Characterizing the Dependence of Satellite-based Augmentation Systems upon the Spatial Distribution of GPS Measurements

Lawrence Sparks
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA, USA

Co-Authors:
A. Komjathy and A.J. Mannucci
Jet Propulsion Laboratory, California Institute of Technology, USA

1 Abstract

Satellite-based augmentation systems (SBAS) such as the United States' Wide Area Augmentation System (WAAS), the European Geostationary Navigation Overlay Service (EGNOS), and the Japanese Global Navigation Satellite System (MSAS) are designed to improve the accuracy and ensure the integrity of user position estimates determined from Global Positioning System (GPS) measurements. In the absence of selective availability, the ionosphere is the largest source of positioning error for single-frequency users of the GPS signals. A critical objective of any SBAS system is to estimate ionospheric delays accurately and to bound the errors in these estimates reliably. Irregularities in the ionosphere that go undetected can result in significant user delay error. Thus, the accuracy of a given delay estimate can depend sharply on the spatial distribution of GPS measurements from which the delay estimate is derived. This paper presents a method for characterizing in terms of a scalar metric the degree to which a given spatial region is sampled densely and uniformly by a given set of GPS measurements. The metric is used to assess the danger posed to WAAS delay estimate accuracy by undersampled ionospheric irregularities.

2 Introduction

In SBAS systems such as WAAS (see RTCA, 1996), an ionospheric delay error and its confidence bound at a user location are derived from vertical ionospheric delay estimates, modeled at a set of regularly-spaced intervals in latitude and longitude, *i.e.*, at *ionospheric grid points* (IGPs). The vertical delay estimate at each IGP is calculated from a planar fit of neighboring slant delay measurements, projected to vertical using the standard *thin shell model*. The points where the measurement raypaths cross the ionospheric shell height are known as *ionospheric pierce points* (IPPs). When the spread of IPPs surrounding a given IGP is highly skewed, a region near or to one side of the IGP may be undersampled, allowing a significant ionospheric disturbance to remain undetected.

To protect the user from the danger posed by undersampled irregularities, WAAS augments the error bounds that are broadcast for each vertical delay estimate at an IGP. The amount of this augmentation is derived from an *ionospheric threat model.* The spatial threat model is derived by systematically excluding data from planar fits of vertical delay measurements within a neighborhood of each ionospheric grid point (IGP). The excluded data are then used to evaluate the maximum fit residual as a function of two metrics that quantify the spatial spread of the measurements: (1) the radius of a circle, centered at the IGP, large enough to encompass the IPPs used in the fit, and (2) the ratio of the IPP centroid radius to the fit radius, where the IPP centroid radius defines the distance from the IGP to the centroid of the fit IPPs. This latter ratio, the *relative centroid,* is meant to characterize the degree to which the IPPs are distributed *uniformly* across the fit region.

The primary purpose of defining a spread metric is to identify IPP configurations that might fail to sample a significant ionospheric irregularity. The extent to which a given metric is successful in achieving this goal has a direct impact on the magnitude of the ionospheric error bounds or *grid ionospheric vertical error* (GIVE). Since each user must be protected from the effects of poor sampling, any sampling metric that tends to confuse "better" and "worse" IPP distributions will cause the "better" distributions to be treated too conservatively, *i.e.,* when the broadcast error bound depends upon the value of the spread metric, a higher error bound will be broadcast than is actually warranted. Using the relative centroid as a spread metric has proven critical to successful WAAS operation.

In this paper, we discuss a new spread metric that offers, in contrast to the relative centroid, a greater sensitivity to the nature of the IPP distribution (note: much of the material presented here has appeared previously – see Sparks *et al.*, 2003). This spread metric is designed to meet the following requirements: (1) the metric should be sensitive to the angular distribution of the IPPs about the IGP; (2) the contribution of a single IPP to the metric defined at a given IGP should decrease with separation distance; (3) the metric should improve monotonically as the number of IPPs increases; and, (4) the parameterization of the metric should permit control of its sensitivity to the variation of a single IPP location. The metric varies in magnitude between 0 and 1, where 0 represents a "good" IPP spread and 1 a "poor" IPP spread. Proper choice of the free metric parameters depends upon the range in the number of points to be included in each fit and the range in the size of fit radii centered at each IGP.

We report the results of an initial study that has generated a spatial threat model using the new metric. We find that the contribution to the GIVE that protects against the threat of an undersampled irregularity can be reduced, on average, by at least 32%. Hence, we conclude that the new spread metric can aid SBAS performance by helping to define error bounds that are safe but not overly conservative.

In the sections that follow, we first review how WAAS currently deals with the threat of undersampled ionospheric irregularities. After discussing the limitations of the relative centroid as a spread metric, we describe the new spread metric and its dependence on two fixed parameters. Finally we provide an example where the new metric has been incorporated into the spatial threat model, and we discuss the consequent reduction in broadcast error bounds.

