Models for the effect of integrity monitors on GNSS user error range distributions

Juan Blanch, Yu-Fang Lai, Todd Walter, *Stanford University*

Pedro Silva, *Hexagon*

ABSTRACT

We develop and evaluate a novel method to characterize the user error range error distribution after it has been monitored. This approach enables a simple characterization of the post-monitor distribution that accounts for the prior probability of a fault, the nominal errors of the prior distribution and the measurement error characteristics. This approach could enable the reduction of error bounds in current systems providing GNSS integrity like SBAS.

INTRODUCTION

GNSS-based safety-of-life systems like Satellite-based Augmentation Systems (SBAS) are designed to provide very reliable position error bounds. To achieve the required level of integrity, it is necessary to, first, characterize the errors that could impact the user position, and second, to determine the effect of combined errors at the user level. In most of these systems, the range errors affecting the user have been checked by a set of monitors (either on the ground or on board the satellite).

In this paper, we develop and evaluate a novel method to characterize the user error range error distribution after it has been monitored. This approach enables a simple characterization of the post-monitor distribution that accounts for the prior probability of a fault, the nominal errors of the prior distribution and the measurement error characteristics.

In many systems, among them the Wide Area Augmentation System and EGNOS, the generation of the corrections and the integrity monitoring of the corrections are two distinct processes. Typically, the focus of the correction generation is to maximize accuracy and availability. The algorithms used in the generation of the corrections include many, potentially complex, features intended to improve accuracy under nominal error conditions. The complexity and volume of the code in the generation of the corrections makes it very challenging to build integrity within the generation of the corrections. For this reason, the integrity is checked by a separate process that uses simpler algorithms, and therefore easier to analyze and certify. The purpose of this paper is to provide a simple method to integrate the prior distribution of the errors in the final error bound.

The prior probability distribution of the range error affecting the user has heavy tails, and this is precisely why a monitor is needed to guarantee integrity. This distribution can be determined offline and described conservatively by a nominal distribution, where the errors will usually be well characterized by a narrow zero mean Gaussian distribution, and a priori probability of fault, which is the probability that the error does not follow the nominal conditions. The prior probability of fault can be either bounded based on the

amount of collected data or read directly from it (still from the offline process). The integrity monitor can be modeled as a noisy real time measurement of the user range error that is then compared to a threshold.

The method developed in this paper is based on a worst-case characterization of the distribution of the error conditioned on the monitor not having tripped. Once the post-monitor distribution has been determined, a Gaussian overbound of the resulting probability distribution is computed. It is this Gaussian overbound that is sent to the user (when the monitor has not tripped).

After deriving the formulas to determine the user error bounds, we will first evaluate the resulting Gaussian overbound for a range of prior probabilities of fault, and second, we will briefly evaluate its effect on a dual frequency SBAS.

OVERVIEW FOR THE EFFECT OF MONITOR ON DOWNSTREAM ERROR DISTRIBUTION

The purpose of this analysis is to derive an overbound of the probability distribution after it has been monitored by a ground monitoring system. Let us assume that the errors follow a nominal distribution with high probability, but that there is a residual probability that the error does not follow that distribution and can be unbounded. Without monitors on the ground, it would be up to the user to mitigate the possibility of a large error (that is an error that is not within what is expected from the nominal distribution). The effect of a threshold monitor will be to clip the tails of the distribution. The amount of clipping will be dependent on the measurement noise in the ground monitors and the characteristics of the prior probability. In this analysis we provide a model for this effect that can then be used to compute the error distribution as experienced by the user.

LIST OF DEFINITIONS

As a reference, here we include a list of variables and parameters used in this analysis.

y monitor measurement

ε error under consideration (usually a component of the user range error)

ε_t measurement noise

p PDF of prior distribution of ε

p_0 PDF of prior distribution under H_0

p_1 PDF of prior distribution under H_1 (impulse at b)

σ_{core} standard deviation of p_0

σ_t standard deviation of measurement noise

α probability of H_0

T_M monitor detection threshold

PRIOR DISTRIBUTION

The prior distribution can be described as the linear combination of two distributions, one for the core and one for the tails. We will assume that the core can be bounded by a Gaussian distribution. For the tails, we can assume that it could be any distribution, for example a delta at a given bias value b. This is expressed by writing

$$p(x) = (1 - \alpha)p_0(x) + \alpha p_1(x) \tag{1}$$

With

$p_0(x)$~N$(0, \sigma_{core})$ (hypothesis H_0)

$p_1(x)$~δ_b (hypothesis H_1)

Figure 1 shows the tail cdf for an example were $\sigma_{core} = 1$ and $\alpha = 10^{-5}$. In this plot, the bias *b* is outside the range of the x axis. If the bias was in the range, we would see instead a vertical line dropping down to the tail cdf of the core distribution.

