Empirical Characterization and Modeling of GPS Positioning Errors Due to Ionospheric Scintillation

Charles S. Carrano[1], Keith M. Groves[2], James M. Griffin[1]

1) Radex, Inc., 131 Hartwell Avenue, Lexington, MA 02421
2) Air Force Research Laboratory, Space Vehicles Directorate, Hanscom AFB, MA 01731

Abstract

Ionospheric scintillation degrades GPS positioning accuracy in a number of ways. Along each scintillating satellite-receiver link, errors are introduced in the measurements of pseudorange and carrier phase. When the scintillation on a given link is sufficiently intense, the link becomes intermittently unavailable for use in the position solution. The loss of each link leads to an increased dilution of precision as the geometry of the overall satellite constellation is degraded. These effects reduce the accuracy of the computed position and result in temporary losses of positioning service when fewer than four links remain available at any given time. The duration of these positioning service "outages" depends on the duration and severity of the scintillation event, the geometry of the satellites in view, and the recovery time of the equipment.

In this work we characterize and model the effects of scintillation on GPS positioning accuracy using data acquired with an Ashtech Z-12 receiver at Ascension Island. This campaign was conducted by the Air Force Research Laboratory (AFRL) during solar maximum conditions in March of 2002. Positioning errors exceeding 75 meters in the horizontal and 150 meters in the vertical were routinely encountered, and positioning outages occurred nearly every evening. We demonstrate the effects of scintillation on the pseudoranges, and model the propagation of these errors into the position solution in terms of an empirical parametrization of the scintillation environment. A key feature of the model is its ability to simulate link-by-link measurement errors and the effective degradation in the constellation geometry, both individually and in concert. Results generated by the model are presented and shown to compare favorably with actual measurements of the positioning error.

The parameters of the model may be tuned for different GPS receiver designs, either through field testing as it was done in this study or through hardware-in-the-loop simulation using tools such as the AFRL Antenna WaveFront Simulator (AWFS). Coupled with a regional specification of the scintillation environment, based either on data or climatology, the model may be used to provide regional GPS position error maps suitable for integration into space weather forecast products.

ACKNOWLEDGEMENTS

The work at Radex, Inc. was supported by AFRL contract F19628-00-0089.

INTRODUCTION

Ionospheric scintillation describes the rapid fluctuations in phase and/or amplitude of radio waves that penetrate the ionosphere caused by electron density irregularities along the signal propagation path. The scintillation of GPS signals occurs most often in the equatorial region of the globe at night and, to a lesser extent, also in the auroral and polar regions. The effects of the ionosphere on GPS are most pronounced during the peaks of the solar cycle [*Groves, et al.* 2000].

Ionospheric scintillation affects GPS receivers in several ways. Amplitude fluctuations manifest themselves as intermittent signal fades and enhancements which result in errors decoding the GPS data messages and also in estimating the ranges. Phase fluctuations stress the ability of the receiver to maintain lock and cause "cycle slips" or breaks in the measured phase that may prevent the receiver

from using the phase to refine its range measurements. When the receiver is unable to maintain lock on at least four or more GPS satellites, a temporary loss of positioning service results. The duration of these service "outages" depends on the duration and severity of the disturbances, the geometry of the satellites in view, and the recovery time of the equipment. The likelihood for a scintillation induced outage to occur can be increased by external factors such as an unfavorable geometry due to a partially obscured sky view.

In this paper, we focus on the effect of ionospheric scintillation on GPS positioning accuracy, as observed during a campaign conducted by Air Force Research Laboratory at Ascension Island March 5-19, 2002 [*Groves et al.*, 2000]. Several GPS receivers were operated during this campaign, but for the sake of brevity we consider only data from a single Ashtech Z-12 receiver in this study. The monthly average sunspot number was 98.4, representative of solar maximum conditions. After characterizing the observed effects of scintillation on satellite availability, range estimation, and overall positioning accuracy, we present an empirical model for simulating GPS positioning errors given a specification of the scintillation environment.

SCINTILLATION EFFECTS ON SATELLITE AVAILABILITY

One of the most important ways in which ionospheric scintillation degrades the performance of GPS is by effectively reducing the number of satellite links available for use in computing the receiver position. The rapid signal fades and enhancements that are characteristic of scintillation are caused by the destructive and constructive interference of the radio wave fronts as they scatter from electron density irregularities in the ionosphere. In the case of GPS these signal fades lead to "intermittent availability," whereby the receiver repeatedly loses and reacquires lock on the link, rendering it temporarily unavailable for use in the position solution.