3 Assessing the threat posed by undersampled irregularities

This section provides brief review of the methodology we use to assess the threat posed by undersampled ionospheric irregularities. Specifically, we review (1) the algorithm used to estimate vertical delay at ionospheric grid points, (2) the metrics WAAS currently uses to characterize IPP distributions, and (3) the threat model used to bound the error posed by undersampled ionospheric irregularities.

3.1 Estimation of vertical delay at IGPs

The current WAAS algorithm for estimating the local vertical delay at each IGP is based upon the thin shell model of the ionosphere (see, for example, Mannucci *et al.*, 1999, Birch *et al.*, 2002). By treating the ionosphere as if it were collapsed into a shell at a specified height h, a slant delay measurement S may be related to an estimate V of the vertical delay at the IPP by a simple geometric factor:

$$S = M(\alpha, h) \cdot V \tag{1}$$

where

$$M(\alpha, h) = \left[1 - \left(\frac{R_e \cos\alpha}{R_e + h}\right)^2\right]^{-1/2} \tag{2}$$

is the thin-shell obliquity factor dependent upon h, the earth radius R_e, and the elevation angle α.

WAAS IGPs are spaced uniformly at 5° intervals in the conterminous United States (10° in Alaska). To estimate the vertical delay at an IGP, measurements with IPPs near the IGP are projected to vertical and fit to a plane. All measurements with IPPs that lie within a minimum fit radius R_{min} are included in the fit. If the number of such measurements is less than N_{pts}, the fit radius is extended until it defines a circle that surrounds N_{pts} points. If a circle with a maximum fit radius of R_{max} fails to encompass N_{pts} points, the fit is performed with fewer points, provided that at least N_{min} lie within R_{max}.

3.2 Characterizing the spatial distribution of IPPs near IGP

The estimate of the local vertical delay at the IGP can be highly inaccurate when the distribution of IPPs in the vicinity of the IGP is poor, *i.e.*, when this distribution contains large gaps or the bulk of the IPPs lie to one side of the IGP (as will often be the case when dealing with IGPs near the edge of the WAAS grid). Thus it proves useful to define metrics that characterize the spatial distribution of the IPPs included in the fit. A useful set of metrics should be highly sensitive to the density and the uniformity of the IPP coverage in the fit domain.

As noted previously, WAAS currently uses two metrics to characterize an IPP configuration: the fit radius, and the ratio of the IPP centroid radius to the fit radius. As the fit radius increases, the IPP density decreases, which permits larger irregularities to escape detection. The fit radius may thus be regarded as a rough indicator of the mean IPP density. The relative

centroid serves to measure the degree of uniformity provided by the IPP sampling. When IPPs are spaced uniformly throughout the fit domain, this metric vanishes; on the other hand, it approaches unity when IPPs congregate near a single point at the edge of the fit domain.

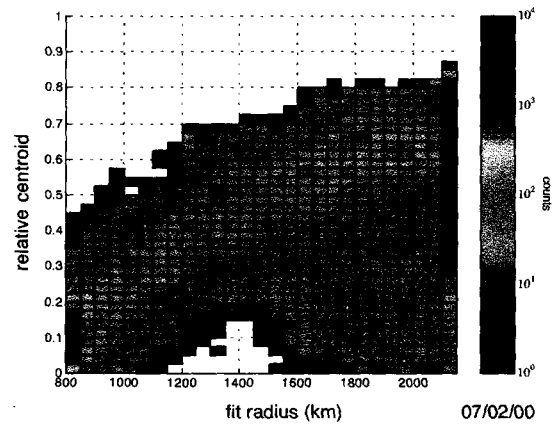

Figure 1. The distribution of relative centroid and fit radius for IPP configurations corresponding to slant TEC measurements recorded by WAAS receivers on July 2, 2000.

Figure 1 shows a typical distribution of metric values for fits at WAAS IGPs over the course of one day. In the analysis N_{pts}, N_{min}, R_{min}, and R_{max} have been set to representative values: 30 points, 10 points, 800 km, and 2100 km, respectively. (Note: optimal values for these parameters have yet to be determined.) The distribution is characterized by broad peaks except when the fit radius is near R_{min}. Note that all tabulated counts concerning fits that include less than N_{pts} points (*i.e.,* the fit radius is R_{max}) are relegated to the final column.

3.3 The spatial threat model

To place an upper bound on the effects of undersampled irregularities, WAAS uses a spatial threat model that is designed to quantify a set of worst-case errors according to their dependence upon the IPP distribution. Since the threat due to undersampled irregularities is likely to be largest under disturbed conditions, such a threat model is developed using data recorded on days when significant ionospheric disturbances have occurred. The procedure for generating a threat model is summarized here (for more details, see Sparks *et al.,* 2001).

