

Automated Detection of the Magnetopause for Space Weather from the IMAGE Satellite

Michael L. Rilee^a and James L. Green^b

^aRaytheon Information Technology & Scientific Services

^bNASA Goddard Space Flight Center, Mailstop 630
Greenbelt, MD 20771 USA

ABSTRACT

The Radio Plasma Imager (RPI) is a low power radar on board the IMAGE spacecraft to be launched early in year 2000. The principal science objective of RPI is to characterize the plasma in the Earth's magnetosphere by radio frequency imaging. A key product of RPI is the plasmagram, a map of radio signal strength vs. echo delay-time vs. frequency, on which magnetospheric structures appear as curves of varying intensity. Noise and other emissions will also appear on RPI plasmagrams and when strong enough will obscure the radar echoes. RPI echoes from the Earth's magnetopause will be of particular importance since the magnetopause is the first region that the solar wind impacts before producing geomagnetic storms. To aid in the analysis of RPI plasmagrams and find all echoes from the Earth's magnetopause, a computer program has been developed to automatically detect and enhance the radar echoes. The technique presented is derived within a Bayesian framework and centers on the construction and analysis of a Likelihood Function connecting magnetospheric structures and RPI plasmagrams. Once this technique has been perfected on archival IMAGE data it will be recoded and used on board the IMAGE spacecraft in a series of tests thereby greatly facilitating organizations like the National Oceanic and Atmospheric Administration's (NOAA) Space Environment Center (SEC) to perform real-time analysis of space weather.

Keywords: IMAGE, Radio Plasma Imaging, RPI, Bayesian analysis, Likelihood

1. INTRODUCTION

The sun has been found to have a very dynamic and turbulent upper atmosphere that continually flows outward and is called the solar wind. Enormous ejections of matter into the solar wind have been found to be the cause of a majority of disturbances in the Earth's magnetosphere, the magnetic region that protects the Earth. These disturbances are called geomagnetic storms, the result of which can produce beautiful auroras as seen in Canada and the Northern United States. But these storms also produce massive power outages and adversely affect a large number of orbiting spacecraft, which perform a variety of telecommunication and global positioning functions. As our society grows more dependent on spacecraft-based technology, it is imperative that we learn how to detect geomagnetic storms as part of our space weather program.

NASA's Imager for Magnetopause-to-Aurora Global Exploration (IMAGE)¹ has been designed to study the global response of Earth's magnetosphere to changes in the solar wind during its two-year mission in space. The satellite carries some of the most sophisticated imaging instruments ever flown in near-Earth orbit and should be able to help researchers predict space weather.

One way of accomplishing NASA's strategic goals in space weather research is to develop techniques that will provide rapid detection of major changes in the configuration of the magnetosphere caused by the solar wind. It is well known that large variations in the solar wind produce significant geomagnetic storms. One must actually detect the motion and deformation of the magnetopause in order to observe the solar wind energy that is actually received at this boundary.

The IMAGE mission coincides with solar maximum, which is a period of intense solar activity that occurs every 11 years. During solar maximum the Earth is continually buffeted by explosive eruptions of plasma from the Sun. Researchers hope to gather significant information about the magnetosphere during this intense period of solar

Technical correspondence: M. L. Rilee, NASA/GSFC Mailstop 931. Other author information: M.L.R.: E-mail: Michael.L.Rilee.1@gsfc.nasa.gov; J.L.G.: E-mail: James.L.Green.1@gsfc.nasa.gov.

