Protection Level Calculation in the Presence of Heavy Tail Errors Using Measurement Residuals

Juan Blanch, Todd Walter, Per Enge

Stanford University

blanch@stanford.edu

ABSTRACT

In safety-of-life applications of satellite navigation, the Protection Level (PL) equation translates what is known about the pseudorange errors into a reliable limit on the positioning error. The current PL equations for Satellite based augmentation systems are based on Gaussian statistics: all errors are characterized by a zero mean Gaussian distribution which is an upper bound of the true distribution in a certain sense. This approach is very practical: the calculations are simple and the receiver computing load is small. However, when the true distributions are far from Gaussian, such characterization forces an inflation of the protection levels that damages performance. This happens for example with heavy tail distributions or errors for which there is not enough data to evaluate the distribution density up to small quantiles. Also, in the certification process, it is very difficult to agree on a given distribution when the statistics are gathered from a multitude of situations (like elevation angle in the case of multipath). With the development of new optimization methods and the increasing computing power, it is worthwhile exploring new ways of computing integrity error bounds.

In this paper we present a way of computing the optimal protection level when the pseudorange errors are characterized by a mixture of Gaussian modes. First, we will show that this error characterization adds a new flexibility and helps account for heavy tails without losing the benefit of tight core distributions. Then, we will state the positioning problem using a Bayesian approach. Finally, we will apply this method to

protection level calculations for the Wide Area Augmentation System (WAAS) using real data from the National Satellite Test Bed and the WAAS network. The results are very promising: Vertical Protection Levels are almost halved in average without increasing the maximum ratio between the actual vertical error and the VPL.

INTRODUCTION

In the next ten years the number of pseudorange sources for satellite navigation and their quality is expected to increase dramatically: The United States is going to add two new civil frequencies (L5 and L2C) in the modernized GPS, and Europe is planning to launch Galileo which should be fully operative before 2015, also with multiple frequencies. By combining two frequencies, users will be able to remove the ionospheric delay which is currently the largest error, thus reducing nominal error bounds by more than 50%. In particular, safety-of-life applications using augmentation systems will be greatly enhanced. However, it will remain a challenge to provide small hard error bounds - Protection Levels – to meet stringent navigation requirements. Airborne multipath, ephemeris error, loss of signal due to scintillation (in equatorial regions), are still a challenge in the path to provide Cat III GNSS augmentation systems. For example, even with dual frequency, it is not obvious that the Wide Area Augmentation System (WAAS) with GPS alone would meet 100% APV II (20 meters vertical) availability over the United States [1].

There are many ways to improve the performance of an SBAS without changing the message standards [2]: by adding satellites, by adding reference stations, by improving the algorithms at the master station (specially the clock and ephemeris algorithms). It is worthwhile however, now that the new L5 MOPS is being developed, to explore possible modifications to the message content to improve performance. The current methodologies to provide integrity to augmentation system users are based on Gaussian overbounding techniques. For every source of error, the user receives a standard deviation that corresponds to the Gaussian overbound of the error. For this reason, every source of pseudorange error needs to be overbound, in a certain sense, by a gaussian distribution up to very small quantiles - on the order of the probability of hazardously misdetection (10^{-7}) [3]

This is a very difficult task: for example, it is not possible to have experimental stationary distributions for the errors because the conditions and environment are always changing. Also, the errors are all mixed together so it is hard to isolate them. As a result, it is necessary to increase the Gaussian overbound to be sure to cover the tails of the individual error distributions. However, by doing so, we ignore the fact that the core of the distribution is usually much tighter than the overbound (by a large factor), thus giving up performance [4].

In this paper, we present an estimation technique where errors are characterized by Gaussian mixtures. By using a bayesian approach, this technique optimally takes advantage of the tight core of the error distributions while accounting for the heavy tails. Although this technique could be used in several places at the master station level in an SBAS, we will focus on its application to the Protection Level calculation at the receiver. The paper is organized as follows. First, we will explain how the pseudorange error can be characterized by a mixture of Gaussian distributions to account for heavy tails while preserving a tight core. Then, we will compute the a posteriori error density and the resulting error bound on the user position. Finally, we will present a possible application of this technique to Protection Level calculation and show its results on real data collected at the WAAS reference stations.

PSEUDORANGE ERROR MODEL

Typically, pseudorange errors look Gaussian at the core of the distribution. At the tails, however, either we do not have enough data points to have a good representation of the distribution, or the points that we have suggest that the tails are worse than Gaussian [5]. There are several ways to account for heavy tails in the pseudorange errors. In this paper, we will use multimodal Gaussian distribution to characterize pseudorange error distributions with heavy tails. Let z be the random variable representing the pseudorange error. The density of z can be written:

$$p(z) = \sum_{i=1}^{q} a_i f_{m_i,\sigma_i}(z)$$

In this equation $f_{m_i,\sigma_i}(z)$ is the density of a Gaussian with mean m_i and standard deviation σ_i. The only requirements on the coefficients a_i are that their sum be one and that the density be positive for all z. Although the equations will be written for more than two terms modes, we will mostly work here with bimodal mixtures.

