

All-sky imaging of meteor trails at 55.25 MHz with the first station of the Long Wavelength Array

J. F. Helmboldt,¹ S. W. Ellingson,² J. M. Hartman,³ T. J. W. Lazio,³ G. B. Taylor,^{4,5} T. L. Wilson,¹ C. N. Wolfe²

A new capability for high-sensitivity, all-sky monitoring of VHF meteor-trail reflections with the first station of the Long Wavelength Array, or “LWA1,” is described. LWA1 is a $\sim 100\text{m}$ -diameter HF/VHF array of 256 crossed-dipole antennas with a unique transient buffer mode that allows it to monitor for meteor trails via all-sky imaging with the same sensitivity as a single dish antenna $\gtrsim 40\text{m}$ in diameter. To demonstrate this capability, we have used a two-hour observing run conducted in August 2012 aimed at detecting and characterizing meteor-trail reflections of analog TV transmissions at 55.25 MHz. The analysis techniques described here allowed for a detection rate of $\sim 10,000$ trails per hour, including the detection of two meteor streams with radiants in the Aries/Perseus and Aquila/Hercules regions that were not previously reported in the literature. In addition, we have found a population of relatively long-duration (~ 1 to a few minutes), typically faint trails and have used high-resolution time series of the brightest trails to characterize decay times over a relatively large geographical area. Potential enhancements that could be enabled by the addition of more LWA stations are also discussed.

1. Introduction

Specular reflections of VHF signals off dense, transient structures in the lower ionosphere caused by meteors, or “meteor trails,” are a well-known phenomenon. These reflections allow for short-duration, over-the-horizon communications in the VHF regime [e.g., *Ellyett and Davies*, 1948; *McKinley*, 1961]. They also provide insights into the meteors themselves, as well as the details of how they interact with Earth’s upper atmosphere [e.g., *Elford*, 2004]. In addition, the effect of neutral winds in the mesosphere/lower thermosphere (MLT) region on meteor

trails has made radar-based observations of them one of the most reliable methods for probing the wind profile in this region of the atmosphere, which can be highly variable [e.g., *Li et al.*, 2012].

Consequently, there have been a large number of radar-based studies of meteor trails. These include monostatic and bistatic systems, with either dish antennas (LOTS OF REFS) or dipole-based arrays (LOTS OF REFS). Dish-antenna systems, such as the ARPA Long-range Tracking and Instrumentation Radar (ALTAIR), offer the ability to make pointed, high-sensitivity observations over a relatively small area of the sky (REFS). The same is true for phased arrays of large numbers of dipole antennas, such as the Jicamarca Radio Observatory, which are operated like a single-dish telescope, but are electronically steered. Such observations allow for the detection of trails associated with the least massive and/or slowest meteors (example REF), as well as observations of trails with remarkably high resolution in time, Doppler frequency, and/or range (example REFS). There are complementary systems that use a small number of dipole antennas like the Canadian Meteor Orbit Radar (CMOR) receiver arrays, which lack the sensitivity of dishes or large arrays, but can

¹US Naval Research Laboratory, Code 7213, Washington, DC 20375, USA.

²Bradley Dept. of Electrical & Computer Engineering, Virginia Tech, Blacksburg, VA 24061.

³Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA.

⁴Department of Physics and Astronomy, University of New Mexico, Albuquerque NM, 87131, USA.

⁵Greg Taylor is also an Adjunct Astronomer at the National Radio Astronomy Observatory.

monitor the entire sky for meteor trails, localizing each detection via interferometry (REFS).

This paper details a new and powerful asset within this field. The first station of the planned Long Wavelength Array (LWA), referred to as LWA1 [Taylor et al., 2012; Ellingson et al., 2013] is a HF/VHF array of 256 dipole antennas with excellent sensitivity in the range 20–80 MHz and an all-sky imaging capability. In addition to the number of antennas within the array, the antenna design, optimized for observations of cosmic sources over the entire sky, enhances LWA1’s sensitivity over similar dipole-based arrays [see Hicks et al., 2012]. However, what makes LWA1 unique as a receiver array for meteor trail observations is its transient buffer mode [see Sec. 2.1 and Ellingson et al., 2013], which allows one to beam-form anywhere on the sky after the fact. Thus, LWA1 can monitor the entire sky for meteor trails, but with sensitivity comparable to a single-dish receiving antenna. Here, we demonstrate this capability using transmitters of opportunity, nearby analog TV transmitters. The details of this all-sky capability are described in Sec. 2. Results from a two-hour observing run conducted in August 2012 are detailed in Sec. 3, including the detection of two meteor streams that appear to have never been observed before. In Sec. 4, we discuss these results and the prospects of future capabilities enabled by the development of additional LWA stations.

2. Observations and Data Processing

2.1. The LWA1 Observatory

LWA1 is intended to be the first of more than 50 stations for the full LWA. While the future of the larger LWA interferometer is uncertain, LWA1 was successfully completed in the fall of 2011 and is currently run as a fully operational, stand-alone observatory. The LWA1 observatory is an HF/VHF array of 256 crossed-dipole antennas arranged in a quasi-random configuration, spanning an area approximately 100 m in diameter. The antennas are optimized for the spectral range 20–80 MHz, but can be used to observe between 10 and 88 MHz. They are beveled downward at a roughly 45° angle to improve sensitivity at low elevations [see Hicks et al., 2012, for a more detailed description of the antennas].

The capabilities and operating modes of LWA1 are described in detail by Taylor et al. [2012] and Ellingson et al. [2013]. Briefly, the observatory has three

main modes, a beam-forming mode (DRX) and two all-sky, transient buffer modes. In DRX mode, a digital beam-former provides up to four separate beams on the sky with up to two different tunings, each with as much as about 16 MHz of usable bandwidth. The transient buffer modes record the output from each antenna in either a wide-band (TBW) or narrow-band (TBN) mode. In either mode, the antenna signals can be beam-formed after the fact, allowing for all-sky images to be made. TBW observations consists of 61-ms captures of the raw antenna outputs that give one access to the entire observable bandwidth. However, a single capture is nearly 10 GB of data and takes roughly 5 minutes to write to disk, implying that one cannot observe continuously in this mode. The complementary TBN mode allows one to observe continuously by tuning each antenna signal to a specified central frequency with a maximum sampling rate of 100 kps and usable bandwidth of about 70 kHz. Because meteor trails reflections can occur throughout the visible sky and can last longer than the 61-ms length of the TBW captures, TBN is the preferred mode for detecting and characterizing these reflections.

It is the combination of the available TBN mode and the antenna design that makes LWA1 a unique instrument for meteor trail observations. The effective collecting area of a single LWA dipole at 55.25 MHz is 4.57 m² [Hicks et al., 2012], nearly two times larger than what one expects from a simple dipole, i.e., $\lambda^2/(4\pi) = 2.34 \text{ m}^2$. This is a result of the structural design of the antennas as well as their active/impedance-matching systems and the ground-screens employed to reduce the influence of the earth below each antenna stand [again, see Hicks et al., 2012, for more details]. This optimized antenna design, coupled with the sheer number of antennas, gives LWA1 the collecting area of a dish antenna with a diameter of $\sim 40\text{m}$ or more at this frequency, comparable to the ALTAIR large-aperture radar (47m in diameter). However, unlike such a dish, LWA1’s TBN mode allows it to continuously monitor the entire sky for meteor activity.

2.2. Meteor Trails at 55.25 MHz

Similar to the Radio Meteor Survey, Extended System [RAMSES; Wislez, 1995] in Europe, LWA1 frequently detects relatively strong signals from meteor reflections of analog TV (NTSC) signals. In particular, the video carrier of the analog Channel 2 (55.25 MHz) appears to be consistently sensitive

to specular meteor-trail reflections. Fig. 1 shows a map of all the stations in the US and in the Canadian and Mexican border regions broadcasting at this frequency according to the FCC. Those stations for which the LWA1 zenith line of sight at 100km altitude is visible, 60 in all, are color-coded by effective radiated power (ERP).

Unlike smaller arrays, LWA1 can be used to make reasonably high-fidelity images of the entire visible sky, rather than relying on phase-fitting techniques to locate meteor trail reflections on the sky [e.g., *Holdsworth et al.*, 2004]. This also allows one to detect, locate, and characterize multiple reflections occurring at the same time, even if they have similar Doppler speeds. In addition, relatively large trails can be resolved within LWA1 images and their evolution can be tracked as they move in the sky. In addition, as stated above, the excellent sensitivity of LWA1 combined with its TBN mode allows for the detection of meteor trails that are too faint to see with typical all-sky monitoring arrays that use a much smaller number of antennas while covering a much larger area (i.e., the entire visible sky) than either dish antennas or phased arrays.

A data processing and imaging pipeline has been developed to work with LWA1 TBN data with a sampling rate of 100 ksp/s, tuned to 55.25 MHz for all-sky meteor trail imaging. The pipeline is discussed in detail in Appendix A. In short, because of the large volume of data involved, it is not practical to make all-sky images at a high temporal sampling rate. Instead, each image is made using 5.079s of data. The data are imaged within 32 frequency intervals, or “channels,” each 6.10-Hz wide (or, 33.1 m s^{-1} in Doppler speed).

Initially, a mean image is made over all channels to identify the brightest meteor trails. For each of these detections, the antenna data are used to generate high-resolution (5.12ms) time series by beamforming toward each source using a single 195.3-Hz wide channel (i.e., $1,060 \text{ m s}^{-1}$). These bright detections are also used to refine the calibration of antenna-based gains to improve the fidelity of images that are then produced for each of the 32 channels, yielding all-sky image cubes. This allows for the identification of much fainter trails that can be separated from extremely bright ones that substantially elevate the image noise, but only within one to a few channels. This form of Doppler discrimination allows for the number of detections to increase by a

factor of ~ 5 . For each significant detection within each image cube, the peak channel, sky position, flux density/intensity, and image RMS are recorded. This is all done separately for the X (north-south) and Y (east-west) polarization antennas.

Fig. 2 shows the number of detections per all-sky image pixel from a two-hour observing run conducted on 16 August 2012, starting at 06:38 UT. While the dominant source of detections are meteor trails, reflections from aircraft as well as the ground wave from the station XEPM in Juarez, Mexico are also observed. However, the most striking features are the two arc-like structures apparent in only the X polarization. These are likely indicative of meteor streams and will be discussed further in Sec. 3.3. In all, there were 22,542 and 21,776 detections at the X and Y polarizations, respectively. Among these, 4,937 (X-polarization) and 4,156 (Y-polarization) were bright trails for which high-resolution time series were generated.

The apparent meteor stream features seen within Fig. 2 offer an opportunity to illustrate the unique capability of LWA1 for meteor-trail monitoring. As described in Appendix A, the ideal image noise for an unresolved source for LWA1 at 55.25 MHz with 6.1 Hz of bandwidth and a 5.079s integration time is 2,230 Jy ($1 \text{ Jy} = 10^{-26} \text{ W m}^{-2} \text{ Hz}^{-1}$). For trails detected within channels with no extremely bright trails that increase the noise within the image, the measured RMS noise is consistent with this (see Appendix A). For a more typical, smaller array, the noise and resulting detection limit is much higher. For instance, the CMOR receiver arrays with 5 crossed-dipole antennas each. The specifications given for the three CMOR receiver arrays imply that they, like LWA1, would be sky-noise dominated near 55 MHz, implying a similar noise temperature per antenna (REF). However, because of the smaller number of antennas and collecting area per antenna (see Sec. 2.1), if one replaced LWA1 with a CMOR array, the noise would increase to 248,600 Jy. Fig. 3 shows all-sky maps of meteor counts if we were to do just that, plotting the counts for only those detections with intensities $> 5 \times 248,600 \text{ Jy}$ (i.e., we use a 5σ detection limit; see Appendix A). One can see immediately that the arc-like, meteor-stream features are not longer visible, due to the predominantly faint nature of the trails making up those streams.

We note that this is not really a fair one-to-one comparison as CMOR is essentially a monostatic system, which can cause a significant increase in re-

ceived power over a bistatic system, depending on the geometry of the observations. On the other hand, according to the FCC, there are at least 60, 55.25 MHz transmitters available to LWA1 for these observations, 5 with ERPs larger than that of the 12 kW CMOR transmitter, ranging from 16 to 100 kW (see Fig. 1). While CMOR also has lower-frequency bands that are more sensitive to meteor trail reflections, the receiver noise increases substantially with decreasing frequency, which is true for any sky-noise dominated system in the VHF regime (i.e., $T_{sky} \sim \nu^{-2.55}$). In addition, the imaging and deconvolution techniques described in Appendix A, made possible by a well-filled synthetic aperture, greatly improve the system detection limits for LWA1 and the trail counts near the horizon are likely larger as the LWA antennas are optimized to increase sensitivity at lower elevations relative to dipoles like those used with the CMOR array which are horizontally mounted. However, the use of a true radar transmitter and larger duty cycle allows CMOR an extra level of discrimination (i.e., range and time) to detect faint trails in the presence of bright ones. Despite these considerations, the stark difference between Fig. 2 and Fig. 3 provides a powerful illustration of the utility of LWA1 as a sensitive all-sky meteor monitor, capable of detecting streams that might otherwise be too faint for other all-sky receiver arrays to see.