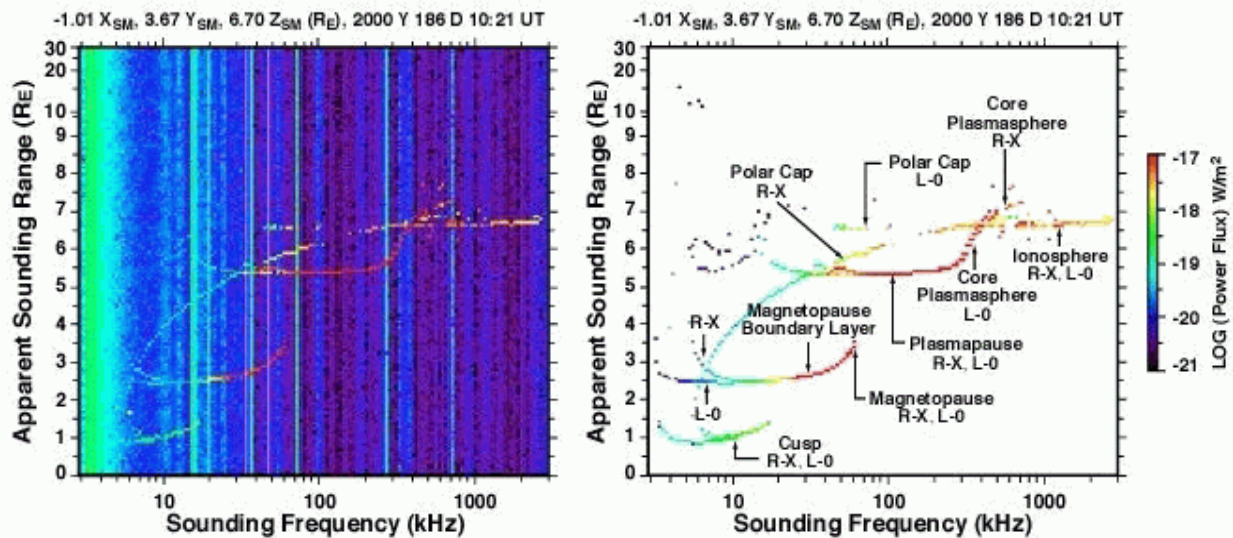


Figure 1. RPI plasmagrams showing the total signal power (left, Panel A) and the simulated echo power and magnetospheric feature identifications (right, Panel B). A plasmagram shows radio sounder echoes that map magnetospheric density as a function of distance from the spacecraft.

activity. NOAA has made arrangement with NASA to receive the IMAGE data in real-time and use it their space weather forecast center in Boulder, Colorado. NOAA's Space Environment Center (SEC) is responsible for sending out space weather alert to customers all over the world.

1.1. Radio Plasma Imager

One of the remote sensing instruments on IMAGE is designed to observe the structure and dynamics of the magnetospheric boundaries during geomagnetic storms using the Radio Plasma Imager or RPI. The RPI instrument consists of a radio transmitter, receiver, and 3 axis antenna systems.² The RPI concept is based on radar, or radio sounding techniques, to remotely measure plasma densities and motion.³ The RPI uses frequencies from 3 kHz to 3 MHz that will measure plasma densities from about 0.1 to 10^5 cm^{-3} .

Just like a radar, RPI transmits short pulses of electromagnetic waves into the space plasma and detects the reflected pulses. The RPI transmitted waves reflect at plasma cutoffs, where the index of refraction of the wave becomes zero. The primary presentation of RPI data will be in the form of plasmagrams. A plasmagram is a color coded plot of the power in an echo as a function of frequency and echo delay. A simulated plasmagram from ray tracing calculations is shown in Fig. 1. Both plasmagrams in Fig. 1 are identical except that the plasmagram in Panel A is the best representation of the RPI observations since it contains, in addition to the simulated echoes, the measured RPI instrument noise level and the local thermal plasma emissions. Panel B of Fig. 1 shows only the simulated echoes of Panel A. The RPI simulated echoes are presented in the form of echo time delay (t), expressed in terms of apparent range (right scale), as a function of the sounder frequency (bottom scale). The apparent range corresponds to $ct/2$ where c is the speed of light. The intensity is color-coded and has been calculated for each echo in this figure.

The simulated echoes in panel B of Fig. 1 are labeled with their appropriate polarizations (right hand-extraordinary or R-X and left hand-ordinary or L-O). The plasmagram contains the simulated echo measurements from a complete RPI instrument cycle when the spacecraft is near apogee.

The key to understanding magnetospheric dynamics is in the analysis of the RPI echoes in Fig. 1. These echoes must be extracted from the other magnetospheric emissions and the instrument noise background. Once the echo

traces are extracted information such as the distance and motion of the magnetopause can be determined. Currently there is no automated process for performing this data extraction. In addition, if the magnetopause is moving rapidly the pre-selected instrument mode may be such that the echo data is degraded beyond recognition. Only on board analysis and instrument echo parameter changes would significantly improve the real-time RPI echo data under certain circumstances.

1.2. Echo Extraction

The technique we use for the extraction of echoes is derived within a Bayesian framework because this allows us to quantitatively assess evidence for magnetospheric signatures in the RPI data. Furthermore, this framework allows the results of our analysis, which may be summarized as a posterior probability defined on the space of competing hypotheses, to be used as a prior or base for further analysis. Conceptually, this framework fits well the needs of space weather forecasting and nowcasting where the possible range of past, current, and future natural parameters is more important than any single most favored value.

Bayesian approaches have achieved good results, particularly in applications with low signal-to-noise ratios (SNRs) or sparse or missing data.⁴⁻⁶ These approaches are based on the Likelihood Principle which states that when making inferences from observations, a probability density known as the likelihood function represents all the knowledge we wish to bring to bear on the problem of connecting data with theory.^{7,8} By constructing a likelihood function that connects echo models with RPI plasmagrams, we can take advantage of methods of reasoning with probabilities to cast quantitative questions about our hypotheses. The key hypothesis we wish to test concerns the existence of magnetopause echo power in plasmagram regions dominated by noise; such a region can be seen in Fig. 1 below ~ 7 kHz. The quantitative test we perform is a likelihood ratio test where we compare models with and without echo power to construct odds for the hypothesis that echoes exist in the noise dominated region.

We now describe the process of extracting echo information from the plasmagram. The treatment of high frequency signals is discussed first. Following that the echo trace model and the likelihood function are presented. Then the application of the echo extraction to plasmagram data is presented. This paper concludes with remarks about the future directions of this research.

