A data-driven testbed for evaluating GPS carrier tracking loops in ionospheric scintillation

Todd E. Humphreys, Mark L. Psiaki, Brent M. Ledvina, Alessandro P. Cerruti, and Paul M. Kintner, Jr.

***Abstract*—A large set of equatorial ionospheric scintillation data has been compiled, used to characterize features of severe scintillation that impact Global Positioning System phase tracking, and used to develop a scintillation testbed for evaluating tracking loops. The data-driven testbed provides researchers a tool for studying, and receiver developers a tool for testing, the behavior of carrier tracking loops under realistic scintillation conditions. It is known that severe equatorial scintillation causes cycle slipping and, in the worst cases, complete loss of carrier lock. Testbed results indicate that cycle slips are primarily caused by the abrupt, near half-cycle phase changes that occur during the deep power fades of severe equatorial scintillation. For a class of standard tracking loops, parameters values that minimize scintillation-induced cycle slipping are identified.**

I. Introduction

A transionospheric radio wave can exhibit temporal fluctuations in phase and intensity caused by electron density irregularities along its propagation path. This phenomenon, called scintillation or fading, is a subject of both scientific interest and practical concern [1]. At Global Positioning System (GPS) frequencies (L band), strong scintillation is manifest in deep power fades (> 15 dB) that are often associated with rapid phase changes. Such vigorous signal dynamics stress a receiver's carrier tracking loop and, as their severity increases, lead to navigation bit errors, cycle slipping, and complete loss of carrier lock [2]–[8].

Increasingly, modern GPS applications are exploiting techniques that require carrier phase measurements. In view of this trend, and of their already widespread use, attention is focused here on phase tracking loops, or phaselock loops (PLLs), as opposed to frequency tracking loops. In particular, attention is focused on squaring-type PLLs designed to track a bi-phase modulated carrier.

Severe L-band scintillation is both infrequent and geographically confined. The type known as equatorial scintillation, or equatorial spread F, generally occurs between local sunset and 2400 local time in the region extending $\pm 15^\circ$ about the magnetic equator [9]. Another common type of scintillation occurs at high latitudes [10]. Significant effects have also been noted in the mid-latitude region, but they occur infrequently [11]. This paper concentrates on equatorial scintillation because it is particularly difficult to track.

Authors' addresses: T.E. Humphreys and M.L. Psiaki, Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca NY, 14853. Email:(teh25@cornell.edu); P.M. Kintner and A.P. Cerruti, Dept. of Electrical and Computer Engineering, Cornell University; B.M. Ledvina, Dept. of Electrical and Computer Engineering, Virginia Tech., Blacksburg VA, 24060.

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

Despite their relative rarity, scintillation effects are a serious concern for certain GPS user communities whose applications place demanding requirements on receiver performance. Consider three examples:

1) Proposals for GPS-based ionospheric research call for large arrays of GPS receivers whose combined phase and amplitude measurements will be used to image fine-scale ionospheric structures during scintillation events. Such studies would be compromised by receivers that repeatedly lose phase lock.
2) Carrier-phase-based GPS techniques that are capable of subdecimeter-level relative positioning accuracy are finding increased use in civilian and military applications [12],[13, Ch. 7]. These techniques rely on accurate resolution of carrier phase ambiguities, a process that may take several minutes. Severe scintillation poses a threat to such carrier-based techniques because, among other things, repeated cycle slipping on the same GPS link, or slipping on more than one link, causes new ambiguities in the carrier phases that must be resolved before the system can recover operational accuracy.
3) The aviation community is anticipating an overhaul of the air traffic control system that would see GPS replace radar as the primary aircraft location sensor [14]. There is concern that severe scintillation effects may prevent the system from meeting its exacting integrity requirements.

As these examples show, there is interest in developing GPS receivers whose tracking loops are specially designed to maintain lock in the presence of severe scintillation. Design and testing of such receivers depends crucially on an accurate understanding of scintillation and its effects on receivers. Several researchers have responded to this need by developing models for scintillation effects on PLLs [3]–[6], [8], [15], but these models tend to underestimate the effects of severe equatorial scintillation by failing to capture its essential feature: large, abrupt phase changes associated with deep power fades.

To put this paper's contributions in context, the following list outlines current approaches to characterizing scintillation effects on GPS receivers. The strengths and limitations of each approach are noted.

Statistical models [4], [5]: In this approach, scintillation intensity, as measured by the S_4 index and the phase spectrum profile, is used to estimate the PLL's phase error variance, σ_φ^2. A hypothetical receiver is said to lose lock when σ_φ^2 exceeds a certain threshold. This approach's generality and simplicity are appealing, but, as implemented in [4], [5], it tends to

underestimate the rate of cycle slipping in actual scintillation. First, it assumes the PLL is operating in its linear regime, which, in the case of strong scintillation, is likely not the case. Second, the value of the lock threshold is chosen based on a comparison with models whose phase noise is uncorrelated in time. Such models are inapt for scintillation, which is strongly time correlated.

Simulation using synthetic phase and amplitude time histories [6], [15]: In this approach, synthetic scintillating phase and amplitude time histories are generated in such a way that their power spectra and probability distributions match well-accepted models [16]–[18]. The phase and amplitude time histories are then fed as inputs to a simulated PLL whose phase error variance σ_φ^2 and rate of cycle slipping are observed. One advantage of this approach is that it requires no assumption that the loop operates in its linear regime. Hence, the cycle slipping and phase error are a true reflection of the simulated loop's response to the synthetic inputs. This approach must be adopted with some care, however, because—as will be discussed in Section III-C—correct phase power spectra and amplitude distributions do not necessarily lead to realistic scintillation.

Simulation using recorded wide-bandwidth scintillation data [3], [19]: In this approach, a non-real-time software GPS receiver processes raw wide-bandwidth digital data that were previously recorded on a digital storage device during a scintillation event. The behavior of the software receiver's PLL is observed as it tracks through the fluctuations in the recorded data. Since no simplifying models or assumptions are introduced, this approach is a very realistic test of receiver performance. It is flexible in the sense that a designer may evaluate different tracking loop strategies and may add noise to the recorded data bitstream to simulate a reduced signal-to noise-ratio. However, the method is computationally expensive—correlations must be computed from the raw data on each run. Furthermore, the "true" phase time history is unknown, so phase error and cycle slips are not defined, though complete loss of carrier lock is obvious.

Hardware-in-the-loop simulation [8], [20]: In this method, empirical [20] or synthetic [8] phase time histories are used to drive phase variations in the RF output of a GPS signal simulator. Likewise, amplitude fades are simulated by modulating the simulator's output signal strength manually [20] or according to a synthetic amplitude time history [8]. A GPS receiver to be tested is connected to the simulator and the behavior of the receiver's tracking loops is observed. To the extent that the input phase and amplitude time histories are realistic and loss of lock is reliably detected, this method is an effective end-to-end test of a full receiver. However, as currently implemented in [20], the "true" phase data is 1st-order low pass filtered at 2 Hz. As will be discussed in Section III-C, this practice artificially softens the effect of severe phase scintillation on the receiver's PLL. The same section will note that the synthetic phase time history in [8] is likely unrealistic for severe equatorial scintillation.

Field testing [7]: Field testing is of course the ultimate confrontation with reality both for receivers and for the scintillation models used to validate them. Thus, the results of a campaign on Ascension Island during the last solar maximum exposed worrisome vulnerabilities in the receivers tested and raised questions about the overly optimistic simulations conducted prior to the campaign [7]. Nonetheless, bench testing is preferred to field testing because it can be repeated under controlled conditions. Repeatable tests are helpful in gaining the insight necessary to guide receiver design.

This paper presents a data-driven simulation testbed for evaluating carrier phase tracking loops that is meant to overcome the foregoing methods' limitations. This tool will be referred to as the scintillation testbed or simply the testbed throughout the remainder of this paper. The testbed is realistic and flexible, and it reveals an intuitive connection between cycle slips and the scintillation patterns that cause them. It is based on a library of empirical phase and amplitude time histories recorded under a wide range of equatorial scintillation conditions, and represents an extension of the work in [3] and [19].

