**Atmospheric
Measurement
Techniques**

Progress in turbulence detection via GNSS occultation data

L. B. Cornman[1], R. K. Goodrich[1,3], P. Axelrad[2], and E. Barlow[2]

[1]Research Applications Laboratory, National Center for Atmospheric Research, Boulder, Colorado, USA
[2]Department of Aerospace Engineering Sciences, University of Colorado, Boulder, Colorado, USA
[3]Department of Mathematics, University of Colorado, Boulder, Colorado, USA

Correspondence to: L. B. Cornman (cornman@ucar.edu)

Abstract. The increased availability of radio occultation (RO) data offers the ability to detect and study turbulence in the Earth's atmosphere. An analysis of how RO data can be used to determine the strength and location of turbulent regions is presented. This includes the derivation of a model for the power spectrum of the log-amplitude and phase fluctuations of the permittivity (or index of refraction) field. The bulk of the paper is then concerned with the estimation of the model parameters. Parameter estimators are introduced and some of their statistical properties are studied. These estimators are then applied to simulated log-amplitude RO signals. This includes the analysis of global statistics derived from a large number of realizations, as well as case studies that illustrate various specific aspects of the problem. Improvements to the basic estimation methods are discussed, and their beneficial properties are illustrated. The estimation techniques are then applied to real occultation data. Only two cases are presented, but they illustrate some of the salient features inherent in real data.

1 Introduction

There is a long and distinguished history in the study of electromagnetic (EM) wave propagation through random media (cf. Tatarskii, 1971; Yeh and Liu, 1982; Ishimaru, 1997). These works have provided a firm theoretical foundation for estimating statistical properties of the neutral atmosphere and ionosphere via the statistical properties of the received EM signals. That is, the characteristics of the turbulent atmosphere can be deduced from, for example, correlation and/or spectral analysis of the phase and/or amplitude of the received signal. In the past, the bulk of the experimental analysis in this area has been performed with ground-based transmitters and receivers (e.g., radars and lidars), as well as with ground-based receivers and space-based transmitters. With the advent of Global Navigation Satellite System (GNSS) constellations, a new avenue has become available to investigate the turbulence properties of the earth's atmosphere. The deployment of an ever-increasing number of Low Earth Orbiting (LEO) satellites – with high quality, high-sampling rate receivers – provides a very valuable new source of turbulence measurements: GNSS-LEO occultations.

Previous efforts to study turbulence with RO signals have been focused on the study of the middle to upper stratosphere and the ionosphere (e.g., Gurvich and Brekhovskikh, 2001; Gurvich and Kan, 2003; Gurvich and Chunchuzov, 2003). On the other hand, very few works have addressed the characterization of turbulence in the upper troposphere and lower stratosphere. Numerous other practical applications can benefit by having turbulence measurements (and resultant climatologies) from GNSS occultations. For example: space weather (Hajj et al., 2000), transionospheric communication links (Secan et al., 1997), aviation safety and navigation systems (Cornman et al., 2004; Conker et al., 2003), the accuracy of ground-based GNSS receivers (Ganguly et al., 2004), as well as the effect of turbulence on measuring atmospheric state variables from GNSS-LEO occultations (Gorbunov and Kirchengast, 2005).

In this paper, a detailed analysis of the problem is presented, including the derivation of a model for the power spectrum of the log-amplitude and phase fluctuations of the

permittivity (or index of refraction) field. The bulk of the paper is then concerned with the estimation of model parameters, such as the intensity and location of the turbulence along the line of sight. A thorough study of the properties of the parameter estimators is presented using simulated data. This includes statistics over a large number of realizations, as well as case studies that illustrate various specific aspects of the problem. The application of the methodologies is then made with real occultation data. As the focus in this paper was the development of methods, the purpose of the real data analysis is to show how the theoretical and practical concepts carry over into the real data. Hence, only two cases are presented, but they illustrate some of the salient features inherent in real data. A more thorough analysis of real data is left to follow-on efforts.

2 Wave propagation theory

2.1 Assumptions

In order to make the problem tractable, a number of simplifying assumptions will be made. Most of these are routinely used in the literature for optical wavelengths (cf. Tatarskii, 1971; Ishimaru, 1997). Very little work has been done to justify many of these assumptions for the decimeter wavelengths in the GNSS problem; however, a few excellent papers exist, including Strohbehn and Clifford (1967), Clifford and Strohbehn (1970), Clifford and Lataitis (1985), and Lee and Harp (1969). We will list all of the assumptions here, but their use will be introduced below in the context of the development.

1. Polarization effects are ignored. This allows for the wave equation to be written in terms of a single scalar component of the electric field vector. This is a reasonable assumption for the decimeter wavelengths considered for the GNSS problem (cf. Clifford and Strohbehn, 1970).

2. Weak and single scattering is assumed. This is reasonable for the GNSS transmitter wavelengths and typical path lengths through turbulence (Clifford and Strohbehn, 1970).

3. The assumption that the transmitter wavelength, λ, is much smaller than the smallest fluctuations in the permittivity field, l_0, $(\lambda \ll l_0)$, allows for a number of simplifications in the development (including the first two items above). While this is not necessarily true for decimeter wavelengths, Clifford and Strohbehn (1970) show that the practical effects of the assumption $\lambda \ll l_0$, can be extended into the regime, $\lambda \geq l_0$, but not $\lambda \gg l_0$.

4. Taylor's Hypothesis is valid. That is, for short time periods advection dominates over internal fluid stresses (Monin and Yaglom, 1975). This is important for a medium that is moving, which in general is the case for the atmosphere. The practical effect of this assumption is that we can assume that the permittivity for a small scattering region is constant as it moves during a short time interval.

5. The Markov Approximation is valid (Tatarskii, 1971, Chapter 5, Section B). This means that there are minimal effects on the received signal due to correlation in the permittivity field along the line-of-sight. Therefore, one can assume that this correlation is represented by a delta-function.

6. The transmitter is a point source (i.e., we do not consider the radiation pattern), and the receiver acts as point aperture along the line-of-sight (i.e., we do not consider the gain pattern). Due to the typical large distances between the transmitter and the scattering medium, the first assumption is quite reasonable. If the distance between the scattering volume and the receiver is also large, and a high-gain antenna is used, the second assumption is also a reasonable one. For an omnidirectional receiver antenna and/or the receiver very close to the scattering volume (e.g., for an airborne receiver in the atmosphere), then these effects do need to be accommodated (cf. Lee and Harp, 1969; Tatarskii, 1971, Sect. 53).

7. Atmospheric attenuation due to absorption is assumed to be smooth and varying slowly during an occultation, and so that its effects can be accommodated by trend removal of the amplitude or phase signal in the time domain. This also means that we may consider the permittivity field to be real-valued.

8. The turbulence patch has compact support. This is required to ensure the existence of various integrals that arise in the calculations.

9. The effects due to an inhomogeneous background (e.g., refraction) are minimal. This means that we may assume straight-line propagation.

10. The medium is purely random. There are no deterministic structures, such as layered phenomena.

Most of the assumptions presented above are reasonable for a first-order analysis of the turbulence problem. Nevertheless, there are two aspects of the problem that are not addressed in this initial study: the effects due to a background or turbulent permittivity field characterized by inhomogeneity and/or anisotropy. These are obviously real-world situations, but as we are trying to keep to a simple analysis at this point, they are not included in what follows.

2.2 The first Rytov approximation

If polarization effects can be ignored, the vector wave equation for monochromatic signals derived from Maxwell's equations can be written in the scalar form,

$$(\nabla^2 + k^2 \varepsilon(\boldsymbol{r})) U(\boldsymbol{r}) = 0 \qquad (1)$$

where U is a component of the electric field (a transverse component for our case), k is the (constant) wavenumber for the transmitted frequency, $k = 2\pi f/c = 2\pi/\lambda$, and ε is the permittivity of the medium. The index of refraction, n, is given by $\varepsilon = n^2$, and free space is characterized by $\varepsilon = 1$. Since we have assumed weak scattering, a perturbation approached is used. Re-write Eq. (1) as

$$(\nabla^2 + k^2) U(\boldsymbol{r}) = -k^2 \varepsilon'(\boldsymbol{r}) U(\boldsymbol{r}) \qquad (2)$$

where $\varepsilon = \langle \varepsilon \rangle + \varepsilon'$, and $\varepsilon' \ll \langle \varepsilon \rangle$. It is assumed that $\langle \varepsilon \rangle = 1$, its free-space value. The angled brackets refer to an ensemble average. Standard approaches to solving this equation are provided in the literature (e.g., Tatarskii, 1971; Ishimaru, 1997) – one of which is the first-order Rytov approach. We return to the Helmholtz wave equation and let

$$U(\boldsymbol{r}) = e^{\psi(\boldsymbol{r})} \qquad (3)$$

where the complex phase, ψ, is given by

$$\psi(\boldsymbol{r}) = \chi(\boldsymbol{r}) + i S(\boldsymbol{r})$$

and χ is the log-amplitude function, and S is the phase function. Next, expand ψ in a power series in terms of the small parameter $\varepsilon'/\langle \varepsilon \rangle$. Plugging Eq. (3) into Eq. (2) and solving for $U_0 \psi_1$ leads to the first-order term (cf. Tatarskii, 1971; Ishimaru, 1997),

$$\psi_1(\boldsymbol{r}) = \frac{k^2}{U_0(\boldsymbol{r})} \int G_0(\boldsymbol{r} - \boldsymbol{r}') \varepsilon'(\boldsymbol{r}') U_0(\boldsymbol{r}') d\boldsymbol{r}' \qquad (4)$$

Where U_0 is the solution to Eq. (2) when ϵ' is zero, and

$$G_0(\boldsymbol{r} - \boldsymbol{r}') = \frac{e^{ik\|\boldsymbol{r} - \boldsymbol{r}'\|}}{4\pi \|\boldsymbol{r} - \boldsymbol{r}'\|} \qquad (5)$$

is the free-space Green's function for the problem (the double vertical lines indicate the vector magnitude). Equation (4) is known as the first Rytov approximation.

2.3 Correlation function and frequency spectrum

We depart from the standard development as presented by, for example, Tatarskii (1971) or Ishimaru (1997), going forward. We consider the situation where the transmitter, receiver, and media are all moving relative to each other, as well as with respect to a fixed coordinate system. Typically, the origin of the coordinate system is set at the transmitter or at a point within the scattering volume. However, in what follows it is assumed that the origin of the coordinate system is at an arbitrary point. For the GNSS-LEO geometry,
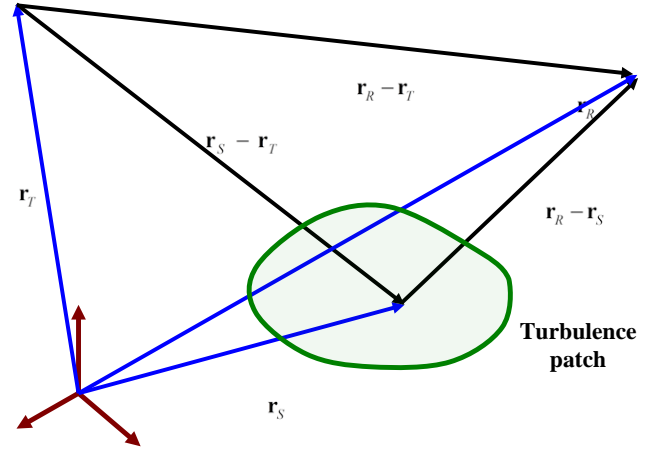

Fig. 1. Scattering geometry at time t_1.

we will want the origin of the coordinate system to be at the (local) center of the earth. In the arbitrary coordinate system, the receiver is at the point $\boldsymbol{r}_{\mathrm{R}}$, the transmitter is at $\boldsymbol{r}_{\mathrm{T}}$, and a given scatterer is at $\boldsymbol{r}_{\mathrm{S}}$ (see Fig. 1). (We are considering a continuum field of permittivity, so the term scatterer is used loosely.) We assume that multiple scattering does not occur, and that the transmitter acts as a point source, (i.e., an outgoing spherical wave), which leads to

$$\psi_1(\boldsymbol{r}_{\mathrm{R}}) = \frac{k^2}{4\pi} \|\boldsymbol{r}_{\mathrm{R}} - \boldsymbol{r}_{\mathrm{T}}\| e^{-ik\|\boldsymbol{r}_{\mathrm{R}} - \boldsymbol{r}_{\mathrm{T}}\|}$$
$$\int \frac{e^{ik\|\boldsymbol{r}_{\mathrm{R}} - \boldsymbol{r}_{\mathrm{S}}\|}}{\|\boldsymbol{r}_{\mathrm{R}} - \boldsymbol{r}_{\mathrm{S}}\|} \frac{e^{ik\|\boldsymbol{r}_{\mathrm{S}} - \boldsymbol{r}_{\mathrm{T}}\|}}{\|\boldsymbol{r}_{\mathrm{S}} - \boldsymbol{r}_{\mathrm{T}}\|} \varepsilon'(\boldsymbol{r}_{\mathrm{S}}) d\boldsymbol{r}_{\mathrm{S}} \qquad (6)$$

