

Modeling the effects of ionospheric scintillation on GPS carrier phase tracking

Todd E. Humphreys, Mark L. Psiaki, and Paul M. Kintner, Jr.

Abstract—A characterization is given for the behavior of Global Positioning System phase tracking loops in the presence of severe equatorial ionospheric scintillation. The purpose of this work is to develop a simple, general, and realistic scintillation effects model that can be used to improve the scintillation performance of phase tracking loops. The new characterization of scintillation effects proposed herein employs a differentially-detected bit error model to predict cycle slipping rates that approximately agree with data-driven simulation tests.

I. INTRODUCTION

There is interest in developing Global Positioning System (GPS) receivers whose carrier phase tracking loops are specially designed to maintain lock in the presence of severe equatorial scintillation [1]. Design and testing of such receivers depends crucially on accurate models of scintillation and its effects on receivers. Several researchers have responded to this need by developing models for scintillation effects on phaselock loops (PLLs) [2]–[7], but these models tend to underestimate the effects of severe equatorial scintillation by failing to capture its essential feature: deep power fades (> 15 dB) accompanied by abrupt, approximately half-cycle phase transitions, so-called canonical fades [1].

In [1], the current authors develop a data-driven scintillation testbed and with it show that canonical fades are the principal cause of carrier unlock for squaring-type PLLs attempting to track through severe equatorial ionospheric scintillation. The current paper extends the work in [1] by employing the scintillation testbed to carry out an extensive characterization of squaring-type PLL behavior in the presence of severe equatorial scintillation. Insight gained from this characterization will lead to a simple model linking certain scintillation statistics to mean time between PLL cycle slips.

To set the stage for interpreting the behavior of PLLs under test, Section II addresses some general effects of scintillation on phase tracking. The testbed is briefly reviewed in Section III. Section IV describes the example PLLs that will be evaluated in the testbed and Section V presents their performance results. The results will demonstrate a close connection between cycle slips and errors in differentially-detected navigation data bits. This connection will suggest a model, developed in Section VI, of scintillation effects on

GPS receiver phase tracking loops. The conclusions follow in Section VII.

II. OVERVIEW OF SCINTILLATION EFFECTS ON PHASE TRACKING LOOPS

Phase tracking loops are affected by scintillation in three related ways: (1) increased phase error variance, (2) cycle slipping, and (3) frequency unlock. This section gives an overview of these effects to set terms and develop intuition before the scintillation testbed is introduced. To preserve a focus on scintillation effects, receiver oscillator noise is neglected here and in the rest of this paper. For a general treatment of phase tracking loops, the reader is referred to [8]–[11].

A. Phase Error Variance

Consider a standard (non-squaring) PLL with true phase input $\theta(t)$ and phase estimate $\hat{\theta}(t)$. When the phase error $\varphi(t) = \theta(t) - \hat{\theta}(t)$ is small enough that the PLL's phase detector can be regarded as linear, then, for zero-mean white driving noise, the PLL's phase error variance $\sigma_\varphi^2 = E[\varphi^2(t)]$ (in rad²) is accurately approximated by [10]

$$\sigma_\varphi^2 = \frac{B_n N_0}{C} \equiv \frac{1}{\rho_L} \quad (1)$$

where C/N_0 is the carrier power to noise density ratio, B_n is the PLL's single-sided noise bandwidth, and ρ_L is the loop SNR. Under normal circumstances, GPS carrier phase tracking requires a squaring (e.g., Costas) PLL, which is insensitive to the half-cycle phase changes induced by the data bit modulation. In a squaring PLL, the actual phase error tracked is 2φ , with the corresponding variance denoted by $\sigma_{2\varphi}^2$. Furthermore, ρ_L is reduced by a squaring loss factor approximately equal to [7]

$$S_L = \left(1 + \frac{N_0}{2T_a C}\right)^{-1} \quad (2)$$

where $1/T_a$ is the pre-detection bandwidth. [Equation (2) is exact for a Costas loop with a multiplicative phase detector.] Thus, for the squaring loop,

$$\sigma_\varphi^2 = \frac{\sigma_{2\varphi}^2}{4} = \frac{1}{\rho_L S_L} \quad (3)$$

is a useful approximation for σ_φ^2 in the linear regime. For analysis of the squaring loop, an equivalent loop SNR is defined as [12, p. 206]

$$\rho_{eq} \equiv \frac{\rho_L S_L}{4} \quad (4)$$

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, 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.

which leads to $\rho_{eq} = 1/\sigma_{2\varphi}^2$ for small 2φ .

At large values of φ , the assumption of PLL linearity breaks down and analysis becomes much more difficult. An exact expression for σ_φ^2 for a 1st-order non-squaring PLL driven by white Gaussian noise is found in [8, Ch. 4]. Precise phase error statistics for all but this standard 1st-order loop are typically obtained via simulation. Fortunately, one can show that the exact phase error variance for the standard 1st-order loop is a reasonable proxy for that of higher order loops. Thus one can identify the region of approximate linear PLL operation by noting that, for the standard 1st-order loop, the linear model [Eq. (1)] is reasonably accurate (within 20%) for $\rho_L > 4$, or $\sigma_\varphi < 28.6^\circ$ [8, Ch. 4]. Likewise, a squaring loop behaves approximately linearly for $\rho_{eq} > 4$, or $\sigma_\varphi < 14.3^\circ$.

Note that the above results assume that all phase errors are due to constant-intensity white measurement noise. Furthermore, as just mentioned, Eqs. (1) – (4) assume PLL linearity. All of these assumptions are violated during severe scintillation: Amplitude fading causes variations in the loop SNR, phase changes are time-correlated, and, when attempting to track through the large, rapid phase changes associated with canonical fading, the PLL cannot be expected to operate in its linear regime. For these reasons, calculating σ_φ^2 for a PLL tracking through strong scintillation does not appear straightforward. The approach taken in [2] and [5] breaks σ_φ^2 down into three components: (1) a measurement noise component that takes amplitude fading into account, (2) a component due to dynamics in $\theta(t)$, and (3) a component due to oscillator noise. But to calculate these components, the approach in [5] invokes PLL linearity, which is violated during severe scintillation.

B. Cycle Slipping

A PLL's phase detector is periodic, meaning that it cannot distinguish between the phase errors φ and $\varphi + 2n\pi$ (non-squaring loop) or φ and $\varphi + n\pi$ (squaring loop), where n is an integer. As a result, an infinite set of stable attractors exists for the nonlinear difference equations that describe the PLL error dynamics. At low loop SNR, or during vigorous phase scintillation, the phase error can slip from one stable attractor to another, leading to infinite σ_φ^2 in the steady-state. This is the familiar cycle slip phenomenon associated with PLLs [11], [13], [10, Ch. 6].

The mean time to first cycle slip T_s is defined as the average time required for the loop phase error to reach $\pm 2\pi$ ($\pm\pi$ for the squaring loop) for the first time, starting from an initial condition of zero phase error. For first order loops, and in other cases where cycle slips occur as isolated events, T_s is the same as the mean time between cycle slips; if cycle slips occur in bursts—as may happen for $\rho_L, \rho_{eq} < 5$ in 2nd- or higher-order loops—then T_s and the mean time between cycle slips are not related simply [13].

As with the calculation of σ_φ^2 , an analytical solution for T_s has only been possible for the simple case of a 1st-order unstressed (zero static phase error) PLL driven by white Gaussian noise, in which case [8, p. 101]

$$T_s = \frac{\pi^2 \rho_L I_0^2(\rho_L)}{2B_n} \quad (5)$$

is the time to first slip/mean time between slips for a non-squaring loop; $I_0(\cdot)$ is a modified Bessel function of the first kind. An approximate T_s for 1st-order squaring loops obtains by substituting ρ_{eq} for ρ_L . Unstressed 2nd- and higher-order loops have lower T_s than unstressed 1st-order loops, and stressed loops are more prone to cycle slipping than unstressed loops; nonetheless, Eq. (5) remains a useful upper bound. For GPS applications, a 2nd- or 3rd-order loop is required to accurately track carrier phase in the presence of Doppler-induced quadratic phase growth. In fact, even the 2nd-order loop experiences significant loop stress ($\sim 1^\circ$ static phase error) during the largest GPS line-of-sight accelerations. Only the 3rd-order loop maintains near-zero static phase error for all GPS geometries. Accordingly, scintillation testbed experiments will focus on 3rd-order loops.