All but the most inexpensive GPS receivers report which satellites they *track* while operating. The Ashtech Z-12 receiver, in addition, records which of the tracked satellites are actually *used* to compute the receiver position. Figure 1 illustrates the effects of intermittent availability caused by scintillation. Shown in the figure is the time history of the number of satellites used to compute the position solution, the number of scintillating satellites with $S_4 \geq 0.3$, and the Circular Error Probable (CEP) associated with the position solution. An azimuth-dependent elevation mask was used to minimize the influence of multipath. The number of scintillating satellites at any given instant was evaluated by counting the tracked satellites with $S_4 \geq 0.3$, together with any untracked satellites that should have been visible above the mask according to ephemeris. The assumption here is that these untracked satellites have been lost due to scintillation and thus should be included in the count.

The data shown in Figure 1 was recorded on March 16, 2002. Prior to approximately 20:25 UT that evening, no scintillation had been measured and the CEP did not exceed a few meters, even at the 95th percentile. Soon afterward, one link after another began to experience scintillation until just before 22:00 UT when all the satellites visible above the elevation mask were scintillating simultaneously. Such extreme conditions, where the signals from all visible GPS satellites experienced scintillation in unison was not uncommon during this campaign. It is clear from the figure that as more satellite links scintillated, the CEP generally increased. Note the rapid fluctuations in the number of satellites used to compute the position solution while scintillation was occurring. These rapid fluctuations are what is meant by intermittent availability due to scintillation. Starting at approximately 21:45 UT and lasting for several minutes, the number of available satellites repeatedly dropped below four. These are GPS service outages during which no position is reported by the receiver.

Figure 2 shows the time history of the position dilution of precision (PDOP) as well as the horizontal and vertical position errors for this same evening. These position errors were evaluated using the receiver-reported position; we did not compute the position ourselves in post-processing. During scintillation, PDOP exhibits brief but large spikes as satellites are intermittently dropped and

reacquired. The timescale of these large fluctuations is on the order of a few seconds. The temporal structure of the horizontal and vertical position errors is characterized by a slowly varying component with superimposed spikes. A careful examination of the data suggests that the rapid fluctuations are primarily due to intermittent availability (satellites being dropped and reacquired), while the slowly varying component is primarily due to ranging errors on the individual links.

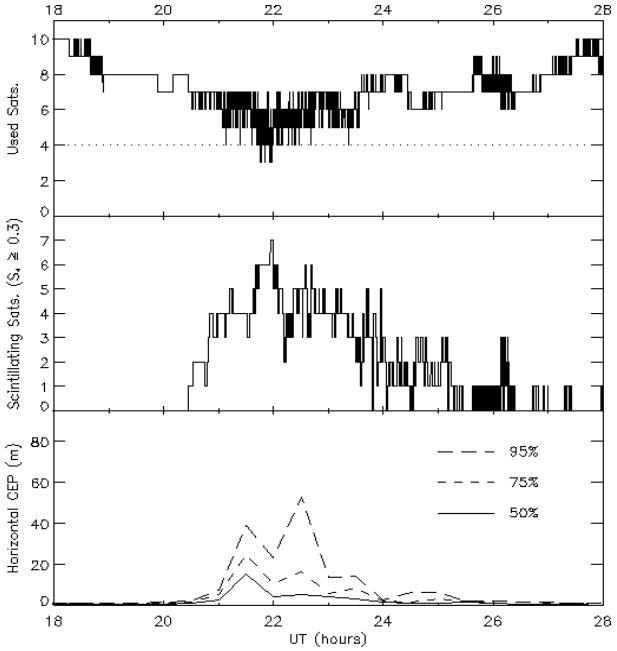

Figure 1. Time versus number of satellites used in the position solution (top), the number of scintillating satellites with $S_4 \geq 0.3$ (middle) and the horizontal CEP at the 50th, 75th and 95th percentiles (bottom) , based on a 30-minute moving window.