In order to ensure the existence of the integral, we assume that the permittivity fluctuations have compact support, and $\varepsilon'(\boldsymbol{r}_{\mathrm{S}})$ is vanishingly small (or has a minimal effect on $\psi_1(\boldsymbol{r}_{\mathrm{R}})$) well away from the line of sight, $\boldsymbol{r}_{\mathrm{R}} - \boldsymbol{r}_{\mathrm{T}}$. If $\varepsilon'(\boldsymbol{r}_{\mathrm{S}})$ is a random field, we assume that Eq. (6) exists in the sense that it is a member of an ensemble of random permittivity fluctuations. We will assume straight-line propagation. Return to Eq. (6). If the scattering angles are small – a reasonable assumption if the ratio λ/l_0 is small – we can make the so-called parabolic approximation in the integral equation (Tatarskii, 1971, Section 45). (l_0 is a length scale on the order of the smallest fluctuations in the medium.) In order to use this approximation, we will consider basis vectors that are parallel or perpendicular to the vector $\boldsymbol{r}_{\mathrm{R}} - \boldsymbol{r}_{\mathrm{T}}$. We will denote the component of vectors parallel to $\boldsymbol{r}_{\mathrm{R}} - \boldsymbol{r}_{\mathrm{T}}$, "$x$" and the perpendicular parts, $\boldsymbol{\rho}$ (i.e., the projection of a given vector onto the plane perpendicular to $\boldsymbol{r}_{\mathrm{R}} - \boldsymbol{r}_{\mathrm{T}}$). Note that $\boldsymbol{\rho}_{\mathrm{R}} - \boldsymbol{\rho}_{\mathrm{T}} = 0$. Consider the quantity,

$$\|\boldsymbol{r}_{\mathrm{R}} - \boldsymbol{r}_{\mathrm{S}}\| = \left[(x_{\mathrm{R}} - x_{\mathrm{S}})^2 + \|\boldsymbol{\rho}_{\mathrm{R}} - \boldsymbol{\rho}_{\mathrm{S}}\|^2 \right]^{1/2}.$$ We will assume that $x_{\mathrm{T}} < x_{\mathrm{S}} < x_{\mathrm{R}}$. If $x_{\mathrm{R}} - x_{\mathrm{S}}$, the distance along the x axis from the scatterer to the receiver, is much greater than the transverse distances (valid for small-angle scattering), then

$\|\boldsymbol{r}_\mathrm{R}-\boldsymbol{r}_\mathrm{S}\|$, can be approximated by

$$\|\boldsymbol{r}_\mathrm{R}-\boldsymbol{r}_\mathrm{S}\| = (x_\mathrm{R}-x_\mathrm{S})\left[1+\frac{\|\boldsymbol{\rho}_\mathrm{R}-\boldsymbol{\rho}_\mathrm{S}\|^2}{(x_\mathrm{R}-x_\mathrm{S})^2}\right]^{1/2}$$

$$\approx x_\mathrm{R}-x_\mathrm{S}+\frac{\|\boldsymbol{\rho}_\mathrm{R}-\boldsymbol{\rho}_\mathrm{S}\|^2}{2(x_\mathrm{R}-x_\mathrm{S})} \tag{7}$$

This is the parabolic approximation. A similar expansion can be made for $\|\boldsymbol{r}_\mathrm{R}-\boldsymbol{r}_\mathrm{T}\|$, which is simply $x_\mathrm{R}-x_\mathrm{T}$. Since $k\|\boldsymbol{r}_\mathrm{R}-\boldsymbol{r}_\mathrm{S}\| = 2\pi\|\boldsymbol{r}_\mathrm{R}-\boldsymbol{r}_\mathrm{S}\|/\lambda \gg 1$, we will use the full expression in Eq. (7) for the exponential terms in Eq. (6); whereas, we only need to keep the leading term for the non-exponential terms. Plugging all these factors back into Eq. (6), we have for the non-exponential terms

$$\frac{\|\boldsymbol{r}_\mathrm{R}-\boldsymbol{r}_\mathrm{T}\|}{\|\boldsymbol{r}_\mathrm{R}-\boldsymbol{r}_\mathrm{S}\|\|\boldsymbol{r}_\mathrm{S}-\boldsymbol{r}_\mathrm{T}\|} \approx \frac{x_\mathrm{R}-x_\mathrm{T}}{(x_\mathrm{R}-x_\mathrm{S})(x_\mathrm{S}-x_\mathrm{T})} \tag{8}$$

which we will denote as $b(x_\mathrm{R},x_\mathrm{S},x_\mathrm{T})$. It can be shown that the exponential term from Eq. (6) is given by

$$\|\boldsymbol{r}_\mathrm{R}-\boldsymbol{r}_\mathrm{S}\|+\|\boldsymbol{r}_\mathrm{S}-\boldsymbol{r}_\mathrm{T}\|-\|\boldsymbol{r}_\mathrm{R}-\boldsymbol{r}_\mathrm{T}\|$$

$$\approx \frac{b(x_\mathrm{R},x_\mathrm{S},x_\mathrm{T})}{2}\|\boldsymbol{\rho}_\mathrm{S}-\boldsymbol{\rho}_\mathrm{T}\|^2 \tag{9}$$

Plugging Eqs. (8) and (9) into Eq. (6) gives the desired result,

$$\psi_1(\boldsymbol{r}_\mathrm{R}) \approx \frac{k^2}{4\pi}\int b(x_\mathrm{R},x_\mathrm{S},x_\mathrm{T})$$

$$\exp\left[ik\frac{b(x_\mathrm{R},x_\mathrm{S},x_\mathrm{T})}{2}\|\boldsymbol{\rho}_\mathrm{S}-\boldsymbol{\rho}_\mathrm{T}\|^2\right]\varepsilon'(\boldsymbol{r}_\mathrm{S})d\boldsymbol{r}_\mathrm{S} \tag{10}$$

If we move the origin to $\boldsymbol{r}_\mathrm{T}$, this is equivalent to setting $x_\mathrm{T}=0$ and $\boldsymbol{\rho}_\mathrm{T}=0$, and since we have assumed that $\boldsymbol{\rho}_\mathrm{R}-\boldsymbol{\rho}_\mathrm{T}=0$, this means that $\boldsymbol{\rho}_\mathrm{R}=0$. Equation. (10) then describes the field,

$$\psi_1(x_R,\boldsymbol{\rho}_\mathrm{R}=0) \approx \frac{k^2}{4\pi}\int\frac{x_\mathrm{R}}{x_\mathrm{S}(x_\mathrm{R}-x_\mathrm{S})}$$

$$\exp\left[ik\frac{x_\mathrm{R}}{2x_\mathrm{S}(x_\mathrm{R}-x_\mathrm{S})}\|\boldsymbol{\rho}_\mathrm{S}\|^2\right]\varepsilon'(\boldsymbol{r}_\mathrm{S})d\boldsymbol{r}_\mathrm{S} \tag{11}$$

which is equivalent to Sect. (49), Eq. (5) in Tatarskii (1971). Therefore, we see that Eq. (10) has the correct limiting form when the origin is placed at the transmitter.

Next, consider the transmitter, receiver and atmosphere moving relative to each other – and all with respect to a fixed origin, e.g., the center of the Earth. Since the longitudinal motions do not change the value of the correlation of the field in an appreciable manner (Klyatskin, 1970), we assume that the "x"-positions are not changing, and so the velocity vectors described below are the projection of the three-dimensional velocity vectors onto the plane perpendicular to $\boldsymbol{r}_\mathrm{R}-\boldsymbol{r}_\mathrm{T}$. We consider a short enough time interval, $\tau=t_2-t_1$, such that $\boldsymbol{\rho}_i(t_2) \approx \boldsymbol{\rho}_i(t_1)+\boldsymbol{V}_i\tau$, (where
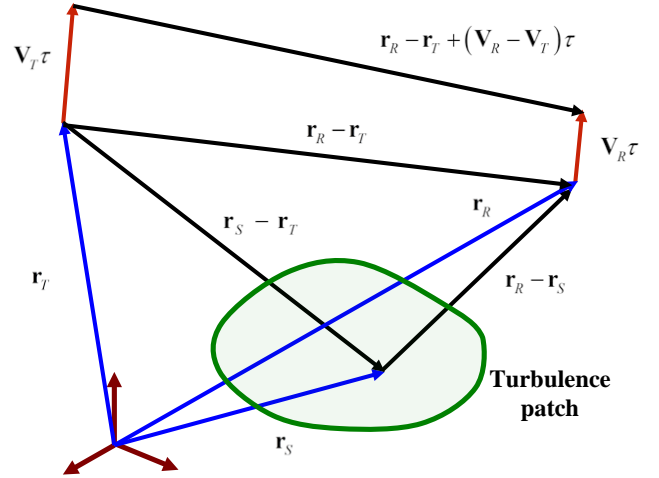

Fig. 2. Scattering geometry at times t_1 and t_2.

"i" stands for R, T, or S). For this work we will assume that the velocities are constant over all (short) time intervals (i.e., $V_i(\tau)=V_i$); that they are the same for all the scatterers (i.e., $V_{\mathrm{S}_i}=V_{\mathrm{S}_j}$); and that they are deterministic quantities (i.e., $\langle V_i\rangle=V_i$). The same procedure used to derive Eq. (10) can be applied to calculating $\psi_1(\boldsymbol{r}_\mathrm{R}(t_2))$, except that we now use basis vectors parallel and perpendicular to $\boldsymbol{r}_\mathrm{R}-\boldsymbol{r}_\mathrm{T}+(\boldsymbol{V}_\mathrm{R}-\boldsymbol{V}_\mathrm{T})\tau$ (see Fig. 2). Unfortunately, this will complicate matters later on when we calculate the correlation functions; which in turn contain products of $\psi_1(\boldsymbol{r}_\mathrm{R}(t_1))$ and $\psi_1(\boldsymbol{r}_\mathrm{R}(t_2))$. That is, we will have to transform the components of vectors in one basis into the other. However, if the two vectors $\boldsymbol{r}_\mathrm{R}-\boldsymbol{r}_\mathrm{T}$ and $\boldsymbol{r}_\mathrm{R}-\boldsymbol{r}_\mathrm{T}+(\boldsymbol{V}_\mathrm{R}-\boldsymbol{V}_\mathrm{T})\tau$ are close to parallel we can then choose the basis vectors perpendicular to them to be identical. This assumption is valid to first order if $\|\boldsymbol{V}_\mathrm{T}\tau\|\ll\|\boldsymbol{r}_\mathrm{R}-\boldsymbol{r}_\mathrm{T}\|$ and $\|\boldsymbol{V}_\mathrm{T}\tau\|\ll\|\boldsymbol{r}_\mathrm{R}-\boldsymbol{r}_\mathrm{T}\|$. In this case, the components of a vector in each basis system will be the same – up to the level of approximation used. Given the distances involved in typical GNSS-LEO occultations, these conditions will be met; hence, we can safely assume that the same basis can be used for calculations involving the components of both $\boldsymbol{r}_\mathrm{R}-\boldsymbol{r}_\mathrm{T}+(\boldsymbol{V}_\mathrm{R}-\boldsymbol{V}_\mathrm{T})\tau$ and $\boldsymbol{r}_\mathrm{R}-\boldsymbol{r}_\mathrm{T}$.

Next, we calculate $\psi_1(\boldsymbol{r}_\mathrm{R}(t_2))$. Consistent with the approximations just discussed, we assume that

$$x_\mathrm{R}-x_\mathrm{S} \gg \|\boldsymbol{\rho}_\mathrm{R}-\boldsymbol{\rho}_\mathrm{S}+(\boldsymbol{V}_\mathrm{R}-\boldsymbol{V}_\mathrm{S})\tau\| \tag{12}$$

and similarly for $x_\mathrm{S}-x_\mathrm{T}$. It can be shown that the exponential terms are given by

$$\|\boldsymbol{r}_\mathrm{R}-\boldsymbol{r}_\mathrm{S}+(\boldsymbol{V}_\mathrm{R}-\boldsymbol{V}_\mathrm{S})\tau\|+\|\boldsymbol{r}_\mathrm{S}-\boldsymbol{r}_\mathrm{T}+(\boldsymbol{V}_\mathrm{S}-\boldsymbol{V}_\mathrm{T})\tau\|$$
$$-\|\boldsymbol{r}_\mathrm{R}-\boldsymbol{r}_\mathrm{T}+(\boldsymbol{V}_\mathrm{R}-\boldsymbol{V}_\mathrm{T})\tau\|$$
$$= b(x_\mathrm{R},x_\mathrm{S},x_\mathrm{T})\|\boldsymbol{\rho}_\mathrm{S}-\boldsymbol{\rho}_\mathrm{T}-\boldsymbol{V}_\mu\tau\|^2/2$$

where,

$$\boldsymbol{V}_\mu = \left(\frac{x_\mathrm{S}-x_\mathrm{T}}{x_\mathrm{R}-x_\mathrm{T}}\right)(\boldsymbol{V}_\mathrm{R}-\boldsymbol{V}_\mathrm{T})-(\boldsymbol{V}_\mathrm{S}-\boldsymbol{V}_\mathrm{T}) \tag{13}$$

For the non-exponential terms in $\psi_1(r_R(t_2))$, using Eq. (7) gives the same result as for $\psi_1(r_R(t_1))$, i.e., Eq. (8).

Hence, we can write Eq. (10) at two different times as

$$\psi_1(x_R, \rho_R(t_1)) \approx \frac{k^2}{4\pi} \int b(x_R, x_{S_1}, x_T)$$
$$\exp\left[\frac{ikb(x_R, x_{S_1}, x_T)}{2}\|\rho_{S_1} - \rho_T\|^2\right]$$
$$\varepsilon'(r_{S_1}(t_1))dr_{S_1} \qquad (14)$$

and

$$\psi_1(x_R, \rho_R(t_2)) \approx \frac{k^2}{4\pi} \int b(x_R, x_{S_2}, x_T)$$
$$\exp\left[\frac{ikb(x_R, x_{S_2}, x_T)}{2}\|\rho_{S_2} - \rho_T - V_\mu \tau\|^2\right]$$
$$\varepsilon'(r_{S_2}(t_2))dr_{S_2} \qquad (15)$$

where the subscripts "1" and "2" on x_S and ρ_S have been used to indicate that we are considering scatterers at the points r_{S_1} and r_{S_2}, respectively.

