

Maximum Likelihood Estimation of Phase Screen Parameters from Ionospheric Scintillation Spectra

Charles S. Carrano, Charles L. Rino, and Keith M. Groves

Institute for Scientific Research, Boston College, 140 Commonwealth Ave., Boston, MA, 02467, USA

Abstract

At the previous Ionospheric Effects Symposium, we presented an extension of the phase screen theory of ionospheric scintillation for the case where the refractive index irregularities follow a two-component power law spectrum. A specific normalization was invoked to achieve a universal scaling, such that different combinations of perturbation strength, propagation distance, and frequency produce the same results. The theory is equally valid in the weak, intermediate, and strong scatter regimes. In this paper we consider the inverse problem, whereby phase screen parameters are inferred from measured scintillation time series as a means of interpreting the irregularities physically. The screen parameters are obtained by fitting the spectral density function (SDF) of intensity fluctuations with a parametrized theoretical model using the Maximum Likelihood (ML) technique. We refer to this fitting procedure as Irregularity Parameter Estimation (IPE) since it provides a statistical description of the refractive index irregularities from the scintillations they produce. In this paper, we introduce an additional rescaling that enables IPE to be applied without a-priori knowledge of the location of the irregularities or their motion.

1. INTRODUCTION

At the 2015 IES Meeting, we presented a phase screen solution for the spectrum of intensity scintillations when the refractive index irregularities follow a two-component power law spectrum (Carrano and Rino, 2016). Here, we consider the inverse problem, whereby phase screen parameters are inferred from measured scintillation time series as a means of interpreting the irregularities physically. We accomplish this by fitting the spectrum of intensity fluctuations with a parametrized theoretical model using the Maximum Likelihood (ML) technique. We refer to this fitting procedure as Irregularity Parameter Estimation (IPE) since it provides a statistical description of the refractive index irregularities from the scintillations they produce. In this sense, IPE may be thought of as stochastic back-propagation.

Previous applications of IPE have been implemented by minimizing the difference between measured and theoretical intensity spectral density functions (SDFs) in a least-squares sense (Carrano et al., 2012; Carrano and Rino, 2016). While accurate estimates of the screen parameters may be obtained via least-squares fitting, the estimates are not optimal because the fitting errors are not normally distributed (as they are assumed to be using this method). It is well known that the periodogram is scattered about its expectation according to a chi-squared distribution of order $2M$, where M is the number of periodograms averaged together to estimate the SDF (Appourchaux, 2003; Barret and Vaughan, 2012). In the limit of large M , the errors become normally distributed and the least-squares and ML techniques become equivalent. By using the known distribution of the errors, the ML technique can provide improved estimates of the screen parameters compared with least-squares estimation, particularly when the opportunity for averaging periodograms is limited. A parametric model of receiver noise is used to mitigate its influence on the screen parameter estimates. Monte Carlo simulations are used to characterize the probability distribution of the errors and verify the usefulness of confidence intervals for the screen parameters provided by the ML technique.

The one- and two-component power law irregularity models are nested, since the latter contains the former as a special case. Therefore, the likelihood will necessarily be equal or larger when fitting with the more general model. We show an example where the Akaike Information Criterion (Akaike, 1973) provides statistical justification for the more two-component irregularity model. We believe this is the first time evidence for a two-component irregularity spectrum has been rigorously substantiated using ground-based ionospheric scintillation measurements.

2. A THEORETICAL FITTING FUNCTION FOR SCINTILLATION

Carrano and Rino (2016) developed a theory for the spectral density function (SDF) of intensity fluctuations resulting from propagation through two-dimensional field-aligned irregularities in the ionosphere. A piecewise power-law phase screen model was considered with SDF

$$\Phi_{\delta\phi}(q) = C'_p \begin{cases} q^{-p_1}, & q \leq q_b \\ q_b^{p_2-p_1} q^{-p_2}, & q > q_b \end{cases} \quad (1)$$

In the above, C'_p is the phase spectral strength, q is wavenumber, q_b is the break wavenumber, and p_1 and p_2 are the phase spectral slopes at low and high wavenumbers, respectively. Let z be the propagation distance past the phase screen, and k be the free-space wavenumber of the propagating wave. As described in (Carrano and Rino, 2016), the propagation problem can be cast in universal form by normalizing quantities with the Fresnel scale $\rho_F = \sqrt{z/k}$. In particular, equations (11) and (12) of that paper give the normalized SDF of intensity fluctuations as a function of normalized wavenumber $\mu = q\rho_F$. The intensity SDF may be written explicitly in terms of the four parameters U , p_1 , p_2 , and μ_b as

$$I(\mu; U, p_1, p_2, \mu_b) = 2 \int_0^\infty \exp\{-\gamma(\eta, \mu)\} \cos(\mu\eta) d\eta \quad (2)$$

where γ is the so-called structure interaction function

$$\gamma(\eta, \mu) = U \cdot \begin{cases} \gamma_1(\eta, \mu) + \mu_b^{p_2-p_1} \gamma_2(\eta, \mu), & \mu_b \geq 1 \\ \mu_b^{p_1-p_2} \gamma_1(\eta, \mu) + \gamma_2(\eta, \mu), & \mu_b < 1 \end{cases} \quad (3)$$

and

$$\gamma_1(\eta, \mu) = 16 \int_0^{\mu_b} \chi^{-p_1} \sin^2(\chi\eta/2) \sin^2(\chi\mu/2) \frac{d\chi}{2\pi}, \quad \gamma_2(\eta, \mu) = 16 \int_{\mu_b}^\infty \chi^{-p_2} \sin^2(\chi\eta/2) \sin^2(\chi\mu/2) \frac{d\chi}{2\pi} \quad (4)$$