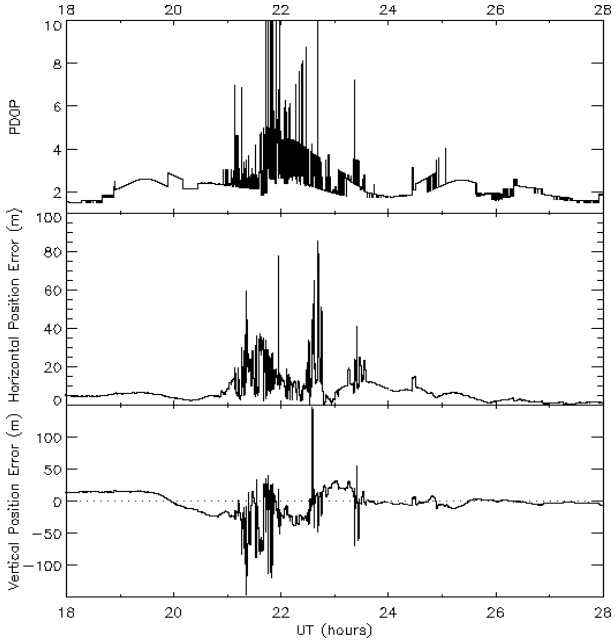

Figure 2. Time versus position dilution of precision (top), horizontal position error (middle), and vertical position error (bottom).

A Statistical Analysis of Satellite Availability

Generally speaking, GPS satellite availability decreases as the scintillation intensity increases, although the relationship between the two is complex and best described statistically. In Figure 3 we present a histogram showing the availability likelihood as a function of the scintillation intensity, S_4. Here we have defined the availability likelihood as the percentage of satellites used in the position solution relative to the total number of satellites visible above our elevation mask. The plot includes data between 20-28 UT for all of the days of the campaign. Note the rapid fall-off in availability likelihood as an S_4 of unity is approached. This may occur because the scintillation index poorly characterizes the true intensity of the turbulence in this regime due to

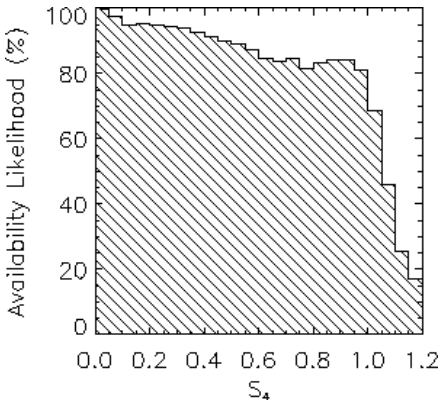

Figure 3. Histogram showing the likelihood of a link being included in the position solution versus the scintillation intensity on that link.

saturation effects. In general, it is expected that the availability likelihood distribution will depend on the hardware and software design of the particular receiver being operated.

SCINTILLATION EFFECTS ON THE PSEUDORANGES

As more satellite links scintillate, causing them to be intermittently dropped and reacquired by the receiver, the overall dilution of precision increases. As one might expect, this increase in dilution of precision generally results in increased positioning errors. Experimental evidence, on the other hand, suggests that positioning errors can also be large even when the dilution of precision is small.

Figure 4 shows a plot of horizontal dilution of precision (HDOP) versus horizontal position error for the night of March 16, 2002. The points are colored according to the number of links that experienced scintillation with $S_4 \geq 0.3$. Notice that while the HDOP and horizontal position errors often increase together (as more links scintillate), many of the largest errors actually occur when HDOP is low. Large position errors that occur despite a favorable viewing geometry (low DOP) are the result of the ranging errors on the individual links, rather than intermittent availability.

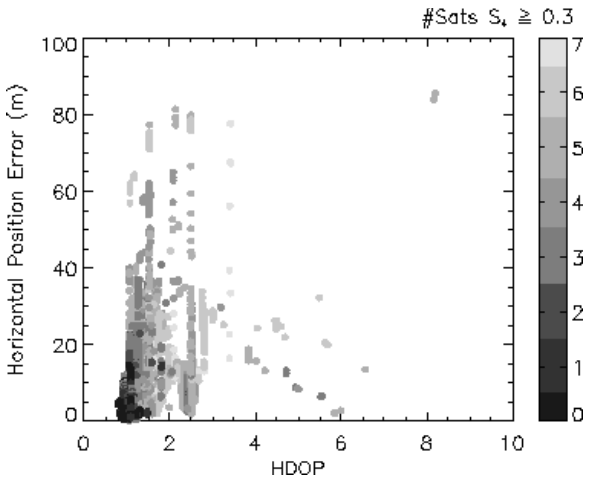

Figure 4. Horizontal position error versus horizontal dilution of precision. The points are colored by the number of scintillating satellite links with $S_4 \geq 0.3$ at that time.