2. TECHNIQUE

The fundamental goal of this work is to enable the detection of magnetopause signatures in RPI data in the presence of strong interference. A first step towards this goal is to extract signatures of echo traces from the plasmagram. To reduce the dimensionality of the problem and thereby facilitate the solution's implementation and performance we treat strong and weak echoes differently. Strong echoes are easily discernable against the background noise; weak echoes are not. For the present calculation, we note that the high frequency portion of a magnetopause echo trace is easily discernable. For lower frequencies we construct a likelihood function that describes the RPI output and plasmagram echo traces.

2.1. The case of strong signals

In the strong signal region, signal detection can be achieved using simple thresholding. In our implementation, suitable low-noise frequencies are chosen and searched; echo candidates are identified by SNRs greater than an arbitrary threshold, here 5. Once identified, a simple notion of continuity is used to trace the echo through adjacent frequencies. Identified echoes that are within user given time-delay and frequency tolerances of each other are considered to be part of the same echo trace.

Once an entire set of echo traces has been collected, the trace is searched for branches; these show up as multiple echoes occurring on the same frequency. These branches are then broken into single valued subtraces, which are then pieced back together to form a set of traces, each element of which is single valued in time-delay and have the correct high frequency behavior. Thus an echo trace that has low frequency R-X and L-O wave branches is identified and stored as two traces which share the same high frequency branch but differ at the low frequencies. In this way, we identify all echo traces meeting provided degrees of time-delay/frequency continuity and signal strength.

Figure 2 shows an example of how these identifications are used. In this portion of a plasmagram, points from a strong echo trace have been chosen as control points for a spline model. From our ray tracing models of RPI sounding, we expect that this echo trace might extend to lower frequencies (compare Panels B of Figs. 1 and 2).

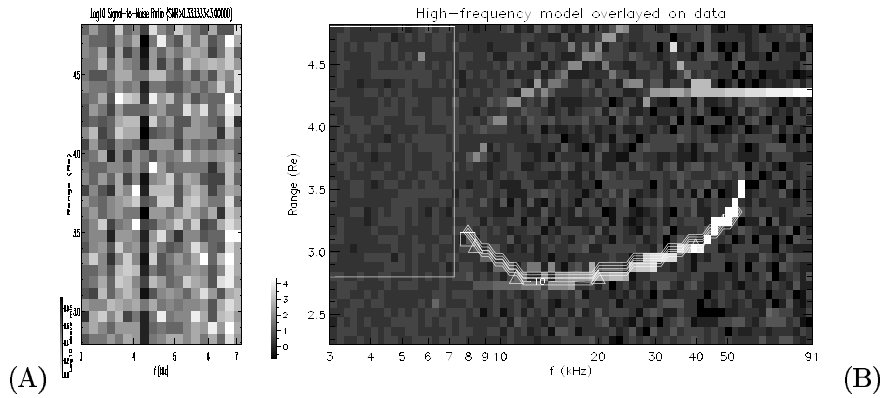


Figure 2. The signal-to-noise ratio at low frequencies (A), and detail from an *IMAGE/RPI* plasmagram (B). Displayed is the observed power $\log_{10} S$ overlotted by the spline-based model \tilde{S} (see Sect. 2.2.1 for details on S and \tilde{S}). The large box in the upper left outlines the region of strong natural noise emissions that we wish to search for evidence of RPI echoes; the SNR in this region is plotted in Panel A. We search this data for continuations of the high signal-to-noise echo visible and marked between 7 and 52 kHz.

However the comparatively strong noise at low frequencies obscures the trace. Figure 2 (A) shows an example of the SNR for low frequencies. In this figure, data outside the range $1/3 < \text{SNR} < 3$ are truncated to the lower and upper bounds. This figure effectively demonstrates that in this region of the plasmagram the signal is hidden from visual inspection by the noise.

2.2. Model based approach

A quantitative model is required to make quantitative assessments of the existence of echo power. Above ~ 7.5 kHz the echo power is strong enough that we may be somewhat cavalier in the formation of our model. Below ~ 7.5 kHz, we must be more careful because the SNR is very small (Fig. 2). From instrument tests, the RPI's response to an input signal was observed to be reasonably well described by a Poisson density about a mean (digitized) signal power. Proceeding from this observation, we can construct a likelihood function that models the probability that a particular data set will be obtained given a particular input power. Our certainty that echoes exist can then be cast as a probability within the Bayesian framework; the likelihood function is to be interpreted within this framework.

Bayes' rule provides a framework for combining models, data, and one's *a priori* expectations to help draw conclusions:

$$p(\tilde{S}|SI) = \frac{\text{likelihood} \times \text{prior}}{\text{evidence}} = \frac{p(S|\tilde{S}I)p(\tilde{S}|I)}{p(S|I)}.$$

The approach used here is based on the work of Loredo.⁵ The symbol I denotes the specific context defined by the models of the situation. It stands for the information required to describe the modeled physical situation, e.g. the distributions used, the physics involved, and other *meta-model* information. Different models are labeled via I and compared using standard probabilistic methods.

The posterior probability $p(\tilde{S}|SI)$ is the measure of one's confidence that a particular hypothesis or set of parameters \tilde{S} describes the reality within which the data S was obtained. In this paper, tildes ($\tilde{\cdot}$) denote model quantities; many quantities are implicitly functions of time (echo delay time or apparent range) and frequency; these relations will be made explicit when necessary.