In practice, one does not measure the spatial distribution of intensity variations, but instead a time series resulting from irregularity and penetration point motion. In this case, temporal frequencies (f) in the observed scintillation spectra are related to spatial wavenumbers according to $\mu = 2\pi f / f_F$, where $f_F = V_{eff}/\rho_F$ is the Fresnel frequency and V_{eff} is the effective scan velocity defined in (Rino, 1979). This can be used to express the temporal spectrum of intensity fluctuations as a function of five fitting parameters U , p_1 , p_2 , μ_b , and f_F :

$$I(f; U, p_1, p_2, \mu_b, f_F) = 2 \int_0^\infty \exp\left\{-\gamma\left(\eta, \frac{2\pi f}{f_F}\right)\right\} \cos\left(\frac{2\pi f \eta}{f_F}\right) d\eta \quad (5)$$

In general, this fitting function must be evaluated numerically. Interestingly, the fitting function lacks an explicit dependence on the propagation geometry and even on Fresnel scale. As a consequence, scintillation spectra may be fit without *a-priori* knowledge of the distance to the phase screen or the details of irregularity and penetration point motion. Of course, if these details are known they can be used to reduce the number of parameters that must be estimated via spectral fitting. The model is equally valid in weak, moderate, and strong scatter conditions to the extent that a 1D phase screen model adequately describes the propagation physics. In practice, this last condition limits the applicability of the model to propagation geometries such that the line of sight scans through the field-aligned irregularities in an approximately transverse manner.

3. MAXIMUM LIKELIHOOD ESTIMATION

We are interested in fitting the parametric spectral model (5) to the spectrum $I^m(f)$ obtained from a measured scintillation time series. For notational convenience, let $\theta = (U, p_1, p_2, \mu_b, f_F)$ be the vector of fitting parameters. Let f_i be the i^{th} frequency in the periodogram and consider the ratio $R_i = I^m(f_i) / I(f_i; \theta)$ as a random variable. For a perfect model, R_i will follow a chi-squared distribution of order d , $R_i \sim \chi_d^2 / d$, where d is twice the number of periodograms M that are averaged together to estimate $I^m(f_i)$. The conditional probability of measuring the i^{th} harmonic, $I_i^m = I^m(f_i)$, given the model evaluated at this frequency, $I_i = I(f_i)$, is

$$p(I_i^m | I_i) \sim \frac{d}{I_i} \chi_d^2 \left(d \frac{I_i^m}{I_i} \right). \quad (6)$$

Assuming each harmonic is statistically independent from the others, the probability of measuring all harmonics in the spectrum given the model is the product of these conditional probabilities

$$p(I^m | I) \sim \prod_{i=1}^N \frac{d}{I_i} \chi_d^2 \left(d \frac{I_i^m}{I_i} \right). \quad (7)$$

The likelihood of the parameters given the measurements $L(\theta | I^m)$ is defined by considering the above as a function of the parameters rather than the measurements, i.e. $L(\theta | I^m) \equiv p(I^m | I(\theta))$. Our goal is to find the estimate $\hat{\theta}$ for the parameters that maximizes the likelihood function, or equivalently, the log likelihood function which can be expressed in convenient form as a summation:

$$\text{Log}(L) \sim \sum_{i=1}^N \frac{d}{I_i} \chi_d^2 \left(d \frac{I_i^m}{I_i} \right). \quad (8)$$

In practice, we minimize $S = -\text{Log}(L)$ numerically using the downhill simplex method. It can be shown that under fairly general conditions S is shaped like a multi-dimensional paraboloid in the vicinity of its minimum ($\hat{\theta}$), and $\Delta S \equiv \text{Log}(L(\theta)) - \text{Log}(L(\hat{\theta}))$ is distributed as χ_v^2 where v is the number of parameters of interest. Confidence intervals on each parameter may be computed in a manner analogous to mapping confidence intervals by the $\Delta\chi^2$ method, i.e. by noting where ΔS achieves a value corresponding to the confidence level of interest ($\Delta S = 2.71$ corresponds to 90 percent confidence limits on one parameter with $v=1$). Formal covariances may be computed by estimating the curvature of the log likelihood function at its maximum:

$$\sigma_{ij}^2 = (F)_{ij}^{-1}, \quad F_{ij} = - \left\langle \frac{\partial^2 \text{Log}(L)}{\partial \theta_i \partial \theta_j} \right\rangle \quad (9)$$

where the brackets indicate ensemble average. If only a single realization is available, we use the observed log likelihood to estimate the covariances and omit the averaging. We find it convenient to estimate the logarithm of non-negative parameters such as U and μ_b , transforming back to the linear domain after the spectral fitting is completed. In such cases, we also transform the confidence intervals accordingly. We experimented with different scale factors for S which has the effect of adjusting the confidence intervals to make them consistent with the results of Monte Carlo simulation over a reasonable range of phase screen configurations.

