Protecting Against Unsampled Ionospheric Threats

Todd Walter, Sriram Rajagopal, Seebany Datta-Barua, and Juan Blanch
twalter@stanford.edu

Stanford University

Introduction

Nominally, the mid-latitude ionosphere sampled by WAAS is smooth and slowly varying. However, during ionospheric storms, small-scale irregularities may form. These irregularities cause rapid changes in Total Electron Content (TEC) over comparatively short distances. WAAS has observed changes greater than 20 meters of vertical delay at the GPS L1 frequency over a few hundred kilometers. Such features are very difficult to model with the single frequency SBAS correction scheme that uses a 5x5 degree grid. Some of these features exhibit significant curvature and may even fit entirely within a single grid cell. Worse yet, some of these features may exist, but not be sampled by any of the ground-based WAAS reference Stations (WRSs). In this case, they would be invisible to the system. The WAAS ionospheric algorithms must protect users under all conditions. The protection bound broadcast to the user, termed the GIVE, must be sufficiently large to guard against such undetectable features. Therefore, it is essential to understand the magnitude of such threats given the gaps in our observations.

The methodology used to study these features is to start with well-observed threats from previous storms. Then data is artificially removed or deprived from the set to create a gap in observability. The delay value is predicted for the removed data points based on the remaining ones. The observed error is then compared against metrics describing our observability. Previous work has focused on geometric deprivation schemes. Ionospheric Pierce Points (IPPs) or WRSs are removed based on their geographic location. We hope to have the resulting hole in coverage line up with the ionospheric feature of interest. However, this process is hit and miss. Additionally, it is arguable that is optimistic in the sense that we will likely fail to line up our worst-case feature with a worst-case hole in coverage.

Instead, we propose a new deprivation scheme that tailors itself to the feature of interest. Instead of stamping out geographic regions without regard to the ionospheric delay values, we remove the IPPs that are in greatest disagreement with the surrounding ionosphere. As these IPPs are the most helpful to alerting the WAAS storm detector to the presence of the feature, they are the most harmful points to remove. Consequently, we refer to this deprivation technique as malicious deprivation. This deprivation scheme leads to a removal of just those IPPs that describe the feature while leaving the undisturbed IPPs in our data set. This results in the smallest possible hole in coverage. We will describe the features that this deprivation scheme has found.

Ionospheric Threats

It has been observed that the ionosphere is not always smoothly varying [Hansen et al.] [Doherty et al.]. There may be steep gradients and isolated features whose delay values are substantially different from the surrounding ionosphere. Such features are a threat to WAAS users as they may experience large position errors if their raypaths pass through a feature while the paths from the reference stations do not. WAAS must broadcast a bound on the error the user may experience. This bound must take into account the possibility of threatening ionospheric features. Therefore, WAAS must determine the magnitude of the threat that can exist in between the measurements that it has.

WAAS uses the thin shell model, where the ionospheric delay is modeled as coming from a thin slab at a height of 350 km above the surface of the Earth. Nominally, this model is quite accurate, but it can break down during disturbed conditions when the ionosphere may create significant delay at different heights. Under this model, the ionosphere can be represented as two-dimensional. WAAS broadcasts corrections to the user as a grid of vertical delays and confidence bounds called Grid Ionospheric Vertical Errors (GIVEs). An Ionospheric Grid Point (IGP) is defined every five degrees. During storms however, the thin shell model is no longer accurate. The grid correction scheme can lead to large errors and the GIVEs must

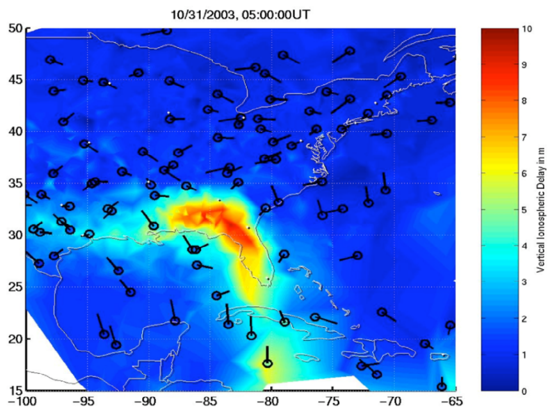

Figure 1. *Data from October 31, 2003 is used to simulate a severe threat to WAAS. Under this scenario, a small blob of ionosphere escapes detection of the system*