As a preamble to the presentation of the scintillation testbed, Section II introduces the empirical scintillation library and Section III presents observations about general characteristics of severe scintillation. Particular attention will be focused on an especially troublesome fading phenomenon herein termed "canonical fading." The testbed is presented in Section IV. Section V illustrates the effects of canonical fading and identifies loop parameters that maximize scintillation robustness. The conclusions follow in Section VI.

II. The Empirical Scintillation Library

A library of empirical equatorial phase and amplitude scintillation phenomena has been compiled for use in the scintillation testbed. This section describes the data on which the library is based and briefly summarizes the processing required to generate the library's individual records.

To cover a broad range of severe scintillation effects, the empirical scintillation library includes equatorial scintillation data from two different sources. The first source is a set of data from the DNA Wideband satellite experiment, a radio beacon experiment that was operational from 1976 to 1979 [17], [21]. From a near-polar 1000-km orbit, the Wideband satellite broadcast 10 coherent continuous-wave signals whose frequencies ranged from VHF to S-band and included an L-band signal at 1239 MHz that is nearly coincident with the GPS L_2 signal. The 2891-MHz S-band signal, which, because of its high frequency was minimally affected by scintillation, was used as the phase reference for the lower frequencies in a coherent, multi-band receiver chain. The receiver chain's baseband phase and amplitude outputs were band-limited to 150 Hz and digitized at 500 Hz. It should be emphasized that these baseband data, not the original high frequency carrier signals, are the Wideband data referred to in this paper.

The Wideband data included in the scintillation library were recorded at two equatorial ground stations, one in Ancon, Peru (11.8° N latitude) and the other on Kwajalein atoll in the Marshall Islands (9.4° N latitude). Each ground station was equipped with a 9.1-m parabolic tracking antenna. The high-gain antennas, together with 12-bit quantization and the

500-Hz sampling rate, yielded high-quality, low-noise complex signal time histories. Moreover, the fact that the received signals were continuous-wave ensured that cycle slips, when they occurred, were whole-cycle slips or multiples thereof; i.e., there are none of the half cycle slips that occur when tracking the bi-phase modulated GPS signals. Whole cycle slips are present during very strong patches of scintillation, as evidenced by a comparison among signals at different frequencies, but these do not diminish the usefulness of the data in an empirical model. This issue will be further discussed in a later section.

Among the 10 signals broadcast by the Wideband satellite, the 1239-MHz L-band signal, which is nearest the 1227.6-MHz GPS L_2 frequency, is of greatest interest for studying scintillation effects on GPS receivers. However, because the Wideband data were recorded roughly two years before solar maximum, the L-band data are not representative of the strongest equatorial scintillation at GPS frequencies. To represent extreme scintillation events, the L-band data are augmented in the scintillation library with unscaled data from the uppermost UHF carrier at 447 MHz. As will be detailed in the following section, it is conjectured that the UHF data provide test cases that exceed the worst L-band scintillation likely to occur at solar maximum.

Ten passes marked by strong scintillation were chosen from the archives of the Wideband experiment for inclusion in the scintillation library. Each recorded satellite pass yields approximately 12 minutes of continuous L-band and UHF data, which are stored in the library as complex time histories at the original 500-Hz sampling rate. No high-pass filtering to remove trends nor low-pass filtering to remove receiver noise were required.

The second source of data for the scintillation library is a set of GPS L_1 C/A code digital data sampled at 5.7 MHz. The data were recorded in December 2003 at Cachoeira Paulista, Brazil (22.7° S latitude) [19], which lies along the southern boundary of the Appelton equatorial anomaly—a region notorious for strong scintillation [9]. The recorded data set spans several hours and includes strongly scintillating signals on multiple GPS links. The data have been extensively processed to generate empirical scintillation library records that are comparable in quality to those from the Wideband data. The goal of this processing has been to eliminate, insofar as possible, all receiver and Doppler effects on the phase and amplitude data, leaving only effects due to scintillation.

The raw 5.7-MHz samples are initially processed through a software GPS receiver to extract carrier phase, amplitude, and navigation bit estimates for each GPS signal present. If bit parity failures are detected—as is often the case during patches of strong scintillation—then a record of correct navigation bits is compiled and used to eliminate, or "wipe off," the data bits from the incoming GPS signal in a reprocessing run, this time using a standard nonsquaring PLL in place of the squaring (Costas) loop. This strategy ensures that, like the Wideband data, the GPS phase time histories are free from half-cycle slips, though they may be affected by whole cycle slips. The carrier phase output is further processed to eliminate, to the extent possible, variations due to receiver clock instability, PLL dynamics, and Doppler. Both the output phase and amplitude data, which are sampled at 100 Hz, are then filtered to reduce receiver noise. However, owing to the weakness of the GPS signal and the low gain of the receiving antenna, the GPS data's signal-to-noise ratio (SNR) remains considerably lower than the Wideband SNR. Details of the GPS data preliminary processing techniques are found in the Appendix. Table I summarizes the contents of the empirical scintillation library.

TABLE I
SUMMARY OF EMPIRICAL SCINTILLATION LIBRARY CONTENTS

Data source	Original Carrier Frequency (MHz)	Combined length of records (hours)
Wideband	447 (UHF)	1.9
Wideband	1239 (near GPS L_2)	1.9
GPS	1575 (GPS L_1)	3.7

III. Characteristics of Severe Scintillation

Short samples from two different scintillation library records are presented in Fig. 1 (Wideband L-band data) and in Fig. 2 (GPS L_1 data). Both records manifest strong scintillation, each with $S_4 \approx 0.9$, where S_4 is the standard scintillation index defined by

$$S_4^2 = \frac{\langle I^2 \rangle - \langle I \rangle^2}{\langle I \rangle^2} \tag{1}$$

in which I is signal intensity (squared amplitude) and $\langle \cdot \rangle$ denotes time average. As mentioned previously, the GPS data's SNR is worse than that of the Wideband data (evidenced by the small, high-frequency ripples on the plots in Fig. 2), but this does not obscure the obvious similarities between the two samples. The most striking features are the deep power fades that occur simultaneously with abrupt, approximately half-cycle phase changes whose sense (downgoing or upgoing) appears random. Such fades are found in patches of severe scintillation throughout the scintillation library. They appear to be a universal feature of strong equatorial scintillation. As will be demonstrated later on, they are the primary cause of phase unlock for PLLs tracking a strongly scintillating signal that carries bi-phase modulated data, as is the case for GPS signals at L_1. They will be designated "canonical fades" for convenience in this paper. The term "canonical fade" is not meant as a hard categorization, but serves as shorthand for "power fade of sufficient depth and rapidity that the associated phase time history exhibits an abrupt, near half-cycle phase change."

A. Scintillation Severity

It will be useful to discuss the contents of the scintillation library in terms of the scintillation index S_4 and of a general measure of scintillation rapidity, the decorrelation time τ_0, defined as the time lag at which a scintillating complex baseband signal's autocorrelation function falls off by a factor $1/e$. For purposes of phase tracking, scintillation severity can be conveniently parameterized by S_4 and τ_0. As argued in [22], a characterization of scintillation severity in terms of S_4 and τ_0

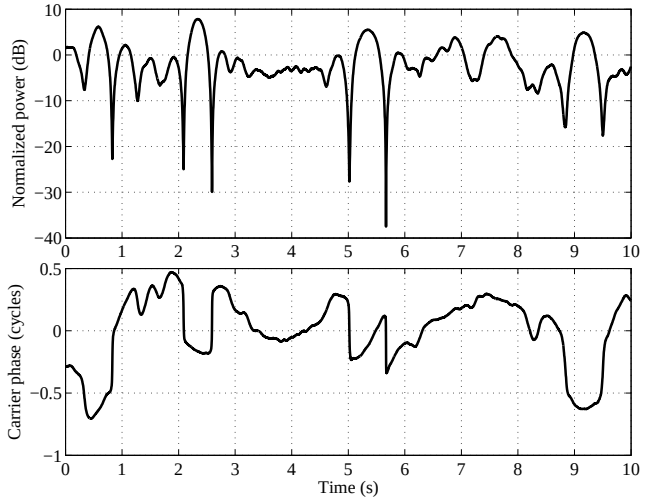

Fig. 1. An excerpt of normalized signal power (upper panel) and carrier phase (lower panel) from a record of Wideband L-band data with $S_4 \approx 0.9$.