For a given epoch and a given IGP, a subset of the measurements is selected to represent possible user measurements, and these data are withheld from the fit. Subsequently the excluded measurements are projected to vertical and compared to the corresponding estimated values based upon the planar fit. Fit errors are tabulated only when *non-storm* (or *near-storm*) conditions exist, that is, only for fits that fail to trigger the storm detector. (The storm detector is based on a χ^2 "goodness-of-fit" statistic: local storm conditions are declared whenever the χ^2 of the planar fit exceeds a specified threshold; see Walter *et al.*, 2000).

This process is repeated for many other choices of subsets of excluded data. Two types of data deprivation domains are used: annular and three-quadrant. Fits that exclude data from annular regions about the IGP allow us to assess the impact of localized spatial irregularities. A second deprivation domain excludes data from three quadrants of a rectangular (latitude-longitude) grid, whose origin lies on a diagonal passing through the IGP in question. Use of

this type of domain serves to quantify errors that arise when the IPP distribution is skewed to one side of the IGP.

The data set used in the analysis consists of six days of slant delay measurements collected by the existing 25 WAAS Reference Stations on the following days: January 11, 2000, April 6, 2000, April 7, 2000, July 15, 2000, July 16, 2000, and March 31, 2001. On each of these days, the ionosphere experienced one or more strong disturbances. These data are first post-processed (1) to eliminate interfrequency biases, (2) to remove the effects of cycle slips in carrier phase measurements, (3) to level the carrier phase measurements to the corresponding range measurements, and (4) to filter spurious measurements by means of the redundancy provided by multiple receivers at each station. The resulting data, designated *supertruth*, contain minimal error due to measurement noise and range multipath.

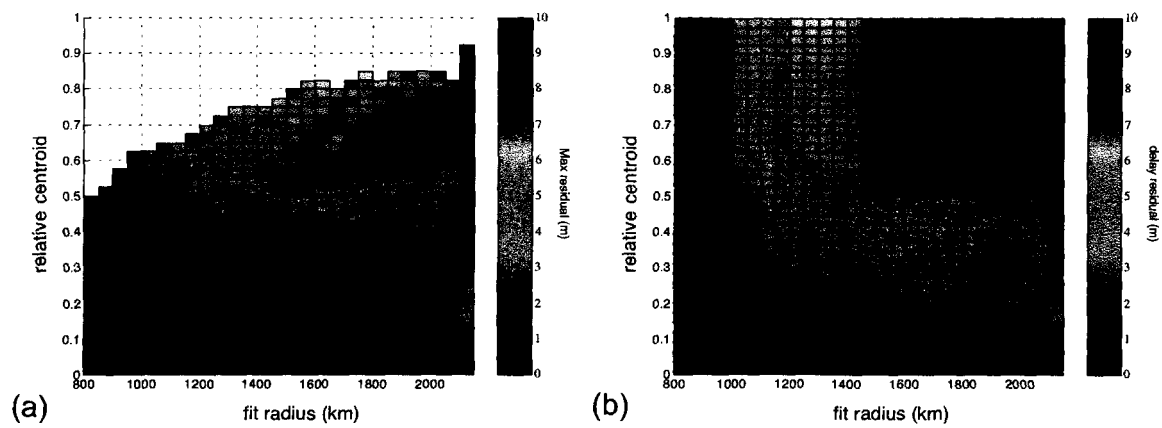

Figure 2. (a) Maximum fit residual and (b) overbound as functions of relative centroid and fit radius using the data deprivation algorithm applied to supertruth data from the following dates: January 11, 2000, April 6, 2000, April 7, 2000, July 15, 2000, July 16, 2000, and March 31, 2001.

Figure 2a displays pixel colors that represent the maximum error that occurs for a given fit radius and relative centroid. Only points that lie within a 5° by 5° *threat domain* centered on the IGP are represented. Again note that all tabulated errors concerning fits that include less than N_{pts} points are relegated to the final column. As expected, the behavior exhibited in the final column indicates that, as the number of points used in the fit drops below N_{pts}, the probability of large irregularities going undetected increases dramatically.

Figure 2b presents an overbound of the previous figure, where the maximum delay error is required to be a monotonically increasing function of each metric and the resulting bounds have been increased by 10% to provide padding. From such an overbound WAAS determines the amount by which the broadcast error bounds must be augmented to protect against the possible presence of undetected irregularities. The mean value of this augmentation is determined by the frequency with which each region of the overbound is accessed, as specified by an IPP distribution such as that given in Fig. 1.

4 New spread metric

Use of the relative centroid as an IPP spread metric has proven critical to permitting WAAS to meet performance specifications. Nevertheless, this metric possesses certain attributes that

restrict its ability to distinguish IPP configurations that may fail to sample significant ionospheric irregularities. After enumerating these attributes, a new metric is proposed in this section that surmounts these limitations.