The important distinction between deterministic and random quantities is that the latter should be quantified via their statistical properties. For our purposes, we will consider second-order correlation functions, and subsequently the associated Fourier frequency spectrum. Since $\psi_1 = \chi_1 + iS_1$, it can be seen that $\chi_1 = \text{Re}\,\psi_1 = (\psi_1 + \psi_1^*)/2$ and $S_1 = \text{Im}\,\psi_1 = (\psi_1 - \psi_1^*)/2i$. It is straightforward to show that the correlation function of the log-amplitude field is given by

$$B_{\chi_1}(\rho_R(t_1), \rho_R(t_2)) = \frac{1}{2}\big[\text{Re}\langle\psi_1(\rho_R(t_1))\psi_1(\rho_R(t_2))\rangle$$
$$+\text{Re}\langle\psi_1(\rho_R(t_1))\psi_1^*(\rho_R(t_2))\rangle\big] \qquad (16)$$

For the correlation function of the phase, the first term in Eq. (16) changes sign; and for the cross-correlation of the log-amplitude and phase, the real parts are replaced by imaginary parts and the second term in Eq. (16) changes sign.

In order to obtain simplified expressions for the double integrals in Eq. (16), two further approximations are invoked. First, we apply Taylor's Hypothesis (Monin and Yaglom, 1975) to the field $\varepsilon'(r_{S_2}(t_2)) = \varepsilon'(r_{S_2}(t_1) + V_{S_2}\tau)$. This hypothesis assumes that due to pure advection, $\varepsilon'(r_{S_2}(t_2)) = \varepsilon'(r_{S_2}(t_1))$. That is, the permittivity fluctuation due to a scatterer at $r_{S_2}(t_1)$ will have the same numerical value after it moves to $r_{S_2}(t_2)$. This means that the correlation function of the permittivity field can be written as

$$B_{\varepsilon'}(r_{S_1}(t_1), r_{S_2}(t_2)) = \langle\varepsilon'(r_{S_1}(t_1))\varepsilon'(r_{S_2}(t_2))\rangle$$
$$= \langle\varepsilon'(r_{S_1}(t_1))\varepsilon'(r_{S_2}(t_1))\rangle \qquad (17)$$

(it is assumed that the permittivity field is a real-valued one). This allows for Eqs. (14) and (15) to be written solely as a function of t_1, so we can drop that notation except for

where it is needed explicitly. Next, we invoke the Markov approximation (Tatarskii, 1971, Section 64). Consider the correlation function of the permittivity field

$$B_{\varepsilon'}(x_{S_1}, x_{S_2}, \rho_{S_1}, \rho_{S_2}) = \langle\varepsilon'(x_{S_1}, \rho_{S_1})\varepsilon'(x_{S_2}, \rho_{S_2})\rangle \qquad (18)$$

For a homogeneous field, the correlation function can only be a function of $\rho_{S_1} - \rho_{S_2}$ and $x_{S_1} - x_{S_2}$, and furthermore if the field is also isotropic it can only be a function of $\|\rho_{S_1} - \rho_{S_2}\|$ and $|x_{S_1} - x_{S_2}|$. Under the Markov approximation it is assumed that there is no correlation of the permittivity field in the longitudinal direction, and hence

$$B_{\varepsilon'}(x_{S_1}, x_{S_2}, \rho_{S_1}, \rho_{S_2}) = \delta(|x_{S_1} - x_{S_2}|)A_{\varepsilon'}(\|\rho_{S_1} - \rho_{S_2}\|) \qquad (19)$$

where $A_{\varepsilon'}$ is the correlation function in the transverse coordinates. Using Eq. (19), the two quantities, $\langle\psi_1(\rho_R(t_1))\psi_1(\rho_R(t_2))\rangle$ and $\langle\psi_1(\rho_R(t_1))\psi_1^*(\rho_R(t_2))\rangle$ can be evaluated as

$$\langle\psi_1(\rho_R(t_1))\psi_1(\rho_R(t_2))\rangle$$
$$= -\frac{k^2\pi^2}{2}\int \exp\left[-\frac{i\|q\|^2}{kb(\eta)}\right]J_0(\|q\|\|V_\mu(\eta)\|\tau)$$
$$\varphi_{\varepsilon'}(0, \|q\|)\|q\|\,d\|q\|\,d\eta \qquad (20)$$

and,

$$\langle\psi_1(\rho_R(t_1))\psi_1^*(\rho_R(t_2))\rangle$$
$$= \frac{k^2\pi^2}{2}\int J_0(\|q\|\|V_\mu(\eta)\|\tau)\varphi_{\varepsilon'}(0, \|q\|)\|q\|\,d\|q\|\,d\eta \qquad (21)$$

In Eqs. (20) and (21), $\varphi_{\varepsilon'}(0, \|q\|)$ is the three-dimensional spectrum of the (assumed) isotropic permittivity fluctuations with $q = (q_y, q_z)$, $\eta = (x_{S_1} + x_{S_2})/2$, and $b(\eta) = \frac{R}{(\eta - x_T)(x_R - \eta)}$, where $R = x_R - x_T$. We need the real and imaginary parts of Eqs. (20) and (21) to calculate the correlations. Equation (21) is real-valued, so its imaginary part is zero. The real and imaginary parts of Eq. (20) will result in a cosine and a sine term, respectively, in the integral. Except for these trigonometric terms, the integrands for Eqs. (20) and (21) are the same. Therefore, using the trigonometric identity, $2\sin^2\theta = 1 - \cos 2\theta$, the correlation function of the log-amplitude fluctuations is given by

$$B_{\chi_1}(\rho_R(t_1), \rho_R(t_1 + \tau)) = 4\pi^2 k^2 \int\limits_{x_T}^{x_R}\int\limits_0^\infty \varphi_{n'}(0, \|q\|)$$
$$J_0(\|q\|\|V_\mu(\eta)\|\tau)\sin^2\left(\frac{\|q\|^2}{2kb(\eta)}\right)\|q\|\,d\|q\|\,d\eta \qquad (22)$$

where we have used the spectrum of the index of refraction field, $\varphi_{n'}$, instead of that for the permittivity field. (For weak scattering, $\varphi_{\varepsilon'} \approx 4\varphi_{n'}$.) We have also indicated the explicit ranges of the integrals. If we set $V_R = V_T = 0$, this means that $V_\mu \to -V_S$, and as above, if we move the origin to the

transmitter we recover Eqs. (19–11) in Ishimaru (1997). So again, we see that our results have the correct limiting form as presented in the literature. Using $2\cos^2\theta = 1 + \cos 2\theta$, the correlation function of the phase fluctuations is the same as that for the log-amplitude case, excepting that the sine-squared term becomes a cosine-squared term. The cross-correlation function is of the same form excepting that the sine-squared term becomes a sine term – without the factor of two in the denominator of the argument – and the coefficient in front is given by $\pi^2 k^2/2$. (After we performed the derivation described above, a reference was found that provides a very similar result: Strohbehn, 1974)

To compute the frequency spectrum, we take the cosine transform of Eq. (22). We define the cosine transform of a function $g(t)$ as

$$g(f) = 4 \int_0^\infty g(t)\cos(2\pi f t)dt \qquad (23)$$

Hence,

$$\varphi_{\chi_1}(f) = 16\pi^2 k^2 \int \varphi_{n'}(0, \|\boldsymbol{q}\|) J_0\left(\|\boldsymbol{q}\| \|\boldsymbol{V}_\mu(\eta)\| \tau\right)$$
$$\sin^2\left(\frac{\|\boldsymbol{q}\|^2}{2kb(\eta)}\right)\cos(2\pi f \tau)\|\boldsymbol{q}\| \, d\|\boldsymbol{q}\| d\eta d\tau \qquad (24)$$

After performing the integration over τ and $\|\boldsymbol{q}\|$ we get (Cornman, et al., 2009)

$$\varphi_{\chi_1}(f) = \int_{\eta_1 - \Delta\eta/2}^{\eta_1 + \Delta\eta/2} \frac{16\pi^{5/2} k^2 A C_{n'}^2(\eta)}{V_\mu(\eta)} \tilde{y}^{-4/3}$$
$$\left\{\frac{2\Gamma(1/3)}{5\Gamma(5/6)} - \mathrm{Re}\left[e^{ia\tilde{x}^2} U\left[1/2, -1/3, -ia\tilde{y}\right]\right]\right\} d\eta \qquad (25)$$

where we have used the isotropic von Karman form for the index of refraction spectrum

$$\varphi_{n'}(0, \|\boldsymbol{q}\|^2) = A C_{n'}^2(\eta)\left[\|\boldsymbol{q}\|^2 + \left(2\pi/L\right)^2\right]^{-11/6} \qquad (26)$$

where A is a constant, $C_{n'}^2(\eta)$ is proportional to the intensity of the turbulence ("structure constant") at the point η, R is the distance along the line of sight from the transmitter to the receiver, k is the transmitter wavenumber, and L is a length scale related to the correlation function for the index of refraction field. Also in Eq. (25), we are assuming a patch of turbulence centered at $\eta = \eta_1$ with extent $\Delta\eta$; and we have denoted $a = \eta(R - \eta)/kR$, $\tilde{x} = 2\pi f/V_\mu$, $\tilde{y} = \tilde{x}^2 + (2\pi/L)^2$; Γ is the gamma function, and U is the confluent hypergeometric function of the second kind. The frequency spectrum of the phase is given by an expression similar to Eq. (25), excepting that the minus sign between the two terms in the curly brackets is replaced by a plus sign. If the turbulence patch is narrow enough along the line of
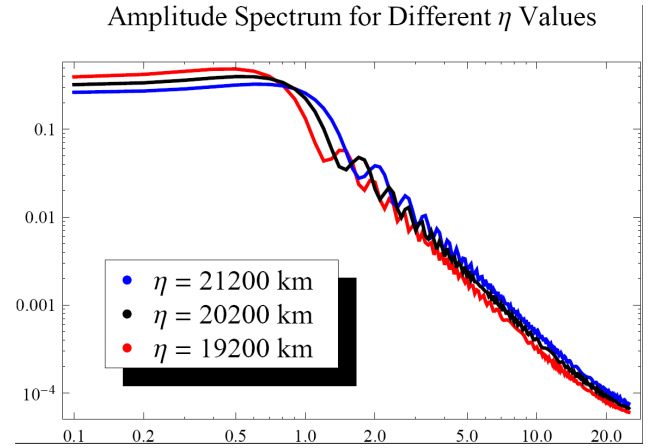

Fig. 3. Amplitude frequency spectrum (in Hz) with varying turbulence location values.

sight, then the integral over η can be approximated by a simple mid-point formula,

$$\phi_{\chi_1}(f) \approx \frac{16\pi^{5/2} k^2 A C_{n'}^2(\eta_1)\Delta\eta}{V_\mu(\eta_1)} \tilde{y}^{-4/3}$$
$$\left\{\frac{2\Gamma(1/3)}{5\Gamma(5/6)} - \mathrm{Re}\left[e^{ia\tilde{x}^2} U\left[1/2, -1/3, -ia\tilde{y}\right]\right]\right\} \qquad (27)$$

Since it is a multiplicative factor, $C_{n'}^2(\eta_1)\Delta\eta$ simply changes the overall level of the spectrum, and changes in the turbulence length scale essentially raises or lowers the spectral levels at the lowest frequencies. The changes in η_1 are more complicated, affecting both the amplitude and the frequency of the oscillations – as can be seen in Fig. 3 between one and two Hz. The variations in these parameters are not distinct, i.e., changes in one parameter can look like changes in another one. This fact makes parameter estimation difficult, as will be discussed subsequently.

3 Parameter estimation

Assuming that the received signal is a Gaussian random process, its spectral values at each frequency will have an exponential distribution. This result comes from the fact that the real and imaginary parts of the Fourier transform of the signal (divided by the variance) are both standard and independent Gaussian random variables. Hence, the spectrum, which is proportional to the sum of the squares of these quantities, is a sample from a Chi-Squared distribution with two degrees of freedom (cf. Theorem 8.8 in Freund, 2004). Furthermore, a Chi-Squared distribution with two degrees of freedom is an exponential distribution with mean and variance equal to two (cf. Definitions 6.3 and 6.4 in Freund, 2004). Therefore, the spectrum is distributed exponentially with mean and variance equal to twice the variance of the signal. It is important to note that this randomness is due to the random permittivity
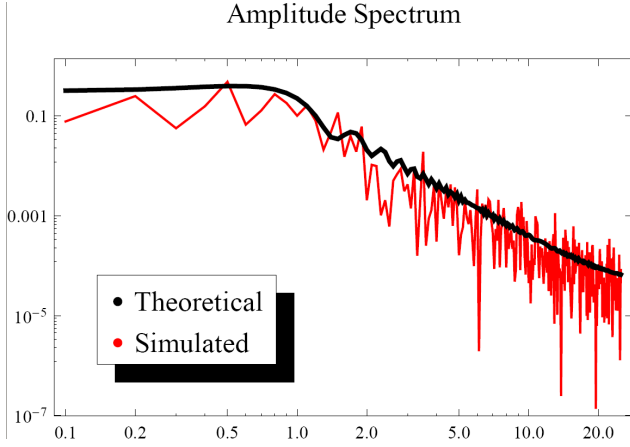
Amplitude Spectrum

Fig. 4. Amplitude frequency spectrum (in Hz) for theoretical and simulated case.

field – not the additive receiver and processing noise. Using the fact that the spectral points are exponentially distributed, numerous realizations of the spectra can be generated by replacing each spectral point with a sample from an exponentially distributed random number sequence. An example of this is shown in Fig. 4. The exponential probability density function is given by,

$$P(x) = \frac{1}{\beta} e^{-x/\beta} \tag{28}$$

where β is the mean value for the distribution. The probability distribution function is given by $P(x \leq b) = 1 - \exp(-b/\beta)$, so that $P(x \leq \beta) \approx 0.63$; and hence, more often than not, samples from this distribution will be less than the mean, as can be seen in Fig. 4. For an exponential distribution, the standard deviation is equal to the mean. Hence, for the simulation, the spectral level at a given frequency is given by replacing that point by a sample from an exponential distribution with its mean value being the theoretical spectral value at that point (i.e., from Eq. 27). The simulation is intended to represent real-world situations; hence, it was assumed that the received signal would be available at 50 Hz, and approximately 10 s worth of data would go into a single spectrum. The parameters used in the simulation are: $C_{n'}^2 \Delta \eta = 4 \times 10^{-9}$, $L = 5 \, \text{km}$, and $\eta_1 = 20\,200 \, \text{km}$. Furthermore, a noise floor, $Q = 3.5 \times 10^{-5}$, is added to the data. The magnitude of this level is derived from COSMIC GNSS-LEO data.