Examples of additional data products yielded by the pipeline for the brightest trails are shown for a single 5.079s interval (Y polarization) in Fig. 4. Within this figure, the mean image over all 32 channels is shown in the upper left panel. This is shown in an l, m projection where l and m are direction cosines defined in this case as $l = \cos e \sin a$ and $m = \cos e \cos a$ where e is the elevation and a is the azimuth measured clockwise from north. The high-resolution, amplitude time series for the four significant detections within the image are plotted to the right. From these, one can see that the object in the lower right of the image is the ground wave from XEPM, as it shows no evidence of a typical meteor trail time series. The southern-most and northern-most objects are short-duration trails, while the object in the upper left of the image appears to be a relatively long-duration trail with a complex decay pattern, indicative of Fresnel oscillations.

The average dynamic spectrum over all antennas is also shown in the lower left of Fig. 4. This shows the mean power as a function of Doppler speed and

time with a temporal resolution of 0.164 seconds. Here, one can see the constant signal from the XEPM ground wave near zero Doppler speed. One can also see evidence of somewhat complex velocity structure associated with the longer-duration meteor trail together with some aliasing at intervals of roughly 325 m s^{-1} caused by the pulse used to mark the end of each video frame of the NTSC TV signal at a rate of $\sim 60 \text{ Hz}$. This complex structure with multiple peaks could indicate real structure within the trail, or could be the result of scattering of signals from multiple transmitters.

3. Meteor Trail Analysis

3.1. Correlation Functions and Long-duration Trails

One way to exploit the excellent all-sky sensitivity of LWA1 is via a statistical description of the spatial and temporal extent of the relatively large number of detected meteor trails. The sensitivity of LWA1 as a receiving array allows one to do this within a single observing run rather than requiring several days of meteor counts. For this analysis, we have chosen to compute the two-dimensional correlation function, ζ , among trail detections as a function of angular and temporal separation, which essentially provides an estimate of the degree to which sources are cluster, in this case in time and on the sky. For this computation, we limited ourselves to those trails above an elevation of 30° where a plane-parallel approximation is valid, making the interpretation of the results more straightforward. Over the entire two-hour run, there were 5,120 detections at X-polarization above this elevation and 5,529 at Y-polarization. We note that this elevation requirement eliminates detections of the XEPM ground wave and, in this instance, all detections of reflections off aircraft.

The correlation function was computed at 51^2 pairs of temporal separation, Δt , and angular separation, $\Delta\theta$, according to

$$\zeta(\Delta t, \Delta\theta) = \frac{DD(\Delta t, \Delta\theta)}{\langle RR \rangle(\Delta t, \Delta\theta)} - 1 \quad (1)$$

where DD is the number of data-data pairs separated in time by Δt and on the sky by $\Delta\theta$, and $\langle RR \rangle$ is the number of random-random pairs within a simulated data-set of randomly generated trails with the same number of trails as the actual data used, averaged over several random data-sets. In other words, ζ provides a quantitative estimate of how clustered

trails tend to be on the sky and in time relative to what one would expect for a completely uncorrelated set of trails.

The random catalogs of trails were generated starting with a population of 2×10^6 trails uniformly/randomly distributed within a thin layer at 100km altitude. To incorporate the effect of the LWA antenna response, the trails within the random catalog were first put into 100 zenith angle bins from 0–90°. Within each of these bins, $\lfloor N \cos(Z)^{1.6} \rfloor$ trails were randomly selected, where N is the total number within the bin and Z is the bin zenith angle. The factor of $\cos(Z)^{1.6}$ closely approximates the LWA antenna response for elevations above 30° [Dowell, 2011], and thus provides a means to simulate the effect of this response on the observed distribution of trails if they were indeed randomly distributed within a thin layer. Among the remaining simulated trails, a number equal to that contained within the actual data-set used to compute DD was randomly selected. We found that the zenith angle distributions for these randomly-generated catalogs closely resembled what was observed for the actual trail data, implying that other factors such as image sensitivity and the inherent distribution of trail intensities/line densities did not significantly impact the observed distribution of trails on the sky. Following this, for each random catalog, the trails were randomly assigned to one of the 5.079s imaging intervals (1361 in all) from the observing run.

For each polarization, 10 random catalogs were used to compute $\langle RR \rangle$ and its associated uncertainty, i.e., the standard deviation with each $\Delta t, \Delta \theta$ bin among the 10 catalogs. The error in DD was assumed to be simple shot noise, or \sqrt{DD} . These quantities and their uncertainties were used to compute ζ and its 1σ error within each of the 51^2 $\Delta t, \Delta \theta$ bins for $\Delta t < 300$ s and $\Delta \theta < 23^\circ$.

We have displayed the results in Fig. 5, plotting ζ as a function of $\Delta \theta$ within each Δt bin. We found that for $\Delta t \gtrsim 90$ s, the correlation functions were essentially noise and are not plotted. One can see that, as might be expected, the correlation is quite strong at the smallest angular separations for the shortest time scales, indicative of relatively short-duration trails. There is some indication of significant correlation on an angular scale of 5° , just larger than the full width at half maximum of the LWA1 beam at 55.25 MHz, 2.9° . This may point to a significant population of trails that are marginally spatially re-

solved. Moving to larger Δt bins, one can see the peak of the correlation function moves to larger values of $\Delta \theta$. These peaks are about an order of magnitude smaller than that seen for the smallest Δt bin, but are nonetheless quite significant. They indicate the presence of a small population of relatively long-duration trails, lasting up to nearly 80s with implied angular speeds between roughly 0.02 – 0.05 $^\circ/\text{s}$. For the typical elevations within the sub-sample used here and an altitude of 100km, this corresponds to transverse speeds of 0.05 – 0.1 km s^{-1} . This is comparable with the typical radial speeds implied by the channels within which the meteor trails within the sub-sample used here were detected, roughly 0.1 km s^{-1} .

Spurred by the results shown in Fig. 5, we searched the data for evidence of individual long-duration trails. We did this by plotting each of the direction cosines versus time for all detected trails, looking for contiguous, or nearly contiguous groups of data, indicating a relatively long-duration trail “streaking” across the sky. Four examples of such trails are shown in Fig. 6. Each panel shows an image of the peak signal-to-noise ratio, S/N, over the duration of each of these trails, zoomed in near the trail itself. Over-plotted are the positions of detections from individual 5.079s images, color-coded by time (UT).

The first example (upper left panel) has the main trail displayed in the upper center part of the frame to highlight the fact that there appears to be one or more additional and relatively faint long-duration trails to the south. The main trail is faint itself when compared to the many shorter-duration trails detected around it, but is clearly distinct, evidenced by its clear propagation path toward slightly west of south, covering nearly 20° in 5.8 minutes (0.06 $^\circ/\text{s}$). The second example (upper right) shows a much brighter trail that lasts for about 2 minutes, traveling about 3° due north in the process. The final two examples show two relatively similar, somewhat faint trails lasting about 5m (bottom left) and 3m (bottom right). The first is slower, moving toward the southwest at roughly 0.04 $^\circ/\text{s}$. The final example trail moved at nearly 0.1 $^\circ/\text{s}$ toward the northwest.

(COMPARISON WITH BOURDILLON ET AL. PAPER)

3.2. Bright Trail Decay Times

The decay time of a meteor trail can provide useful information about the meteor itself, as well as

the surrounding atmosphere. For those trails with relatively smooth, exponential decays, the e-folding time, τ , is related to both the meteor mass and the coefficient of ambipolar diffusion, D_A (REFS). However, conditions above XX km typically cause the exponential decay of a meteor trail’s light curve to be significantly disrupted, owing to wind-driven dynamics within that region, and τ can be completely uncorrelated with D_A . Therefore, it is advantageous to isolate those meteor trails within our LWA1 data that have relatively smooth exponential decays so that they can potentially be used to constrain, for example, the properties of the MLT region over a relatively large area.

To do this, we started with the high-resolution time series generated for the 4,000–5,000 bright trails, discussed in Sec. 2.2. For each of these time series, we used a non-linear least-squares fitting routine to fit the following function

$$A(t) = \begin{cases} A'_p e^{-u/(t-t_0)} e^{-(t-t_0)/\tau} + A_N, & \text{if } t > t_0 \\ A_N, & \text{if } t \leq t_0 \end{cases} \quad (2)$$

where A_N is the noise floor of the amplitude time series, t_0 is the time when the trail signature starts to appear, u is the time scale for the initial (typically rapid) increase in amplitude, $A'_p = A_p \exp(2\sqrt{u/\tau})$, and A_p is the peak amplitude. This function was taken from (REF) who developed it as a way to characterize the duration of gamma-ray bursts, which have similar temporal profiles.

Examining many time series and their best fitting versions of equation (2), we found that the bright detections could be put into three categories: (1) those that were not meteor trails, but either ground waves from XEPM or reflections from airplanes, or “non-trails” (2) trails with smooth exponential decays, and (3) trails with substantial deviations from smooth decays, or “Fresnel trails.”

These are typically easily identifiable by eye, but this is impractical for the data-set used here with 4,000–5,000 bright detections from each polarization. To categorize detections automatically, we developed a simple algorithm. First, any time series for which the difference between the minimum and maximum values of the best model fit was less than five times the median absolute deviation (MAD) of the fit residuals was classified as a non-trail.

Next, for actual trails, a high-order polynomial fit applied to the log (base ten) of the ratio of the amplitude profile to the model fit was used to charac-

terize the amount of structure due to relatively long temporal-scale fluctuations, and not noise or small-scale structure inherent to the TV signal, which would indicate Fresnel-like oscillations. Through trial and error, we found that a 5th order polynomial worked well for this purpose. We also found that flagging trails where the maximum absolute value of this polynomial fit was more than three times the MAD of its residuals identified all of the Fresnel trails identified by eye using 245 test cases. The reliability of this was greatly improved by isolating the polynomial fit to the 1.27s period (i.e., one fourth the full 5.079s) closest to the fitted value for t_0 . In some instances, trails that looked like exponential trails were also flagged by this criterion, but in general, it provides a conservative estimate of the population of true exponential profiles among the larger sample. Examples of detections put into each of the three groups by this algorithm are shown in Fig. 7 where the amplitude temporal profiles are plotted along with their model fits. In all, 1,666 and 1,035 bright trails were identified as exponential trails by this algorithm at X and Y polarizations, respectively.

Using the fitted values for τ for exponential trails only from both polarizations, we computed the median value of τ with $1^\circ \times 1^\circ$ bins in latitude and longitude over the entire two-hour observing run. These positions were computed assuming a reflection altitude of 100km and full spherical geometry. Only those bins with at least 10 exponential trails were used. A map of these τ values is shown in the upper panel of Fig. 8. There is obvious structure within this plot, with τ increasing toward the southwest. However, much of this structure may be due simply to the geometry of the observations. For the case of forward scattering, τ is not only inversely proportional to D_A , but also proportional to $\sec^2 \phi$, where ϕ is half the angle between the vector pointing from the receiver to the trail and that pointing from the transmitter. Since we cannot know a priori which transmitters are being reflected by which trails (some may even be reflecting multiple transmitters), we cannot explicitly correct the τ map shown in Fig. 8 for this effect. However, assuming that the broadcasting habits of the TV stations involved remain consistent, this geometrical effect should be constant and can be empirically determined using numerous, similar observing runs, perhaps utilizing a principal component analysis.

In the lower panel of Fig. 8, we have also plotted the median value of τ among exponential trails as a

function of UT, within bins that are 5m wide. This plot shows the typical τ to be about 0.05s with significant fluctuations, most notably a 15–20m “dip” near 07 UT and a 20m “plateau” starting at 08 UT. An examination of the latitude/longitude coverage of the exponential trails within each 5m bin implied that these features are not the result of biases toward any particular geographical location where τ may be different due simply to observational geometry.

3.3. Meteor Streams

As mentioned in Sec. 2.2, there are two features apparent within the all-sky meteor count maps from the 16 August 2012 observing run (see Fig. 2) that are consistent with what one expects for meteor streams. Both are only noticeable within the X-polarization data, which is likely indicative of the locations of the radiants and the transmitters being reflected. Scattering by meteor trails (REF) and more generally, any case of total internal reflection tends to favor the polarization component that is perpendicular to the trail/medium boundary. Since the monopole transmitters used are typically horizontally polarized, one would then expect stronger detections at north-south (X) polarization from a meteor stream if the radiant and the transmitter were both roughly due east or west.

