A Comprehensive Ionosphere Storm Data Analysis Method to Support LAAS Threat Model Development

Alexandru Ene, Di Qiu, Ming Luo, Sam Pullen, and Per Enge

Department of Aeronautics and Astronautics, Stanford University

ABSTRACT

Several recent papers have analyzed the potential threat of severe ionosphere spatial gradients on LAAS users [1-4]. To support this work, a comprehensive methodology has been developed to analyze WAAS Supertruth data as well as both raw and JPL-processed data from the IGS/CORS receiver network to search for anomalous gradients. Anomalous gradients that result from this method are used to populate and validate the LAAS ionosphere spatial gradient "threat model" that is used to determine the potential impact of ionosphere anomalies on LAAS users.

This paper explains the ionosphere data analysis method in detail. This method includes two separate chains of analysis. One chain uses dual-frequency (L1 and L2) data that is post-processed by JPL to remove receiver inter-frequency biases as well as obvious receiver glitches. A search for the largest ionosphere spatial gradients (in terms of the apparent gradient "slope" in mm/km) is then performed, and the resulting observations are put through a screening process that confirms that the observed gradients are due to actual ionosphere gradients as opposed to other causes (e.g., incorrect L1-L2 bias estimates; unexplained measurement "jumps"). Surviving observations from this chain are then validated by an analysis of ionosphere rate of change using L1 code-minus-carrier only to insure that the apparent gradients are not due to receiver errors after recovery from loss of lock on L2. Observations confirmed to be due to ionosphere gradients are then analyzed further (using additional nearby receivers) to estimate the gradient width and propagation speed to the extent possible.

This paper presents results for analyses conducted using this method for the October 29 - 30 and November 20, 2003 ionosphere storms over CONUS. These days appear to represent the worst ionosphere storms experienced in CONUS since WAAS and IGS/CORS data became available. Preliminary results for several earlier storms in CONUS that might also be of concern to LAAS are also presented.

1. INTRODUCTION

The ionosphere is a region of the atmosphere located between about 50 – 1000 km above the Earth's surface. In this region, solar radiation produces free electrons and ions that cause phase advance and group delay to radio waves [6]. If uncorrected, the error introduced by the ionosphere into the GPS signal can be as high as tens of meters. Ionosphere delays on GPS pseudorange measurements normally are very highly correlated over short distances (to within 2 – 3 mm/km, one sigma). As a result, local-area differential GPS systems such as the Local Area Augmentation System (LAAS) (see [LAAS refs]) remove almost all ionosphere delay errors under nominal conditions. However, unusual solar-geomagnetic events such as Coronal Mass Ejections (CME) from the Sun can cause the ionosphere to behave anomalously, as was observed during the April 2000, October 2003, and November 2003 ionosphere storms [2-5].

The original motivation for this study is the ionosphere anomaly observed from Wide Area Augmentation System (WAAS) "Supertruth" data on April, 6, 2000. An apparent ionosphere "front" was observed running roughly East-West and produced a large gradient (5 – 8 m) in vertical ionosphere delays over a short distance (20 – 100 km). The resulting gradient is two orders of magnitude higher than the typical one-sigma ionosphere vertical gradient value identified previously [7]. Several recent papers have analyzed the potential threat of severe ionosphere spatial gradients on LAAS users [1, 2]. These papers hypothesized the possible extent of ionosphere anomalies but did not have a complete database of the known ionosphere events to support their assumptions

A comprehensive methodology for ionosphere anomaly data assessment was developed to automatically search for anomalies and speeds within the Conterminous United States (CONUS) region during days of severe ionosphere activity, also known as "ionosphere storm days". Because the measurements of ionosphere delay come from imperfect receivers whose ability to track GPS signals is also affected by severe ionosphere behavior, each

apparent anomaly must be examined in order to determine whether it is caused by a real ionosphere event or, instead, results from faulty measurements.

Section 2 of this paper describes the existing LAAS spatial ionosphere threat model as revised in mid-2004, before this work was completed (see [2]). Section 3 describes the data sources that were used and how ionosphere anomaly parameters were estimated from the data. Section 4 describes how large apparent gradients were screened to insure that they were not due to errors in the dual-frequency GPS receivers that generated the raw data. Section 5 tabulates the results and identifies the most severe gradients found that survived the screening process. Section 6 discusses how these results affect our understanding of the threat that ionosphere anomalies pose to LAAS.

2. LAAS THREAT MODEL

From the point of view of a user approaching a LAAS-equipped airport, an anomalous ionosphere gradient can be modeled as a semi-infinite wave front. The gradient is assumed to be a linear change in vertical ionosphere delay between maximum and minimum delays. There are two types of gradients, in temporal domain and in spatial domain. A temporal gradient results from the differential ionosphere delay between one receiver's measurements (of the same satellite) at two different epochs. At the same time spatial gradient results from the differential delay between measurements of two receivers to the same satellite at the same epoch. The spatial gradient (slope), along with width and speed of the wave front, are parameters in the LAAS ionosphere threat model. Note that the maximum ionosphere delay difference is the product of slope and width. Figure 1 illustrates this simplified model of an ionosphere wave front [2].

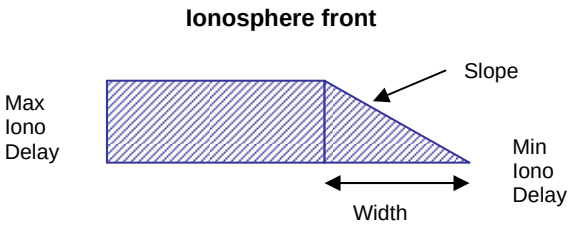

Figure 1: Ionosphere Wave Front Model

The LAAS ionosphere threat model attempts to place upper bounds on the wave front model parameters so that the "worst-case" ionosphere anomaly impact can be determined. A two-dimensional view of the current threat space is illustrated in Figure 2. The red area indicates high elevation "stationary" scenarios, while the green area is for high elevation moving scenarios, and the area circled by the blue lines is associated with low elevation satellites. A "stationary" ionosphere front is considered to be a front with a speed less than 70 m/s, which is a typical jet aircraft final-approach speed. More details of how this threat model was developed can be found in [2].

The objective of this study is to validate the LAAS threat model using actual data points resulted from the data processing.

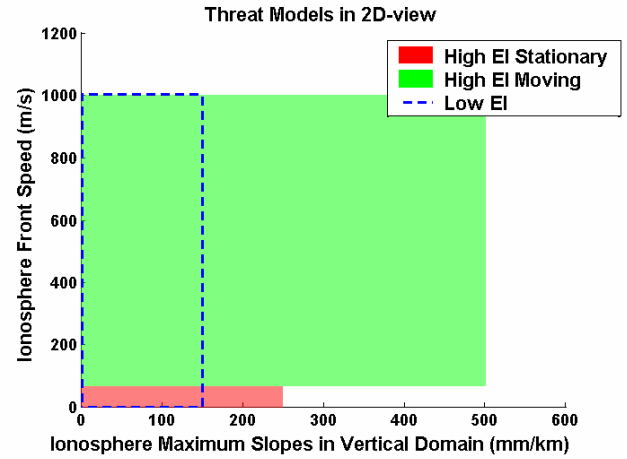

Figure 2: LAAS Threat Model (2-D View)

3. DATA ANALYSIS

The data processing approach described in this paper involves several algorithms to both generate gradient estimates and to validate that gradients are "real" ionosphere events as opposed to receiver "glitches". To make this section easier to flow, a block diagram is shown in Figure 3. The diagram only represents the data analysis for the CORS data set. The WAAS Supertruth data is used to gain insight of this study due to its better quality; however, its low resolution will not provide accurate estimate on threat model parameters (slope and speed). Even though WAAS Supertruth data can not be used to validate the threat model, it provides learning experiences to progress further data analysis. The procedure of the data processing in Figure 3 will be discussed in detail in later sections.

WAAS Supertruth data and data from the National Geodetic Survey Continuously Operating Reference Stations (NGS-CORS) network of receivers are studied for this analysis. This data is obtained from the WAAS network of 25 stations, in which each station is equipped with three redundant dual-frequency receivers. Raw WAAS data is conditioned by post-process carrier leveling and the removal of satellite and receiver biases. Supertruth data is the result of a voting process between the three receivers to remove any glitches present on a single one of them (see [7, 8]).

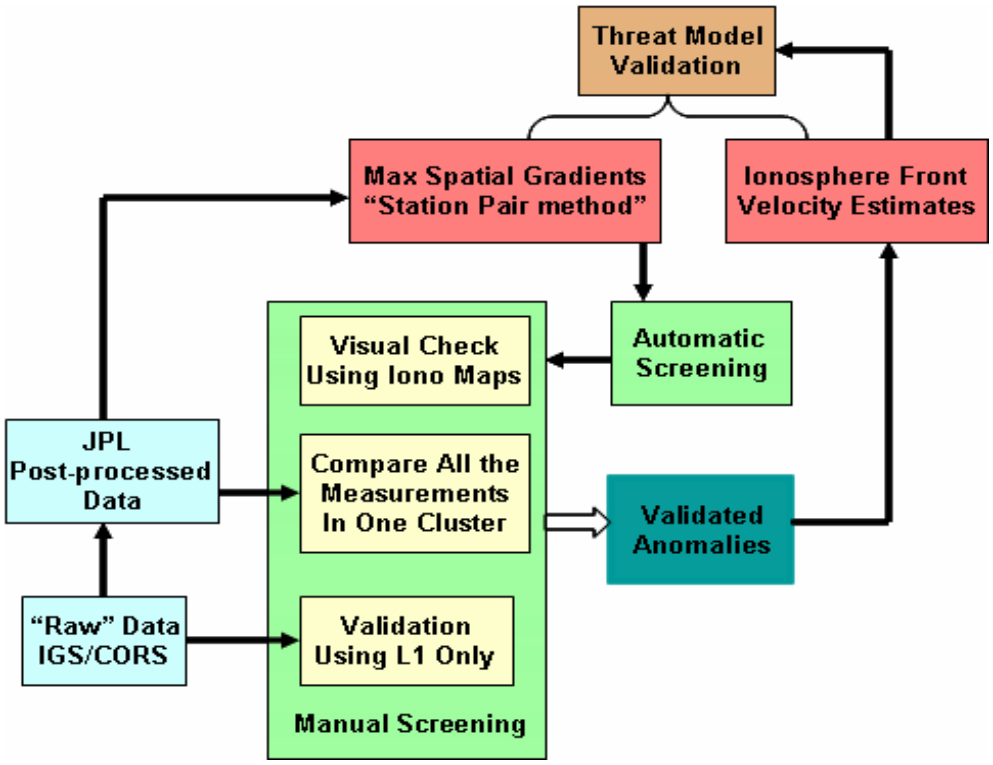

Figure 3: IGS/CORS Data Processing Flow Diagram

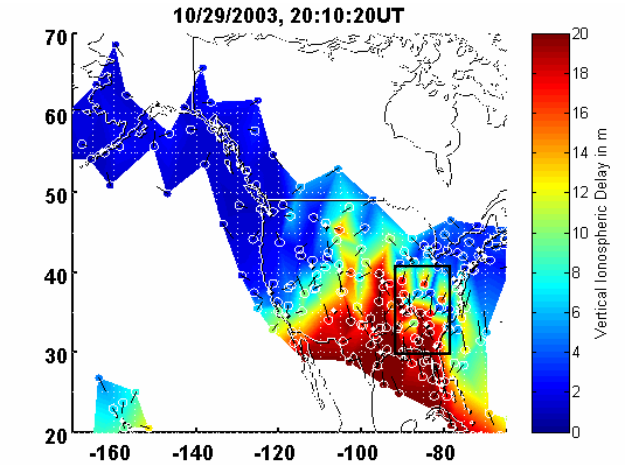

Figure 4: Ionosphere Spatial Anomaly on October 29th, 2003

The second data set was obtained from the CORS network. The CORS network consists of approximately 400 receivers across the US, and it is part of the International GPS Service (IGS), which spans the world and has a total of 980 stations (these numbers are as of 2003). Since the CORS network has a higher density of receivers within CONUS than does WAAS, it provides a basis for obtaining more detailed information on severe ionosphere gradients and for ionosphere front speed estimation. Both raw and JPL-processed CORS data are used in this analysis. The method JPL uses to process this data is called the "WAAS planar fit" ionosphere model algorithm. The thin-shell model is assumed to compute the residual difference between the slant measurements and estimated slant delays based on this planar fit [8]. The "thin shell" assumption, which models the ionosphere as a thin layer at a fixed altitude (assumed to be 350 km for WAAS Supertruth data and 450 km for JPL-processed CORS data) is also necessary for the threat model. The intersection of the GPS signal path from the satellite to the station, called line of sight (LOS), with the thin shell is called the Ionosphere Pierce Point (IPP), which can be thought of as the effective location of each GPS ionosphere measurement, if a thin-shell model is assumed.