Referring back to Eq. (27), the parameters to be estimated are $C_{n'}^2(\eta_1)\Delta \eta$, L, η_1 and Q: the turbulence intensity, the turbulence length scale, the location of the center of the turbulence patch from the transmitter, and the noise level. As the functional form of Eq. (27) is highly non-linear in L and η_1, a simultaneous solution for all four parameters is complicated. The preferred solution is via the maximum likelihood

(ML) method (Freund, 2004). Consider a model of the form,

$$\mu_i = c\varphi_i(\eta_1, L) + Q \tag{29}$$

where "i" refers to the frequency, f_i, $c = C_{n'}^2(\eta_1)\Delta \eta$, and φ_i is the right-hand side of Eq. (27), divided by c. Consider a measured spectrum, denoted X_i. It is assumed that the X_i are independent exponential random variables with expected values $\langle X_i \rangle = \mu_i$. The likelihood function is then given by

$$L = \frac{e^{-\sum\limits_{i=1}^{N} \frac{X_i}{\mu_i}}}{\prod\limits_{i=1}^{N} \mu_i} \tag{30}$$

where N is the number of frequency values. (It is assumed that by context, the reader will not confuse the use of L for the likelihood function as well as the turbulence length scale.) The negative log-likelihood function is

$$-\ln L = \sum_{i=1}^{N} \left[\frac{X_i}{\mu_i} + \ln \mu_i \right] \tag{31}$$

The ML estimates for the parameters are found by solving the coupled set of non-linear equations, $\partial(-\ln L)/\partial p_j = 0$, where the p_j are the parameters, c, Q, L and η_1. Instead of dealing with this formidable problem, a simpler two-dimensional problem will be discussed to motivate the techniques. It will be assumed that the length scale L is known. Since the model, μ_i, is linear in c and Q, a linear weighted least squares approach can be employed to estimate those two parameters. Once that is done, these values can then be used in a one-dimensional ML solution for η_1. Some variations on this approach can also be used, and will be discussed below. Consider the least squares problem for c and Q. It will be the simultaneous solution to the coupled equations, $\partial R/\partial p_j = 0$, $(p_j = \{c, Q\})$, where $R = \sum\limits_{i=1}^{N} (\mu_i - X_i)^2 w_i$, and the weights w_i are independent of c and Q. If the weights are chosen such that $w_i = \varphi_i^{-1}$, then this is similar to a Chi-Squared minimization problem. The solution for the parameters can then be solved via Cramer's method for a 2×2 system of equations, giving

$$c = \frac{\sum\limits_{i,j=1}^{N} \frac{X_i}{\varphi_j} - N \sum\limits_{i=1}^{N} \frac{X_i}{\varphi_j}}{\sum\limits_{i,j=1}^{N} \frac{\varphi_i}{\varphi_j} - N^2}, \quad Q = \frac{\sum\limits_{i,j=1}^{N} \frac{\varphi_i X_j}{\varphi_j} - N \sum\limits_{i=1}^{N} X_i}{\sum\limits_{i,j=1}^{N} \frac{\varphi_i}{\varphi_j} - N^2} \tag{32}$$

An alternative is to calculate Q directly from the high-frequency parts of X_i. Denoting this estimate, $\bar{Q}$, a one-dimensional weighted least squares solution then gives

$$c = \frac{\sum\limits_{i=1}^{N} (X_i - \bar{Q})}{\sum\limits_{i=1}^{N} \varphi_i} \tag{33}$$

Fig. 5. Negative log-likelihood function as a function of η_1 in km, with varying turbulence location.
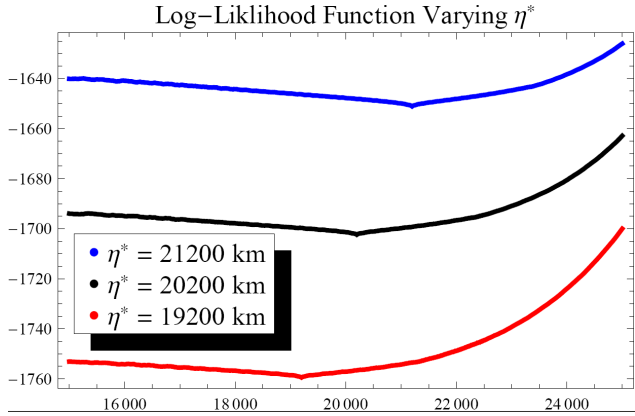

Fig. 6. Negative log-likelihood function as a function of η_1 in km, for varying turbulence length scale.

A further alternative – the one that will be employed in the following – is to calculate c via a "pseudo" ML method. That is, we assume that the data X_i is replaced by $\bar{X}_i = X_i - \bar{Q}$, and then solve for c via the ML formula, $\partial(-\ln L)/c = 0$. Letting $\bar{\mu}_i = c\varphi_i$

$$\frac{\partial}{\partial c}\sum_{i=1}^{N}\left[\frac{\bar{X}_i}{\bar{\mu}_i} + \ln\bar{\mu}_i\right] = -\sum_{i=1}^{N}\left[\frac{\bar{X}_i}{\varphi_i}\right] + Nc \qquad (34)$$

Equating this to zero, gives the desired ML-like estimate for c,

$$c(\eta) = \frac{1}{N}\sum_{i=1}^{N}\frac{\bar{X}_i}{\varphi_i(\eta)} = \frac{1}{N}\sum_{i=1}^{N}\frac{X_i - \bar{Q}}{\varphi_i(\eta)} \qquad (35)$$

Plugging this back into Eq. (31) gives

$$-\ln L(\eta) = N[1 + c(\eta)] + \sum_{i=1}^{N}\ln\varphi_i(\eta) \qquad (36)$$

Where the explicit dependence on η has been indicated (recall that we are not including dependence on the length scale L at this point). Note that the last term is independent of the data and can thereby be computed a priori. Equation (36) is now a one-dimensional function and the solution for η_1 is given by minimizing it with respect to η. For the results presented below, we merely take the minimum value of Eq. (36) as a function of η as the estimate for η_1. Once the estimate for η_1 is determined, this value is plugged back into Eq. (35) to estimate c.

Different values of c will shift the negative log-likelihood curves up or down (for simplicity in terminology, the phrase "likelihood" will mean "negative log-likelihood."). Figures 5 and 6 show the results of Eq. (36) when varying η_1, and L. Changes in η_1 and L change the likelihood function in complicated fashions, not only varying the level of the curves, but also their shape. The flatness of the likelihood function at the smaller η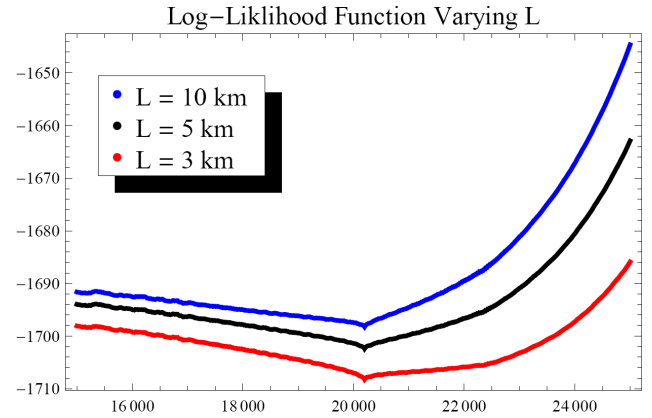 poses a problem for the η_1 estimation – one would like a steep bowl leading down to the correct minimum. The flatness means that variations in the likelihood function – due to the exponentially distributed random errors in the spectral values – can cause a false minimum to be smaller than one for the true minimum. A simple method was developed to deal with this problem: using a weighting function in the denominator of Eq. (35). This function was chosen to be solely dependent on η. The only portions of the line of sight that are of interest are those that go through the atmosphere or at least portions of it where turbulence can be strong enough to disturb the signal significantly. For example, excluding outer regions of the ionosphere where the permittivity perturbations due to fluctuations in the electron density field can be small. Hence, the weighting function can be chosen such that it is close to one for those portions of interest along the line of sight, and drops smoothly to zero away from those regions. The effect of the weighting function on the likelihood function is illustrated in Fig. 7. It can be seen that the likelihood function is unchanged essentially, between approximately 18 000 to 22 000 km. This is important, as the purpose here is not to bias the estimates, but to try to eliminate non-physical estimates.

In order to examine the statistical properties of the c and η_1 estimators, 1000 realizations of the simulated spectra were generated. Before discussing these results, however, some preliminary concepts are presented.

3.1 Theoretical error analysis

Since the random errors in η_1 are not easily described theoretically, these will be analyzed in the context of the simulation results. However, the error statistics of c can be studied analytically. Consider then the estimator for c as given by Eq. (35). Recall that this was referred to as a pseudo-ML estimator. This is because the correct ML estimate is given by
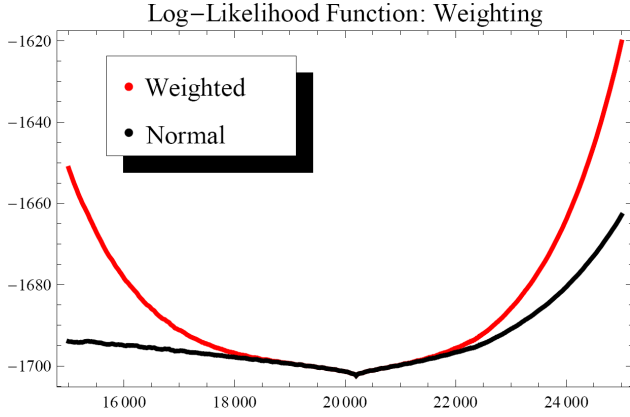
Fig. 7. The effect of the weighting function on the log-likelihood function as a function of η_1 in km.

using

$$-\ln L = \sum_{i=1}^{N} \left[\frac{X_i}{c\varphi_i + Q} + \ln(c\varphi_i + Q) \right] \tag{37}$$

Then, c is found by solving

$$0 = \frac{\partial(-\ln L)}{\partial c} = \sum_{i=1}^{N} \left[-\frac{X_i}{(c\varphi_i + Q)^2 \varphi_i} + \frac{\varphi_i}{c\varphi_i + Q} \right] \tag{38}$$

which is now non-linear in c, thus requiring a numerical solution. Therefore, we restrict our analysis to the "pseudo-ML" estimate. First, consider the expected value of c,

$$\langle c \rangle = \frac{1}{N} \sum_{i=1}^{N} \frac{\langle X_i \rangle - \langle \bar{Q} \rangle}{\varphi_i(\eta_1)} = \frac{1}{N} \sum_{i=1}^{N} \frac{\mu_i - Q}{\varphi_i(\eta_1)}$$

$$= \frac{1}{N} \sum_{i=1}^{N} \frac{c\varphi_i(\eta_1) + Q - Q}{\varphi_i(\eta_1)} = c \tag{39}$$

Which shows that this is an unbiased estimator of c (where it has been assumed that $\bar{Q}$ is an unbiased estimator of Q, and η_1 is the true value for the parameter η). Next, consider the variance of c – with the same assumptions made above,

$$\text{Var}(c) = \text{Var}\left(\sum_{i=1}^{N} \frac{\bar{X}_i}{N\varphi_i(\eta_1)} \right) \tag{40}$$

It can be shown that the variance for a linear combination of independent random variables, γ_i, is given by (Freund, 2004)

$$\text{Var}\left(\sum_i a_i \gamma_i \right) = \sum_i a_i^2 \text{Var}(\gamma_i) \tag{41}$$

Hence, with $a_i = 1/(N\varphi_i)$, Eq. (40) becomes

$$\text{Var}(c) = \frac{1}{N^2} \sum_{i=1}^{N} \frac{\text{Var}(\bar{X}_i)}{[\varphi_i(\eta_1)]^2} \tag{42}$$

Since it has been assumed that the X_i are independent, exponentially distributed random variables, their variance is equal to the square of their mean, and Eq. (42) becomes

$$\text{Var}(c) = \sum_{i=1}^{N} \frac{(c\varphi_i(\eta_1) + \bar{Q})^2}{N^2 \varphi_i^2(\eta_1)} \tag{43}$$

Note that if $\bar{Q} = 0$, $\text{Var}(c) = c^2/N$.

Things are more complicated when η_1 is not the true value. Denote this, $\bar{\eta}_1$. The estimated value for c (denoted $\bar{c}$) is now

$$\langle \bar{c} \rangle = \frac{c}{N} \sum_{i=1}^{N} \frac{\varphi_i(\eta_1)}{\varphi_i(\bar{\eta}_1)} \tag{44}$$

and hence the bias, $b = \langle c \rangle - \langle \bar{c} \rangle$ is given by

$$b = c \left[1 - \frac{1}{N} \sum_{i=1}^{N} \frac{\varphi_i(\eta_1)}{\varphi_i(\bar{\eta}_1)} \right] \tag{45}$$

Therefore, it can be seen that the bias is a function of the sum of the ratios of the spectrum evaluated at the true and estimated η_1 values, respectively. Note that if $\bar{\eta}_1 = \eta_1$, then the bias will be zero, as expected.