To better constrain both the radiants and the likely transmitter(s) involved, we did the following computation. The arc-like pattern associated with radar detections of meteor streams represents the region on the sky where specular reflection is most likely. For a monostatic system, this is an arc where the trails within the stream, all moving with roughly the same orientation, are perpendicular to the line of sight. The fact that the trails are confined to a relatively thin layer is what creates the arc-like structure (i.e., rather than a full plane). The situation with a bistatic system is only slightly more complicated. In this case, specular reflection occurs when the trail is perpendicular to the normal vector required for reflection between the transmitter and the receiver according to the law of reflection.

For each stream, we first isolated the trails likely associated with it by eye with a polynomial region of interest. We then varied the location of the radiant and its associated unit vector, \hat{s} , over the same all-sky image grid used within the imaging process, using each of the transmitters from the FCC list. For each transmitter, stream, and \hat{s} , we computed

the normal vector required for specular reflection, \hat{n} , assuming an altitude of 100km. We used this to compute the magnitude of the dot product $|\hat{s} \cdot \hat{n}|$, which is zero for specular reflection, at each trail location within the stream. We then computed the χ^2 difference between this and the expected value (zero) for each radiant location, mapped over the whole sky for each transmitter and stream. Each of these χ^2 maps was then used to identify the minimum χ^2 and best-fitting radiant location.

For the longer of the two streams, we found that there were eight transmitters with similar minimum χ^2 values and similar best-fitting radiants. The quality of the fit for each transmitter was confirmed by making an all-sky map of $|\hat{s} \cdot \hat{n}|$ using the best-fitting radiant and comparing by eye the shape of the locus where $|\hat{s} \cdot \hat{n}| = 0$ to the stream itself. This was also done for a few other transmitters with slightly higher χ^2 minimums to verify that their fits were unsatisfactory. Examples of this are shown in Fig. 9 for three of the eight best-fitting transmitters with the highest ERPs. Note that these are also the only three among the eight with $\text{ERP} > 1$ kW. We repeated this procedure for the second, shorter stream for which there were 18 transmitters that seemed to provide equally good fits. In this case, there were likewise only three transmitters stronger than 1 kW, and the results for these three are shown in Fig. 9 as well.

Using these best-fit cases, we computed estimates of the radiant locations for both streams. The longer trail has a radiant almost due east and all of the best-fitting transmitters are east of LWA1, consistent with what was expected for a pure X-polarization detection. Similarly, the shorter stream has a radiant to the west and the best-fitting transmitters are west of LWA1. Accounting for precession and time of observation using the middle of the observing run, 07:35 UT, the J2000 coordinates for the two radiants are at $\alpha = 2.6 \pm 0.2^{\text{h}}$, $\delta = 30.0 \pm 0.9^{\circ}$ for the longer stream and $\alpha = 18.7 \pm 0.2^{\text{h}}$, $\delta = 11 \pm 1^{\circ}$ for the shorter stream. These locations are both significantly different from the estimated radiant for the Perseid shower of $\alpha = 3.1^{\text{h}}$, $\delta = 58^{\circ}$ (shown as a “P” in the panels of Fig. 9), which is purported to be active until roughly 20 August. The radiant of the longer stream is close to the boundary between Aries and Perseus; the second radiant is near the Aquila/Hercules border region. A search of the relevant literature revealed no mention of streams/showers emanating from these regions of the sky during this time of year.

These two streams provide a unique opportunity to computed estimates of the line density distributions of their meteor trails. Because the transmitters involved and trail orientations are well constrained, assuming a particular altitude, we can use the detected intensity of each trail to estimate the line density, q , according to

$$I = q^2 \frac{\tau}{2t_{int}} \frac{g_r \lambda^3 r_e^2}{32\pi^2 |\vec{r}_r|} \times \sum_{t=1}^{N_T} \frac{P_t g_t \sin^2 \gamma_t}{|\vec{r}_t| (|\vec{r}_t| + |\vec{r}_r|) (1 - \sin^2 \phi_t \cos^2 \beta_t)} \quad (3)$$

Here, I is the trail intensity, t_{int} is the integration time (5.079s), τ is the *amplitude* decay time (see Sec. 3.2), λ is the observing wavelength, r_e the classical electron radius, \vec{r}_r is the vector pointing from the receiver to the trail, and g_r is the receiver gain computed using the flux density calibration described in Appendix A and the LWA dipole response [$\sim \cos(Z)$ ^{1.6}; *Dowell, 2011*]. The factor of $\tau/(2t_{int})$ accounts for the fact that the observed intensity is diluted by the 5.079s integration time. In other words, the received power for an exponentially decaying trail is $P = P_{max} \exp(-2t/\tau)$, so that the mean power integrated over t_{int} is $P_{max} \tau / (2t_{int})$ for $t_{int} \gg \tau$. Since we do not have high-resolution time series for these trails, we simply assume the median value of $\tau = 0.05$ s obtained from the model fits to the bright-trail time series presented in Sec. 3.2 for all trails within the streams.

For transmitter t (out of N_T total), \vec{r}_t is the vector pointing from the transmitter to the trail, P_t is the transmitted power, g_t is the transmitter antenna gain, ϕ_t is half the angle between \vec{r}_r and \vec{r}_t , β_t is the angle between the trail and the plane occupied by both \vec{r}_r and \vec{r}_t , and γ_t is the angle between the transmitted electric field vector and \vec{r}_r at the reflection point [*Wislez, 1995*]. These angles are labeled in a schematic shown in Fig. 10. The transmitter antenna gains were computed assuming a simple monopole antenna, i.e., $g_t = \cos^2[\cos(\theta)\pi/2] / \sin^2(\theta)$, where θ is the zenith angle (from the transmitter's point of view) and $g_t = 0$ for $\theta > 90^\circ$. Since all of the transmitters are horizontally polarized, γ_t was computed assuming the electric field vector is parallel to $\vec{r}_t \times \vec{R}_t$, where \vec{R}_t is the vector pointing from the center of the earth to the transmitter location.

Line densities were computed separately for trails within each of the two streams. The q distributions measure for each are shown in Fig. 11, each normalized by the total number of trails. Both distribu-

tions turn over below a particular line density, likely owing to incompleteness due to chiefly to variability in image sensitivity/noise. The total power available from the western transmitters that illuminated the Aquila/Hercules trail is significantly larger than those associated with the Aries/Perseus stream, and thus, its q distribution turns over at a lower line density.

It is possible to perform a kind of completeness correction to these distributions by dividing each bin by the number of “detectable” trails for that bin, rather than the total number of trails. By detectable, we mean a trail whose image noise and sky position yield a minimum detectable value of q , or q_{min} , that is small enough such that a trail within the q bin in question would also be detectable within the same image and at the same sky location. Because of the finite widths of the q bins, one needs to know the shape of the underlying q distribution to do this properly. Alternatively, one can compute upper and lower limits to the completeness-corrected distribution by normalizing each bin by the number of trails with q_{min} less than the lower boundary of the bin (upper limit) and then by the number with q_{min} less than the upper boundary (lower limit). These upper and lower limits are shown as shaded regions for the two streams in Fig. 11 with the 1σ errors, computed assuming simple shot noise, added to the upper limits and subtracted from the lower limits.

From the completeness-corrected distributions, one can see that the entire distribution for the Aquila/Hercules stream as well as that for the Aries/Perseus stream for $q > 10^{12.2} \text{ cm}^{-1}$ are consistent with a single power-law. A linear fit to these implies that for a line density distribution of the form $dN/dq \sim q^{-s}$, $s = 2.02 \pm 0.07$. This is consistent with other radar-based constraints on the shape of meteor-trail line density/mass distributions (REFS). The Aries/Perseus stream shows a clear break near $q \sim 10^{12.1} \text{ cm}^{-1}$, which is similar to the limit usually adopted for the boundary between under-dense and over-dense trails ($\sim 10^{12} \text{ cm}^{-1}$). Below this limit, the completeness-corrected Aries/Perseus stream distribution follows a different power-law that implies $s = 1.84 \pm 0.13$. This is shallower than the power-law slope for the over-dense trails and all the Aquila/Hercules trails, but only marginally so (they differ by $\sim 1.2\sigma$).

4. Discussion

We have demonstrated the unique ability of LWA1 to provide high-sensitivity, all-sky monitoring of meteor trail reflections in the VHF regime. The large collecting area of LWA1, coupled with its TBN mode offer the “best of both worlds,” providing the sensitivity of a $\gtrsim 40\text{m}$ -diameter dish (at 55 MHz) with the all-sky capability of a dipole-based array. Using transmitters of opportunity, analog TV broadcasts at 55.25 MHz, we have shown the potential of LWA1 as a meteor radar receiver. With this multi-static system and an automated processing pipeline, we have shown that $\sim 10,000$ trails can be detected per hour. A statistical analysis of the detected trails revealed a significant population of relatively long-duration trails, lasting ~ 1 to a few minutes. Many of these are relatively faint with portions of them at or near the detection thresholds of their respective images. (MORE DETAILED COMPARISON WITH BOURDILLON ET AL. PAPER).

The brightest trails were used to determine exponential decay times with high-resolution time series (5.12ms sampling), yielding a relatively large map, $10^\circ \times 7^\circ$ in longitude and latitude, of τ with just under two hours of data. With many additional, similar observing runs, we will be able to correct such maps for repeatable effects caused by the geometry of the multistatic system, yielding a series of maps of the dominant physical contribution to τ , the coefficient of ambipolar diffusion, D_A . This in turn provides information about the state of the MLT region that may be incorporated in assimilative models of the upper atmosphere (e.g., NOGAPS-ALPHA; REF).

We have also found evidence within our example observing run conducted on 16 August 2012 of two meteor streams. The streams were relatively faint, and likely would not have been detected with a smaller, less sensitive array (see Fig. 2 and 3). Our best estimates for the radiant for these streams are in the boundary regions between Aries and Perseus and between Aquila and Hercules. Both are substantially different from the radiant for the Perseid shower, which should have still been active at the time of the observations. There appears to be no mention of these two showers in the literature, implying that this may be the first time either has ever been observed. The line density distributions for these streams are generally consistent with those observed for other streams, with mass in-

dexes ~ 2 (REFS). However, the distribution for the Aries/Perseus stream shows a clear break at the approximate boundary between over- and under-dense trails, with the under-dense trails have a somewhat shallower mass index. These results further highlight the utility of LWA1 as a meteor radar receiver and as a potentially powerful asset in the search for previously unknown streams that are relatively difficult to detect.

These results make the development of new LWA stations a tantalizing proposition for future meteor trail studies. First, with multiple stations observing the same trail from different locations, parallax can be used to obtain a distance estimate. In the event that a dedicated radar transmitter is implemented, multiple stations could be used to accurately map out the orbits of detected trails, similar to what is done with the CMOR arrays, but with vastly superior sensitivity. With such a transmitter, one could also use an array of LWA stations to map the wind profile of the MLT region over a relatively large geographical area and, with potentially > 100 detections per minute, on relatively short time scales. Interferometry among multiple stations may also be used to constrain the physical extent of meteor trails and, depending on the configuration of the stations, to even produce high-angular resolution images of such trails.

Appendix A: Data Processing and Imaging Pipeline

As stated in Sec. 2, for the 55.25-MHz meteor-trail observations, LWA1 was operated in TBN mode, that is, the data from each individual antenna was recorded at a sampling rate of 100 kbps after being tuned to a central frequency of 55.25 MHz. While it is possible to make all-sky images with such data by treating LWA1 as a normal phased array and beamforming the TBN data over the entire sky, doing so at a rate of 10^5 images every second is computationally impractical. Even filtering the data to a narrower band (see below) only marginally reduces the computational burden.

A more practical way to handle this large data volume is to use LWA1 as a “multiplying” interferometer rather than a phased array, or “adding” interferometer. The former uses correlations of signals from pairs of antennas, or “baselines,” averaged over a chosen time interval to produce images, rather than

the complex voltages themselves. In other words, a “visibility” for the baseline consisting of antennas i and j is given by $V_{i,j} = \langle \varepsilon_i \varepsilon_j^* \rangle$, where ε is the complex voltage and $\langle \rangle$ denotes an average over time. It is this time averaging that makes the use of visibilities potentially much more computationally efficient than beam-forming voltages because one image can be made for an entire time interval instead of having to generate a total-power, beam-formed image for each temporal sample, and then averaging. In other words, the process of correlating voltages over a fix time period and then imaging them can be much faster than constructing many beam-formed images over that same time period.