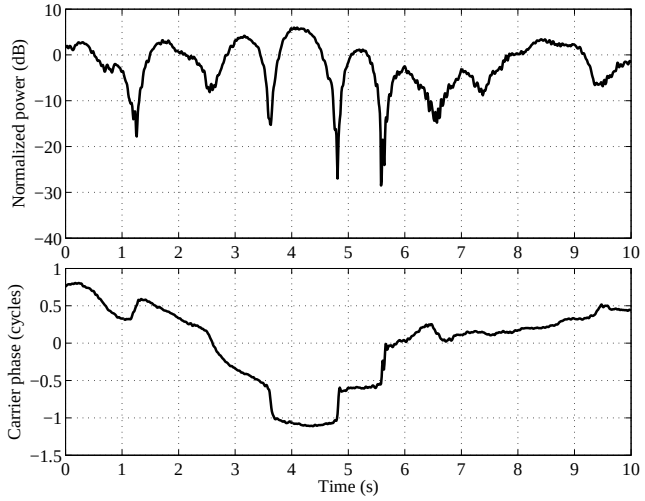

Fig. 2. As Fig. 1 for GPS L_1 data. Again, $S_4 \approx 0.9$.

is preferable to a characterization in terms of S_4 and the phase error variance σ_φ^2, which are the usual parameters employed in the literature on scintillation and tracking loops, because τ_0 can be calculated without invoking the PLL linearity assumption and because the pair (S_4, τ_0) is a better predictor of PLL cycle slipping than the pair (S_4, σ_φ^2). The scintillation library includes a reasonably comprehensive sampling of equatorial scintillation phenomena, with S_4 indices ranging from near zero to approximately 1.2 and values of τ_0 from 0.1 to 2 seconds.

During intervals of weak to moderate scintillation ($S_4 \lesssim 0.5$), the Wideband 447-MHz data manifest S_4 indices that are approximately 3.7 times greater than those for the corresponding interval at 1239-MHz. This is roughly consistent with the $S_4 \propto f^{-1.5}$ scaling predicted by weak scintillation theory [1]. As scintillation severity increases, multiple scattering effects cause the S_4 indices at all scintillation library frequencies (447, 1239, and 1575 MHz) to saturate near 1. Thus, when considered only in terms of the S_4 index, patches of strong UHF scintillation do not appear more severe than patches of strong L-band scintillation.

In contrast, the rapidity of strong scintillation at UHF can exceed that at L-band for equivalent line-of-sight dynamics. This is because the multiple-scattering effects associated with severe scintillation lead to spectral broadening (decreased τ_0), and these effects are stronger at lower frequencies [1]. Thus, values of τ_0 for the scintillation library's Wideband UHF and L-band data recorded simultaneously during patches of severe scintillation reach lows of 0.1 and 0.4 seconds, respectively, which confirms the increased severity of the UHF scintillation.

Because of the different ionospheric conditions along the various GPS links represented in the GPS L_1 data, the data include a wide range of τ_0 values, from ~ 2 seconds during weak to moderate scintillation to 0.4 seconds during severe scintillation. That the minimum values of τ_0 for the Wideband L-band and GPS data are not more disparate is surprising given the large difference in the radio line of sight scanning velocity in the ionosphere's F layer. (Recall that the F layer is the dense ionospheric layer at approximately 350 km altitude where the irregularities that cause scintillation tend to form. Thus, the scanning velocity mentioned here is the velocity of the line of sight puncture point through a hypothetical shell around the earth at approximately 350 km altitude.) Over the intervals shown in Figs. 1 and 2, the Wideband and GPS τ_0 values were 0.4 and 0.8 seconds, respectively. By comparison, the respective scanning velocities were about 4000 m/s and 100 m/s, and typical background ionospheric drift velocities (the velocity of ionospheric structures with respect to a frame fixed to the earth) range from 100 to 200 m/s [23]. Hence, excluding other factors, one would expect the Wideband and GPS τ_0 values for these segments to differ by at least a factor of 10 instead of 2. This paradox can be resolved by considering the geometry of ionospheric structures and their drift rates relative to the scan velocity vector. As noted in [23], time scales of equatorial scintillation are determined by the horizontal velocity of the scan point measured with respect to the drifting ionosphere and in a direction perpendicular to the local magnetic field. For the GPS data, this relative velocity is determined by a combination of ionospheric drift and the east-west scanning of the line-of-sight vector. For the Wideband data, the near-polar orbit of the transmitting satellite implies that most of the scan velocity is in the north-south direction—nearly aligned with the local magnetic field. This is believed to be the reason why the large scanning velocity does not lead to greatly reduced values of τ_0.

It is shown in [22] that complex signal scintillation is, on average, more difficult to track as S_4 increases and as τ_0 decreases. Therefore, for the present study, the upshot of the foregoing S_4 and τ_0 considerations is that the Wideband L-band data can be considered a reasonable proxy for severely scintillating signals at the GPS L_2 frequency, whereas the Wideband UHF data is only useful in the sense that it is more vigorous than all conceivable extreme cases of L-band scintillation, and can thus be used to test the limits of GPS tracking loop capability. Slow scintillation (large τ_0) at high S_4 can also challenge tracking loops when deep power fades persist for several seconds. Such cases are represented in the

library's GPS L_1 data by several instances of power fades exceeding 15 dB and lasting ~ 2 seconds.

Note that the above time scales apply only to the static receivers that were used to collect the data contained in the scintillation library. Receivers mounted on dynamic platforms may see significantly faster or slower scintillation. For example, when a dynamic platform causes the line-of-sight vector to move "in resonance" with irregularity patterns, deep power fades may last significantly longer than 2 seconds [23]. By redefining the scintillation library's sampling interval, the library can easily accommodate more extreme scintillation time scales, but this has not been done for the present study.

B. First-order Statistics and Phase Power Spectra

Besides S_4 and τ_0, other signal statistics useful for scintillation modeling are shown in Figs. 3 and 4. Empirical first-order amplitude and phase rate distributions are plotted in Fig. 3. In panel (a) the full 300-second GPS record from which the excerpt in Fig. 2 is taken has been used to generate an empirical joint distribution, expressed in density of samples, in the variables $\Delta\theta$, the phase change over 20 ms (the duration of a GPS navigation data bit), and α, the normalized fading amplitude. Strongly scintillating Wideband data are similarly distributed. The distribution reveals the effect of canonical fading as an increased spreading of $\Delta\theta$ for small values of α. Some spreading in $\Delta\theta$ is also due to an increase in phase errors at low loop SNR, but this is a lesser effect. The marginal distribution of α [panel (c)] agrees well with both the Rice distribution [24, Ch. 2], and with the popular Nakagami-m distribution, given by [16]

$$p(\alpha) = \frac{2m^m \alpha^{2m-1}}{\Omega^m \Gamma(m)}\, e^{-m\alpha^2/\Omega}, \quad \alpha \geq 0 \tag{2}$$

where the parameter $m = 1/S_4^2 \geq 1/2$ and $\Omega \equiv E[\alpha^2]$ is the mean-square of the fading amplitude. Here, amplitude normalization implies $\Omega = 1$. In practice, the ensemble expectation $E[\cdot]$ is computed as the time average $\langle\cdot\rangle$ over intervals of approximate stationarity. The marginal distribution of $\Delta\theta$ [panel (b)] appears benign, but a close examination reveals long tails relative to the Gaussian distribution. Section V will show that these tails are responsible for most scintillation-induced cycle slipping.