In the previous paragraph we have written the density for a single pseudorange. Here we show how we can derive the model for a set of independent pseudorange errors. From now on the random variable z is a vector. Let us consider n pseudorange sources and label z_k the error on each of them (now z is a vector). Each error is characterized by a Gaussian mixture:

$$p(z_k) = \sum_{i=1}^{q_k} a_{i,k} f_{m_{i,k},\sigma_{i,k}}(z_k)$$

The joint density is given by:

$$p(z_1,\ldots,z_n) = \prod_{k=1}^{n} p_k(z_k) = \prod_{k=1}^{n}\left(\sum_{i=1}^{q_k} a_{i,k} f_{m_{i,k},\sigma_{i,k}}(z_k)\right)$$

If one develops this expression, we see that the joint distribution is a mixture of multivariate Gaussians. The covariance matrices are given by each possible combination of the modes in each pseudorange error. Let us label C_j the covariance for a given mode and p_j the probability of that mode. The density of the random variable z is given by:

$$p(z) = \sum_{j=1}^{N_{\mathrm{mod}\,es}} p_i f_{M_i,C_i}(z)$$

This is equivalent to saying that the covariance of z is C_i with probability p_i.

ERROR DISTRIBUTION CALCULATION

Now that we have a characterization of the error, we can derive an estimator adapted to it. We will start by computing the probability density of the position x given the measurements y:

$$p(x \mid y)$$

It is rare to consider this expression because, when the errors are Gaussian and a least squares estimator with the proper covariance is used, the error density is given by a multivariate Gaussian centered a the estimated position, whose covariance does not depend on the measurements (one can compute the covariance of the position estimate without knowing the actual measurements). But now the situation is different, as the errors are no longer Gaussian.

It is assumed here that the linear model for GPS measurement holds (G is the geometry matrix):

$$y = Gx + z$$

We start by writing Bayes formula:

$$p(x \mid y) = \frac{p(x, y)}{p(y)}$$

Now let us develop $p(x, y)$:

$$p(x, y) = p(y \mid x)\, p(x)$$

Here *p(x)* designates the a priori distribution of the position. The expression for $p(y \mid x)$ is easy to compute, because *z* is a mixture of Gaussian distributions:

$$p(y \mid x) = \sum_{j=1}^{N_{\mathrm{mod}}} p_j \left|W_j\right|^{\frac{1}{2}} e^{-\frac{1}{2}(y-Gx)^T W_j (y-Gx)}$$

where W_j is the inverse of the covariance C_j matrix corresponding to the j[th] mode:

$$W_j = C_j^{-1}$$

To compute $p(y)$, we integrate over all possible positions:

$$\begin{aligned}
p(y) &= \int_x p(y \mid x)\, p(x)\, dx \\
&= \int_x \sum_{j=1}^{N_{\mathrm{mod}}} p_j \left|W_j\right|^{\frac{1}{2}} e^{-\frac{1}{2}(y-Gx)^T W_j (y-Gx)} p(x)\, dx \\
&= \sum_{j=1}^{N_{\mathrm{mod}}} p_j \left|W_j\right|^{\frac{1}{2}} \int_x e^{-\frac{1}{2}(y-Gx)^T W_j (y-Gx)} p(x)\, dx \\
&= \sum_{j=1}^{N_{\mathrm{mod}}} p_j \left|W_j\right|^{\frac{1}{2}} e^{-\frac{1}{2} y^T \left(W_j - W_j G \left(G^T W_j G\right)^{-1} G^T W_j\right) y} \int_x e^{-\frac{1}{2}\left(x - \left(G^T W_j G\right)^{-1} G^T W_j y\right)^T \left(G^T W_j G\right)\left(x - \left(G^T W_j G\right)^{-1} G^T W_j y\right)} p(x)\, dx
\end{aligned}$$

Let *p(x)* tend to a uniform distribution over the whole space in both the numerator and the denominator (it is possible to include an a priori in the position of x, but to be consistent with the assumptions of current methods we make the a priori tend to a uniform distribution over the whole space). There is an analytic expression for the integral term:

$$\int_x e^{-\frac{1}{2}\left(x - \left(G^T W_j G\right)^{-1} G^T W_j y\right)^T \left(G^T W_j G\right)\left(x - \left(G^T W_j G\right)^{-1} G^T W_j y\right)} dx = \sqrt{2\pi} \left|G^T W_j G\right|^{-\frac{1}{2}}$$

