

Simulating Ionosphere-Induced Scintillation for Testing GPS Receiver Phase Tracking Loops

Todd E. Humphreys, Mark L. Psiaki, Joanna C. Hinks, Brady O'Hanlon and Paul M. Kintner, Jr.

Abstract—A simple model is proposed for simulating equatorial transionospheric radio wave scintillation. The model can be used to test Global Positioning System phase tracking loops for scintillation robustness because it captures the scintillation properties that affect such loops. In the model, scintillation amplitude is assumed to follow a Rice distribution, and the spectrum of the rapidly-varying component of complex scintillation is assumed to follow that of a low-pass 2nd-order Butterworth filter. These assumptions are justified, and the model validated, by comparison with phase-screen-generated and empirical scintillation data in realistic tracking loop tests. The model can be mechanized as a scintillation simulator that expects only two input parameters: the scintillation index S_4 and the decorrelation time τ_0 . Hardware-in-the-loop tests show how the model can be used to test the scintillation robustness of any compatible GPS receiver.

I. INTRODUCTION

Increased dependence on the Global Positioning System (GPS) and other satellite navigation systems makes users vulnerable to signal loss or degradation caused by ionospheric effects. Radio wave scintillation, the temporal fluctuation in phase and intensity caused by electron density irregularities along the propagation path, stresses a GPS receiver's carrier tracking loop, and, as severity increases, can lead to navigation bit errors, cycle slipping, and complete loss of carrier lock [1]–[9].

In anticipation of the 2011 solar maximum, when scintillation effects will be more severe, there is interest in testing civilian and military GPS receivers for scintillation robustness. Such testing entails subjecting a receiver's tracking loops to realistic phase and amplitude scintillation. This can be done by passing scintillation time histories through a software model of the tracking loops [1], [2], [4], [8], [10], [11]; or by forcing phase and amplitude variations in the output of a GPS signal simulator [9], [12]; or, in the ultimate confrontation with reality, by field testing receivers in a region prone to strong scintillation [5]. The first two of these testing strategies can give misleading results if the scintillation time histories are not realistic. For example, in field testing on Ascension Island during the 2000 solar maximum, researchers noted receiver performance degradations much worse than those predicted by simulations conducted prior to the campaign [5], [9].

Authors' addresses: T.E. Humphreys, M.L. Psiaki, and J.C. Hinks, Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca NY, 14853, Email: (teh25@cornell.edu); B. O'Hanlon and P.M. Kintner, Dept. of Electrical and Computer Engineering, Cornell University.

This work has been supported in part by the NASA Office of Space Science through grant No. NNX06AC34G. Madhulika Guhathakurta is the grant monitor.

This is a preprint of an article of the same title published in IEEE Journal of Selected Topics in Signal Processing, Vol. 3, No. 4, August 2009

In a previous paper [1], the current authors propose a receiver testing strategy that is based on drawing scintillation time histories from a large library of empirical equatorial scintillation data. The scintillation library includes severe complex signal scintillation from the Wideband Satellite experiment [13] and from specially-processed GPS data. The data reveal a universal feature of strong equatorial scintillation: deep power fades (> 15 dB) accompanied by abrupt, approximately half-cycle phase transitions. These "canonical fades" are shown in [1] to be the primary cause of loss of carrier lock in GPS phase tracking loops.

Although a receiver testing strategy based on empirical data is attractive for its realism, it nonetheless has several drawbacks: (1) A test engineer is only at liberty to adjust the scintillation behavior insofar as this behavior is represented in the recorded data; (2) thermal noise in the receiver that was originally used to record the data can leave high-frequency variations that make it difficult to precisely specify the carrier-to-noise ratio of a given test (such is the case in [1] for the scintillation library's GPS data); and (3) the empirical scintillation data is only stationary over short time intervals, making impossible extended testing under consistent scintillation statistics.

These limitations can be overcome by generating synthetic scintillation via computer simulation. Techniques for synthesizing scintillation include first-principles physics-based ionospheric models [14]; phase screen models [15]–[17]; and statistical models [8]–[10]. For testing carrier tracking loops, one seeks the simplest scintillation model—in terms of number of parameters and computational expense—that faithfully retains the scintillation properties that are relevant to carrier tracking. This goal favors statistical models over the computationally expensive and parameter-laden first-principles and phase screen models.

Because statistical models are abstractions of the physics that inspire them, extra effort must be made to ensure that their outputs are realistic. As noted in [1], the methods used in [8]—and likely in [9] and [10], though details are not provided—shape the phase and amplitude spectra independently. This practice tends to produce scintillation time histories that are artificially easy to track because they do not manifest realistic canonical fades. As demonstrated in the present paper, the key to synthesizing realistic scintillation is to focus on properly shaping the spectrum of the entire complex scintillation signal, not the amplitude and phase data taken independently. The proper spectral shape of the complex scintillation signal and the general structure of the scintillation model proposed in this paper are inspired by the model of equatorial scintillation

effects on GPS phase tracking loops developed in [2]. The result is a simple and computationally efficient technique for simulating realistic equatorial scintillation. (References [1], [2], and the current paper focus on equatorial scintillation because it is particularly difficult to track.)

The next section develops the scintillation model. Development is aided by analyses of empirical scintillation amplitude distributions and power spectra. Thereafter, in Section III, the model is validated by comparing its effect on phase tracking loops with the effects of phase-screen-generated and empirical scintillation. An example application of the scintillation model (hardware-in-the-loop testing) is described in Section IV. The proposed scintillation model is based on a connection between cycle slips and data bit errors. This connection is only relevant for squaring-type phase tracking loops; that is, for loops designed to track bi-phase modulated signals. Such loops are common in modern GPS receivers. The applicability of the model to non-squaring (full-cycle) tracking loops is discussed in Section V. The conclusions follow in Section VI.

II. SCINTILLATION MODEL

The scintillation model proposed here is premised on the notion, advanced in [2], that the scintillation properties that tend to induce cycle slips in squaring-type GPS phase tracking loop designs also tend to cause bit errors in differential detection of the 50 bps binary GPS navigation message that is phase modulated onto the L_1 carrier signal. In other words, if the noise and the scintillation-induced phase change between adjacent data bits (each 20 ms long) is so severe that one cannot decide correctly whether a bit sign change occurred, then a cycle slip is also likely. More precisely, it was shown in [2] that, over a broad range of operating conditions, the mean time between differentially-detected bit errors T_e acts as a lower bound to the mean time between cycle slips T_s for squaring-type phase tracking loops with good scintillation performance in the sense that on average $T_s/2 < T_e \leq T_s$. By accepting T_e as a rough proxy for T_s , one trades the difficult problem of cycle slip prediction for the more tractable problem of bit error prediction.