3.1 Use of WAAS Supertruth Data

The WAAS Supertruth data on October 29th, 2003, was used to examine the severity of the ionosphere storm on that day. A snapshot of the ionosphere delay map in the vertical (zenith) domain over CONUS at 20:10:20 UTC is shown in Figure 4.

In Figure 4, the x-axis and y-axis denote the longitude and latitude in degrees, with positive values east of the reference meridian and north of the equator. The color scale indicates the magnitude of the vertical ionosphere delay in meters. The dark red region indicates a 20 meter delay, while the dark blue region indicates a minimum delay of 2 meters. The white circles represent IPPs. Transitions from the dark red region to the dark blue

region represent sharp ionosphere spatial gradients. This ionosphere phenomenon appears to be an actual event on a macroscopic scale. Further analysis of the data is needed to understand the severity of this phenomenon at the local scale that is relevant to LAAS.

Due to the large size of the data, which is collected every 5 seconds for each receiver, it is necessary to divide up the data into subsets based on the locations of the stations, for more efficient processing. A data subset thus contains the information about the IPPs located in a particular region. After formation of a subset, we examine all IPP pairs and search for ionosphere spatial gradients large enough to be hazardous to LAAS. All of the IPPs in the same region are paired up and the differential ionosphere delay (DI) between each pair is calculated. The ionosphere spatial gradient is the differential delay divided by the separation distance between this IPP pair. The maximum slope for this particular subset is then found.

The initial results of the October 29th, 2003 data analysis appear to exhibit some extremely high slopes, which physically do not seem to be possible. After a more detailed examination, it was found that all these high maximum ionosphere slopes occur at extremely small IPP separations, often less than 1 km. Even though the apparent spatial gradients are high, it cannot be concluded that these are caused by ionosphere anomalies solely. Some small separations are caused by the geometry of the two satellites. For instance, the lines-of-sight of the two IPPs may "cross" each other (lie at a much smaller separation than the stations themselves). To avoid this, a filtering process was added to eliminate cases of intersections between two lines of sight. This elimination is done by only pairing up the IPPs that fall within a certain limited band of azimuth and elevation. An illustration of this constraint is shown in Figure 5.

The plot in Figure 5 is a visualization of the filtering process in East-North-Up coordinates. The center of the pyramid shape denotes the reference LOS, and it is also marked by the crosshair in the two-dimensional square projection. In this particular example, the "limit band" is set to 45 degrees for both azimuth and elevation. All IPPs inside the pyramid shape can be paired up with the reference IPP (the center IPP). Gradients from all remaining IPP pairs are computed after this filtering process is done. As expected, this filter removes the artificial gradients with very small station separations.

After the preliminary search of ionosphere spatial gradients with IPP pair configuration, it is necessary to further investigate the results to make sure the anomalies are not caused by faulty measurements. The investigation algorithm will be described in section 4. The anomalies from WAAS Supertruth data analysis show that the ionosphere phenomenon on October 29th, 2003 is an actual event, which may have threatened to LAAS users. However, it is difficult to estimate velocity of the ionosphere front, one of the threat model parameters due to the limited number of reference stations in this data set. Therefore, CORS data set is relied upon for further study.

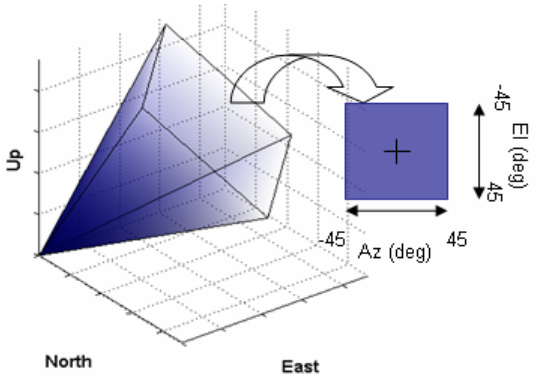

Figure 5: Illustration of Filtering Process for WAAS Supertruth Data

3.2 Use of CORS Data

As mentioned earlier, the CORS network, with approximately 400 stations within CONUS has far more stations than does WAAS and forms a much denser lattice. Before the data analysis process, the CORS database is split up by delimiting "clusters" in the highest-density receiver regions on the US map. Denser spots provide the best possible spatial resolution sampling of the ionosphere. The procedure of separating data into clusters of receivers helps avoid thin shell model assumptions. Since the electron density is distributed at a broad range of altitudes, and the IPP separation distance is sensitive to the assumed shell height, grouping nearby receivers into clusters will produce a better approximation than the thin shell model on a local level. In addition, since all receivers in a cluster have approximately parallel lines of sight towards a particular satellite, the distance between IPPs at any height within the ionosphere is very nearly equal to the straight line station separation.

The implementation of the above considerations results in grouping nearby stations into eight clusters over CONUS of around 21 square degrees average area. An illustration of these clusters is shown is Figure 6. Each of these clusters contains 10 – 30 receivers with a baseline separation of 15 – 200 km.

3.2.1 Spatial Gradient Search

Given the focus on clusters of nearby stations, CORS data processing is somewhat different from WAAS data processing. One difference is that the station pair configuration rather than IPP pair configuration was used

to estimate spatial gradients. In this way, the error introduced by the thin-shell model approximation can be reduced because the IPP separation distance is sensitive to the assumed shell height. Another difference is to only pair up the stations or IPPs which look at the same satellite, thus the fact that their lines-of-sight are approximately parallel also minimizes the thin-shell model effects.

Since satellite geometry plays an important role in the ionosphere gradients, it is important to look at how satellite elevations affect the perceived ionosphere gradients. Therefore, the data in each cluster is separated into five different elevation bins: $0 – 12^o$, $12 – 20^o$, $20 – 30^o$, $30 – 45^o$, and $45 – 90^o$. Data in each different elevation bin is then arranged in "station pair" configurations, and the differential ionosphere delay between each station pair is calculated.

Automated analysis of all station pair configurations is implemented using Matlab. First, it is necessary to automatically extract the relevant data for each cluster from the larger set of data provided by JPL. Data from all receivers was grouped in a single data file, covering a period from an hour to the entire period of active ionosphere during the storm day, an amount depending on the available computer resources. Next, for each LOS elevation range, receivers in each cluster were grouped into pairs, and the difference in ionosphere delay observed by each pair was recorded. Each such DI between two receivers looking at the same satellite was estimated. Subsequently, the slope of ionosphere delay between all pairs of receivers looking at each satellite was estimated by dividing DI by the baseline distance between the two receivers. Finally, a search is performed over the resulting gradients, in order to find the largest ones within each cluster of CORS stations.

The data processing using station pairs configuration, as described above, only provides preliminary results. Further investigation and validation are needed, as shown in Figure 3. The two processes that follow the computation of slopes are automatic and manual screening, and these are described in detail in Section 4.

3.2.2 Ionosphere Front Speed Estimation

After investigating the anomalous ionosphere spatial gradients, estimating the ionosphere front speed associated with these anomalies is needed to populate the threat space. An ionosphere front can be classified as either a moving or stationary front based on the estimated speeds. According to previous work (see [2]), a slower wave front could have a more severe impact on LAAS users. The faster the wave front moves, the better the chance that a LAAS Ground Facility (LGF) can detect the anomaly early on. This distinction between "fast" and "slow" wave fronts is relative to the aircraft speed during a typical precision approach.

A mathematical model has been derived to estimate ionosphere front speeds. Before the discussion of the implementation of the model, a linear assumption of an ionosphere front is made to simplify the analysis, as

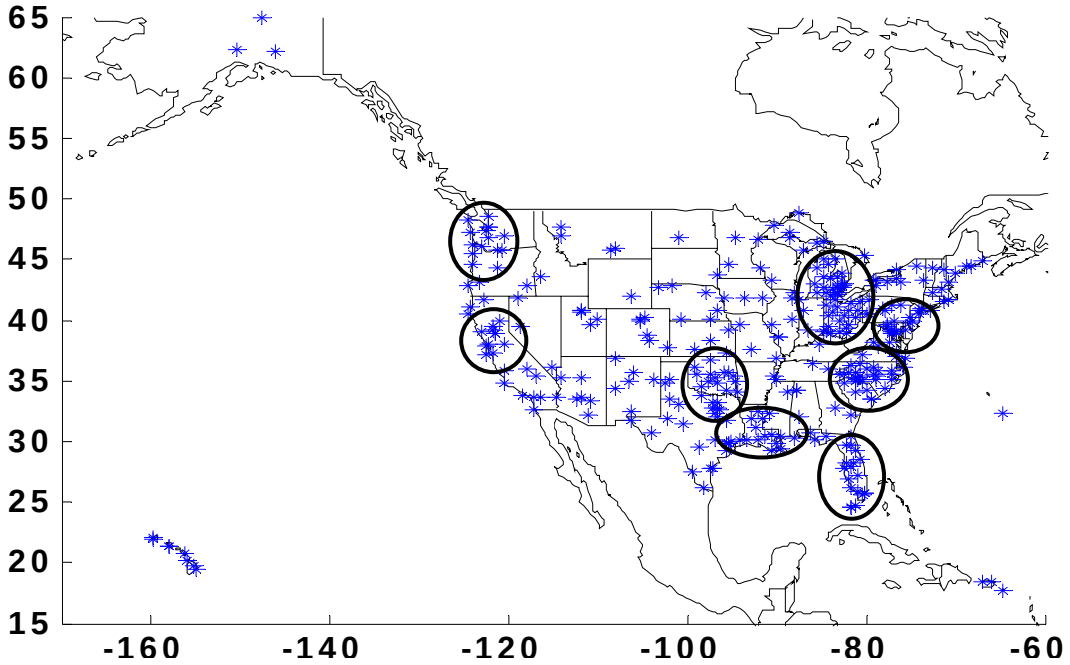

Figure 6: NGS-CORS Clusters. Asterisks represent station locations

described in the introduction. To make it clear, the term "linear" in this analysis refers to the semi-infinite assumption of an ionosphere front illustrated in Figure 1, while the linear assumption used to estimate spatial gradients is the linear slope of the ionosphere ramp. When IPPs cross the ionosphere front, the delay measurements should change from high to low, as show in Figure 7. The peaks represent the maximum ionosphere delays. As the IPPs cross the wave front, the delay measurements start decreasing until they reach minimum, approximately constant delays. The change is approximately linear, as modeled. Four stations are chosen in this case. The reason that four stations were chosen is that they are located nearby inside the same cluster. When the stations are too far apart, the linear model assumption would not apply.

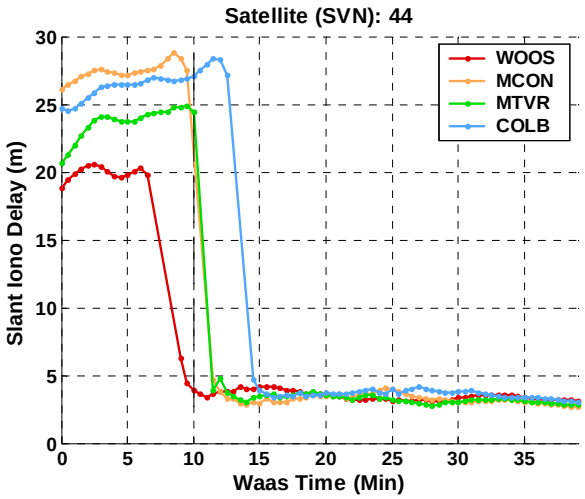

Figure 7: Steep Gradient Measurements for SVN 44

To deduce the wave front speed itself, it is necessary to calculate the speed and direction of IPP motion due to the orbital motion of GPS satellites. The speed (V_{ipp}) and direction (α) of the IPP can be obtained using data measurements, which contains the IPPs' locations every 30 seconds. The ionosphere delay measurements provide the combination of ionosphere wave front and IPP (satellite) movements.

