Suggestion of coherent radio reflections from an electron-beam induced particle cascade

S. Prohira,[1,2] K. D. de Vries,[3] D. Besson,[4,9] A. Connolly,[1,2] C. Hast,[5] U. Latif,[4] T. Meures,[6] A. Nozdrina,[4] J. P. Ralston,[4] Z. Riesen,[7] D. Saltzberg,[8] J. Torres,[1] S. Wissel,[7] and X. Zuo[8]

[1]*The Ohio State University, Columbus OH 43210, USA*
[2]*Center for Cosmology and AstroParticle Physics, The Ohio State University, Columbus OH 43210, USA*
[3]*Vrije University Brussels, Dienst ELEM, IIHE, Pleinlaan 2, Brussel 1050, Belgium*
[4]*University of Kansas, Lawrence, KS 66045, USA*
[5]*SLAC National Accelerator Laboratory, Menlo Park CA 94025, USA*
[6]*University of Wisconsin-Madison, Madison, Wisconsin 53706, USA*
[7]*California Polytechnic State University, San Luis Obispo CA 93407, USA*
[8]*University of California, Los Angeles, Los Angeles, California 90095, USA*
[9]*National Research Nuclear University, Moscow Engineering Physics Institute, Moscow 115409, Russia*

(Received 23 January 2019; published 14 October 2019)

Testbeam experiment 576 at the SLAC National Accelerator Laboratory sought to make the first measurement of coherent radio reflections from the ionization produced in the wake of a high-energy particle shower. The >10 GeV electron beam at the SLAC End Station A was directed into a large high-density polyethylene target to produce a shower analogous to that produced by an EeV neutrino interaction in ice. Continuous wave radio was transmitted into the target, and receiving antennas monitored for reflection of the transmitted signal from the ionization left in the wake of the shower. We detail the first run of the experiment and report on preliminary hints of a signal consistent with a radio reflection at a statistical significance of 2.36σ.

I. INTRODUCTION

A particle shower in a medium produces high-energy particles that traverse that medium, ejecting ionization electrons from atoms in the bulk as the shower evolves. Two in-nature scenarios, high-energy cosmic ray interactions in the atmosphere, and neutrino interactions in the ice, have been considered [1,2], with the primary difference in shower development being the density of the medium. The Telescope Array RAdar (TARA) [3] project was the first dedicated experiment to attempt detection of the extensive air shower [4] from a cosmic ray interaction in the atmosphere using the radar method. TARA reported no signal [5], but placed a strong experimental limit on the extant model of in-air radar reflections [6]. Subsequent models [7,8] have further disfavored the in-air detection method. For dense media (like ice) several experiments have sought to detect radar reflections from ionization deposits in a laboratory setting [9–11]. Chiba *et al.* [9]

reported positive results for reflections from ionization deposits in dense material, albeit not from particle-shower-induced ionization. The Testbeam 576 (T576) experiment at SLAC was designed to make the first direct measurement of radar reflection from the ionization produced by a particle shower in a dense material.

The $\mathcal{O}(1$–10 GeV$)$ electron beam at the SLAC has a nominal bunch number of 10^9 electrons. Directing this beam into a target of high-density polyethylene (HDPE) produces a shower equivalent to that produced by a 1 EeV primary neutrino, which can be interrogated with radio in an effort to quantify the ionization parameters of a true neutrino-induced cascade. To that end, the first run of T576 ran in May 2018, and is the focus of this article. A second run occurred in October of 2018 with several improvements, and the results from the analysis of those data is forthcoming.

II. EXPERIMENTAL SETUP

The End Station Test Beam facility provides users a $\mathcal{O}(1$ Hz$)$ bunch of high-energy electrons switched from the main linear accelerator (linac) over into End Station A (ESA). ESA is a "parasitic" user facility at the SLAC; i.e., the parameters of the electron bunch (energy, beam current) are selected by the main linac user, rather than the End

FIG. 1. The T576 experimental setup. The large white rectangular polyhedron at the center is the HDPE target. The beam enters from the left, with the entry point shielded by aluminum sheeting in an effort to mitigate transition radiation (TR). The circles (red) indicate the receiver/transmitter antennas. Second from left is the transmitter; the others are receivers.

Station user. For our purposes this was actually advantageous, as a scan of energies and currents allowed investigation of how a putative signal depends on those parameters. For T576 the beam current was typically ~250 pC, corresponding to roughly 10^9 electrons per bunch. The run-time variation of the beam current is shown in Fig. 6. The primary electron energy varied from 10 to 14.4 GeV throughout the experiment, with most of the data accumulated at 14.4 GeV. At the point where the beam exits the beam pipe at the end of the ESA, the bunch is highly collimated, occupying less than a cubic centimeter in volume.

Figure 1 shows the target assembled on site at End Station A at the SLAC. The HDPE target was initially constructed for the T510 [12] experiment, for which it was used to study the geomagnetic emission from a particle shower created within the plastic target. For T576, the

HDPE target was aligned with the beam by placing it on top of large concrete blocks. Transmitting and receiving antennas were positioned around the target in various configurations throughout the experiment, as described in detail below. Two different types of antennas were used: a log-periodic dipole antenna (LPDA) having a voltage standing-wave ratio less than 3.0 over a 1–18 GHz bandwidth, and a Vivaldi antenna with a 0.6–6 GHz bandwidth. The transmitter and receiver amplification was varied throughout the experiment as well, in order to quantify and mitigate backgrounds and also investigate the scaling properties of observed signals.

A typical signal chain and the data acquisition system (DAQ) configuration are presented in Fig. 2. In the later analysis section, more detail will be given on the amplification and filtration choices made for various "runs" during the experiment. The DAQ was a Tektronix TDS-694C four-channel, 10 GS/s digital oscilloscope connected to a laptop via a GPIB-USB adapter. This laptop was remotely accessible via a network link from the "counting house," which allowed for control of all scope parameters and real-time readout of the data. The transmitter, a Rhode and Schwarz SMHU signal generator, was also controlled remotely via the same computer with another GPIB-USB adapter, allowing real-time frequency and output level tuning. The final piece of equipment (also controlled via GPIB cable) was an Instruments for Industry SMCC100 power amplifier for the transmitter permitting output level variation, as well as automatic leveling control and queries for forward and reflected power. As no personnel are allowed inside of the End Station during operation, having such a high degree of remote control over the parameters of the experiment was critical for minimizing downtime for hardware adjustments, thereby allowing accumulation of as much data as possible. An integrating
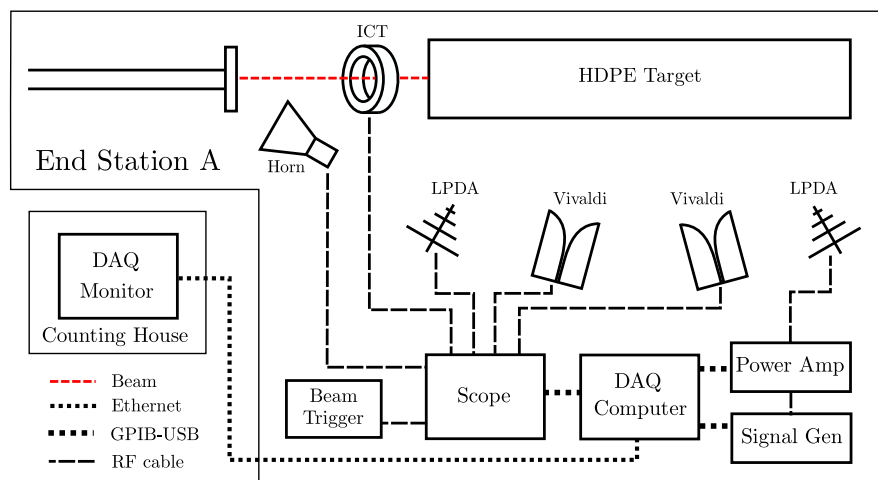

FIG. 2. The T576 signal chain. The DAQ system and transmitter resided in the End Station A, and were remotely monitored via an Ethernet link from the counting house, a remotely accessible location for users while the beam is on. The various components shown are described in the text.

current toroid (ICT) was used to monitor the beam current for every event, and occupied the fourth channel of the scope for the duration of the experiment.

The scope was triggered by either (a) a logic pulse from the accelerator itself or (b) a sharp transition radiation signal from an s-band horn (indicated by "horn" in Fig. 2), depending on the run. The TR horn signal was very sharp and consistent, but most of our data were taken using the beam logic pulse as a trigger since it could be modified remotely to allow precise time shifts of trigger point relative to the true arrival of the beam. Later runs substituted a third receiver antenna for the s-band horn, to better characterize the expected reflection signal as a function of the angle.

Partway through the run, the reported power amplifier output level began to drift by approximately 20% compared to the actual output power (determined by observing the signal strength in the scope). In what follows, we therefore assume a 20% systematic error on the transmitter output power.

III. SPECIFIC MEASUREMENTS AND GOALS

Radar observables can be classified as either model dependent or model independent. Model-dependent measurements in this case correspond to those observables which are dependent on multiple parameters, e.g., the plasma lifetime, the time required for overall ionization in the shower to decrease by $1/e$, and the microscopic scattering physics. The spectral content, temporal duration, and angular distribution of signal are important observables that will ideally either falsify or confirm different models. Simple measurements of coherence [13], that is, whether the received power in the signal region scales with distance as $\sim R^{-4}$, and linearly with transmitter power, are considered model-independent measurements.[1] For T576, we attempted as many different combinations of measurements as possible to test our results in both model-dependent and -independent manners.