Phase and amplitude scintillation cause cycle slipping via two different mechanisms. The first of these—and the most common—was discussed briefly in the previous section: a squaring-loop PLL is often unable to track through the abrupt phase change associated with a canonical fade. This makes sense: in the limit as the fade depth increases, the abrupt, nearly π -rad phase transition looks like bi-phase data modulation, to which the squaring-loop PLL is insensitive by design. Hence, the PLL detects no phase shift and a half-cycle slip occurs. In marginal cases, where the PLL might be capable of distinguishing a scintillation-induced phase transition from a data-bit-induced phase transition, the sudden drop in loop SNR increases the likelihood of a cycle slip. One can demonstrate this by shifting the phase and amplitude time histories relative to one another so that the signal power fades and the rapid phase changes are no longer aligned. Experiments show that such shifting leads to an approximate 20% decrease in cycle slips—evidence that simultaneous power fades and abrupt phase changes are a particularly challenging combination. Scintillation testbed results will show that canonical fading accounts for over 90% of equatorial-scintillation-induced cycle slips.

Prolonged amplitude fading is the second mechanism by which scintillation causes cycle slipping. This phenomenon may be considered a special case of canonical fading in which the fading time scale is elongated so that the amplitude fade is accompanied by phase dynamics that are slow compared to a typical 10-Hz PLL noise bandwidth. In this case, broadband measurement noise dominates and Eq. (5) applies. Cycle slips occur rarely by this mechanism. To see why, consider that the overwhelming majority of deep fades (> 15 dB) in the scintillation library's GPS data last less than 2 seconds. Assume that, during a rare 2-second fade, the received C/N_0 drops to 23 dB-Hz. For a pre-detection bandwidth of $1/T_a = 100$ Hz and a PLL noise bandwidth equal to $B_n = 10$ Hz, the resulting equivalent loop SNR is $\rho_{eq} = 4$. Substituting ρ_{eq} into Eq. (5) yields $T_s \approx 22.3$ seconds for a 1st-order loop, which implies a slip probability of less than 1/10 over the 2-second fade. Consistent with this analysis, cycle slips due to prolonged amplitude fading are rare, amounting to less than 10% of the cycle slips in the scintillation testbed results.

C. Frequency Unlock

The general term “phase unlock” refers to single or successive cycle slips. At very low loop SNR or in extreme scintillation a PLL may never recover phase lock after a long succession of cycle slips. This phenomenon, called “drop lock” in the PLL literature, is related to the PLL’s frequency pull-in range. For reasons that will become clear, the term “frequency unlock” is a more precise descriptor than drop lock for the phenomenon as it relates to the discrete-time PLLs used in modern GPS receivers.

A PLL’s frequency pull-in range is the maximum frequency step input that a PLL is able to “pull in” and eventually achieve phase lock. For example, a continuous-time 1st-order non-squaring PLL has a pull-in range equal to the loop gain K [11]. For higher-order PLLs, the frequency pull-in range can be thought of as the maximum tolerable mismatch $\Delta\omega = |\omega_c - v|$ between the carrier frequency ω_c and the PLL’s internal estimate of carrier frequency v , assuming that higher order loop filter states (e.g., the estimate of carrier frequency rate) are relaxed, where applicable.

Continuous-time PLLs whose loop filters contain one or more perfect integrators have an infinite frequency pull-in range [10, Ch. 8]. On the other hand, the frequency pull-in range of 2nd- and higher-order discrete-time PLLs is limited by the loop update (accumulation) interval T_a . When the frequency mismatch $\Delta\omega$ exceeds a certain threshold $\Delta\omega_m$, then v is attracted toward a stable equilibrium value that satisfies $T_a\Delta\omega = n\pi$ (non-squaring loop) or $T_a\Delta\omega = n\pi/2$ (squaring loop), $n = 1, 2, 3, \dots$. Intuitively, these equilibrium values exist because the loop cannot detect a phase error change of $2n\pi$ (non-squaring loop) or $n\pi$ (squaring loop) between loop updates. The value of $\Delta\omega_m$ is a function of the particular loop configuration. It can be surprisingly small for PLLs common in GPS receivers: for a 3rd-order Costas loop with $T_a = 10$ ms and $B_n = 10$ Hz, $\Delta\omega_m = 81$ rad/s ≈ 13 Hz. At very low loop SNR or during vigorous scintillation, cycle slips can occur in bursts as noise and phase dynamics force v momentarily away from ω_c [13]. If, due to such forcing, $\Delta\omega$ exceeds $\Delta\omega_m$, then there is a high probability that v will become trapped at one of the incorrect stable equilibrium values. Thus the PLL experiences frequency unlock.

Frequency unlock and momentary phase unlock have rather different practical consequences. Unlike momentary phase unlock (i.e., cycle slipping), frequency unlock often leads to complete loss of the GPS signal link—a result of signal attenuation due to frequency detuning. If v settles on an equilibrium value such that $n \geq 2$ (non-squaring loop) or $n \geq 4$ (squaring loop), then the baseband signal power drops by more than 13 dB, making it likely that the PLL will experience further frequency detuning and eventually lose the signal entirely. Worse yet, re-acquisition may not be possible at low SNR or during scintillation. Such complete signal loss can occur even in high-quality GPS receivers during severe scintillation [14], with sobering implications for GPS-dependent systems.

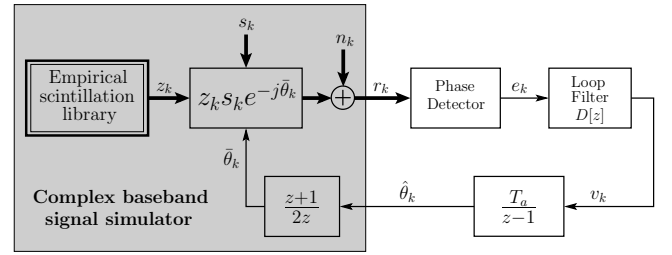


Fig. 1. Block diagram of the scintillation testbed. Thick lines denote complex signal routing.

III. 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 derives input phase and amplitude time histories from a library of empirical scintillation data (referred to hereafter as the scintillation library), feeds these to a PLL under test, and observes the PLL’s phase error variance and phase lock behavior. A complete description of the testbed is given in [1]. This section gives a brief overview and describes an additional feature beyond those detailed in [1].

The scintillation testbed is illustrated in block diagram form in Fig. 1. Within the gray box in Fig. 1 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 :

$$r_k = \frac{1}{T_a} \int_{t_{k-1}}^{t_k} [z(t)s_k + n(t)] e^{-j\hat{\theta}(t)} dt \quad (6)$$

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; $s_k = \sqrt{E_a} \exp[j(\hat{\theta}_k + \theta_c)]$, with $\hat{\theta}_k \in \{0, \pi\}$, is the value over the k th 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. The quantities $z_k, \theta_k = [\hat{\theta}(t_k) + \hat{\theta}(t_{k-1})]/2$, and n_k in Fig. 1 are the respective averages of $z(t), \hat{\theta}(t)$, and $n(t)$ over the accumulation interval; n_k is modeled as an element of a zero-mean complex Gaussian noise sequence with variance $E[n_k^* n_j] = N_0 \delta_{kj}$.

The three blocks to the right of the gray box in Fig. 1 are particular to the PLL under test and will be treated in the next section. In addition to the functions depicted in Fig. 1 and detailed in [1], the scintillation testbed has been augmented to evaluate the bit error probability of binary differential phase shift keying (binary DPSK) bit detection. Binary DPSK makes a decision about the sign of a data bit based on the phase change between two adjacent data bit intervals. It is useful for the testbed to produce binary DPSK bit error rate statistics because, as will be shown, there exists a close connection

between binary DPSK bit errors and cycle slips for GPS signal tracking during severe equatorial scintillation. A ‘‘Fast DPSK’’ scheme (proposed in [15]) that does differential detection using only one sub-bit-length accumulation interval on each side of a data bit edge is also evaluated. Both DPSK techniques operate on the baseband signal r_k , which implies knowledge of the carrier frequency ω_c . In practice, ω_c can be estimated using a low-bandwidth frequency tracking loop.

For the case of an additive white Gaussian noise channel, the testbed’s binary DPSK bit error probability estimates have been validated against theoretical values (given, for example, in [16, Ch. 7]).