According to previous work, one straightforward way is to find a pair of stations which are crossed by the front in a direction approximately perpendicular to the line connecting the two stations. In this ideal case, the average front speed is just the station separation over the time duration across the ionosphere ramp [4]. However, due to the limited number of stations and the non-linearity of wave fronts, this method does not apply to most observed gradients. A more general method to estimate ionosphere wave front speed and its direction is developed here. As illustrated in Figure 8, after knowing the IPP speed, there still are three other unknowns: inclination (i), speed (V_{iono}) and direction of motion (θ) of the wave front. Solving for these three unknowns requires three equations; therefore at least three stations "swept" by this wave front are needed. V_{iono} is not necessarily perpendicular to the linear wave front, but since only its speed in the normal direction is of concern for LAAS threat model, the velocity can be decomposed after all three unknowns are computed.

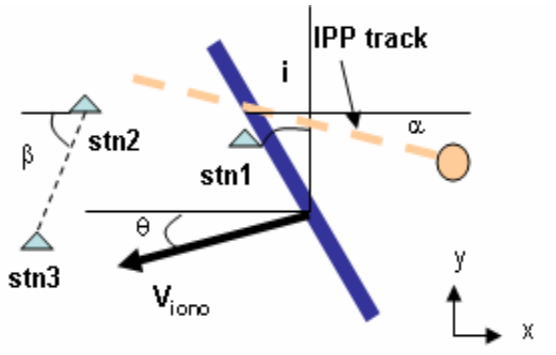

Figure 8: Wave Front Speed and Direction Model

Figure 8 is a picture to illustrate how we define these unknowns. As shown, the ionosphere wave front is inclined and moves southwest as indicated by the black arrow. It sweeps over station 1 first, then station 2, then station 3. The satellite here is moving southeast, indicated by the IPP track (the orange dotted line). Let t_1 be the particular epoch or instant time sweeps station 1, while t_2 and t_3 are the times corresponding to stations 2 and 3. The points (x_1, y_1), (x_2, y_2) and (x_3, y_3) indicate the locations of the three stations in local coordinates. From all these given or calculated parameters, the next two equations follow:

$$(x_2 - x_1) - (y_2 - y_1)\tan i = (V_{iono}\cos\theta + V_{ipp}\cos\alpha)(t_2 - t_1)$$

$$(x_3 - x_1) - (y_3 - y_1)\tan i = (V_{iono}\cos\theta + V_{ipp}\cos\alpha)(t_3 - t_1)$$

To avoid redundancy, the last equation is rotated into the normal direction to the ionosphere front.

$$d\cos\beta_n = (V_{iono}\cos\theta_n + V_{ipp}\sin\alpha_n)dt$$

This last equation applies to any two of the three stations, where β_n, θ_n and α_n are the angles between the line connecting two stations and the normal direction of the front, the moving direction of the wave front and the normal direction of the front, IPP moving direction and the normal direction of the front. These three angles can be represented by β, θ, α and i, where β is a known parameter describing the angle between the line connecting the two stations and the x-axis, as illustrated in Figure 8. The delay measurements in Figure 7 are used as an example to explain this method.

The four stations in the OH/MI cluster are used in order to create different combinations of three stations each. For each such triplet, the normal velocity of the front is calculated, and the results are compared at the end. As shown in the plot, the sharp gradient represents the crossing of an ionosphere wave front. The peak of each curve represents the time when the wave front first affects this station. The maximum ionosphere spatial slope is between station WOOS and MCON and is 313 mm/km in the slant domain. The black dashed vertical line indicates the time when this maximum slope occurs. Implementing the above method, the final result is:

V_{ipp} (m/s)	82.699
α (deg)	-12.051
i (deg)	11.561
θ (deg)	0.42262
V_{iono} (m/s)	100.62
V_n (m/s)	74.725

The negative sign for α means the satellite moves southeast. The wave front inclines to the left from the vertical about 11.2 degrees and moves to the left approximately horizontally. An ionosphere map can be used to examine the results visually.

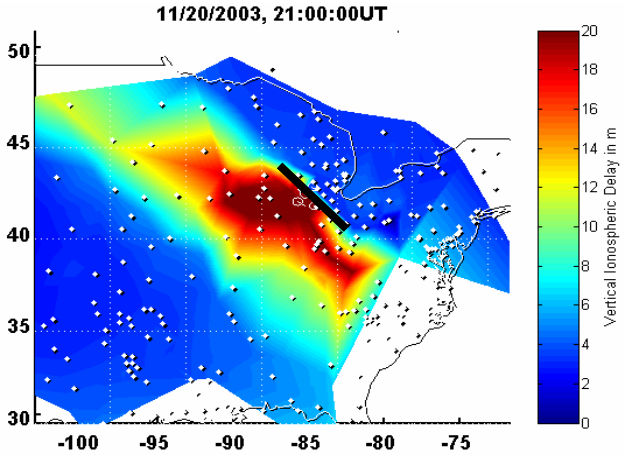

Figure 9: Ionosphere Map at Epoch when Maximum Slope Occurs

As presented in Figure 9, the wave front is approximately linear locally, showing as a black solid line. The white circles on the plot are IPPs. At the epoch shown, the wave front has already passed the IPPs in the blue region northeast of the black line. The normal ionosphere speed is about 75 m/s, which is very close to a jet aircraft speed of 70 m/s.

This method has its limitations. First, the model for the ionosphere wave front is a straight line, whereas real anomalies have many back-and-forth "bends" in them. In addition, the sampling rate of the ionosphere data is 30 seconds, which limits the accuracy of the observed times at which the wave front sweeps stations.

4. SCREENING PROCESS FOR CORS DATA

The goal of the project described in this paper was to isolate the most extreme cases of ionosphere activity, since these are the ones most relevant to the LAAS ionosphere threat model. In order to select these cases from the total amount of available data, a working procedure needed to be established, consisting of both an algorithm for automatic screening of the bulk of the data and a manual examination of the worst points exposed by the computer program. As a starting point for developing such a screening procedure, the two ionosphere storms in autumn 2003 were chosen. In the first phase, CORS data from the respective storm days has been used to develop an analysis algorithm for detecting the most important storm-day events, from a threat model point of view. The computer program implementing this algorithm was written in Matlab, and it was designed to additionally gather supporting evidence that its findings are actual ionosphere disturbances in a way that is easy to examine and understand by a human operator.

The JPL-processed data for the ionosphere storm days is still filled with receiver errors and other artifacts, which is the reason why an event can only be confirmed as real after a careful analysis. The screening algorithms used at JPL are based on an error detection threshold. However, when this threshold is set too high, a high number of false error detections will occur, thus making it impossible to output a sufficient amount of data points at a 30-second sampling rate. Therefore, it was preferable to request that JPL allow more data points by setting a lower threshold, since a robust manual screening procedure was in place to discover the potential errors. Part of that procedure is the comparison against raw L1 RINEX data, which provides a parallel estimate of the spatial ionosphere gradients.

The screening method presented here consists of two stages. These correspond to the "Automatic Screening" and "Manual Screening" blocks in the organizational flowchart [Figure 3]. In the first phase, a computer program isolates the most credible events above a chosen slope threshold (here 200mm/km). It is necessary to have an automatic process for doing this, as the total number of events above the slope threshold in the data is too large for a human to inspect them one by one. Subsequently, once a small number of anomalous gradients have been isolated, it is more beneficial for a human user to inspect these manually for details. This is the case because a variety of trends that can be present in the ionosphere delay measurements, which would be extremely difficult if not impossible to summarize into a set of general rules

programmable into a computer for further screening. At this point, there is also the issue of a computer program not being transparent enough regarding the events it eliminates from consideration based on a more complicated set of rules, hence the underlying fear that real ionosphere events could be eliminated as false negatives. If a human operator would need to review the events eliminated from consideration by the computer algorithm, in order to find these false negatives, this would be equivalent to the amount of work needed for the independent manual screening of all these events. Thus, the need for further automated screening in the second phase is obviated.

It is desirable to check the ionosphere delay as a function of time manually, in order to decide whether the maximum slopes found are caused by a real ionosphere anomaly or by receiver biases and other errors. The procedure used here does not make any assumptions on what types of errors might occur in a GPS receiver. It only tries to differentiate between ionosphere-like and non-ionosphere-like trends or gradients in the recorded measurements. Furthermore, the relation between maximum gradients and the elevation angles is also interesting for the analysis. According to the October 28-31 2003 data, most maximum gradients occurred in the low elevation range (below 45 degrees). Additionally, when looking at slant ionosphere delays, it is useful to compare only data from the same narrow elevation range, so that differences between slope values are determined by the actual ionosphere and not just by variations in the LOS obliquity angles.

4.1 Automatic Processing of Storm Data

The main step of automatic screening is to examine all slopes that are over a certain threshold (e.g. 200 mm/km in slant measurements), and eliminate those gradients apparently due to receiver bias and the isolated data points, for which no trend of ionosphere activity can be inferred. Slopes are considered to be mainly due to receiver bias when two receivers are collocated (i.e., less than 100 m apart), or no ionosphere activity seems to be detected in one of the receivers (i.e., readings from this receiver do not vary in time), or their DI does not vary in time (i.e., they exhibit a constant bias over the period of observation). In what regards the minimal temporal variation used in the algorithm, only ionosphere delay variability of over 15 mm/s is considered as significant. This threshold has been set to a value an order of magnitude lower than the extreme value observed by Datta-Barua [5] for the 10/29/2003 storm day, such that the most significant rapid changes in the delay value are above the threshold, while most quiet ionosphere variability is beneath. An important source of error for the algorithm is, consequently, the case when data contains cycle slips, causing the delay value to jump suddenly between two successive measurement points. These cases have to be eliminated manually from the screening results, as will be seen in the next section.

This selection algorithm might be found to be aggressive and biased against detecting slow-moving ionosphere fronts. However, given the high activity on 11/20/2003, a robust method was necessary, to insure that only the events that look the most realistic are selected. The selection algorithm does a good job at sifting out receiver artifacts, and retains only the data caused by strong, moving ionosphere fronts. This will be exemplified in section 5, which discusses selected events. Particularities encountered in the data dictate the constraints, with which the development and validation of this method need to negotiate. Completely stationary, small scale ionosphere features would be indistinguishable from a constant receiver bias in this approach. Yet, they could be confirmed by having several receivers exhibit similar levels of "bias" or by charting the measurements on a map of ionosphere delays, on which none of the receiver measurements are eliminated from consideration.

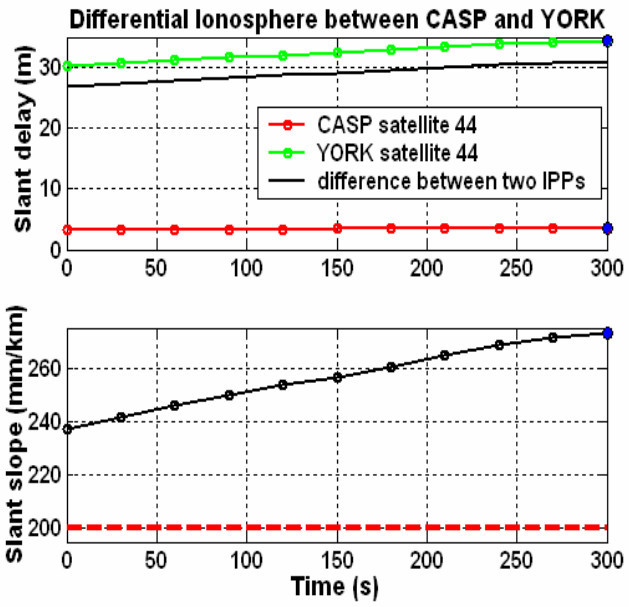

Figure 10: The CASP receiver does not record any variation in ionosphere delay, and the slope is basically driven by the other receiver in the pair. This effect is probably caused by a receiver problem.