The likelihood $p(S|\tilde{S}I)$ is the model that connects experimental observations to parameters associated with important features of the physical world. It is one's certainty that the data set S would be obtained when the physical situation $\tilde{S}I$ prevails. The variable parameters that determine the nature of the physical system are denoted by \tilde{S} .

Finally, *a priori* information is written as $p(\tilde{S}|I)$ and can include information \tilde{S} from previous experiments as well as an educated guess. One popular method of determining priors has been to try to minimize their information

content and hence their impact on the final analysis. This is the basis of the Maximum Entropy Method.⁹ In the present work, only the likelihood or likelihood ratios are used while we neglect (explicit) priors. Hence the evidence or marginal density $p(S|I)$, which acts as a normalization constant and depends on the prior, is neither available nor required.^{6,7}

2.2.1. Instrument and signal model

In this section we present the model of the RPI echoes. This model was developed independently of the models used to construct the simulated plasmagram shown in Fig. 1. Beyond the observation of Poisson statistics there is little physics involved in the model developed here. Instead, the basic assumption is that echo traces on plasmagrams can be described as relatively continuous arcs. Detailed models based on other more physical models of the Earth’s magnetosphere can be handled in a way similar to what is presented here, but such modelling is beyond the scope of the present work.

Within the context of RPI, we assume the signal and noise to be additive

$$\tilde{S} = \sum_x \tilde{S}_x = \tilde{S}_{echo} + \tilde{S}_{noise},$$

where x covers the domain of phenomena that add to the RPI response. Here x is restricted to echo and noise power. The echo trace is factored into portions responsible for the echo strength W_{echo} and apparent range C_{echo} as functions of frequency f :

$$\tilde{S}_{echo} = W_{echo} C_{echo}.$$

W_{echo} has units of power or energy as appropriate and C_{echo} is dimensionless. After examining several different forms, the form of the echo strength is a log-log-polynomial with coefficients η_k^w :

$$\ln W_{echo} \equiv \ln W_{echo}(f) = \sum_{k=0}^2 \eta_k^w (\ln f)^k.$$

The form of the apparent range is

$$\ln C_{echo} \equiv \ln C_{echo}(t, f) = -a (t - \tilde{t}(f))^2.$$

The parameter a is the reciprocal width of the modeled echo trace, and for the work here a is set so that the model traces are not resolved in time. Here and hereafter times t are presented as apparent ranges measured in Earth radii. The maxima of C_{echo} occur when the right hand side of the previous equation is zero. The locations of these zeros as functions of frequency are given by

$$\tilde{t}(f) = \text{Spline}_k(\eta_k^c, \ln f).$$

The notation “ $\text{Spline}_k(\eta_k^c, \ln f)$ ” denotes the spline function through the N^c points $\eta_1^c, \eta_2^c, \dots, \eta_{N^c}^c$; the domain of the function is $\ln f$. For the implementation used, the model parameters η_k^c are the spline’s control points (t_k, f_k) . We typically used cubic splines to control the model’s echo return time. Thus the parameters which define this model echo trace are $\eta_{echo} = \{\{\eta_k^w\}_{k=0}^2, \{\eta_k^c\}_{k=1}^{N^c}, a\}$.

The noise model is written

$$\tilde{S}_{noise} = W_{noise} \equiv W_{noise}(f),$$

where W_{noise} is only a function of frequency for this work, but could be a more general function. For most purposes, W_{noise} can be determined during passive RPI observations. For consistency with the echo model, we write the parameters that define the noise model as η_{noise} ; in this case we can simply take $W_{noise}(f) = \eta_{noise}(f)$, independent of time. On the other hand, if the noise environment has more structure in time (and frequency) we can include these affects in our models at the expense of increasing the dimensionality of the model parameter space. With suitable prior information the technique of marginalizing such extra parameters can help limit the dimensionality of the problem.

2.2.2. Likelihood function

Since the RPI output appears adequately modeled by Poisson densities, the Poisson density can model low strength digitized signals and approaches Gaussian statistics for high strength signals. The observations of RPI output are analog-to-digital (A/D) digitizer output, written as $S\delta t$, above some prescribed minimum average power and arriving in a time interval δt . For this work we suppose that the prescribed minimum is set so that the background noise power input causes the digitizer to emit $S\delta t = 3$, independent of frequency. Parenthetically, in this work $S\delta t$ is real valued, not integer valued.* It is expected that new models will be constructed to adequately describe RPI observations once it is in orbit.

The model connecting observed data with physical parameters is the likelihood:

$$p(S|\tilde{S}I) = \frac{(\tilde{S}\delta t)^{S\delta t} e^{-\tilde{S}\delta t}}{(S\delta t)!} = \frac{(\sum_x S_x \tilde{S}_x \delta t)^{S\delta t} e^{-\sum_x \tilde{S}_x \delta t}}{(S\delta t)!}$$

The duration of a given observation δt can be a function of time and frequency and is absorbed into the signal and noise powers in what follows. It is more convenient to work with the log-likelihood:

$$\ln p(S|\tilde{S}I) = S \ln \tilde{S} - \tilde{S} - \ln(S)! \approx -S \ln \frac{S}{\tilde{S}} - \tilde{S},$$

where Stirling's approximation has been used.