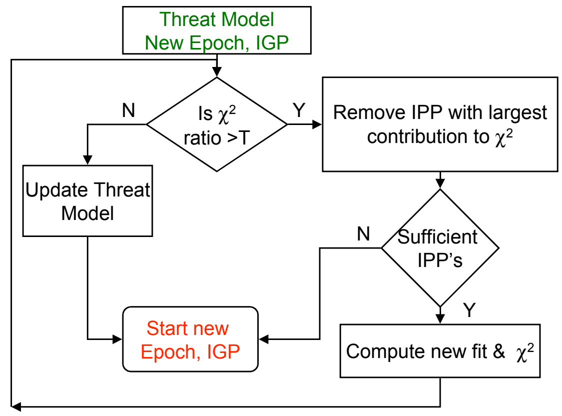

Figure 2. *This flowchart shows the process for malicious data deprivation where points are removed based on their deviation from a planar fit.*

be increased accordingly. WAAS employs an irregularity detector to determine when the ionosphere is in this disturbed state [Walter et al.]. This detector uses the chi-square value of a planar fit of the nearest IPPs. When the thin shell model is accurate, the chi-square value tends to be small, as the approximation breaks down, the chi-square value increases. When it crosses a threshold, WAAS sets the GIVEs to their maximum value (45 m) and waits for the ionosphere to recover. We have observed that it is in this time frame that the most significant unobserved ionospheric events are likely to occur.

Figure 1 is an example of an ionospheric threat to WAAS. The color information is derived from the CORS network [NGS CORS] using about 400 receivers. The black circles represent the samples of the ionosphere from the WAAS reference Stations (WRSs). Each circle depicts one IPP. The lines or "tails" in the circles point back to the WRS and the length of the line provides an indication of elevation angle (the longer the line the lower the angle). We have removed five WAAS IPPs to better illustrate the concern. This feature, which is more than eight meters different from the background ionosphere, can fit entirely within a gap in coverage. Although the gap in this instance was artificially created, equally sized holes exist naturally elsewhere. Since the remaining IPPs do not sample the feature, they are entirely unaware of its existence. The predicted ionospheric delay for this region would be below two meters and the confidence would nominally be very tight as there appears to be good sampling and all of the observed IPPs are consistent with the two meter prediction. Only the addition of an undersampled threat term [Sparks et al. 2001] will create a bound that is sufficiently large to protect against this case. It is necessary for WAAS to add a term that characterizes the quality of the IPP sampling of the ionosphere to the final GIVE calculation. Several metrics have been used to quantify the sampling of the IPP [Sparks et al. 2003]. This paper is not concerned with the metric but rather in finding and defining the threat.

Data Deprivation

The only way that WAAS can evaluate a threat is to examine data that has been affected by it. If an actual threat had occurred between all our measurements, we would have no way of knowing. Therefore, we must first sample the threat to be aware of its existence. Once it has been sampled, we need to simulate the effect of undersampling. We achieve this through data deprivation. IPPs that sample the feature are removed from the data. Previous methods have focused on either removing IPPs geometrically, or by reference station and satellite. An example of geometric deprivation would be to remove all IPPs west of the IGP in question. This would be done to examine the effect of a large gradient moving in to the edge of coverage. However, such a deprivation scheme is not representative of an actual IPP distribution. At the edge of coverage the number of IPPs taper off more gradually. In addition, the gradient may not line up with the deprivation. It may occur several degrees to the East or West, or it may have a different alignment or shape. There is no guarantee that the remaining set of IPPs will fail to sample the feature.