An example of what the computer program is set to eliminate is the measurement taken by the CASP receiver during the November 20, 2003 storm [Figure 10]. When plotted against the simultaneous measurement from a nearby receiver, YORK, a large gap is visible between the indications of the two stations. While YORK seems to indicate the presence of active ionosphere, causing a slant

delay of over 30 meters, the CASP receiver appears unaffected by the storm, indicating a low value for the delay, close to 5 meters, which is more characteristic of a quiet ionosphere. Intuitively, one of the two receivers has to be wrong and the high value of the difference in delay, producing a large spatial gradient, is probably unrealistic. The algorithm outlined above would eliminate this event from consideration, as the absence of any variability, even on a small scale, in the measurements at CASP, is indicative that the recordings of ionosphere delay are not real. As supporting evidence, a movie caption is presented in Figure 11, illustrating this apparent high slope event. In the zoomed-in bottom picture, the corresponding IPPs for CASP and YORK are marked by a white and a black dot, respectively. The red color represents the presence of high delays due to active ionosphere in that region of the sky, while the small white patch is caused by the erroneously low reading from the CASP receiver. There is a clear disagreement between delays from all the other stations in the area and what CASP measured.

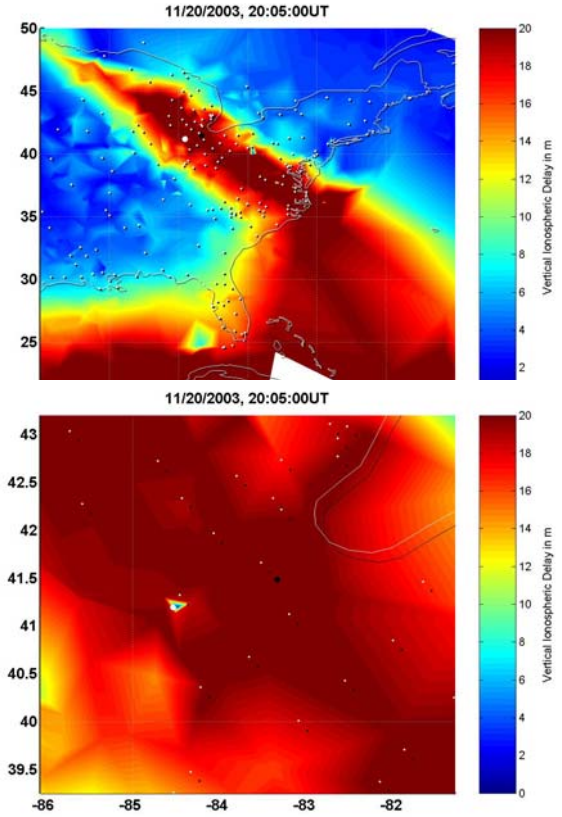

Figure 11: Map of ionosphere delay around IPP for receiver CASP. The bottom image of the November 20 storm is the magnified version of the upper plot around the area of interest.

4.2 Manual Analysis of Selected Cases

The resulting cases after the computer program has been used to screen the entire data set were screened manually. These may be real events above the slope threshold. For each of these events, our own (manual) analysis has decided whether the magnitude of the gradients, together with the length of time that these gradients manifest themselves look reasonable. There are a series of checks which can be used to verify that a high slope event is indeed caused by ionosphere activity, or, conversely, produce evidence as why the observation is unlikely to be linked with an ionosphere event.

4.2.1 Comparison with Neighboring Cluster Pairs

The first way to verify an event is comparison against simultaneous observations from receivers in the same cluster. The assumption here is that, if the majority of the receivers agree to a certain degree with the pair that observed the most extreme values, the event is very probably real. On the other hand, if the observations were influenced by receiver artifacts in the data, it is expected that a majority of the receivers in the cluster do not contain those artifacts. Instead, data from most of the stations represents an image of the event that is closer to what happened in reality at the ionosphere level. For comparison purposes, all ionosphere delay curves for a given cluster are plotted together, such that it is conspicuous which trends are similar, versus which particular detections were likely caused by receiver errors (e.g. jumps, gaps in data, noise). These plots extend across the entire period of time for which data was extracted, such that the behavior of the receiver data can be analyzed both in the presence of disturbed ionosphere and also during calm ionosphere weather at the beginning or the end of the storm. As such, the presence of constant bias in certain channels is more easily dissociated from the detections of delay produced by moving or stationary ionosphere fronts. While the electron content within the ionosphere is expected to vary wildly during the actual storm, delay measurements from the different receivers are expected to converge at times of nominal ionosphere activity or at night time.

4.2.2 Parallel Slope Estimation using L1 Only

The next step in checking the validity of conclusions about the existence of extreme events is comparing the JPL-provided data against raw L1-only RINEX data for the station pair looking at a given satellite, at the same time when a large gradient has been observed. The resulting estimations of these slopes are however not exact. This data is not processed, such that a number of biases in the receivers, which can be otherwise characterized, have not been eliminated. Time-averaged code measurement is used to correct the integer ambiguity

of the carrier phase measurement, such that the final ionosphere delay is:

$$I_{L1\rho\phi} = \frac{\rho_{L1} - \varphi_{L1}}{2}$$

As such, the method is not exact, but provides a robust qualitative picture of gradients that is not subject to fragile L2 tracking loops.

Delay measurement data, packed in the RINEX format, is freely available directly from the NGS-CORS website (http://www.ngs.noaa.gov/CORS/cors-data.html). A point in time is picked, where the measurements from the two receivers are aligned, by subtracting a constant bias from all the measurement points by one of the receivers, thus equalizing the value they measure at the chosen instant. Since the DI between a pair of receivers is observed, this is equivalent to choosing a reference zero-slope point at a particular instance in time and then expressing all the other slope measurements relative to it. Since all L1 slope values are dependent on the point chosen for the alignment, this procedure only provides a lower bound on the actual ionosphere gradient.

The reason why the raw L1 measurements can only provide a lower bound estimate on the slope is that there is no systematic way to remove biases from these measurements. The bootstrap procedure for removing large biases existing in RINEX measurements is to first choose a point in time where we align each receiver measurement in L1 to its corresponding value in the available JPL L1/L2 corrected data. Furthermore, when evaluating the spatial ionosphere gradient between a pair of receivers, a point during the quiet ionosphere part of the data (i.e. far away from the peak of the storm, usually during local night time) is selected, for which the L1 measurements of the two receivers are again aligned, in order to remove any residual bias, as it is known that quiet ionosphere exhibits a much lower spatial variability over distances of the order of the cluster diameter. Since the high slopes are caused only by differences in receiver measurements, usually during a peak of ionosphere activity, only the differences between the L1-only measurements contribute to the estimation, and not the absolute values of the delays, which cannot be determined precisely by the current method.

It is wished to believe that the above procedure is useful for eliminating the majority of inter-receiver biases between two stations, while leaving only the gradient determined by actual ionosphere variability in the L1 slope estimation. However, it is impossible to be confident that these biases were completely removed instead of just partially, or to be certain that delays caused by slowly moving ionosphere were not removed as biases. If the observations do not span long enough in time, until all major ionosphere fluctuations disappear, as assumed, there exists the risk of misdetection of a stationary disturbed ionosphere. The amount of risk is proportional in magnitude to the commonality of large stationary disturbances. Thus, it is only safe to affirm that the verification using raw L1 CORS measurements provides a lower bound for the spatial gradient of an event. It is hoped that this procedure allows the separation of biases present in the receivers from differential delays introduced by the ionosphere. However, it is more of an empirical rather than a rigorous method, and it is rather used to provide a ballpark point of comparison between how much of the previously observed gradient might be due to an ionosphere feature, versus pure constant bias in one of the receivers.

4.2.3 Generation of Ionosphere Delay Maps

The final check that helps validate an ionosphere feature accompanied by a steep gradient is putting together overview pictures of the ionosphere activity in a certain region in the form of ionosphere maps. These maps are composed of the delays observed by all CONUS stations and can be assembled in sequence to produce movies that show the time evolution of the ionosphere disturbances.

Furthermore, if the map consists of data from all receivers in a particular region looking at the same satellite, a consistent picture of the ionosphere structure is obtained, because LOS are near-parallel and do not intersect each other. The real ionosphere is a 3D structure, and a number of concepts that are used in this paper (e.g. gradient, velocity, width) would be ill-defined unless radiography of this structure is performed along a unique direction. Therefore, it is preferable to produce single-satellite images at every time epoch (e.g. data points are taken every 30 seconds), in order to avoid looking through the ionosphere at a broad range of elevation angles simultaneously. For example, if two LOS make a large angle (in elevation and/or azimuth) with each other, even speaking of an "ionosphere front" is ambiguous, as the projections of the ionosphere total electron content (TEC) in section planes perpendicular to the two rays could yield two different front profiles, even at the same location (IPP latitude and longitude, but not necessarily altitude) on the sky.

There are two ways to compose these maps, by interpolation between given measurement points for each instant in time, and using a fixed mesh with all active receivers as vertices [Figure 12]. In the first method, measurements from all active receivers are interpolated to produce an estimate of the ionosphere delay over the entire geographic region that is covered. The measured ionosphere delays can be mapped alternatively, both at the IPPs corresponding to a given shell height, or at the

GPS stations themselves, located at Earth's surface.

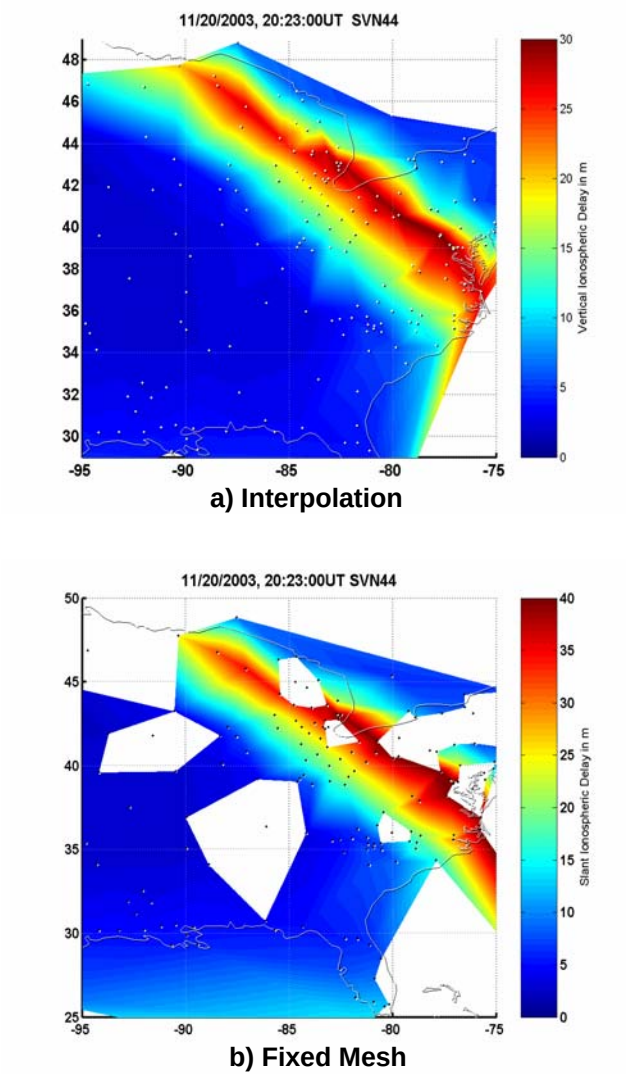

Figure 12: Maps of ionosphere delay for the November 20, 2003 storm day, for which the delay values have been mapped at the station locations. Stations are marked by dots, superimposed over the contour of the Eastern U.S. a) Interpolation; b) Fixed

Alternately, the measurement points comprising a network of vertices, between which intermediate delay values are interpolated, can be fixed to form a unique mesh, which is used for all the frames of a movie. In practice, all receivers that recorded data for at least one time point during the period covered in the movie are set as vertices of the mesh. When data received from a particular station has missing time points, or these were eliminated during JPL processing, a corresponding blank area appears on the map, surrounding the position of that receiver, instead of interpolating from more distant available receivers and producing potentially inaccurate intermediate values. Obviously, this second method, though more exact, is only applicable for mapping delay measurements directly to the receiver locations, because IPPs are never stationary points on the map. As discussed above, this treatment is consistent if applied to localized clusters, for which quasi-parallel LOS to each satellite obviate the need for using a thin shell model for the ionosphere or making reference to the notion of an IPP.