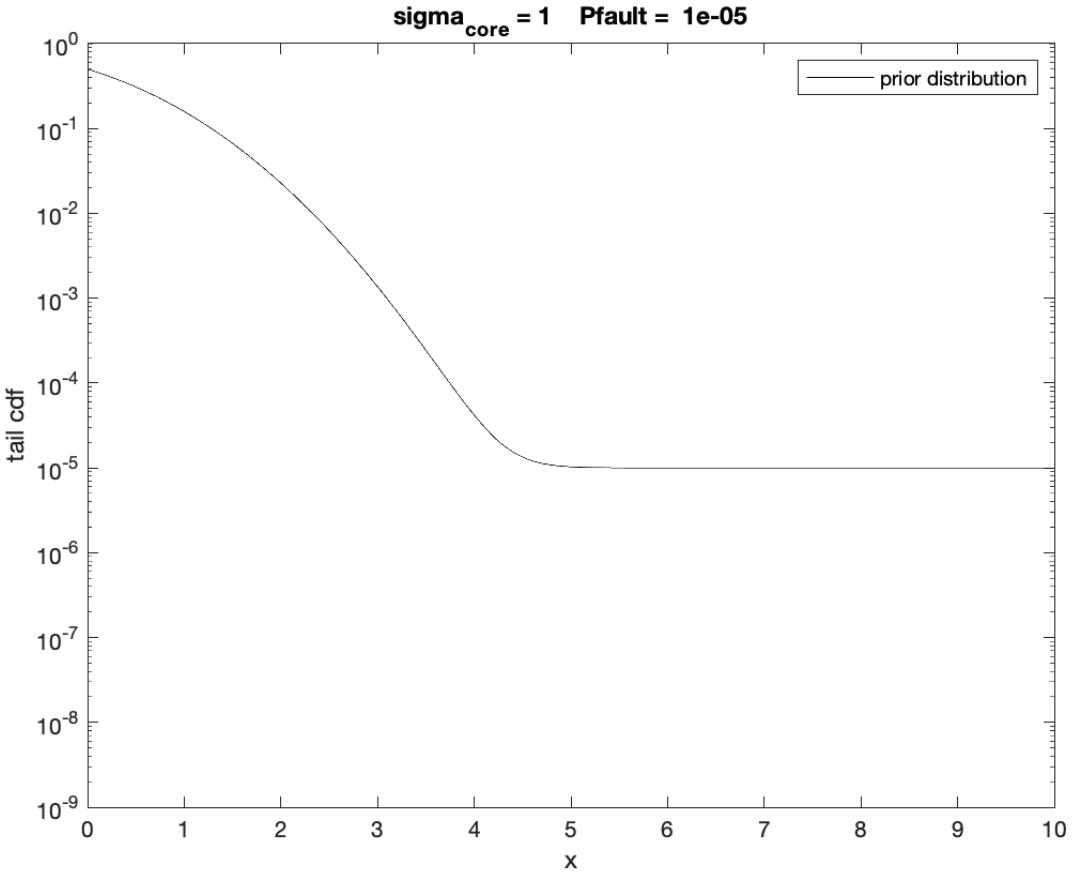

Figure 1. Example prior distribution. The horizontal line starting at around 5 is caused by the bias located outside the shown range.

OFFLINE DETERMINATION OF PARAMETERS

The parameters for this decomposition are determined through an offline process (an example for GPS clock and ephemeris errors is described in the references (Walter 2018, Walter 2019, Liu 2022)**.** In the case where an empirical distribution is representative of the errors, one possible approach consists in

1. Choosing a threshold T. This threshold separates what is considered nominal and what is faulted.
2. Compute the probability of fault as the fraction of the data that is outside the interval [-T T].
3. Gaussian bound the data that is within T (re-normalized so that total probability sums to one)

Note that even if there are no outliers, it is often the case that the amount of data collected is not sufficient to exclude the likelihood of outliers. In that case, one must estimate the probability of fault (that is, of large outliers well outside the core distribution), as done for example in Walter 2019.

