer Berlin Heidelberg.

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

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

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

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

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

McAdams, K. L. (1999). *Sounding rocket based investigations of hf waves in the auroral ionosphere* (Doctoral thesis, Dartmouth College). doi: 10.1349/ddlp.2820

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

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

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

Samara, M. (2005). *Sounding rocket based investigations of whister, upper hybrid and langmuir waves in the auroral ionosphere* (Doctoral thesis, Dartmouth College). doi: 10.1349/ddlp.81

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

Samara, M., & LaBelle, J. (2006b). Structured waves near the plasma frequency observed in three auroral rocket flights. *Annales Geophysicae*, *24*(11), 2911–2919. Retrieved from http://www.ann-geophys.net/24/2911/2006/ doi: 10.5194/angeo-24-2911-2006

Samara, M., LaBelle, J., Kletzing, C. A., & Bounds, S. R. (2004, Nov). Rocket observations of structured upper hybrid waves at fuh = 2fce. *Geophys. Res. Lett.*, *31*(22). Retrieved from http://dx.doi.org/10.1029/2004GL021043