Reference station and satellite removal have the advantage that they result in more realistic IPP distributions. By removing the westernmost stations, for example, the edge of coverage can be simulated as moving eastward relative to the ionosphere. Ideally, we would like to create a complete three-dimensional model of the ionosphere at each instant in time that can be shifted and rotated against the IPP geometries we are likely to have. Reference station and satellite deprivation only achieve a coarse approximation to this goal. The reference stations are hundreds of kilometers apart. A worst-case feature might occur by shifting a feature only tens of kilometers or by rotating the storm features relative to the sampling. Although this type of deprivation is preferable to geometric removal, it also may fail to correctly line up with a severe disturbance resulting in the worst-case sampling.

The stated goal of FAA certification was to bound the ionospheric error given an arbitrary sampling. Thus for any set of IPPs, no matter how we arrived at them, we want to predict the ionosphere in a specified region and then bound the possible error on this prediction. Returning to Figure 1, if we had the sampling of the ionosphere shown by the black circles then we need to make sure that the error bound would be sufficient to cover the eight meter position error.

Malicious Deprivation

The deprivation scheme that we seek is to remove only the points that sample ionospheric threat. Such points alert us to its presence. The worst case deprivation would leave as many quiet IPPs as possible in the fit such that we would appear to have a good sampling, but remove the points that would trip the chi-square detector. The previous deprivation schemes are specified independently of the measured ionospheric delay value. What we will do with malicious deprivation is remove the points that have the greatest residual from the planar fit. These points increase the chi-square value the most. As we remove such points, the remaining IPPs are more consistent with the resultant planar fit and we are more likely to believe that we have a good fit to quiet ionosphere.

Figure 2 shows the flow diagram for how malicious data deprivation is implemented. At every epoch, and for every IGP, the planar fit is calculated, as is the chi-square value. If the chi-square value is over the storm detector threshold, the IPP with the largest normalized residual is removed from consideration. For WAAS it is then replaced with the next nearest IPP, provided it is not too distant from the IGP. The planar fit and chi-square values are recomputed. This process continues until the chi-square value falls below the threshold or there are insufficient remaining IPPs to compute a fit. In this manner, the IPPs that are least consistent with the background ionosphere are removed from the dataset. Malicious data deprivation automatically forms a glove fit around whatever disturbance exists. The resulting IPP distribution is essentially the worst possible sampling of the ionosphere: the maximum number of IPPs that fail to detect the presence of a disturbance.

We have explored some variations on malicious deprivation. One concern that arose was that the IPPs removed might not be physically connected by an ionospheric event. To address this issue we added a constraint that once three or more IPPs we excluded by malicious all IPPs within the convex hull of the excluded IPPs were removed as well. This would force a region to be removed rather than assorted intermixed IPPs. However, we found that only rarely did malicious deprivation experience this problem, and these cases were always minor compared to the more serious threats such as illustrated in Figure 1. Since the convex hull calculation is time consuming we do not typically employ it, however we do manually inspect each significant threat to ensure that it remains unnecessary.

Another variation is called continued malicious deprivation. Here we do not stop the loop in Figure 2 once the chi-square value falls below the threshold. Instead, we always continue until there are not sufficient IPPs to perform the fit. The rationale behind this is that some algorithms use the chi-square value to continuously adjust the GIVE, rather than just make a binary decision of good or bad. In this case removing a few more points and lowering the chi-square value may result in a worse threat than keeping them.

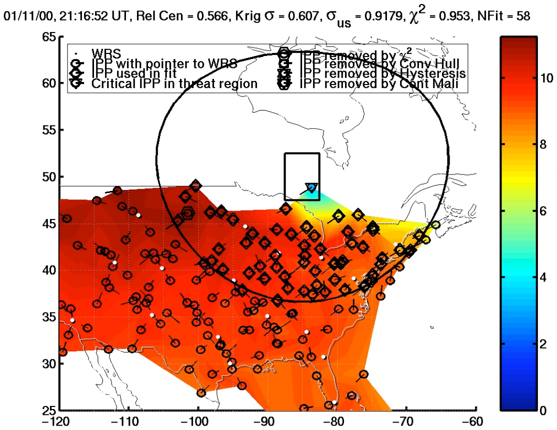

Figure 3. *January 11, 2000: Most of CONUS sees elevated TEC (between 8 - 10 m). However, to the North, there is sharp drop to below 5 m. This feature is sampled by only one IPP.*