As outlined in [2], estimating T_e for a given interval of complex signal scintillation requires a model of the complex channel response function $z(t)$. One can think of $z(t)$ as the scintillation time history; its phase and magnitude are the phase and magnitude changes imposed on the GPS carrier signal by the scintillating communications channel.

The conjecture that underpins the proposed model can be summarized as follows: The close connection between T_e and T_s implies that, if a model for $z(t)$ accurately predicts T_e , then the same model can be used to synthesize scintillation because it necessarily captures the scintillation properties that cause cycle slipping, an important scintillation effect on GPS phase tracking loops. It will be shown in Section III that this conjecture is borne out by experiment.

Let $\alpha(t) \equiv |z(t)|$ and assume that the $z(t)$ is normalized so that $\Omega \equiv E[\alpha^2(t)] = 1$. Then, when no scintillation is present, $z(t) = 1$. An estimate of T_e depends on two properties of $z(t)$, namely, its amplitude distribution $p(\alpha)$ and a measure of the

rapidity with which $z(t)$ wanders about the complex plane [2]. This characterization becomes intuitive when one considers that the abrupt phase shifts in $z(t)$ that cause false sign changes between adjacent data bits tend to occur when $z(t)$ transits rapidly through a small neighborhood of the origin. Thus, if one knows the probability of being near the origin [given by $p(\alpha)$] and the average speed of $z(t)$, one can predict T_e .

In deriving T_e , $z(t)$ is typically assumed to have the form

$$z(t) = \bar{z} + \xi(t) \quad (1)$$

where \bar{z} is the direct component, modeled as a complex constant, and $\xi(t)$ is the time-varying multipath component (see, for example, [18] where the form is explicit and [19] and [20, Sec. 8.2.5.2] where it is implicit). The multipath component $\xi(t)$, also referred to as the fading process, has an autocorrelation function defined by $R_\xi(\tau) = \frac{1}{2}E[\xi^*(t)\xi(t+\tau)]$. The channel decorrelation time $\tau_0 > 0$ is defined as the value of τ for which $R_\xi(\tau)/R_\xi(0) = e^{-1}$. A narrow $R_\xi(\tau)$ (small τ_0) implies a scintillating channel that changes rapidly with time. Hence, the shape of $R_\xi(\tau)$ —most importantly, the width of its main peak—defines the measure of $z(t)$ rapidity mentioned above.

To further develop the scintillation model, the form of $p(\alpha)$ and of $R_\xi(\tau)$ must be specified. This is the subject of the next two subsections.

A. Amplitude Distribution

There does not exist at present a rigorous theory that predicts the probability distribution of $z(t)$. Only in certain limiting cases have sufficient moment calculations been carried out to completely characterize a first-order distribution [21], [22]. It can be shown, for example, that when scintillation severity enters the so-called saturation regime, the amplitude distribution $p(\alpha)$ approaches a Rayleigh distribution. In general, however, only numerical or empirical approaches have proven tractable for defining the distribution of $z(t)$.

In [22], several candidate distributions for $z(t)$ were evaluated against scintillation data from the Wideband Satellite experiment. The results of hypothesis test calculations identified the Nakagami-m distribution as the best fit to the empirical intensity (and amplitude) distributions. However, the analysis is not definitive for severe equatorial scintillation at frequencies near or above UHF because the data sets used were short (20 seconds) and few (20 sets at UHF and 9 sets at L-band). Furthermore, the exposition in [22] does not make clear how time correlation in the empirical $z(t)$ is handled in the hypothesis testing analysis.

Drawing on the extensive collection of scintillation data in the empirical scintillation library (introduced in [1]), a new hypothesis testing analysis has been carried out to determine an appropriate model for the amplitude distribution $p(\alpha)$. The new analysis focuses on strong equatorial scintillation at UHF and L-band frequencies, and limits the candidate models to the popular Nakagami-m and Nakagami-n (Rice) distributions,

given respectively by [20, Ch. 2]

$$p_m(\alpha) = \frac{2m^m \alpha^{2m-1}}{\Omega^m \Gamma(m)} e^{-m\alpha^2/\Omega} \quad (2)$$

$$p_n(\alpha) = \frac{2\alpha(1+K)}{\Omega} I_0 \left(2\alpha \sqrt{\frac{K+K^2}{\Omega}} \right) e^{-K-\alpha^2(1+K)/\Omega} \quad (3)$$

where $\alpha \geq 0$, $\Gamma(\cdot)$ is the gamma function, $K \geq 0$ is the Rician parameter, and $m = 1/S_4^2 \geq 1/2$. The quantity S_4 is the standard scintillation index defined by

$$S_4^2 = \frac{\langle I^2 \rangle - \langle I \rangle^2}{\langle I \rangle^2} \quad (4)$$

in which $I = \alpha^2$ is signal intensity and $\langle \cdot \rangle$ denotes time average. The Rician K parameter is related to m and S_4 by

$$K = \frac{\sqrt{m^2 - m}}{m - \sqrt{m^2 - m}} = \frac{\sqrt{1 - S_4^2}}{1 - \sqrt{1 - S_4^2}}, \quad S_4 \leq 1 \quad (5)$$

It should be noted that the Nakagami- m distribution is defined for $m \geq 1/2$, which corresponds to $S_4 = 1/\sqrt{m} \leq \sqrt{2}$. At $S_4 = m = 1$, the Nakagami- m and Rice distributions both converge to the Rayleigh distribution. For $S_4 > 1$, the Nakagami- m distribution is defined whereas the Rice distribution is not. This limitation of the Rice distribution is not restrictive in practice because, aside from rare ‘‘focusing’’ behavior, S_4 takes on values near or below unity [21]. At values of S_4 less than unity, the Nakagami- m and Rice distributions are similar, as illustrated in Fig. 1, where the two distributions are shown to agree closely with a histogram (thick solid line) of representative Wideband UHF data from the scintillation library.

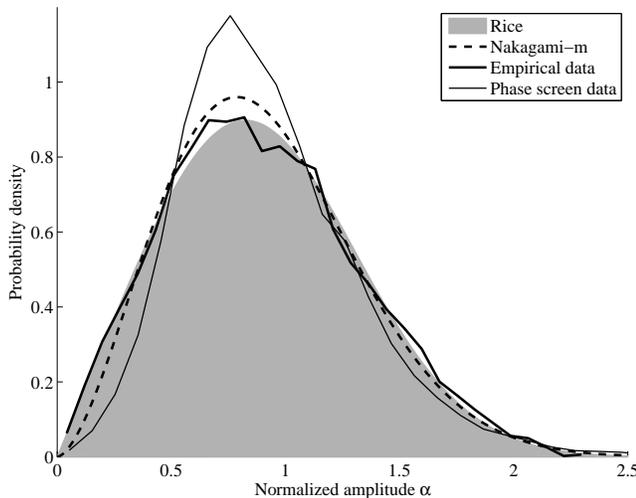


Fig. 1. Amplitude distribution of empirical scintillation data from the scintillation library (thick solid line) compared with the Nakagami- m and Rice distributions and with the amplitude distribution of scintillation data generated by a phase screen model. For all distributions $S_4 = 0.87$.