The denominator is then:

$$p(y)=\sqrt{2\pi}\sum_{j=1}^{N_{\mathrm{mod}}} p_j \left|W_j\right|^{\frac{1}{2}} \left|G^T W_j G\right|^{-\frac{1}{2}} e^{-\frac{1}{2}\chi_j^2}$$

where we have:

$$\chi_j^2 = y^T\left(W_j - W_j G\left(G^T W_j G\right)^{-1} G^T W_j\right) y$$

Notice that this expression would be chi-square distributed if the measurements followed the j$^{\text{th}}$ mode. The numerator can be written (where *p(x)* is canceled out):

$$\begin{aligned}
p(y\mid x) &= \sum_{j=1}^{N_{\mathrm{mod}}} p_j \left|W_j\right|^{\frac{1}{2}} e^{-\frac{1}{2}(y-Gx)^T W_j (y-Gx)} \\
&= \sum_{j=1}^{N_{\mathrm{mod}}} p_j \left|W_j\right|^{\frac{1}{2}} e^{-\frac{1}{2}y^T\left(W_j - W_j G\left(G^T W_j G\right)^{-1} G^T W_j\right) y} e^{-\frac{1}{2}\left(x-\left(G^T W_j G\right)^{-1} G^T W_j y\right)^T \left(G^T W_j G\right)\left(x-\left(G^T W_j G\right)^{-1} G^T W_j y\right)} \\
&= \sqrt{2\pi}\sum_{j=1}^{N_{\mathrm{mod}}} p_j \left|W_j\right|^{\frac{1}{2}} \left|G^T W_j G\right|^{-\frac{1}{2}} e^{-\frac{1}{2}\chi_j^2} \left\{\frac{1}{\sqrt{2\pi}}\left|G^T W_j G\right|^{\frac{1}{2}} e^{-\frac{1}{2}\left(x-\hat{x}^{(j)}\right)^T \left(G^T W_j G\right)\left(x-\hat{x}^{(j)}\right)}\right\}
\end{aligned}$$

where:

$$\hat{x}^{(j)} = \left(G^T W_j G\right)^{-1} G^T W_j y$$

is the position estimate using a least squares algorithm assuming that the measurements have the covariance W_j^{-1}. We notice now that the term :

$$\left\{\frac{1}{\sqrt{2\pi}}\left|G^T W_j G\right|^{\frac{1}{2}} e^{-\frac{1}{2}\left(x-\hat{x}^{(j)}\right)^T \left(G^T W_j G\right)\left(x-\hat{x}^{(j)}\right)}\right\}$$

is the density of a multivariate Gaussian centered on $\hat{x}^{(j)}$ and with covariance $\left(G^T W_j G\right)^{-1}$ which we will note:

$$p_{\hat{x}^{(j)},\left(G^T W_j G\right)^{-1}}\left(x\right)$$

With these notations, the density of the a posteriori distribution of *x* is given by:

$$p\left(x \mid y\right) = \sum_{j=1}^{N_{\text{mod}}} c_j p_{\hat{x}^{(j)},\left(G^T W_j G\right)^{-1}}\left(x\right)$$

where the coefficient a_j is defined by:

$$c_j = \frac{p_j \left|W_j\right|^{\frac{1}{2}} \left|G^T W_j G\right|^{-\frac{1}{2}} e^{-\frac{1}{2}\chi_j^2}}{\sum_{i=1}^{N_{\text{mod}}} p_i \left|W_i\right|^{\frac{1}{2}} \left|G^T W_i G\right|^{-\frac{1}{2}} e^{-\frac{1}{2}\chi_i^2}}$$

This expression gives the a posteriori distribution of the position location given the measurements. This density appears as a linear combination of Gaussian densities associated with the optimal least square estimate corresponding to each mode. One can see that the expression depends heavily on the measurements themselves through the chi-square statistic for each of the covariances in the mixture, and through each corresponding estimate.

ERROR BOUND CALCULATION

In the previous section we determined how to compute the density of the position error. Here we explain how to translate the density in an error bound for a given probability *ε*. Let us suppose that we want to compute an error bound in the vertical domain. The problem is to find a position estimate $\hat{x}_1$ and an error bound (the Vertical Protection Level (VPL)) such that:

$$\text{Prob}\left(\left|x_1-\hat{x}_1\right|>VPL\mid y\right)<\varepsilon$$

From the density of the position, it is easy to derive the density for each coordinate:

$$p\left(x_1\mid y\right)=\sum_{j=1}^{N_{\text{mod}}}a_j\,p_{\hat{x}_1^{(j)},\left(G^TW_jG\right)_{1,1}^{-1}}\left(x_1\right)$$

To determine $\hat{x}_1$ and VPL, we need to find an interval I such that:

$$\int_{x_1\in I}p\left(x_1\mid y\right)\geq 1-\varepsilon$$

This interval is not unique and can be adapted to different requirements. In this work we chose to determine it by setting the probability of being on each side of the interval to ε/2. This can be easily implemented using a slicing algorithm (it takes very few iterations) to determine independently each bound. Once we have the upper and the lower bound of the interval it is easy to compute $\hat{x}_1$ and VPL.