Example 1 - Fitting a Simulated Intensity Spectrum using a One-Component Irregularity Model

We demonstrate the technique by fitting simulated scintillation data with known phase screen parameters. In this example, we conducted a phase screen simulation of scintillation for a one-component power law spectrum with $p=p_1=p_2=3.0$. We specified the scattering strength $U = 0.6$ and Fresnel frequency $f_F = 2$ Hz. Figure 1a shows the intensity fluctuations from the phase screen simulation for which the scintillation index is $S_4 = 0.58$. We add complex white noise to the complex amplitude of the received wave to simulate the effects of receiver (thermal) noise, which affects one's ability to extract phase screen parameters from the observed signal fluctuations. We computed the spectrum of the simulated scintillation by averaging four periodograms taken over 60-second data segments (so that $M=4$, and $d=8$). This spectrum is shown as the black curve in Figure 1b.

We chose to estimate scattering strength and spectral index with the remaining parameters held fixed, so that $\theta = (U, p)$. Following Ruddick et al. (2000) we fit the spectral data with the sum of the theoretical SDF plus a parametric model for the noise (in this case, our noise model is simply a constant with known spectral density). The dashed red curve in Figure 1a shows the

theoretical SDF plus noise, while the solid red curve shows the theoretical SDF only. The S_4 index corresponding to the best fit theoretical SDF plus noise, while the solid red curve shows the theoretical SDF alone. The integral of the solid red curve over wavenumber gives $S_4 = 0.57$, which is very close to the measured S_4 . The best fit parameters $U = 0.7$ and $p = 3.05$, and their associated 90% confidence intervals obtained by the $\Delta\chi^2$ method are labeled on the plots. The confidence intervals for U and p contain their truth (specified) values in this case, although this cannot be guaranteed in general. Figures 1c and 1d respectively show the probability density function and cumulative density function (CDF) of the ratio R_i computed using the best fit solution. These compare very well with those of the chi-squared distribution (shown in green for comparison), which suggests that the assumed distribution of the errors from the MLE fit is correct. The Kolmogorov-Smirnov (K-S) test, which measures the largest deviation between the observed and theoretical CDFs, can be used to provide a quantitative goodness of fit measure. The K-S for this fit was small (0.045), suggesting a high-quality fit.

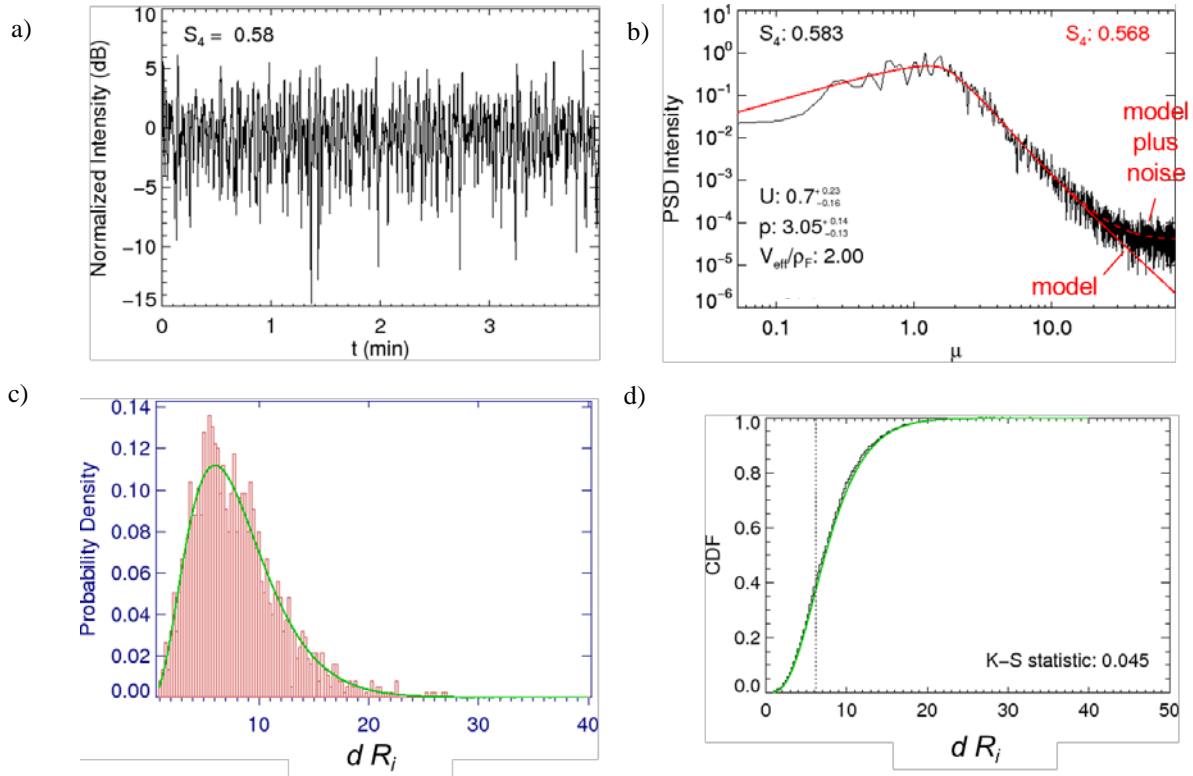


Figure 1. Fitting a simulated intensity spectrum using a one-component irregularity model. The phase screen parameters were $U = 0.6$, $p = 3.0$, $f_F = 2$ Hz.

Next we repeated this experiment multiple times to generate statistics of the fitting errors via Monte Carlo simulation. Figures 2a and 2b show the resulting MLE estimates for scattering strength U and spectral index p , respectively. The blue horizontal bars delineate the 90th percentile confidence intervals while the blue dots indicate the best estimates. The truth values for the parameters are indicated with vertical dashed lines. For all of the 100 Monte Carlo realizations, the confidence intervals contained the truth values for each parameter. This indicates that our confidence intervals are over-conservative, at least for this particular screen configuration. We find that this is not the case for steeply sloped spectra when strong focusing effects are present.