The signals are measured and modeled on a set of discrete times and frequencies, here denoted as a set of elements tf . If we take the statistics of measurements obtained at different tf to be independent, we can write the log-likelihood as

$$\ln p(S|\tilde{S}I) = \sum_{tf} \ln p(S_{tf}|\tilde{S}_{tf}I).$$

For this work we neglect priors, normalization and deal only with the log-likelihood, $\ln p(S|\tilde{S}I)$:

$$\begin{aligned} \ln p(S|\tilde{S}I) &= \\ &= \ln p(S|\eta_{echo}, \eta_{noise}, I) \\ &\approx - \sum_{tf} \left(S_{tf} \ln \frac{S_{tf}}{\tilde{S}_{tf}} + \tilde{S}_{tf} \right) \end{aligned} \quad (1)$$

where $\tilde{S}_{tf} \equiv \tilde{S}_{tf}(\eta_{echo}, \eta_{noise})$. Equation (1) is the foundation of our method. A Maximum Likelihood (ML) approach would entail a search for the maximum of Eq. (1). The analyses presented below are not strict ML approaches. Instead, we present convenient maps of the likelihood that show where RPI echo traces may be hidden in noisy data. Within the Bayesian framework, we would multiply the likelihood by an a priori density to obtain the posterior density which would then summarize our quantitative conclusions. Therefore, the likelihoods presented are proportional to the posteriors and can likewise show which parameters, and hence models, are favored by the data.

3. APPLICATION TO MODELED IMAGE/RPI DATA

3.1. The case of strong signals

As mentioned above, the echo power at higher frequencies does not require special analysis; echo traces with strong SNRs are readily extracted from the data (Sect. 2.1). In fact, we use the high SNR data available above ~ 7.5 kHz to aid in the recovery of the low SNR data below that frequency. A convenient representation is obtained if we allow only the parameters determining the low frequency endpoint of the model echo trace to vary; all other parameters are set to values determined from the high SNR data. If the spline control parameters $\eta_k^c = (t_k, f_k)$ are ordered by increasing frequency, then $\eta_1^c = (t_1, f_1)$ will be the low frequency endpoint of the model echo trace and the only parameters allowed to vary. We determine the form of the echo power η_{echo} , i.e. $\{\eta_k^w\}_{k=0}^2$, from the high SNR data and hold these values fixed for the low SNR data. We have set W_{noise} by time-averaging data that does

*Results of tests with integer valued measurements $S\delta t$ do not differ qualitatively from what is presented here.

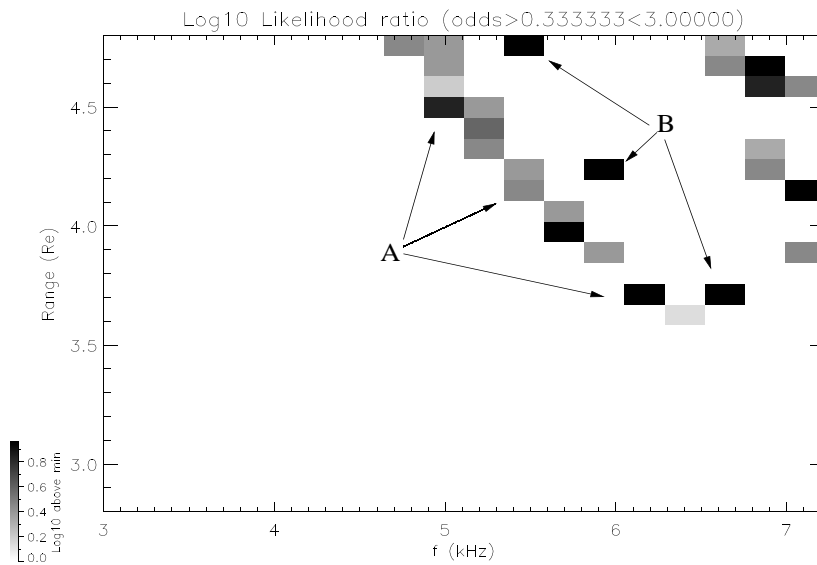


Figure 3. The likelihood-ratio: Odds that the RPI input contains echo+noise vs. solely noise. Possible echo traces are labeled A and B for comparison with figures below. The gray scales shown here are truncated below and above odds of 1/3 and 3 respectively.