PROTECTION LEVEL EQUATION AND MESSAGE CONTENT

Although there is an analytical expression for the error density position, strictly, there is no PL equation. Instead, there is an algorithm that determines the error bound based on the error density. As an example, we suggest here a small modification to the SBAS message content that would allow SBAS users to take advantage of this technique. The current SBAS message (defined in the Minimum Operational Performance Standards for SBAS (MOPS)) allows users to compute at each time the standard deviation of each pseudorange error [6] σ_i^2. The user treats the pseudorange error as if it was a Gaussian random variable with zero mean and standard deviation σ_i^2. We can simply account for the fact that the errors have a tight core and heavy tails by splitting this single mode in two modes: one describing the core and another one describing the tails. The random error with density $p_{0,\sigma_i^2}(z_i)$ is replaced by:

$$a_{core} p_{0,\gamma_{core}\sigma_i^2}(z_i) + (1 - a_{core}) p_{0,\gamma_{tails}\sigma_i^2}(z_i)$$

To perform this split the user needs three additional scalars:

- the probability of being in the core of the distribution a_{core},
- the ratio between the standard deviation of the core and σ_i^2 (smaller than one) γ_{core}
- the ratio between the standard deviation of the tails and σ_i^2 (smaller than one) γ_{tails}

One could choose to make this parameters depend on the satellite, but it would be also possible to set a conservative set of parameters valid for all satellites, and this is the approach that will be taken here. Such an addition to the MOPS would be backwards compatible: a user without the capability to apply the new algorithm would simply use the current PL equations.

EXPERIMENTAL AND SIMULATION RESULTS

In this section, two types of results are presented. The first set of results is intended to test the integrity of the error bound calculation as well as its general behavior. The second set of results is a simulation intended to evaluate the effect on availability of this new technique.

The algorithm was first tested with simulated data by generating errors that were distributed as independent Gaussian mixtures matching the theoretical assumptions. However, it is the performance under real error distributions that do not match exactly the estimator assumptions that is critical. For that purpose, we used ionospheric delay measurements collected at the WAAS reference stations on October 31st, 2003 – one of the worst ionospheric storms observed by the WAAS reference station network. On that day, WAAS service was limited, because the ionospheric disturbance detector had triggered in many sectors. Here, each station was successively treated as a user. For a given station the ionospheric delays were corrected using an algorithm that closely mimics the actual WAAS algorithm. The main differences here were that only the ionospheric delay was considered and that the standard deviations of the errors were not quantized. Figure 1 shows the histogram of the 167599 normalized pseudorange errors residuals:

$$\text{normalized residuals} = \frac{\text{pseudorange error}}{\text{sigma}}$$

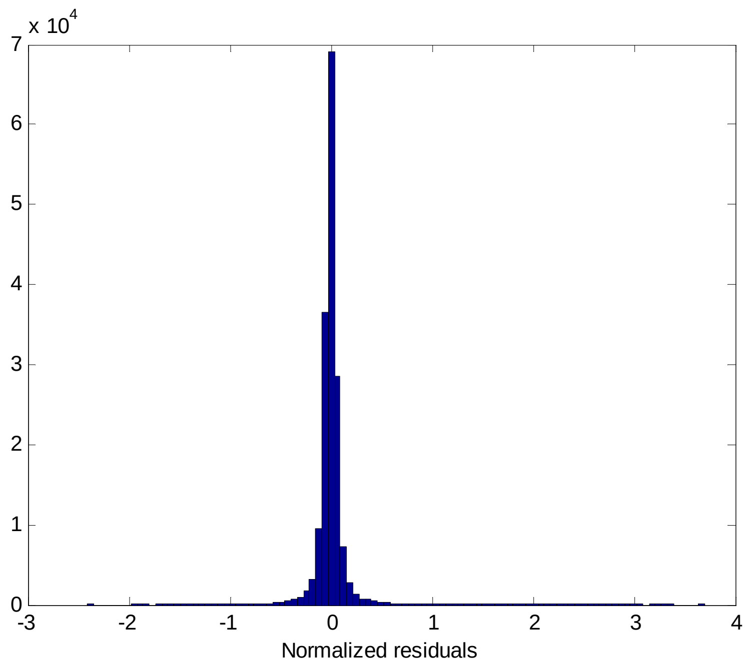

***Figure 1**. Histogram of the normalized pseudorange errors*

The above histogram is an extreme example of the kind of error distributions that are seen in satellite navigation: the core of the distribution is Gaussian and the rest is worse than Gaussian. To see that, it is practical to plot the quantile-quantile plot (qq-plot) of the histogram: for several percentiles we plot the quantile of the empirical distribution and as a function of the quantile of a normal distribution. Figure 2 shows the qq-plot of the data shown in Figure 1.