Detrending the Pseudoranges

Before one can examine the effects of scintillation on the pseudoranges, the dominant contributions due to the geometric range to the satellite and the clock biases must be removed. While this detrending may be accomplished by differencing or least squares techniques, we do not do so here because these methods combine the ranging errors from multiple satellite links in an indistinguishable way. Similarly, we do not form the two frequency difference for scintillating links because we want to investigate the errors on each frequency separately.

For each satellite, k, we model the pseudorange measurement on L1 at the receiver's time of reception as:

$$P_1^k = \rho^k(t, t^k) + c\,\tau^k - c\,\tau + S^k + I^k + T^k + E^k. \qquad [1]$$

In equation [1] above, t represents the true time of reception while t^k represents the true time of reception. The term $\rho^k(t, t^k)$ represents the range from the satellite at the true time of transmission to the receiver at the true time of reception, τ^k is the satellite clock bias, τ is the receiver clock bias, and c is the speed of light in a vacuum. The term S^k represents scintillation induced ranging errors, I^k is the ionospheric delay, T^k is the tropospheric delay, and E^k represents unmodeled errors including multipath, hardware delays, ephemeris errors, and thermal noise. This formulation is novel only in the explicit addition of the scintillation induced ranging error term. An analogous expression exists for the pseudorange on L2.

To investigate the variation of the scintillation induced ranging error, S^k, we detrend the measured pseudorange by estimating and subtracting, or otherwise neglecting, all the other terms in [1]. We do not estimate the terms T^k and E^k but instead attempt to minimize them by the application of an elevation mask. The geometric range, $\rho^k(t, t^k)$, is expanded in a Taylor series about the time of reception [*Blewitt*, 1999] and evaluated using precise ephemeris [*Remondi*, 1991] together with the

known receiver position. In principle, the motion of the Earth during the signal transit time should be accounted for in the geometric range, but the effect was neglected here as the error incurred (about 60 meters) is small compared to the scintillation induced ranging errors to be examined. The satellite clock bias, τ^k, is also provided by precise ephemeris. The receiver clock bias, τ, is common to all satellite links and is estimated here using the pseudorange measurement from a quiet (non-scintillating) link, n, as:

$$\tau = [\, P_1^n - \rho^n(t, t^n) + \tau^n + I^n \,]/c. \qquad [2]$$

The ionospheric delay on the quiet link, I^n, is evaluated as follows:

$$I^n = (P_2^n - P_1^n)/(\gamma - 1), \qquad [3]$$

where $\gamma = (f_1/f_2)^2$ is the squared ratio of the L1 and L2 frequencies. In practice, the pseudoranges are detrended in piecewise segments of time, and for each segment one link serves as the quiet link to estimate the receiver clock bias. In this way several different quiet links may be used to detrend the pseudoranges over extended periods of time.

Once the receiver clock bias been determined, the pseudorange on a given link may be detrended by subtracting the geometrical range and the receiver and clock biases, leaving only the contributions from S^k and I^k. As the maximum value for the ionospheric delay is physically limited to about 100 meters, variations in the detrended range greater than this are attributable to scintillation. Due to the relatively slow variation of the ionospheric delay and Earth rotation effects, the these effects may be largely removed in practice by subtracting the median value of the detrended range over piecewise segments.

Observations of Scintillation Induced Ranging Errors

The signal-to-noise ratio and the detrended L1 and L2 pseudoranges for the night of March 16, 2002 are shown in Figure 5. To save space, only four satellites are shown although more than this were tracked by the receiver during this period. The ranging errors appear as intermittent kilometer-scale excursions from zero that are correlated with the deep fades in the signal. Note that, qualitatively, the range error excursions on L1 and L2 tend to occur simultaneously and in the same direction. However, forming dispersive and non-dispersive linear combinations of P_1 and P_2 (not shown) reveals that, both components are actually important. This is in stark contrast to the ionospheric delay, which is a wholly dispersive phenomenon. That the scintillation induced ranging errors have a non-dispersive component should not be surprising, since they are errors made by the receiver in response to the ionospheric turbulence, as opposed to measures of some physical property of the ionosphere itself. Speculations as to the cause of these ranging errors might include incorrect decoding of the transmit time from the navigation message or signal distortion leading to secondary correlation peaks in the receiver's code tracking loop. It is clear, in any case, that for a substantial fraction of the time that scintillation was occurring the receiver operated with its code lock loop broken and was therefore forced to depend on its internal clock. More work will be needed to properly identify the mechanisms that produce scintillation induced ranging errors. Despite the fact that the term S^k may lack a physical significance, however, it is clear that ionosphere can have two impacts on the measured pseudoranges —it introduces a dispersive delay which may be corrected using the dual frequency approach, and can also introduce non-dispersive ranging errors caused by scintillation which cannot be corrected in this way.