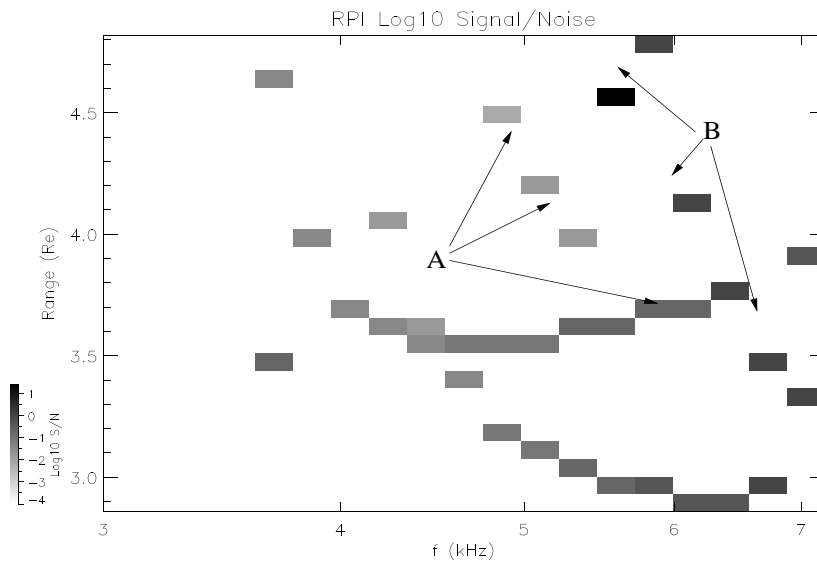


Figure 4. The actual signal-to-noise ratio calculated from the echo and noise power models used to construct the plasmagram. Note the location of the ridges of high likelihood labeled A and B in figure 3 are coincident with actual signals, some with SNR < 1.

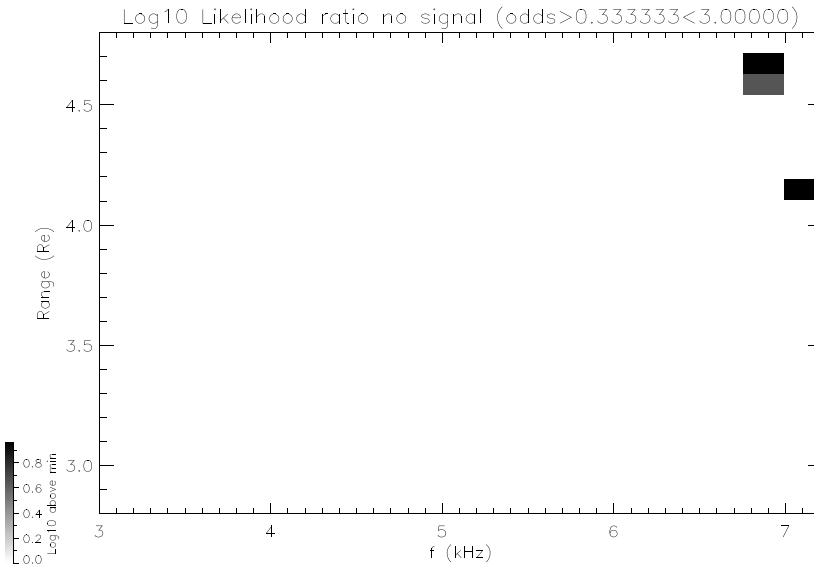


Figure 5. The odds ratio with the echo signatures artificially removed from the data. This test shows that the likelihood approach is actually sensitive to the presence of the weak echoes and is not serendipitously matching the traces.

not contain echoes. This is likely adequate for present purposes because noise environment will likely change slowly during observations at any one frequency.

The likelihood mapped as a function of time-delay and frequency, η_1^c , resembles a filtered plasmagram, a feature that greatly aids interpretation. Therefore, aside from the endpoint η_1^c , all parameters defining the echo and noise components of the signal are determined from the high SNR data or during echo-free observations.

3.2. Application to low frequency data

Figure 2 shows that simple methods of signal detection based on estimated SNRs do not provide much information for very weak signals. With the likelihood based method, echo detections appear as enhancements in regions of favored parameters. Furthermore, because of our probabilistic interpretation, we can calculate the odds that a given input signal is composed of echo+noise power or solely noise power. To calculate this odds, we form the ratio of the likelihood with to that without echo power. This is essentially a Bayesian model comparison. Finally, by forming the ratio of the odds to the SNR we obtain a rough and ready gauge of the SNR enhancement possible with this technique.

Figures 3 through 7 summarize the effectiveness of our method for this model problem. In this case the trace to be found is the upper branch of the magnetopause echo. Figure 3 shows the likelihood-based odds ratio. Data outside the range 1/3 to 3 have their gray scales truncated. The application of our method dramatically shows where evidence supporting the model is to be found. We have thus found evidence for echo traces whose SNRs are below one (Fig. 4). It is natural to be concerned about the reality of these detections.

To gauge the trustworthiness of these detections, we artificially removed the signal and repeated our calculations with only noise power existing below ~ 7.5 kHz. The results are presented in Fig. 5 where detections are to be considered as false positives. If the method were merely fitting the noise, then the structures visible in Fig. 3 should still be apparent. Therefore, we conclude that our method's rate of false positives is low and its signal to noise enhancement can detect a certain class of signals with $\text{SNR} < 1$.

3.3. Benefit of the method

As mentioned above, a simple way to gauge the benefit of this technique is to form an effective SNR enhancement factor from the odds and the SNR. We do not claim that this estimate is an optimal measure of the risks of this method, but it is convenient and suits our present purpose. From this viewpoint, Fig. 6 shows the estimated

enhancement in the SNR that our method provides over the plasmagram data. Figure 7 shows how SNR data points are mapped to odds. The large effective enhancements for regions of strong echo power (e.g. trace B) in such noisy data may lead to the identification of physical phenomena that would otherwise be missed. An additional important benefit is the identification of regions of less, but non-negligible odds (e.g. trace A) which may also point to important but hard to detect magnetospheric structures.