EFFECT OF MONITORING

The distribution of the errors will be modified by the real time monitor. Roughly speaking, the monitor will remove the large errors. For example, with a monitor without measurement noise, we would naturally expect the distribution to be clipped exactly at the monitor threshold value. However, because the monitor itself has measurement noise, the effect will be more complex. The measurement at the monitor can be written

$$y = \varepsilon + \varepsilon_t$$

where the error to be bounded is $\varepsilon \sim p$ (see Equation (1)) and the measurement error is $\varepsilon_t \sim N(0, \sigma_t)$.

MONITOR TEST STATISTIC

The monitor compares the measurement to a threshold T_M. We note that this threshold does not need to be related to the threshold *T* that was used in the offline process. If $|y| \leq T_M$, the monitor does not trip, and the user can use the range error that includes the error ε. If not, a flag is sent, and the measurement is considered unhealthy.

The choice of threshold will be dependent on the tradeoff between the probability of alert and the magnitude of the error (the tightness of the downstream distribution). It is practical to express the threshold as a function of the standard deviation of the monitor measurement under the assumption that the error follows the core distribution

$$T_M = K_{FA}\sigma_y$$

where $\sigma_y^2 = \sigma_{core}^2 + \sigma_t^2$.

DISTRIBUTION TO BE CHARACTERIZED: DOWNSTREAM MONITOR DISTRIBUTION

We want to compute or bound the distribution of the error given that the monitor test statistic has not tripped (given by $p(\varepsilon||y| \leq T_M)$). To this purpose, we will compute an upper bound of the tail CDF given by

$$\mathrm{P}(\varepsilon||y| \leq T_M)$$

Derivation of posterior distribution

We start by writing the Bayes rule that relates the probability of two concurrent events with the probability of one given the other

$$\mathrm{P}(\varepsilon > x||y| \leq T_M) = \frac{\mathrm{P}(\varepsilon > x, |y| \leq T_M)}{\mathrm{P}(|y| \leq T_M)}$$

Denominator

Using the formula of total probability, the denominator can be written

$$\mathrm{P}(|y| \leq T_M) = \mathrm{P}(|y| \leq T_M|H_0) \cdot (1-\alpha) + \mathrm{P}(|y| \leq T_M|H_1) \cdot \alpha$$

For the first term, we have

$$\mathrm{P}(|y| \leq T_M|H_0) = 1 - 2Q\left(\frac{T_M}{\sigma_y}\right)$$

where $\sigma_y^2 = \sigma_{core}^2 + \sigma_t^2$ and Q is the tail CDF of a normal distribution. For the second term, it is conservative to assume that

$$\mathrm{P}(|y| \leq T_M|H_1) \approx 0$$

because this only inflates the probability (although not by much).

Numerator

For the numerator, we proceed the same way by writing

$$\mathrm{P}(\varepsilon > x, |y| \leq T_M) = \mathrm{P}(\varepsilon > x, |y| \leq T_M|H_0) \cdot (1-\alpha) + \mathrm{P}(\varepsilon > x, |y| \leq T_M|H_1) \cdot \alpha$$

For the first term, we have the inequality

$$\mathrm{P}(\varepsilon > x, |y| \le T_M | H_0) \approx \mathrm{P}(\varepsilon > x | H_0) = Q\left(\frac{x}{\sigma_{core}}\right)$$

It may be possible to refine this and exploit the detection threshold for the core distribution, but we will not be doing it here. The second term is given by

$$\mathrm{P}\big(\varepsilon > x, |y| \le T_M | H_1(b)\big) = H(b-x)Q\left(\frac{b-T_M}{\sigma_t}\right) \quad (2)$$

where *H* is the step function.