Figure 2c shows the parameter estimates for each realization, while the dashed lines indicate the truth values for each parameter. Note that the parameter estimates are roughly centered on the “cross-hair” indicating that the MLE estimates are relatively unbiased. The solid ellipse encloses the 90th percentile region computed from the covariance of the parameter estimates, while the dashed ellipse is computed from the formal covariances. The shape of the confidence region is a distorted, rather than true, ellipse because we estimate $\text{Log}(U)$ and p and map the confidence interval (which is a true ellipse in in terms of these coordinates) back to linear coordinates. Note that the estimates of scattering strength and spectral index are positively correlated; there is a tendency to under- or over-estimate the values of both of these parameters simultaneously.

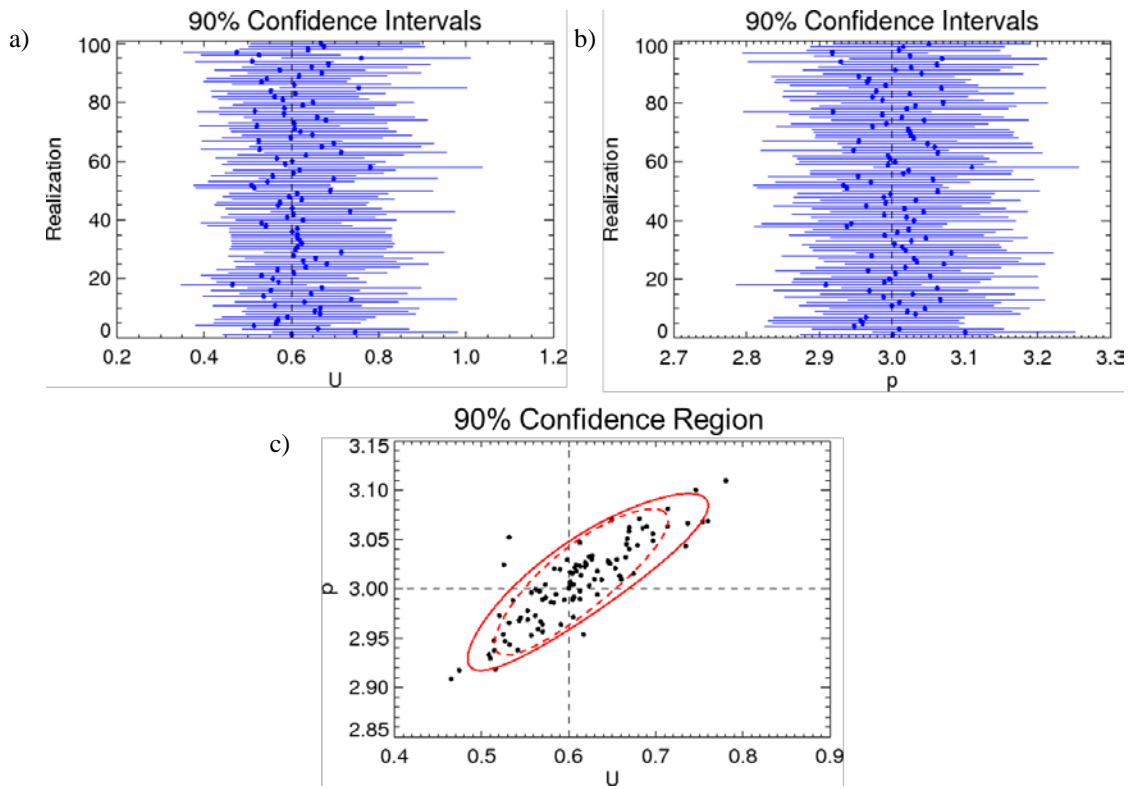


Figure 2. Parameter estimates and confidence intervals (top) and confidence regions (bottom) for the simulation shown in Figure 1. In the bottom plot, the solid curve is a covariance ellipse computed from the parameter estimates while the dashed curve is computed from the formal covariances provided by the ML fit. Dashed black lines indicate truth values for the parameters.

Example 2 - Fitting a Simulated Intensity Spectrum using a Two-Component Irregularity Model

Next we apply the technique to a simulated scintillation time series generated using a two-component power law irregularity model with $U = 0.6$, $p_1 = 3.0$, $p_2 = 3.5$, $\mu_b = 5$, and $f_F = 2$ Hz. Figure 2a shows the intensity fluctuations from the phase screen simulation. The scintillation index is $S_4 = 0.59$. As before, we add complex white noise to the complex amplitude of the propagating wave to simulate the effects of receiver (thermal) noise. We computed the spectrum of the simulated scintillation by averaging four periodograms taken over 60-second data segments (so that $M=4$, and $d=8$). This spectrum is shown as the black curve in Figure 2b.