The variance of $\bar{c}$ is given by Eq. (43), but with $\bar{\eta}_1$ in the denominator. Note that if $Q = 0$, the variance of $\bar{c}$ is

$$\text{Var}(\bar{c}) = \left(\frac{c}{N} \right)^2 \sum_{i=1}^{N} \left[\frac{\varphi_i(\eta_1)}{\varphi_i(\bar{\eta}_1)} \right]^2 \tag{46}$$

Since the X_i are exponentially distributed random variables, the random error is significant – as can be seen in Fig. 4. In order to reduce the random error, averaging spectra is helpful. Consider the statistics of the estimator of c with averaged spectra (denoted $\hat{c}$). The averaged spectral values are given by

$$\hat{X}_i = \frac{1}{M} \sum_{j=1}^{M} (X_i)_j \tag{47}$$

where j indexes the elements that go into the average – of which there are M. The estimator, $\hat{c}$, is then given by

$$\hat{c} = \frac{1}{N} \sum_{i=1}^{N} \frac{\hat{X}_i}{\varphi_i} = \frac{1}{NM} \sum_{i=1}^{N} \sum_{j=1}^{M} \frac{(X_i)_j}{\varphi_i} \tag{48}$$

Interchanging the order of summation and summing over i gives, $\hat{c} = \sum_{j=1}^{M} c_j/M$. Taking the expected value gives

$$\langle \hat{c} \rangle = \frac{1}{M} \sum_{j=1}^{M} \langle c_j \rangle = \langle c \rangle \tag{49}$$

where it was assumed that the $\langle c_j \rangle$ are identical. This shows that the estimator from the averaged spectra is unbiased. Using Eq. (41), the variance of $\hat{c}$ is given by

$$\text{Var}(\hat{c}) = \text{Var}\left(\sum_{j=1}^{M} c_j/M \right) = \sum_{j=1}^{M} \frac{\text{Var}(c_j)}{M^2} = \frac{\text{Var}(c)}{M} \tag{50}$$

where it was assumed that the variances of c_j are identical. This shows that averaging the spectra reduces the variance of c by a factor of $1/M$.

As seen in Eqs. (44) and (46), the ratio $\varphi_i(\eta_1)/\varphi_i(\bar{\eta}_1)$ is important in determining the error properties of the estimator $\bar{c}$. In order to study this further, we will need the following concepts. The probability density function for a gamma distribution is given by

$$P(x) = \frac{x^{\alpha-1} e^{-x/\beta}}{\beta^\alpha \Gamma(\alpha)} \tag{51}$$

If X is gamma-distributed with parameters α and β, and ν is a positive constant, then νX is gamma-distributed with parameters α and $\nu\beta$. Since the spectral values X_i are assumed to be independent exponentially distributed random variables – or gamma-distributed with $\alpha = 1$ and $\beta = \mu_i$ – it follows that the ratio X_i/μ_i will be exponentially distributed with mean one (use the property just discussed, with $\nu = \mu_i$). If the μ_i are evaluated with the estimates for Q, c and η_1, then the ratio should be approximately distributed as an exponential distribution with mean one. The same holds for averaging the spectra, but here the ratio $\hat{X}_i/\mu_i$ will be gamma-distributed with $\alpha = M$ and $\beta = 1/M$ (Freund, 2004) (recall that $\hat{X}_i$ is given by Eq. 47). These properties can then be used to test the quality of the estimation methods via a goodness-of-fit analysis.

4   Simulation studies

As mentioned above, 1000 independent realizations of the log-amplitude spectrum were generated. Equation (36) was used to calculate estimates of η_1, and then those are in turn are used to estimate c via Eq. (35), with $\eta = \eta_1$. Recall that the true values for these parameters are $\eta_1 = 20\,200$ and $c = 4 \times 10^{-9}$. There are four different estimation methodologies that were employed: the pseudo-ML method, one in which a weighting function is used, and then the same methods applied to averaged spectra. Five spectra were used in the averaged-spectrum analysis. A discussion of the overall estimator performance is presented first, followed by a set of case study analyzes that deal with more specific aspects.

4.1   Overall statistical performance

Figures 8 and 9 illustrate the overall simulation results. These are hybrid plots, with both the coloring and widths being related to the number of samples at that specific parameter value. Furthermore, the smaller purple squares indicate outliers. Tables 1 and 2 present some descriptive statistics for the results. Most of the statistics used in the tables are order statistics, i.e., quantiles and inter-quantile ranges. These metrics are more robust in the presence of outliers, and so they are more appropriate here than other dispersion metrics such as skewness and kurtosis. It can be seen that
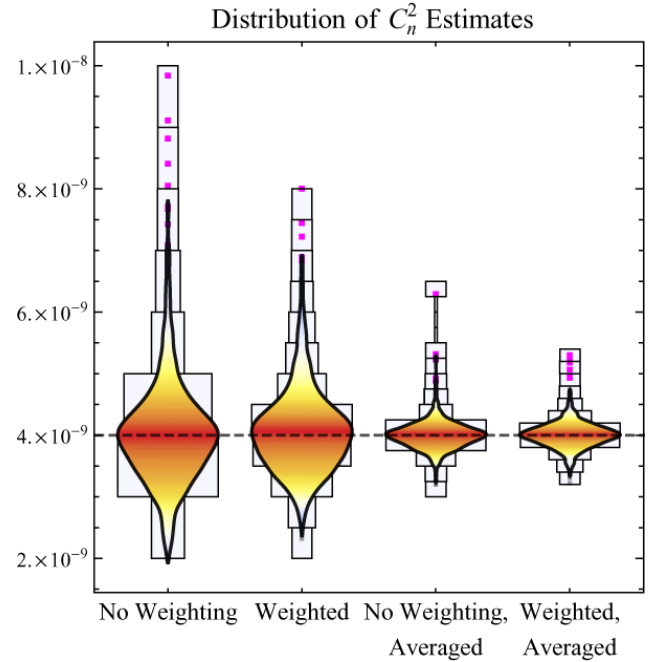

Fig. 8. Distribution of $C_\eta^2 \Delta\eta$ estimates over 1000 realization. The dashed horizontal line is the true value.

the median values are all quite close to the expected value, reflecting the comments above about how the estimators are unbiased. Those comments referred to the estimator for c, but that fact is also represented in the median value of the estimator for η_1. The use of the weighting function has the desired effect: it eliminates the most egregious outliers. However, from the tables and the figures, it can be seen that the decrease in outliers comes at the price of increasing the errors in the "middle" statistics (e.g., the 40–60 and 30–70 quantile ranges). It is clear that spectral averaging is a powerful method for reducing the random errors – even with averaging only five spectra. The beneficial effects of spectral averaging are far greater than using a weighting function. A definitive bias towards larger outliers in c can be seen in Fig. 8. This is more clearly exemplified in Fig. 10, a contour-scatterplot of the c estimates versus the η_1 estimates. The horizontal and vertical lines indicate the true value of those parameters. The solid black curve comes from the sum in Eq. (44) and subsequently embodied in the equation for the bias, Eq. (45). This plot show how the bias is a fundamental aspect of the problem, not an artifact of the processing methodology. Recall that the bias comes from the sum of the ratios of the true log-amplitude spectrum (i.e., with the true value of η_1) divided by the log-amplitude spectrum evaluated at a different η_1. As can be seen from Fig. 3, excepting for the lowest frequencies, the spectrum for the $\eta_1 = 19\,200$ curve is always lower than that for $\eta_1 = 20\,200$; hence, the sum will be greater than one, thus producing an overestimate of c. Alternatively, if the estimated value of η_1 is greater than the true value, the sum of
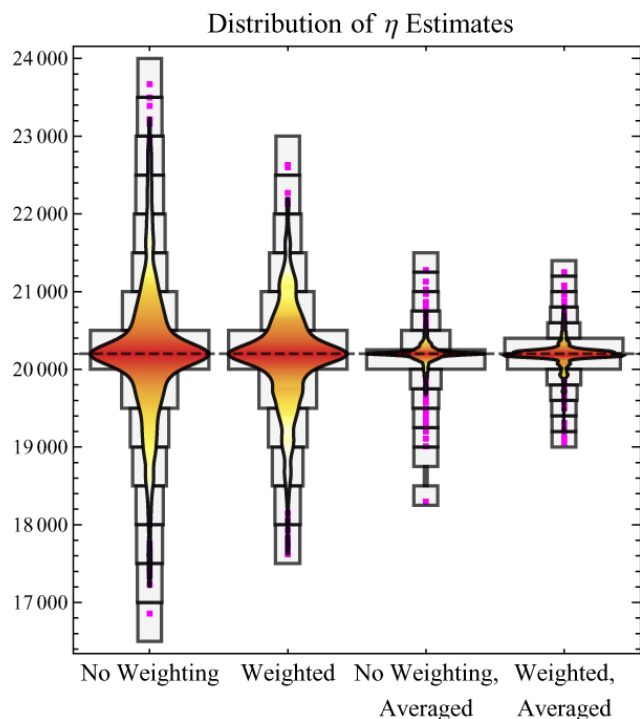
Fig. 9. Distribution of η_1 estimates. The dashed horizontal line indicates the true value.

Fig. 10. Contour-scatterplot of $C_\eta^2 \Delta\eta$ estimates versus η_1 estimates. The horizontal and vertical lines indicate the true value values for those parameters, respectively. The solid black curve is from Eq. (45).

the ratios will be less than one and will give an underestimate of c. This explains the biased structure for the estimates of c, as seen in Figs. 8 and 10.

A significant concern in any estimation algorithm is the identification and elimination (or mitigation) of outliers. In the results just presented, it could be seen that the overall statistical performance of the estimators is quite good, but in a real world application one wants to know if a given parameter estimate is accurate or not. This is a difficult problem in general and even more so in this application because the parameters do not have independent effects on the estimators. As mentioned above, the ratios of the true (in this case, simulated) spectra to the model spectra with the estimated parameters should be approximately distributed as an exponential random variable with mean one. The collection of ratios in question is taken over the set of frequencies samples for the given realization. An investigation was made into whether fitting an exponential distribution to the ratios could be used to determine the accuracy of the parameter estimates. This is a standard technique in the category of "goodness-of-fit" tests. A number of tests were employed: Anderson-Darling, Cramer-von Mises, Kolmogorov-Smirnov, Kuiper, Pearson chi-squared, and the Watson U^2 test (Lindgren, 1968). Each test produces a so-called test statistic, which then leads to the probability (or "p-value") that the given data comes from the specified distribution. For each realization, the ratios were calculated and then tested as to how likely they were to come
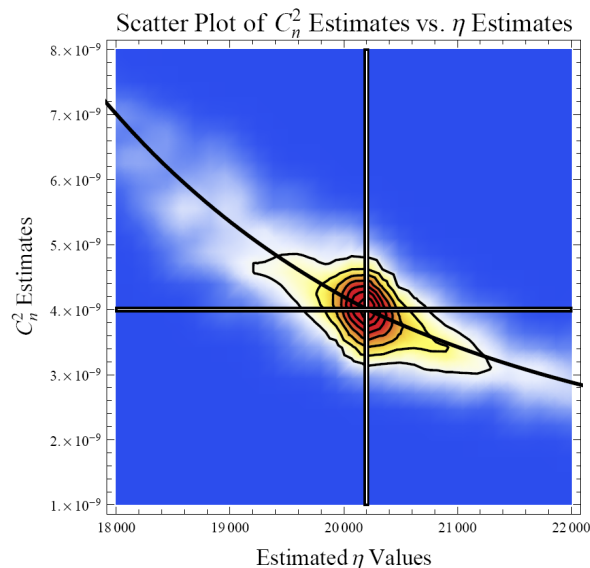
from a mean-one exponential distribution. Figure 11 illustrates the results using the chi-squared test, giving the p-value as a function of the c estimate for each of the 1000 realizations. The color-coding is proportional to the number of counts at that (c, p-value) grid cell. An optimal result would produce something like an upside-down "V" pattern (though not uniform in the number of counts), with the base of the two legs at the zero p-value axis, and the apex at the true c value and a p-value of one. That is, the higher the p-value (i.e., closer to one), the better the estimate – and vice-versa. As can be seen from the figure, the chi-squared test is clearly not optimal. There is skill, but there are too many mid to low p-values associated with good c estimates, as well as some poorer c estimates with high p-values. In other words, the p-value for this test could not be used as a single criterion for identifying outliers. The Pearson chi-squared test is used so frequently that a comment is warranted as to one reason that it does not perform well in this case. As discussed in Lindgren (1968, p. 486), the Pearson test is not well-suited to continuous distributions, the difficulty laying in choosing the appropriate binning strategy. The other tests, such as the Kolmogorov-Smirnov and Anderson-Darling, which use the distribution functions directly, do not have this limitation. However, it turned out that none of the individual tests had sufficient skill to be used in outlier detection. Thus, a simple hybrid method was used as the goodness of fit metric – the maximum p-value over all the tests. The results of this method for the c parameter estimation are shown in Fig. 12.

Table 1. Statistics for maximum likelihood estimates of $C_n^2 \Delta \eta$.