3.4. Computational performance

Though the current implementation is an experiment in feature extraction, a few words about the current optimization state of the code are appropriate because the IMAGE spacecraft is a prospective target platform for these or similar feature extraction functions. A likelihood ratio map such as Fig. 3 takes about 30 minutes to produce on a Digital 166 MHz Pentium MMX system with 144 MB of memory running Red Hat Linux (5.2). The current implementation is written in Research System Incorporated's Interactive Data Language (IDL) which is byte-compiled, interpreted, and has an extensive library of tuned object code. The IDL implementation of our method is in no sense optimally fast, but with sufficient effort, the efficiency of this method can likely be increased by a factors tens or more. Sources of optimization include algebraic simplifications of the underlying likelihood ratios, inlining of function calls, and recoding to make use of an optimizing compiler. The use of integer instead of floating point operations may enhance efficiency as well. Furthermore, a dynamically adaptive algorithm that concentrates compute resources on more likely regions of parameter space would be less computationally expensive than the grid based searches presented here. The excellent performance of this method at extracting echoes and its adaptability to other probabilistic analyses prompts us to seek ways to overcome the challenge of developing an on board implementation.

4. CONCLUSIONS

This paper provides evidence that useful information can be extracted from RPI simulation data that is apparently dominated by noise. This information extraction is possible because (1) the echo traces place a sufficient mark on the data and (2) the likelihood model describing the plasmagram data is sufficiently accurate. If either of these two features were lacking, then the performance of this method would be reduced. Furthermore, the use of the strong echo data reduces the effective dimensionality of the parameter space on which the likelihood is defined. The method's ability to extract echo traces exhibited in this paper is likely better than will be obtained during the method's application to RPI operations because the echo and noise models will likely require more parameters to describe reality less well.

Fortunately, we have so far found this method of analysis surprisingly robust to inadequacies or even errors in the model. For example, during one set of runs the echo and noise power levels were wrongly set a factor of nine too low. The SNR was still correct, but the probability densities describing the expected response of the instrument were quite far off. Still, the method returned qualitatively similar results: the main difference being a much less dramatic response of the method to echo trace B.

4.1. Including improved models

During RPI operations, we expect to take better advantage of our understanding of the instrument, its environment, and the interaction of the two. A straightforward extension of the echo extraction method would be the inclusion of knowledge about the location of magnetospheric structures via prior probabilities (Sect. 2.2). These priors could be constructed from archived or current spacecraft data and/or models coupling indicators such as the Disturbed Storm-Time index (DST)¹⁰ to magnetospheric structure location.

Also in the current work, we have not taken advantage of important information about the instrument. For example, our results here used model data for only one of RPI's antenna-receiver subsystems. RPI contains two other such subsystems and an operational version of our analysis would include a model describing all three subsystems and their response. Likelihood analyses using all three subsystems will likely outperform analyses using only one.

4.2. The next step

The technique discussed provides the ability to automate the extraction of echoes from the natural noise and instrument background. Once these echoes have been identified in the data then, to a first order approximation, the distance to the target (the magnetopause in this case) can simply be calculated as $R = ct/2$, where R is the distance (in km), c the speed of light (in km/sec), and t the total time of the echo (in sec). With the long three-axis RPI




Figure 6. The ratio of the odds to the actual signal-to-noise ratio as a function of frequency. This is a crude measure of the SNR improvement possible for regions with weak signals, provided the appropriate models are available. When the model well describes the situation, then dramatic inferences are possible. Boxes at points A and B correspond to the echo traces A and B in Fig. 3. Boxes are points with odds above $1/3$; triangles between 10^{-3} and $1/3$; crosses below 10^{-3} . Symbols along top axis do not have a matching nonzero SNR.