Figure 3 shows the configuration for one run. In this run, all three receivers are positioned on one side of the target. One receiver was positioned at the specular reflection point for shower maximum (calculated to be roughly 3 m longitudinally into the target), with the other two set off at either side. For large plasma lifetimes, the reflecting region is large, which will concentrate reflected power at the specular point. For a short plasma lifetime, the scattering becomes more isotropic. This provides a model-dependent measurement, though the effect is not expected to be very strong.

---

[1]The R^{-4} scaling derives directly from the radar equation, which more specifically prescribes a scaling $\propto R_1^2 R_2^2$, where R_1 is the distance from transmitter to reflector and R_2 is the distance from reflector to receiver. This trend may not be observable at the SLAC because of the short baselines of our receivers. For a fixed baseline, the received power should scale linearly with transmitter power.
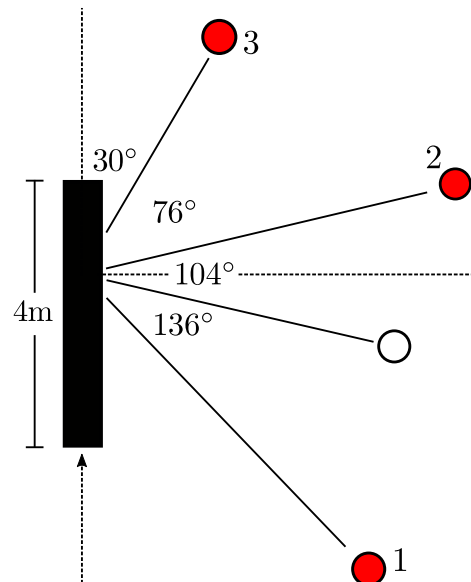

FIG. 3. The setup for run 11 of T576, viewed from above, and drawn to scale. Closed circles are receivers labeled by their DAQ channel number; open circle is the transmitter. One receiver is at the specular reflection point relative to the calculated shower maximum; the others are separated from the specular angles by 30–40 deg. The angles are measured from the vertical dashed line.

The overall goal of T576 was to measure an unambiguous radar reflection from a particle shower. The next sections will discuss the challenges of making such a measurement and the analysis of the data. As will be discussed, the extremely high backgrounds made many of these goals difficult to attain. However, after applying a particularly sensitive method for small-signal detection in large backgrounds, we present a suggestion of a signal that warrants further investigation.

IV. BACKGROUNDS

There were several backgrounds at ESA during T576 data taking. The typical radio-frequency (RF) backgrounds, anthropogenic and generic low-level electromagnetic interference, were relatively low within the thick concrete bunker-style building of the ESA. Occasional bursts of communications radio were observed, but well below trigger threshold. The two most pernicious backgrounds were observed to be room reflections and a very strong rf signal from the beam/target interaction itself.

A. Spurious reflections

The ESA is characterized by many sharp angles, reinforced concrete, and randomly placed metallic equipment accumulated over decades of previous experimentation. For most particle physics applications, this is irrelevant, but for radio, each surface is a potential reflector that can affect the signal seen at the receiver. The reflections in the room were
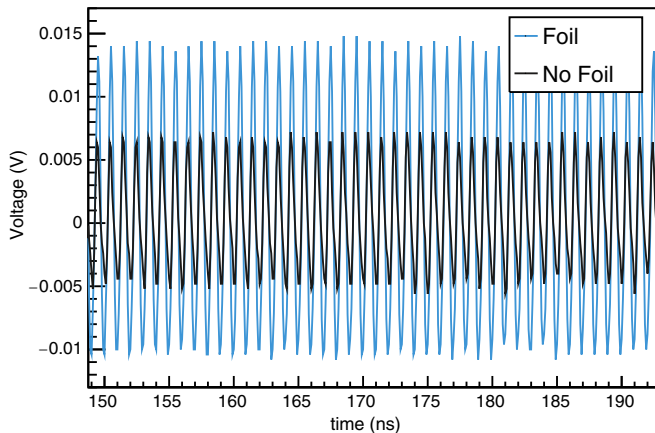
FIG. 4. A foil test, for which a conductor (here 0.3 m × 1 m) is placed at the expected ionization maximum point on the target, to estimate how a reflection should be observed in the receiver. In this test, the foil reflection is in phase with the ambient background, indicating a good receiver position.

so pervasive that moving a receiving antenna relative to the transmitter by several centimeters could, in extreme cases, reduce the received amplitude of a continuous wave (CW) signal by an order of magnitude. Typically such reduction is achieved through active carrier cancellation (a procedure whereby the transmitted signal is split and one half is fed directly into the line of the receiver, to be combined with the signal arriving at the receiving antenna. With proper alignment, the phase of the combined signal cancels the otherwise-large carrier in the receiver completely, and thus allows for smaller signal-to-noise ratio (SNR) signals to be seen in the receiver stream), but at ESA, reflections from myriad surfaces required scanning for receiver nulls empirically. Once a receiver was positioned at such a null, an additional "foil test" calibration was performed to verify that the addition of a reflecting surface near the expected location of the reflecting shower would result in a clear signal enhancement (compared to the no-foil configuration) at that receiver. For some configurations, it was observed that the foil test would result in a further nulling of the signal, indicating a poor receiver location for that particular frequency. For others, such as the one shown in Fig. 4 the amplitude of the carrier with the foil in place is approximately twice larger than without, indicating a favorable receiver position.

The foil tests were also quite useful from a simple physics standpoint—reflections with a piece of foil on the order of the expected dimension of the ionization plasma gave a crude approximation of the signal amplitude to be expected during the run. In good agreement with prerun simulation, amplitudes of $\mathcal{O}(1 \text{ mV})$ were observed.

B. Beam splash

The second, far more challenging background was the so-called ~100 mV, several hundred nanosecond duration "beam splash," which likely is the result of somewhat complicated physics at the point at which the beam strikes the target. This background likely combines sudden appearance [14] (a special type of transition radiation), transition radiation [15–17], and Askaryan radiation [18,19], plus myriad reflections from the room and from within the target itself. Beam splash was observed at all values of θ, as shown in Fig. 5, but was more pronounced in the forward beam direction, as expected (additional details on beam splash will be presented in the analysis section).

It is worth noting that beam splash would not be present for an experiment seeking to use this technique to detect in-ice neutrinos. The only background to the radar signal, from the shower itself, would be the Askaryan signal, over a very restricted solid angle.

V. DATA ANALYSIS

This section describes the data analysis for one of the cleanest runs, towards the end of the experiment, when most of the backgrounds had been at least partially characterized. We follow a technique which employs several different methods of matrix decomposition to filter background and extract signal [13]. An alternative method using similar techniques is presented in the Appendix. We present a suggestion of a signal, and describe a follow-up beam test with slightly different parameters to definitively establish this signal.

A. Setup

For this run (run 11), antennas were aligned vertically (VPol), and there was no active carrier cancellation.[2] The present analysis will focus on data taken using a transmit frequency of 1 GHz and 5–25 W output power. The layout of the receivers is given in Fig. 3, and the plots to follow are based on data taken from channels 1 and 2, which were both Vivaldi receiver antennas. There was no filtration or amplification on the input of the receivers, to avoid possible saturation effects, and to initially maximize receiver bandwidth.

B. Raw data

Figure 5 shows an event from run 11 taken in the counting house at run-time. The four panels on the left are the oscilloscope time traces uncorrected for cable delays and time of flight. From top to bottom these are CH1 (Vivaldi), CH2 (Vivaldi), CH3 (LPDA), and CH4 (ICT), respectively; corresponding power spectral densities (PSD) are presented in the right panel. As is evident from the figure, the amplitude of the beam splash is greater than 100 mV. In channel 3, downstream of the

---

[2]Carrier cancellation can reduce the amplitude of the carrier in the receiver, but it is only useful at the single nulled frequency. We opted to remove the cancellation in order to maximize our sensitivity to a wide range of reflected frequencies.
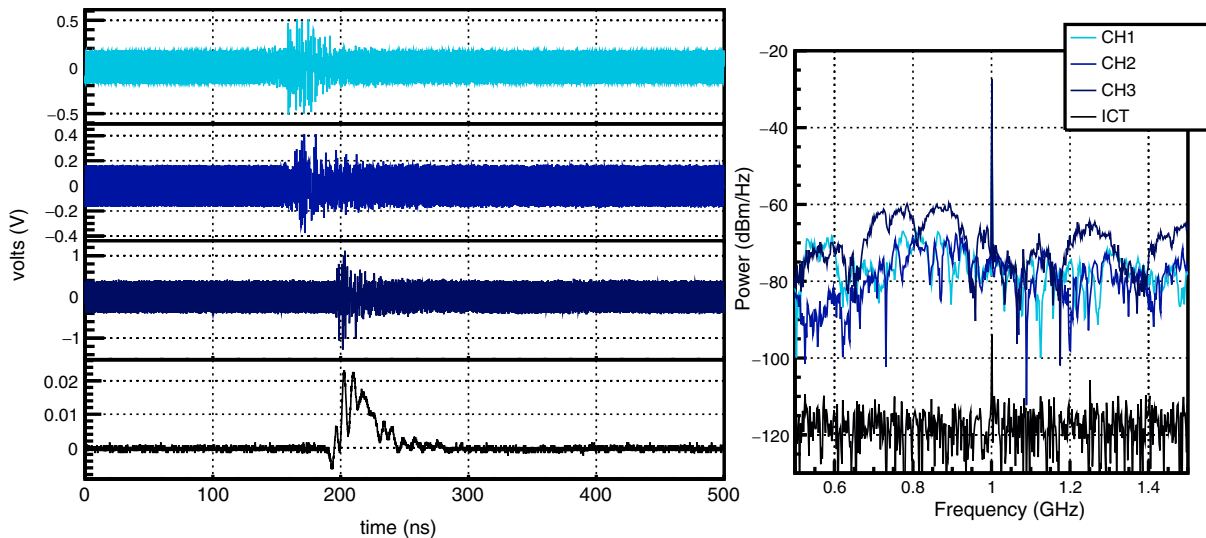
FIG. 5.   A typical T576 event. The four left panels are channels 1–4 from top to bottom, respectively, and the right panel shows their associated PSD. The offsets on the x-axis are due to air and cable propagation delays.

beam, the amplitude exceeds 1 V. The heavy peaking in the spectrum is likely a combination of system response and the room itself, with natural nulls at certain antennas for certain frequencies, as observed during the foil tests. The carrier is evident in the PSD.