IV. EXAMPLE PLLS

The scintillation testbed is a generalized tool for evaluating the performance of carrier tracking strategies in the presence of scintillation. Any carrier tracking loop that operates on the baseband signal r_k can be tested for robustness and evaluated—under controlled scintillation conditions—against an array of alternative strategies. To illustrate its usefulness and to gain insight into the effects of scintillation on typical PLLs used in GPS receivers, a set of example PLLs was tested. Each PLL in the set has a different combination of phase detector and loop filter.

The phase detectors considered in the set of example PLLs are given in Table I. All detectors are squaring-type detectors that make use of the in-phase $I_k = \text{Re}(r_k)$ and quadrature $Q_k = \text{Im}(r_k)$ components of r_k , which are simulations of the usual accumulations produced by the baseband mixer and integrate-and-dump accumulators of a digital GPS receiver. The two-quadrant arctangent (AT), decision directed (DD), and conventional Costas (CC) detectors are standard phase detectors used in GPS receivers [17], [18]. The AT and CC detector make direct use of the I_k and Q_k components whereas the DD detector attempts to wipe off data bit modulation by estimating the current data bit value as $\text{sign}(I_{m,k})$, where $I_{m,k}$ is the sum of the in-phase components of all accumulations up to time t_k within the current (m th) data bit interval. The decision-directed four-quadrant arctangent (DD-AT) detector is a hybrid of the DD and AT detectors in which data bit wipeoff via $\text{sign}(I_{m,k})$ enables full four-quadrant phase detection. If the accumulation interval T_a equals the GPS data bit interval $T_b = 20$ ms, then the AT and DD-AT detectors are equivalent.

TABLE I
PHASE DETECTORS FOR THE SET OF EXAMPLE PLLS

Phase Detector	Abbr.	Definition
Two-quadrant arctangent	AT	$e_k = \text{atan}(Q_k/I_k)$
Decision-directed four-quadrant arctangent	DD-AT	$e_k = \text{atan2}[Q_k \cdot \text{sign}(I_{m,k}), I_k \cdot \text{sign}(I_{m,k})]$
Conventional Costas	CC	$e_k = I_k \cdot Q_k$
Decision directed	DD	$e_k = Q_k \cdot \text{sign}(I_{m,k})$
Dot-product four-quadrant arctangent	DP-AT	$e_k = \text{atan2}(Q_k d_{m,k}, I_k d_{m,k})$

The dot-product four-quadrant arctangent (DP-AT) detector is identical to the DD-AT detector except that data bit wipeoff

is based on differential bit detection. Like binary DPSK bit detection, the DP-AT phase detector makes a decision about the sign of the current data bit by measuring—via a dot product—the phase change between the previous and the current data bit interval. Suppose it is known that \tilde{d}_{m-1} is the transmitted data bit over the $(m-1)$ th data bit interval. Then the DP-AT phase detector’s estimate $d_{m,k}$ of the value of the data bit during the accumulation ending at time t_k within the m th data bit interval is

$$d_{m,k} = \begin{cases} -\tilde{d}_{m-1}, & \text{Re}(r_{m,k}^* \tilde{r}_{m-1}) < 0 \\ \tilde{d}_{m-1}, & \text{Re}(r_{m,k}^* \tilde{r}_{m-1}) \geq 0 \end{cases} \quad (7)$$

where $r_{m,k} = I_{m,k} + jQ_{m,k}$ is the sum of all r_k up to time t_k within the m th data bit interval and \tilde{r}_{m-1} is the sum of all r_k within the $(m-1)$ th data bit interval. The value $d_{m,k}$ is used to wipe off data bit modulation from each r_k during the m th data bit interval, as indicated in Table I. At the end of the m th data bit interval, $d_{m,k}$ becomes \tilde{d}_m and $r_{m,k}$ becomes \tilde{r}_m . For added stability, the DP-AT detector is designed to revert to the DD-AT detector when $\text{Re}(r_{m,k}^* \tilde{r}_{m-1}) \approx 0$. It will be shown that the DP-AT detector is well suited for making phase measurements during severe equatorial scintillation.

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

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

where the sequence extends to the K_N term for an N th-order loop filter, which yields N th-order closed-loop PLL dynamics. The loop constants K_n are determined according to the controlled-root formulation for digital PLL design introduced in [19]. Loops of order 1 through 3 can be selected, with the design parameter $T_a B_n$ ranging from 0.01 to 0.4. The loop filter’s output v_k passes through the integration block $[T_a/(z-1)]$ to generate the phase estimate $\hat{\theta}_k$.

All PLLs in the example set are designed to have a constant bandwidth, meaning that their noise bandwidth B_n remains constant for small phase errors despite changes in the fading amplitude $\alpha_k = |z_k|$. When connected to a loop filter $D[z]$ with fixed coefficients, the AT-type phase detectors (AT, DD-AT, DP-AT) automatically deliver a constant B_n because, in the absence of noise, the arctangent function is insensitive to changes in α_k . The outputs e_k of the DD and CC phase detectors, on the other hand, are sensitive to changes in α_k . They must be normalized by α_k (DD) or α_k^2 (CC) to yield a B_n that, for small phase errors, is equal to the constant B_n of the AT-type detectors. In the scintillation testbed, such normalization is accomplished by estimating α_k^2 as

$$\hat{\alpha}_k^2 = \langle r_k^* r_k \rangle - N_0 \quad (9)$$

where the time average $\langle r_k^* r_k \rangle$ is mechanized using a first-order low-pass filter with a time constant of 0.1 seconds.

V. TESTBED RESULTS AND DISCUSSION

The scintillation testbed has been used to carry out an extensive set of tests to characterize squaring-type PLL behavior in the absence of scintillation (measurement noise only)

and in the presence of scintillation (both measurement noise and scintillation-induced phase and amplitude variations). The latter results will indicate a connection between cycle slips and DPSK-detected bit errors. This connection will be exploited in subsequent sections to develop a model that can be used to approximately predict cycle slips for the broad class of typical GPS PLLs represented by the set of example PLLs.

Performance tests of two elaborate variable-bandwidth phase tracking strategies have also been conducted using the scintillation testbed. The results of these tests, given in Section V-C, will suggest a limit on the cycle slip performance of any unaided PLL in severe equatorial scintillation.

The following tests assume the standard PLL architecture shown in the three blocks on the right-hand side of Fig. 1. In view of the parameter comparison tests conducted in [1], a baseline PLL configuration with loop order 3, $T_a = 10$ ms, and $B_n = 10$ Hz is assumed.

A. PLL Behavior in the Absence of Scintillation

It is instructive to contrast the behavior of the example PLLs in the absence of scintillation (white measurement noise only) with their behavior during scintillation. It will be shown that scintillation-free and scintillation tests that produce equivalent phase error variance σ_φ^2 are not equally challenging for the PLL under test.

To test PLL behavior in the absence of scintillation, the scintillation library's complex outputs z_k are set to unity for all k ; the white measurement noise simulated by n_k remains as described in [1]. Runs of 100-second length have been iterated 3000 times with C/N_0 ranging from 21 to 30 dB-Hz. Cycle slip results are presented in Fig. 2, which plots the normalized mean time between cycle slips $T_s B_n$ vs. the equivalent loop SNR ρ_{eq} (recall that ρ_{eq} is the loop SNR for squaring-type PLLs such as those found in the example set). For convenience, the horizontal axis is also expressed in terms of C/N_0 , which is related to ρ_{eq} via Eqs. (1), (2), and (4), and in terms of empirical values of σ_φ for the DD-AT PLL. Note that large phase errors at these low values of ρ_{eq} cause the empirical σ_φ values to be larger than predicted by the linear approximations expressed in Eqs. (3) and (4). Also note that to obtain sensible values of σ_φ in the presence of cycle slips, the phase errors φ_k are taken modulo π . The same is true for all values of σ_φ presented hereafter.

The results in Fig. 2 assume the baseline loop order and update interval (3rd order and $T_a = 10$ ms), but apply generally to $B_n \lesssim 20$ Hz. Second-order loops and loops with $T_a = T_b = 20$ ms have also been tested (recall that T_b is the GPS data bit interval). All loops demonstrated a marginal reduction in cycle slips with a change in T_a from 10 to 20 ms. No significant reduction was noted, over the range of C/N_0 values tested, due to a change from 3rd-order to 2nd-order loops.