4.1 Deficiencies of relative centroid as a spread metric

A spread metric based upon the centroid of the IPP distribution will not be particularly sensitive to the angular distribution of IPPs about the IGP. For example, Fig. 3 shows six distinct IPP configurations that all have the same centroid (an IGP is assumed to coincide with the origin of each set of coordinate axes). Clearly the sampling of the region in the vicinity of the IGP is very different in each case, and the size of an irregularity that could go unsampled varies considerably from one case to another.

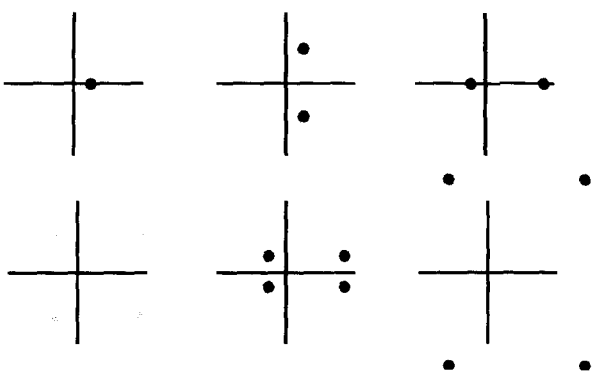

Figure 3. Six distinct IPP configurations that all have the same centroid.

Another limitation of a spread metric based upon the IPP centroid is that the contribution of a single IPP to the metric is complex and depends upon the number and placement of the other IPPs in the fit. In particular, adding a single IPP can make the coverage appear worse, *i.e.*, it can increase the magnitude of the centroid radius. Finally, there are no free parameters that allow one to adjust the sensitivity of the metric to the movement of a single IPP.

4.2 Definition of new spread metric

To overcome these limitations, we have defined a new metric to which each IPP in the fit contributes a distinct term. Let each IPP determine a pie-shaped wedge such that the line passing through the IPP and the IGP bisects the wedge (see Fig. 4). Specify that the width of the wedge diminish as the IPP moves away from the IGP. The metric then consists of the fraction of the area of the unit circle surrounding the IGP that is *not* covered by any of the pie-slices that constitute the fit. Good IPP coverage will then correspond to a metric value of zero, while poor coverage will have a metric value closer to unity. For example, the coverage represented in Fig. 4 corresponds to a metric value of approximately 0.3.

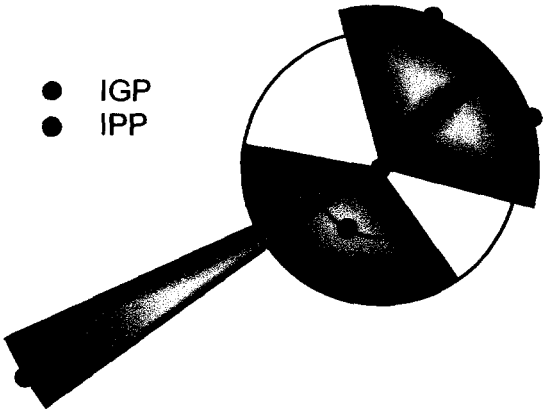

Figure 4. Schematic representation of the new spread metric. Each IPP defines a pie-shaped wedge that covers a portion of the unit circle. The metric is proportional to the area of the unit circle that remains uncovered.

To make this definition more quantitative, define the angular spread of each pie-slice to be

$$s_i = \frac{2\pi}{A_{spd} + r_i^2 / R_{spd}^2}, \tag{3}$$

where r_i is the distance from the IGP to the ith IPP (in km) and A_{spd} and R_{spd} are constants (A_{spd} is dimensionless; R_{spd} has units of km). We have chosen the asymptotic dependence of the metric to vary inversely with the square of the distance from the IGP. The spread metric at a given IGP is then

$$M_{spd} = 1 - \frac{1}{2\pi} \sum s_i, \tag{4}$$

where the summation is performed exclusive of overlap. Note that M_{spd} lies within the range

$$0 \le M_{spd} < 1. \tag{5}$$

Proper choice of the constants A_{spd} and R_{spd} will be problem-dependent. The primary considerations that govern the choice of these constants are (1) the range of the number of points used in fits, and (2) the range of the fit radius. An inappropriate assignment of these constants can lead to wedge widths that are either too small and, consequently, metric values will tend to congregate near unity, or too large causing metric values to vanish. Ideally a proper choice of these constants will give rise to metric values that vary over a significant fraction of the interval between zero and one.

A_{spd} controls the width of the wedge for an IPP that lies near the IGP. R_{spd} controls how rapidly the wedge width shrinks as the IPP moves away from the IGP. For two different choices of A_{spd} and three different choices of R_{spd} we have displayed in Fig. 5 the variation in the width of the pie slice as the IPP moves from the IGP to a distance of 2100 km from the

IGP. Notice that making A_{spd} and R_{spd} larger tends to make the spread metric less sensitive to IPPs near the IGP.