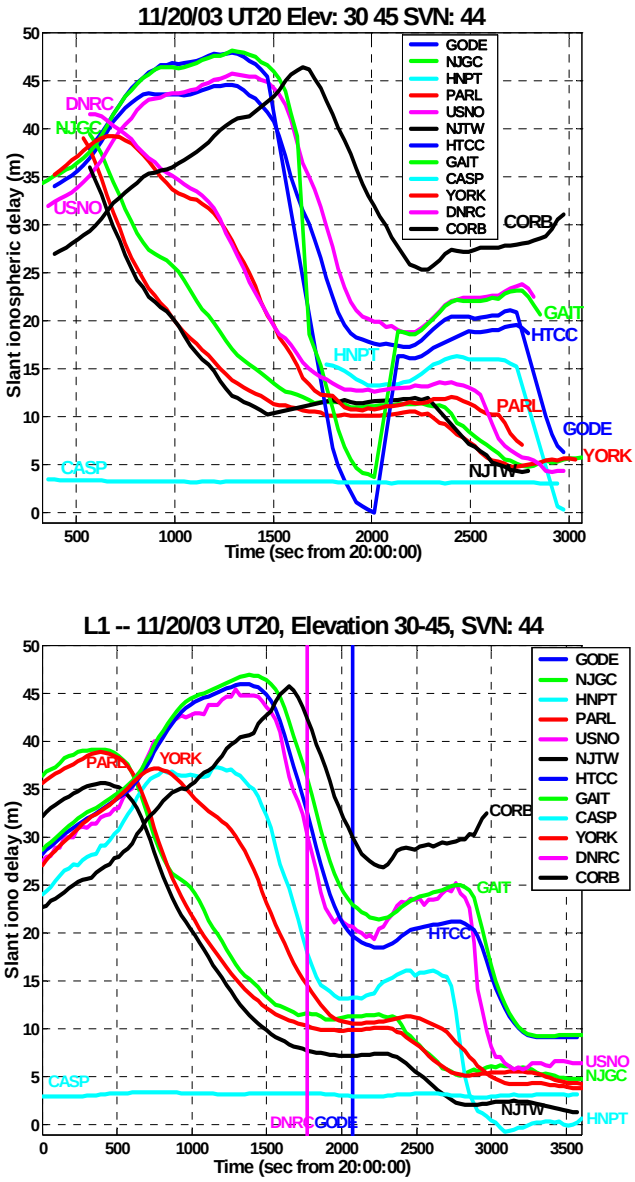

Figure 13: Visualization of the high slope event between USNO and GAIT receivers in the Washington DC cluster on November 20, 2003, in the context of the observations of all receivers in the same cluster.

4.3. Validation of the Screening Process

The following example illustrates how the screening method outlined in this section is applied to real data. This example illustrates a case in which an apparent

gradient was flagged by automatic screening but was not validated by manual screening. The selected event is the large anomaly (587.01 mm/km in slant) recorded between USNO and GAIT along a LOS towards SVN 44 (PRN 28), at 20:30:30 UT on November 20, 2003 [Chart 1]. In this case, the L1-only data yielded a considerably different slope from the L1/L2 pre-processed data. The L1-estimate was only about 200 mm/km slant slope, which is a lower bound on what the actual gradient might have been, but gives a reason to believe that the actual ionosphere slope was only 200 - 300 mm/km in reality.

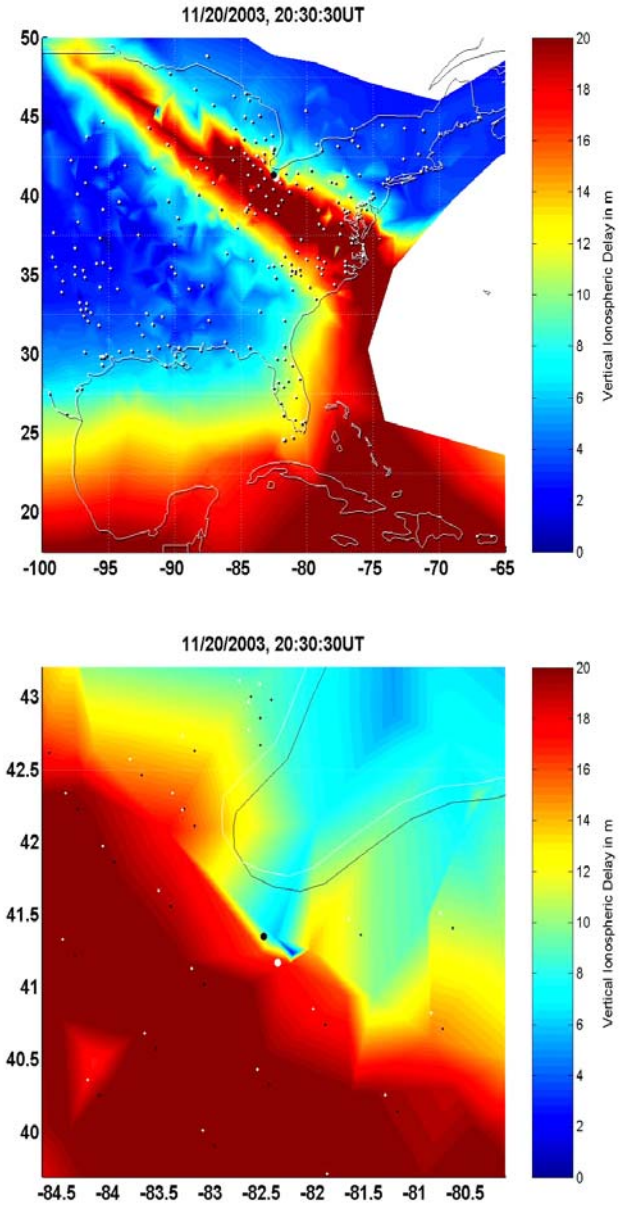

Figure 14: Visualization of the high slope event between USNO and GAIT receivers in the Washington DC cluster on November 20, 2003.

This value would be comparable with observations by other receiver pairs active in the Washington DC cluster at the time [Figure 13]. The additional difference in slope was likely caused by a loss of lock in the L2 band, triggered by the extreme ionosphere conditions.

Figure 14 is an instant from a movie of the ionosphere delay, showing that, indeed, an ionosphere gradient is present in between the LOS for the two receivers. The black dot marks the IPP corresponding to the receiver GAIT, and the white dot that for USNO. Although the loss of lock for the HTCC and GAIT receivers distorts slightly the reality of the map, it is clear that the observations were made towards the edge of a large ionosphere feature, about the time when the transition line from high to low delay was intersecting the LOS of the receivers in the Washington DC cluster toward SVN 44.

5. RESULTS AND ALGORITHM VALIDATION

This section presents some of the most relevant results of this analysis, illustrating the main points made in this paper and the specific categories of ionosphere events encountered in the data. A visual depiction of the automatic elimination process of unrealistic gradients is presented in Figure 16. Most of the high slopes corresponding to a small receiver separation in the histogram above end up being eliminated, as large receiver biases cause the highest apparent spatial gradients when divided by small baseline separations. Figure 17 exemplifies how the initial set of data is screened for the cases that could periclitate the bounds of the existing LAAS threat model, which are be eventually selected by the computer running the screening program. All the selected cases at the bottom are further inspected as described in Section 4.2.

5.1 Summary of Oct.-Nov. 2003 Results

The resulting high-slope events are summarized by the cluster (geographic area) and satellite elevation range in which they occurred, along with details identifying the event. The way clusters were delimited on the CONUS map [Figure 6] is summarized in Figure 15. The tables included at the end of the paper contain the events with the highest slope encountered in each cluster and elevation range during that particular day, if any events over 200mm/km slant delay gradient have successfully passed the automatic screening process [Tables 1 – 7]. Each of the tabulated events has been verified against L1-only RINEX raw data, except the cases where it is mentioned that accurate L1-only data could not be obtained.

Number	Name	Longitude range	Latitude range	Number of receivers	
				2000	2003
1	Wash DC	-77.75 -74.75	38.15 40.25	6	15
2	NC	-83 -75.25	33.4 37.6	6	33
3	Florida	-83 -80	24 31	4	19
4	LA/TX	-99 -88.75	29.2 31.2	2	18
5	OK/TX	-99.5 -94	31.5 37	7	23
6	OH/MI	-86 -79.31	38.5 45.5	5	47
7	NorCal	-123.1 -120.35	37 40	8	11
8	Sea WA	-122.95 -122	46 48.55	6	7
9	NE	-81 -66	35 45	27	-

Figure 15: **Definition of Receiver Clusters**

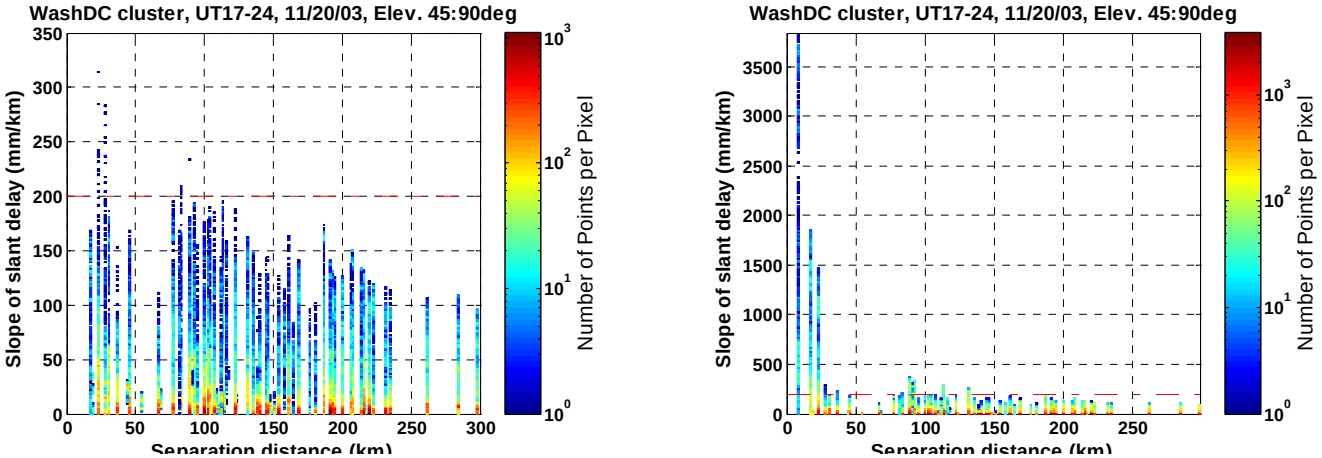

Figure 16: Histograms of slant slope measurements before (left) and after (right) automatic screening with a slant slope threshold of 200 mm/km.

Number of analyzed receiver pairs* per elevation range (Nov 20, 2003)

Slope	*0-12*		*12-20*		*20-30*		*30-45*		*45-90*	
Total	133,594	100.0000%	956,433	100.0000%	950,088	100.0000%	1,265,094	100.0000%	2,859,410	100.0000%
0-100	130,832	97.9325%	945,506	98.8575%	945,705	99.5387%	1,255,597	99.2493%	2,840,624	99.3430%
100-200	1,888	1.4132%	8,676	0.9071%	1,939	0.2041%	4,906	0.3878%	10,810	0.3781%
>200	874	0.6542%	2,251	0.2354%	2,444	0.2572%	4,591	0.3629%	7,976	0.2789%
selected	20	0.0150%	10	0.0010%	3	0.0003%	38	0.0030%	86	0.0030%

** a pair represents* **two** *different* **receivers** *looking at the* **same satellite** *at the* **same time**

Figure 17: Data Analysis Flowchart for November 20th, 2003. Example illustrating the total number of identified pairs, versus the final number of events selected by automatic screening.

From the results tables 1-7, a few examples have been chosen to illustrate in more detail the slope validation process. The first example represents a validated event, which was supported both by automatic screening and supplementary manual screening. The event is a large slope observed by receivers ERLA and GRTN along a line of sight to the same SV 44 [Chart 2]. Although the L1/L2 combined measurements exhibit a sudden jump, about half an hour after the event presented in Section 4.3, the L1-only data contains some intermediary points during this jump, which is reassuring that the measurements came from an actual event at the ionosphere level. It is also visible from this example how the codeless tracking on the L2 frequency is more vulnerable to cycle slips during strong ionosphere disturbances, and thus a known C/A-type signal is needed at all frequencies if the LAAS is to be deployed in a multi-frequency format.