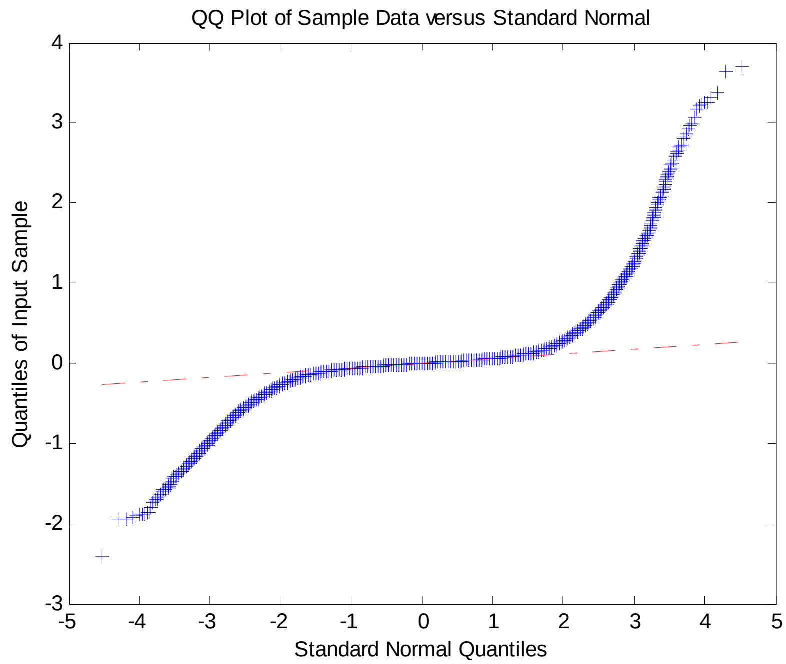

Figure 2. *qq-plot of the normalized pseudorange errors*

If a qq-plot is a straight line, it means that the data is well characterized by a Gaussian. In Figure 2, we see that the Gaussian assumption breaks down at the 10^{-2} quantile. A rough estimate of the Gaussian overbound is given by the slope of a straight line going through the origin and a point in the curve. One can see that there is a large ratio between the standard deviation of the core and the Gaussian overbound.

We compared two options to compute the Vertical Protection Level:

- Gaussian overbound: the VPL is computed using the current SBAS VPL equation using a model for broadcast sigmas (which leads to the previous histogram)

- New algorithm: as said before, the pseudorange error distribution is derived from the previous one by splitting the gaussian mode in a mixture of two gaussian modes (these parameters were set arbitrarily based on the results of Figure 2):

$$a_{core} = 97.5\% \qquad \gamma_{core} = .3$$
$$a_{tails} = 2.5\% \qquad \gamma_{tails} = 1.5$$

Figures 3 and 4 show the vertical position errors divided by the Vertical Protection Levels for the Gaussian overbound method and the new algorithm (respectively). The maximum ratios were respectively .62 and .6.

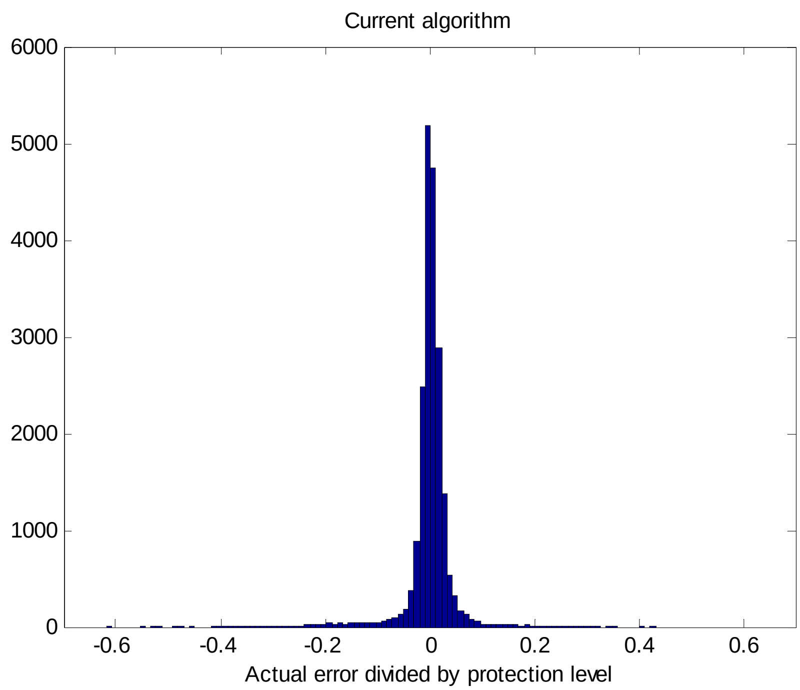

Figure 3. *Vertical position error normalized by the Protection Level computed using the Gaussian overbound method*