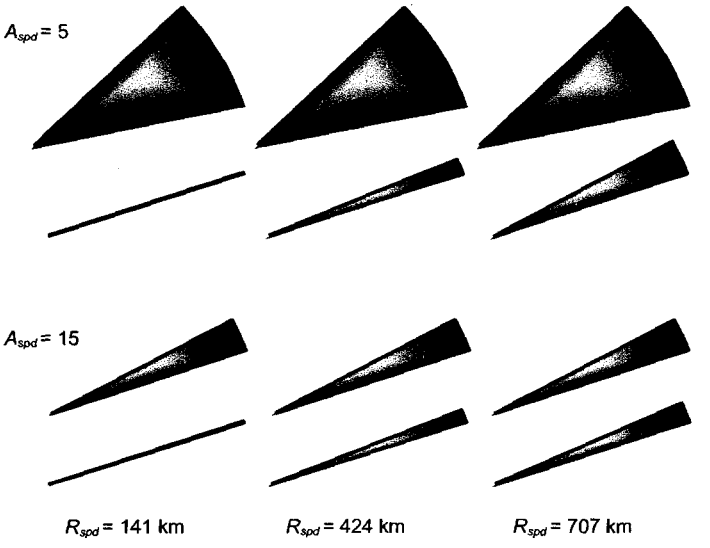

Figure 5. Ranges of IPP wedge widths for representative A_{spd} and R_{spd}. In each pair of wedges, the top wedge refers to an IPP near the IGP, and the lower wedge refers to an IPP at 2100 km.

By design this spread metric overcomes the limitations associated with the relative centroid. The metric is more sensitive to the angular distribution of the IPPs about the IGP than is the relative centroid. Furthermore, the contribution of a single IPP to the metric decreases with distance from the IGP, and the metric varies monotonically as the number of IPPs increases. Finally the parameterization of the metric permits control of its sensitivity to the variation of a single IPP location.

We use a simple but quick algorithm for computing the new spread metric, an algorithm that provides an approximation of its analytic value. First, we define a grid of N_{spd} spokes emanating from the IGP and spaced uniformly. (In practice we set N_{spd} = 360.) For each IPP, a pie-shaped wedge is calculated, and any spoke lying within the wedge is eliminated from the set. After all IPPs have been processed, the metric is computed as 1 minus the fraction of spokes that remain in the set.

5 Spatial threat model parameterized by new metric

The motivation for defining a new spread metric has been to identify more reliably IPP configurations that allow potentially significant ionospheric irregularities to go undetected. There are two distinct aspects to this problem. First, the choice of spread metric affects how IPP configurations are characterized (*e.g.,* see Fig. 1). Second, it affects how the fit residual overbound is characterized (*e.g.,* see Fig. 2b). Changing the spread metric does not change the magnitudes of the fit residuals to be tabulated (*e.g.,* Fig. 2a), but it does redefine the bins in which the counts are accrued. From the point of view of optimal SBAS performance, a *good* spread metric will tend to move the larger fit residuals to bins that are accessed less often. Such considerations will influence the optimal specification of constants in the new spread metric.