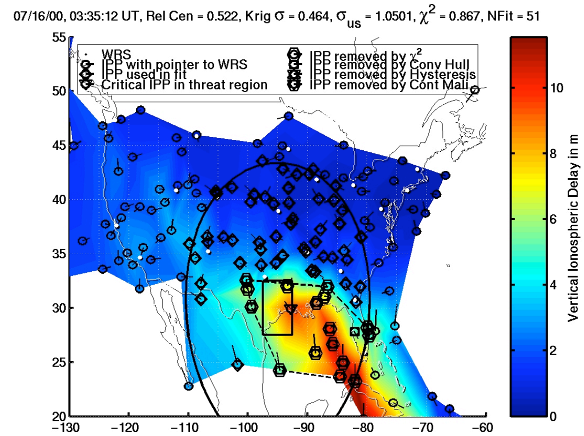

Figure 4. *July 16, 2000: The nighttime ionosphere is quiet for the majority of CONUS, except an isolated region near the Gulf of Mexico experiences a 10 m increase.*

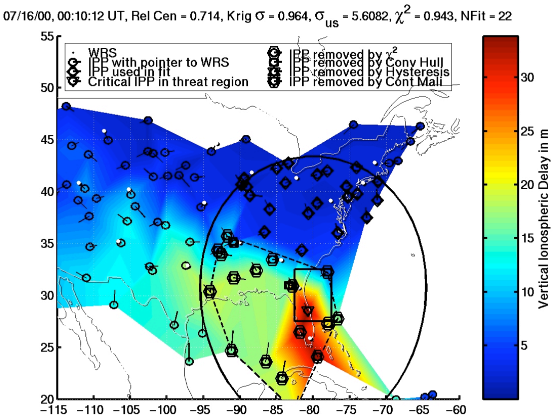

Figure 5. *July 16, 2000: A region of ionosphere with vertical TEC greater than 30 m (185 TECU) is seen in the Southeastern U.S. The Northern part of the country sees delays below 5 m.*

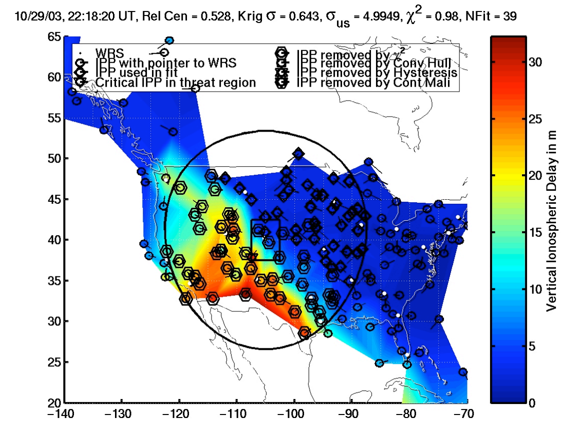

Figure 6. *October 29, 2003: A very strong gradient can be seen in the central part of CONUS. Vertical TEC changes by nearly 30 m over 5 degrees.*

Threats Identified

Malicious deprivation is very good at identifying the types of threats that are most harmful to an SBAS. Finding and identifying such threats are essential to then designing an algorithm to protect against them. We identify here some of the worst threats WAAS has found in its storm data sets. These sets span the timeframe from October 1999 to October 2004 and include the most recent solar peak and its declining side. These figures are only a representative set, as each storm contains many more threats that are similar

Figure 3 shows data from January 11, 2000 at 21:17 UTC. The daytime peak is just starting to decline when a depletion becomes just observable to the north over Canada. There is a greater than 5 m drop in just a few degrees of latitude. This threat demonstrates that extrapolating ionospheric delays beyond the measurements requires a sufficient increase in the GIVE. It also shows the reality of some of these threats as only a single IPP in the data set sampled the feature at this time.