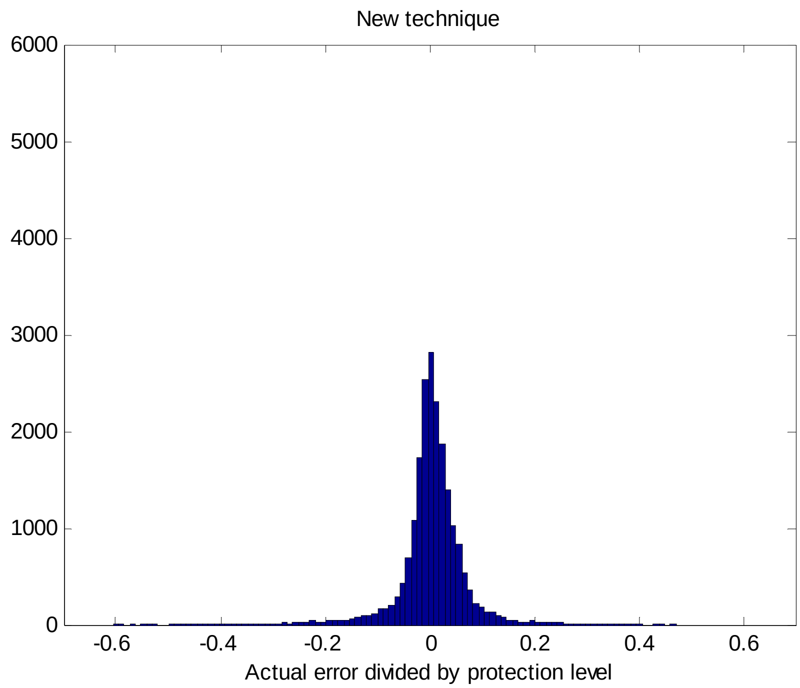

Figure 4. *Vertical position error normalized by the Protection Level computed using the new method*

One can see that, while the maximum ratios between actual error and Protection Level are similar, the residuals for the new algorithm are larger, suggesting that we are being less conservative while preserving integrity. This is made more obvious in the triangle charts shown in Figures 5 and 6: the correlation between actual error and PL is larger with the new algorithm. We see that in average, the Protection Level is divided by two, and that integrity seems to be preserved.