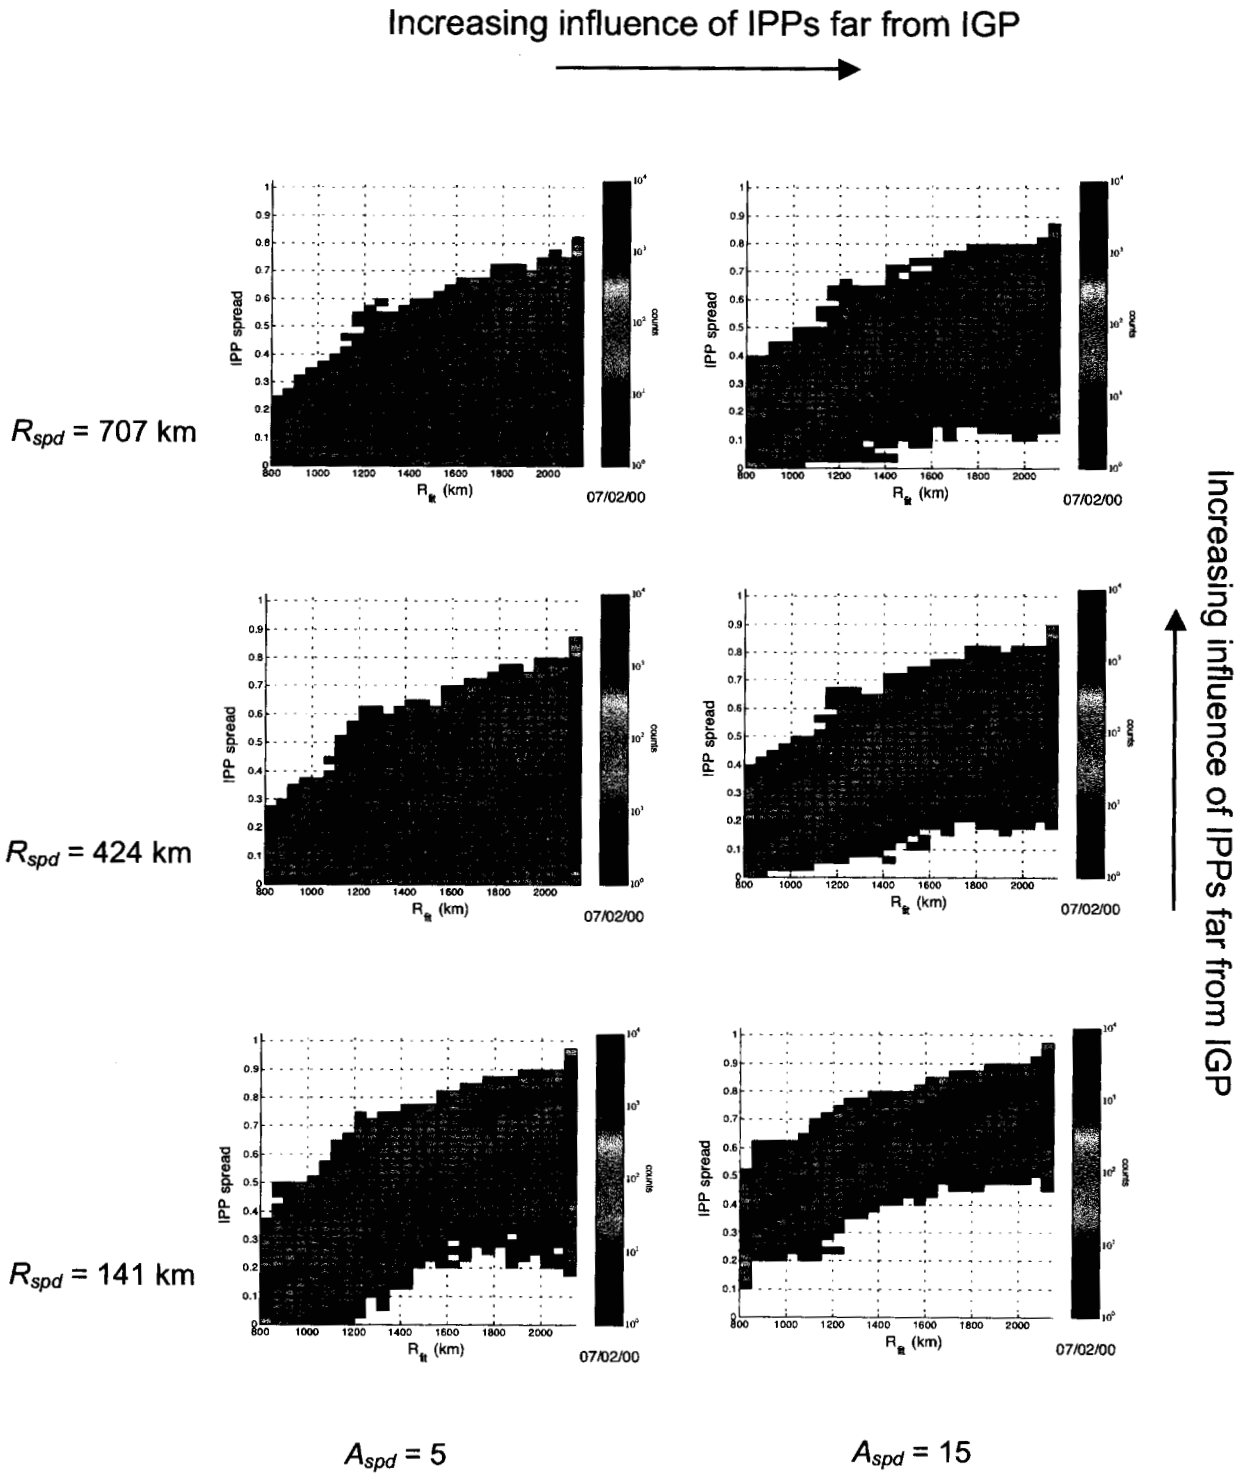

Figure 6. Distributions of IPP spread (M_{spd}) and fit radius for IPP configurations corresponding to slant TEC measurements recorded by WAAS receivers in the course of one day (July 2, 2000), using different values for the spread metric constants A_{spd} and R_{spd}. Arrows indicate changes in the metric constants that lead to an increasing influence of IPPs far from the IGP.