Though the beam splash is large, it is exceedingly stable, which will later be extremely important to the background-subtraction procedure. The shot-to-shot variation depends on the amount of charge in the bunch, as seen in Fig. 6, where the energy in the beam splash scales with the beam charge measured by the ICT. This is useful for building up a background "template" and for constructing "null" data, to train the analysis techniques. Previous experiments [16] have made measurements of transition radiation which
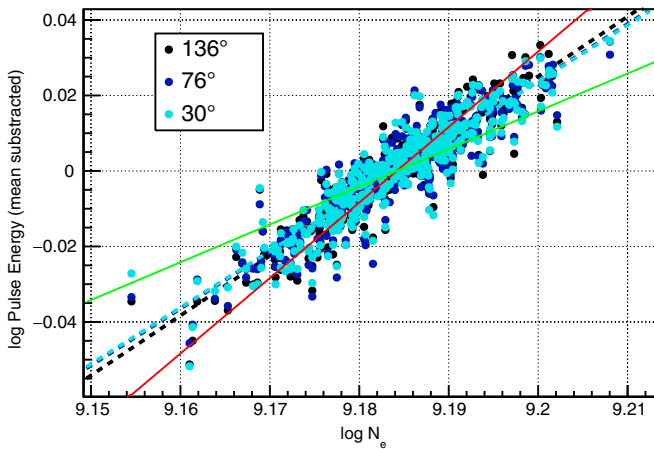

FIG. 6.   The ICT-measured electron number per bunch versus pulse energy, as measured by the antennas indicated by their angle from the beam momentum direction. The mean has been subtracted from all distributions in order to highlight the trend. Total coherence would correspond to a slope of 2 (red line), incoherence to a slope of 1 (green line).

show a quadratic scaling of TR energy with beam charge indicating coherence. The electron number N_e only varied by roughly 20% during our run, but our fit in log-log space has a slope roughly halfway between the expectation for complete incoherence (slope = 1, corresponding to the green line in Fig. 6) and complete coherence (slope = 2, corresponding to the red line). Interestingly, as the receiver is moved relative to the shower, the coherent contribution of the beam splash increases, albeit only slightly. To improve our signal sensitivity, the data at this point are up-sampled by a factor of 5 and then filtered at ±300 MHz from the carrier using a time-domain software bandpass filter.

C. Null data

To train the background-subtraction procedure, we developed a routine for building up what will be called null data, which are devoid of signal. These consist of carrier-only (beam OFF) data events added to beam-splash-only (transmitter OFF) data events, and constructed as follows:

(1)   A real event is selected from the data file.

(2)   A carrier-only event (beam OFF, transmitter ON) is selected from a carrier-only file with the same frequency and output power settings. It is matched to the real event in both amplitude (via scaling) and phase (via cross-correlation with the first 100 ns of the real event).

(3)   A template of beam-only events is produced by averaging over a beam-only run of 90 events. This template is then scaled using the measured value from the ICT and aligned, in time, with the real event via cross-correlation, windowed around the beam onset.

(4)   Now that the carrier-only event is aligned with the real carrier in the presignal region, and the beam-only event is aligned with the real beam splash in the
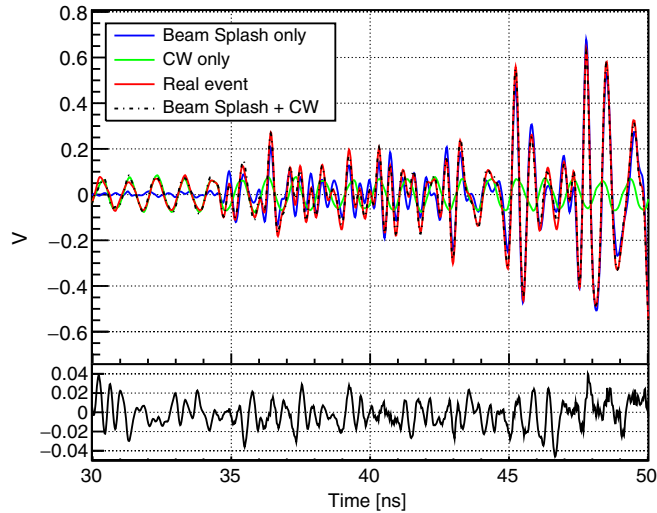
FIG. 7. An example of null data construction. A cw-only event and a beam-splash-only event are aligned to a real event and summed. This produces an event which mimics the real event in phase and amplitude. The residuals (real-null) is shown in the bottom panel.

signal region, the carrier-only and beam-only events are summed together to produce a null event, which contains no signal.

An example of this procedure is shown in Fig. 7. Indicated in the figure are the cw-only and beam-splash-only events used to make the null event, along with the real event and the resultant null event. This method of construction of the null data is important because the phase

relationship between the carrier and the beam splash changes from event to event, so any analysis technique needs to treat this variation carefully. It is essential that our null dataset has identical carrier/beam phase relationships as the real data. Many techniques for background reduction, such as simple averaging, will fail in this case, due to the lack of fixed phase in the carrier. Similarly, monitoring for power scaling at the signal region is not possible, as sometimes the carrier and beam splash add constructively, and sometimes destructively.

In what follows, every analysis step was carried out concurrently on two sets of data: real and null. The lack of a signal appearing in the null data gives confidence that any signal observed in the real data is not an artifact of our signal-extraction procedure.

D. Overview of analysis techniques

Figure 8 shows the full dataset for this analysis (e.g., all events from the cleanest run) as well as the associated null set. The analysis methods employed are based closely on [13], with necessary modifications specific to T576; we will use the vocabulary of that reference here. The procedure resembles pattern-matching routines for signal processing such as the Karhunen-Loeve technique, and is particularly suited to low-SNR data. It primarily involves decomposition of data into a basis of *patterns*, which are orthogonal modes (analogous to Fourier modes) that describe the data. The power of the process built on singular value decomposition (SVD) is that instead of predefined modes, as in Fourier or wavelet decomposition,
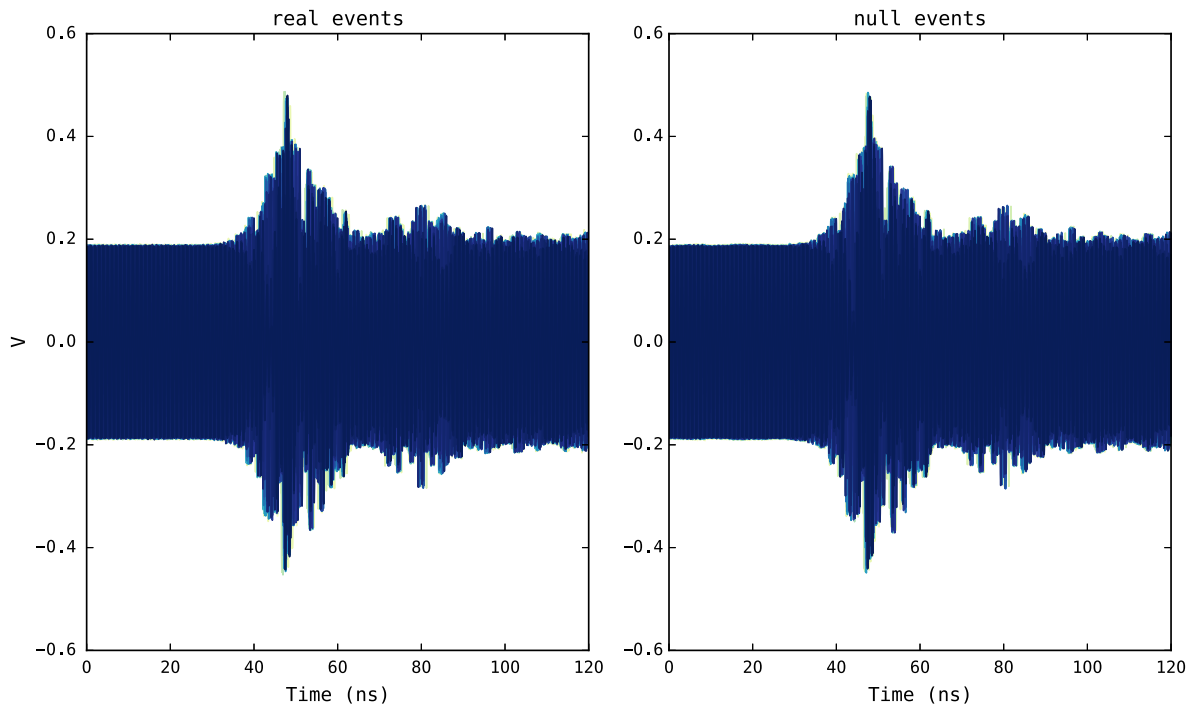

FIG. 8. The datasets used in this analysis. Left: Real data from the run. Right: Null data produced via the procedure described in the text. All events from both real and null sets are overlaid (~180 events).

the SVD method finds an orthogonal basis within the data itself to describe the data. This basis of patterns, or "eigenpatterns," or "modes" (these terms will be used interchangeably) is ordered in significance by corresponding singular values, or eigenvalues (i.e., weights). The relative scale of the weight is a measure of how well the data are described by that corresponding pattern. For T576, we expect that the beam background will occupy the most significant patterns in the decomposition, and by removing these, we can then reconstruct the reflected signal event, evident as less significant patterns in the decomposition.