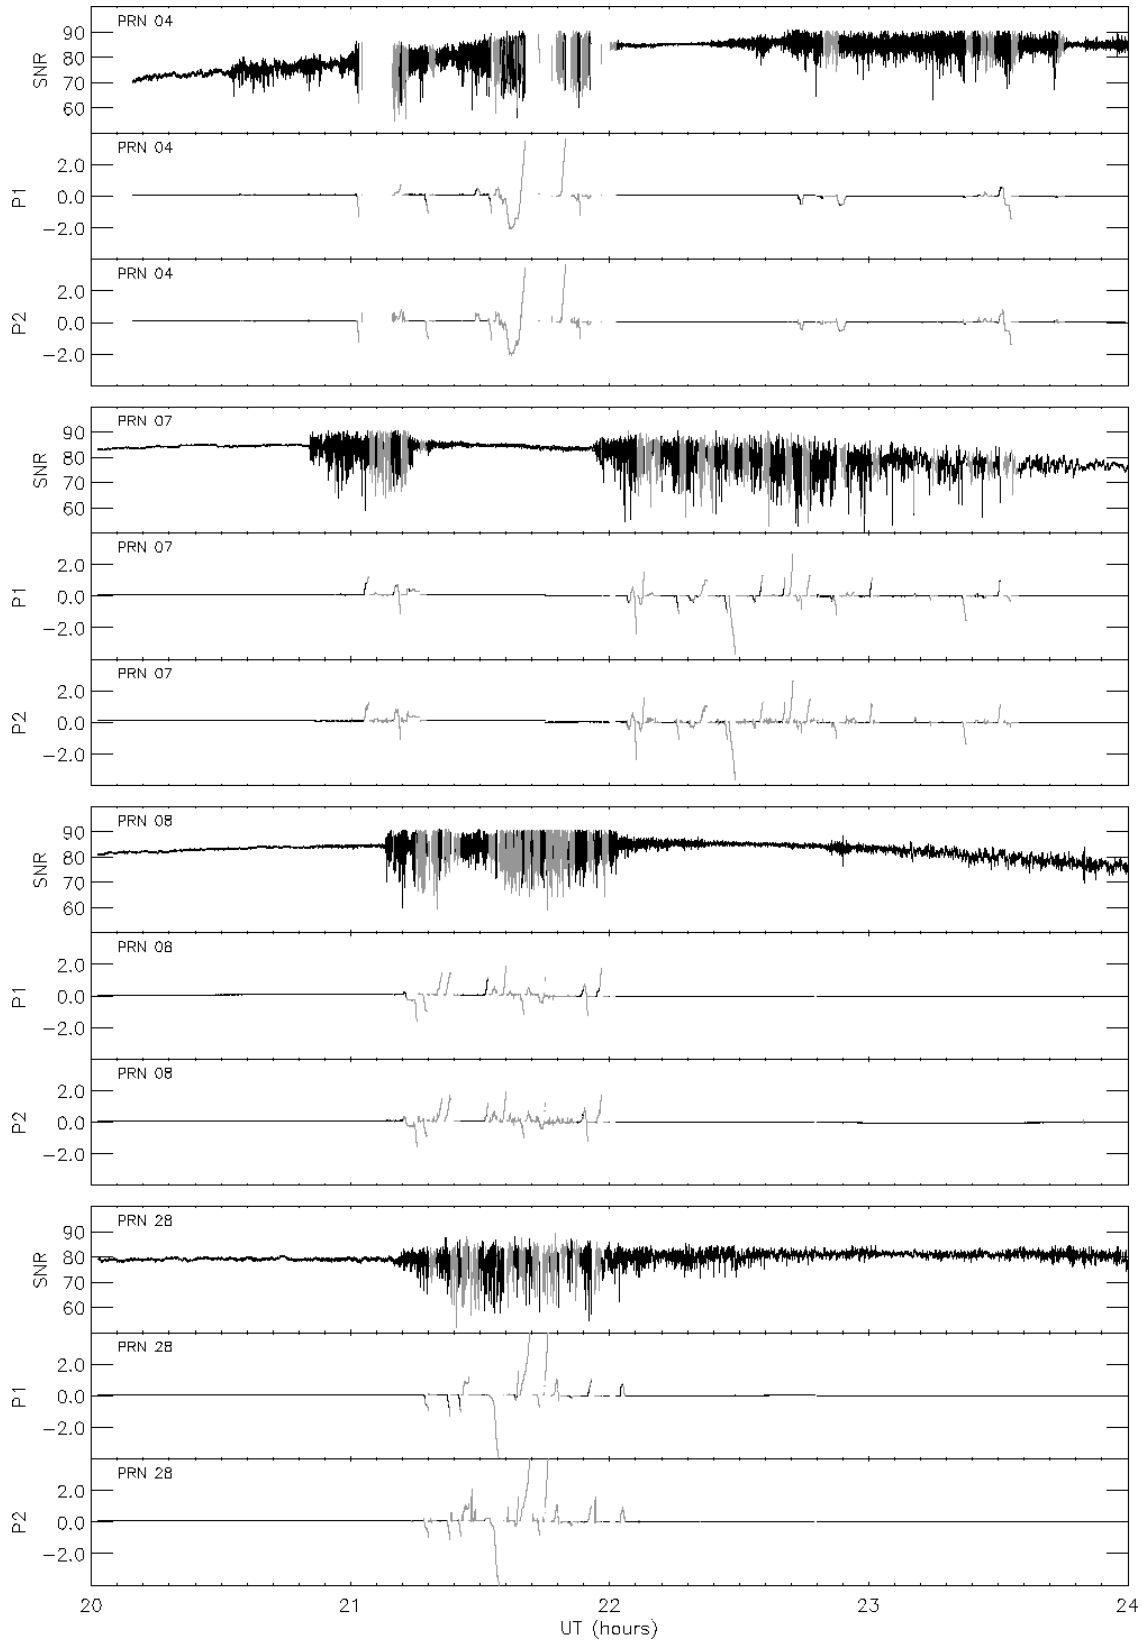

Figure 5. Signal-to-noise ratio on the L1 carrier (dB-Hz), and detrended L1 & L2 pseudoranges (km). The excursions from zero in the detrended pseudoranges are principally due to scintillation induced ranging errors. Black and gray indicate nominal and anomalous tracking status, respectively.

The Ashtech Z-12 also reports a tracking status byte that indicates the occurrence of an irregularity in the processing of the signal on that link. Such irregularities include questionable measurement of the code or carrier phase, an unsettled code phase loop, a possible cycle slip, and a possible loss of lock since the last epoch. The coloring of the plots in Figure 5 is such that black indicates nominal status and gray indicates that at least one such irregularity was detected by the receiver. Note that while the ranging errors can be quite large (up to 8 km errors were observed), the largest errors are accompanied by a warning that the receiver may use for the purpose of quality control.

To quantify the statistical behavior of scintillation induced ranging errors, we devised an automated event detection algorithm to extract and record each excursion. According to this algorithm, each excursion in the detrended pseudorange that exceeds 100 meters in absolute value defines an event. Once an event has been found, its temporal extent is determined by tracing the detrended pseudorange backward and forward in time until it remains less than 100 meters for at least 15 seconds, or until a data dropout occurs. We measured the duration and maximum range extent for each event, first while including all samples and then again while including only those samples with nominal tracking status. These parameters are shown plotted against the maximum (peak-to-peak) signal fade in Figure 6. The figures include data for all days of the campaign. A twenty degree elevation mask was employed to minimize the effects of multipath.

Note that while the largest excursions occur during deep fades of the signal, the relationship is not linear. In particular, there appears to be a threshold fade of about 10 dB (shown with a dotted line in the figure) beyond which the amplitude of the ranging error increases dramatically. This threshold is presumably the fade beyond which the receiver's code lock loop is critically stressed. These plots were also generated using S_4 rather than maximum fade (not shown) as the dependent variable. It appears that scintillation induced ranging errors correlate slightly better with maximum fade than with S_4, and the event duration may correlate with fade depth, but more data would be needed to confirm these observations. The median event duration throughout the campaign was approximately 40 seconds, although this number may be strongly influenced by the design our event extraction algorithm.