For the same set of IPP configurations (*i.e.*, the same set of fits) used to generate Fig. 1, Fig. 6 shows distributions of metric values corresponding to different values for A_{spd} and R_{spd}. Note that there is a stronger correlation between the value of the fit radius and the spread metric, especially when A_{spd} = 15, than there was when the relative centroid was used as the spread metric (Fig. 1). In other words, as the fit radius increases, the IPP coverage as measured by the new spread metric tends to get worse. In contrast, the relative centroid can remain good (*i.e.*, a metric value near zero) even as the fit radius increases, if the IPPs are spread around the IGP. Arrows are drawn in Fig. 6 to indicate that the influence of IPPs far from the IGP grows as A_{spd} and R_{spd} increase.

For several values of A_{spd} and R_{spd}, we have generated spatial threat models. To compare results, we evaluate the mean overbound that would be accessed by a given set of IPP configurations for a representative day (July 2, 2000). This mean value, $\overline{B}$, is calculated as follows:

$$\overline{B} = \frac{\sum_{i,j} n_{ij} b_{ij}}{\sum_{i,j} n_{ij}}, \tag{6}$$

where n_{ij} is the number of counts in the bin of the ith spread metric and jth fit radius (*e.g.*, see Fig. 1), and b_{ij} is the corresponding overbound (*e.g.*, see Fig. 2b).

The lowest value of $\overline{B}$ achieved to date has been for A_{spd} set to 15 and R_{spd} set to 424 km. Figure 7 displays the distributions of fit residuals in the generation of the spatial threat models, tabulated as a function of (a) relative centroid and (b) the new spread with these values for A_{spd} and R_{spd}. Notice that the IPP configurations that produce large fit residuals are more sharply confined to the region of high metric values in the latter case; that is, the latter spread metric is better at distinguishing IPP configurations that are more likely to produce large residuals.

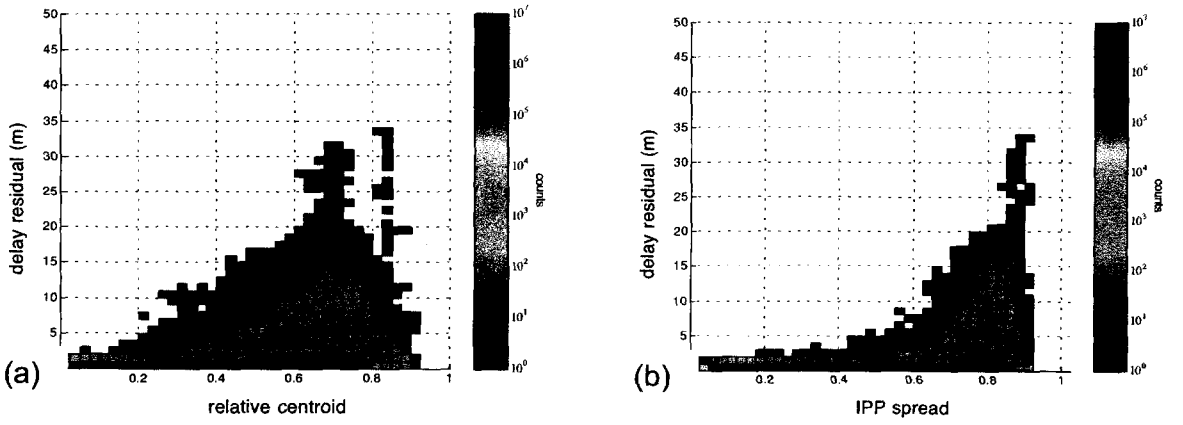

Figure 7. Distribution of fit residuals tabulated for threat models parameterized by (a) relative centroid and (b) a spread metric with A_{spd} = 15 and R_{spd} = 424 m. The data and processing are identical to those used to generate Fig 2.