$C_n^2 \Delta \eta$	Median ($\times 10^{-9}$)	STD ($\times 10^{-9}$)	40–60 Quantile Range	30–70 Quantile Range	20–80 Quantile Range	10–90 Quantile Range
Basic	3.98	0.90	0.31	0.67	1.13	2.00
Weighted	4.01	0.70	0.29	0.60	0.96	1.64
Spectral Average	4.00	0.24	0.09	0.18	0.32	0.50
Weighted, Spectral Average	4.00	0.22	0.08	0.18	0.29	0.48

Table 2. Statistics for maximum likelihood estimates of η_1.

η_1 (km)	Median	STD	40–60 Quantile Range	30–70 Quantile Range	20–80 Quantile Range	10–90 Quantile Range
Basic	20 200	903	15	400	1010	2045
Weighted	20 190	653	90	360	845	1560
Spectral Average	20 200	197	10	40	120	280
Weighted, Spectral Average	20 200	176	30	60	100	270

While certainly not optimal, this approach is a significant improvement over the individual tests. For the η_1 estimates, there was more skill than in the c estimation scenario, but as a stand-alone outlier detection method, this goodness-of-fit method is still not adequate. In the next section, a number of case studies are discussed, which leads to a better understanding of the difficulties in the estimation and outlier identification problem.

4.2 Simulation case studies

Four cases out of the 1000 realizations are analyzed in the subsequent material: realizations 64, 127, 177, and 356. These four were chosen to study the relationship between p-values and parameter estimates. Specifically, the four cases illustrate the combinations of high and low p-values with both good and bad parameter estimates, (i.e., high-good, high-bad, low-good and low-bad). Realization 127 is also used to show how statistical comparisons between the log-likelihood functions derived from the simulated data and parameter estimates can be misleading.

The first case studied is realization number 64. Figure 13 through Fig. 15 show the spectra, likelihood function and distribution of ratios for this realization, respectively. Figure 13 shows the true spectrum (black), simulated spectrum (red), and the spectrum with the estimated c and η_1 values (green). Another spectrum function is plotted, (labeled "fit," in blue), where the parameters were set "by-eye" to get a better fit to the likelihood function (more on this below). Figure 14 illustrates the likelihood functions derived from the true spectrum (black); the simulated spectrum (red); the spectrum from the estimated parameters (green); and the spectrum from the fitted parameters (blue). Figure 15 shows a histogram of the spectral ratios using the spectrum derived from the estimated parameters, a curve representing a smoothed histogram (red), and a curve which is a mean-one exponential density function (black). Table 3 gives the estimated values of c and η_1, a "low" and "high" rating of the p-values, and a chi-squared type of error function, computed as

$$E_E = \sum_{i=1}^{P} \frac{[-\ln L_D(\eta_i) - (-\ln L_E(\eta_i))]^2}{\ln L_E(\eta_i)} \tag{52}$$
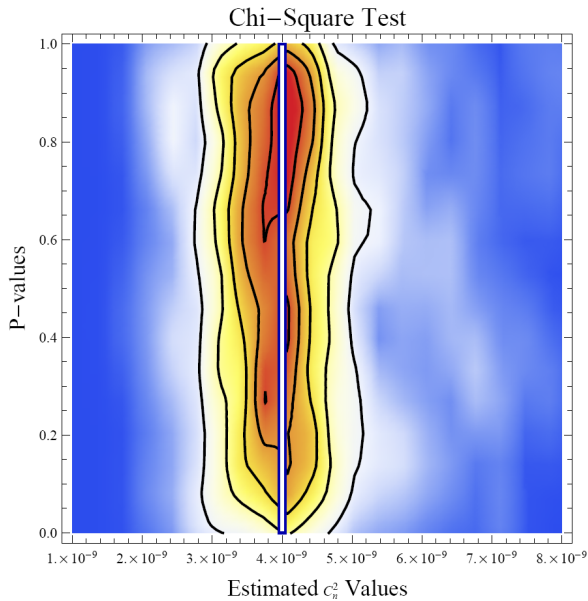
Fig. 11. Probability values as a function of $C_\eta^2 \Delta\eta$ estimates: Chi-squared test.
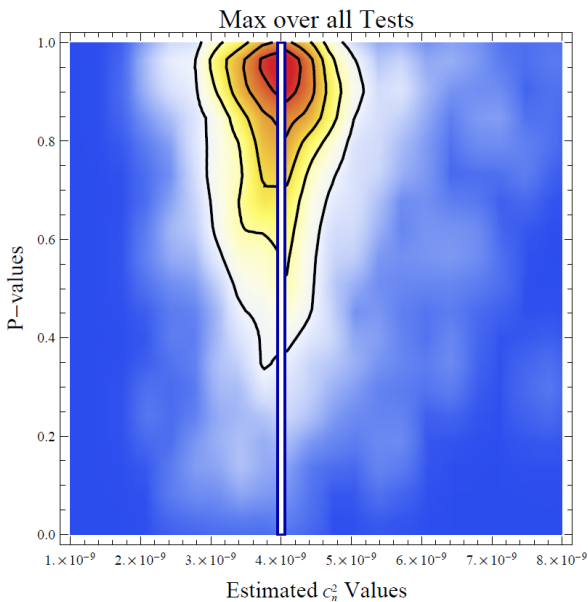

Fig. 12. Probability values as a function of $C_\eta^2 \Delta\eta$ estimates: maximum over all tests.

where, the sum is over all η values (15 000 to 25 000 km with a resolution of 10 km, i.e., $P = 1000$), $-\ln L_D$ is the negative log-likelihood calculated from the data, and $-\ln L_E$ is that derived from the estimated parameters. Another statistic, E_F, is calculated using the log-likelihood function from the fitted

Fig. 13. Amplitude frequency spectrum (in Hz) for realization 64, showing fit ($\eta = 20\,200$, $L = 5\,\mathrm{km}$, and $C_\eta^2 \Delta\eta = 3.84 \times 10^{-9}$), compared to estimated ($\eta = 17\,210\,\mathrm{km}$, $L = 5\,\mathrm{km}$, and $C_\eta^2 \Delta\eta = 8.4 \times 10^{-9}$).

Fig. 14. Log-likelihood functions, as a function of η_1 in km for realization 64, showing hand-fitted as well as estimated and true functions.

parameters instead of that from the estimated ones. Note that these statistics are not goodness of fit metrics for the fit to the spectrum, but rather for the fit to the likelihood function. The reason for this choice is that the parameter estimates come from the likelihood function – not the spectrum.
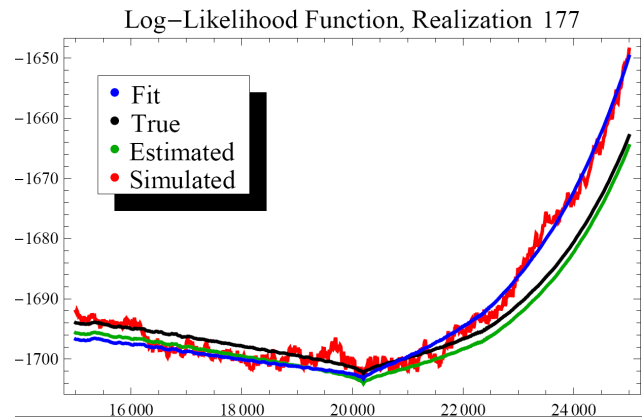
From Table 3, it can be seen that the parameter estimates are quite poor for realization 64, and the error statistics are moderately large (relative to the other realizations). This implies that in this case, neither the E_E and E_F statistics, nor the difference between them would indicate that the parameters are significantly in error. The p-values for this realization are very high, which indicates that they are also not helpful in determining whether the parameters are outliers. This is reinforced by how well the distribution of the ratios matches the exponential curve – as seen in Fig. 15. As mentioned previously, the random errors for an exponential distribution

Table 3. Parameter estimates and statistical measures for selected realizations.

Realization Number	η_1	$C_n^2 \Delta\eta$	P-values	Log-Likelihood Error Statistic (E_E)	Log-Likelihood Error Statistic (E_F)
64	17 210	8.4×10^{-9}	High	100	54
127	20 200	4.0×10^{-9}	High	45	8
177	20 200	4.0×10^{-9}	Low	16	2
356	17 320	9.1×10^{-9}	Low	284	10

Fig. 15. Distributions plots of the ratios in Eq. (45), for realization number 64.

Fig. 16. Log-likelihood functions, as a function of η_1 in km for realization 177, showing hand-fitted as well as estimated and true functions.

are large, and hence the log-likelihood function can have undulations, which can result in a global minimum offset from the "true" minimum. It was also mentioned that this issue is especially problematic for smaller η values, where the log-likelihood function has a shallow incline. This is exactly what occurs with this realization, as clearly seen in Fig. 14. The model parameters chosen "by eye" produce a somewhat better fit to the empirical log-likelihood function. This was a surrogate for a more sophisticated algorithm which might find the best fit to the empirical log-likelihood by varying the c and η_1 (as well as L) parameters as part of a three-dimensional minimization algorithm. This gives hope for more accurate parameter estimation by using a more sophisticated minimization method, e.g., estimating L in addition to c and η_1. Since the values for c and η_1 are more important for practical applications, the error induced by lowering the value for L, in exchange for more accurate c and η_1 estimates is acceptable. However, as seen in Figs. 5, 6, and Eq. (36), the effect on the log-likelihood function due to variations in these parameters are not independent. This is quite apparent in the effect that changing η (cf. Fig. 5) has on the log-likelihood function. That is, both the level and minimum value change with varying η.

Realizations 356 and 177 both have low p-values with poor and excellent parameter estimates, respectively. E_E is quite large for realization 356; however, adjusting c in a minor fashion (to 4.3×10^{-9}) and L in a more significant way (to 9), produced an excellent fit to the empirical likelihood function ($E_F = 10$). Another interesting aspect of this case is that the fitted log-likelihood function, which produced an excellent fit to the empirical one, is still off in level to the true one. This is due to the random errors that produce a simulated log-likelihood function that is off in level. This means that minimizing the E_F function can still produce parameters that are in error. Realization 177 resulted in excellent parameter estimates, as well as a very good E_E value. The fit did not change the parameter estimates appreciably, although E_F is better than E_E. This case is interesting in that the estimated log-likelihood function is well matched to the true one – as seen in Fig. 16 – whereas the fit log-likelihood function is well matched to the simulated one. Of course, in a real world scenario, the true log-likelihood function would not be available so one would have to depend on the fit to give all the information.

This case also brings up an interesting point, one that is even more apparent in realization 127 (cf. Fig. 17): that good
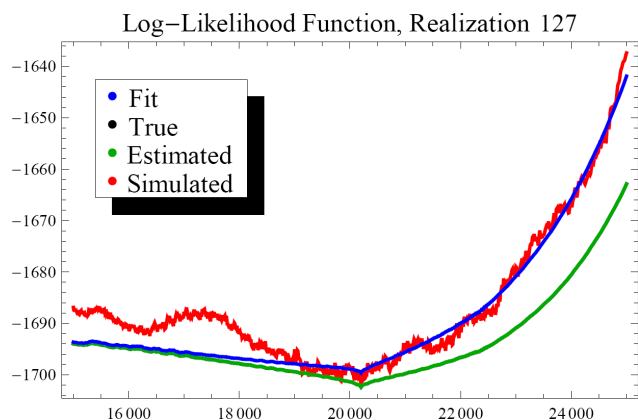
Fig. 17. Log-likelihood functions, as a function of η_1 in km for realization 127, showing hand-fitted as well as estimated and true functions.

Fig. 18. Log-likelihood functions (as a function of η_1 in km) – empirical and estimated – for the basic and weighted parameter estimation methods. The data are from realization 64.

parameters estimates only require that the estimates of the level (c) and location of the minimum (η_1) match those of the true data, i.e., at a single point in the ($\eta_1, -\ln L_D$) domain. On the other hand, getting a good fit to the log-likelihood function requires reasonable matches at many points. No combination of parameters might produce a good fit to the "measured" (real or simulated) log-likelihood function; however, if the fit and true curve matched at just the right point, the estimates would be perfect. This implies that the fitting method should be more robust in parameter estimation – or at a minimum it could provide information regarding data quality. The hope that p-values could provide some measure of confidence in the parameter estimation was not fulfilled. Nevertheless, it was found that distributions with numerous outliers (i.e., large ratios) tended to produce lower p-values. This leads to two possible augmentations for the future: re-calculating the p-values with the outliers removed, and using that as a revised confidence measure; and removing the outlier points from the parameter estimation process.

4.3 Weighting and spectral averaging

As mentioned above, other techniques that can improve parameter estimation entails the use of a weighting function in calculating the log-likelihood function, and spectral averaging. Figure 18 shows the benefit in using a weighting function. The red and brown curves are the un-weighted empirical and estimated likelihood curves, respectively, for realization 64. Recall that the parameter estimates for this case were poor ($c = 8.4 \times 10^{-9}$ and $\eta_1 = 17210$). Since the η_1 estimate is on the low end of the $\eta-$ range, the use of a weighting function should improve the estimates. As can be seen from the figure (blue and green curves for the weighted empirical and estimated likelihood curves, respectively), the new estimates are superior to the un-weighted method: $c = 3.98 \times 10^{-9}$ and $\eta_1 = 20150$. Tables 1 and 2
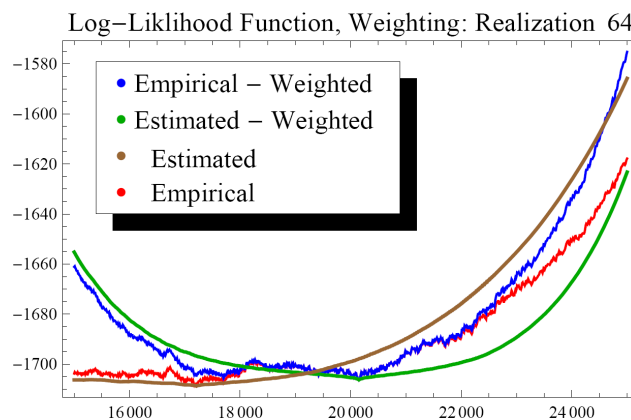
give the statistical results over 1000 realizations. Next, we turn to spectral averaging. A small number of realizations (5) are averaged to show the benefits of this method. Of course, for a GNSS-LEO occultation geometry, the vertical extent of the atmosphere covered over the period required to average five spectra can be large. Running averages could be used, but then the samples would not be independent and the effect of averaging on the variance – as given by Eq. (50) – would be lessened. Figure 19 illustrates the benefit of averaging on reducing the random errors. The red curve is the averaged spectrum and the black curve is the true spectrum. Comparing this case with, for example, the case shown in Fig. 4, it can be seen that the averaged spectrum is matched better to the true spectrum, and hence the parameter estimation should be improved similarly. Figure 20 shows the log-likelihood functions from the true and averaged spectrum. Again, the beneficial effects of averaging can be seen, i.e., the log-likelihood function generated from the averaged spectrum is better matched to the true log-likelihood function. Tables 1 and 2 give the statistical results over 1000 realizations. It can be seen that by reducing the random error, spectral averaging improves the parameter estimation much more significantly than the weighting scheme. It was shown above that when averaging spectra, the distribution function for the ratios becomes a gamma distribution with parameters $\alpha = M$ and $\beta = 1/M$; where M is the number of spectra going into the average (here, five). This can be seen in Fig. 21 with the simulated data.