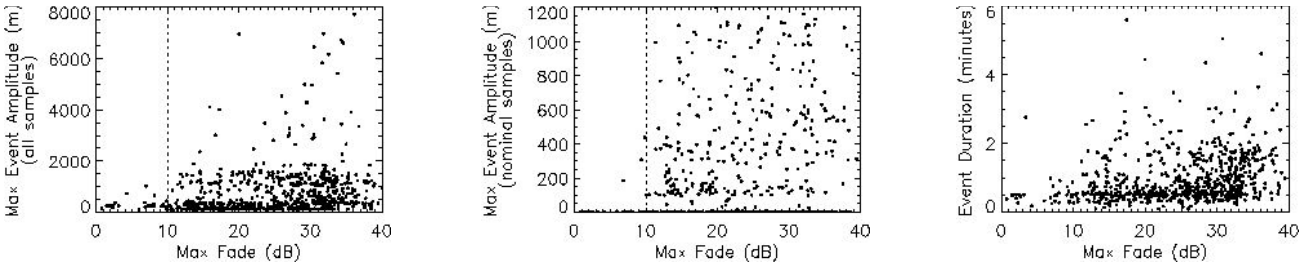

Figure 6. Fade depth of the SNR on the L1 carrier versus the following: maximum absolute value of range error events including all samples (left), maximum absolute value of range error events including only samples with nominal tracking status (middle), and the event duration (right).

MODELING GPS POSITION ERRORS DUE TO SCINTILLATION

The modifications of the signal structure caused by ionospheric scintillation are quite complex and cannot be modeled or mitigated as straightforwardly as the ionospheric delay. This is because the spatial distribution of the turbulent irregularities along the propagation path is unknown. Hence, we employ a simple empirical model that attempts to relate positioning errors to the scintillation intensity on each satellite link. The model requires the location of the receiver to be known, but requires no actual data from the receiver other than the scintillation intensity index on each link. The complete model consists of two parts: a tracking model and a ranging error model.

Given the scintillation intensity for a satellite in view, the tracking model predicts whether or not

the link will be included in the position solution. Our *tracking model* may be expressed as the rule:

$$\text{if } x > \alpha f(S_4), \text{ then exclude this link.} \quad [4]$$

In the above, the function $f(S_4)$ is the empirically determined availability function plotted in Figure 3, α is a dimensionless weighting factor, and x is a uniformly distributed random number between zero and one.

Next, we model the scintillation induced ranging error as a random perturbation whose amplitude depends linearly on the scintillation intensity. The *ranging error model* may be expressed as:

$$S = (x - \tfrac{1}{2})\, \beta\, S_4, \quad [5]$$

where β is a scaling factor with units of meters. We note that this functional form has been chosen for simplicity, but it is clearly oversimplified in light of the data shown in Figure 6.

Given the receiver position, we evaluate the L1 pseudorange for each link according to equation [1]. The geometric range to the satellite is estimated using precise ephemeris, as described earlier. We employ the simplifying assumption that I^k, T^k, and E^k are zero for each link so that only the effects of scintillation are modeled. Without loss of generality, we set the satellite and receiver clock biases to zero (although we will still solve for the receiver clock bias as an unknown). Once we have evaluated the modeled pseudoranges for those satellites that the tracking model indicates should be included in the solution, we solve for the receiver position using standard techniques [*Blewitt*, 1999]. The position error is then computed by comparing the results to the assumed receiver position.

The free parameters in the model, α and β, as well as the availability function are intended to be determined experimentally for each receiver type. Figure 7 shows the results of the simulation for the night of March 16, 2002. The weighting parameters were chosen as α=1.06 and β=60 meters, in order to most closely reproduce the observations for this night (shown in Figure 2). The availability function used was the one shown in Figure 3. A ten degree elevation mask was employed.

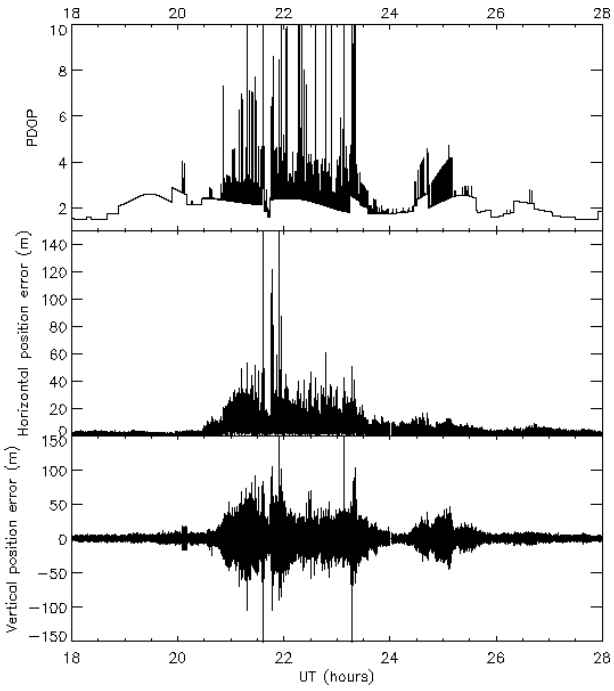

Figure 7. Simulated PDOP, simulated horizontal position error and simulated vertical position error versus time.