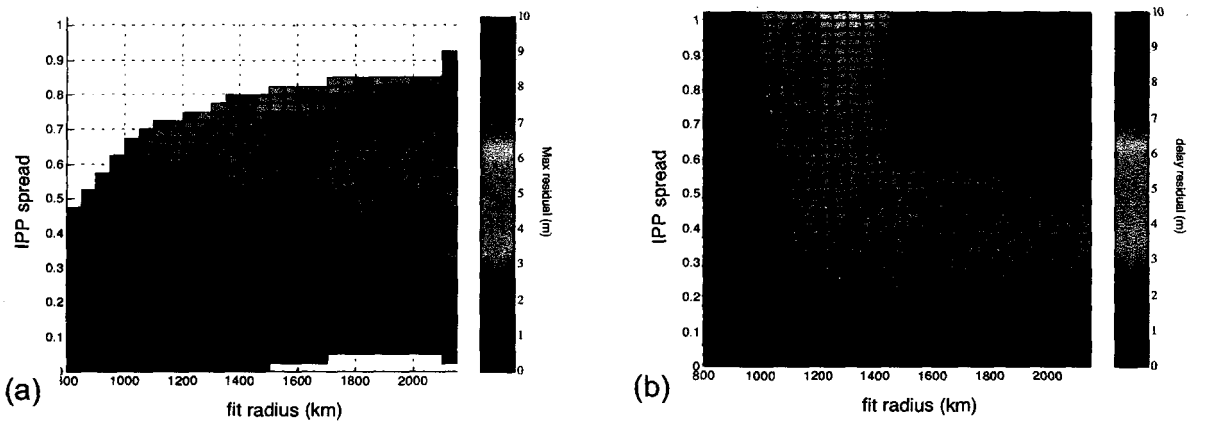

Figure 8. (a) Maximum fit residual and (b) as functions of IPP spread (M_{spd}) and fit radius for threat model parameterized by a spread metric with A_{spd} = 15 and R_{spd} = 424 m. The data and processing are otherwise identical to those used to generate Fig 2.

Figure 8 displays the maximum fit residuals and their overbound, tabulated as a function of the new spread metric and fit radius, when A_{spd} is set to 15 and R_{spd} is set to 424 km. Comparing Fig. 8b to Fig. 2b, we find that overbounds at a given level (*i.e.,* color) have tended to move to higher spread metric values. Comparing Fig. 1 and the distribution for A_{spd} = 15 and R_{spd} = 424 km in Fig. 6, we find distributions that are roughly comparable (except that the former contains some counts in the lower right-hand corner that are missing in the latter distribution, as discussed previously). Thus we should expect that, on average, the new distribution of spread metric values accesses regions of large overbounds less often, and this is indeed what we find. Using the new spread metric reduces the mean overbound accessed from 10.7 ± 10.0 meters to 7.3 ± 6.1 meters (where the uncertainty represents one standard deviation), a decrease of 32%. The contribution to the GIVE that protects against the threat of an undersampled irregularity is proportional to the overbound. Thus, we should expect that it will also be reduced, on average, by a similar amount, depending upon precisely what data sets are used to generate the spatial threat model.

6 Conclusions

We have defined a new IPP spread metric designed to identify readily GPS measurement configurations that might fail to sample a significant ionospheric irregularity in the vicinity of an SBAS IGP. This metric exhibits the following attributes: (1) the metric is sensitive to the angular distribution of the IPPs about the IGP; (2) the contribution of a single IPP to the metric defined at a given IGP decreases with separation distance; (3) the metric improves monotonically as the number of IPPs increases; and, (4) the parameterization of the metric permits control of its sensitivity to the variation of a single IPP location.

Using this metric to parameterize the spatial threat model, we have shown that the magnitude of the error bounds broadcast by WAAS can be reduced. In particular, the contribution to the GIVE that protects against the threat of an undersampled irregularity has been reduced, on average, by 32%. Thus, the new spread metric can aid SBAS performance by helping to define error bounds that are safe but not overly conservative.

7 Acknowledgments

The research described in this paper was performed by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

8 References

Birch, M.J., J.K. Hargreaves, and G.J. Bailey, "On the use of an effective ionospheric height in electron content measurement by GPS reception", *Radio Sci.*, *37*(1), 1015, doi:10.1029/2000RS002601, 2002.

Mannucci, A.J., B.A. Iijima, U.J. Lindqwister, X. Pi, L. Sparks, B.D. Wilson, "GPS and Ionosphere" in *URSI Reviews of Radio Science, 1996-1999*, Oxford University Press, 1999.

RTCA Special Committee 159, *Minimum Operational Performance Standards for Global Positioning System/Wide Area Augmentation System Airborne Equipment, Doc. No. RTCA/DO-229*, RTCA, Inc., Washington, DC, January 1996.

Sparks, L., X. Pi, A.J. Mannucci, T. Walter, J. Blanch, A. Hansen, P. Enge, E. Altshuler, R. Fries, "The WAAS Ionospheric Threat Model," in *Proceedings of the International Beacon Satellite Symposium 2001*, Boston College, Boston, MA, June 3-6, 2001.

Sparks, L., A. Komjathy, and A.J. Mannucci, "The Dependence of WAAS Ionospheric Error Bounds upon the Spatial Distribution of GPS Measurements," in *Proceedings of the ION GPS NTM 2003*, Anaheim, CA, January 22-24, 2003.

T. Walter, A. Hansen, J. Blanch, P. Enge, A.J. Mannucci, X. Pi, L. Sparks, B. Iijima, B. El-Arini, R. Lejeune, M. Hagen, E. Altschuler, R. Fries, A. Chu, "Robust Detection of Ionospheric Irregularities," in *Proceedings of the ION GPS 2000, 13th International Technical Meeting of the Institute of Navigation*, Salt Lake City, UT, September 11-14, 2000.