Following [13], we use the following terminology:

(1) A *vector V* is synonymous with an event captured by the DAQ.

(2) A *pattern* is a basis mode from the decomposition, or an eigenvector, weighted by its associated eigenvalue.

(3) A *filter f* is a combination of one or more patterns which can be used to isolate, and, if desired, subtract components of the data. It can be thought of as a weighted sum over normal modes (again, the analogy to Fourier modes is useful here, in that a signal is built up of a sum of weighted normal modes).

The singular value decomposition is symbolically defined as

$$M = u\Lambda v*, \tag{1}$$

where M is a matrix to be decomposed, and u and v are matrices containing the singular vectors of M. These are the patterns which describe the data in M, and are ordered by the matrix Λ, which has the singular values along its main diagonal. The singular values are the weights of the corresponding patterns in u and v.

E. Carrier subtraction

Careful removal of the carrier from the T576 data is useful in isolating the signal. Removal of the carrier via successive sine-subtract filtration [20] is possible, but not problem-free. First, fitting a sine wave is susceptible to fitting errors; small errors in the subtraction can mean incomplete removal of an enormous background. Fits can be improved with prefiltering the data, but this costs information content. Second, the amplitude of the carrier in sine subtraction must be fixed to one value. If not, the amplitude envelope must also be extracted by a fit to the data, for which the amplitude may vary. Third, the presence of harmonics requires further fitting and subtraction, each potentially removing too much or too little information, and possibly introducing artifacts.

Using decomposition to remove the carrier solves these problems, if, for example, the modulations in amplitude are periodic or in any way repetitive, and the harmonics are stable. This is because the decomposition will yield the most significant mode of the data, which is not necessarily a Fourier mode, and may be some complicated

(but correlated, e.g., from the same source) structure. It can then be removed in whole or in part by the filtration method described below.

Because a carrier naturally has some periodicity, it is useful to break the data up into bins, to see if there exists an optimal binning for background subtraction. To find such a binning, we can use a decomposition. First we construct a vector V out of the *presignal* region and partition it into bins of length D, then we build a matrix out of these chunks, with each chunk corresponding to a row in this matrix

$$M_{ij} = V_{(i*D+j)}. \tag{2}$$

If the vector has length N, such that there are $d = N/D$ chunks in V, then the matrix M has dimensions $d \times D$. We can then perform SVD on this matrix and examine the "orderliness" of the singular values Λ as a function of the bin size D. This orderliness can be quantified by calculating the von Neumann entropy [21] $S = -\sum_{ij} \Lambda_{ij} \log(\Lambda_{ij})$ for the singular values in Λ and plotting that quantity against the bin size. If there are no significant patterns (e.g., if the vector is uncorrelated noise) the entropy will exhibit the standard logarithmic dependence, and vary as $\log(D)$. A downward excursion from the $\log(D)$ curve indicates that the data are more orderly with that binning. Figure 9 shows such a plot for this analysis. As evident in the figure, the values are all well below $\log(D)$, with a strong down-going excursion at $D = 25$, corresponding to one half period of the transmitted 1 GHz CW signal, indicating an optimal binning for this run.

When we subsequently bin with $D = 25$, the resultant distribution of singular values indicates that the carrier can be described fully by a small number of modes n. We can then zero out the remainder of the singular values,
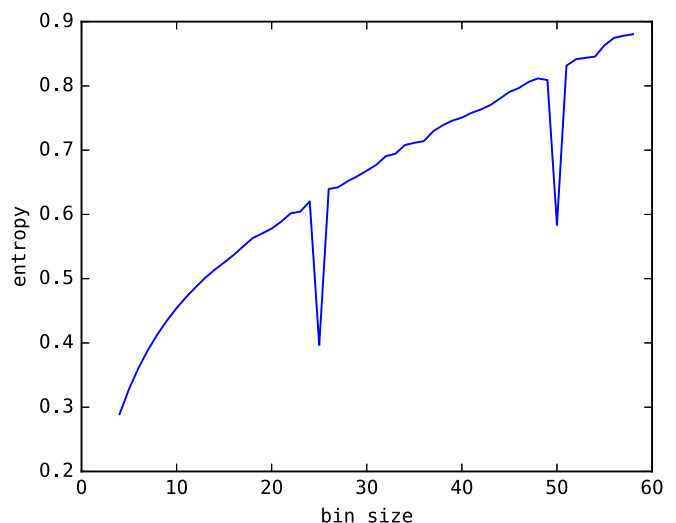

FIG. 9. The von Neumann entropy as a function of the bin size D, as described in the text.
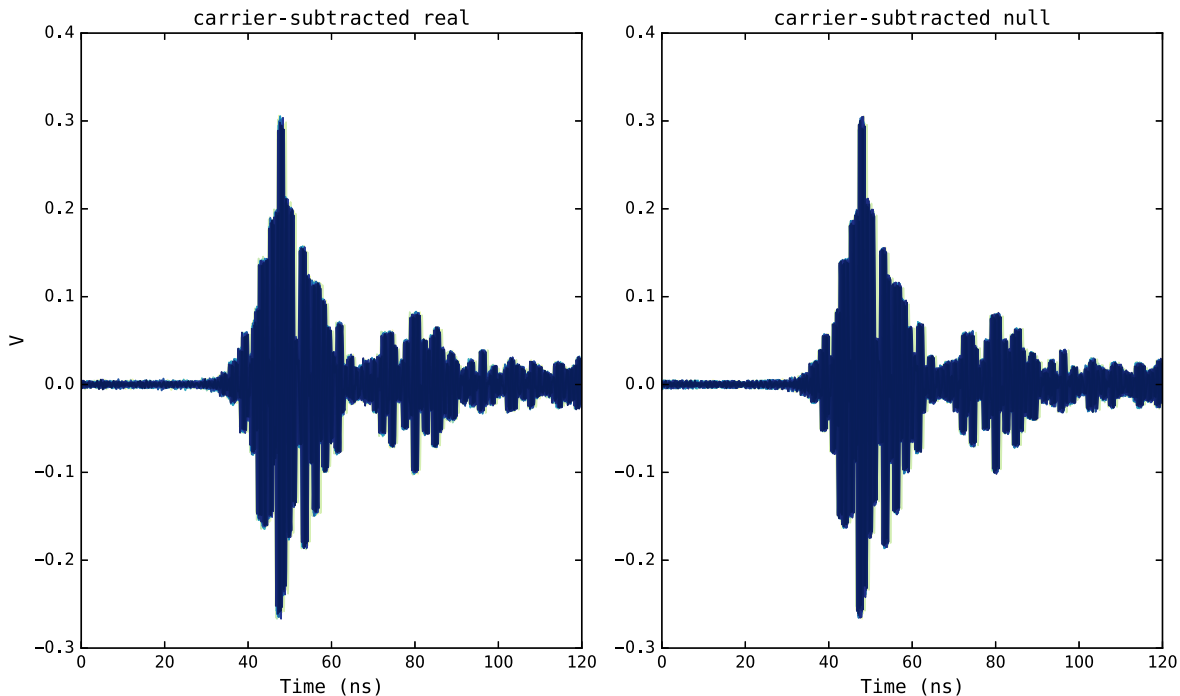
FIG. 10. The datasets with the carrier removed through the process described in the text. Left: Real data. Right: Null data. All events from both real and null sets are overlaid (∼180 events).

$$\Lambda'_{ij} = \Lambda_{ij}\Theta(n-i), \tag{3}$$

where Θ is the step function, and reconstruct the matrix M' by reversing the SVD, using the truncated matrix of singular values Λ'. The indices of the new matrix can then be flattened to recover a filter f_{carrier} which can then be subtracted from the signal region of the same event. The results of this filtration are shown in Fig. 10 for both real and null data, in which the ∼200 mV carrier has been reduced to the level of noise by this procedure. The carrier removal was performed on each event individually using the above procedure.

F. Alignment of the sets

At this point we have two carrier-subtracted sets, real and null, although some trigger-point jitter from the experiment still remains. We therefore subsequently align all of the events in both the real and null datasets, to the same, single reference event (selected arbitrarily, and subsequently discarded from the analysis) via a cross-correlation routine. The aligned, carrier-subtracted waveforms (all overlaid) are shown in Fig. 11, enlarged slightly to better illustrate the quality of the alignment.

G. Extraction of the signal

Now that we have carrier-removed and aligned waveforms, we perform decomposition of the remaining waveforms, removing the most prominent modes, co to the beam splash, and the least prominent modes, corresponding to

uncorrelated noise, and then, finally, performing a careful average on what remains to accentuate any possible signal.