Another characterization commonly used for scintillation effects modeling is the phase power spectrum, shown for several example data sets in Fig. 4. The usual filtering of the GPS L_1 data (see the Appendix) has been omitted so that the frequency content of the raw scintillating signal is properly exhibited. The figure's three spectra are all derived from strongly scintillating data ($S_4 > 0.9$). Over frequencies of interest for phase tracking, the spectra conform to the profile predicted by weak-scatter phase screen theory [18]:

$$S_\theta(f) = \frac{T}{(f_0^2 + f^2)^{p/2}} \tag{3}$$

where T is a strength parameter, f_0 is the temporal frequency corresponding to the ionospheric outer scale, and p is the phase power law index. In practice, the value of f_0 is not well-defined. In any case, it is much smaller than frequencies relevant to PLL phase errors, which are on the order of the PLL noise bandwidth (~ 10 Hz). Accordingly, it is assumed that $f_0 \ll f$ over frequencies of interest, and Eq. (3) is approximated by the power law $S_\theta(f) = Tf^{-p}$, where T is reinterpreted as the power density at 1 Hz.

The pronounced bend in the GPS L_1 spectrum near 10 Hz marks the GPS data's noise floor. No such floor is present in the Wideband data out to the 250-Hz Nyquist cutoff frequency, a reflection of the Wideband data's high SNR. Accordingly, when used in the scintillation testbed, the Wideband library records are modeled as having no noise component (infinite SNR) whereas each GPS record is modeled as having a known, irreducible noise component.

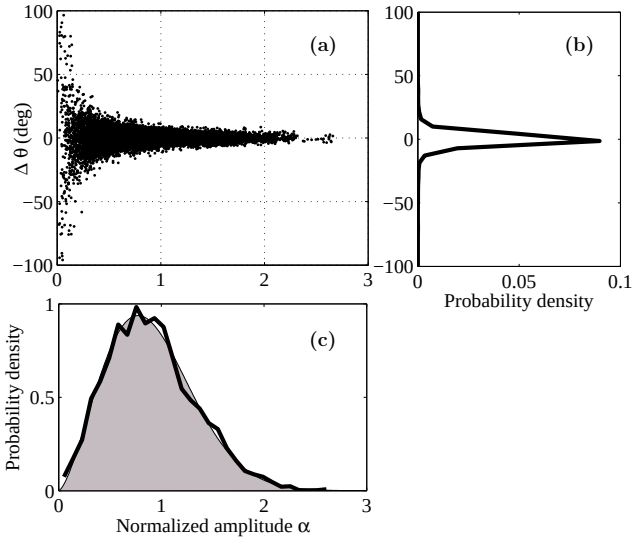

Fig. 3. (a) Empirical joint probability distribution in the variables $\Delta\theta$ and α. (b) Marginal empirical distribution in $\Delta\theta$. (c) Marginal empirical distribution in α (solid line) compared with the Nakagami-m distribution (shaded area) corresponding to the empirical S_4 value.

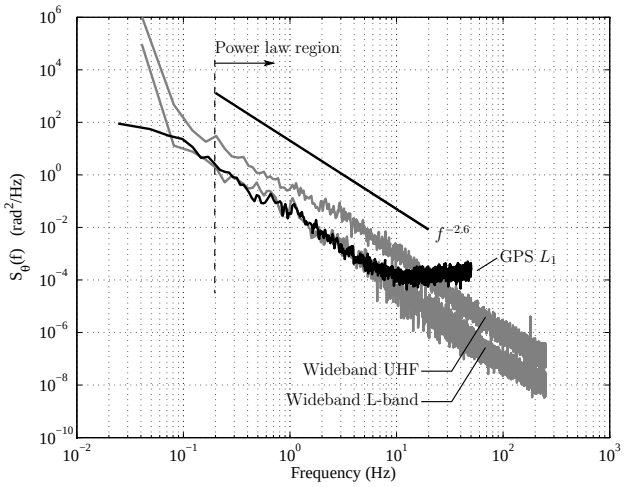

Fig. 4. Phase power spectra of various data samples. For frequencies to the right of the dashed line, the spectra are well approximated by the power law $S_\theta(f) = Tf^{-p}$. For reference, the solid line shows the spectral slope corresponding to $p = 2.6$.

Until diverging at the GPS noise floor, the GPS and Wideband data have similar p values, in this case $p \approx 2.6$. A survey of all library records shows 2.6 to be the empirical limit of p at high S_4. Such convergence is a consequence of the canonical-fade-induced abrupt phase changes, which, as they become more frequent and discontinuous, drive p toward 2, the high frequency asymptote for any discontinuous signal.

The important point of Fig. 4 as regards phase tracking is that phase scintillation is not confined below any particular frequency. Hence, some phase power density will unavoidably fall outside a practical PLL's tracking bandwidth, resulting in phase tracking errors. Such errors lead to cycle slips when the phase changes are sufficiently large and abrupt, as is often the case with canonical fading.

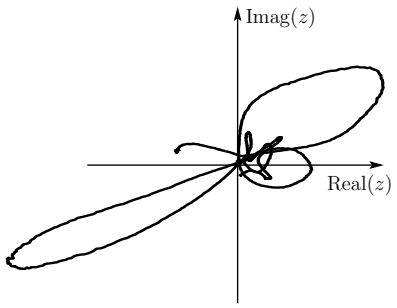

Fig. 5. A trace on the complex plane of the complex baseband signal $z(t)$ corresponding to the first three seconds of the amplitude and phase time histories in Fig. 1.

C. Canonical Fading

It is instructive to make a closer examination of the canonical fading phenomenon. Even before conducting experiments, one would expect canonical fading to have a strong adverse effect on PLL performance. Consider an analogy for the effect of canonical fading on a PLL's ability to track carrier phase: Suppose that, while driving at night, a driver's headlights dimmed each time the car came upon a curve in the road. Suppose, further, that the curve could be hard to the left or hard to the right with equal probability. Under these conditions, the driver would find it difficult to maintain the car on the road. In the analogy, the dimming headlights represent the fading-induced low loop SNR, which makes phase much harder to measure. Occasional cycle slips during canonical fades are inevitable as the PLL "misses the curve." This hypothesis is borne out by testbed experiments.

At first glance, it appears that the rapid phase changes associated with the canonical fades in Figs. 1 and 2 might be phase errors that arise at low loop SNR rather than genuine shifts in the incoming carrier phase. However, the two effects can be distinguished by noting that, once the nominal SNR level is recovered after a fade, the net phase change should be small under the hypothesis that the true carrier phase is changing smoothly. Instead, a net phase change approaching $1/2$ cycle is observed. Furthermore, given that both the GPS receiver (with data bit wipeoff) and the Wideband receiver are incapable of fractional cycle slips, the approximately half-cycle phase shifts cannot be attributed to receiver cycle slipping.

The canonical fading phenomenon becomes intuitive when a receiver's scintillating complex baseband signal $z(t)$ is represented on the complex plane (Fig. 5). The signal's instantaneous phase and amplitude can be associated with a phasor that wanders about the plane with a speed inversely proportional to τ_0. As strong fading drives the phasor's magnitude toward zero, it's motion through a small neighborhood centered at the origin tends to produce swift, half-cycle phase changes, the hallmark of canonical fades.

Canonical fades have been noted in earlier studies of scintillation [15], [18] (though not by this name), and have been reproduced in phase screen scintillation models [25], [26], but they have not been accounted for in models of scintillation effects on GPS tracking loops [4]–[6], [8], [15]. These latter models successfully reproduce the first-order scintillation statistics described in [17] and the phase and amplitude spectra derived from phase screen theory in [1] and [18]; nevertheless, their phase and amplitude time histories are unconvincing because they fail to reproduce the canonical fading phenomenon at high S_4.