Figure 4 shows data from July 16, 2000 at 3:35 UTC. An enhancement over the Gulf of Mexico can be seen with a delay value in excess of 10 m while the quiet background ionosphere is closer to a nighttime normal of one or two meters. Although some quiet IPPs extend around the feature, it requires a fairly large hole in coverage to escape detection.

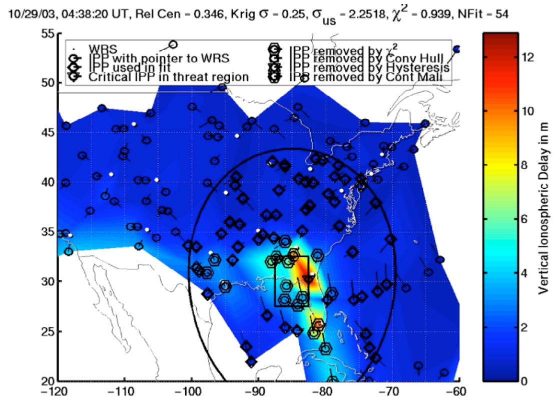

Figure 7. *October 31, 2003: The worst ionospheric threat to SBAS. Within a single 5 degree cell the is an isolated 10 m blob of ionosphere*

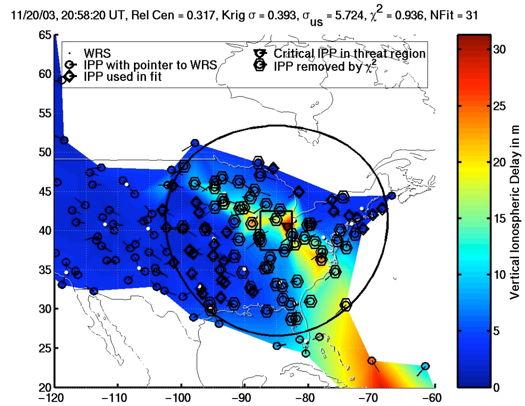

Figure 8. *November 20, 2003: A finger of very significant delay extends through CONUS with relatively quiet ionosphere on either side.*

Figure 5 shows data from a little over two hours earlier in the day on July 16, 2000. Although the feature is now larger geographically, and harder to miss, it also is substantially larger in magnitude with a 30 m vertical gradient over just a few degrees. Without the measurements in the south, the user could end up with a very significant positioning error.

Figure 6 displays a very strong gradient through the center of CONUS. Again a 30 meter vertical difference occurs over a few degrees of longitude. Although the disturbance extends slightly into the lower delay values, there is a fairly sharp transition from the quiet ionosphere of a few meters to over thirty. This emphasizes the danger of extrapolating the planar fit outside of where one has measurements.

Figure 7 displays the worst ionospheric threat identified so far. This is the same feature as shown in Figure 1, but about 20 minutes earlier. Although it is also similar in appearance to Figures 4 and 5, here the feature is much more localized. The transition from low to high to low occurs in just over 5 degrees of longitude. Additionally, there are quiet measurements very close in to the feature on nearly all sides. As Figure 1 illustrates, it does not take a very large hole in coverage to miss a feature this size. The change in vertical delay is more than ten meters, which means that it cannot be protected by a 6 m GIVE value. Due to message quantization the next interval is 15 m. This has a devastating impact on availability if it is assumed to be present at any time.

Figure 8 shows the ionosphere from November 20, 2003 at 21:00 UTC. There is a finger of enhanced density extending from the southeast coast to the northern border. Again, the change in delay is very large over a short distance. Although malicious deprivation found it is possible to take IPPs from either side of the finger to form a smooth planar fit, the size of the excluded region in the center is quite large. This feature is not nearly as threatening as the previous one.