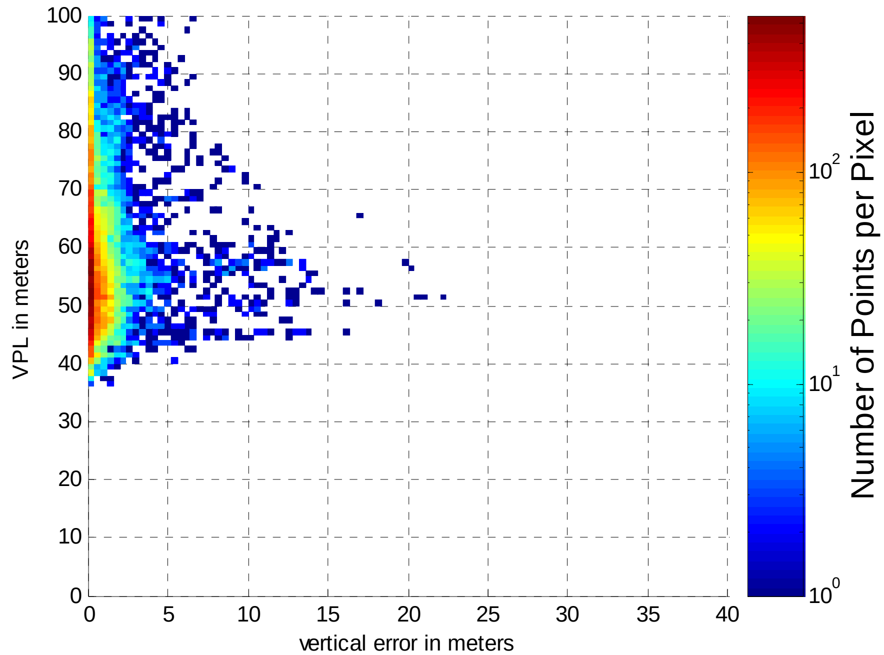

Figure 5. *Triangle chart for the Gaussian overbound method*

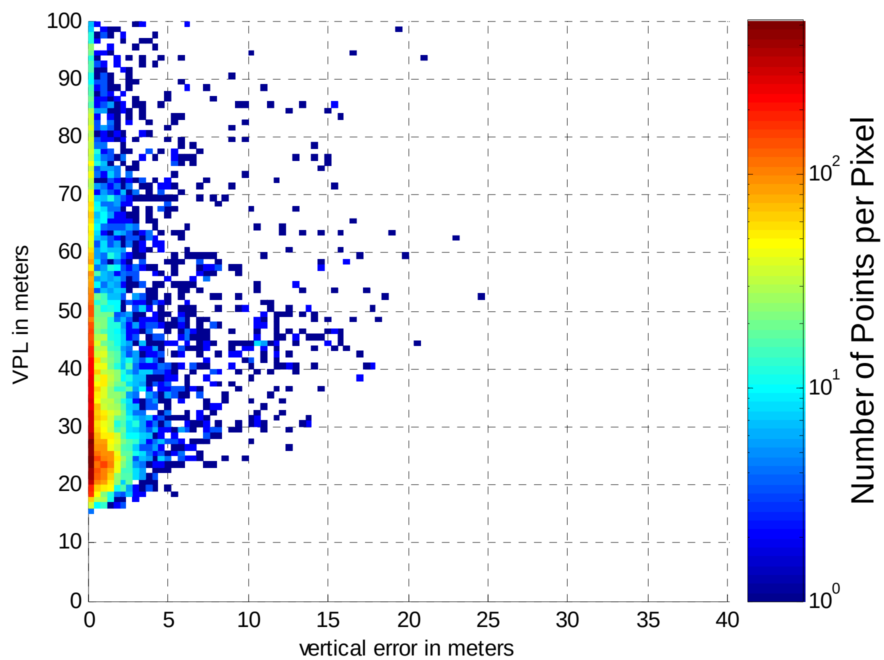

Figure 6. *Triangle chart for the new algorithm*

To evaluate the performance of the new algorithm in terms of VPL, we used the MATLAB based service volume analysis tool MAAST. This software evaluates the performance of WAAS by computing the protection level for users placed on a regular grid over a given period of time [7]. For these simulations we have assumed:

- a constellation of 24 optimal GPS satellites (which is the constellation specified by the MOPS) and 2 Geostationary satellites (POR, AOR-W)
- that all satellites are dual frequency L1-L5 so there is no ionospheric error other than the uncertainty on the ionospheric delay estimate;
- the current network of 25 WAAS reference stations.

The purpose of these simulations was to compare the performance of a dual frequency WAAS using the current overbounding approach to on using the algorithm proposed in this work. Figure 7 shows the 99% VPL percentile over the course of a day using the current approach.

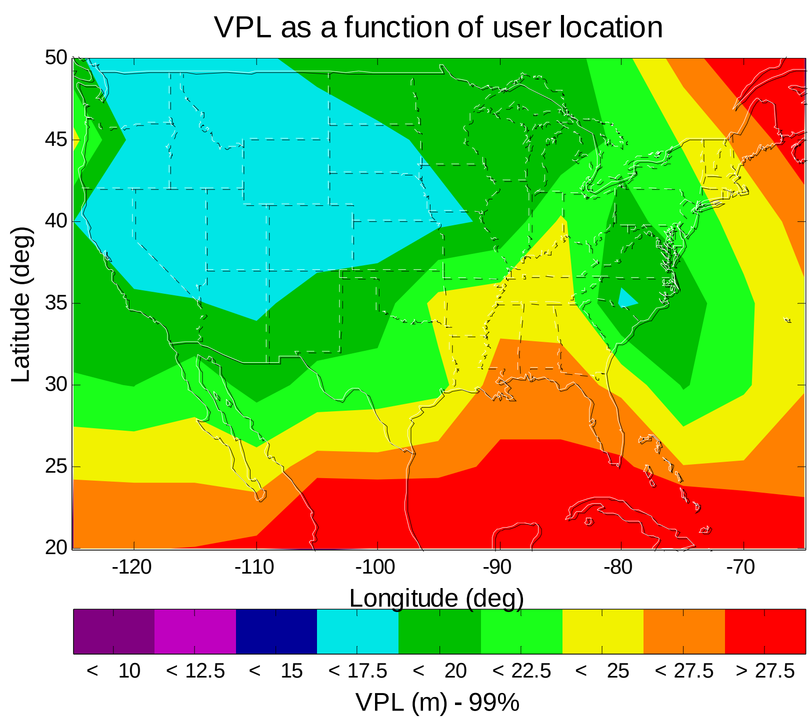

Figure 7. *Dual Frequency WAAS 99% VPL percentile with current overbounding approach*

To produce the equivalent plot for the new technique it was necessary to simulate measurement errors, so a solution could be computed. Since we wanted to evaluate the performance under nominal conditions, the simulated errors were taken to be random Gaussian variables with a standard deviation being a fraction (.3) of the WAAS user range error. Figure 8 shows the 99% VPL percentile over the course of a day using the new technique. One can see that VPL's are reduced between 40% and 50%, which is a very promising figure.