If the source(s) being imaged is(are) far enough away relative to the size of the longest baseline, the far-field approximation often used in radio astronomy can be applied to relate the intensity on the sky to the measured visibilities according to

$$V(u, v, w) = \int \int I(l, m) e^{-2\pi i(ul + vm + w\sqrt{1-l^2-m^2})} d\Omega \quad (\text{A1})$$

where Ω denotes solid angle, l and m are direction cosines, and u , v , and w are the normalized baseline coordinates, i.e., the difference in antenna positions toward the east, north, and observing field center, respectively, in units of wavelengths [Thompson *et al.*, 1991]. Since LWA1 is roughly 100m in diameter and the meteor trails being imaged are at altitudes $> 80\text{km}$, this far field approximation is valid. The LWA1 antennas lie nearly in a plane, making $w \approx 0$ when the field center is set to the local zenith. This means that an image can be generated by computing visibilities, gridding them in the u,v -plane, and performing a two-dimensional (inverse) fast Fourier transform (FFT). Here, the integration time used for computing the visibilities was set to roughly 5 seconds. This was chosen to be similar to that used by the LWA1 Prototype All-Sky Imager [PASI; Ellingson *et al.*, 2013]. The PASI integration time was chosen to be short enough that the cosmic sources being imaged would have moved an insignificant amount in the sky compared with the size of the LWA1 beam. Because these cosmic sources must be subtracted from the 55.25 MHz data when imaging meteor trail reflections, this seemed an appropriate choice for the meteor-trail imaging pipeline. While this can have the effect of diluting the intensity of short-duration trails ($< 1\text{s}$), the large volume of data necessitates the use of as long an integration time as possible to

reduce the amount of imaging required. For instance, the pipeline describe below, using a reasonably well-equipped desktop computer, takes ~ 3 minutes per image with a very weak dependence on the choice of integration time.

In practice, images made with this visibility/FFT-based approach are convolved with the Fourier transforms of both the sampling function in the u,v -plane, i.e., the point spread function (PSF), and the products of the antennas’ complex gains. For the LWA1 data, the effect of the complex gains was mitigated by applying the calibration derived by Ellingson *et al.* [2013] to the complex voltages before correlating them to produce visibilities. This calibration was obtained and applied using python-based packages made available to the public as the LWA Software Library [Dowell *et al.*, 2012]. This calibration, however, is not perfect and small, short time-scale variations in antenna gains also exist, both of which affect image fidelity. In addition, even with a perfect calibration, the first sidelobes of the LWA1 PSF are roughly -18 dB below the peak, making confusion from sidelobes from bright sources a limiting factor in detecting faint ones.

Often times, the technique used to grid the visibilities also has a significant effect on image fidelity, especially for sparse arrays. In these cases, a gridding function designed to mitigate aliasing is often employed [see, e.g., Thompson *et al.*, 1991]. However, the configuration of LWA1 was optimized for straightforward beam-forming, and LWA1 is consequently far from sparse, having a reasonably well-behaved beam with no aliased versions of the main lobe, even beyond the approximate Nyquist limit. We have therefore used a simple gridding function that just sums the visibilities within each u,v grid cell (i.e., “natural weighting”) that gives the closest approximation to beam-forming that is possible within the FFT-based approach.

To reduce the effects of the PSF and antenna-based gain errors, standard imaging deconvolution techniques were applied. These consist of applications of the CLEAN algorithm followed by self-calibration, both of which are explained in detail by Cornwell *et al.* [1999] and Cornwell and Fomalont [1999], respectively. In short, CLEAN approximates the image with a series of point sources convolved with the PSF. This is done by identifying the pixel with the largest absolute value, shifting the PSF to be centered on that pixel, subtracting a scaled version of the shifted PSF from the image (this scale fac-

tor is called the CLEAN “gain”), and then repeating the process until the actual value of the pixel with the largest absolute value is negative or a maximum number of iterations is reached. The delta functions, or “CLEAN components” can then be restored to the residual image, convolving them with a Gaussian beam with a width similar to the main lobe of the PSF.

The resulting set of CLEAN components is then used to compute model visibilities which are used to fit for the complex antenna gains, usually using a gradient search, which is referred to as self-calibration. Since there are $N_A(N_A - 1)/2$ baselines for an interferometer with N_A antennas, this is generally an over-constrained problem, especially for LWA1 where 32,640 baselines are used to solve for 256 complex gains. Here, the gradient-search method employed within the software package DIFMAP [Shepherd *et al.*, 1995] was adapted to perform self-calibration within the pipeline. Only the phases of the complex gains were solved for (i.e., “phase-only” self-calibration) because when both the amplitudes and phases are solved for, there is a tendency to “fit the noise,” that is, to compensate for gain fluctuations due to noise as though they were real changes in antenna gain. This is especially true when such a large number of baselines is used.

The steps included in the pipeline were applied separately to each polarization (X/north-south and Y/east-west) as follows:

1. *Initial Processing:* Read in 507,904 TBN samples from each antenna (5.07904 seconds of data) in 31 groups of 16,384 samples. For each group, filter the data by multiplying by a Hamming window and applying an FFT, keeping only the inner 121 channels (i.e., ± 2000 m s⁻¹ in Doppler speed). For each 5.079 interval, this yields a complex spectrum of voltages for each antenna with 121 channels and 31 time steps. During this, a set of 992 single-channel voltages is also produced for each antenna by performing an FFT (with Hamming smoothing) for every 512 samples (5.12ms) and keeping only the central 195.2-Hz wide channel to be used later in step 4.

2. *Correlation:* Correlate the complex voltage spectra to produce a $32,640 \times 121$ array of visibilities. Subtract the mean of the first 32 and last 32 channels from each visibility spectrum to remove the contribution of cosmic sources, which have relatively flat, continuum spectra across the full band of 121 channels (738 Hz).

3. *Initial Imaging of Bright Sources:* Average the continuum-subtracted visibilities over the innermost 32 channels, grid them in the u,v -plane, and perform FFT-based imaging. Perform two rounds of CLEAN and phase-only self-calibration followed by a final round CLEAN to produce an integrated, CLEANed image. Here, a relatively large CLEAN gain of 0.5 is used with a maximum of 100 iterations, producing a relatively shallow application of CLEAN. This is because the purpose of this step is to identify the brightest sources in the field for (1) refining the antenna gains via self-calibration and (2) generating high-resolution time series for the brightest trails using the 195.2-Hz bandwidth voltages generated in Step 1.

4. *High-resolution Time Series:* Locate significant detections of sources within the final image produced in Step 3 by identifying peaks within the noise-free, restored CLEAN image that are more than five times the RMS of the difference image (i.e., the image with the CLEAN components, convolved with the PSF, subtracted). Beam-form the 195.2-Hz bandwidth voltages generated in Step 1 toward each of these sources to generate time series with 5.12ms sampling.

5. *Full Image Cubes:* Apply the gain corrections determined in Step 3 via self-calibration with bright sources to the visibilities within each of the innermost 32 channels. Separately image each of these channels with a deeper application of CLEAN, specifically, a gain of 0.05 and a maximum of 1,000 iterations.

6. *Final Source Identification:* Normalize each channel within the image cube by the RMS of its difference image. The noise level can vary substantially from channel to channel because extremely bright sources, which may only be present in 1–3 channels, can artificially raise the noise level, making such images dynamic-range limited. For example, we have found that the maximum achievable image signal-to-noise (S/N) with our pipeline is about 200, regardless of the peak intensity. After applying this normalization, construct a two-dimensional image of the peak S/N over all channels. Identify peaks within this maximum S/N image that are (1) larger than five (i.e., $> 5\sigma$ detection) and (2) associated with at least one CLEAN component from within the image cube to minimize false detections. Finally, for each of these detected sources, identify the channel where its S/N is the largest using the normalized image cube.

This pipeline was run on a two-hour observing run conducted on 16 August 2012, starting at 06:38 UT.

The number of detections per image pixel from the entire run are shown separately for each polarization as images in Fig. 2. An example image of bright sources and their associated high-resolution time series from a single 5.079s period are shown in Fig. 4. Both of these figures are discussed in Sec. 2.2. To illustrate the benefit of using Doppler discrimination to identify relatively faint sources in the presence of brighter ones (i.e., Step 6 above), we show in Fig. 12 images produced in Step 3 for the first 5.079s of the 16 August 2012 observing run for X and Y polarizations as well as maximum S/N images produced in Step 6 for comparison. One can see that the maximum S/N images reveal numerous sources not visible within the images produced by averaging over all channels. This is why the source detection based on the full normalized image cubes made within Steps 5–6 was able to identify roughly five times as many potential trails as that based on the mean images generated in Step 3 (see Sec. 2.2).

To better assess the sensitivity of the LWA1 observations and to allow for meteor trail line densities to be computed in special cases (see Sec. 3.3), we derived a flux density calibration for the image cubes and the individual sources detected within them. We did this using a strong cosmic source, Cygnus A (or, “Cyg A”) with a well-known spectrum within the VHF regime. We did this using roughly the first five minutes of data from the 16 August 2012 observing run. Within this window of time, the data were correlated within each 5.079s period (59 in all), using 6.1-Hz channels as was done within the main pipeline. However, to maximize the contribution of Cyg A and minimize that of meteor trails, the inner 60 kHz of the band was used. This amounts to nearly 10,000 channels, and with over 30,000 baselines, it is not practical to construct a full two-dimensional array of visibilities of this size to be imaged. Instead, as each baseline was correlated, the resulting visibilities were median combined into a single channel, thus minimizing the impact of the extremely narrow-band meteor-trail signals.

The median image among the 59 all-sky images produced in this way is shown in the upper panel of Fig. 13. Within this, one can clearly see Cyg A as well as Cassiopeia A (Cas A), the Galactic Center (GC), and the Galactic Plane. Note that while similar in intensity to Cyg A, the flux density of Cas A is somewhat variable at these frequencies [see *Helmboldt and Kassim*, 2009, and references therein] and

not as reliable for calibration purposes. To derive a flux density calibration, the peak intensity near Cyg A was measured on each image. We then estimated the contribution from the Galactic background using a second-order polynomial fit to an annular region centered on Cyg A to account for the non-uniform background caused by Cyg A’s close proximity to the Galactic Plane. We then adjusted the expected flux density of Cyg A at 55.25 MHz, 20,474 Jy [*Baars et al.*, 1977] for the LWA antenna response at the zenith distance of Cyg A at the time of observation, 14.38°, amounting to an expected flux density of 19,478 Jy. This allowed us to convert the arbitrary intensity units produced by our pipeline to real flux density units. Based on the variation from image to image in the intensity of Cyg A and the uncertainty in the polynomial fit used to derive the background intensity, we estimate this calibration has an uncertainty of about 3%.

As a kind of sanity check, we assembled the distributions, one for each polarization, of the calibrated image RMS values for all sources detected from the normalized image cubes generated in Step 6 above and compared them with the theoretical noise limit for an LWA1 image made at 55.25 MHz with our imaging parameters. As noted above, many of the images within the image cubes are dynamic range limited due to relatively bright sources within the image, with a maximum achievable S/N of about 200. Such images will have RMS values well above the theoretical limit. Thus, we expect these distribution to have large tails extending toward higher RMS values, but to peak roughly near the theoretical limit.

This limit was computed according to

$$\sigma = \frac{2k_b T_{sys}}{A_{eff} \sqrt{\frac{1}{2} N_A (N_A - 1) N_{pol} \Delta \nu t_{int}}} \quad (\text{A2})$$

where k_b is Boltzmann’s constant, T_{sys} is the system temperature, A_{eff} is the effective area of a single LWA antenna, N_A is the number of antennas, N_{pol} is the number of polarizations, $\Delta \nu$ is the bandwidth, and t_{int} is the integration time [*Thompson et al.*, 1991]. We assumed the system temperature was $T_{sys} = 350K + T_{sky}$ where T_{sky} is the average sky temperature, which is approximately $T_{sky} \simeq 45(\nu/300 \text{ MHz})^{-2.55}$ K. Treating each polarization separately, our imaging parameters give $N_{pol} = 1$, $N_A = 256$, $\Delta \nu = 6.1 \times 10^6$ Hz and $t_{int} = 5.079$ s. Recall from Sec. 2.1 that for an LWA antenna at 55.25 MHz, $A_{eff} = 4.57 \text{ m}^2$. This gives a theoretic noise

value of 2,230 Jy. The distributions of calibrated image noise are shown in the lower panel of Fig. 13. One can see that they both have large positive tails, as expected, and both peak near this theoretical limit, with the Y polarization having a slightly higher peak RMS than the X polarization. This result serves as a partial validation of our flux density calibration and of the performance of our imaging pipeline.