We first build up two matrices, M^R and M^F for real and null data respectively, which have the following structure:

$$M_{ki} = V_i^k. \tag{4}$$

Here k is a label for identifying the event (e.g., $k = 1, 2 \ldots N$ for N events) and i is the index of the data within V^k. Therefore, M is a matrix in which each row is an event.
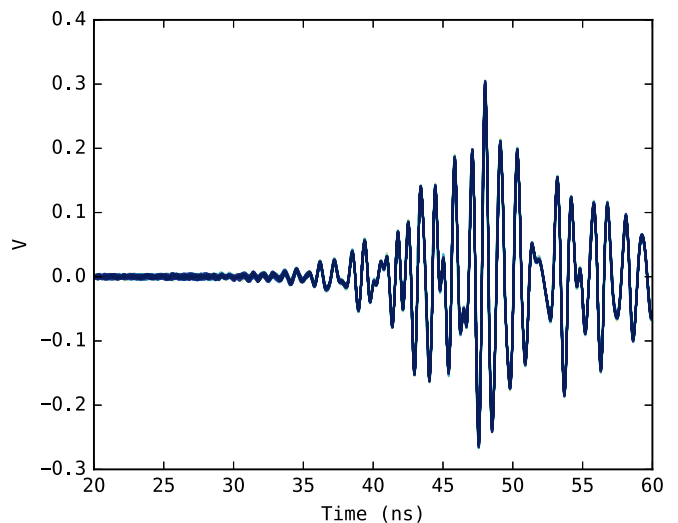

FIG. 11. Carrier-removed data aligned via cross-correlation. All events from both real and null sets are overlaid (∼180 events).
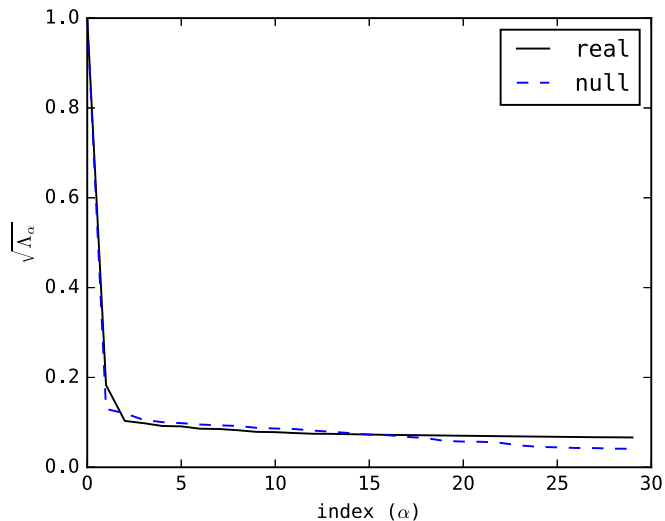
FIG. 12. The normalized distributions of singular values for real and null sets after decomposition.

We next make a decomposition of each matrix, which will simultaneously decompose all events into a basis for each full set, real and null. We then examine the normalized distribution of singular values $\Lambda_{\alpha\alpha}$, i.e., the entries of the diagonal matrix Λ. As shown in Fig. 12, where we have plotted $\sqrt{\Lambda_{\alpha\alpha}}$ to emphasize the shape of the curve, the two sets follow the same trend (the distribution is truncated above $n = 30$ for readability).

We now present the results of this analysis. The frequency-versus-time plots (spectrograms) that follow have units of power spectral density in $V^2\text{GHz}^{-1}$. The color bar scale is the same for real and null, and varies from plot pair to plot pair.

By inspection (via comparison to the original waveforms), it is evident that the beam splash, as expected,

corresponds to the first singular values in both real and null sets. This is shown in Fig. 13, for which we plot the average spectrogram of all events in the set, real and null, after reversing the decomposition, but prior to removal of any patterns (hence $M^{R\prime} = M^R$, $M^{F\prime} = M^F$). This is useful as a reference in what follows. We then truncate the singular values in the opposite way as before, that is, we zero the first three (most significant) singular values, and also remove the $n > 40$ modes to further suppress noise. We then reverse the decomposition to recover the filtered events. Next, we construct a spectrogram of each event, and, finally average the spectrograms. The result of this procedure is shown in Fig. 14, in which we observe a clear difference in the time-spectral content between the real and null sets. There is a clear scaling at the time when the reflected radar signal is expected (roughly 42–45 ns into the trace) in the real set, but not in the null set. Moreover, if we selectively examine each pattern and reverse the decomposition for each one singularly, in no case do we observe scaling at the signal onset point in the null data (excepting, of course, the first two beam-splash patterns).

The scaling in the signal region, which is only observed in the real data, is a suggestive hint of a signal. Of particular interest is the timing of the signal onset. It is clear by comparing Figs. 13 and 14 that the peak strength of the scaling in the signal region of the filtered data occurs before the peak strength of the beam splash in the unfiltered data. Calculations (based on cable delays and the known time of the beam on target) similarly predict arrival of the reflection, in the receivers, 5–10 ns before the peak of the beam splash.

Figure 15 shows a comparison of the real data to simulated signal, which has been produced using the RadioScatter code [22,23] using the exact specifications of this run and a plasma lifetime of 10 ns. The agreement is
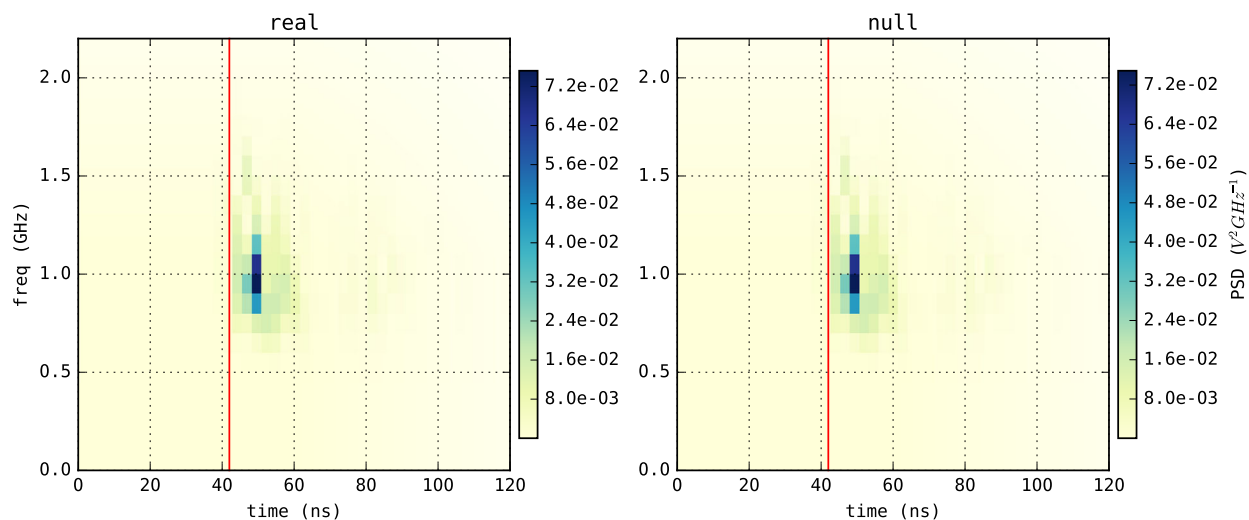

FIG. 13. The average of all events, real and null, after reversing the decomposition, without removing any patterns. The solid vertical line indicates the approximate expected signal onset point, in time.
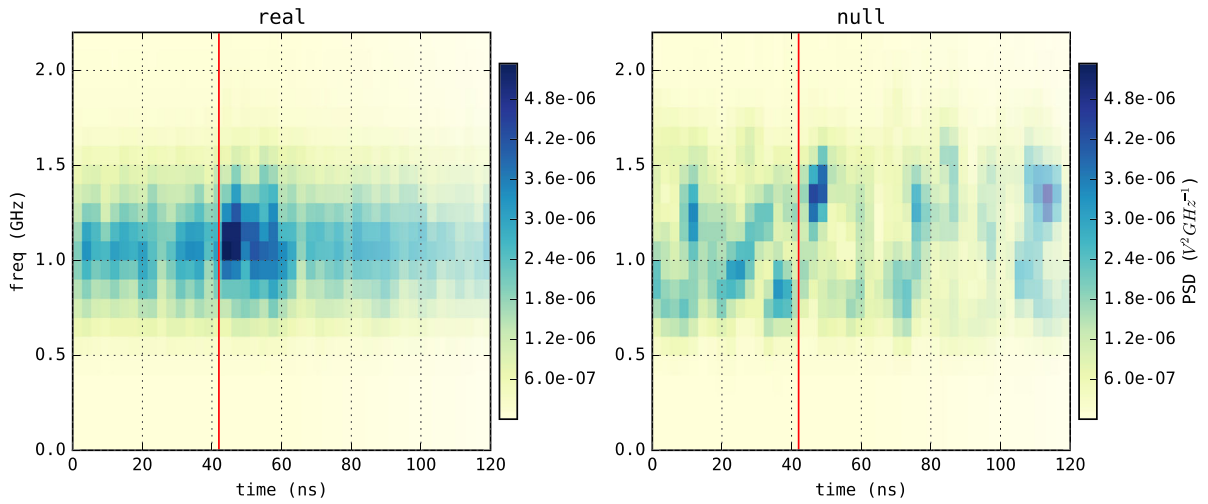
FIG. 14.   The average of all events, real and null, after reversing the decomposition and removing the most significant patterns, which correspond to the beam splash. The solid vertical line indicates the approximate signal onset point.
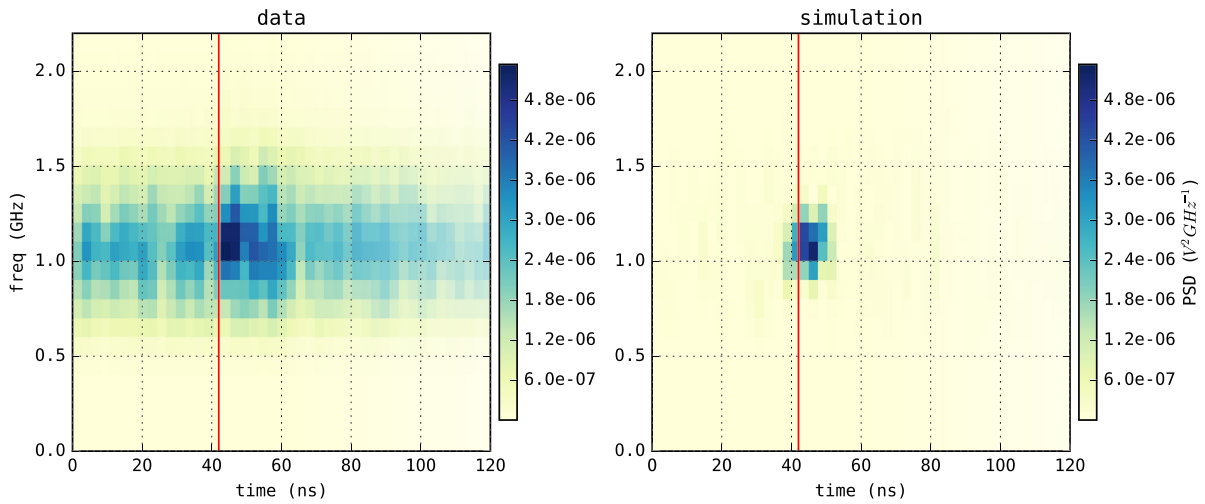