The structure of PDOP from the simulation (Figure 7) is structurally similar to that which was measured (Figure 2), indicating that our tracking model performed well. The position errors for the simulation are largest during scintillation and the envelope of the errors show reasonable qualitative agreement with the observations in Figure 2. The large spikes occur when the constellation geometry changes due to scintillation induced loss of lock on one or more satellites. Numerical experimentation conducted with the model suggests that both intermittent availability and link-by-link ranging errors are important effects. The model provides a unique means for exploring these two effects, both individually and in concert.

The rapid fluctuations within the envelope of the simulated position errors occur because our random numbers (x) are uncorrelated from one sample to the next, whereas the actual ranging errors

are correlated over many samples (about 40 seconds, on average). We intentionally run the simulation at the faster rate of one sample per second in order to increase the number of realizations available for computing statistics (such as the CEP) from the output.

It is difficult to reconcile why a ranging error perturbation of only ±30 meters is needed to produce position errors similar to the observations when the observations themselves contain kilometer scale ranging errors. This discrepancy may suggest that the receiver is able to minimize the effects of large ranging errors using filtering techniques or other methods of quality control. Our position error model does not include the stabilizing effects of filtering. Recently, we have begun to detrend the pseudorange data from other GPS receiver models that were operated during this campaign. While all of the receivers exhibited position errors similar in magnitude to those of the Ashtech Z-12 receiver presented here, not all of the receivers exhibited ranging errors this large. In fact, a preliminary analysis suggests that some receiver models indeed show ranging errors on the order of the 60 meter perturbation used in the simulation. Much new work needs to be done to resolve this issue. In addition, a more realistic functional form for the ranging error perturbation should be developed for use with the model.

CONCLUSIONS

Ionospheric scintillation contributes to GPS positioning errors in at least two ways. It causes intermittent availability which degrades the effective GPS constellation geometry, and it introduces ranging errors along each scintillating satellite-receiver link. A statistical analysis of data from Ascension Island in 2002 leads to the conclusion that both effects are important. An interesting result from the analysis is that the effects of scintillation are fundamentally different than those due to ionospheric delay. In particular, scintillation-induced ranging errors have a non-dispersive component that can not be corrected using the dual frequency approach.

A new GPS error model is presented that couples the effects of scintillation on each satellite, together with the geometrical influences of their individual locations within the available constellation. The parameters of the model are simple and can be tuned for different receiver types. The model can simulate positioning errors and receiver outages from a single GPS receiver, given only the location of the receiver and the scintillation intensity on each link. If coupled with a regional specification of scintillation, the model could be used to generate regional GPS error maps suitable for integration into space weather forecast products. While the analysis and simulation results presented here are specific to a particular receiver type, the methodology itself is general and may be readily appied to other receiver models. Efforts to better understand the effects of scintillation on GPS positioning accuracy are in currently underway. These include simulation and verification via hardware-in-the-loop experiments conducted with the AFRL Antenna WaveFront Simulatior (AWFS).

REFERENCES

Blewitt, G, “Basics of the GPS Technique: Observation Equations,” in Geodetic Applications of GPS, p. 10-54, ed. B. Johnson, Nordic Geodetic Commission, Sweden, ISSN 0280-5731, 1997.

Groves K., Basu, S., Quinn, J., Pedersen, T., Falinkski, K., Brown, A., Silva, R., and Ning, P., “A Comparison of GPS Performance in a Scintillation Environment at Ascension Island,” Proceedings of ION GPS 2000, Salt Lake City, UT, September 2000.

Remondi, B. W., “NGS Second Generation ASCII and Binary Orbit Formats and Associated Interpolation Studies,” Proceedings of the Twentieth General Assembly, International Union of Geodesy and Geophysics, Vienna, Austria, August 11-24, 1991.