Also shown in Fig. 1 is the amplitude distribution of scintillation data that were generated by a phase screen scintillation model. In the phase screen approach to scintillation

modeling, the field fluctuations below the disturbed ionosphere are calculated as if produced by a thin phase-changing screen [21], [23], [24]. The model employed here is a modified version of the phase screen model presented in [15]. The modified version accommodates the development of amplitude fluctuations within the irregularity slab by replacing the single phase screen of [15] with a set of two phase screens separated by 100 km, with the uppermost screen located 350 km above the observation plane. The phase variations across each screen are Gaussian and have a power-law-type spatial power spectrum with an outer scale of 7 km and a spectral index of $\nu = -3.4$ (a rationale for this choice of ν will be given in Sec. II-B). A wavelength corresponding to the GPS L_1 frequency (1575.42 MHz) is assumed.

The phase screen scintillation model has been validated in several studies and has gained wide acceptance among theoreticians and experimentalists [24]. In this paper, it is exploited as an aid in studying the fading process spectrum (next subsection) and in validating the proposed scintillation model (Section III). However, for intermediate values of the perturbation strength of the screen (defined by the parameter C_s in [15]), that is, perturbation strength values which lead to $0.6 < S_4 < 1.0$, the phase screen model tends to produce scintillation time histories whose amplitude distributions depart markedly from empirical amplitude distributions, as illustrated in Fig. 1. For perturbation strength values well into the saturation regime ($S_4 \approx 1$) or in the weak scintillation regime ($S_4 < 0.4$), the phase-screen-generated and empirical distributions agree closely. This behavior reinforces the preference for a statistical scintillation model over the phase screen model for testing carrier tracking loops.

To evaluate the goodness-of-fit of the Nakagami- m and Rice distributions, average chi-square values were computed for 79 sets of Wideband UHF data and for 33 sets of GPS L_1 data drawn from the scintillation library. The data sets, whose lengths range from 50 to 300 seconds, correspond to relatively stationary intervals of strong scintillation ($S_4 > 0.6$). Intervals were chosen by inspection from the library records. Only a subset of the data samples from each interval was used in the chi-square calculations. Each sample in the subset is separated in time from adjacent samples by twice the decorrelation time τ_0 of the full interval. This ensures that data samples are independent from each other, as assumed by the chi-square technique. The length of each data set was chosen such that each chi-square test operated on approximately 100 samples. Thus, for slowly-varying scintillation (long τ_0), longer data sets were required to produce 100 independent samples. Partitioning of the samples led to 8 chi-square degrees of freedom (DOF) for the Wideband UHF data and 7 for the GPS L_1 data (recall that the number of chi-square DOF is equal to the number of bins used to partition the data in the calculation of the chi-square statistic). In the results of the chi-square tests, presented in Table I, the chi-square values are near the respective chi-square DOF for both the Nakagami- m and Rice distributions. This indicates that both the Nakagami- m and the Rice distributions provide a good fit to the data, though the Rice distribution retains a slight advantage for both data sources.

TABLE I
CHI-SQUARE VALUES FOR FITS TO NAKAGAMI-M AND RICE
DISTRIBUTIONS

Data Source	Sets, DOF	Nakagami-m	Rice
Wideband UHF	79, 8	11.8 ± 8.8	9.0 ± 4.3
GPS L ₁	33, 7	8.42 ± 5.9	7.7 ± 5.7

While the goodness-of-fit results in Table I justify either a Nakagami-m or a Rice model for the amplitude distribution, the Rice model is easier to implement in practice: one has only to assume that the fading process $\xi(t)$ in Eq. (1) is Gaussian, and a Rice distribution naturally results for the amplitude $\alpha(t) = |z(t)|$. Hence, in the proposed model, $\xi(t)$ is assumed to be a complex, zero-mean, stationary, Gaussian random processes with autocorrelation function $R_\xi(\tau)$. The Rician K parameter relates the magnitude of \bar{z} to $\sigma_\xi^2 \equiv R_\xi(0)$:

$$K = \frac{|\bar{z}|^2}{2\sigma_\xi^2} \quad (6)$$

B. Autocorrelation Function

To complete the proposed scintillation model, the form of the autocorrelation function $R_\xi(\tau)$ must be specified. Equivalently, one may specify $S_\xi(f)$, the power spectrum of the complex fading process $\xi(t)$, which is related to $R_\xi(\tau)$ by the Fourier transform. One would hope to bring the vast theory developed over many years in the scintillation literature to bear on the problem of modeling $S_\xi(f)$. But whereas the scintillation model in Eq. (1) is perfectly adequate for purposes of channel and receiver effects modeling, it is not valid in general for scintillation because it fails to capture the naturally-occurring low-frequency wander in the mean \bar{z} . Thus, the scintillation literature generally treats the complex scintillation signal $z(t)$ as a whole. (For example, the theoretical power spectra of the scintillation amplitude $\alpha(t) = |z(t)|$ and intensity $I(t) = \alpha^2(t)$ are, for weak scintillation, well understood [21].) Of course, for strong scintillation, the direct component \bar{z} is negligible and $\xi(t)$ becomes equivalent to $z(t)$. Unfortunately, in the case of strong scintillation the scintillation theory is on unsure footing [21]. Therefore, it appears most practical to adopt an experimental approach to specifying $S_\xi(f)$.

The following procedure has been used to determine the structure of $S_\xi(f)$ based on empirical and phase-screen-generated scintillation time histories. The phase of a complex scintillation time history $z(t)$ is high-pass filtered to remove its low-frequency components. A zero-phase filter is employed to preserve alignment of the phase and amplitude data. For convenience, let $z_f(t)$ represent the filtered version of $z(t)$. The filter cutoff frequency is adjusted so that the mean of $z_f(t)$, denoted \bar{z}_f , is near the value of \bar{z} predicted by the Rician K parameter, itself derived from the measured S_4 [cf. Eqs. (4)-(6)]. Further small adjustments to the filter cutoff frequency are made to roughly equalize the variances of the quadrature components of $\xi_f(t) = z_f(t) - \bar{z}_f$. Typical final cutoff frequencies range from 0.05 to 0.4 Hz. It should be emphasized that only the phase of $z(t)$ is altered by such

filtering; the amplitude time history—and the amplitude and intensity spectra—remain unchanged. Furthermore, the filtered and the original scintillation time histories are equivalent from the perspective of a carrier tracking loop. This is because the naturally-occurring variations in the mean \bar{z} are much slower than ordinary tracking loop response times. This has been demonstrated by passing both $z(t)$ and $z_f(t)$ through the software tracking loops introduced in [2] and noting that the loops perform equivalently (in terms of cycle slips and phase error variance) when operating on either time history.