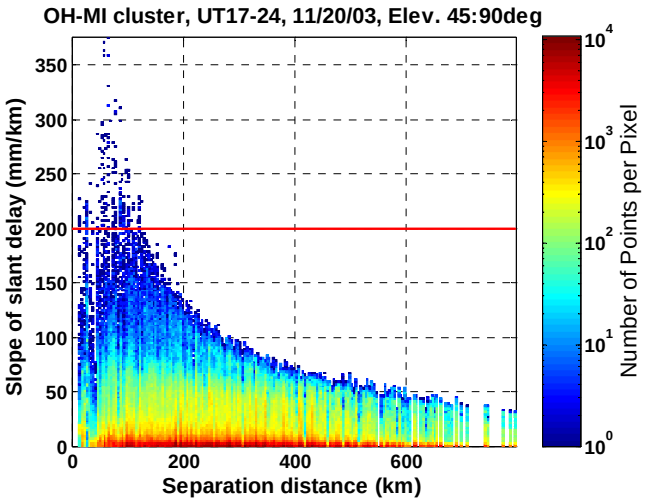

Figure 18: Histogram of ionosphere spatial gradients versus baseline separation for the OH/MI receiver cluster.

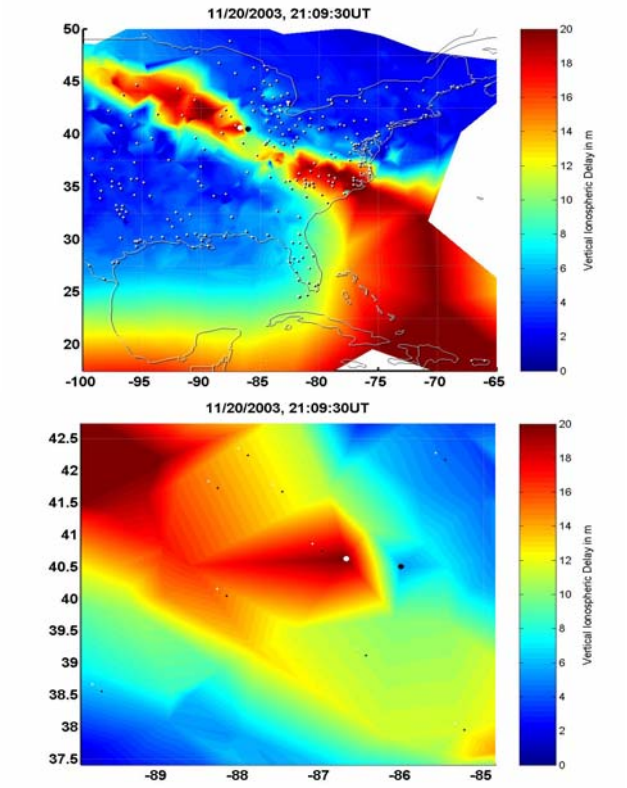

Figure 19: Visualization of the high-slope event between ERLA and GRTN receivers in the OH/MI cluster on November 20th, 2003.

The measured size of the gradient is very similar between the L1/L2 and L1-only measurements in this case. In addition, numerous pairs of receivers located in the vicinity, within the Ohio/Michigan cluster, record similar abnormally high slopes [Figure 18]. In a larger perspective, the event has been also captured in an overall picture of the ionosphere at that time [Figure 19], which is a strong confirmation that the event was not caused by a glitch in the receiver data.

Another example, this time from the October 29, 2003 storm day, seems to confirm this idea, that the slopes derived from raw L1-only data are slightly lower than reality, if we are to trust the JPL-processed data for better accuracy in the absence of receiver artifacts. A gradient of over 200 mm/km is observed between receivers LEBA and ERLA in the direction of SVN 46 at 21:05:30 UT [Chart 3]. Here, looking at the L1/L2 data first, a visible decrease in the overall delay is observed by both receivers successively over an interval of 30 seconds to 1 minute. Though the duration of the event is very brief, this is very clearly a case of an ionosphere-induced spatial gradient. Nevertheless, the sudden, unexpected drop makes a check against L1-only data desirable. The reason why significantly more data is present in the L1-only measurements (i.e. there is no gap in the middle, and data is available for earlier times as well), is that the L1/L2 data was split into elevation ranges (e.g. present here is data for satellites at elevations between 30 and 45 degrees), while the L1 data was not. However, in the time period around the anomalous event both data sets are consistent with each other, and the two manually aligned sets of L1-only measurements confirm a slope over the slant slope threshold.

Figure 20 summarizes the most severe events discovered in the October and November 2003 storms by plotting the estimated slope and speed of each point on top of the ionosphere threat model shown previously in Figure 2. The triangles represent the data points from L1 and L2 measurements that did not appear to be caused by faulty receivers, but for which a parallel estimation of the slope by the means of raw L1-only data was not available. The squares represent L1/L2 data points validated by L1 code-minus-carrier measurements. It must also be mentioned that an estimation of the front velocity was not always available for the highest spatial gradients discovered by each cluster. Oftentimes a propitious alignment between multiple stations and a moving ionosphere front does not occur, thus rendering the front speed estimation impossible.

Two important conclusions can be drawn from Figure 20. First, none of the events witnessed during the October and November 2003 storms (the worst on record for which we have detailed GPS data) approach the boundaries of the "high-elevation/moving" component of the threat model. The maximum slope in the vertical domain was about 300 mm/km, while the maximum speed was about 600 m/s. This suggests that the outer boundaries of the "high-elevation/moving" threat model (500 mm/km and 1000 m/s, respectively) may be too conservative. Some degree of conservatism is warranted,

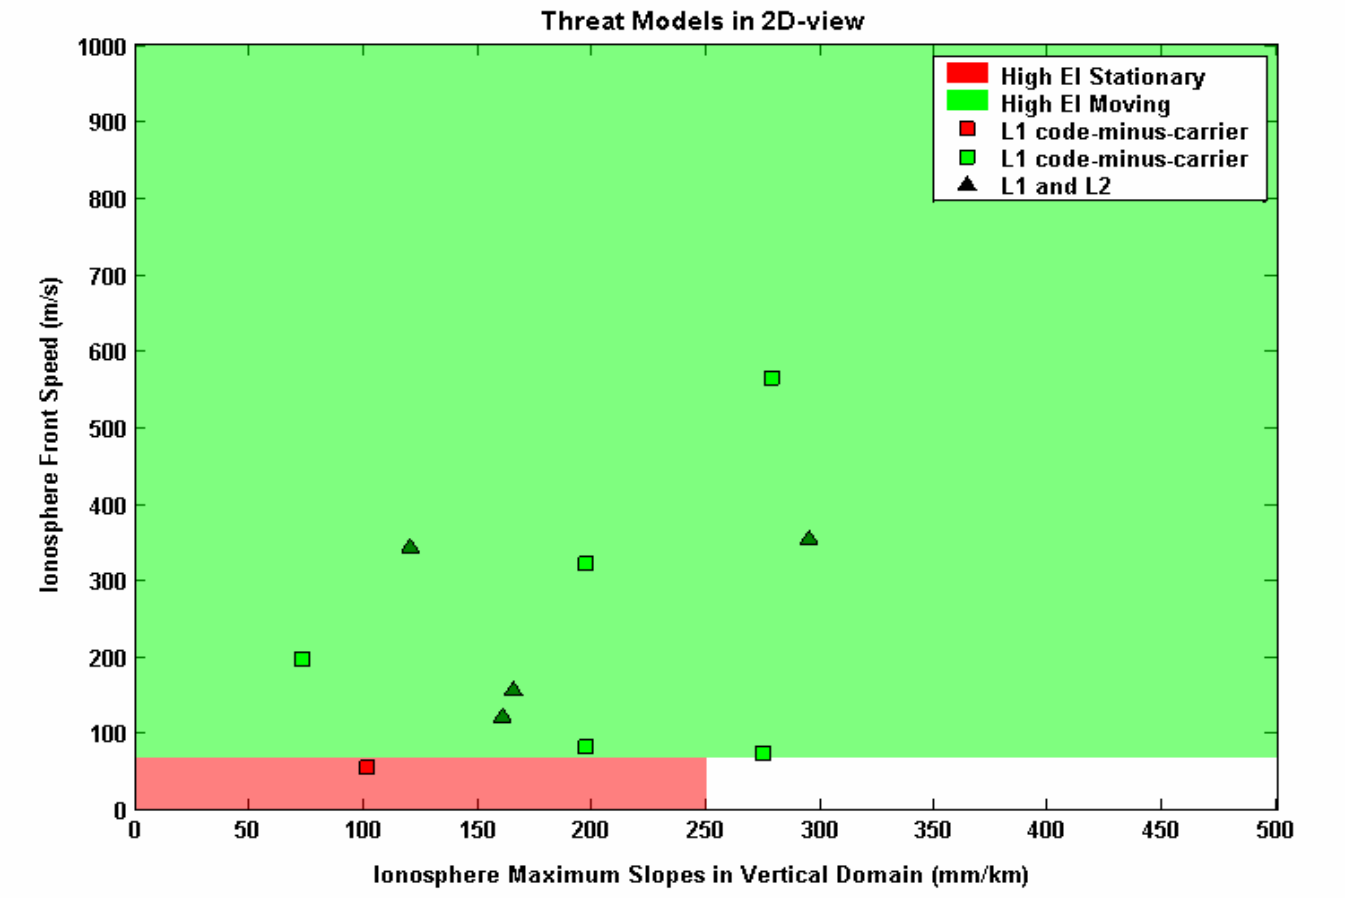

Figure 20: LAAS Threat Model with Validated Data Points

however, because we only have detailed GPS ionosphere data for the 2000 – 2004 period. Second, one of the anomalies verified by L1 code-minus-carrier analysis just barely falls into the "moving" part of the threat model because its apparent speed (about 85 m/s) slightly exceeds that of an approaching aircraft (about 70 m/s). Since its estimated slope of about 270 mm/km exceeds the upper bound of the "stationary" component of the threat model, and because the speed estimation method described in Section 3.2.2 only allows for rough estimates, further analysis of this event (as well as the other two low-speed cases) is needed to confirm that this particular event really was moving faster than an approaching airplane from the point of view of a hypothetical LGF installed nearby.

5.2 Summary of Results for Other Storm Days

The data sets that constituted the basis of this investigation covered the periods of most intense ionosphere activity during the past five years, as these are the times when it is most likely for an anomalous spatial or temporal gradient to occur. More precisely, the data included observations made over the following periods of time (hour is given as Universal Time): 4/6/2000 17:00-03:00, 4/7/2000 14:00-23:00, 7/15/2000 20:00-03:00, 10/29/2003 18:00-02:00, 10/30/2003 19:00-01:00, 11/20/2003 17:00-24:00, and 7/17/2004 18:00-01:00.

The algorithm described by the block diagram of Figure 3 was applied to additional storm days signaled by the literature in the field (see [4, 5, 7]) and identified as well by the monitoring stations across the US. The experience gained from examining data for October/November 2003 was applied to the other known days of abnormally high ionosphere activity over CONUS.

Because not all stations were available in the year 2000 data, a slightly different set of clusters than those shown in Figure 6 was used for this analysis. This particular adjustment had to be made in order to insure that the search area over CONUS is relevant for the storms in year 2000. Since the CORS network consisted of a much smaller number of stations in 2000 compared to 2003 (approximately 150 versus 400, respectively), an alternation had to be made to the division of the US map into clusters of GPS stations. It is known that the 2000 storms were stronger on the East Coast [4, 5, 7] and thus, in order not to miss any potential events, a Northeast (NE) cluster was created, including receivers from Washington D.C. area to Maine. This cluster, which is much larger than the others, was appropriate for the case where

stations were spread farther apart, but still denser than in much of the rest of the US at the time.