This is because, conditioned on $H_1(b)$, the two probabilities are independent and

$$\mathrm{P}\big(\varepsilon > x | H_1(b)\big) = H(b-x)$$

$$\mathrm{P}\big(|y| \le T_M | H_1(b)\big) = Q\left(\frac{b-T_M}{\sigma_t}\right)$$

Maximization over range of biases

The product in Equation (2) is dependent on the fault bias b. To compute the worst-case tail CDF, we maximize it over all possible values of b. Because the first term is a step function that becomes one at x and the second one is a decreasing function of b, we have

$$\max_b \mathrm{P}\big(\varepsilon > x, |y| \le T_M | H_1(b)\big) = \max_b H(b-x)Q\left(\frac{b-T_M}{\sigma_t}\right) = Q\left(\frac{x-T_M}{\sigma_t}\right)$$

An upper bound of the numerator is then given by

$$\max_b \mathrm{P}(\varepsilon > x, |y| \le T_M) = Q\left(\frac{x}{\sigma_{core}}\right) \cdot (1-\alpha) + Q\left(\frac{x-T_M}{\sigma_t}\right) \cdot \alpha$$

Considering the lower bound on the denominator and the upper bound on the numerator, we get the following upper bound on the tail CDF of the posterior distribution

$$\max_b \frac{\mathrm{P}(\varepsilon > x, |y| \le T_M)}{\mathrm{P}(|y| \le T_M)} \le \frac{Q\left(\frac{x}{\sigma_{core}}\right) \cdot (1-\alpha) + Q\left(\frac{x - T_M}{\sigma_t}\right) \cdot \alpha}{\left(1\text{-}2Q\left(\frac{T_M}{\sigma_y}\right)\right) \cdot (1-\alpha)}$$

This upper bound is a right CDF overbound of the user range error distribution. The left CDF overbound is identical. For this reason, it can be used in the convolution of the different error sources when computing the user position error.

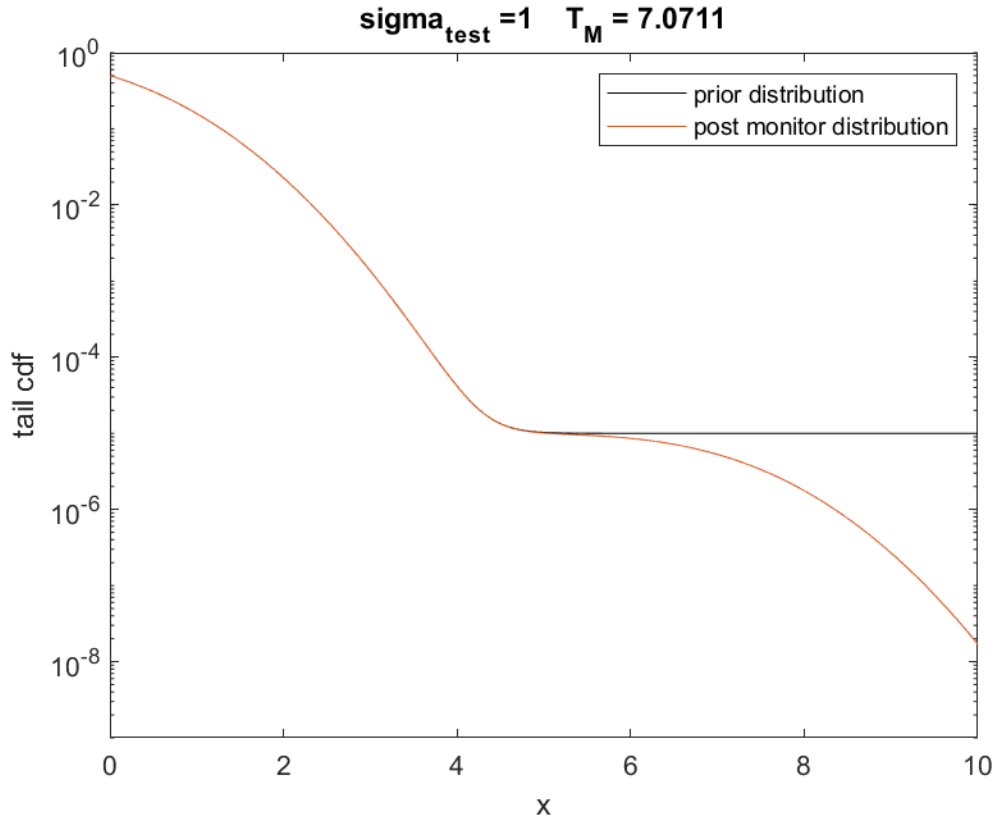

Figure 2. Upper bound on the post monitor distribution for $\sigma_{core} = 1, \alpha = 10^{-5}, T_M = 7.07$, and $\sigma_t = 1$.