Using the above filtering technique, empirical and phase-screen-generated scintillation time histories can be made to conform to the model in Eq. (1), wherein the direct component \bar{z} is constant. The form of $S_\xi(f)$ then becomes apparent in the power spectra of the filtered time histories. Fig. 2 plots $S_\xi(f)$ for scintillation generated by the phase screen model described previously. Cases of weak and strong scintillation are considered. For the weak scintillation case, a single phase screen with a low perturbation strength ($C_s = 10^{19}$ in the parameterization of [15]) was used. For the strong scintillation case, two phase screens, each with a high perturbation strength ($C_s = 10^{22}$), were used. The phase screen parameters were otherwise equivalent for the two cases.

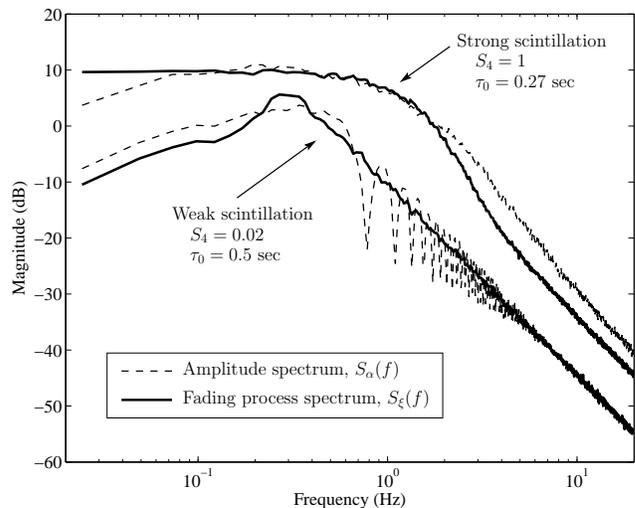


Fig. 2. Comparison of amplitude and fading process power spectra in the weak and strong scintillation regimes. All spectra are based on phase-screen-generated scintillation data. For visual clarity, the strong and weak scintillation spectra have been offset from each other by 10 dB.

For comparison with $S_\xi(f)$, Fig. 2 also plots the amplitude spectrum $S_\alpha(f)$ derived from the same data for both the weak and strong cases. The weak-scintillation amplitude spectrum agrees closely with the model from weak scintillation theory (given, for example, in [25]). At low frequencies, $S_\alpha(f)$ is relatively flat up to a break point known as the Fresnel frequency, f_F . The value of f_F increases with decreasing wavelength of the propagating radio wave, with decreasing height of the ionospheric irregularities, and with increasing velocity of the radio line-of-sight ionospheric pierce point relative to the irregularities. After breaking at f_F , the weak-scintillation amplitude spectrum manifests “Fresnel oscillations” and finally rolls off toward a high frequency asymptote

f^ν , where ν is the spectral index of the phase screen and is related to the spectral index p of the 3-dimensional irregularity medium by $\nu = 1 - p$. Empirical values for p are approximately 4.4 [21], which explains the choice of $\nu = -3.4$ for the phase screen spectral index. For strong scintillation (upper plots in Fig. 2), the amplitude spectrum also approaches f^ν at high frequencies, but the spectrum is wider and the Fresnel frequency is less distinct than for weak scintillation. This “spectral broadening” of $S_\alpha(f)$ with increasing scintillation strength is well documented in the literature [25].

One notes from Fig. 2 that, although $S_\xi(f)$ manifests no Fresnel oscillations (in either the weak or strong case), and, for strong scintillation, initially rolls off more steeply than $S_\alpha(f)$, it nonetheless generally conforms to the shape of $S_\alpha(f)$. This is consistent with the empirical findings presented in [26], where it was shown that during strong scintillation [for which $z(t) \approx \xi(t)$], amplitude fluctuations dominate the spectral shape of the quadrature components of $z(t)$.

The plots in Fig. 2 suggest that $S_\xi(f)$ can be modeled approximately as the frequency response of a low-pass filter with a 2nd-order rolloff. For simplicity, let $S_\xi(f)$ be modeled as the frequency response of a 2nd-order low-pass Butterworth filter. In this case, $R_\xi(\tau)$ is given by [18]

$$R_\xi(\tau) = \sigma_\xi^2 e^{(-\beta|\tau|/\tau_0)} [\cos(\beta\tau/\tau_0) + \sin(\beta|\tau|/\tau_0)] \quad (7)$$

where the factor $\beta = 1.2396464$ ensures that $R_\xi(\tau_0)/R_\xi(0) = e^{-1}$. Such a model falls within the two limiting forms for scintillation spectra—the so-called Gaussian and f^{-4} spectra—given in [27].

Panel (a) of Fig. 3 shows a representative empirical spectrum $S_\xi(f)$ that is based on data from the scintillation library ($S_4 = 0.87$), and compares this to the proposed 2nd-order Butterworth model and to the Gaussian and f^{-4} models. The corresponding normalized autocorrelation functions $R_\xi(\tau)/R_\xi(0)$ are given in panel (b). All models have been matched to the decorrelation time τ_0 of the empirical data.

The inset in panel (b) gives an expanded view of the first 40 ms of the autocorrelation plots. This critical window, equal in length to two GPS data bits, determines the degrading effect of bit-to-bit fluctuations in $z(t)$ on differential bit detection [2]. Together with the amplitude distribution $p(\alpha)$, the shape of $R_\xi(\tau)$ over the first 40 ms completely specifies the probability of error for GPS differential binary bit detection. Hence, it is this 40-ms window that must be accurately captured in the proposed model. In the frequency domain, this amounts to accurately modeling the shape of $S_\xi(f)$ in the neighborhood of the Fresnel frequency.

It can be shown that with a proper choice of τ_0 , any one of the models in Fig. 3 can be made to closely match the first 40 ms of the autocorrelation function of a given Wideband data record in the scintillation library. (The GPS data, which contain some high frequency noise that slightly distorts $R_\xi(\tau)$, were excluded from these tests.) On average, however, the 2nd-order Butterworth filter provides the best fit; moreover, it is the easiest to implement. This justifies its adoption in the proposed scintillation model.