FIG. 15.   A comparison of the resultant filtered data to the RadioScatter simulation, for the same geometry and transmitter settings as the real run, with a plasma lifetime of 10 ns. The solid vertical line indicates the approximate signal onset point.

quite good, with the difference in power between data and simulation less than 10% and a very similar spectrogram shape. We note that we have not yet fully incorporated the full system response into the analysis chain; this is currently in progress. Accounting for cable losses (small, given the modest cable runs) and antenna inefficiencies is expected to reduce the signal power in the simulation by a few percent. Consistent with the noisy environment, the data trace is somewhat "messier" than the simulation. A more careful selection of patterns, e.g., eliminating some of the less significant noise modes, would likely clean up this background a bit. Although a full analysis of the noise modes has not yet been performed, we present an example noise mode (representative of all modes above $n = 10$ or so) in Fig. 16.

We follow the same procedure outlined above for data taken when the output power of the transmitter was reduced

from ~25 W to a nominal value of ~5 W.[3] For the latter data, the power of the carrier in the presignal region was actually observed to be smaller by a factor of 3.6 rather than 5. The resultant summed spectrogram is shown in Fig. 17, where there is a suggestion of scaling in both signal regions commensurate with the different outputs, although we observe only a factor of ~1.5 difference between the peak power in these two spectrograms rather than the factor of 3.6 cited above. This discrepancy is likely due to the aforementioned instability in the power amplifier, but warrants further investigation. Nevertheless, the shape of

---

[3]Data were obtained for only two power levels due to experimental constraints. With a limited number of triggers available during a run, it was decided that it was more valuable to collect a large number of triggers at a small number of power settings, than the other way around.
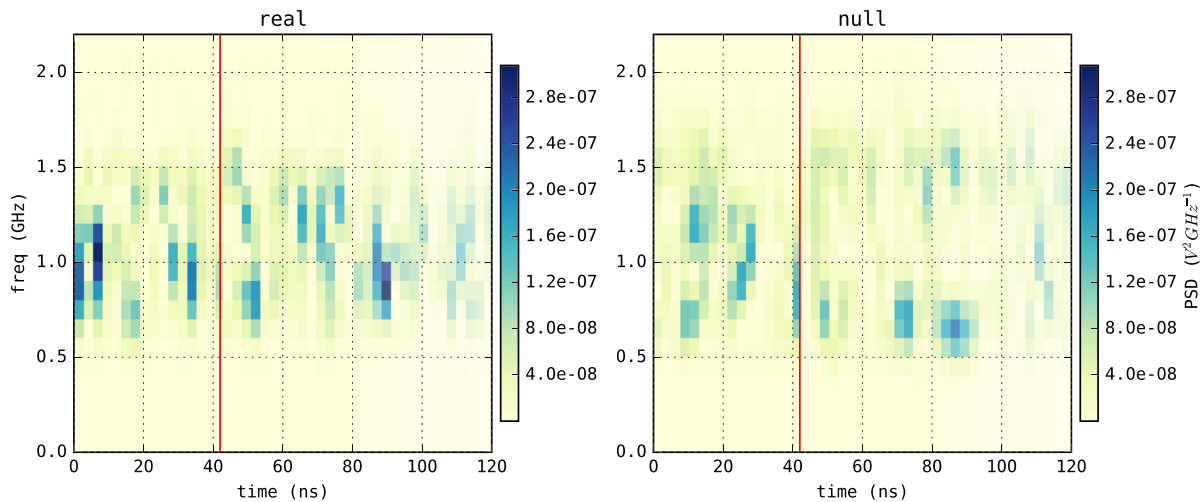
FIG. 16.   An example of a higher-order mode with singular value $n = 15$. The solid vertical line indicates the approximate signal onset point.
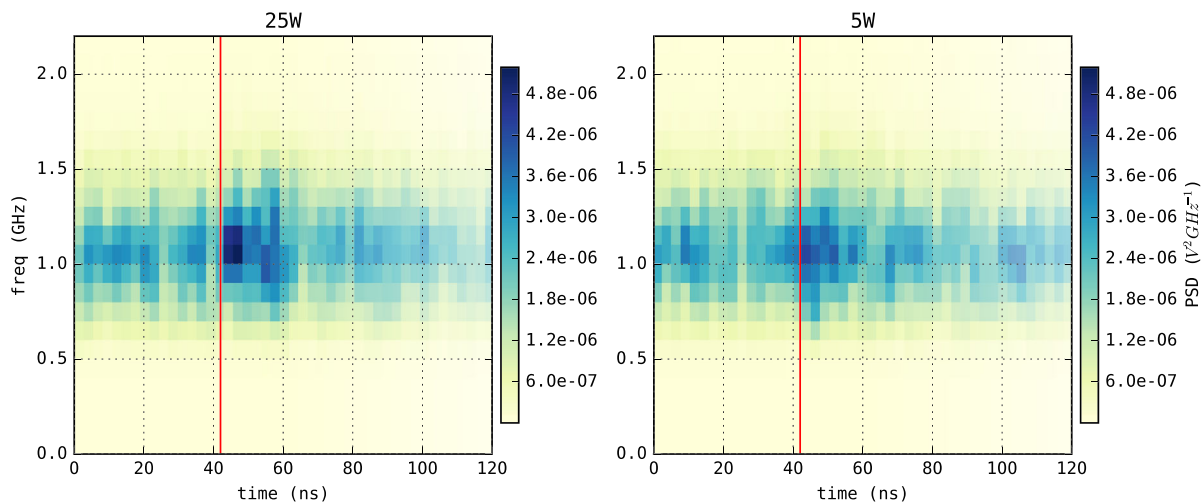

FIG. 17.   The average of all events, for 25 W output and 5 W output, after reversing the decomposition and removing the most significant patterns. The solid vertical line indicates the approximate signal onset point.

the signal is similar and the time onset identical to within one bin. Simulations predict that this configuration should produce radar reflections with a SNR of about 2, in fair agreement with the data.

VI. SIGNIFICANCE

We now present a quantitative assessment of the significance of the signal hint presented here, based on our analysis of the real and null datasets. To assign a significance to the observed excess, we employ a 2D sideband subtraction technique shown diagrammatically in Fig. 18. The adjacent sidebands in x and y are averaged and subtracted from the signal region. This ensures that the

apparent excess in the signal region is not simply a sum of the backgrounds from any residual beam splash plus carrier, at the point where they cross in the signal region.