5 GPS-COSMIC occultation case studies

The theory, techniques, and observations presented in the previous sections would be inadequate if they did not pertain directly to real data. In this section, an analysis of two case studies from GPS-COSMIC occultations is presented. The purpose of these case studies is to show how the methods and
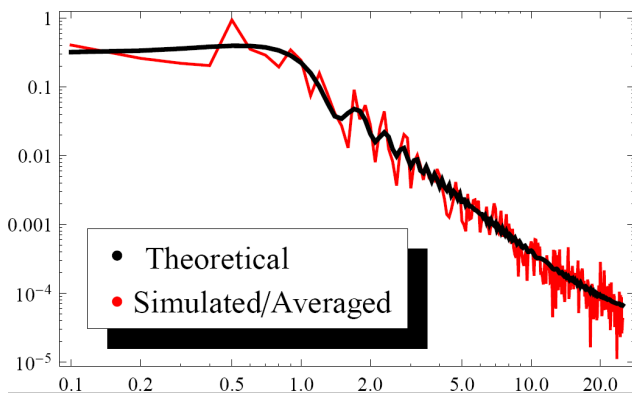
Amplitude Spectrum: Spectral Averaging
Realization 25

Fig. 19. Averaged and theoretical frequency spectra (in Hz). Five spectra, beginning at realization 25, were averaged.
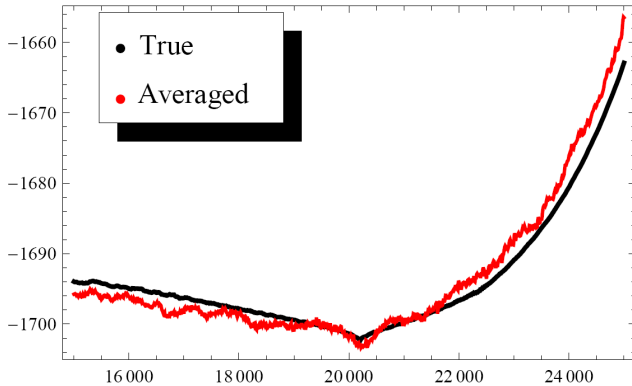
Log−Liklihood Function: Spectral Averaging
Realization 25

Fig. 20. Log-likelihood function, as a function of η_1 in km, derived from averaged and theoretical amplitude frequency spectra. Five spectra, beginning at realization 25, were averaged.

Distribution Plot: Spectral Averaging
Realization 25

Fig. 21. Distribution of ratios for the five averaged amplitude frequency spectra shown in Fig. 19, above. Note that the theoretical density function is from a gamma distribution.

SNR Data − Trend Removal
COSMIC Occultation 2, (2300,2800)

Fig. 22. SNR data for COSMIC occultation number 2 (black). Trend curve is in red and the vertical, dashed blue lines indicate the analysis window.

results from above apply to the real world scenario – not as an in-depth analysis of the real data. These cases were selected since they were from a region in which there was vigorous convective activity, as well as turbulence reports from pilots in the vicinity. In each case, a 500-sample analysis window was selected and processed using the same methodologies as used with the simulated data. There is one difference, however. A trend was removed from the amplitude time series data prior to the analysis. The trend was computed via a running median filter with a width of 2001 points, or approximately 40 s. The appropriateness of the filter window was checked by verifying that the residuals had a mean that was close to zero. The noise level was computed by taking the median value of the log-amplitude spectrum over the last 50 points (i.e., the high-frequency region).
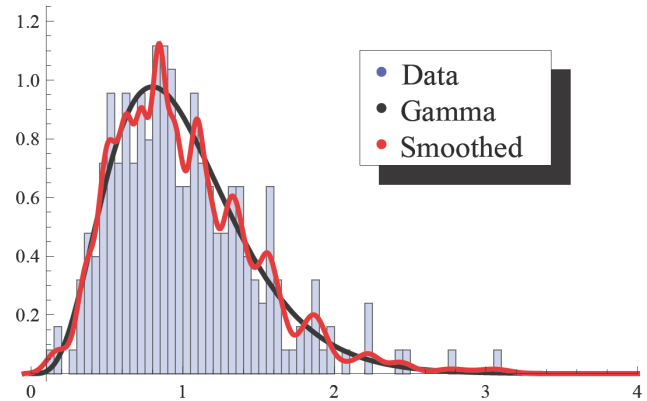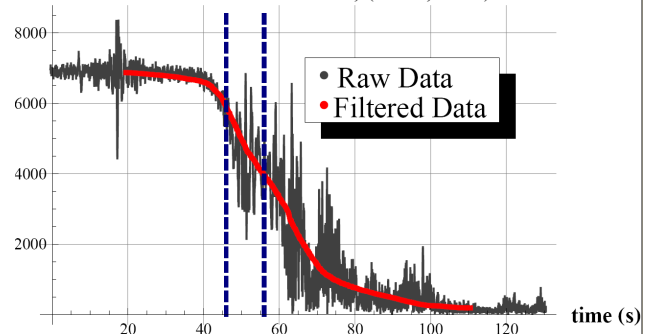
Figure 22 shows the amplitude time series over the course of occultation 2 (arbitrary nomenclature). The black curve is the raw data, the red curve shows the trend, and the two vertical dashed blue lines delineate the analysis window. This window was chosen because it had large-amplitude fluctuations, and could be used to test subjectively whether some of the assumptions presented in Sect. 2.1 were violated; specifically whether inhomogeneity and/or layered structures would affect the results. Figure 23 shows the spectra of the data (red), one using the estimated parameters (in green, with $c = 5.1 \times 10^{-10}$, $\eta_1 = 19010$, $L = 5.0$ and $Q = 7.0 \times 10^{-6}$), and one using the fitted parameters (in blue, with $c = 5.19 \times 10^{-10}$, $\eta_1 = 19010$, $L = 3.1$ and $Q = 6.0 \times 10^{-6}$). The likelihood functions for the data (red), the estimated parameters (green) and fitted parameters (blue) are presented in Fig. 24. For this case, $E_E = 31.6$ and $E_F = 2.1$. Figure 25
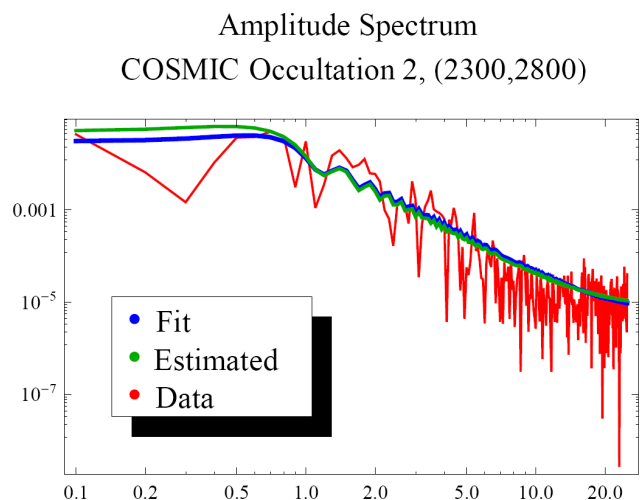
Fig. 23. Amplitude frequency spectrum (in Hz) for occultation 2 (red), model spectrum using estimated parameters (green), and "by-eye" fitted parameters (blue).
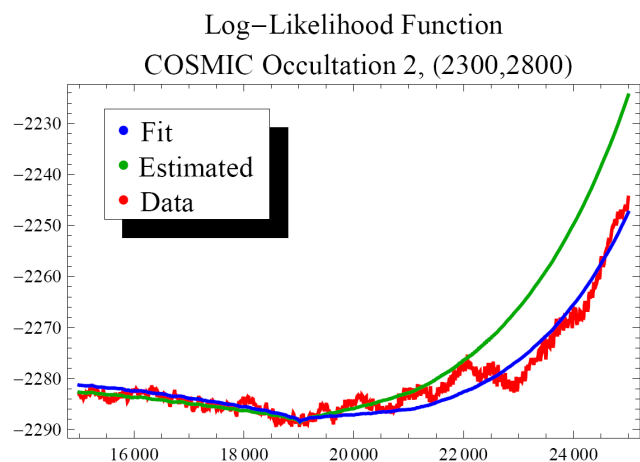

Fig. 24. Log-likelihood functions, as a function of η_1 in km, from data (red), estimated parameters (green), and fitted parameters (blue) for occultation 1.

illustrates the distribution of ratios. From these three figures it can be seen that the characteristics of the real data are similar to those of the simulated cases (e.g., realization 64). The p-values for this case are mid-level to high, as might be expected by how well the distribution from the real data matches an exponential one.

The amplitude time series for occultation 1 is shown in Fig. 26. It can be seen that there are significant fluctuations within the window. This window was chosen because of that fact – to see how well the theory applies to such a dynamic case. Figure 27 shows the spectra for the data (red), the spectra derived from the estimated parameters (green), and that using a set of fitted parameters (blue). The spectrum of the data shows features that were not seen in the simulated data. For example, the "hump" from around 3 to 6 Hz. In fact,
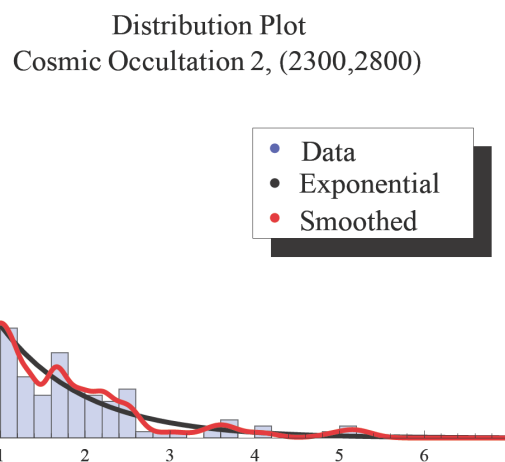

Fig. 25. Distribution of ratios for occultation 2.

if we look at the amplitude (red) and phase (black) time series over the analysis window (Fig. 28), a significant ramp can be seen in the amplitude data at around 54 s, with what appears to be a damped oscillation in the subsequent 1.5 s. The phase data shows a step-like structure during the same 1.5 s, with the steps in concert with those in the amplitude data. (It should be noted that the de-trending used with the phase data was very different than that used with the amplitude data. A polynomial fit over an extended time segment, centered at the analysis window was employed. For this example, the fit was performed over a window extending 150 points before and after the analysis window.) The source of these oscillations is unknown, but their periods are approximately 10–16 samples, which at 50 Hz. gives frequencies from around 3–5 Hz – which matches what is seen in the frequency spectrum. The hump in the spectrum generates magnified oscillations in the likelihood function, seen in the empirical (red) curve in Fig. 29, with the unfortunate by-product that one of these oscillations happens to be the global minimum, in turn skewing the parameter estimation (see the green curve in this figure). A by-eye fit was performed (blue curve) which provided far superior parameter estimates – or at least a better fit to the empirical likelihood function – since the true values are unknown. The estimated and fitted error statistics are $E_E = 299.2$ and $E_F = 2.4$, respectively. For the estimated parameters, $c = 2.76 \times 10^{-9}$, $\eta_1 = 17730$, $Q = 1.8 \times 10^{-5}$ and $L = 5$, and for the fitted parameters, $c = 1.31 \times 10^{-9}$, $\eta_1 = 20\,800$, $Q = 6.0 \times 10^{-6}$, and $L = 7.0$. It is encouraging that even with the data quality issues, physically reasonable parameter estimates can be obtained with fits to the log-likelihood function. The distribution of the ratios for this data also contains artifacts from the oscillations. In fact, these are directly related to the hump in the spectrum around 5 Hz, as those values – being much larger than the model spectrum – in turn produce large ratios. These large ratios in turn produce poor p-values.
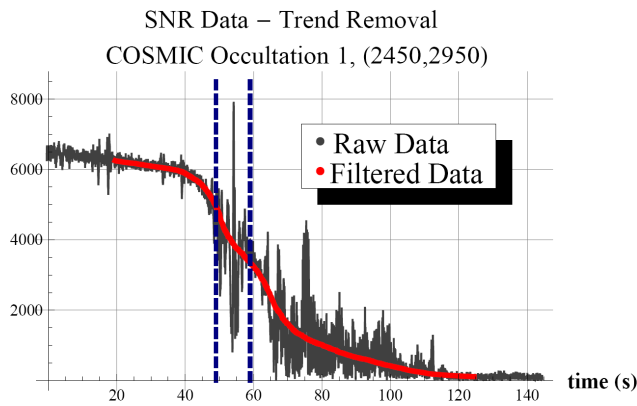
Fig. 26. SNR data for COSMIC occultation number 1 (black). Trend curve is in red and the vertical, dashed blue lines indicate the analysis window.
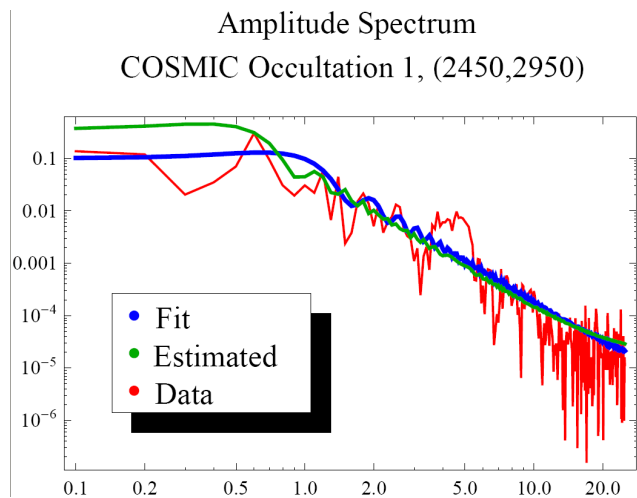

Fig. 28. De-trended amplitude and phase time series (in seconds) from the analysis window for occultation 1. (Between the vertical, dashed blue lines in Fig. 26.) Note spike just after 54 s, and oscillations ("ringing") for the subsequent 1.5 s.