We chose to estimate scattering strength, spectral indices, and break wavenumber with the Fresnel frequency held fixed, so that $\theta = (U, p_1, p_2, \mu_b)$. Once again we fit the spectral data with the sum of the theoretical SDF plus a constant representing the noise spectral density. The dashed red curve in Figure 2a shows the theoretical SDF plus noise, while the solid red curve shows the theoretical SDF only. The S_4 index corresponding to the best fit result was 0.55. The best fit parameters $U = 0.7$, $p_1 = 2.38$, $p_2 = 3.53$, $\mu_b = 3.8$ and their associated 90% confidence intervals obtained by the $\Delta\chi^2$ are labeled on the plots. The confidence intervals for each estimated parameter contain their truth (specified) values. Figures 2c and 2d respectively show the probability density function and cumulative density function (CDF) of the ratio R_i computed using the best fit solution. These compare very well with those of the chi-squared distribution (shown in green for comparison). The Kolmogorov-Smirnov (K-S) test was 0.087, which indicates a good fit.

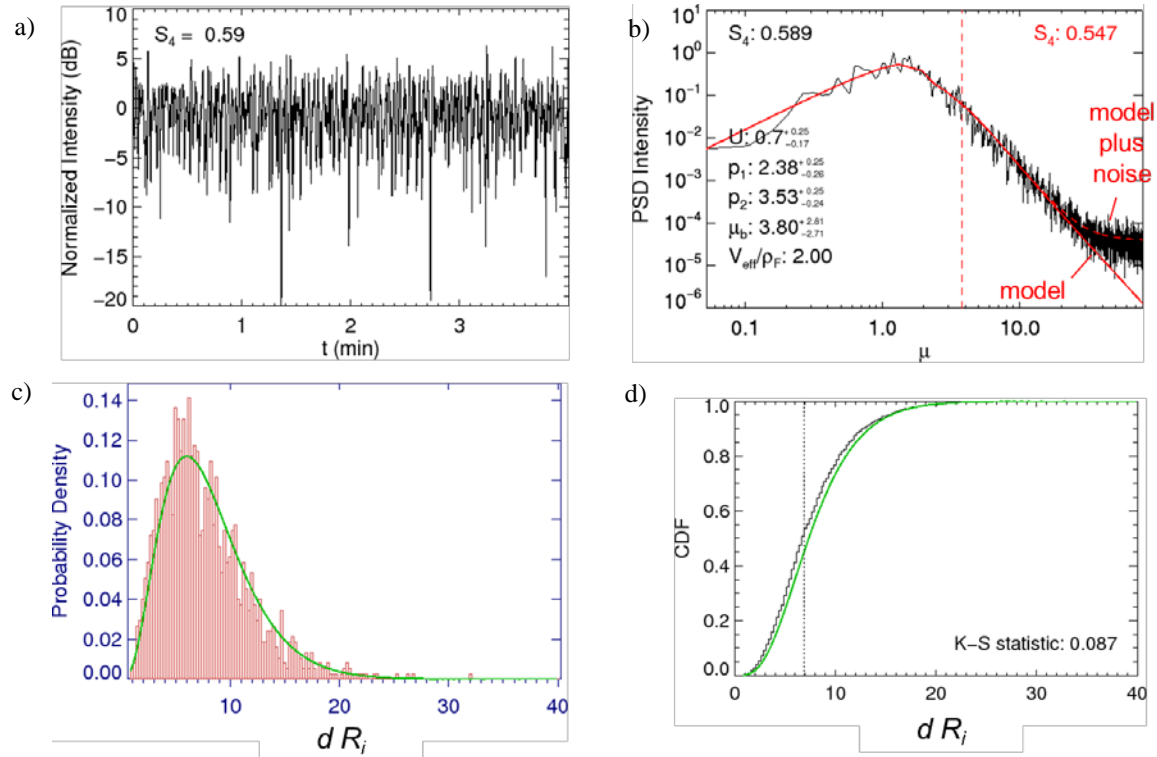


Figure 3. Fitting a simulated intensity spectrum with a two-component irregularity model. The phase screen parameters were $U = 0.6$, $p_1 = 3.0$, $p_2 = 3.5$, $\mu_b = 5$, and $f_F = 2$ Hz.

Next we repeated this experiment multiple times to generate statistics of the fitting errors via Monte Carlo simulation. Figures 4a-4d show the resulting MLE estimates for U , p_1 , p_2 , and μ_b . The blue horizontal bars delineate the 90th percentile confidence interval while the blue dots indicate the best estimates. The truth values for the parameters are indicated with vertical dashed lines. For most of the 100 Monte Carlo realizations conducted, the confidence intervals contained the truth values for each parameter. The parameter whose confidence interval most frequently did not contain the true value was the break wavenumber μ_b . In this sense, the spectral break is the most difficult parameter to estimate using MLE.

Figure 4e shows the parameter estimates for each realization, while the dashed lines indicate the truth values for each parameter. In general, the parameter estimates are more or less centered on the “cross-hairs” indicating that the MLE estimates are relatively unbiased. The solid ellipses enclose the 90th percentile region computed from the covariance of the parameter estimates, while the dashed ellipse is computed from the average covariance obtained from the curvature of the log likelihood function. Note that the estimated scattering strength is negatively correlated with break wavenumber, whereas the spectral indices are positively correlated with break wavenumber. Not surprisingly, we found it difficult to estimate the break wavenumber accurately when $p_1 \approx p_2$ since this parameter becomes degenerate when the spectrum has a single slope. While the specific value of the break wavenumber itself becomes immaterial in this case, an accurate spectral fit may still be found.