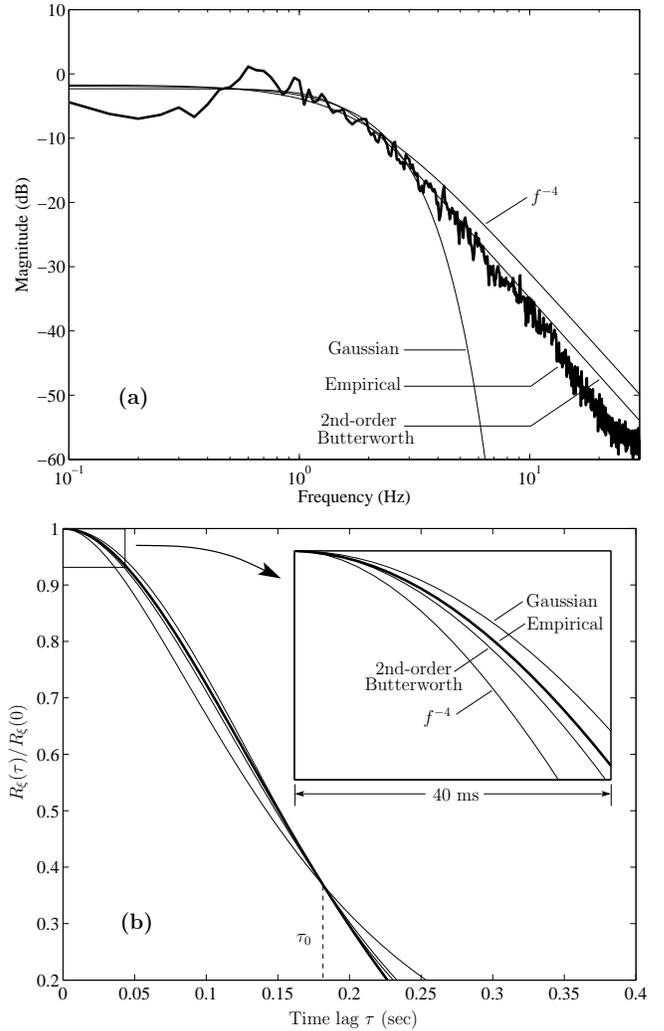


Fig. 3. (a) Empirical fading process spectrum (thick line) compared with the f^{-4} , Gaussian, and 2nd-order Butterworth spectral models. (b) Autocorrelation functions corresponding to the spectra in panel (a).

C. Scintillation Simulator Mechanization

The proposed scintillation model can be mechanized as shown in Fig. 4. In this form, the model becomes a scintillation simulator capable of generating realistic scintillation time histories $z(t)$. Although the model is presented here as a continuous-time system, it is easily recast as a discrete-time system for computer-based implementation. The model can be implemented as either a batch or sequential process. For simplicity, the discussion below assumes that the model’s input parameters S_4 and τ_0 are constant over the simulation interval. Besides being simple to describe and implement, this is also the most useful case since it leads to output scintillation data with stationary statistics. Tracking loop robustness to scintillation is best evaluated with long, statistically stationary scintillation time histories.

The scintillation simulator is driven by a stationary zero-mean complex white Gaussian noise process $n(t)$ with (two-sided) power spectral density $N_0/2$. The process $n(t)$ passes through a 2nd-order low-pass Butterworth filter with amplitude

response function

$$|H(f)| = \frac{1}{\sqrt{1 + \left(\frac{f}{f_n}\right)^4}} \quad (8)$$

where $f_n = \beta/(\sqrt{2\pi}\tau_0)$ is the filter corner frequency, with β as in Eq. (7) and τ_0 being the desired decorrelation time. Let the resulting zero-mean filtered noise process be denoted $\tilde{\xi}(t)$, with steady-state variance $\sigma_{\tilde{\xi}}^2 \approx f_n N_0$. The (constant) value of the direct component \tilde{z} is computed as

$$\tilde{z} = \sqrt{2\sigma_{\tilde{\xi}}^2 K} \quad (9)$$

where K is derived from the desired $S_4 \leq 1$ via Eq. (5). The direct component \tilde{z} is summed with $\xi(t)$, and the resulting process $\tilde{z}(t)$ is normalized by $\tilde{\alpha} = E[|\tilde{z}(t)|]$ to produce the synthetic scintillation time history $z(t) = \tilde{z} + \xi(t)$.

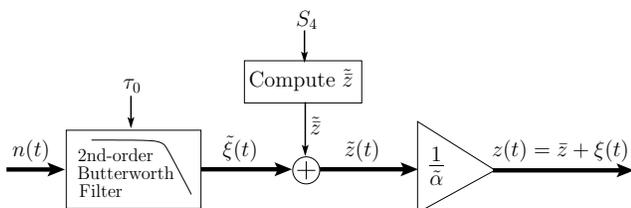


Fig. 4. Block diagram illustrating a straightforward mechanization of the proposed scintillation model. Thick lines denote complex signal routing.

The 2nd-order Butterworth filter in Fig. 4 is most conveniently implemented as a discrete-time filter, and the noise $n(t)$ as a sequence of independent samples from a random number generator. In such a discrete-time implementation, the statistics $\sigma_{\tilde{\xi}}^2$ and $\tilde{\alpha}$ can easily be calculated as a sample variance and a sample average, respectively, over a time interval $T \gg \tau_0$. To ensure proper steady-state statistics of the filtered noise process $\tilde{\xi}(t)$, one must either allow the filter to settle before using its output or implement an initialization procedure based on a discrete Lyapunov computation of the steady-state covariance matrix corresponding to the filter's two states.

The severity of the scintillation time history $z(t)$ is determined by the values of τ_0 and S_4 . In general, higher S_4 and lower τ_0 lead to more severe scintillation. For reference, the GPS data in the scintillation library—recorded using a stationary L₁ receiver—manifest $0 < S_4 \lesssim 1$ and $0.5 \leq \tau_0 < 2$ seconds [1]. Receivers mounted on dynamic platforms will see a broader range of τ_0 . Extreme scintillation in the library's Wideband UHF records manifests τ_0 values as low as 0.09 seconds.

A convenient lumped scintillation severity index can be obtained by calculating T_e as a function of S_4 , τ_0 , and nominal C/N_0 , as described in [2]. Then, when testing a typical GPS phase tracking loop using a simulated $z(t)$, a test engineer can expect the mean time between cycle slips T_s to be roughly equal to T_e .

III. MODEL VALIDATION

The following strategy was adopted for validating the proposed scintillation model. First, time histories of “truth”

scintillation were obtained. These were either generated using the previously described phase screen model (with two phase screens) or were drawn from Wideband UHF records in the scintillation library. In the latter case, relatively stationary scintillation intervals exceeding 150 seconds in length were chosen. The scintillation library's GPS data were excluded from the validation tests because of their thermal noise contamination. Nonetheless, as discussed in [1], there is no qualitative difference between the Wideband UHF and GPS scintillation data, apart from the noise contamination. In other words, the shape of the autocorrelation function $R_{\xi}(\tau)$ for given values of S_4 and τ_0 is the same for the Wideband UHF data as for the GPS data. Owing to the strengthening of scintillation effects with decreasing carrier frequency [21], UHF scintillation is, for the same line of sight, more severe (higher S_4 and lower τ_0) than L-band scintillation. Hence, the Wideband UHF data may be thought of as extreme L-band scintillation and remains useful for testing receiver tracking loops.