Acknowledgments. Basic research in astronomy at the Naval Research Laboratory is supported by 6.1 base funding. Construction of the LWA has been supported by the Office of Naval Research under Contract N00014-07-C=0147. Support for operations and continuing development of the LWA1 is provided by the National Science Foundation under grant AST-1139974 of the University Radio Observatory program. Part of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

- Baars, J. W. M., R. Genzel, I. I. K. Pauliny-Toth, and A. Witzel (1977), The absolute spectrum of CAS A - an accurate flux density scale and a set of secondary calibrators, *Astron. Astrophys. Supp.*, *61*, 99–106
- Cervera, M. A., W. G. Elford, and D. I. Steel (1997), A new method for the measurement of meteor speeds: The pre- t_0 phase technique, *Rad. Sci.*, *32*, 805–816
- Cornwell, T., R. Braun, and D. S. Briggs (1999), Deconvolution, in *Synthesis Imaging in Radio Astronomy II*, *ASP Conference Ser.*, vol. 180, edited by G. B. Taylor, C. L. Carilli, and R. A. Perley, pp. 151–170, ASP, San Francisco, Calif.
- Cornwell, T. and E. B. Fomalont (1999), Self-Calibration, in *Synthesis Imaging in Radio Astronomy II*, *ASP Conference Ser.*, vol. 180, edited by G. B. Taylor, C. L. Carilli, and R. A. Perley, pp. 187–199, ASP, San Francisco, Calif.
- Dowell, J. (2011), Parametric Model for the LWA-1 Dipole Response as a Function of Frequency, Long Wavelength Array Memo 178, Dec. 20, 2011 <http://www.ece.vt.edu/swe/lwa/>
- Dowell, J., D. Wood, K. Stovall, P. S. Ray, T. Clarke, G. Taylor (2012), The Long Wavelength Array Software Library, *J. Astro. Inst.*, *1*, id. 1250006, doi: 10.1142/S2251171712500067
- Elford, W. G. (2004), Radar observations of meteor trails, and their interpretation using Fresnel holography: a new tool in meteor science, *Atmos. Chem. Phys.*, *4*, 911–921
- Ellingson, S. W., et al. (2013), The LWA1 Radio Telescope, *IEEE Trans. Ant. Prop.*, in press
- Ellyett, C. D. and W. Davies (1948), Velocity of Meteors Measured by Diffraction of Radio Waves from Trails during Formation, *Nature*, *161*, 596–597
- Greenhow, J. and J. Neufeld (1955), The Diffusion of Ionized Meteor Trails in the Upper Atmosphere, *J. Atmos. Terr. Phys.*, *6*, 133–140
- Helmholtz, J. F. and N. E. Kassim (2009), The Evolution of Cassiopeia A at Low Radio Frequencies, *Astronom. J.*, *138*, 838–844
- Hicks, B. C., et al. (2012), A Wide-Band, Active Antenna System for Long Wavelength Radio Astronomy, *Pub. Ast. Soc. Pac.*, *124*, 1090–1104
- Holdsworth, D. A., I. M. Reid, and M. A. Cervera (2004), Buckland Park all-sky interferometric meteor radar, *Rad. Sci.*, *39*, RS5009
- Li, G., B. Ning, L. Hu, Y. Chu, I. M. Reid, and B. K. Dolman (2012), A comparison of lower thermospheric winds derived from range spread and specular meteor trail echoes, *J. Geophys. Res.*, *177*, A03310
- McKinley, S. W. R. (1961), *Meteor science and engineering*, McGraw-Hill, New York, N.Y.
- Shepherd, M. C., T. J. Pearson, and G. B. Taylor (1995), Difmap: an Interactive Program for Synthesis Imaging, *Bull. Amer. Astro. Soc.*, *27*, 903
- Taylor, G. B., et al. (2012), First Light for the First Station of the Long Wavelength Array, *J. Astro. Inst.*, *1*, id. 1250004, doi: 10.1142/S2251171712500043
- Thompson, A. R., J. M. Moran, G. W. Swenson (1991), *Interferometry and synthesis in radio astronomy*, Krieger Pub., Malabar, Fla.
- Wislez, J.-M. (1995), Forward scattering of radio waves off meteor trails, in *Proceedings of the International Meteor Conference*, edited by P. Roggemans and A. Knöfel, pp. 99–117, Brandenburg, Germany

J. F. Helmboldt, US Naval Research Laboratory, Code 7213, 4555 Overlook Ave. SW, Washington, DC 20375 (joe.helmholtz@nrl.navy.mil)

(Received _____.)

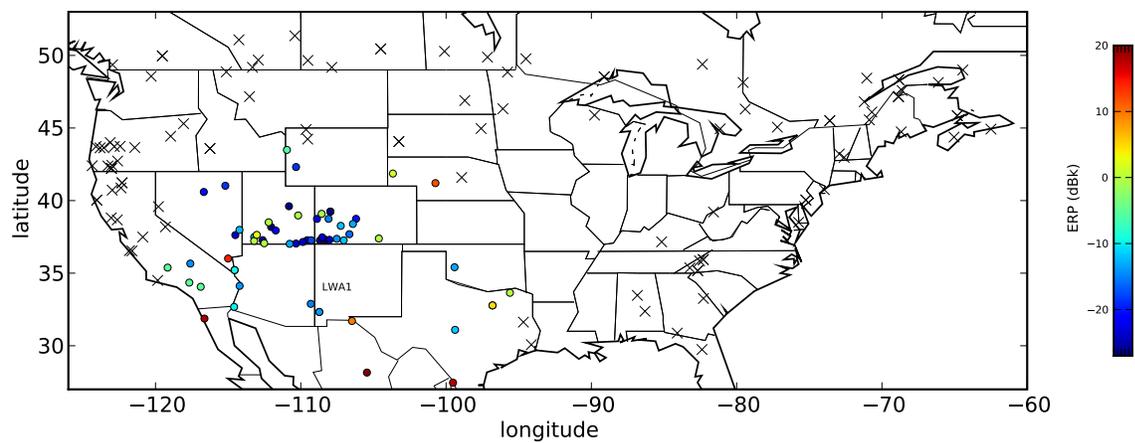


Figure 1. According to FCC records, the locations of all transmitters broadcasting analog Channel 2 (NTSC) in the US and in the border regions of Canada and Mexico. The 60 stations for which LWA1's zenith line of sight at an altitude of 100km is visible are color-coded by ERP; the remaining stations are plotted as black \times 's.

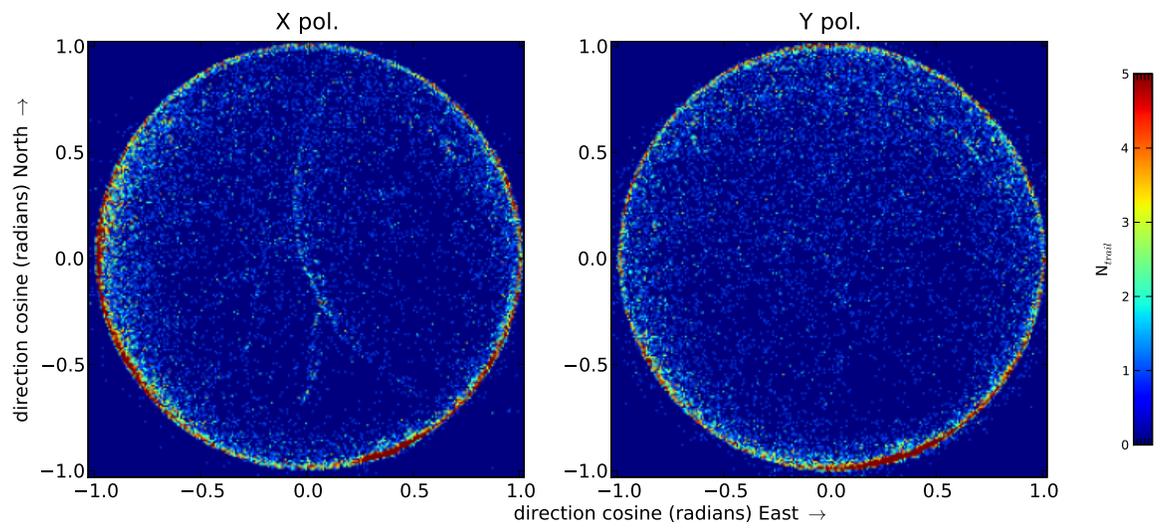


Figure 2. A two-hour, 55.25-MHz observing run conducted on 16 August 2012. Shown are images of meteor trail counts per all-sky image pixel for each polarization, X (north-south linear) and Y (east-west linear).

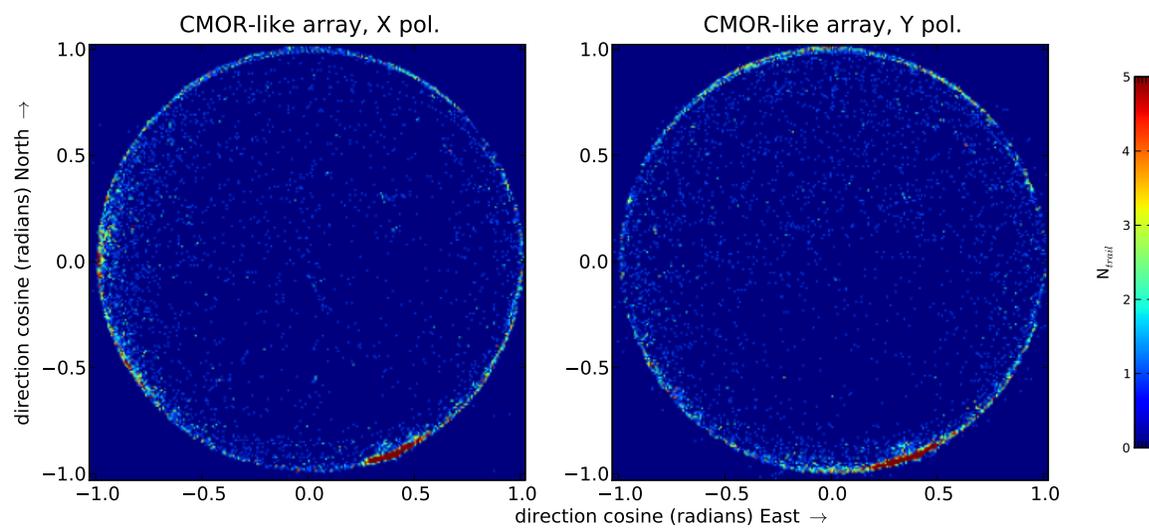


Figure 3. The same as Fig. 2, but excluding those trails that would not be detected if a 5-antenna receiver array similar to those used by CMOR was used instead of LWA1.

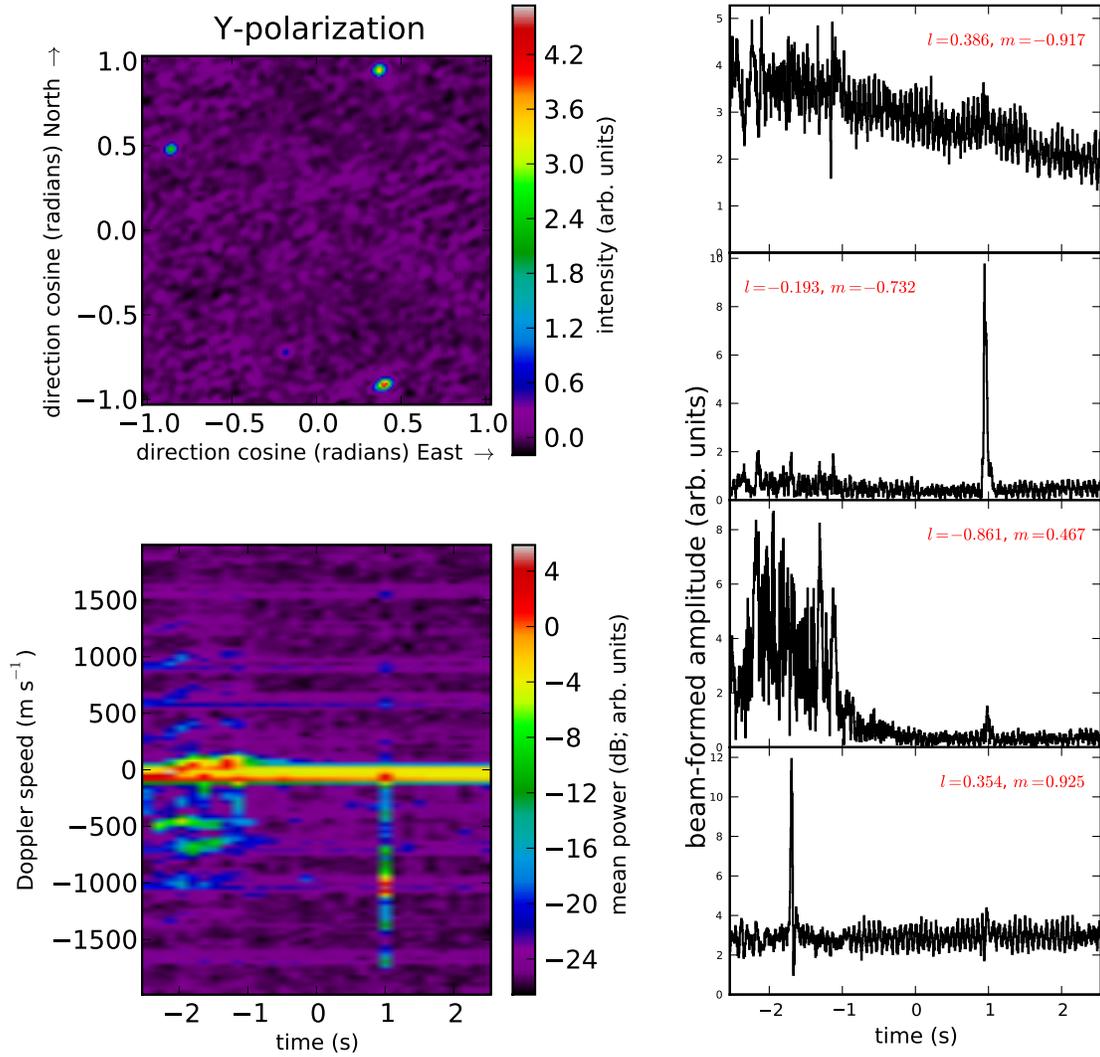


Figure 4. (upper left) An all-sky image made at 55.25 MHz (195.2 Hz bandwidth) with the Y (east-west) linear polarization LWA1 dipoles with 5.079s of data; (lower left) the dynamic spectrum averaged over all LWA1 Y-polarization dipoles at a temporal resolution of 0.16384 s; (right) amplitude time series for significant detections at a resolution of 5.12ms.