Malicious data deprivation was run on 28 storms days identified as likely to create threats. On nominal days, with no ionospheric disturbances the malicious data deprivation finds no threats that are not already covered by the formal error. Even on nine of the minor storm days investigated, no threats were found. These days are:

- October 22, 1999
- February 12, 2000
- June 8, 2000
- August 11, 2000
- April 1, 2001
- September 4, 2002
- September 7, 2002
- September 11, 2002
- May 31, 2003

Of the remaining storms, only the first five periods listed below had extreme threats. The remaining days contain significantly reduced threats compared to these first five. The threatening periods in approximate order of reducing importance are as follows:

- October 29-31, 2003
- November 20, 2003
- July15-16, 2000
- April 6-7, 2000
- March 31, 2001
- October 21, 2001
- November 6, 2001
- November 24, 2001
- January 11, 2000
- May 29-30, 2003
- September 8, 2002
- November 22, 2003
- April 18, 2002

Thus the ionosphere over CONUS is quiet the vast majority of the time, even for solar maximum years. However, a few times a year during solar max, storms arise which can create very large threats. The worst of these for the most recent peak are identified here.

Conclusions

During the worst ionospheric storms localized ionospheric irregularities can occur. In the very worst case, these features can escape detection altogether. The WAAS user must be protected from all such threats. The first step is to identify what features can exist. Malicious deprivation is a tool that, given a set of IPPs, can automatically identify the worst-case sampling. By examining the resulting threats, we can determine the worst-case features and finally specify the extent of the threat model. With this knowledge the ionospheric estimation algorithms and, more importantly, the GIVEs can be determined.

One limitation of malicious deprivation is that it is a snapshot approach. It takes no credit for prior observations. It also does not address the question of how likely the feature is to be missed. Instead it addresses the more pessimistic question of how much protection does the worst-case sampling require. The GIVE algorithm should try to take advantage of previous measurements, as the result shown in Figure 7 would otherwise preclude availability in any but the most densely sampled region. Such algorithms are beyond the scope of this paper.

A listing of the most severe recent storms is presented along with an approximate ranking. In addition, several minor storms are shown to produce no threats at all. Except for a few hours on each of a few days, no unsampled threats exist. With this new deprivation scheme we can very confidently rule out days with no features as we know that even under the worst case sampling, no threat is present. Malicious deprivation represents a very useful tool in sorting out data sets and ordering the few remaining threats.

Acknowledgements

This work is sponsored by the FAA Satellite Product Team. The authors would also like to thank Attila Komjathy of JPL for processing the CORS data.

References

Doherty, P., Delay, S., Walter, T., Dehel, T., McHugh, T., and Bunce, D., *Space Weather Effects on WAAS: A Performance and Status Report*, in Proceedings of the Beacon Satellite Symposium, Trieste, Italy, October 2004.

Hansen, A., T. Walter, J. Blanch, and P. Enge, *Ionospheric Spatial and Temporal Correlation Analysis for WAAS: Quiet and Stormy*, in Proceedings of ION GPS-2000, Salt Lake City, UT, September 2000, pp. 634-42.

National Geodetic Survey (NGS) Continuously Operating Reference Stations (CORS) http://www.ngs.noaa.gov/CORS/

Sparks, L., Pi, X., Mannucci, A.J., Walter, T. Blanch, J., Hansen, A., Enge, P., Altshuler, E., and Fries, R., *The WAAS Ionospheric Threat Model* in Proceedings of the Beacon Satellite Symposium, Boston, MA, June 2001.

Sparks, L., Komjathy, A., Mannucci, A.J., *The Dependence of WAAS Ionospheric Error Bounds upon the Spatial Distribution of GPS Measurements*, in proceedings of ION National Technical Meeting, Anaheim, CA, January 2003.

Walter, T., Hansen, A., Blanch, J., Enge, P., Mannucci, T., Pi, X., Sparks, L., Iijima, B., El-Arini, B., Lejeune, R., Hagen, M., Altshuler, E., Fries, R., and Chu, A., *Robust Detection of Ionospheric Irregularities*, Navigation: Journal of The Institute of Navigation, Vol. 48, No. 2, Summer 2001, pp. 89-100.