In the next step, the truth data were detrended as described in Section II-B, and the S_4 and τ_0 parameters for each truth set were estimated. The estimated value of τ_0 is the value that optimizes—in the least squares sense—the fit of the proposed autocorrelation model [Eq. (7)] to the first 40 ms of the empirical autocorrelation function.

Next, for each truth data set, 10 independent realizations of synthetic scintillation were generated, each equal in length to the truth set and matched to the S_4 and τ_0 parameters estimated from the truth set. The truth set and the 10 synthetic realizations were then passed as inputs to the scintillation testbed described in [1] and [2]. The dot-product four-quadrant arctangent, decision-directed four-quadrant arctangent, two-quadrant arctangent, decision-directed, and conventional Costas phase detectors were evaluated. A phase tracking loop noise bandwidth $B_n = 10$ Hz and an accumulation interval $T_a = 10$ ms were assumed. These values were found in [1] to lead to good tracking behavior. For all tests presented, the additive white noise that is typically used to simulate receiver thermal noise was eliminated, effectively making the tracked signals' carrier-to-noise ratios infinite. This was done so that the effects of scintillation could be studied in isolation from the effects of thermal noise.

The response of each phase tracking loop to the scintillation was noted in terms of the phase error variance σ_{φ}^2 and the number of cycle slips N_s over each test interval. Means and standard deviations of σ_{φ} and N_s were computed for the synthetic data from tests using the 10 independent realizations. Results are presented in Table II for two phase-screen-generated truth sets, and in Table III for eight empirical truth sets. The results presented are for the decision-directed four-quadrant arctangent phase detector, but are representative of the results for all the phase detectors tested. In the tables, T denotes the length of the test in seconds, and the spread about the means of σ_{φ} and N_s for the synthetic scintillation represents one standard deviation.

In general, the tracking loops' responses to the truth and to the synthetic scintillation are very similar. Two exceptions are the second phase-screen-generated test, whose poor

TABLE II
SCINTILLATION EFFECTS COMPARISON: PHASE SCREEN TRUTH DATA

Parameters			Truth Scint.		Synthetic Scint.	
S_4	τ_0 (s)	T (s)	σ_φ (deg)	N_s	σ_φ (deg)	N_s
1.0	0.28	328	16.3	60	16.8 ± 0.2	70.7 ± 7.3
0.79	0.33	328	5.2	3	9.2 ± 0.3	12.1 ± 3.3

TABLE III
SCINTILLATION EFFECTS COMPARISON: EMPIRICAL TRUTH DATA

Parameters			Truth Scint.		Synthetic Scint.	
S_4	τ_0 (s)	T (s)	σ_φ (deg)	N_s	σ_φ (deg)	N_s
0.87	0.18	200	16.4	32	17.5 ± 0.5	35.9 ± 4.7
1.0	0.36	265	14.1	37	15.0 ± 0.5	41.6 ± 5.9
0.69	0.18	174	12.7	12	11.8 ± 0.9	5.6 ± 1.6
0.87	0.26	225	11.6	23	12.7 ± 0.5	19.2 ± 4.6
0.61	0.47	162	3.96	0	3.63 ± 0.2	0.10 ± 0.3
0.96	0.09	81	28.5	60	32.7 ± 1.0	69.4 ± 5.8
0.95	0.26	123	14.1	21	15.6 ± 0.5	19.8 ± 3.7
0.51	0.71	138	2.12	0	1.60 ± 0.1	0 ± 0

correspondence can be attributed to the distortion of $p(\alpha)$ mentioned in Section II-A, and the cycle slip count in the third empirical test. In all other cases (besides the $N_s = 0$ cases for which a percentage difference is not defined but the correspondence is obviously good), the σ_φ and N_s values for the synthetic scintillation are within 25% of those for the true scintillation. On the whole, these results suggest that, for testing the response of GPS phase tracking loops to scintillation, the proposed scintillation model is useful and the scintillation it produces is realistic.

IV. EXAMPLE APPLICATION

A hardware-in-the-loop test was conducted to demonstrate how the proposed scintillation model can be used to test GPS receivers for scintillation robustness. For this test, a Matlab implementation of the scintillation simulator presented in Section II-C was used to generate synthetic scintillation under two scenarios: (1) $S_4 = 0.8$ and $\tau_0 = 0.8$ seconds, and (2) $S_4 = 1.0$ and $\tau_0 = 0.5$ seconds. These scenarios represent, respectively, moderately strong (more common) and severe (less common) L-band equatorial scintillation such as occurs after sunset during solar maximum in the neighborhood of the geomagnetic equator, as evidenced by a survey of GPS data in the empirical scintillation library introduced in [1].

Synthetic phase and amplitude time histories for each scenario were averaged over 10-ms intervals, formatted as phase and amplitude offsets in a time-stamped User Actions File, and loaded into a Spirent GSS7700 GPS signal simulator. The Spirent simulator was configured to generate a GPS signal from only one GPS satellite, which, for convenience, was set to behave as if it were geostationary so that phase variations due to satellite motion would be eliminated. Once armed, the Spirent simulator brought in data from the User Actions File to drive 100-Hz phase and amplitude variations in its RF output. Each scenario lasted 655 seconds.

Two GPS receivers, the Cornell GNSS Receiver Implementation on a DSP (Cornell GRID receiver) [28], and a com-

mercial scintillation monitor, were connected to the Spirent simulator and tested for scintillation robustness under this setup. Each receiver is equipped with a stable oven-controlled crystal oscillator whose phase jitter is negligible compared to the phase variations caused by the synthetic scintillation. Due to differences in the RF front-ends of the two receivers, the nominal C/N_0 value was higher for the GRID receiver (51.8 dB-Hz) than for the commercial receiver (43.8 dB-Hz). Testing was only conducted at the GPS L_1 frequency.

Figure 5 shows the effects of the first 300 seconds of Scenario 2 on the phase tracking behavior of the two receivers. The jagged phase output of the commercial receiver (top panel) reflects its tracking strategy of re-initializing its carrier phase estimate whenever a phase anomaly (a phase error that exceeds a predefined value) is detected. Not every phase anomaly would have resulted in a cycle slip; hence, by dividing the test interval by the number of times the phase estimate was re-initialized, one arrives at a pessimistic estimate of the mean time between cycle slips T_s over the interval. The bottom panel of Fig. 5 shows the difference between the “truth” phase generated by the scintillation simulator and the phase measured by the Cornell GRID receiver (the plot has been detrended to remove the effects of reference oscillator offset). Several half-cycle slips are evident in the plot. Each slip can be associated with a canonical fade in the synthetic scintillation time history.