Mathematically if $P(x, y)$ is the power in the spectrogram,

$$\langle x \rangle = \frac{\int_{x1}^{x2} \int_{y2}^{y3} P dx dy + \int_{x3}^{x4} \int_{y2}^{y3} P dx dy}{2}, \tag{5}$$

and similarly,

$$\langle y \rangle = \frac{\int_{x2}^{x3} \int_{y1}^{y2} P dx dy + \int_{x2}^{x3} \int_{y3}^{y4} P dx dy}{2}, \tag{6}$$

and

$$\langle b \rangle = \frac{\int_{x1}^{x2} \int_{y1}^{y2} P dx dy + \int_{x1}^{x2} \int_{y3}^{y4} P dx dy + \int_{x3}^{x4} \int_{y1}^{y2} P dx dy + \int_{x3}^{x4} \int_{y3}^{y4} P dx dy}{4}. \tag{7}$$
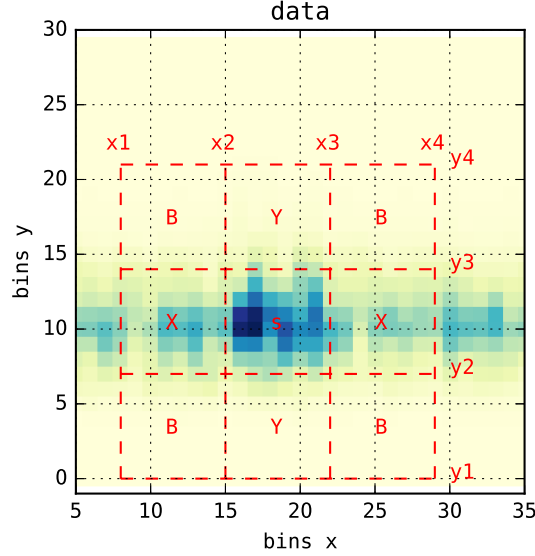
FIG. 18. The sideband subtraction method. Regions with similar labels (*x*, *y*, *b*) are integrated, and these integrals are averaged, then subtracted from the signal region.
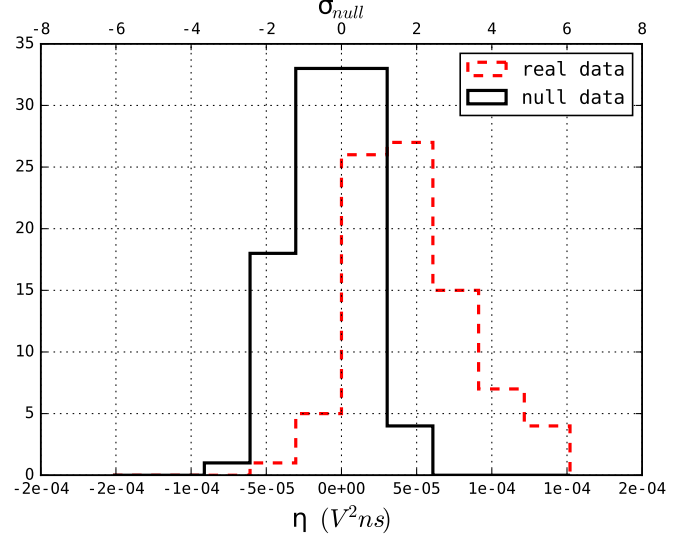

FIG. 19. The significance of the signal using the 2D sideband subtraction routine. x-axis units are given in V^2ns and also in terms of standard deviations of the null data distribution denoted σ_{null}.

The regions "*b*" are selected as "ambient" background, or an overall level which sits below the beam-splash remnant ("*y*") and the carrier-subtraction remnant ("*x*") such that e.g., the carrier remnant is $\langle x \rangle - \langle b \rangle$. The signal-region excess η is then,

$$\eta = \int_{x2}^{x3} \int_{y2}^{y3} P dx dy - (\langle x \rangle - \langle b \rangle) - (\langle y \rangle - \langle b \rangle) - \langle b \rangle \tag{8}$$

$$= \int_{x2}^{x3} \int_{y2}^{y3} P dx dy - \langle x \rangle - \langle y \rangle + \langle b \rangle. \tag{9}$$

We justify this procedure as follows: If the beam splash is sufficiently broadband, as it seems to be, then the amplitude of the remainder of the beam splash in the signal region (whatever has not been fully removed by the SVD method) will be well approximated by an average of the regions colocated in time, but with frequencies above and below the signal region. Similarly, along the time axis, the remnant of the carrier (leftover from subtraction) should not prefer any particular region; therefore the average of the regions before and after the signal region should approximate the remainder of the carrier within the signal region. These backgrounds are then subtracted from the signal region. Because both of these backgrounds ($\langle x \rangle$ and $\langle y \rangle$) also contain an overall ambient background ($\langle b \rangle$), we must add this background back in to avoid oversubtraction, as in Eq. (8).

We perform a sideband subtraction for each event in both the data and background sets, and plot the result in Fig. 19, where the x-axis is presented in units of the standard deviation of the background distribution σ_{null}. There is a clear excess in the integrated power η for the

signal events. The mean of the real data sideband-subtracted distribution is $2.36\sigma_{\text{null}}$ from the mean of the null distribution, at $-.28\sigma_{\text{null}}$. By inspection, some of the events in the real data distribution are consistent with the null data, while some show a significance greater than $5\sigma_{\text{null}}$. For the purposes of this analysis, we use the mean of the real data excess distribution to estimate a significance of $2.36\sigma_{\text{null}}$.

The process of performing the same analysis on the real and null sets eliminates analysis systematics, which would be present in both. Therefore, the main systematic is in the construction of the null data.

A. Model-dependent constraints resulting from this signal

Assuming that the excess described here is a true radar signal, it allows us to set some bounds on free parameters of radar models. The main parameter of interest is the ionization electron lifetime, the time it takes for overall ionization in the shower to decrease by $1/e$. In the RadioScatter simulation code this lifetime is a user-defined parameter. We find that to produce example signals such as Fig. 15, the plasma lifetime must be no longer than 10 ns (otherwise the excess is too long in duration), and seems to be consistent with a lifetime of 5–10 ns. The amplitude of the reflected signal is largely a function of the overall plasma density (which is a function of the energy of the primary) and this lifetime, with amplitude increasing with longer lifetime up to the point where the peak ionization density can persist for a full oscillation of the interrogating field. The effective volume for a future detector, as proposed in [22], was simulated using a plasma lifetime of 1 ns.

It must be noted, however, that the HDPE is not chemically the same as ice, meaning the ionization lifetime may not be identical to that for ice. A detailed study of this is

forthcoming, but the time quoted here is in good agreement to the extant measurements of plasma lifetime in ice [24].

VII. DISCUSSION AND NEXT STEPS

The results presented here, based on analysis of the receiver at the specular reflection point, comprise the majority of the usable data from the first run of the experiment, given the challenges outlined above. Given the strength of this 2.36σ result, a second run was performed in October 2018 with several improvements. The primary aim of the second run was to improve the signal-to-noise ratio of the experiment by mitigating the influence of the beam splash, which was anticipated [16], but not at as high amplitude as observed. These improvements were

  (i) Increased passband filtration. The putative signal from this analysis, in good agreement with simulation, shows a radar return near the carrier frequency. By tightly filtering, much of the beam-splash power can be removed, thus allowing a higher-resolution measurement on the oscilloscope.

  (ii) A second DAQ was used, based on the LAB-v4 chip [25], with 12 channels at 12-bit resolution and 3.2 GS/s sampling speed per channel.

  (iii) A different power amplifier with a wider frequency range (up to 6 GHz) and a stable output. Since the beam splash has a falling spectrum, the SNR can be improved by increasing the transmit frequency.

In addition to these improvements, a list of "smoking gun" signatures formed the basis of our experimental program for the second run. These were

  (i) An $\sim R_1^{-2} R_2^{-2}$ dependence on the putative signal in the real data, where R_1 is the transmitter-shower baseline and R_2 is the shower-receiver baseline. This signature may not be observable, given that our antennas are not strictly in the diffractive far field.

  (ii) A scaling of the return signal duration as a function of frequency.

  (iii) A frequency shift of the return signal at receivers displaced from the specular point.

  (iv) A scaling of the return signal as a function of the azimuth, with signal amplitude trending differently than the pure carrier amplitude.

  (v) A scaling of the return signal as a function of transmitter output power, to have more statistics on the power scaling observed in this analysis.

Data were taken to test for these signatures with both data acquisition systems, with sizeable sets acquired for each. At the time of this writing, the analysis for run 2 is ongoing.

ACKNOWLEDGMENTS

APPENDIX: EXPANSION OF DATA IN A BASIS

We can take a dataset and expand it in a basis. This basis can be a decomposition, as in the text, of any matrix. We can build up a matrix of, for example, null data, then decompose this matrix into a basis, and expand real data into this basis. This expansion can then be used as a filter for the real data. In what follows we have two datasets, a real set, and a null set.

The general procedure is as follows: We take the null set and decompose it via singular value decomposition into a basis. We then take both sets and perform the same carrier subtraction described in Sec. V E. At this point the real set contains beam splash, whatever remains of the carrier after subtraction, and the putative signal. The null set will, by definition, contain only the beam splash and whatever remains of the carrier after subtraction. Then, we take each real data event and expand it into the null basis. This expansion will contain only the elements in the real data which resemble the patterns in the null set; i.e., this is our filter. This expansion (filter) is then subtracted from the original event, which leaves only the components of the real event which do not resemble the null basis.