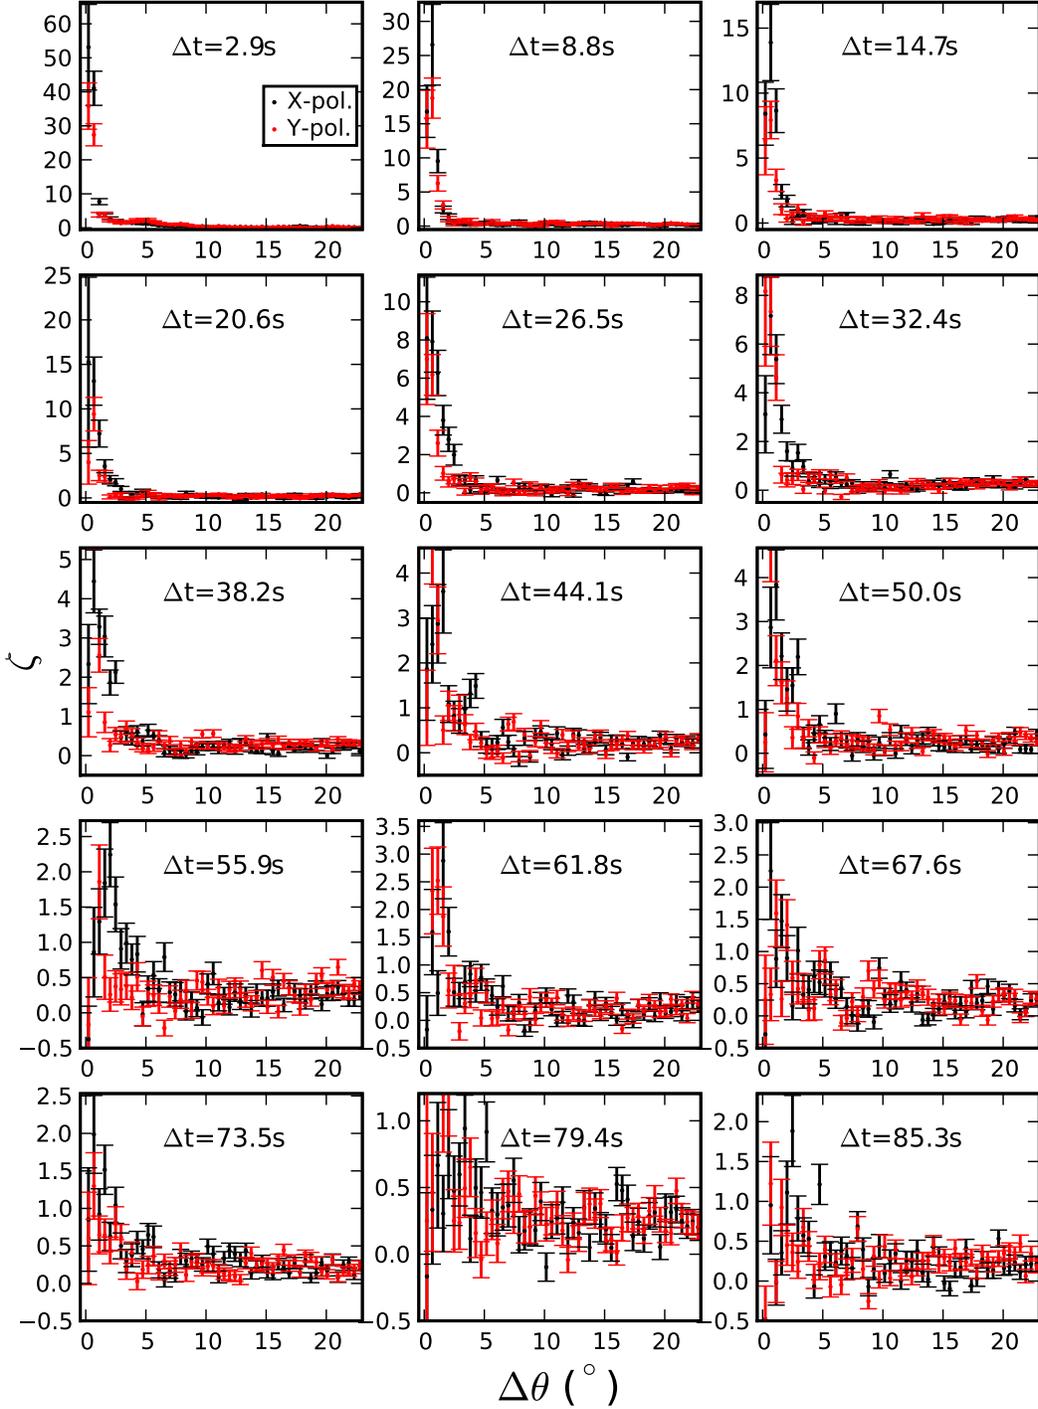


Figure 5. Correlation functions, ζ , as functions of angular separation, $\Delta\theta$, for different fixed temporal separations, Δt , for all trails above 30° elevation. Details of the computations are given in Sec. 3.1. The two polarizations are plotted separately with X (north-south) in black and Y (east-west) in red.

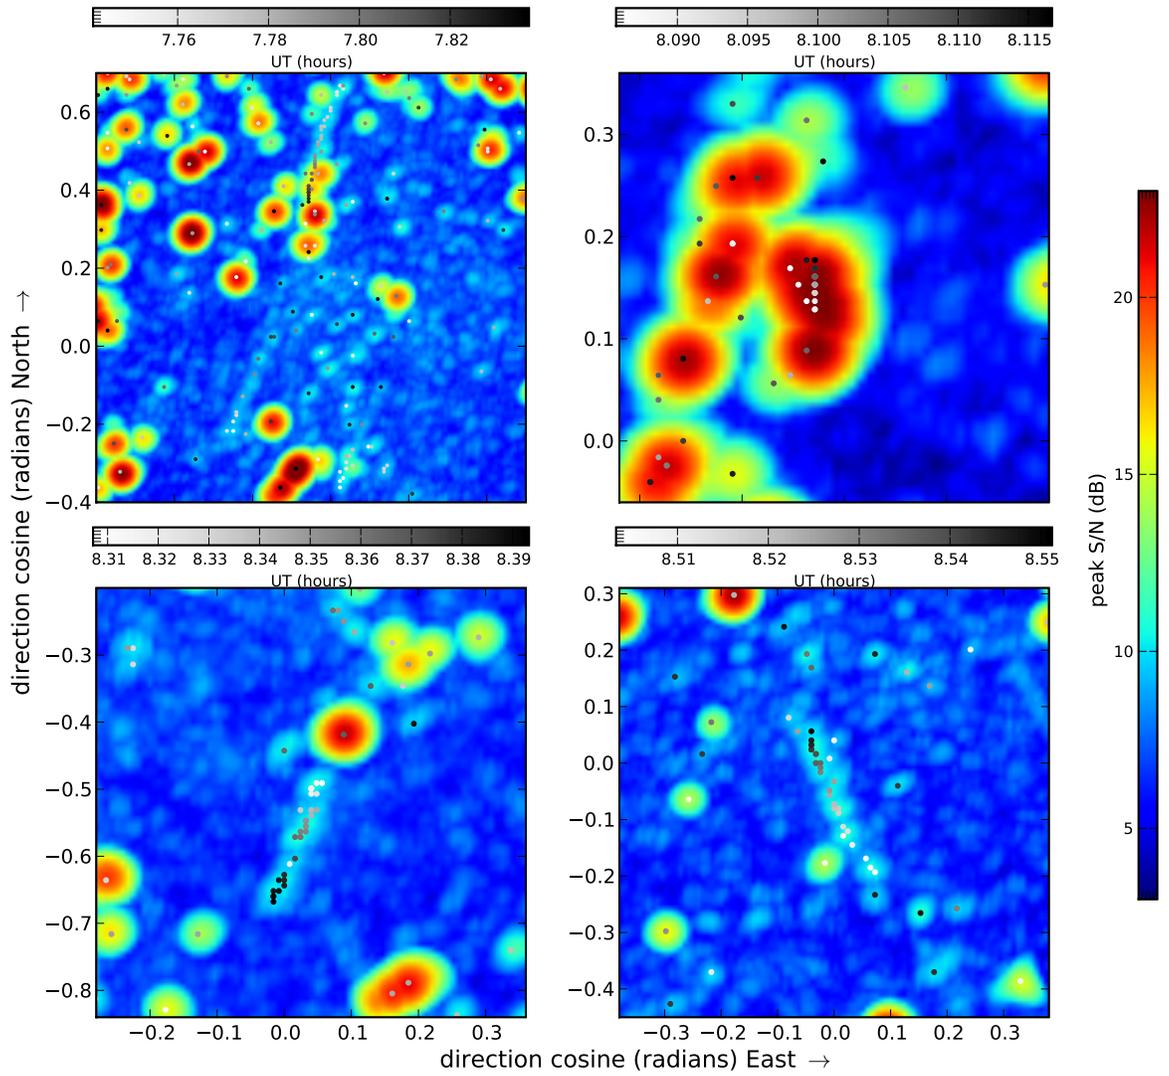


Figure 6. Four examples of long-duration meteor trails described in more detail in Sec. 3.1. In each panel, the peak S/N over the full duration of the trail is displayed with the position of detections from individual 5.079s images plotted as points, color-coded by time (UT).