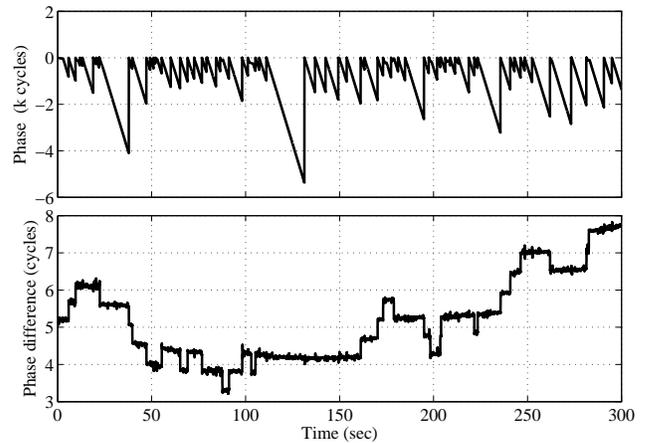


Fig. 5. First 300 seconds of the phase time histories resulting from hardware-in-the-loop scintillation simulation under Scenario 2. Top panel: Raw phase output of a commercial scintillation monitor with nominal $C/N_0 = 43.8$ dB-Hz. Bottom panel: Difference between the “truth” phase generated by the scintillation simulator and the phase measured by the Cornell GRID receiver with nominal $C/N_0 = 51.8$ dB-Hz.

TABLE IV
HARDWARE-IN-THE-LOOP TEST RESULTS

Quantity	Scenario 1		Scenario 2	
	GRID	GSV4004B	GRID	GSV4004B
S_4	0.8	0.8	1.0	1.0
τ_0 (s)	0.8	0.8	0.5	0.5
C/N_0 (dB-Hz)	51.8	43.8	51.8	43.8
T_e (predicted, s)	109	30	15	9
T_s (actual, s)	164	26	10	5

Table IV summarizes the results of the hardware-in-the-loop tests. Here, T_e refers to the mean time between differentially-detected navigation bit errors that would be expected to occur during each scenario, based on S_4 , τ_0 , and the nominal C/N_0 . As discussed before, T_e can be thought of as a rough predictor of T_s for properly configured squaring-type phaselock loops. The GRID results in Table IV confirm this connection between T_e and T_s . Also, the GRID T_s values closely match the mean time between cycle slips detected in off-line tests in which the scintillation testbed introduced in [1] was applied to a phaselock loop identical to that in the GRID receiver and was driven by the original Scenario-1 and Scenario-2 synthetic scintillation time histories (mean T_s values over 10 runs were 139 ± 26 seconds for Scenario 1 and 10 ± 0.6 seconds for Scenario 2). One can conclude from this that the hardware-in-the-loop testing technique does not introduce any significant effects not already modeled in the scintillation testbed.

V. APPLICABILITY OF THE SCINTILLATION MODEL TO NON-SQUARING PHASE TRACKING LOOPS

The scintillation model proposed here is premised on a connection between cycle slips and errors in differential detection of the 50 bps binary GPS message that is modulated onto the L_1 C/A, L_1 P, and L_2 P signals. However, the new civilian GPS signals at L_1 , L_2 , and L_5 include “pilot” signals that are free of navigation data bit modulation. Inclusion of these pilot signals aids acquisition by allowing coherent integration beyond 20 ms and aids tracking by allowing use of non-squaring (full-cycle) phase tracking loops, which, as noted in [2], extends the mean time between cycle slips during severe scintillation by roughly factor of 20. One might naturally inquire whether the scintillation model proposed here is applicable for testing non-squaring loops.

Validation tests like those presented in Section III have been conducted to determine whether synthetic scintillation data generated with the proposed scintillation model are realistic for non-squaring tracking loops. Again, a phase tracking loop bandwidth $B_n = 10$ Hz and an accumulation interval $T_a = 10$ ms were assumed, and “truth” empirical scintillation data were taken from Wideband UHF records in the scintillation library. Just as in Section III, it was found that when the synthetic scintillation data are tuned to have S_4 and τ_0 values close to those of a statistically stationary interval of “truth” scintillation, the simulated and “truth” data lead to similar phase error variance σ_φ^2 and mean time between cycle slips T_s .

It is conjectured that the scintillation model applies equally well to squaring and non-squaring GPS tracking loops because the model matches the autocorrelation function $R_\xi(\tau)$ of empirical scintillation data over lag (τ) values relevant to tracking with typical GPS loop bandwidths ($B_n \approx 10$ Hz). In other words, matching the shape of $R_\xi(\tau)$ over $0 \leq \tau \leq 40$ ms not only leads to a model that correctly reproduces the bit error rate for differential binary bit detection, but also leads to a fading process spectrum whose shape is realistic in the neighborhood of a 10-Hz loop’s corner frequency and above. This latter condition, along with the scintillation

model’s other properties, seems to ensure that the synthetic scintillation appears realistic even to non-squaring tracking loops. Longer matching intervals would be appropriate for testing GPS receivers with B_n significantly less than 10 Hz.

It should be noted that to properly evaluate the scintillation robustness of a multi-frequency GPS system requires processing multiple scintillation time histories—one for each signal frequency. The multiple time histories should reflect the coherence bandwidth of the disturbed ionosphere. In particular, the pairwise complex cross-correlation coefficients of the scintillation time histories should fall off with frequency separation and with increased perturbation strength as established in [29]. A scintillation model capable of generating realistic multi-frequency scintillation will be the subject of future work.

VI. CONCLUSIONS

A scintillation model has been proposed that exploits the close connection between differentially-detected data bit errors and cycle slips in squaring-type GPS phase tracking loops. The connection implies that, from the perspective of a tracking loop trying to maintain lock during scintillation, the critical scintillation properties are the amplitude distribution and the autocorrelation function of the scintillation’s rapidly-varying complex component. Studies presented here using phase-screen-generated and empirical scintillation data show that the amplitude distribution is well-modeled as a Rice distribution and the the autocorrelation function is well-modeled as corresponding to the frequency response of a low-pass 2nd-order Butterworth filter. Such models can be easily combined and mechanized as a scintillation simulator that takes as inputs the scintillation index S_4 and the decorrelation time τ_0 , and outputs realistic scintillation. The mean time T_e between errors in differentially-detected data bits can be calculated from the simulator inputs and the nominal C/N_0 , and serves as a convenient lumped index of scintillation severity.

The proposed scintillation model has been validated in comparison tests with phase-screen-generated and empirical scintillation data. It has also been demonstrated how the model can be used to evaluate the scintillation robustness of any compatible GPS receiver in hardware-in-the-loop testing. Even though the model is motivated by a connection between data bit errors and cycle slips in squaring-type phase tracking loops, it produces scintillation that is also realistic for non-squaring (full-cycle) tracking loops such as those appropriate for tracking the new data-free civilian GPS signals.