To understand why, recall from Fig. 3 that canonical fading shows up as a correlation between signal amplitude and the absolute value of phase rate—a correlation that is not detected by first-order phase and amplitude statistics. Moreover, although the phase time histories shown in Figs. 1 and 2 lead to power-law phase spectra (Fig. 4), the inverse is not generally true; that is, a power-law phase spectrum does not guarantee a realistic phase time history. The method used to generate a synthetic phase time history in [6] (and likely in [8], though details are not provided) consists of passing white Gaussian noise through a low-pass shaping filter. While this method can indeed produce a phase time history with an arbitrarily shaped spectrum, it does not tend to generate the large, nearly discontinuous phase jumps associated with canonical fades. Moreover, because the methods in [6] and [8] shape the phase and amplitude spectra independently, whatever abrupt phase changes are present in the synthesized phase time history are not, in general, coincident with deep amplitude fades. As a result, the synthetic phase time history is unrealistically easy to track.

Failing to account for canonical fading can also compromise scintillation models based on empirical data. For example, in [20] a scintillation monitor's 50-Hz phase measurements are used to drive phase variations in a hardware-in-the-loop simulation. In a preprocessing step, the phase data are first-order low-pass filtered with a cutoff frequency of 2-Hz. This is done under the assumption that the filtering reduces the high-frequency noise in the measured phase data without significantly distorting the phase variations due to scintillation. The problem with this approach is that it artificially rate-limits the genuine phase changes associated with canonical fades, thereby making the phase history easier to track. To avoid such rate limiting, the empirical scintillation models advanced in the present paper are either not low-pass filtered (Wideband data) or are low-pass filtered with a generous 20-Hz cutoff

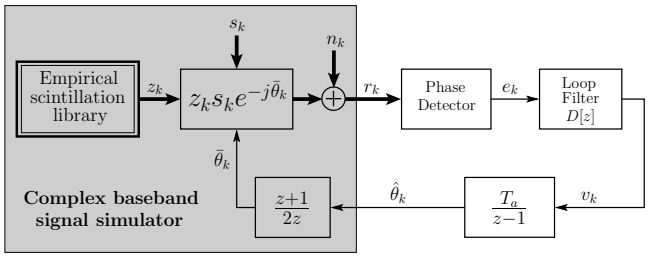

Fig. 6. Block diagram of the scintillation testbed. Thick lines denote complex signal routing.

frequency (GPS data).

IV. The Scintillation Testbed

The scintillation testbed is a collection of software routines used to test various phase tracking strategies under realistic equatorial scintillation conditions. The testbed draws input phase and amplitude time histories from the empirical scintillation library, feeds these to a PLL under test, and observes the PLL's phase error variance and phase lock behavior.

A. Description

The scintillation testbed is illustrated in block diagram form in Fig. 6. Its PLL model is an adaptation of the standard baseband model for a discrete-time PLL [27]; a GPS-specific baseband model is developed in [3]. Use is made of the z-transform to denote delays in the discrete-time signals. Routes for complex and scalar quantities are denoted respectively by thick and thin lines.

Within the gray box in Fig. 6 are simulated the mixing and accumulation operations typical in a GPS receiver. The simulation is based on the following model for the complex baseband signal r_k:

$$r_k = \frac{1}{T_a}\int_{t_{k-1}}^{t_k} \left[z(t)s_k + n(t)\right] e^{-j\hat{\theta}(t)} dt \tag{4}$$

In this model, $z(t) = \alpha(t)\exp[j\theta(t)]$ is the complex channel response function, with $\alpha(t)$ and $\theta(t)$ representing scintillation-induced amplitude and phase variations; it is assumed that $z(t)$ is normalized so that $E[|z(t)|^2] = 1$. The quantity $\hat{\theta}(t)$ is the PLL's phase estimate; $T_a = t_k - t_{k-1}$ is the accumulation interval; $n(t)$ is receiver noise represented as complex zero-mean additive white Gaussian noise; $s_k = \sqrt{E_a}\exp[j(\tilde{\theta}_k + \theta_c)]$, with $\tilde{\theta}_k \in \{0, \pi\}$, is the value over the kth accumulation interval of the 50-Hz binary data that commonly modulates GPS signals; θ_c is the carrier phase, which is assumed constant; and E_a is the energy per accumulation. Note that $\tilde{\theta}_k$ does not change from one accumulation interval to the next except when adjacent intervals straddle a data bit transition. Note also that $z(t)$, $n(t)$, and $\hat{\theta}(t)$ are represented here as continuous-time signals whereas they exist as discrete-time signals in modern GPS receivers. The continuous-time approximation is justified because their discrete-time sampling interval is much smaller than T_a.

Equation (4) is a baseband model, meaning that the nominal carrier frequency and all motion-induced Doppler shift effects are assumed to have been perfectly removed via baseband mixing. This is a realistic assumption for 3rd-order PLLs, whose static phase error is zero for a constant carrier frequency rate. Of course, the testbed could easily be modified to include the effects of motion-induced Doppler shift or receiver clock drift. Also, to preserve a focus on scintillation effects, the effects of receiver oscillator noise are neglected in Eq. (4) and throughout this paper.

The phase estimate $\hat{\theta}(t)$ varies linearly and, for practical PLL bandwidths, slowly over the accumulation interval. Hence r_k may be approximated as

$$\begin{aligned} r_k &\approx \frac{e^{-j\bar{\theta}_k}}{T_a}\int_{t_{k-1}}^{t_k} \left[z(t)s_k + n(t)\right] dt \\ &= z_k s_k e^{-j\bar{\theta}_k} + n_k \end{aligned} \tag{5}$$

where z_k and $\bar{\theta}_k = [\hat{\theta}(t_k) + \hat{\theta}(t_{k-1})]/2$ are the respective averages of $z(t)$ and $\hat{\theta}(t)$ over the accumulation interval, and n_k is an element of a zero-mean complex Gaussian noise sequence with variance $E[n_k^* n_j] = N_0\delta_{kj}$. Under the assumption that $E[|z(t)|^2] = 1$, the SNR per accumulation in the absence of scintillation (i.e., when $S_4 = 0$) is given by $E_a/N_0 = T_a C/N_0$, where C/N_0 is the usual carrier power to noise density ratio. The values of E_a/N_0 and C/N_0 when $S_4 = 0$ are referred to as the nominal E_a/N_0 and C/N_0. For the GPS L_1 data, the effective nominal carrier-to-noise ratio is less than C/N_0 because of additional noise power that is already present in the data. The final value of C/N_0 is calculated based on N_0 and on the effective C/N_0 of each GPS L_1 record, an estimation strategy for which is described in the Appendix. Also, in general, when $S_4 > 0$ the variations in $z(t)$ over the accumulation interval lead to values of $\Omega \equiv E[|z_k|^2]$ that are slightly less than unity. Thus the average SNR per accumulation $\bar{\gamma} \equiv \Omega E_a/N_0$ is slightly less than E_a/N_0.

To generate the averages z_k, the scintillation library's 500-Hz Wideband data samples z_{Wj} and 100-Hz GPS L_1 data samples z_{Gj} are treated slightly differently. Both data types are initially normalized so that $E[|z_{Wj}|^2] = E[|z_{Gj}|^2] = 1$. The Wideband data samples are then directly averaged to approximate the z_k, i.e.,

$$z_k \approx \frac{1}{N_W}\sum_{j=(k-1)N_W}^{kN_W-1} z_{Wj} \qquad k = 1, 2, 3, ... \tag{6}$$

, where $N_W = 500 \cdot T_a$ is the number of Wideband samples per accumulation. Since the scintillation testbed only considers $T_a \in \{10, 20\}$ ms, at least 5 samples are used to generate each average, making Eq. (6) a close approximation.