Here we present the mathematical formulation of the basis production and data expansion. We start by building a matrix $M_{ik} = V_i^k$ where each column k of the matrix M is a vector from the null set V. We perform SVD on this matrix,

$$M = u\Lambda v*, \tag{A1}$$

and then systematically zero all but one of the singular values (diagonal entries of Λ), which is set to 1,

$$\Lambda_{ij}^{\alpha} = \delta_{i\alpha}\delta_{j\alpha}. \tag{A2}$$

For example, Λ^3 is a matrix with a 1 in the (3, 3) position and zeros everywhere else. Then, for each index α, a basis vector e^{α} is produced by reversing the decomposition using the new matrix Λ^{α} and summing over the N columns of the reconstructed matrix M^{α},

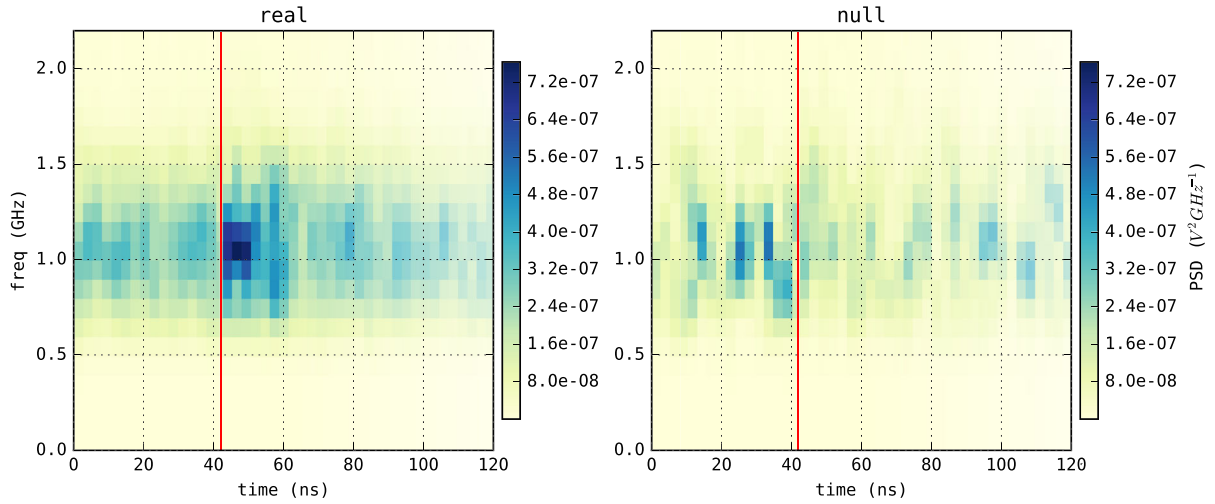$$M^{\alpha} = u\Lambda^{\alpha}v*, \tag{A3}$$

FIG. 20.   A comparison of the real (left) and null (right) data after the expansion in the null basis has been subtracted from each. These spectrograms are the averages of all normalized events in the sets.

$$e_i^\alpha = \sum_{k=1}^{N} M_{ik}^\alpha. \qquad (A4)$$

We can then take a data vector V^d and expand it in this basis (summation over repeated indices implied),

$$c^\alpha = V_i^d e_i^\alpha; \qquad (A5)$$

e.g., the expansion coefficient c^α is the inner product of the signal vector with the normalized basis vector e^α. The expansion of the real event into the null basis is the filter $f_i = c^\alpha e_i^\alpha$. This filter can be subtracted from the data vector $V_{\text{filtered}}^d = V^d - f$, leaving any unfiltered excess in the data vector. For this specific case, what should remain after this filtration is the putative scattered signal, which

was not present in the null set. The result of this procedure, for real and null sets, is presented in Fig. 20.

This figure is qualitatively similar to Fig. 14 though the absolute amplitude is not, as here the procedure is performed on normalized vectors. The signal-region excess is evident between real and null sets, which have the same color scale.

To investigate the effect of the basis in which the filter is constructed, we can reverse the procedure, build a basis out of the real set, and repeat the above procedure using this basis. That is, we expand real and null in the real basis, and use this expansion as a filter. This should remove everything from the null set.

However, since a real scatter is phase unstable on an event-by-event case, any remainder in the signal set would possibly be an indication of a signal. When building a basis,
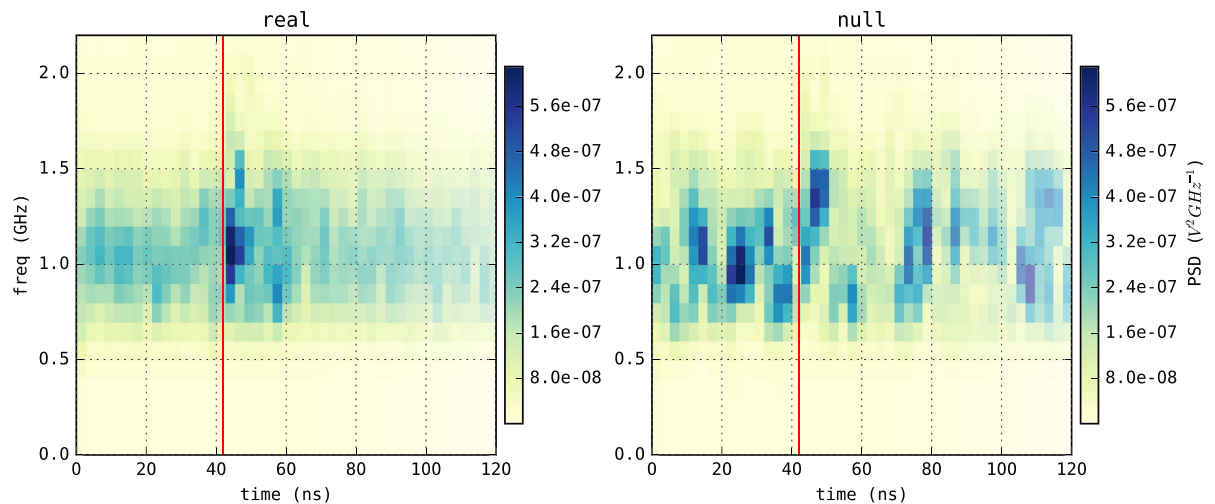

FIG. 21.   A comparison of the real (left) and null (right) data after the expansion in the real basis has been subtracted from each. These spectrograms are the averages of all normalized events in the sets.

the features which are most similar event to event (e.g., the beam splash) will be most prominent, but phase unstable features will be diminished.[4] In this case, the production of the basis vectors averages the reconstructed vectors

[see Eq. (A4)], which results in destructive interference for anything which is not phase stable.

The results of this procedure are shown in Fig. 21. We see that the null set indeed appears to be dominated by noise, and the real set is very quiet except for a small excess in the signal region. This is again explained by the lack of phase stability in the signal region. We note that this signal excess is smaller than the excess after filtration using the null basis (Fig. 20). Importantly, this procedure does not introduce signal-like artifacts into the null set.

---

[4]Contrast this with the method of Sec. V G, in which the overall sum was made on the spectrogram of the reconstructed events after the beam-splash patterns had been removed event by event. In that case, the individual phases of the signal modes did not add destructively because they were being added in power.

[1] P. Blackett and C. Lovell, Proc. R. Soc. A **177,** 183 (1941).

[2] K. D. de Vries, K. Hanson, and T. Meures, Astropart. Phys. **60,** 25 (2015).

[3] R. Abbasi *et al.*, Nucl. Instrum. Methods Phys. Res., Sect. A **767,** 322 (2014).

[4] P. Auger, Rev. Mod. Phys. **11,** 288 (1939).

[5] R. U. Abbasi *et al.*, Astropart. Phys. **87,** 1 (2017).

[6] P. Gorham, Astropart. Phys. **15,** 177 (2001).

[7] J. Stasielak, R. Engel, S. Baur, P. Neunteufel, R. Šmída, F. Werner, and H. Wilczyński, Astropart. Phys. **73,** 14 (2016).

[8] M. I. Bakunov, A. V. Maslov, A. L. Novokovskaya, and A. Kryemadhi, Astropart. Phys. **33,** 335 (2010).

[9] M. Chiba *et al.*, arXiv:0710.4186v1.

[10] K. de Vries, A. Murchadha, K. Hanson, and T. Meures, Proc. Sci., ICRC2015 (**2016**) 1168.

[11] K. DeVries *et al.*, Proc. Sci., ICRC2017 (**2017**) 1049.

[12] K. Belov *et al.* (T-510 Collaboration), Phys. Rev. Lett. **116,** 141103 (2016).

[13] A. Bean, J. P. Ralston, and J. Snow, Nucl. Instrum. Methods Phys. Res., Sect. A **596,** 172 (2008).

[14] K. D. de Vries *et al.*, Phys. Rev. D **98,** 123020 (2018).

[15] V. L. Ginzburg and I. M. Frank, Zh. Eksp. Teor. Fiz. **16,** 15 (1946) [J. Phys. (USSR) **9,** 353 (1945)].

[16] P. W. Gorham, D. P. Saltzberg, P. Schoessow, W. Gai, J. G. Power, R. Konecny, and M. E. Conde, Phys. Rev. E **62,** 8590 (2000).

[17] K. D. de Vries, S. Buitink, N. van Eijndhoven, T. Meures, A. Murchadha, and O. Scholten, Astropart. Phys. **74,** 96 (2016).

[18] G. A. Askar'yan, Zh. Eksp. Teor. Fiz. **41,** 616 (1961) [Sov. Phys. JETP **14,** 441 (1962)].

[19] D. Saltzberg, P. Gorham, D. Walz, C. Field, R. Iverson, A. Odian, G. Resch, P. Schoessow, and D. Williams, Phys. Rev. Lett. **86,** 2802 (2001).

[20] P. W. Gorham *et al.* (ANITA Collaboration), Phys. Rev. D **98,** 022001 (2018).

[21] J. Von Neumann, *Mathematical Foundations of Quantum Mechanics: New Edition* (Princeton University Press, Princeton, NJ, 2018).

[22] S. Prohira and D. Besson, Nucl. Instrum. Methods Phys. Res., Sect. A **922,** 161 (2019).

[23] S. Prohira, https://github.com/prchyr/RadioScatter, 2017.

[24] M. P. De Haas, M. Kunst, J. M. Warman, and J. B. Verberne, J. Phys. Chem. **87,** 4089 (1983).

[25] J. M. Roberts, E. Oberla, P. Allison, G. S. Varner, S. Spack, B. Fox, and B. Rotter, Nucl. Instrum. Methods Phys. Res., Sect. A **925,** 92 (2019).