Fig. 27. Amplitude frequency spectrum (in Hz) for occultation 1 (red), model spectrum using estimated parameters (green), and "by-eye" fitted parameters (blue).

Fig. 29. Log-likelihood functions, as a function of η_1 in km, from data (red), estimated parameters (green), and fitted parameters (blue) for occultation 2.

6 Conclusions

In this paper, an overview of the derivation used to produce the model spectrum was presented. This derivation started from the Helmholtz wave equation, and using a reasonable set of assumptions, resulted in the model spectrum, Eq. (25), and its mid-point approximation, Eq. (27). This latter model was used in the subsequent analysis. The assumptions used in the derivation were spelled out clearly. It was pointed out that, as we considered this a first-order solution, two real world atmospheric situations, inhomogeneity (including layered structures) and anisotropy, were not considered. The model derivation was based on standard weak scattering theory which can be found in many references. However, in our case we needed to modify the standard methods to include the motions of the transmitter, receiver, and atmosphere, as
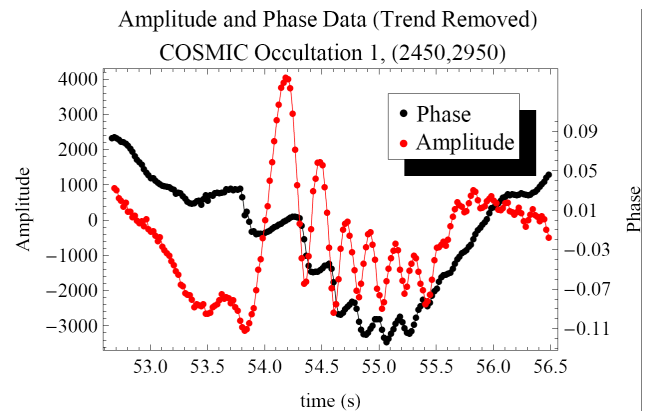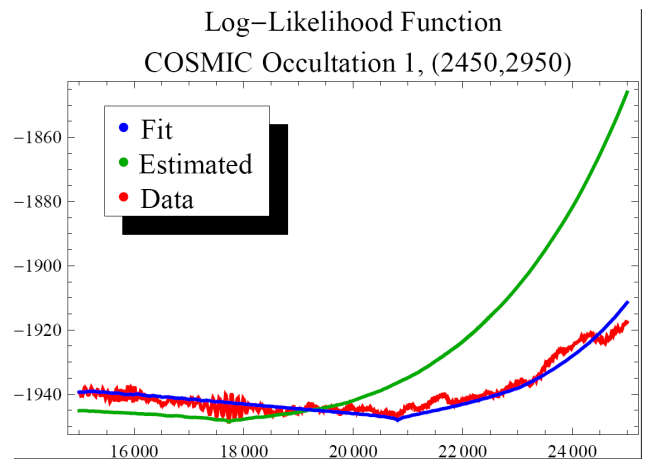
well as placing the origin of the coordinate system at an arbitrary point. It was shown that our results match those in the literature when making the appropriate simplifications.

A parameter estimation methodology was presented, followed by a discussion of the statistical properties of the estimators. A detailed study using simulated data was presented after the theoretical discussions. A number of metrics were examined to evaluate the accuracy of the estimates, e.g., median, standard deviation, and various inter-quantile ranges. It was found that the estimates of the turbulence intensity and the location of the turbulence along the line of sight could be retrieved in an accurate manner. In this instance, the term "accurate" refers to the global statistics over the 1000 simulated realizations. In a practical application, one would like to have accurate estimates for each measurement, or at least a notion as to when the estimates are unreliable.

Goodness-of-fit metrics (p-values) were used to see if the distribution of the ratios of the measured/simulated spectrum values to those from the true one could be used as an indicator of poor estimates. It was found that this technique was not reliable in that regard. Four cases were shown that demonstrated the problem, one where the p-values were high and the estimates were poor, and vice versa – as well as one where the p-value was high and the estimates were good, and vice versa. This was also borne out with a scatterplot of the parameter estimates and p-values. Even when choosing the maximum p-value over the different goodness-of-fit tests, the correlation of p-value and quality of the parameter estimate was reasonable – but not great.

Since the parameter estimation technique used the global minimum of the log-likelihood function, random errors in the spectrum could result in significant errors when trying to estimate the location of the turbulence along the line of sight. An important step in improving the parameter estimation was fitting the model-based likelihood function to the empirical one. In addition to incorporating the turbulence length scale parameter into the fitting, a weighting function was introduced into the likelihood estimation methodology. The weights were close to one for most of the η values deemed to be in the earth's atmosphere, and decreased rapidly outside those bounds. This not only allowed for the elimination of the most egregious outliers, but it also lead to improved parameter estimates from the data interior to the weighting function's effective region. Finally, it was shown that, due to reducing the random errors, spectral averaging greatly improved parameter estimation. These smaller errors then propagate through the parameter estimation resulting in more accurate estimates. It was seem that spectral averaging was far superior to the weighting function in improving the parameter estimates. The methods developed for the analysis of the simulated data were then applied to the analysis of real occultation data. In the two cases presented, it could be seen that the general characteristics seen in the simulated data were also seen in the real data – which is encouraging.

Future work includes the implementation of a three-parameter estimation algorithm including the length scale, and a more thorough analysis of real occultation data. This is especially needed for phase data. Some analysis has been performed with simulated and real phase data, with encouraging results, but the extensive analysis presented here with the amplitude data has not been performed. There are distinct issues regarding trend removal with the phase time series. A local polynomial fitting method was discussed briefly; however, the frequency content in the data resulting from polynomial fitting is highly dependent on the order of the polynomial used and the window over which the fit is performed. This indicates that a more robust methodology needs to be developed.

As mentioned above, future work should include the effects of inhomogeneities (including layered structures) and anisotropy. The reason that the inclusion of layered media is difficult is because large-scale layers will refract the incoming wave, which then means that the straight-line approximation for the propagation path needs to be modified. The permittivity field would be separated into background and random parts. (Note that in this context, "background" refers to all aspects of the deterministic permittivity field.) One could then re-write the original differential equation in a non-Cartesian coordinate system, e.g., a local one which has as its axes the tangent, normal and binormal vectors to the "main" propagation path – for example that described by geometrical optics. These are so-called trajectory coordinates (see for example: Hill, 1985; Mazur and Felsen, 1987). Another technique is to use a path integral approach to determine the Green's function in a multiple scattering context (Mazur, 2002). Yet another approach is to use the so-called Distorted-Wave Born (or Rytov) Approximations (see for example: Beylkin and Oristaglio, 1985; Devaney, 1979). In this method, one assumes that the permittivity field is separated into background and random parts, however the heuristic model is that the incoming waves are refracted by the background and are then scattered by the random field. The contribution of all such scattered fields is then what is measured at the receiver. That is, it is a single-scattering process, but the electric field at the scatterer is now "distorted" by the background field. Of the three methods, this is probably the most tractable one.

References

Beylkin, G. and Oristaglio, M .L.: Distorted-wave Born and distorted-wave Rytov approximations, Opt. Commun., 53, 213–216, 1985.

Clifford, S. A. and Lataitis, R. J.: Mutual coherence function for line-of-sight microwave propagation through a turbulent atmosphere, Radio Sci., 20, 221–227, 1985.

Clifford, S. A. and Strohbehn, J. W.: The theory of microwave line-of-sight propagation through a turbulent atmosphere, IEEE T. Antenn. Propag., AP-18, 264–274, 1970.

Conker, R. S., Bakry El-Arini, M., Hegarty, C. J., and Hsiao, T.: Modeling the effects of ionospheric scintillation of GPS/Satellite-based Augmentation System availability, Radio Sci., 38, 1001, doi:10.1029/2000RS002604, 2003.

Cornman, L. B., Frehlich, R., and Praskovskaya, E.: The detection of upper-level turbulence via GPS occultation methods, in: Occultations for Probing Atmosphere and Climate, edited by: Kirchengast, G., Springer-Verlag Berlin Heidelberg, 2004.

Cornman, L. B., Weekley, A., Goodrich, R. K., and Frehlich, E.: Using airborne GNSS receivers to detect atmospheric turbulence, in: New Horizons in Occultation Research, edited by: Steiner, A., Spring-Verlag Berlin Heidelberg, 2009.

Devaney, A. J.: The inverse problem for random sources, J. Math. Phys. 20, 1687–1691, 1979.

Freund, J. E.: Mathematical Statistics with Applications, 7 Edn. Prentice-Hall, New Jersey, 658 pp, 2004.

Ganguly, S., Jovancevic, A., Brown, A., Kirchner, M., Zigic, S., Beach, T., and Groves, K. M.: Ionospheric scintillation monitoring and mitigation using a software GPS receiver, Radio Sci., 39, RS1S21, doi:10.1029/2002RS002812, 2004.

Gorbunov, M. E. and Kirchengast, G.: Processing X/K band radio occultation data in the presence of turbulence, Radio Sci. 40, RS6001, doi:10.1029/2002RS002812, 2005.

Gurvich, A. S. and Brekhovskikh, V. L.: Study of the turbulence and inner waves in the stratosphere based on the observations of stellar scintillations from space: a model of scintillations spectra, Wave Random Media, 11, 163–181, 2001.

Gurvich, A. S. and Chunchuzov, I. P.: Parameters of the fine density structure in the stratosphere obtained from spacecraft observations of stellar scintillations, J. Geophys. Res., 108, 4166, doi:10.1029/2002JD002281, 2003.

Gurvich, A. S. and Kan, V.: Structure of air density irregularities in the stratosphere from spacecraft observations of stellar scintillations: 1. Three dimensional spectrum model and recovery of its parameters, Izv. Atmos. Ocean. Phy., 39, 300–310, 2003.

Hajj, G. A., Ibanez-Meier, R., Kursinski, E. R., and Romans, L. J.: Imaging the ionosphere with the Global Positioning System, Int. J. Imag. Syst. Tech., 5, 174–184, 2000.

Hill, R. J.: A stochastic parabolic wave equation and field-moment equations for random media having spatial variation of mean refractive index, J. Acoust. Soc. Am., 77, 1742–1753, 1985.

Ishimaru, A.: Wave propagation and scattering in random media, IEEE Press, New York, 574 pp., 1997.

Lee, R. W. and Harp, J. C.: Weak scattering in random media, with applications to remote probing, Proc. IEEE, 57, 375–406, 1969.

Lindgren, B.: Statistical Theory, 3rd Edn., Macmillian Publishing, New York, 1968.

Mazur, R.: Modeling of high-frequency propagation in inhomogeneous background random media, J. Acoust. Soc. Am., 111, 809–822 doi:10.1121/1.1430692, 2002.

Mazur, R. and Felsen, L. B.: High frequency coherence functions propagated along ray paths in the inhomogeneous background of a weakly random media: I – formulation and evaluation of the second moment, J. Acoust. Soc. Am., 81, 925–937, 1987.

Monin, A. S. and Yaglom, A. M.: Statistical Fluid Mechanics, Vol. II: Mechanics of Turbulence, MIT Press, Boston, 886 pp, 1975.

Secan, J. A., Bussey, R. M., Fremouw, E. J., and Basu, S.: High-latitude upgrade to the Wideband ionospheric scintillation model, Radio Sci., 32, 1567–1574, 1997.

Strohbehn, J. W.: Covariance functions and spectra for waves propagating in a turbulent medium to or from moving vehicles, IEEE T. Antenn. Propag., AP-22, 303–311, 1974.

Strohbehn, J. W. and Clifford, S. A.: Polarization and angle-of-arrival fluctuations for a plane wave propagated through a turbulent medium, IEEE T. Antenn. Propag., AP-15, 1967.

Tatarskii, V. I.: The Effects of the Turbulent Atmosphere on Wave Propagation, (translated from Russian by the Israel Program for Scientific Translations Ltd, ISBN 0 7065 0680 4), Reproduced by National Technical Information Service, US Department of Commerce, Springfield, Va. 22151, 1971.

Yeh, K. C. and Liu, C.-H.: Radio wave scintillations in the ionosphere, Proc. IEEE, 70, 324–360, 1982.