Theoretical values of $T_s B_n$ vs. ρ_{eq} for a continuous-time 1st-order PLL (dashed line in Fig. 2) agree closely with the values for the DD-AT and AT PLLs, which demonstrates that Eq. (5) (with ρ_{eq} substituted for ρ_L) is a useful approximation for these PLLs over the range of ρ_{eq} shown. In terms of cycle

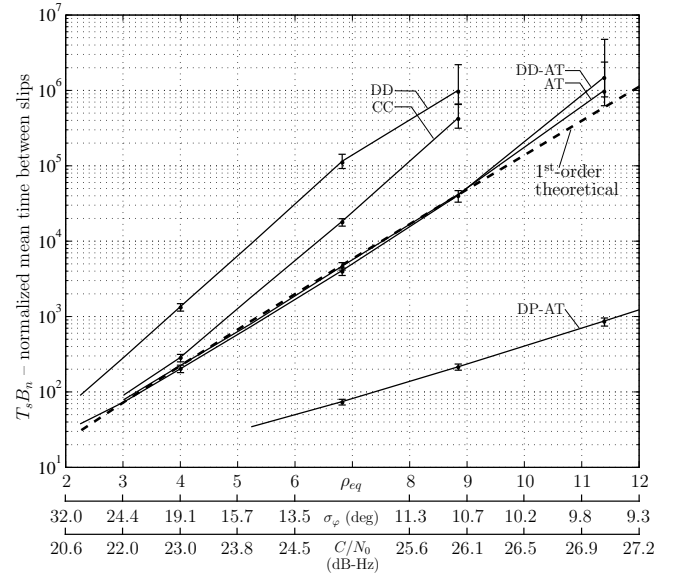


Fig. 2. Experimental normalized mean time between cycle slips as a function of ρ_{eq} (linear scale) for several 3rd-order PLLs with $T_a = 10$ ms operating in the presence of complex white Gaussian noise. The horizontal axis scale is also expressed in terms of the standard deviation σ_φ of the phase error φ_k modulo π , and in terms of the carrier-to-noise ratio C/N_0 , assuming $B_n = 10$ Hz. The values of σ_φ correspond to the DD-AT phase detector but also hold approximately for the CC and AT detectors. The dashed line shows the theoretical $T_s B_n$ for a continuous-time 1st-order PLL [Eq. (5)]. Error bars indicate two standard deviations about the estimated values.

slips, the DD and CC phase detectors have a decided advantage for the measurement-noise-only case. This result can be understood by noting that large phase errors cause the nonlinear (sinusoidal) phase detector characteristic of the DD and CC detectors effectively to reduce the associated PLL's noise bandwidth relative to PLLs based on the AT-type detectors, which have a linear characteristic over $-\pi/2 \leq \varphi_k < \pi/2$, with sharp nonlinear breaks at $-\pi/2$ and $\pi/2$. In the absence of actual variations in the received phase, PLLs with a lower effective B_n experience fewer cycle slips.

Among the phase detectors tested, the DP-AT detector suffered the most cycle slips at each value of ρ_{eq} . This result is intuitive: it can be shown that for the additive white Gaussian noise case the differentially-detected bit estimates $d_{m,k}$ that are used to do data bit wipeoff in the DP-AT phase detector are more prone to errors than the coherently-detected bit estimates $\text{sign}(I_{m,k})$ used in the DD and DD-AT detectors.

B. PLL Behavior in the Presence of Scintillation

Tests of PLL behavior in the presence of scintillation are conducted using the standard scintillation library outputs z_k . Scintillation data are divided into 30-second records so that the signal statistics over each record are reasonably stationary. Measurement noise n_k is added to simulate nominal C/N_0 values ranging from 35 to 55 dB-Hz. Each 30-second record at each value of C/N_0 is iterated 30 times with different measurement noise realizations to accumulate enough data for a statistical analysis of the results. Note, however, that the GPS L_1 results have measurement noise realizations that change

only in part between different runs because the data's intrinsic noise component is invariant.

In all, tests have been conducted on 885 30-second scintillation records with S_4 values ranging from 0.02 to 1.3, where S_4 is the standard scintillation index defined by

$$S_4^2 = \frac{\langle I^2 \rangle - \langle I \rangle^2}{\langle I \rangle^2} \quad (10)$$

in which I is signal intensity (squared amplitude) and $\langle \cdot \rangle$ denotes time average. Wideband and GPS L₁ results are presented separately to avoid conflating results from inputs with different noise properties. On the other hand, Wideband UHF and L-band results have been lumped together despite the inputs' different scintillation properties. This was done because there are not enough intervals of severe scintillation in the Wideband L-band data for clear trends to emerge, and because, for the relationships to be presented in this section, the L-band data generally fit trends established by the UHF data. The results given here apply to the baseline loop configuration (3rd-order loop, $T_a = 10$ ms, and $B_n = 10$ Hz).

The relationships between S_4 and the primary PLL performance metrics, namely, cycle slips and σ_φ^2 , are shown in Figs. 3 and 4. These results are for the DD-AT phase detector but are approximately representative of the other phase detectors tested. Each point in Figs. 3 and 4 corresponds to a different 30-second record from the scintillation library and represents an average over 30 test iterations. Fig. 3 presents results in terms of cycle slip rate on the left vertical axis, and, for convenience, in terms of T_s on the right vertical axis. As would be expected, a general increase in the rate of cycle slips accompanies increasing S_4 . The lack of cycle slips below $S_4 \approx 0.4$ suggests that, whatever its other characteristics (e.g., fading time scales), scintillation with $S_4 \lesssim 0.4$ can be considered benign. Also, for convenience in the following, severe scintillation will be roughly identified with $S_4 > 0.6$. The wide spread in cycle slip rate for $S_4 > 0.6$ indicates that records with equivalent measured S_4 are not necessarily equally challenging to track. This is true within the GPS L₁ records but is especially striking in the comparison between GPS and Wideband data. The implications of this spreading for scintillation effects modeling will be discussed in Section VI.

It was found in initial processing that GPS data bit parity failures—which can be considered a rough proxy for cycle slips—were predominantly associated with canonical fades. Experiments with the scintillation testbed confirmed this result: over 90% of the slips underlying the data in Fig. 3 were induced by canonical fades.

Fig. 4 shows how σ_φ , the standard deviation of the phase measurement error modulo π , increases with increasing S_4 , a dependence that is due both to the fade-induced reductions in loop SNR and to phase scintillation with frequency components that exceed the PLL's B_n . The difference in σ_φ at low S_4 between the Wideband and GPS data reflects the latter's irreducible noise component, which the analysis treats as a dynamically-varying true phase that the PLL does not exactly track. The large values of σ_φ at high S_4 —in some cases

exceeding 30 deg for Wideband data and 13 deg for GPS L₁ data—contribute to the degradation of carrier-phase-dependent GPS systems during strong scintillation.

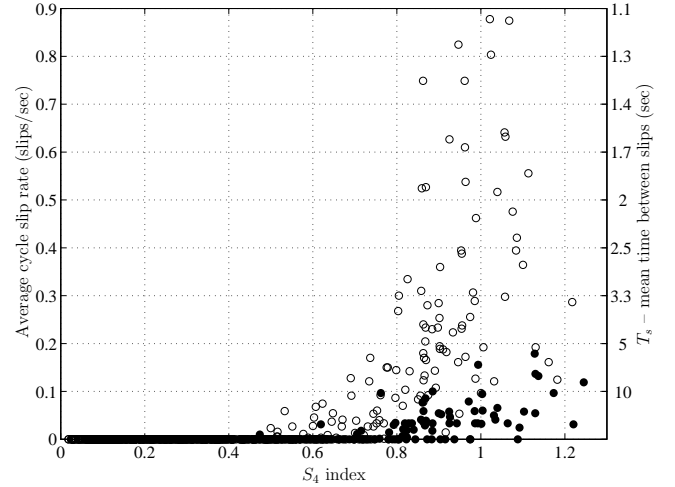


Fig. 3. Average DD-AT PLL cycle slip rate over each 30-second test record vs. S_4 for the Wideband data at $C/N_0 = 43$ dB-Hz (open circles) and for the GPS L₁ data within $40 < C/N_0 < 44$ dB-Hz with mean $C/N_0 = 43$ dB-Hz (filled circles). The right vertical axis expresses the cycle slip rate in terms of T_s .