The search over JPL-provided data for the above storm days yielded no slant gradients over the 200 mm/km threshold for the storm in 2004. In the meantime, a number of extreme events appeared during 2000; however none of these was deemed to be real. The reason is that all the selected cases were marred by receiver errors in both L1/L2 and L1-only, or, alternatively, appropriate L1 data for checking these points was lacking. An example from the first category is the apparent high slope between USNO and GAIT [Chart 4] contains a cycle slip which is corrected by the L1-only data. For the cases where the RINEX information is not accurate enough to determine a L1-only slope, either better methods have to be developed to read and interpret the RINEX format files from CORS, or an alternate data source should be sought that can provide appropriate verification of the L1/L2 results.

6. CONCLUSIONS

The results presented here show that the 28-31 October and 19-20 November 2003 storms were similar in intensity and had the largest number of extreme ionosphere events of all the storms observed since WAAS was fielded in year 2000. Thus, a future LAAS threat model should be designed to protect against similar events. The other storms investigated, while potentially disruptive to civil aviation, fell well within the boundaries of the same threat model, as no events as critical as what was observed during the Autumn 2003 storms were found and confirmed by this screening process during any of the other days investigated over any of the receiver cluster areas. The only events for which further investigation is needed are the three points in Figure 20 whose estimation of ionosphere front velocity falls close to the 70-m/s aircraft approach speed.

An interesting result of this analysis is the discovery of one of the largest values of ionosphere delays of the GPS signal ever recorded. A validated measurement of over 45 meters vertical delay was discovered within CONUS on November 20, 2003 (in addition, delays of over 50 meters were found at low latitudes). This result is larger than the value of 36 meters presented by Misra and Enge [6] as the highest measured delay at any latitude.

The key challenge in this work was deriving the maximum amount of validated (cross-checked to the extent possible) severe ionosphere gradient information from data that is unavoidably limited (in both temporal and spatial density) and corrupted by receiver errors. We believe that the method developed in this paper is close to optimal for the existing WAAS and IGS/CORS networks. A related but somewhat different approach was developed for the Japan GEONET database, which is much denser spatially and allows more-aggressive screening to occur without missing potentially "real" events (see [9] for details).

ACKNOWLEDGMENTS

The authors would like to thank Seebany Datta-Barua, Todd Walter, Jason Rife, and Hiroyuki Konno for their help during this research. The advice and interest of many other people in the Stanford GPS research group is appreciated, as is funding support from the FAA LAAS Program Office (AND-710). The opinions discussed here are those of the authors and do not necessarily represent those of the FAA or other affiliated agencies.

REFERENCES

[1] M. Luo, et al., "Ionosphere Spatial Gradient Threat for LAAS: Mitigation and Tolerable Threat Space," Proceedings of ION NTM 2004. San Diego, CA, Jan. 26-28, 2004, pp. 49 -501.

[2] M. Luo, *et al*., "Ionosphere Threat to LAAS: Updated Model, User Impact, and Mitigations," *Proceedings of ION GPS 2004*. Long Beach, CA, Sept. 21-24, 2004, pp. 2771-2785.

[3] M. Luo, *et al*, "Assessment of Ionospheric Impact on LAAS Using WAAS Supertruth Data", *Proceedings of the ION 58th Annual Meeting*. Albuquerque, NM, June 24-26, 2002, pp. 175-186.

[4] T. Dehel, *et al*., "Satellite Navigation vs. the Ionosphere: Where Are We, and Where Are We Going?" *Proceedings of ION GPS 2004*. Long Beach, CA, Sept. 21-24, 2004, pp. 375-386.

[5] S. Datta-Barua, "Ionosphere Threats to Space-Based Augmentation System Development," Proceedings of ION GNSS 2004. Long Beach, CA, Sept. 21-24, 2004.

[6] P. Misra, P. Enge, *Global Positioning System: Signals, Measurements, and Performance*. Ganga-Jamuna Press, 2001.

[7] S. Datta-Barua, *et al*., "Using WAAS Ionospheric Data to Estimate LAAS Short Baseline Gradients," *Proceedings of the ION 2002 National Technical Meeting*. Anaheim, CA, January 28-30, 2002, pp. 523-530.

[8] A. Komjathy, *et al*., "The Ionospheric Impact of the October 2003 Storm Event on WAAS," *Proceedings of ION GPS 2004*. Long Beach, CA, Sept. 21-24, 2004, pp. 1298-1307.

[9] H. Konno, *et al*., "Analysis of Ionosphere Gradients using Japan GEONET Data," *Proceedings of the ION 2005 National Technical Meeting*. San Diego, CA, January 24-26, 2005.

04-06-2000 17:00-03:00 UT

Cluster	Time (UT) Duration (min)	El Band	Lat	Lon	STN	SVN #	d_stn (km)	Vert Slope (mm/km)	Slant Slope (mm/km)	L1 only (slant)
SEA/WA	00:31:30	12~20	47.69	-122.26	SEAW	40	5.43	161.17	359.72	Bad data
	4.5		47.65	-122.31	SEAT				Cycle slip	SEAT
NE	19:50:30	0~12	39.02	-76.83	GODE	32	36.31	107.11	254.60	Bad data
	6.5		39.13	-77.22	GAIT				Cycle slip	GODE
	01:04:30	12~20	39.02	-76.83	GODE	39	36.31	124.40	287.56	Bad data
	1.5		39.13	-77.22	GAIT				Cycle slip	GODE

07-15-2000 20:00-03:00 UT

Cluster	Time (UT) Duration (min)	El Band	Lat	Lon	STN	SVN #	d_stn (km)	Vert Slope (mm/km)	Slant Slope (mm/km)	L1 only (slant)
NE	00:21:30	0~12	39.02	-76.83	GODE	22	36.31	769.77	1848.94	Bad data
	6.5		39.13	-77.22	GAIT				Rcvr bias	GODE
	00:28:30	12~20	39.02	-76.83	GODE	22	36.31	703.08	1607.83	Bad data
	5.5		39.13	-77.22	GAIT				Cycle slip	GODE
	22:30:00	20~30	39.02	-76.83	GODE	25	23.6765	332.56	609.51	Bad data
	8.5		38.92	-77.07	USNO				Rcvr bias	GODE

10-29-2003 18:00-02:00 UT

Cluster	Time (UT) Duration (min)	El Band	Lat	Lon	STN	SVN #	d_stn (km)	Vert Slope (mm/km)	Slant Slope (mm/km)	L1 only (slant)
Wash. DC	00:38:00	0~12	40. 20	-75.06	PARL	38	30.34	108.01	246.02	26.96
	3		39.94	-74.95	NJTW					
	20:17:30	12~20	39.78	-75.12	NJGC	46	22.55	128.86	270.58	73.35
	13.5		39.94	-74.95	NJTW				Lock loss	
NC	21:4:00	30~45	35.79	-78.64	RALR	46	45.92	212.94	304.00	278.29
	0		35.40	-78.82	LILL					
Florida	20:17:30	20~30	29.30	-81.11	ORMD	38	69.21	174.61	341.10	41.93
	24		29.66	-81.69	PLTK				Rcvr bias	
	21:11:00	30~45	29.30	-81.11	ORMD	38	69.21	261.17	348.50	44.06
	300		29.66	-81.69	PLTK				Rcvr bias	
	23:40:00	45~90	29.30	-81.11	ORMD	38	69.21	347.1882	425.03	108.42
	156.5		29.66	-81.69	PLTK				Rcvr bias	
OH/MI	23:17:00	0~12	43.06	-82.69	AVCA	26	18.01	127.89	308.78	154.89
			42.91	-82.59	OKEE				Lock loss	Lock loss
	20:51:30	20~30	41.919	-84.02	ADRI	46	56.92	133.84	242.79	154.30
			41.61	-83.48	TLDO				Lock loss	Lock loss
	21:05:30	30~45	39.43	-84.28	LEBA	46	53.12	143.34	224.62	206.17
			39.02	-84.61	ERLA					

10-29-2003 18:00-02:00 UT

Cluster	Time (UT) Duration (min)	El Band	Lat	Lon	STN	SVN #	d_stn (km)	Vert Slope (mm/km)	Slant Slope (mm/km)	L1 only (slant)
LA/TX	23:50:00	0~12	29.79	-95.59	ADKS	29	14.89	108.78	240.25	Bad data
	2		29.78	-95.43	TXHU				Cycle slip	TXHU
	20:43:30	12~20	29.79	-95.59	ADKS	33	24.38	107.72	227.21	21.42
	6.5		29.79	-95.33	NETP				Lock loss	
	20:19:00	20~30	29.78	-95.43	TXHU	33	9.64	165.43	293.16	Bad data
	19.5		29.79	-95.33	NETP				Cycle slip	TXHU
	21:20:00	30~45	29.78	-95.43	TXHU	31	9.64	204.41	298.18	Bad data
	24		29.79	-95.33	NETP				Cycle slip	TXHU
	23:45:00	45~90	29.78	-95.43	TXHU	38	9.64	249.89	313.98	Bad data
	1.5		29.79	-95.33	NETP				Cycle slip	TXHU
OK/TX	21:52:00	0~12	33.85	-98.51	TXWF	26	154.34	91.29	212.25	Bad data
	1.5		34.98	-97.52	PRCO					TXWF
	22:15:00	12~20	34.5728	-98.41	OKLW	26	93.2039	109.51	227.85	190.46
	29		34.98	-97.52	PRCO					
	01:47:00	30~45	32.37	-96.86	TXES	44	46.9924	176.68	273.54	Bad data
	0.5		32.76	-97.06	TXAR					TXES
SEA/WA	20:59:00	12~20	47.69	-122.26	SEAW	43	5.4253	160.88	356.25	Bad data
	10		47.65	-122.31	SEAT				Lock loss	SEAT

10-30-2003 19:00-01:00 UT

Cluster	Time (UT) Duration (min)	El Band	Lat	Lon	STN	SVN #	d_stn (km)	Vert Slope (mm/km)	Slant Slope (mm/km)	L1 only (slant)
Florida	21:35:00	20~30	25.73	-80.16	AOML	37	0.29	16324.61	31864.91	4410.03
	35		25.73	-80.16	MIA3				9.5m bias	Rcvr bias
LA/TX	21:31:30	30~45	29.78	-95.43	TXHU	44	9.64	185.25	253.27	Bad data
	61		29.79	-95.33	NETP				Lock loss	TXHU
OH/MI	19:34:00	12~20	43.06	-82.69	AVCA	38	16.47	227.49	473.82	183.26
	2.5		43.04	-82.49	FRTG					
	19:41:00	20~30	43.04	-82.49	FRTG	38	16.27	104.74	205.88	297.38
	0		42.91	-82.59	OKEE				Lock loss	Lock loss
	19:56:30	30~45	43.06	-82.69	AVCA	43	16.47	183.619	245.34	87.78
	1.5		43.04	-82.49	FRTG					
SEA/WA	23:25:00	20~30	46.84	-122.26	CPXF	39	94.09	136.63	232.91	Bad data
	4		46.12	-122.90	KELS					(both)
	22:19:30	45~90	47.69	-122.26	SEAW	44	5.43	492.05	495.88	Bad data
	131		47.65	-122.31	SEAT				Lock loss	SEAT

11-20-2003 17:00-24:00 UT

Cluster	Time (UT) Duration (min)	El Band	Lat	Lon	STN	SVN #	d_stn (km)	Vert Slope (mm/km)	Slant Slope (mm/km)	L1 only (slant)
OH/MI	21:17:00	0~12	40.80	-81.96	WOOS	26	74.51	153.18	355.74	273.82
	7.5		41.42	-81.61	GARF					
	20:50:30	12~20	40.63	-83.61	KNTN	26	59.07	155.58	351.90	424.61
	17		40.31	-84.17	SIDN					
	19:51:30	20~30	40.55	-79.70	PIT1	44	25.14	233.55	396.06	197.12
	30.5		40.44	-79.96	PAPT				Lock loss	
	18:56:30	30~45	42.28	-83.34	HRUF	38	15.14	368.69	539.66	514.10
	0		42.17	-83.24	SIBY				Lock loss	
	21:02:00	45~90	40.38	-82.51	MTVR	44	65.42	319.17	374.86	299.39
	1		39.96	-83.05	COLB					
LA/TX	19:30:30.1	12~20	29.91	-95.15	LKHU	32	22.70	87.58	205.47	40.51
	0		29.79	-95.33	NETP				Lock loss	

11-20-2003 17:00-24:00 UT

Cluster	Time (UT) Duration (min)	El Band	Lat	Lon	STN	SVN #	d_stn (km)	Vert Slope (mm/km)	Slant Slope (mm/km)	L1 only (slant)
NC	18:49:00	20~30	35.56	-77.06	WASR	38	30.10	117.96	216.04	74.97
	7		35.83	-77.03	NCWI					
	21:45:30	30~45	35.56	-77.06	WASR	37	30.10	191.10	256.96	180.19
	0		35.83	-77.03	NCWI				Lock loss	
	21:00:30	45~90	35.40	-78.82	LILL	44	32.193	254.66	320.02	201.38
	0		35.47	-79.16	SNFD				Lock loss	
Wash. DC	20:36:30	0~12	39.78	-75.12	NJGC	29	77.23	119.53	282.07	Bad data
	1		39.16	-75.52	DNRC				Lock loss	DNRC
	20:53:30	20~30	39.78	-75.12	NJGC	37	22.55	138.44	254.45	70.90
	15		39.94	-74.95	NJTW				Lock loss	
	20:30:30	30~45	38.92	-77.07	USNO	44	27.38	408.87	587.01	197.44
	5.5		39.13	-77.22	GAIT				Lock loss	
	20:56:00	45~90	39.02	-76.83	GODE	38	23.68	294.79	314.37	Bad data
	5.5		38.92	-77.07	USNO				Lock loss	GODE

Tables 1-7. Summary of screening results for all known storm days. For each cluster and elevation range, the maximum spatial gradient yielded by the automatic screening for that day is presented, if above 200mm/km in slant differential delay.