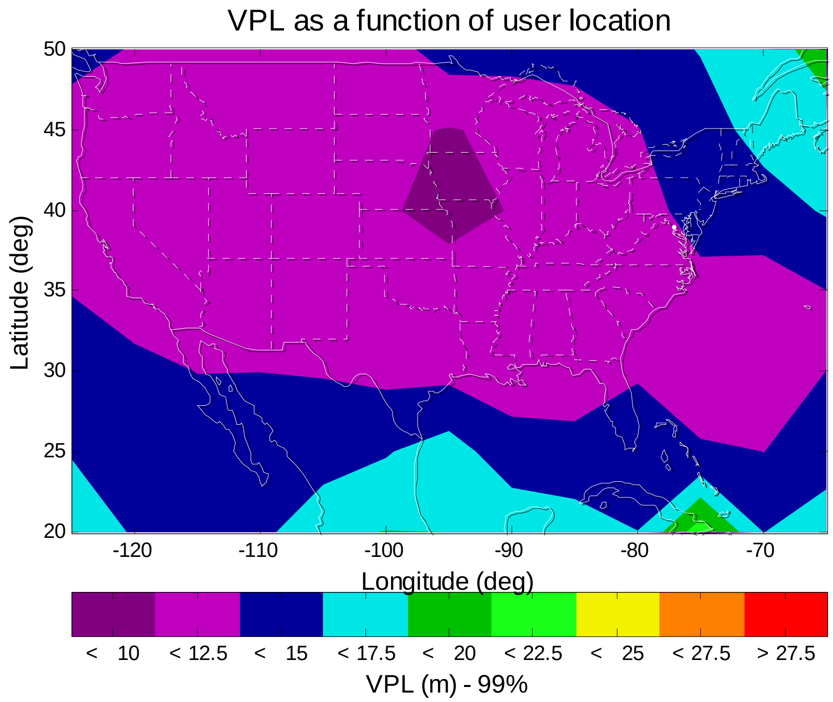

Figure 8. *Dual Frequency WAAS 99% VPL percentile with the new algorithm*

COMPUTATIONAL AND INTEGRITY ISSUES

As it has been presented, this algorithm has some theoretical weaknesses:

- the theory assumes that errors are independent (although correlation could be introduced, if we knew what the appropriate correlation is)
- biases are not accounted (however, very large biases would be detected)
- using a very spread distribution for the failed mode could cause an underestimation of the error (this risk is small, but that needs to be proved)

The two first points are already present in the current overbounding approach, but because of the large margin, they are not a major concern. The last point is specific to the new algorithm and needs to be well understood and accounted for. By looking at the expression for the a posteriori density, one can see that if, for a given set of measurements, we increase the standard deviation of the outer mode the PL will increase, reach a maximum, and decrease again. As a consequence, using a larger sigma for the outer mode is not always a conservative approach.

The computational load comes mainly from the large number of matrix inversions. The number of matrix inversions needed for one position fix is not one, like in the current PL equation, but several hundreds or thousands. This is not an issue for a PC: a position fix and PL computation with 10 satellites took .1 s. This was achieved:

- by exploiting the fact that one can go from one inverted matrix to another by rank one updates, which are far less demanding than full matrix inversions,
- by excluding terms in the density that are known to be small before computing them

However, even with these improvements, it will be challenging to apply this algorithm in a certified airborne receiver.

CONCLUSION

We have developed a formula for the a posteriori position distribution when the pseudorange errors are distributed according to a mixture of gaussian modes. This means that we can theoretically compute the position distribution for any kind of pseudorange error distribution well characterized by a finite mixture of gaussian modes. The most remarkable feature of the formula is its dependence on the actual measurements.

Before this method can be applied in an actual system it will be necessary to better understand its behavior, as the dependence of the error bounds on the measurements is more complex than in current methods. Also, this method has a larger computational load. Future work will need to address these issues. While systematically evaluating the new approach using WAAS NSTB data we will study the possible integrity issues and behavior under different error models. Also, we will try to simplify the formula as much as possible in order to reduce its complexity and computational load.

The method was applied to compute positions and protection levels in the presence of ionospheric delay errors (only as an example of heavy tail distribution). The error bounds appeared to be 50% smaller on average than with the current methods, without affecting the maximum ratio between actual error and protection level. Also, a simulation using a service volume analysis tool showed that the VPLs in a dual frequency WAAS could be reduced from 40% to 50% during nominal conditions. It is therefore worthwhile studying its application in safety of life positioning systems requiring small error bounds.

ACKNOWLEDGEMENTS

This work was sponsored by the FAA GPS product team (AND-730).

REFERENCES

[1] T.Walter, P. Enge, P. Reddan. "Modernizing WAAS," *Proceedings of the Institute of Navigation GNSS 2004*, Long Beach, CA, 2004.

[2] RTCA (2001) Minimum Operational Performance Standards for Global Positioning System/Wide Area Augmentation System airborne equipment, RTCA publication DO-229C.

[3] B. DeCleene. "Defining Pseudorange Integrity – Overbounding," *Proceedings of the Institute of Navigation GPS-00*. Salt Lake City, UT, 2000.

[4] J. Rife, S. Pullen, B. Pervan, P. Enge. "Core overbounding and its implications for LAAS integrity," *Proceedings of the Institute of Navigation GNSS 2004*, Long Beach, CA, 2004.

[5] WAAS PAN report: http://www.nstb.tc.faa.gov/

[6] T. Walter, P. Enge, A. J. Hansen. "A Proposed Integrity Equation for WAAS MOPS," *Proceedings of the Institute of Navigation GPS-97*, pp. 475-84, September 1997.

[7] S. S. Jan, W. Chan, T. Walter, P. Enge. "Matlab Simulation Toolset for SBAS Availability Analysis," *Proceedings of the Institute of Navigation GPS-01*, Salt Lake City, UT, 2001.