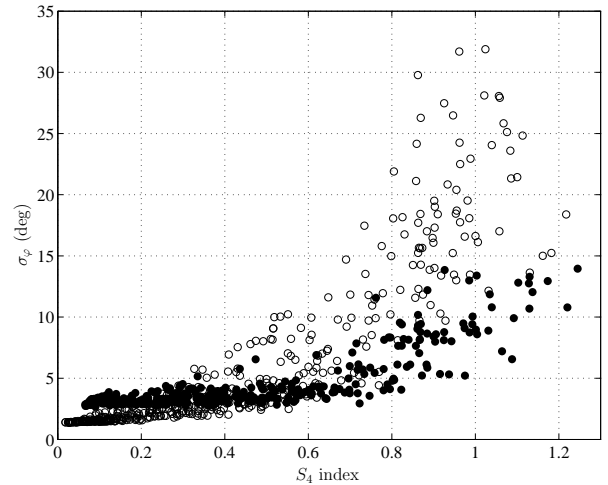


Fig. 4. Standard deviation of DD-AT PLL phase error modulo π over each 30-second test record vs. S_4 for the Wideband data at $C/N_0 = 43$ dB-Hz (open circles) and for the GPS L₁ data within $40 < C/N_0 < 44$ dB-Hz with mean $C/N_0 = 43$ dB-Hz (filled circles).

Fig. 5 plots T_s (solid lines) and the mean time between bit errors T_e (dashed lines) as functions of nominal C/N_0 for Wideband data in the severe scintillation regime ($S_4 > 0.6$). For clarity of presentation, error bars are not shown. Standard deviations for all curves are less than 0.4 seconds at $C/N_0 = 35$ dB-Hz and 1 second at $C/N_0 = 55$ dB-Hz. The range of nominal C/N_0 values considered in Fig. 5 is meant to span the worst to best case C/N_0 values that a terrestrial receiver would see in open-sky conditions. Empirical σ_φ values for the DD-AT PLL are given along the upper horizontal axis. One might think that these values appear too small to be associated with such high cycle slipping rates, but one must bear in mind that

it is the intermittent abrupt phase changes, not just an overall high σ_φ , that drive the slipping. This also helps explain another counterintuitive feature of Fig. 5, namely, the flattening of the T_s (solid) curves with increasing C/N_0 . Such flattening indicates that increasing the nominal tracking loop SNR has little effect because over this range of C/N_0 the cycle slipping rate is dominated not by uncorrelated thermal noise but by the intermittent abrupt phase changes associated with canonical fades.

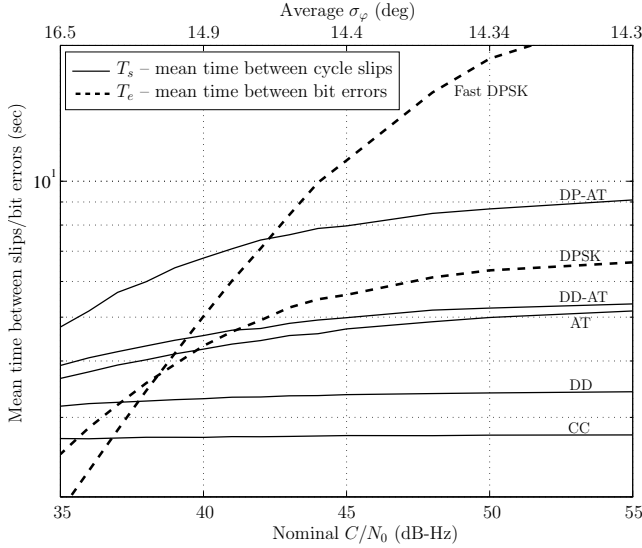


Fig. 5. Experimental mean time between cycle slips T_s as a function of nominal C/N_0 (dB-Hz) for several PLLs operating in severe scintillation ($S_4 > 0.6$). Input scintillation data are Wideband UHF and L-band records from the empirical scintillation library. The horizontal axis scale is also expressed in terms of the standard deviation σ_φ (deg) of the phase error φ_k modulo π (upper axis). The values of σ_φ correspond to the DD-AT phase detector but also hold approximately for the AT and DP-AT detectors. The dashed lines give the mean time between bit errors T_e for binary DPSK and Fast DPSK bit detection.

Compared with the measurement-noise-only case (Fig. 2), the phase detector performance ordering in Fig. 5 is nearly inverted, with the DP-AT phase detector now experiencing fewer cycle slips than the other detectors. Such a reversal in performance is best understood in terms of the PLL's effort to drive the quadrature component Q_k of r_k to zero. During an extreme canonical fade, the phase of the complex scintillation z_k can change by nearly π rad from one data bit interval to the next. In this case, none of the phase detectors can distinguish between a scintillation-induced phase change and a data-bit-induced phase change: a cycle slip results. In less extreme (slower and shallower) canonical fades, however, the DP-AT detector has an advantage. During this kind of fade, the baseband phasor r_k is dislodged from its nominal tracking orientation (aligned with the real axis in the complex plane) and rotates toward the imaginary axis; but as long as phase scintillation and noise do not cause r_k to rotate more than $\pi/2$ rad in one data bit interval, then the DP-AT detector decides the current bit sign correctly and thus avoids a cycle slip. In contrast, the non-differential phase detectors (DD-AT, AT, DD, and CC) almost always slip a cycle if, due to phase scintillation and noise, r_k rotates across the imaginary axis—no matter

how slowly the rotation occurs. Hence, the non-differential phase detectors rely on a wide noise bandwidth B_n to drive r_k safely to its nominal tracking orientation, whereas the DP-AT detector allows r_k to rotate up to π rad away from the real axis, so long as it rotates no more than $\pi/2$ rad per bit interval. In short, the DP-AT phase detector experiences fewer cycle slips because it is better able to tolerate large instantaneous phase errors.

The testbed's demonstration that the DP-AT PLL performs well during severe scintillation is consistent with the claim advanced in [15] that differential bit detection techniques are expected to outperform coherent detection techniques during scintillation. The claim is based on the hypothesis that scintillation can be classified as fast fading relative to the GPS data bit interval $T_b = 20$ ms. The flattening of the T_e curves in Fig. 5 at high nominal C/N_0 indicates an irreducible bit error probability. Such limiting behavior is indeed consistent with fast fading [20], [12, Ch. 8].

Besides its role as a relatively robust phase tracking strategy, the DP-AT PLL offers an important insight into cycle slipping: Errors in the differentially-detected bit estimates $d_{m,k}$ tend to cause the DP-AT PLL to slip cycles. Conversely, if the $d_{m,k}$ are error-free, then the DP-AT PLL tends to avoid cycle slips. This connection is evident in Fig. 5 where, as an accessible proxy for errors in $d_{m,k}$, errors in DPSK-detected data bits are used. The dashed DPSK curve indicates that the mean time between bit errors T_e for DPSK detection serves as a useful lower bound for the DP-AT PLL's T_s . Likewise, the Fast DPSK curve suggests that further cycle slip immunity at high nominal C/N_0 can be attained by reducing the interval over which phase changes are detected. The connection between DPSK-detected bit errors and cycle slips will be developed into a scintillation effects model in later sections.

Only Wideband data were used to generate the curves in Fig. 5 because of the difficulty in setting precise C/N_0 values for tests with the GPS data. Nonetheless, useful results for T_s and T_e are obtained from the GPS data by averaging over all 30-second records of severely scintillating GPS L₁ data ($S_4 > 0.6$) whose nominal C/N_0 values fall within the range $40 < C/N_0 < 44$ dB-Hz. The mean value of C/N_0 for records within this range is $C/N_0 = 43$ dB-Hz. As summarized in Table II, such averaging of the GPS data yields the same phase detector performance ordering as was noted with the Wideband data, albeit with an approximately 4-fold improvement in T_s and T_e . Also, the mean time between DPSK bit errors (36.8 seconds) is consistent with its previously noted role as lower bound to the DP-AT PLL's T_s (43.3 seconds).

TABLE II
AVERAGE T_s AND T_e FOR SEVERELY SCINTILLATING GPS L₁ DATA

Phase Detection		Bit Detection	
Detector	T_s (s)	Detector	T_e (s)
DP-AT	43.3	Fast DPSK	106.3
DD-AT	37.0	DPSK	36.8
AT	33.6		
DD	22.1		
CC	15.7		

Frequency unlock statistics have also been collected over all 30-second records with $S_4 > 0.6$. At a nominal $C/N_0 \sim 43$ dB-Hz, the DP-AT PLL's mean time between frequency unlock was 1.6 and 2.2 hours for the Wideband and GPS data, respectively. The AT and CC phase detectors fared slightly worse than this, whereas the DD and DD-AT phase detectors never experienced frequency unlock in 55 hours of iterated runs. The DD-AT PLL appears to be a good alternative to the DP-AT PLL if one is willing to trade a slight decrease in T_s for a much higher resistance to frequency unlock.