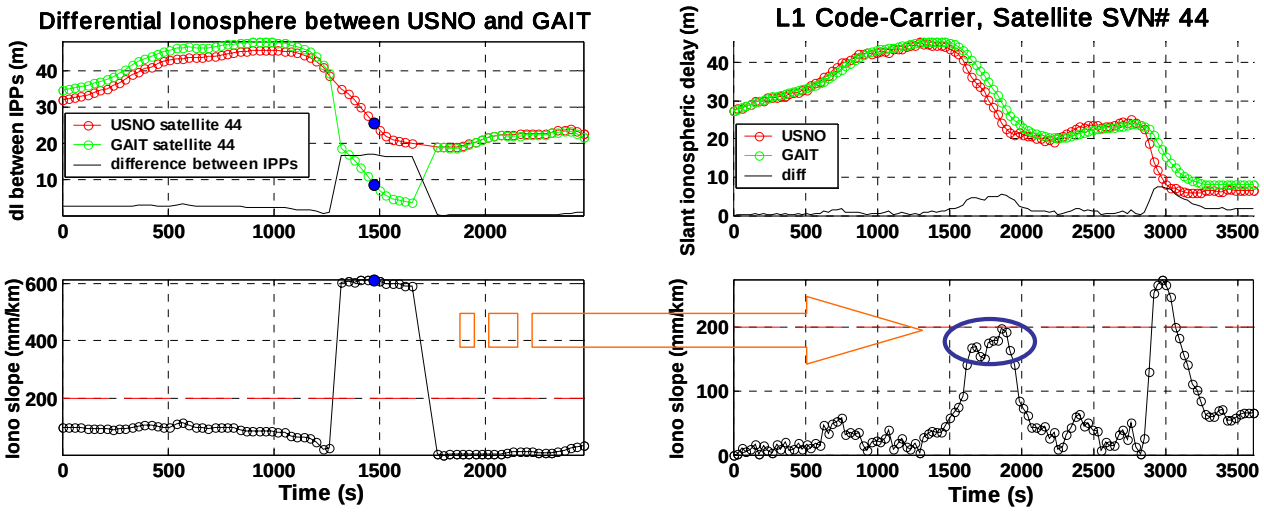

Cluster	Time (UT) Duration (min)	El Band	Lat	Lon	STN	SVN #	d_stn (km)	Vert Slope (mm/km)	Slant Slope (mm/km)	L1 only (slant)
Wash. DC	20:30:30	30~45	38.92	-77.07	USNO	44	27.38	408.87	587.01	197.44
Nov. 20	5.5		39.13	-77.22	GAIT				L2 Lock loss	

Chart 1. Using the L1-only measurements from RINEX data to eliminate the "jump" in the L1/L2 observations.

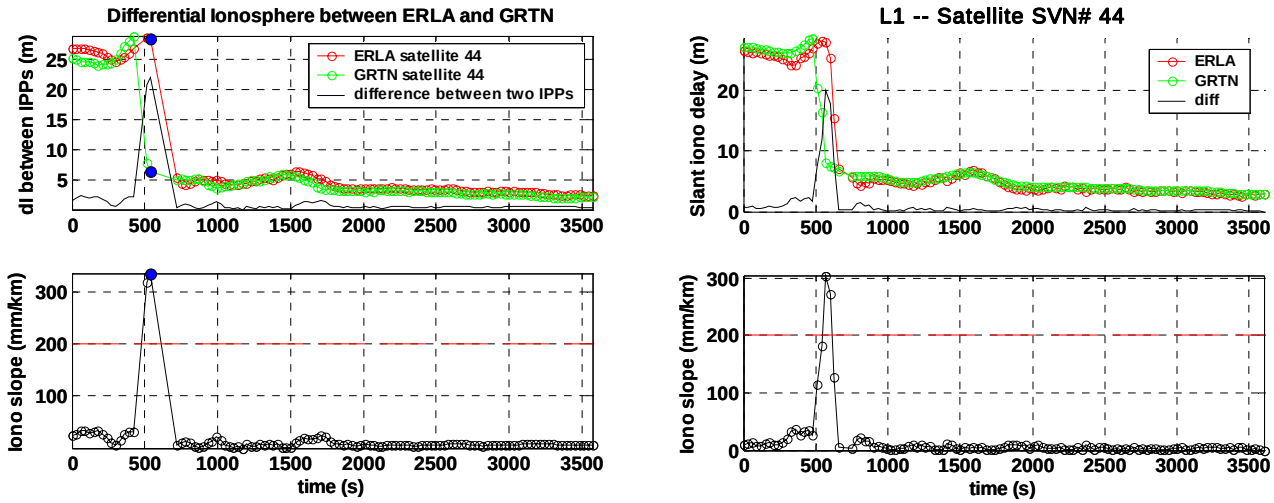

Cluster	Time (UT) Duration (min)	El Band	Lat	Lon	STN	SVN #	d_stn (km)	Vert Slope (mm/km)	Slant Slope (mm/km)	L1 only (slant)
OH/MI	21:09:30	45~90	39.02	-84.61	ERLA	44	65.58	289.98	335.57	304.58
Nov. 20	0.5		38.86	-83.88	GRTN				L2 Lock loss	

Chart 2. Use of L1-only raw data to confirm the presence of a large ionosphere gradient observed in L1/L2 (Nov. 20).

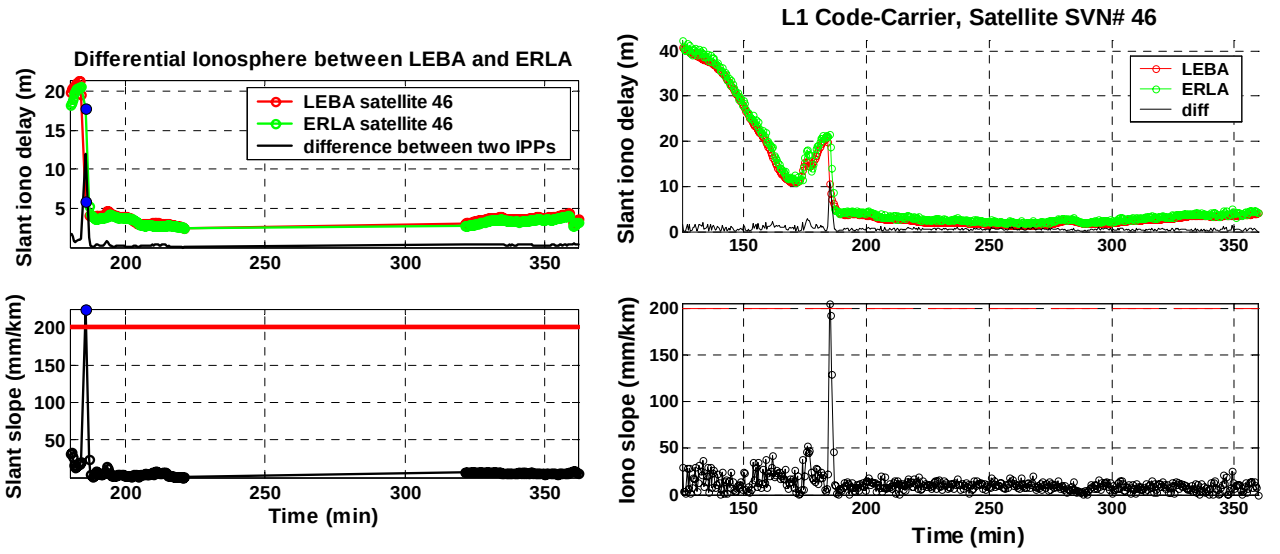

Cluster	Time (UT) Duration (min)	El Band	Lat	Lon	STN	SVN #	d_stn (km)	Vert Slope (mm/km)	Slant Slope (mm/km)	L1 only (slant)
OH/MI	21:05:30	30~45	39.43	-84.28	LEBA	46	53.12	143.34	224.62	206.17
Oct. 29	0.5		39.02	-84.61	ERLA				Lock loss	

Chart 3. Use of L1-only raw data to confirm the presence of a large ionosphere gradient observed in L1/L2 (Oct. 29).

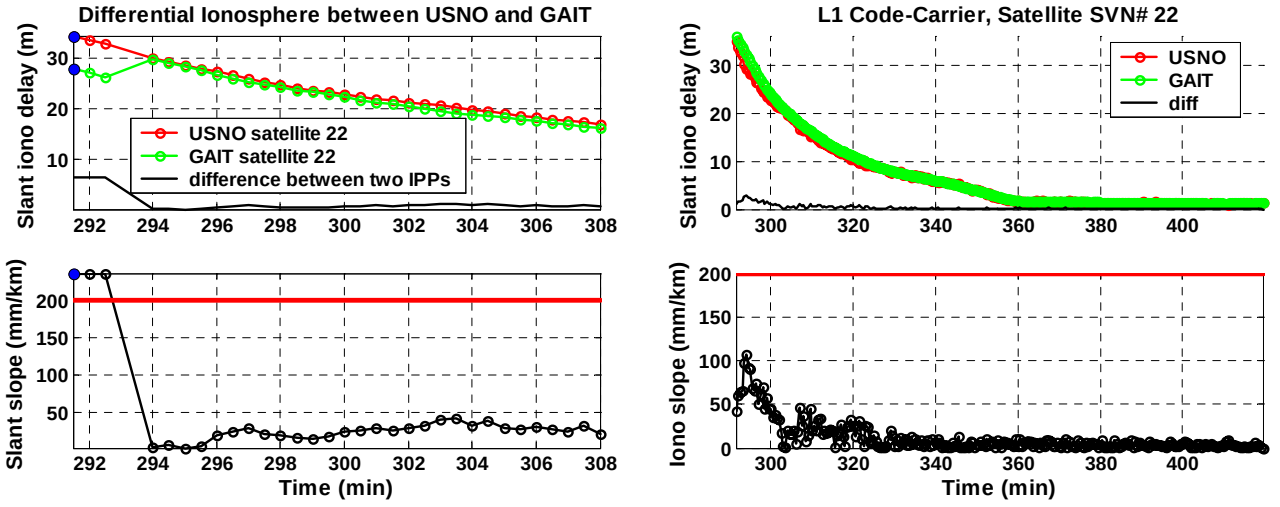

Cluster	Time (UT) Duration (min)	El Band	Lat	Lon	STN	SVN #	d_stn (km)	Vert Slope (mm/km)	Slant Slope (mm/km)	L1 only (slant)
NE	00:28:30	20~30	38.92	-77.07	USNO	22	27.38	128.00	237.22	106.70
Jul. 15	5.5		39.13	-77.22	GAIT				Cycle slip	

Chart 4. Example of cycle slip in L1/L2 causing an apparently anomalous gradient during year 2000 storms.