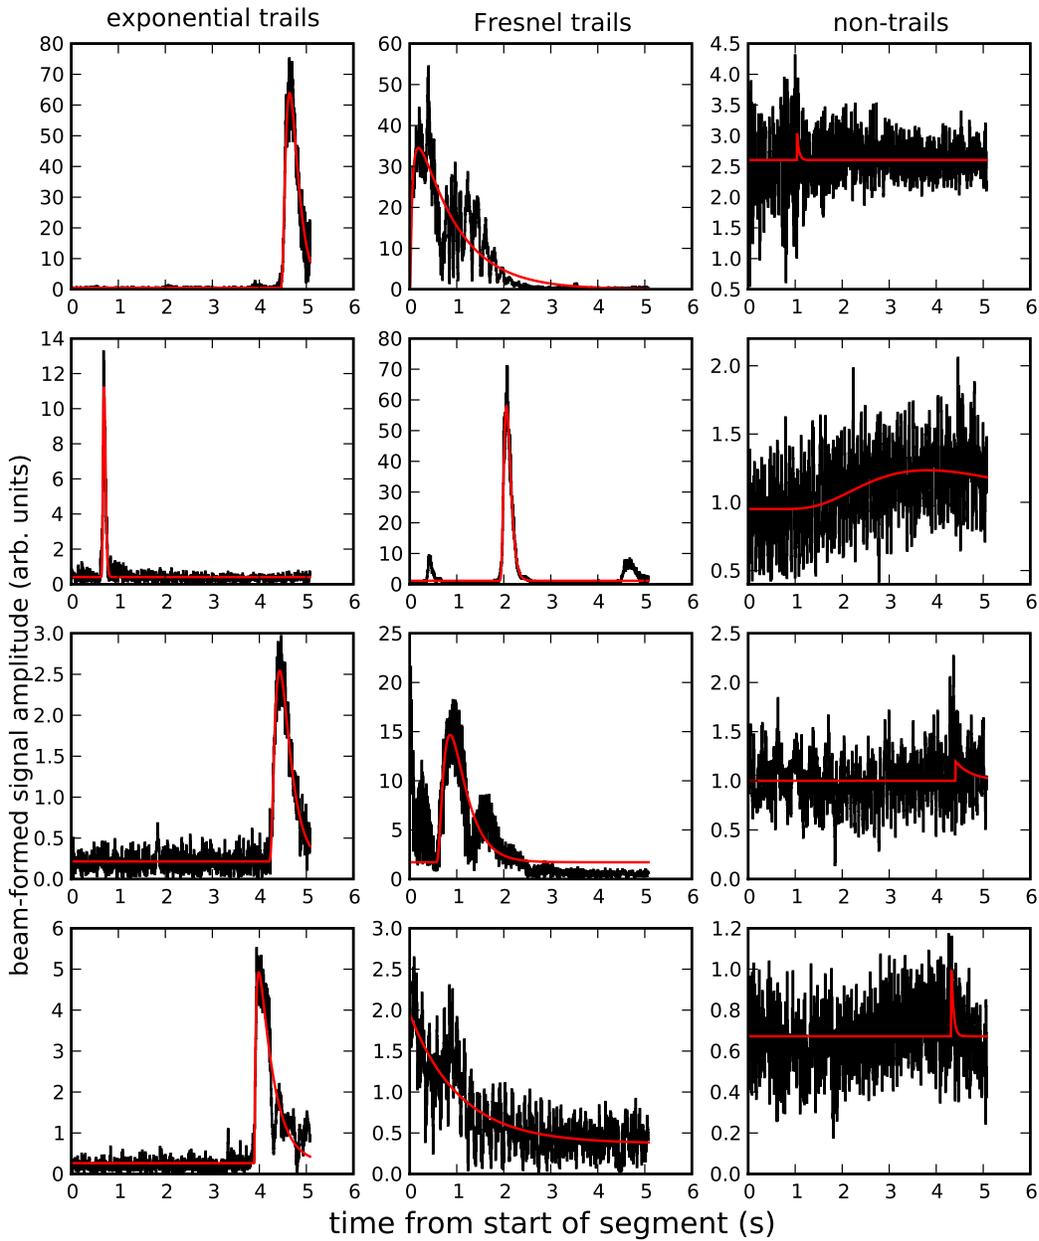


Figure 7. Example high-resolution amplitude time series for bright detections within three different categories: (left) “pure” exponential trails, (middle) Fresnel trails, and (right) non-trails. Model fits are also plotted in red. These fits and the algorithm used to automatically classify amplitude time series are described in detail in Sec. 3.2.

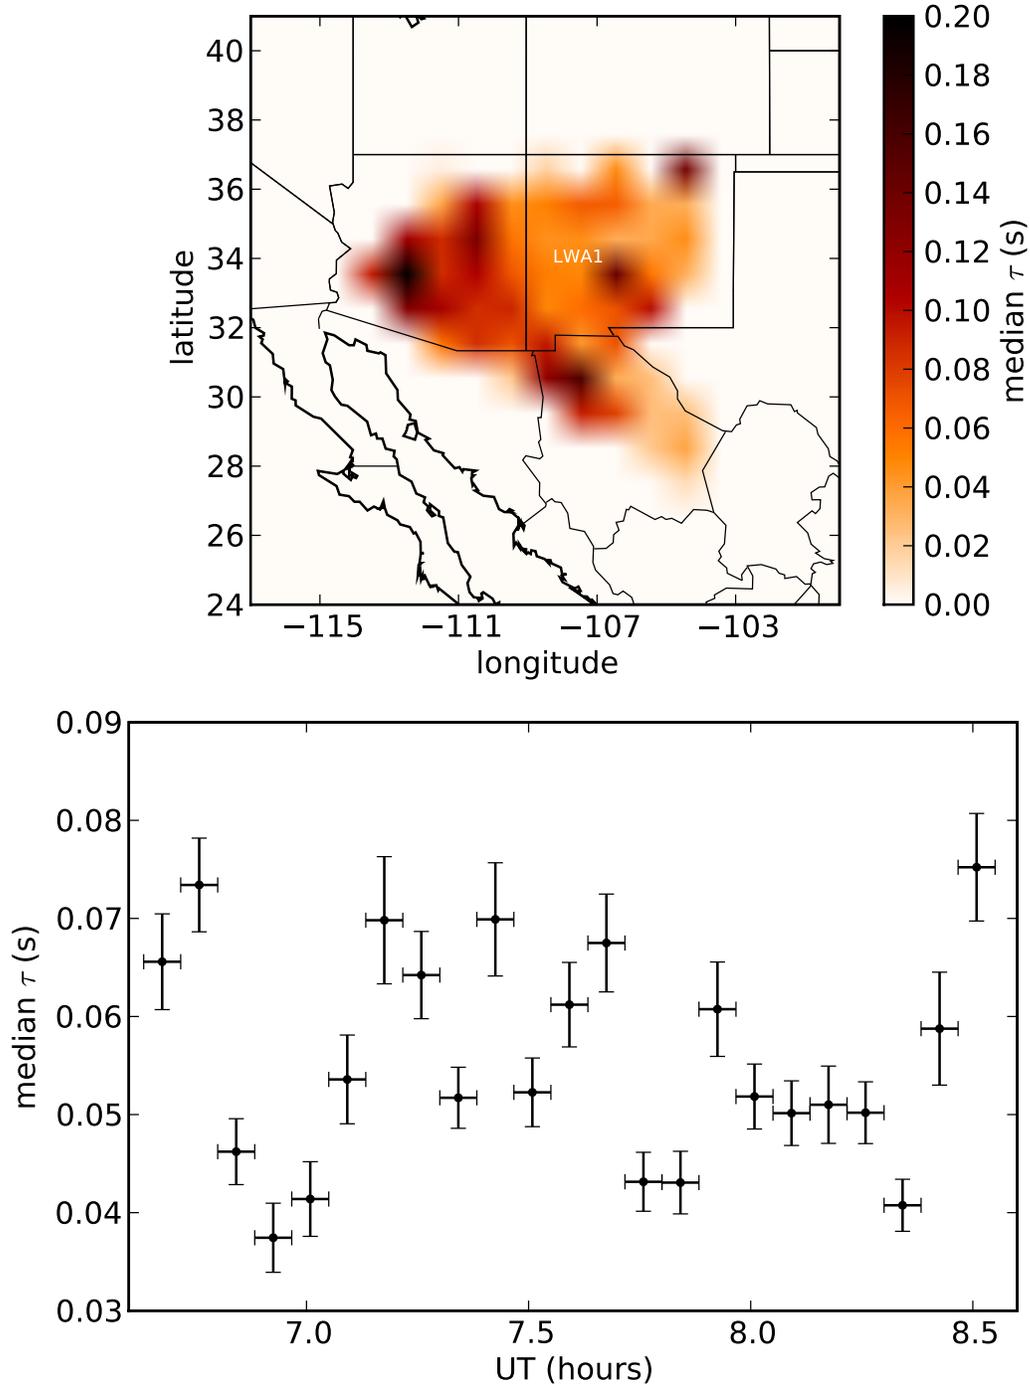


Figure 8. (upper) A map of the median (amplitude) e-folding time, τ , among bright exponential trails over the entire two-hour observing run. (lower) The median τ over all bright exponential trails within 5m intervals.

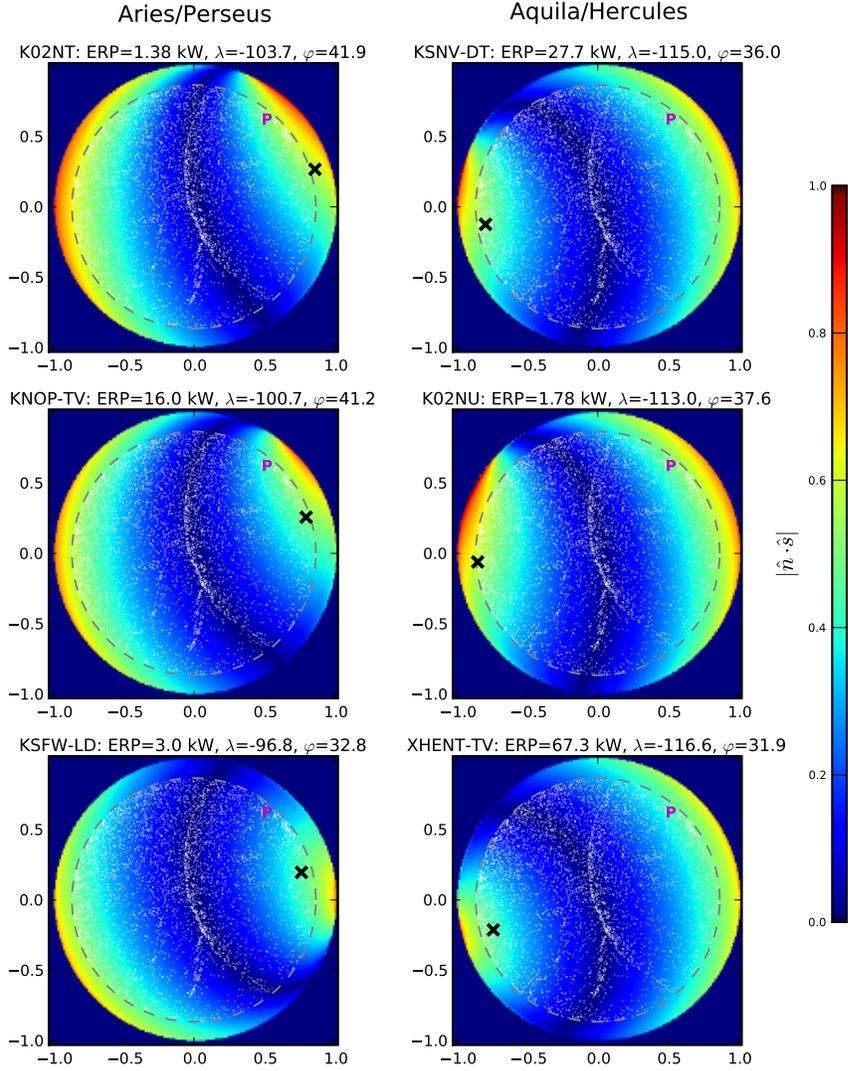


Figure 9. An illustration of the identification of the radiant and main illuminating transmitters for each of the two meteor streams detected at X-polarization. Each image shows the absolute value of the dot product between the best-fitting radiant unit vector, \hat{s} , and the normal vector, \hat{n} , required for a specular reflection between a particular transmitter and LWA1. This dot product is zero along the curve where a stream with radiant \hat{s} produces specular reflections. The call-sign, ERP, and location of each example transmitter are given above each panel. Trail positions are plotted as white dots for elevations above 30° (marked with a grey dashed line). The best-fitting radiants are plotted as black \times 's and the radiant for the Perseid shower is indicated with a magenta P.

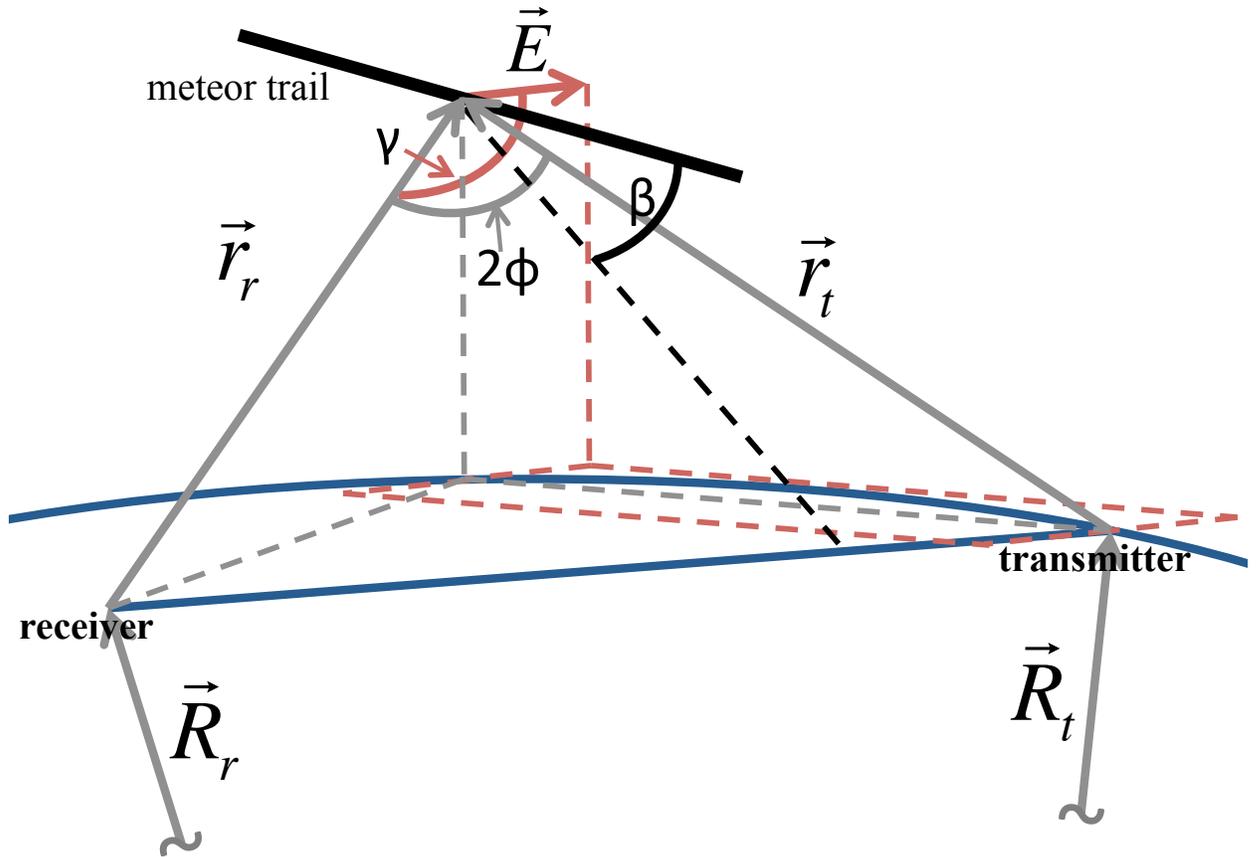


Figure 10. A schematic (not to scale) of the geometry assumed for forward scattering by a meteor trail. The labeled vectors and associated angles are discussed in the text and are included in equation (3). The red dashed lines trace the horizontal plane as viewed from the transmitter; the horizontally polarized electric field, \vec{E} , is parallel to this plane and perpendicular to \vec{r}_t .