C. Variable-bandwidth PLLs and Apparent PLL Performance Limitations

Other phase tracking strategies more exotic than the example PLLs were also evaluated in the scintillation testbed. The first of these is the Kalman-Filter based PLL (KFPLL) developed in [21] and tested in [7]. The KFPLL's functionality encompasses the phase detector and loop filter blocks of Fig. 1. It achieves good performance at low C/N_0 by optimally adapting its loop bandwidth to C/N_0 and by using a hypothesis testing approach for resolving the data-bit-induced phase ambiguity. As demonstrated in [7], the KFPLL's adaptive bandwidth scheme makes it good at avoiding frequency unlock. On the other hand, the scintillation testbed showed that the KFPLL's cycle slip performance is generally worse than that of standard constant-bandwidth PLLs (like those in the example set) because, by reducing B_n during a canonical fade, the KFPLL is less able to track the rapid phase change.

A modified KFPLL was also considered, in which a simple scintillation model expressed in terms of interfering phasors is incorporated into the filter dynamics. When the modified KFPLL detects the onset of a canonical fade, it switches from a quiescent to a fading signal model. After the fade, it employs a hypothesis test to decide whether the phase change was up- or down-going. This method showed only marginal improvement in T_s compared to the much simpler DP-AT PLL from the example set. Here again, canonical fading is the problem: in deep canonical fades, the costs associated with the up- and down-going hypotheses are nearly equal, making a decision based on hypothesis testing only slightly better than random.

To date, it appears that unless a PLL is supplied with additional information, it cannot be expected to perform much better during scintillation than the example set's DP-AT PLL for typical nominal values of C/N_0 . One obvious approach to improving cycle slip performance is navigation data bit aiding. In this approach, the true navigation data bit sequence is obtained by some means and used to flawlessly wipe off the data modulation, just as was done in after-the-fact processing to generate the GPS L₁ records in the scintillation library. Slips can now occur only in full-cycle increments, meaning that the half-cycle canonical-fade phase shifts can be tracked. Moreover, the PLL now operates as a non-squaring loop; hence, its loop SNR increases by approximately 6 dB [cf. Eq. (4)], leading to a decrease in cycle slips due to measurement noise [cf. Eq. (5)]. Testbed results indicate that in severe scintillation ($S_4 > 0.6$) and at a nominal $C/N_0 \approx 43$ dB-Hz, perfect data bit wipeoff extends T_s by a rough factor 20.

VI. SCINTILLATION EFFECTS MODELING

The measure of a good model for scintillation effects on phase tracking is its ability to predict PLL behavior; specifically, to predict σ_φ^2 and T_s . Emphasis will be focused here on predicting T_s because, whereas an increase in σ_φ^2 gradually degrades performance, a cycle slip represents a sudden—and likely more consequential—upset to a carrier-phase-dependent GPS system. The scintillation effects modeling challenge can be posed as follows: Given a time history $z(t)$ of complex signal scintillation, extract the simplest statistical description sufficient to predict the expected cycle slip performance within a class of standard PLLs.

A. Various Approaches to Scintillation Effects Modeling

As a first solution to the modeling challenge, one might consider developing a cycle slip prediction model based solely on the most common scintillation statistic, namely, S_4 . A glance at Fig. 3 reveals the problem with this approach: the wide spread in cycle slips for $S_4 > 0.6$ implies that, by itself, S_4 cannot be made to explain nor accurately predict cycle slipping. Large differences in the measured slip rate at equivalent measured S_4 levels imply that there exist time series statistics relevant to phase tracking that are not captured by S_4 . In other words, the S_4 index doesn't capture the fact that the Wideband records—and some GPS L₁ records—have fast fading time scales (low decorrelation time τ_0) that make the data hard to track.

Another strategy would be to augment S_4 with a bulk phase scintillation statistic such as the variance of the true phase time history σ_θ^2 or—even more descriptive—with the spectral strength and slope parameters T and p . References [2] and [5] adopt this approach: values for S_4 , T and p are used to estimate the phase error variance σ_φ^2 of a PLL with a certain order and B_n ; σ_φ^2 , in turn, is used to predict T_s .

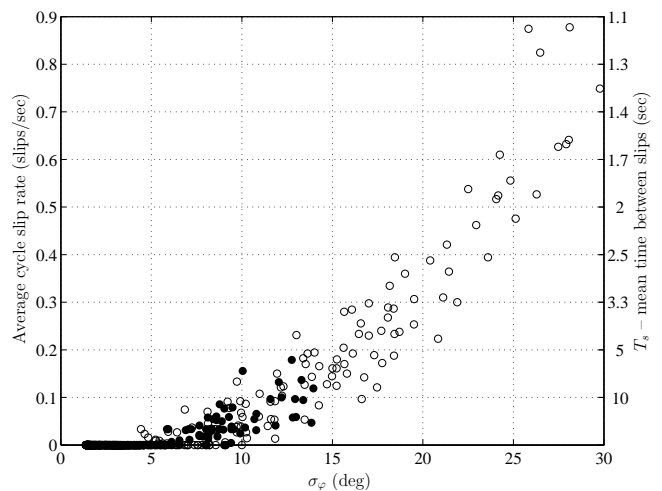


Fig. 6. Average DD-AT PLL cycle slip rate over each 30-second test interval vs. σ_φ for the Wideband data at $C/N_0 = 43$ dB-Hz (open circles) and for the GPS L₁ data within $40 \leq C/N_0 \leq 44$ dB-Hz with mean $C/N_0 = 43$ dB-Hz (filled circles). The right axis expresses the cycle slip rate in terms of T_s .

This approach seems sensible. Consider Fig. 6, which plots the average cycle slip rate vs. σ_φ over each 30-second

record in the scintillation library. Clearly, there exists a strong correlation between σ_φ and cycle slips for $\sigma_\varphi > 5$ deg. What is more, the Wideband data and the GPS data appear to follow the same trend, meaning that, unlike the S_4 index, σ_φ captures the time series statistics that make a particular record hard to track. Unfortunately, it is difficult to model σ_φ^2 as a function of S_4 , T , and p without invoking the PLL linearity assumption, which, as mentioned in Section II-A, is likely invalid for strong scintillation. Moreover, even if a good estimate of σ_φ^2 could be obtained, it is not clear how to connect σ_φ^2 and cycle slips in a way that explains the empirical relationship apparent in Fig. 6. The method adopted in [5], whereby σ_φ^2 is related to T_s on the basis of simulation results that assume Gaussian white measurement noise, leads to unreasonable predictions for T_s —especially for low values of σ_φ^2 . Using the data underlying Fig. 6, one can show that, for the DD-AT PLL, Wideband and GPS data records that lead to $8 < \sigma_\varphi < 10$ deg yield mean T_s values of 29 and 27 seconds, respectively. By comparison, the DD-AT line in Fig. 2 (the non-scintillating case) indicates that for $\sigma_\varphi < 10.2$ deg, $T_s > 5$ hours. Hence, within this range of σ_φ , the method adopted in [5] overestimates T_s by over 600 times.

A more satisfactory modeling approach can be developed around the close connection between differentially-detected data bit errors and cycle slips. Recall from the discussion of Fig. 5. and Table II that, for both the Wideband and GPS data, the mean time T_e between DPSK-detected bit errors acts as a lower bound on the best-performing (DP-AT) PLL's value of T_s . Fig. 7, which plots the DP-AT PLL's average cycle slip rate for all 30-second scintillation library records vs. the DPSK-detected bit error rate over each record, shows that the relationship between slip rates and bit error rates is strongly correlated, and that both Wideband and GPS data obey the same trend. A straight-line fit to the data falls below the $T_s = T_e$ dashed line, indicating that T_e underbounds T_s . Thus, if one can predict T_e for a given time history of complex signal scintillation, then one has a useful lower bound for the DP-AT PLL's T_s .

The connection between T_e and T_s that is apparent in Fig. 7 was explained previously in terms of the differentially-detected bit estimates $d_{m,k}$. It can now be restated more generally as follows: the signal properties that tend to cause DPSK bit errors also tend to induce cycle slips in many PLL designs. By tying DPSK-detected bit errors to PLL cycle slips, one trades the difficult problem of cycle slip prediction for the more tractable problem of bit error prediction.