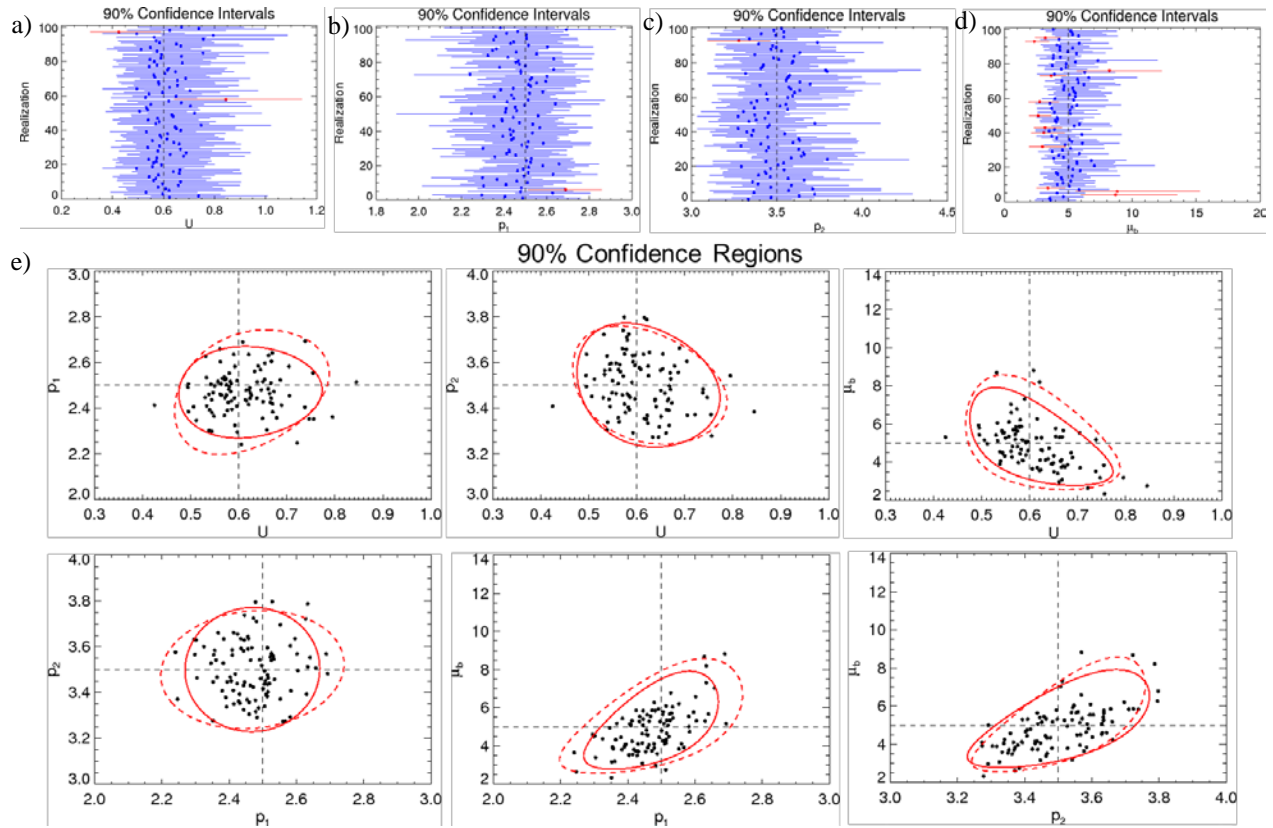


Figure 4. Parameter estimates and confidence intervals (top) and confidence regions (bottom) for the simulation shown in Figure 3. In the bottom plots, the solid red curves are covariance ellipses computed from the parameter estimates while the dashed curves are computed from the formal covariances provided by the ML fit. Dashed black lines indicate truth values for the parameters.

Example 3 – Fitting a Measured Intensity Spectrum when the Irregularities Likely Have a One-Component Spectrum

In this example, we consider scintillation intensity data collected by Air Force Research Laboratory (AFRL) at Ascension Island on 22 March 2000 at 21:47 UT. The L-band signal (1535 MHz) was transmitted by a geostationary communication satellite. The time series of the intensity fluctuations is shown in Figure 5a. Using the spaced-receiver technique to measure the zonal irregularity drift, along with the propagation and magnetic field geometry, we estimated the Fresnel frequency to be $f_F=1.42$ Hz. We applied the ML fitting using both one- and two-component irregularity models. The results using a one-component model are shown in Figure 5b, whereas the results using a two-component model are shown in Figure 5c. The gray shaded area indicates the wavenumber range over which the fits were performed (this range was selected to avoid harmonics distorted by receiver noise). We compare these models using the Akaike Information Criterion (AIC; Akaike, 1973):

$$AIC_i = -2\text{Log}(L) + 2k_i \quad (10)$$

where k_i is the number of free parameters for model i . The model which minimizes the AIC is considered to be the best. The idea is to penalize a model that requires more free parameters to fit the data. While $\log(L)$ is slightly larger for the two-component model (as it must be since the models are nested), the AIC is slightly smaller for the one-component model. The message is that, given the data, there is little statistical justification for using the more complex model in this case. We conclude that the scintillation was probably caused by irregularities with a one-component power law spectrum.