Because the GPS L_1 samples z_{Gj} are derived from actual 100-Hz baseband accumulations, they can be assumed to represent averages over $T_a = 10$ ms. Thus, no averaging is required to represent the z_k corresponding to $T_a = 10$ ms. However, the z_{Gj} must be attenuated to reflect the condition $\Omega \equiv E[|z_k|^2] < 1$ mentioned above. For this, Ω is estimated based on its dependence (detailed in [22]) on S_4 and τ_0. Then for $T_a = 10$ ms, the z_k are approximated as

$$z_k \approx \sqrt{\hat{\Omega}}\; z_{Gk} \tag{7}$$

where $\hat{\Omega}$ is an estimate of Ω in the neighborhood of z_{Gk}. For $T_a = 20$ ms, two of the 10-ms z_k from Eq. (7) are averaged to approximate the 20-ms z_k. The outputs z_k of the empirical scintillation library block in Fig. 6 are taken to represent the approximations derived from Eqs. (6) and (7).

The three blocks to the right of the gray box in Fig. 6 are particular to the PLL under test. They represent the target tracking loop's phase detector, loop filter, and phase integrator. The phase integrator is assumed to be discretized as a simple Euler integration operation.

The testbed also includes routines for detecting discrete cycle slips and for detecting a continuous succession of cycle slips, a condition termed "frequency unlock." The correct way to measure the mean time to first slip T_s empirically for the case of white phase noise is demonstrated in [28] where, after each slip, the phase tracking loop is re-initialized in perfect steady-state lock so that it experiences no transients. However, mean time between cycle slips, not mean time to first slip, is the quantity of interest for PLLs operating in the presence of scintillation. The two quantities are not equivalent when slips occur in bursts, as can be the case for 2nd- and higher-order loops during severe scintillation or at low loop SNR. Thus, the testbed's cycle slip detection routine measures mean time between slips, and T_s will be assumed to refer to this quantity in testbed results.

The measurement of T_s depends on a somewhat arbitrary definition of a cycle slip. The following procedure, which leads to cycle slip counts that agree with a visual inspection, has been adopted for the scintillation testbed. A phase error time history $\varphi_k = \theta_c + \theta(t_k) - \hat{\theta}(t_k)$ is first constructed where for the true scintillation time history $\theta(t_k)$ is substituted the phase of the original z_{Gj} or z_{Wj} samples taken at times nearest each t_k. The sequence φ_k is then 1.5-Hz low-pass filtered to produce the time history $\tilde{\varphi}_k$, from which the differences $b_k = |\tilde{\varphi}_k - \tilde{\varphi}_{k-N}|$ are computed, where N is chosen such that the differenced elements are separated by the rise time of the low-pass filter. A cycle slip is declared when b_k exceeds $2\beta\pi$ rad (non-squaring loops) or $\beta\pi$ rad (squaring loops), with $\beta = 3/5$. Frequency unlock is declared after 5 successive seconds of cycle slipping. The time interval of each testbed run is chosen to be long enough that loop re-initialization at the beginning of each run has a negligible effect on the mean time between cycle slips or frequency unlock. Values of T_s greater than the time interval of each individual testbed run are calculated by averaging cycle slip counts over repeated runs.

It was noted previously that whole cycle slips, although rare, are present in both Wideband and GPS records in the scintillation library. These occur during particularly deep canonical fades where, for example, instead of tracking an abrupt half-cycle upgoing phase change, the original PLL tracks an abrupt half-cycle downgoing phase change, resulting in a whole cycle slip. From the point of view of a PLL being evaluated in the scintillation testbed, the true (upgoing) phase change and the tracked (downgoing) phase change are equally challenging to track. Hence, the whole cycle slips do not diminish the usefulness of the data in the scintillation library.

It should be noted that, owing to GPS modernization, future GPS systems will increasingly be multi-frequency systems, with a separate tracking loop assigned to track the carrier at each frequency. To properly evaluate the scintillation robustness of a multi-frequency GPS system requires processing multiple scintillation time histories—one for each signal frequency. As has been established in the scintillation literature, the pairwise complex cross-correlation coefficients of scintillation time histories taken at different frequencies fall off with frequency separation and with increased perturbation strength [29]. In other words, there is a narrowing of the coherence bandwidth with increased scintillation intensity. Although the empirical scintillation library contains concurrent data at two different frequencies (Wideband UHF and L-band data), and narrowing of the coherence bandwidth is obvious in the data, the L-band-to-UHF frequency ratio is too large (2.8) to be representative of the frequency ratios in modernized GPS and other modern satellite navigation systems (approximately 1.4 or less). A synthetic scintillation model such as the one presented in [30], when augmented to generate realistic multi-frequency scintillation, will likely be the most practical means for testing multi-frequency systems.

B. First-order Validation

Besides agreement with visual inspection, the scintillation testbed has been checked against Viterbi's theoretical results for the first-order continuous-time PLL operating in white Gaussian noise [31]. For these tests, T_a was set to 1 ms (much slower than the loop time constant) and a 1st-order $D[z]$ was designed using a constant update approximation so that the PLL's closed-loop behavior is nearly equivalent to that of the analog PLL assumed in Viterbi's analysis. Tests were conducted using simulated white Gaussian noise (no phase or amplitude scintillation). The test results showed close agreement with theory: the experimental distribution of φ_k modulo 2π was practically indistinguishable from its theoretical counterpart, and the experimental T_s at each SNR tested was within 10% of the theoretical value when compared on a logarithmic plot.

V. Testbed Demonstration and Results

This section reports on a preliminary, qualitative study of scintillation effects based on results of testing with the scintillation testbed. A more comprehensive study can be found in [22].

The results that follow assume the standard PLL architecture shown in the three blocks on the right-hand side of Fig. 6. The phase detectors considered are standard squaring-type detectors used for bi-phase modulated data: the two-quadrant arctangent [$e_k = \mathrm{atan}(Q_k/I_k)$], the conventional Costas [$e_k = I_k \cdot Q_k$], and the decision-directed [$e_k = Q_k \cdot \mathrm{sign}(I_{m,k})$] detectors, where I_k and Q_k denote the in-phase and quadrature components of the complex baseband signal r_k and $I_{m,k}$ is the sum of the in-phase components of all accumulations up to time t_k within the current (mth) data bit interval.

The digital loop filter $D[z]$ takes the phase detector outputs e_k and estimates the phase rate v_{k+1} for the $(k+1)$th

accumulation interval according to

$$v_{k+1}T_a = K_1 e_k + K_2 \sum_{i=1}^{k} e_i + K_3 \sum_{i=1}^{k} \sum_{j=1}^{i} e_j + ... \qquad (8)$$

where the sequence extends to the K_N term for an Nth-order loop filter, which yields Nth-order closed-loop PLL dynamics. The loop constants K_n are determined according to the controlled-root formulation for digital PLL design introduced in [28]. For the decision-directed and the conventional Costas phase loops, the outputs e_k have been normalized to enforce a constant loop bandwidth. For the arctangent phase detector, such normalization is automatic.

A. Effects of Canonical Fading

To appreciate the effects of canonical fading on phase tracking loops, consider Fig. 7, which shows the results of testing a 3rd-order PLL with noise bandwidth $B_n = 10$ Hz, update interval $T_a = 10$ ms, and a two-quadrant arctangent phase detector. Normalized signal power (top trace) and carrier phase (middle trace) are from a library record of severely scintillating GPS L_1 data with $S_4 = 0.91$ and $\tau_0 = 0.43$ seconds. The nominal C/N_0 value over the testbed simulation was approximately 44 dB-Hz.

Several cycle slips are evident in the difference between the library and the estimated carrier phase time histories (bottom trace), leading to a mean time between cycle slips of $T_s \approx 20$ seconds for this run. This is a typical value of T_s for severe scintillation. To allow examination of the conditions that caused the slips, two excerpts have been expanded in the upper two panels. The slips shown in these panels are clearly caused by canonical fades. This is also true for the rest of the slips in the interval, and, indeed, is true for the vast majority of slips that occur when tracking data from the scintillation library.