B. Binary DPSK Bit Error Probability in the Presence of Scintillation

The goal of this subsection is to derive an expression for the binary DPSK bit error probability for a scintillating communications channel. This expression will enable DPSK T_e to be estimated for a given time history $z(t)$ of complex signal scintillation. The formulation proposed here extends the closed-form solution for fast Ricean fading given in [12, Sec. 8.2.5.2] to the specific case of complex signal scintillation with a 2nd-order Butterworth-type power spectrum.

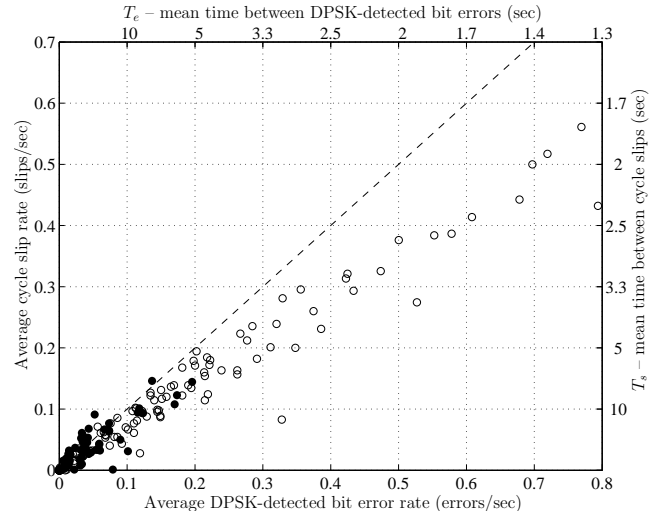


Fig. 7. Average DP-AT PLL cycle slip rate vs. the average DPSK-detected bit error rate over each 30-second test interval for the Wideband data at $C/N_0 = 43$ dB-Hz (open circles) and for the GPS L₁ data within $40 \leq C/N_0 \leq 44$ dB-Hz with mean $C/N_0 = 43$ dB-Hz (filled circles). The right and top axes express the cycle slip rate and bit error rate in terms of T_s and T_e , respectively. The dashed 45° line indicates the condition $T_s = T_e$.

Consider again the model for the complex baseband signal r_k given in Eq. (6). Assume that $T_a = T_b$ and that the phase estimate $\hat{\theta}(t)$ is generated by a low-bandwidth frequency tracking loop so that, like the carrier phase θ_c , it can be considered constant over adjacent bit intervals. For convenience, and without loss of generality, set $\theta_c = \hat{\theta}(t) = 0$ over the interval $[t_{k-2}, t_k)$. Then $r_k = z_k \sqrt{E_b} \exp(j\hat{\theta}_k) + n_k$ is the complex baseband signal, where z_k now represents the average of $z(t)$ over the k th bit interval and E_b is the energy per bit. As before, the noise element n_k is zero-mean Gaussian with $E[n_k^* n_j] = N_0 \delta_{k,j}$. Let the channel response function $z(t)$ be defined as $z(t) = \bar{z} + \xi(t)$, where \bar{z} is the direct component, modeled as a complex constant, and $\xi(t)$ is the multipath component, modeled as a complex stationary zero-mean Gaussian random process with autocorrelation function $R_\xi(\tau) = \frac{1}{2} E[\xi^*(t) \xi(t + \tau)]$. 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 scintillation channel that changes rapidly with time.

The binary DPSK decision variable is $u_k = \text{Re}(r_k^* r_{k-1})$. In the absence of scintillation and noise, $u_k = 1$ when $\hat{\theta}_k = \hat{\theta}_{k-1}$, and $u_k = -1$ when $\hat{\theta}_k = \hat{\theta}_{k-1} + \pi$. Without loss of generality, suppose that $\hat{\theta}_k = \hat{\theta}_{k-1}$. Then $P_e = P(u_k < 0)$ is the binary DPSK bit error probability. Calculation of P_e requires the following statistics of z_k :

$$\sigma^2 \equiv \frac{1}{2} E[|z_k - \bar{z}|^2] = \frac{1}{T_b^2} \int_0^{T_b} \int_{t-T_b}^t R_\xi(\tau) d\tau dt \quad (11)$$

$$\rho\sigma^2 \equiv \frac{1}{2} E[(z_{k-1} - \bar{z})^* (z_k - \bar{z})] = \frac{1}{T_b^2} \int_0^{T_b} \int_{t-2T_b}^{t-T_b} R_\xi(\tau) d\tau dt \quad (12)$$

Assume that the channel response function $z(t)$ is normalized so that $E[|z(t)|^2] = 1$. Then for saturated (Rayleigh)

fading where $\bar{z} = 0$, the quantity $2\sigma^2 \leq 1$ can be interpreted as the loss in signal power due to fluctuations in $z(t)$ over the integration interval. Likewise, $\rho \leq 1$ represents the degrading effect of bit-to-bit fluctuations in $z(t)$ on the DPSK decision variable u_k .

Under the assumption of Ricean fading, the magnitude of z_k , namely $\alpha_k = |z_k|$, has a Ricean probability distribution with mean-square value $\Omega \equiv E[\alpha_k^2] = 2\sigma^2(1 + K)$, where K is the Ricean K -parameter. Defining the average SNR per bit as $\bar{\gamma} \equiv \Omega E_b/N_0 = \Omega T_b C/N_0$, one has, together with ρ and K , all the quantities necessary to evaluate the closed-form expression for P_e given in [12, Sec. 8.2.5.2]:

$$P_e = \frac{1}{2} \left[\frac{1 + K + \bar{\gamma}(1 - \rho)}{1 + K + \bar{\gamma}} \right] \exp\left(\frac{-K\bar{\gamma}}{1 + K + \bar{\gamma}}\right) \quad (13)$$

In practice, estimation of P_e can be approached in two different ways. The first approach assumes that the averages z_k are available. In this case, ρ can be estimated directly by evaluating the expectation operations in Eqs. (11) and (12) as time averages, then dividing Eq. (12) by Eq. (11). Likewise, Ω is estimated by evaluating $\langle |z_k|^2 \rangle$. Assuming that the underlying $z(t)$ is normalized ($E[|z(t)|^2] = 1$), then $\Omega < 1$ for any $S_4 > 0$. Next, the Ricean K parameter is estimated. Under the Ricean model, K is defined as the ratio of the direct component power to the multipath component power, i.e., $K \equiv |\bar{z}|^2/2\sigma^2$. In practice, invoking this definition to estimate K is problematic because low-frequency phase scintillation causes the instantaneous \bar{z} to wander. Instead, it is best to estimate K based on its relationship with S_4 , given by

$$K = \frac{\sqrt{1 - S_4^2}}{1 - \sqrt{1 - S_4^2}}, \quad S_4 \leq 1 \quad (14)$$

where S_4 is evaluated according to Eq. (10) with $I = \alpha_k^2$. If the measured S_4 exceeds unity, K is set to 0. Substituting K , ρ , and $\bar{\gamma} = \Omega E_b/N_0$ into Eq. (13) yields the desired P_e . Tests have shown that the values of P_e calculated according to this approach agree closely with the bit error probability averages observed in the scintillation testbed for a given interval of scintillation.

The second approach, which is more appropriate for theoretical work, calculates P_e based on the τ_0 and S_4 of a hypothetical scintillation time history $z(t)$. In this case, σ^2 and ρ are estimated indirectly via a model of $R_\xi(\tau)$. The model, which closely fits the autocorrelation functions of records in the scintillation library over a broad range of scintillation strength, corresponds to the power spectrum of a 2nd-order Butterworth filter [22]:

$$R_\xi(\tau) = \sigma_\xi^2 e^{(-\beta|\tau|/\tau_0)} [\cos(\beta\tau/\tau_0) + \sin(\beta|\tau|/\tau_0)] \quad (15)$$

where the factor $\beta = 1.2396464$ ensures that $R_\xi(\tau_0)/R_\xi(0) = e^{-1}$. Substituting this expression for $R_\xi(\tau)$ into Eqs. (11) and (12) leads to

$$\sigma^2 = (\sigma_\xi^2/q^2) [2q + f(q) - 1] \quad (16)$$

$$\rho\sigma^2 = (\sigma_\xi^2/2q^2) [f(2q) - 2f(q) + 1] \quad (17)$$

where $f(q) = \exp(-q)(\cos q - \sin q)$ and $q = \beta T_b/\tau_0$.