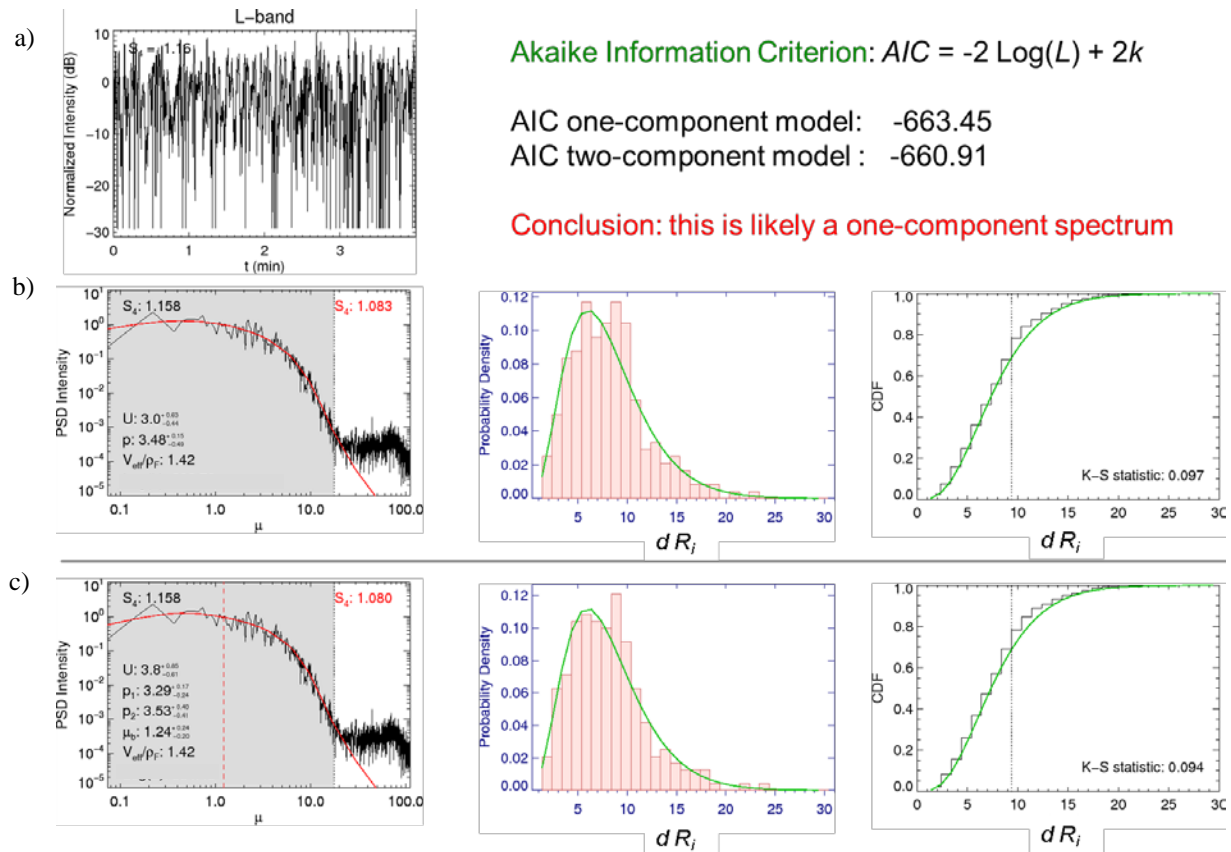
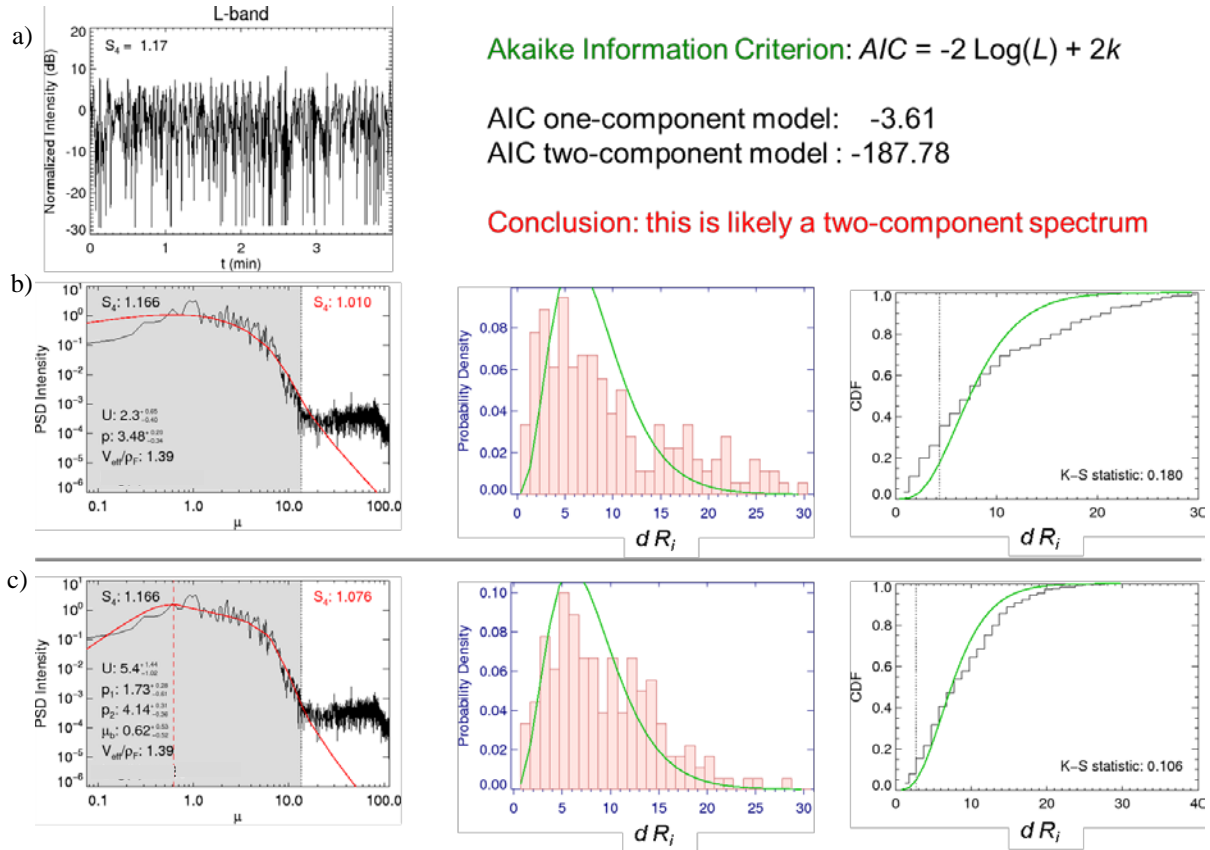