Noting that canonical fading involves two effects, namely, amplitude fades and abrupt phase changes, one might naturally wonder whether one of the two effects dominates the cycle slipping rate. The scintillation library allows one to weigh the contribution of each effect by nulling the other; that is, for a given library record of complex scintillation, one can manipulate the data to eliminate the amplitude fades and retain the phase variations, or vice versa.

The contribution of each effect depends on the nominal C/N_0 value and on the configuration of the PLL under test. For typical nominal C/N_0 values (near 43 dB-Hz for GPS tracking) and for a typical loop configuration (2nd- or 3rd-order loops with B_n near 10 Hz, T_a near 10 ms, and either a two-quadrant arctangent, decision-directed, or conventional-Costas phase detector) the following results hold: When only amplitude fades are retained, the cycle slipping rate plummets to less than 10% of the native rate. Conversely, when only the phase changes are retained, the cycle slipping rate remains at roughly 50% of the native rate. Clearly, then, rapid phase changes dominate the cycle slipping rate under typical conditions. Nonetheless, it is the combined effects of amplitude fades and sharp phase changes that make tracking through scintillation particularly difficult.

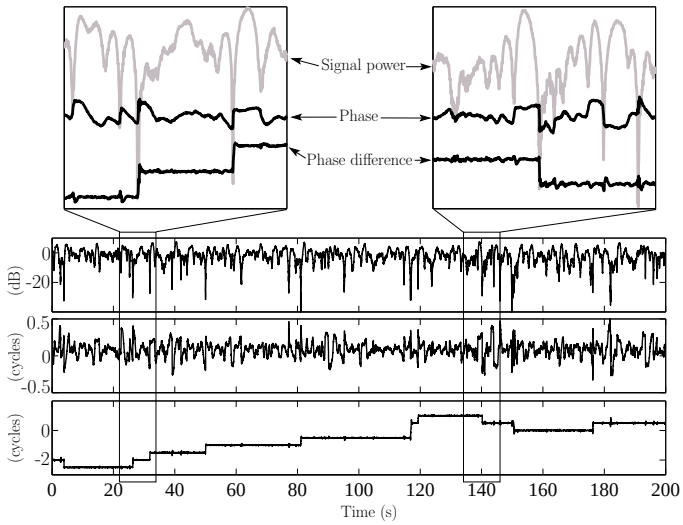

Fig. 7. Example testbed results, showing normalized signal power (top trace) carrier phase (middle trace) and phase error (bottom trace).

B. Loop Order, Update Interval, and Noise Bandwidth Comparisons

The scintillation testbed has been used to compare the performance of PLLs with various loop orders, noise bandwidths, and update intervals. Such comparison tests are used to identify scintillation-robust loop parameters for the class of standard-architecture constant-bandwidth PLLs described above. Tests have been performed by varying one parameter (e.g., loop order) at a time from a baseline configuration with loop order 3, update (accumulation) interval $T_a = 10$ ms, and noise bandwidth $B_n = 10$ Hz. The following results apply generally to the two-quadrant arctangent, decision directed, and conventional Costas phase detectors. A performance comparison among phase detectors, which requires more thorough testing, is found in [22]. Measurement noise n_k was added to simulate a nominal C/N_0 value of 43 dB-Hz for the Wideband data, and nominal C/N_0 values between 40 and 44 dB-Hz for the GPS data, with a mean of 43 dB-Hz. Each PLL configuration was tested on the entire 7.5-hour scintillation library.

Loop order: PLLs of order 2 and 3 were tested. No significant difference in T_s was found. This was true whether all scintillation library data or only GPS L_1 data were used. On the other hand, the 2nd-order loop did show a reduced rate of frequency unlock across all phase detectors (recall that frequency unlock refers to a continuous succession of cycle slips). However, the 3rd-order loop's ability, and the 2nd-order loop's inability, to maintain zero steady-state phase error when tracking quadratic phase growth (not simulated for the present testbed results) would tend to neutralize this advantage, especially for dynamic platforms. Therefore, 2nd- and 3rd-order loops appear to be roughly equally suited for GPS phase tracking in scintillation for static receivers, whereas 3rd-order loops are likely to be preferable to 2nd-order loops for dynamic receivers.

Loop update interval: PLLs with $T_a = 10$ ms and $T_a = 20$ ms were tested. It was found that for all phase detectors an increase in T_a from 10 to 20 ms worsened (shortened)

T_s. Moreover, the incidence of frequency unlock increased substantially—consistent with the reduced pull-in threshold of discrete-time loops with longer update intervals [22]. Thus, $T_a = 10$ ms appears preferable to $T_a = 20$ ms for phase tracking in scintillation.

Loop noise bandwidth: PLLs with $B_n = 7.5$, $B_n = 10$ and $B_n = 15$ Hz were tested. In general, PLLs with higher B_n experience fewer cycle slips in severe scintillation at high nominal C/N_0 because they are better able to track the vigorous phase fluctuations than lower-bandwidth PLLs. As nominal C/N_0 decreases below 40 dB-Hz, however, measurement noise effects become important and a wide bandwidth loop becomes comparatively more prone to cycle slipping and frequency unlock. In this regime, B_n values at or below 7.5 Hz are appropriate. Near a nominal C/N_0 of 43 dB-Hz, an increase in B_n from 7.5 to 10 Hz resulted in a marginal improvement (increase) in T_s but also caused an increase in the incidence of frequency unlock for all phase detectors. As B_n was further widened to 15 Hz, T_s improved still further but the rate of frequency unlock increased dramatically. Hence, it appears that B_n values from 7.5 to 10 Hz are a good choice for operating in severe scintillation at typical nominal values of C/N_0.

VI. Conclusions

Time histories of severe equatorial phase and amplitude scintillation from the Wideband Satellite experiment and from specially-processed GPS data have been compiled into a scintillation library. 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" have been noted before by scintillation researchers, but they have not been expressly accounted for in existing models for scintillation effects on GPS receiver phase tracking loops.

A scintillation effects testbed that derives inputs from the scintillation library has been developed and used to test several standard GPS phase tracking loop designs. Testing results show that scintillation-induced cycle slips in squaring-type phase tracking loops are predominantly caused by canonical fades. Results also indicate that to minimize cycle slipping when tracking through severe equatorial scintillation, a 2nd- or 3rd-order loop with an update interval of 10 ms and a noise bandwidth from 7.5 to 10 Hz is best within the class of standard-architecture constant-bandwidth phase tracking loops.

Appendix

This appendix describes the preliminary processing used to prepare the GPS L_1 data for inclusion in the scintillation library. In an initial processing pass, a GPS L_1 C/A-code software receiver operates on raw 5.7-MHz data samples that were recorded during a scintillation event. The software receiver produces navigation data bit, phase, and amplitude estimates for all GPS signals present in the data. Its PLL is a $B_n = 10$-Hz, 3rd-order loop with an arctangent discriminator and a 100-Hz pre-detection bandwidth; its loop filter is based on the controlled-root formulation [28]. If navigation bit parity failures are detected in the output data, then the raw data samples are reprocessed, this time with the bi-phase navigation data bits wiped off. Bit wipeoff requires reconstruction of an entire 12.5-minute navigation data superframe, which proceeds by replacing corrupted sections with their error-free counterparts. Such a reconstruction is possible because almost all the navigation data bits are either repeated across GPS links or are repeated in time during a 2-hour window. Exceptions to this rule occur infrequently and are limited to the 14-bit TLM message of the TLM word. If scintillation-induced bit errors occur here, they can often be corrected by hand. Since only certain patterns of four-bit errors are undetectable under the GPS parity check code, a reconstructed superframe is considered intact, for purposes of data bit wipeoff, if it passes the parity check.