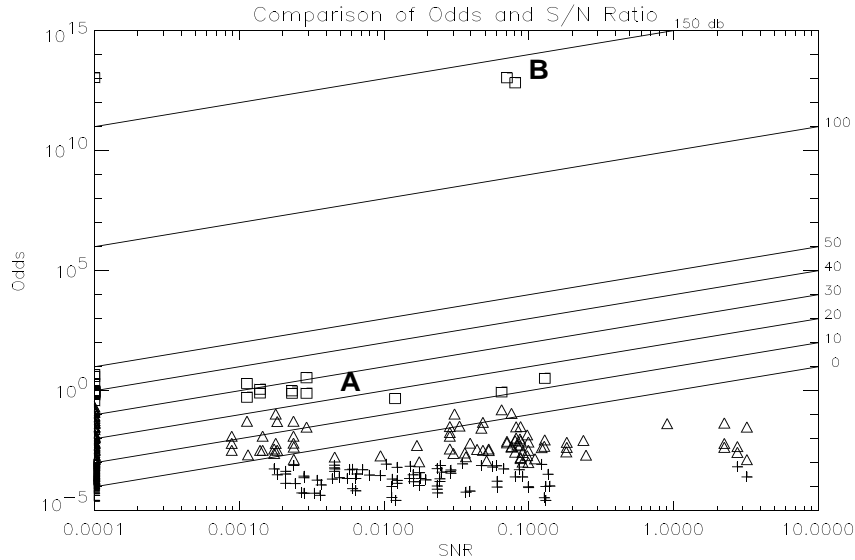


Figure 7. The odds ratio compared with the actual signal to noise ratio. The diagonal lines mark the ratio of the odds to the SNR as measured in decibels (db). Note that odds ratio shows that the evidence for an echo at trace B is overwhelmingly strong. The analysis recommends much less, but still non-negligible, certainty regarding trace A. Points in the lower right, i.e. those that have a strong SNR but weak odds, are not favored by the method for inclusion in an echo trace. Symbols signify the same quantities as in Fig. 6. Symbols along left axis correspond to points that do not have a matching nonzero SNR.

antennas on IMAGE the direction of arrival can also be calculated. It is expected that the echo extraction, range calculation, and direction of arrival will be calculated on board IMAGE and telemetered to the ground to use by NOAA's SEC. These telemetered results will then be used to infer magnetospheric structure and to generate space weather alerts.

At the time of this writing the IMAGE mission has not yet been launched and this analysis has been done with simulated data using all measured RPI instrument characteristics. Once RPI becomes operational, approximately 40 days after launch, this technique will be applied to the archival data then re-coded and loaded onto the on board computers on IMAGE. As noted above, we have identified a number of ways in which the efficiency of the present implementation can be greatly enhanced, thereby enabling on board echo extraction. It is expected that the first end to end test of this type of space weather alert will occur after the main mission of IMAGE has been completed (April 2002).

4.3. Application to Space Weather research

One way of accomplishing NASA's strategic goals in Space Weather research is to develop techniques that will provide rapid detection of major changes in the configuration of the magnetosphere by the solar wind. It is well known that large variations in the solar wind produce significant geomagnetic storms, but solar wind measurements alone are not always enough to predict these storms. The deformation and motion of the Earth's magnetopause are major components of the structure and behavior of the magnetosphere. Therefore measurements of the magnetopause are critical to our understanding and prediction of magnetospheric phenomena. The RPI instrument is a highly flexible multi-mode instrument designed to extract by remote sensing the signatures of the magnetopause and other major magnetospheric boundaries and plasmas. The first instrument of its kind flown to illuminate the magnetosphere, RPI will provide a new and unique view of magnetospheric dynamics. The goal of this research is to develop techniques for the on-board analysis of RPI echoes from the magnetopause in real-time that can be used by NOAA's SEC for generating space weather alerts. The first step in this process, the extraction of echoes from RPI data using simulated data, has been completed.

ACKNOWLEDGMENTS

The authors wish to gratefully acknowledge W. W. L. Taylor, S. Boardsen, and L. Stone for many helpful discussions. This work has been supported by the Cross Enterprise Program (632-67 Thinking Space Systems) and performed under contracts NAS5-30960 and NASW-97002. Also acknowledged is support from the Remote Exploration and Experimentation Project of NASA's High Performance Computing and Communications Program which is funded by the Office of Space Science with the work performed under contract NAS5-32350.

REFERENCES

1. J. L. Burch, "IMAGE Mission Overview," to be published in *Space Science Reviews*, February, 2000. IMAGE special issue.
2. B. W. Reinisch *et al.*, "The Radio Plasma Imager investigation on the IMAGE spacecraft," to be published in *Space Science Reviews*, February, 2000. IMAGE special issue.
3. J. L. Green *et al.*, "Radio Plasma Imager Simulations and Measurements," to be published in *Space Science Reviews*, February, 2000. IMAGE special issue.
4. T. J. Loredo, "From Laplace to Supernova SN 1987a: Bayesian Inference in Astrophysics," in *Maximum Entropy and Bayesian Methods*, P. F. Fougère, ed., pp. 81–142, Kluwer Academic Publishers, 1990.
5. T. J. Loredo, "The Promise of Bayesian Inference for Astrophysics," in *Statistical Challenges in Modern Astronomy*, E. D. Feigelson and G. J. Babu, eds., pp. 275–297, Springer Verlag, 1992. Unabridged version.
6. J. K. Ó Ruanaidh and W. J. Fitzgerald, *Numerical Bayesian Methods Applied to Signal Processing*, Springer-Verlag, 1996.
7. J. O. Berger, *Statistical Decision Theory and Bayesian Analysis*, Springer-Verlag, second ed., 1985.
8. L. D. Stone, C. A. Barlow, and T. L. Corwin, *Bayesian Multiple Target Tracking*, Artech House Inc., 1999.
9. E. T. Jaynes, "Information theory and statistical mechanics," *Phys. Rev.* **106**, p. 620, 1957.
10. M. Sugiura, "Hourly values of equatorial Dst for IGY," in *Annals of the International Geophysical Year*, vol. 35, pp. 945–948, Pergamon Press, (Oxford), 1964.