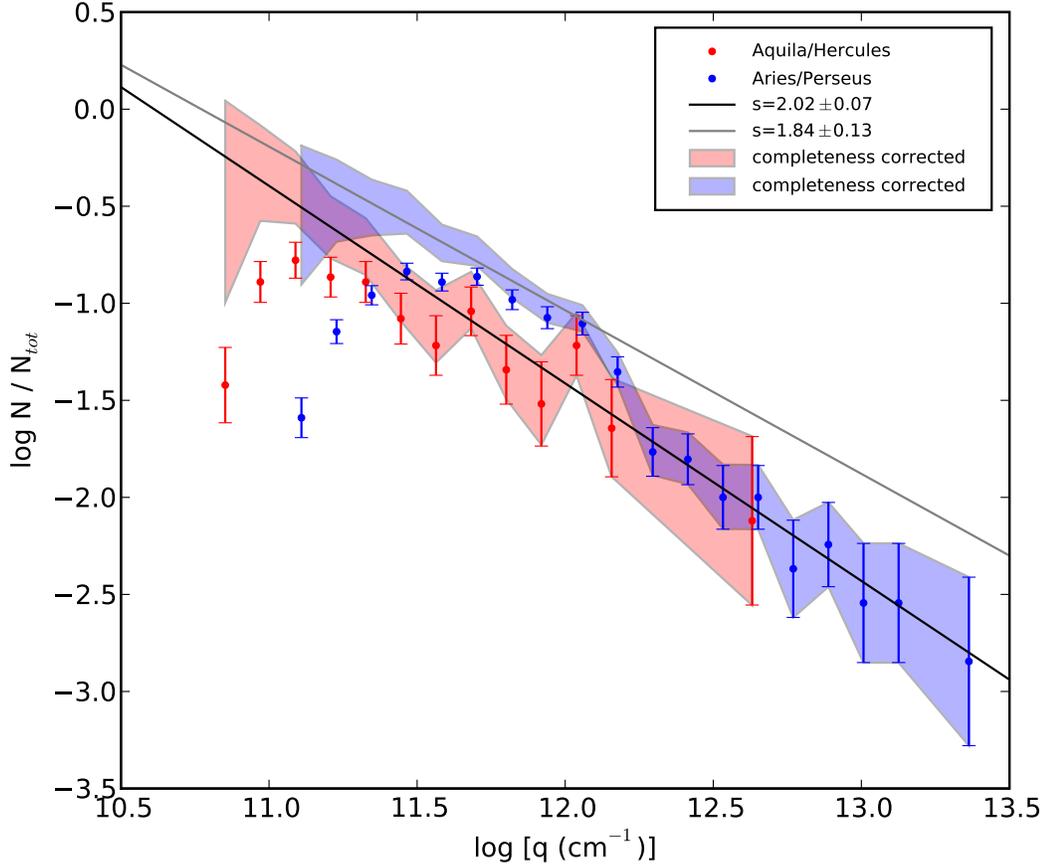


Figure 11. The line density, q , distributions estimated for the trails detected within the two meteor streams described in Sec. 3.3 and shown in Fig. 9. A separate histogram is shown for each stream, each being normalized by the total number of trails (red and blue points). The red and blue shaded regions show the completeness-corrected distributions (upper and lower limits, including 1σ errors) as described in Sec. 3.3. The result of a power-law fit to the combined data from the Aquila/Hercules stream and the $q > 10^{12.2} \text{ cm}^{-1}$ region of the Aries/Perseus stream is shown as a black line. A fit to the under-dense region ($q < 10^{12.1} \text{ cm}^{-1}$) of the Aries/Perseus stream distribution is plotted in grey.

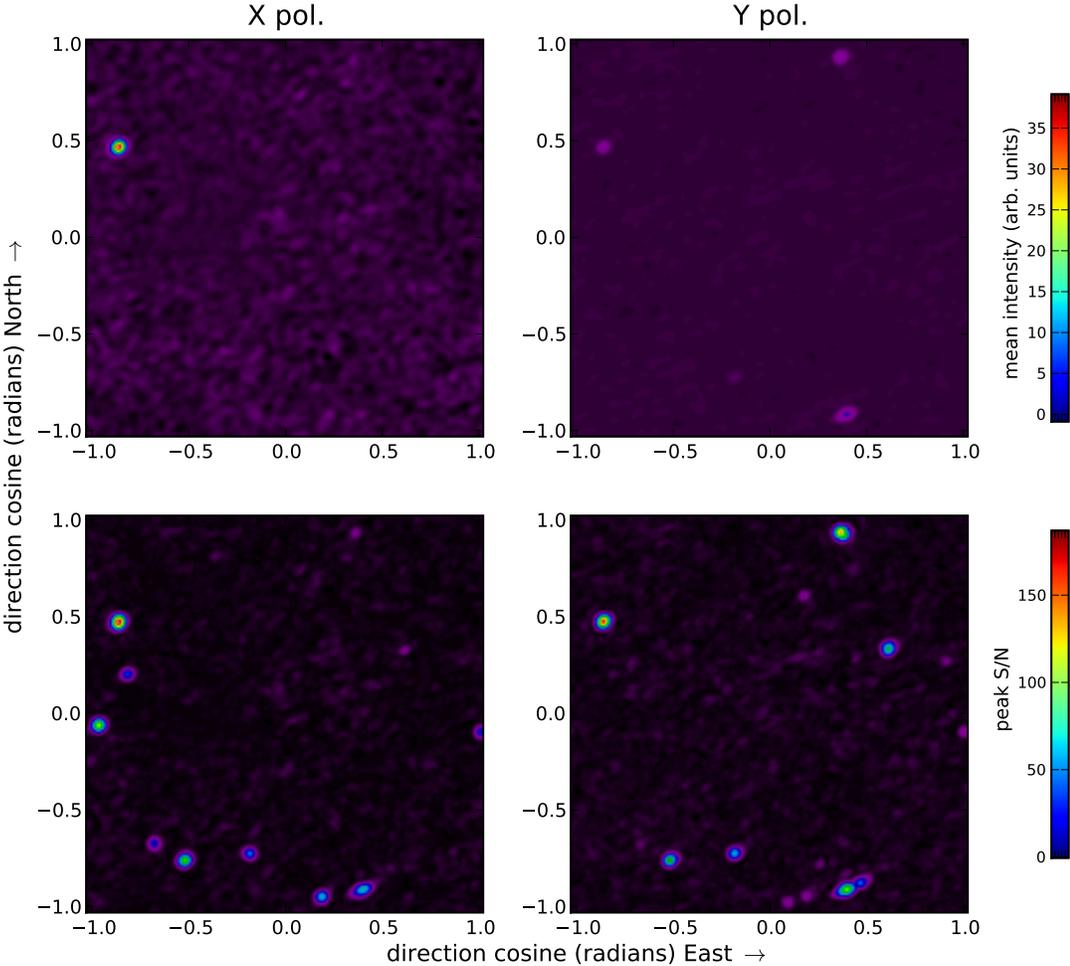


Figure 12. For a single 5.079s period, (upper) the mean images over all channels in X and Y polarization as compared with (lower) images of the maximum signal-to-noise (S/N) over the entire image cube, with the image from each channel normalized by its own estimated noise level.

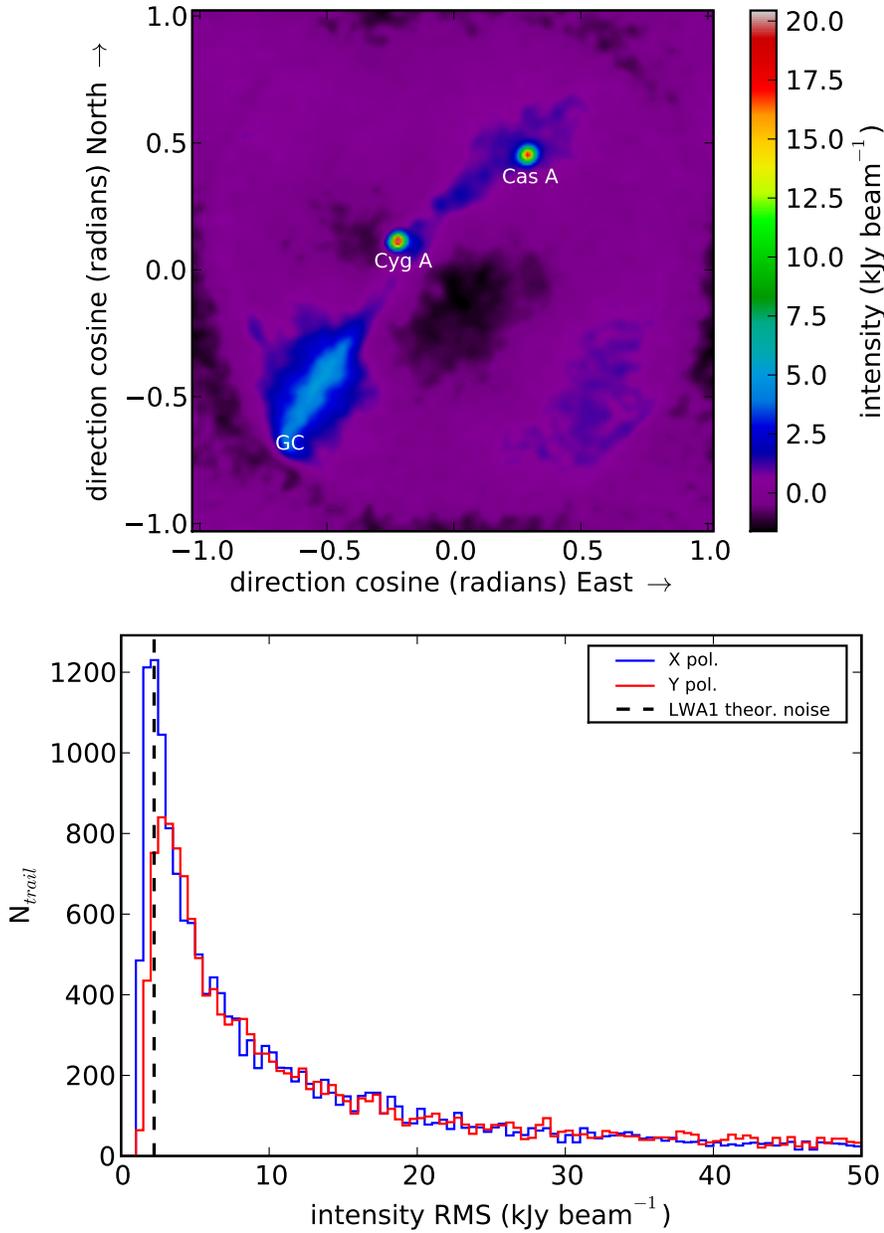


Figure 13. (upper) The median all-sky image (Stokes I) from five minutes of observing with an integration time of 5.079s and 60 kHz of bandwidth. The data were channelized into 6.1 Hz channels and median combined to eliminate the influence of meteor trails. Bright cosmic sources are labeled; the intensity of Cygnus A within these image data was used to flux calibrate the image cubes generated by the pipeline (see Appendix A). (lower) The distribution of image noise among the sources detected within the image cubes. Results are shown separately for the X and Y polarizations; a vertical dashed line represents the theoretical noise limit based on LWA antenna properties and the imaging parameters used.