The approach for a hypothetical $z(t)$ proceeds as follows. As before, a K value corresponding to $z(t)$ is derived from S_4 . Call this value K' to distinguish it from the K value associated with the averages z_k ; the two are in general slightly different. Again assuming $E[|z(t)|^2] = 1$, K' is related to σ_ξ^2 and \bar{z} by $2\sigma_\xi^2(1 + K') = 1$ and $\bar{z}^2 = 2K'\sigma_\xi^2$ (without loss of generality, \bar{z} can be assumed to be real). Solving for σ_ξ^2 and substituting σ_ξ^2 and τ_0 into Eqs. (16) and (17) yields values for σ^2 and $\rho\sigma^2$. These, in turn, are used to estimate the K and Ω for the hypothetical averages z_k via $K = \bar{z}^2/2\sigma^2$ and $\Omega = \bar{z}^2 + 2\sigma^2$. Finally, the K , ρ , and $\bar{\gamma} = \Omega E_b/N_0$ corresponding to the z_k are used in Eq. (13) to solve for P_e .

This approach to estimating P_e has been validated using estimates of τ_0 obtained from the scintillation library's empirical $z(t)$ intervals. Table III summarizes the results of four test cases in which the predicted and actual number of bit errors N_e for four different testbed experiments are compared. For each case, an approximately statistically stationary interval of scintillation data was chosen from the scintillation library. The nominal C/N_0 for each test case was 43 dB-Hz. The length of the interval T , the S_4 index, the value of τ_0 , and the predicted and actual number of bit errors N_e are given for each case. The error range on the actual N_e values represents two standard deviations based on 10 runs with the same scintillation record but with independent simulated receiver noise. In general, the agreement between the predicted and actual N_e values is good, meaning that the model described in this section accurately predicts binary DPSK P_e during severe scintillation.

TABLE III
COMPARISON OF PREDICTED AND ACTUAL DPSK BIT ERRORS

Case	Parameters			N_e (pred.)	N_e (actual)
	T (s)	S_4	τ_0 (s)		
1	100	0.41	1.30	0.002	0 ± 0
2	114	0.97	0.25	31.4	30.4 ± 2.3
3	108	0.88	0.21	16.2	20.6 ± 3.3
4	127	0.97	0.36	18.4	20.7 ± 3.1

Once P_e is obtained by whatever means, one calculates a useful lower bound for the achievable mean time between cycle slips T_s over the interval of scintillation considered by computing the mean time between DPSK-detected bit errors $T_e = T_b/P_e$. Thus, in response to the scintillation effects modeling challenge posed in the introduction to Section VI, one can offer the triple $\{S_4, \tau_0, C/N_0\}$ (or $\{K, \rho, \bar{\gamma}\}$, which contains equivalent information) as a scintillation channel characterization sufficient to approximately predict the cycle slip performance of a class of standard GPS PLLs. In this characterization, S_4 gives a measure of the scintillation intensity, τ_0 describes the speed of the fluctuations, and C/N_0 defines the channel's nominal carrier-to-noise ratio.

C. Applicable Domain of the Scintillation Effects Model

It is necessary to outline the range of parameter values over which the proposed scintillation effects model is applicable. First note that the method presented above for estimating T_e imposes no restrictions on the range of τ_0 or C/N_0 , and that

the condition $S_4 \leq 1$ implied by $K \geq 0$ is not significantly restrictive because, aside from rare “focusing” behavior, S_4 takes on values near or below unity [23]. Thus, T_e can be estimated as described above for nearly all possible S_4 , τ_0 , and C/N_0 values.

On the other hand, the model of T_e as a useful approximation for T_s is only valid within a limited range of τ_0 , C/N_0 , T_a , and B_n values. For example, in the limit as nominal $C/N_0 \rightarrow \infty$, one can always reduce cycle slips by shortening T_a and widening B_n . In this extreme case, T_e and T_s are no longer usefully related. Of course, not all parameter values are of practical interest. For a terrestrial GPS receiver, values of nominal C/N_0 range from 38 to 55 dB-Hz. Within this range and in the presence of severe scintillation, widening B_n much beyond 10 Hz results in a steep increase in frequency unlock. Likewise, there exist practical limitations on the accumulation interval T_a : short T_a lead to increased squaring loss [cf. Eq. (2)], and long T_a tend to destabilize the tracking loop. Finally, realistic decorrelation times of L-band scintillation for a stationary receiver can be generally limited to $1 > \tau_0 > 0.2$ seconds. The following list summarizes practical ranges of τ_0 , C/N_0 , T_a , and B_n for GPS carrier phase tracking in the presence of scintillation:

$$\begin{aligned} 1 > \tau_0 > 0.2 \text{ sec} & \quad 38 \leq C/N_0 \leq 55 \text{ dB} - \text{Hz} \\ 10 \leq T_a \leq 20 \text{ ms} & \quad 5 \leq B_n \leq 20 \text{ Hz} \end{aligned}$$

Simulation tests have shown that the proposed model is applicable over this domain in the sense that on average $T_s/2 < T_e \leq T_s$.

VII. CONCLUSIONS

A scintillation effects testbed that derives inputs from an empirical scintillation library has been used to test several standard—and some exotic—GPS phase tracking loop designs. Testbed results demonstrate that deep power fades (> 15 dB) accompanied by abrupt, approximately half-cycle phase transitions—jointly termed canonical fades—cause the loops to skip cycles at a rate that is significantly faster than is predicted by existing scintillation effects models.

Testbed results also show that the mean time between binary DPSK bit errors serves as a useful lower bound for the mean time between cycle slips experienced by a prototype PLL that uses differentially-detected data bit wipeoff. It has been argued that the prototype PLL is broadly representative of all PLLs that track bi-phase modulated signals in the sense that, without additional information (e.g., data bit aiding), other phase tracking schemes are unlikely to perform significantly better in severe scintillation than does the prototype PLL.

Inspired by the strong correlation between binary DPSK bit errors and cycle slips, a simple model for cycle slip prediction has been developed. The model, which reasonably approximates the rate of cycle slipping observed in the testbed results, is based on the scintillation index S_4 , the complex scintillation signal decorrelation time τ_0 , and the nominal carrier-to-noise ratio C/N_0 .

ACKNOWLEDGMENTS

Thanks to A. Kumar from Cornell’s department of Theoretical and Applied Mechanics for useful discussions on trapping regions in discrete-time nonlinear PLL dynamics.

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*, vol. 46, no. 4, Oct. 2010.
- [2] M. Knight and A. Finn, “The effects of ionospheric scintillation on GPS,” in *Proceedings of ION GPS 1998*. Nashville, TN: Institute of Navigation, 1998.
- [3] 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.
- [4] 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.
- [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] 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.
- [8] A. J. Viterbi, *Principles of Coherent Communication*. New York: McGraw-Hill, 1966.
- [9] W. C. Lindsey and M. K. Simon, *Telecommunication Systems Engineering*. New Jersey: Prentice-Hall, 1973.
- [10] F. M. Gardner, *Phaselock Techniques*, 3rd ed. Hoboken, NJ: Wiley, 2005.
- [11] S. Gupta, “Phase-locked loops,” *Proc. IEEE*, vol. 63, no. 2, pp. 291–306, 1975.
- [12] M. K. Simon and M. Alouini, *Digital Communications over Fading Channels*. New York: Wiley, 2000.
- [13] G. Ascheid and H. Meyr, “Cycle slips in phase-locked loops: A tutorial survey,” *IEEE Transactions on Communications*, vol. COM-30, no. 10, pp. 2228–2241, Oct. 1982.
- [14] 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.
- [15] 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.
- [16] J. G. Proakis and M. Salehi, *Communication Systems Engineering*. Upper Saddle River, NJ: Prentice Hall, 1994.
- [17] P. Ward, *Understanding GPS: Principles and Applications*. Boston: Artech House, 1996, ch. 5: Satellite Signal Acquisition and Tracking, pp. 119–208.
- [18] A. J. Van Dierendonck, *Global Positioning System: Theory and Applications*. Washington, D.C.: American Institute of Aeronautics and Astronautics, 1996, ch. 8: GPS Receivers, pp. 329–407.
- [19] 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.
- [20] 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.
- [21] M. L. Psiaki and H. Jung, “Extended Kalman filter methods for tracking weak GPS signals,” in *Proceedings of ION GPS 2002*. Portland, Oregon: Institute of Navigation, 2002, pp. 2539–2553.
- [22] 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.

- [23] 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.