Figure 5. Fitting L-band scintillation observations at Ascension Island at 21:47 UT: a) intensity fluctuations, b) fitting results for a one-component irregularity model, c) fitting results for a two-component irregularity model.

Example 4 – Fitting a Measured Intensity Spectrum when the Irregularities Likely Have a Two-Component Spectrum

Finally, we consider a second sample of scintillation intensity data collected at Ascension Island later this same evening, at 21:55 UT. The time series of the intensity fluctuations is shown in Figure 6a. Using the spaced receiver technique to measure the zonal irregularity drift, along with the propagation and magnetic field geometry, we estimated the Fresnel frequency to be $f_F=1.39$ Hz. Once again we applied the ML fitting technique using both one- and two-component irregularity models. The results using a one-component model are shown in Figure 6b, whereas the results using a two-component model are shown in Figure 6c. For this case, $\log(L)$ is significantly larger for the two-component model and the AIC is significantly smaller. Our conclusion is therefore that the observations justify the use of the more general two-component irregularity model.



Akaike Information Criterion: $AIC = -2 \log(L) + 2k$

AIC one-component model: -3.61

AIC two-component model: -187.78

Conclusion: this is likely a two-component spectrum

Figure 6. Fitting L-band scintillation observations at Ascension Island at 21:55 UT: a) intensity fluctuations, b) fitting results for a one-component irregularity model, c) fitting results for a two-component irregularity model.

4. CONCLUSIONS

We have introduced a technique called Irregularity Parameter Estimation (IPE) for inferring phase screen parameters from scintillation observations. We began by developing a theoretical model suitable for fitting intensity scintillation spectra. Next, we demonstrated how this function may be fit to simulated and actual ionospheric scintillation measurements using the maximum likelihood (ML) method to extract the phase screen parameters. We used Monte-Carlo simulation to demonstrate that the technique yields approximately unbiased results with useful confidence intervals for the estimated parameters. We employed the Akaike Information Criterion (AIC) to establish whether the use of the more general two-component irregularity model is statistically justified given the scintillation data.

We caution that maximum likelihood parameter estimates and their confidences are not meaningful if the theoretical model is invalid. We believe our proposed model is valid under weak, moderate, and strong scatter conditions to the extent that a 1D phase screen model can adequately describe the propagation physics. This condition limits the applicability of the model to propagation geometries such that the line of sight scans through the field-aligned irregularities in an approximately transverse fashion. Furthermore, for the parameter estimates to be unique, the theoretical model must depend uniquely on the parameters over the entire measurement range (i.e. wavenumber range over which the spectrum may be fit). As the scattering strength increases, the intensity spectrum tends to flatten and develop a Gaussian-like high frequency roll off that becomes increasingly insensitive to the screen parameters over the practical measurement range (which is limited by system noise) typical of a scintillation monitoring system. Thus, it may not always be possible to reliably extract phase screen parameters from a very strongly scintillating time series using this technique. A possible solution to this problem may be to incorporate scintillation observations from multiple frequency carriers, if available, and fit their spectra simultaneously.

REFERENCES

- Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle, Proc. 2nd Inter. Symposium on Information Theory, 267-281, Budapest.
- Appourchaux, T. (2003), On maximum likelihood estimation of averaged power spectra, A&A 412, 903-904, DOI: 10.1051/0004-6361:20034401.
- Barret, D., and S. Vaughan (2012), Maximum likelihood fitting of X-ray power density spectra: application to high-frequency quasi-periodic oscillations from the neutron star X-binary 4U1608-522, ApJ , vol. 746 pg. 131.
- Carrano, C. S., C. E. Valladares, and K. M. Groves (2012), Latitudinal and local time variation of ionospheric turbulence parameters during the Conjugate Point Equatorial Experiment in Brazil, International Journal of Geophysics, vol. 2012, Article ID 103963, 16 pages, doi:10.1155/2012/103963.
- Carrano, C., and C. Rino (2016), A theory of scintillation for two-component power law irregularity spectra: Overview and numerical results, Radio Sci., 51, 789–813, doi:10.1002/2015RS005903.
- Rino, C. (1979a), A power law phase screen model for ionospheric scintillation, 1 Weak scatter, Radio Sci., 14, 1135-1145.
- Ruddick, B., A. Anis, and K. Thompson (2000): Maximum likelihood spectral fitting: The Batchelor spectrum, J. Atmos. Oceanic Technol., 17, 1541–1555, doi:10.1175/1520-0426(2000)017<1541:MLSFTB.2.0.CO;2.
- Vaughn, S. (2005), A simple test for periodic signals in red noise, A&A 431, 391-403, DOI 10.1051/0004-6361:20041453.