This process of checking for parity errors and, if detected, reprocessing with the data bits wiped off ensures that, like the Wideband data, no fractional cycle slips remain in the GPS software receiver's output phase time history. However, other effects remain in the phase data, namely, phase variations due to (1) receiver clock instability, (2) carrier Doppler due to line-of-sight motion, (3) PLL dynamics, and (4) receiver noise. An attempt is made to remove each of these effects as follows.

Receiver clock instability: The frequency reference of the digital storage receiver on which the raw 5.7-MHz GPS data were recorded is a temperature-compensated crystal oscillator (TCXO). It is known that the power spectrum of phase variations due to the TCXO's frequency instability significantly overlaps the power spectrum of phase scintillation [32]. Fortunately, the two effects can be separated in post-processing because, whereas phase scintillation is different from one GPS link to the next, receiver clock noise is common to all links. Removal of clock noise proceeds as follows: Over each interval of recorded GPS data, a relatively strong, minimally scintillating GPS link is chosen as the reference signal. The reference signal's phase time history is subtracted from the time history of a scintillating signal of interest. The resulting phase time history is free of phase variations due to receiver clock instability. Naturally, an unavoidable consequence of this approach is an approximate doubling of the spectrally flat measurement noise content in the differenced phase data.

Motion-induced carrier Doppler: Even after removal of receiver clock effects, a large low-frequency component due to the difference in motion-induced carrier Doppler shift between the scintillating and reference signals remains in the differenced phase time history. This component is easily eliminated by subtracting a low-order curve fit from the differenced phase time history.

PLL dynamics: The scintillating signal's true phase variations are filtered through the transfer function of the software receiver's PLL. Such filtering band limits (~ 10 Hz) and adds second-order dynamics (i.e., overshoot and oscillation) to the desired scintillation time history. To eliminate these effects, the instantaneous phase error $\phi(t_k)$, as measured using the known navigation bits and the four-quadrant arctangent function, is added to the PLL's phase estimate $\hat{\theta}(t_k)$. The resulting signal is band limited only by the software receiver's 100-Hz pre-

detection bandwidth.

Receiver noise: Widening the bandwidth of the software receiver's phase estimate as described above increases its broadband measurement noise component. Ideally, one would like to separate measurement noise from genuine scintillation-induced phase variations, but there is no way to distinguish the two in practice. The best one can do is identify a frequency beyond which phase scintillation has a negligible effect on phase tracking; the phase data is then low-pass filtered at this cutoff frequency. Unfortunately, the scintillating phase's power-law spectrum does not present any obvious candidates for a cutoff frequency. Upon examining the active GPS spectra in Fig. 4, one might be tempted to choose the frequency at which the phase spectrum breaks at the noise floor (~ 10 Hz) as the low-pass cutoff, but testbed runs using data with 10-Hz and 20-Hz cutoffs showed that, for the same effective C/N_0, the 20-Hz data were significantly more challenging to track. Also, at 20-Hz, the canonical-fade phase variations in the GPS data resembled much more closely the sharp phase changes in the low-noise Wideband data. Therefore, a conservative 20-Hz cutoff has been chosen. Both phase and amplitude data are 6th-order low-pass Butterworth filtered at this cutoff; hence, time shifts in the filtered amplitude and phase data are equal. For the amplitude data, filtering is done separately for the real and imaginary components of the baseband complex signal. The magnitude of the filtered result becomes the scintillating amplitude time history.

For use in the scintillation testbed, a measure of the noise component remaining in the amplitude and phase time histories is required so that approximate nominal C/N_0 values can be assigned to tests with the data. A scheme for estimating the intrinsic C/N_0 value of each GPS scintillation record has been developed based on (1) the known noise floor of the digital storage receiver's RF front-end, (2) the scintillation record's nominal amplitude, estimated as $\bar{A} = \sqrt{\Omega}$, where Ω is defined in Section III-B, and (3) a model for how broadband noise is added in the phase differencing process and subtracted in the filtering process. The resulting C/N_0 values assigned to each GPS record ranged from 38 to 44 dB-Hz.

ACKNOWLEDGMENTS

The authors give special thanks to J. Secan of NorthWest Research Associates for providing raw Wideband experiment data.

REFERENCES

[1] 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.

[2] 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.

[3] 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.

[4] M. Knight and A. Finn, "The effects of ionospheric scintillation on GPS," in *Proceedings of ION GPS 1998*. Nashville, TN: Institute of Navigation, 1998.

[5] 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.

[6] 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.

[7] 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*. Salt Lake City, Utah: Institute of Navigation, 2000.

[8] 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.

[9] J. Aarons, "Global morphology of ionospheric scintillations," *Proceedings of the IEEE*, vol. 70, no. 4, pp. 360–378, 1982.

[10] ——, "Global positioning system phase fluctuations at auroral latitudes," *Journal of Geophysical Research*, vol. 102, pp. 17,219–17,231, 1997.

[11] B. M. Ledvina, J. J. Makela, and P. M. Kintner, "First observations of intense GPS L1 amplitude scintillations at midlatitude," *Geophysical Research Letters*, vol. 29, no. 14, pp. 4–1—4–4, 2002.

[12] N. Luo and G. Lachapelle, "Relative positioning of multiple moving platforms using GPS," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 3, pp. 936–947, July 2003.

[13] P. Misra and P. Enge, *Global Positioning System: Signals, Measurements, and Performance*. Lincoln, Massachusetts: Ganga-Jumana Press, 2006.

[14] A. Cameron, "Next-generation air traffic controlled by GPS," GPS World, July 2006.

[15] M. A. Cervera and M. F. Knight, "Time series modelling of intensity and phase scintillation at GPS frequencies," *Acta Geodaetica et Geophsysica Hungarica*, vol. 33, no. 1, pp. 25–40, 1998.

[16] M. Nakagami, "The m-distribution: A general formula of intensity distribution of rapid fading," in *Statistical Methods in Radio Wave Propagation*, W. C. Hoffman, Ed. New York: Pergamon, 1960, pp. 3–36.

[17] 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.

[18] C. L. Rino, "A power law phase screen model for ionospheric scintillation (1. Weak scatter)," *Radio Science*, vol. 14, no. 6, pp. 1135–1145, 1979.

[19] 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.

[20] 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.

[21] 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.

[22] 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*, vol. 46, no. 4, Oct. 2010.

[23] P. M. Kintner, B. M. Ledvina, E. R. de Paula, and I. J. Kantor, "Size, shape, orientation, speed, and duration of GPS equatorial anomaly scintillations," *Radio Science*, vol. 39, pp. 2012–+, April 2004.

[24] M. K. Simon and M. Alouini, *Digital Communications over Fading Channels*. New York: Wiley, 2000.

[25] 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.

[26] V. E. Gherm, N. N. Zernov, and H. J. Strangeways, "Rate of phase change and fade depth on GPS paths due to ionospheric scintillation," in *Proc. of Twelfth International Conference on Antennas and Propagation (ICAP 2003)*, vol. 1. IEE, 2003, pp. 413–416.

[27] W. Lindsey and C. M. Chie, "A survey of digital phase-locked loops," *Proceedings of the IEEE*, vol. 69, no. 4, pp. 410 – 431, April 1981.

[28] S. A. Stephens and J. B. Thomas, "Controlled-root formulation for digital phase-locked loops," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 31, no. 1, pp. 78–95, Jan. 1995.

[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.

[30] T. E. Humphreys, M. L. Psiaki, J. C. Hinks, B. O'Hanlon, and P. M. Kintner, Jr., "Simulating ionosphere-induced scintillation for testing GPS receiver phase tracking loops," *IEEE Journal of Selected Topics in Signal Processing*, vol. 3, no. 4, pp. 716–725, Aug. 2009.

[31] A. J. Viterbi, *Principles of Coherent Communication*. New York: McGraw-Hill, 1966.

[32] A. J. Van Dierendonck, "How GPS receivers measure (or should measure) ionospheric scintillation and TEC and how GPS receivers are affected by the ionosphere," in *Proc. 11th International Ionospheric Effects Symposium*, Alexandria, VA, 2005.