REFERENCES

- [1] T. E. Humphreys, M. L. Psiaki, B. M. Ledvina, A. P. Cerruti, and P. M. Kintner, Jr., “A data-driven testbed for evaluating GPS carrier tracking loops in ionospheric scintillation,” *IEEE Transactions on Aerospace and Electronic Systems*, 2009, to be published.
- [2] T. E. Humphreys, M. L. Psiaki, and P. M. Kintner, Jr., “Modeling the effects of ionospheric scintillation on GPS carrier phase tracking,” *IEEE Transactions on Aerospace and Electronic Systems*, 2009, to be published.
- [3] J. A. Klobuchar, *Global Positioning System: Theory and Applications*. Washington, DC: American Institute of Aeronautics and Astronautics, 1996, ch. 12: Ionospheric Effects on GPS, pp. 485–515.

- [4] T. E. Humphreys, M. L. Psiaki, B. M. Ledvina, and P. M. Kintner, Jr., "GPS carrier tracking loop performance in the presence of ionospheric scintillations," in *Proceedings of ION GNSS 2005*. Long Beach, CA: Institute of Navigation, Sept. 2005.
- [5] K. M. Groves, S. Basu, J. M. Quinn, T. R. Pedersen, K. Falinski, A. Brown, R. Silva, and P. Ning, "A comparison of GPS performance in a scintillating environment at Ascension Island," in *Proceedings of ION GPS 2000*. Institute of Navigation, 2000.
- [6] M. Knight and A. Finn, "The effects of ionospheric scintillation on GPS," in *Proceedings of ION GPS 1998*. Nashville, TN: Institute of Navigation, 1998.
- [7] R. S. Conker, M. B. El-Arini, C. J. Hegarty, and T. Hsiao, "Modeling the effects of ionospheric scintillation on GPS/Satellite-Based Augmentation System availability," *Radio Science*, vol. 38, Jan. 2003.
- [8] C. Hegarty, M. B. El-Arini, T. Kim, and S. Ericson, "Scintillation modeling for GPS-wide area augmentation system receivers," *Radio Science*, vol. 36, no. 5, pp. 1221–1231, Sept.–Oct. 2001.
- [9] G. Bishop, D. Howell, C. Coker, A. Mazzella, D. Jacobs, E. Fremouw, J. Secan, B. Rahn, C. Curtis, J. Quinn, K. Groves, S. Basu, and M. Smitham, "Test bed for evaluation of GPS receivers' performance in ionospheric scintillation—a progress report," in *Proceedings of ION GPS 1998*. Long Beach, CA: Institute of Navigation, 1998.
- [10] M. A. Cervera and M. F. Knight, "Time series modelling of intensity and phase scintillation at GPS frequencies," *Acta Geodaetica et Geophysica Hungarica*, vol. 33, no. 1, pp. 25–40, 1998.
- [11] T. E. Humphreys, B. M. Ledvina, M. L. Psiaki, and P. M. Kintner, "Analysis of ionospheric scintillations using wideband GPS L1 C/A signal data," in *Proc. ION GNSS 2004*. Long Beach, California: Institute of Navigation, 2004, pp. 399–407.
- [12] T. N. Morrissey, K. W. Shallberg, A. J. Van Dierendonck, and M. J. Nicholson, "GPS receiver performance characterization under realistic ionospheric phase scintillation environments," *Radio Sci.*, vol. 39, pp. 1–18, 2004.
- [13] E. J. Fremouw, R. L. Leadabrand, R. C. Livingston, M. D. Cousins, C. L. Rino, B. C. Fair, and R. A. Long, "Early results from the DNA Wideband Satellite experiment - Complex-signal scintillation," *Radio Science*, vol. 13, pp. 167–187, Feb. 1978.
- [14] M. J. Keskinen, "GPS scintillation channel model for the disturbed low-latitude ionosphere," *Radio Science*, vol. 41, pp. 4003–+, July 2006.
- [15] C. L. Rino and J. Owen, "Numerical simulations of intensity scintillation using the power law phase screen model," *Radio Science*, vol. 19, pp. 891–908, June 1984.
- [16] A. Pidwerbetsky, "Simulation and analysis of wave propagation through random media," Ph.D. dissertation, Cornell University, 1988.
- [17] T. L. Beach, "Global positioning system studies of equatorial scintillations," Ph.D. dissertation, Cornell University, Ithaca, New York, 1998.
- [18] L. J. Mason, "Error probability evaluation for systems employing differential detection in a Rician fast fading environment and Gaussian noise," *IEEE Transactions on Communications*, vol. COM-35, no. 1, pp. 39–46, Jan. 1987.
- [19] P. A. Bello and B. D. Nelin, "The influence of fading spectrum on the binary error probabilities of incoherent and differentially coherent matched filter receivers," *IRE Transactions on Communication Systems*, vol. CS-10, pp. 160–168, June 1962.
- [20] M. K. Simon and M. Alouini, *Digital Communications over Fading Channels*. New York: Wiley, 2000.
- [21] K. C. Yeh and C. H. Liu, "Radio wave scintillations in the ionosphere," *Proceedings of the IEEE*, vol. 70, no. 4, pp. 324–360, 1982.
- [22] E. J. Fremouw, R. C. Livingston, and D. A. Miller, "On the statistics of scintillating signals," *Journal of Atmospheric and Terrestrial Physics*, vol. 42, pp. 717–731, Aug. 1980.
- [23] E. N. Bramley, "The accuracy of computing ionospheric radio-wave scintillation by the thin-phase-screen approximation," *Journal of Atmospheric and Terrestrial Physics*, vol. 39, pp. 367–373, 1977.
- [24] C. L. Rino, "On the application of phase screen models to the interpretation of ionospheric scintillation data," *Radio Science*, vol. 17, no. 4, pp. 855–867, July–Aug. 1982.
- [25] R. Umeki, C. H. Liu, and K. C. Yeh, "Multifrequency spectra of ionospheric amplitude scintillations," *Journal of Geophysical Research*, vol. 82, pp. 2752–2760, July 1977.
- [26] W. J. Myers, R. J. Gjeldum, C. H. Liu, and K. C. Yeh, "A study of ionospheric scintillations of phase and quadrature components," *Journal of Geophysical Research*, vol. 84, pp. 2039–2048, May 1979.
- [27] R. A. Dana, "Effects of ionospheric scintillation on differential demodulation of GPS data," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 33, no. 3, pp. 893–902, July 1997.
- [28] T. E. Humphreys, B. M. Ledvina, M. L. Psiaki, and P. M. Kintner, Jr., "GNSS receiver implementation on a DSP: Status, challenges, and prospects," in *Proceedings of ION GNSS 2006*. Fort Worth, TX: Institute of Navigation, 2006.
- [29] C. L. Rino, V. H. Gonzalez, and A. R. Hessing, "Coherence bandwidth loss in transionospheric radio propagation," *Radio Science*, vol. 16, pp. 245–255, 1981.