Figure 2 shows the resulting tail CDF bound (in orange) for an example with $\sigma_{core} = 1$, $\alpha = 10^{-5}$, $T_M = 7.07$, and $\sigma_t = 1$. The effect of the monitor is to reduce the tails. A smaller σ_t will result in a bigger reduction.

GAUSSIAN BOUND ON POSTERIOR DISTRIBUTION

From the previous section, we have the upper bound on the posterior distribution

$$F(x) = \frac{Q\left(\frac{x}{\sigma_{core}}\right)\cdot(1-\alpha)+ Q\left(\frac{x-T_M}{\sigma_t}\right)\cdot\alpha}{\left(1\text{-}2Q\left(\frac{T_M}{\sigma_y}\right)\right)\cdot(1-\alpha)} \tag{3}$$

To find a gaussian overbound of this tail CDF, we first compute the point β at which *F* is equal to 0.5. This represents the smallest bias for which there exists a Gaussian tail CDF overbound of F. The Gaussian overbound is computed by centering the Gaussian bound at β

$$\sigma_{ob} = \max_{\beta<x<L} \frac{x-\beta}{Q^{-1}(F(x))} \tag{4}$$

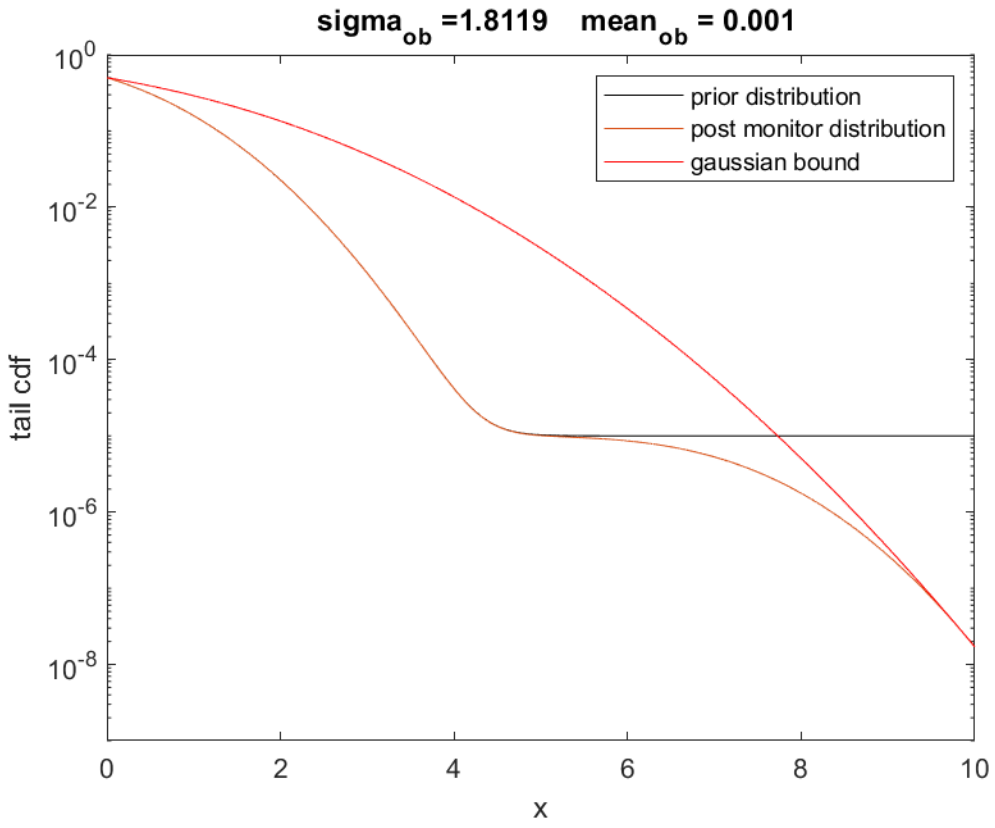

Figure 3. Upper bound on the post monitor distribution for $\sigma_{core} = 1$ *,*$\alpha = 10^{-5}$, $T_M = 7.07$, and $\sigma_t = 1$ *with Gaussian overbounding tail CDF (in red)*

Note that we must choose an upper bound of the interval over which the maximization is performed in (4). This range should be chosen so that the residual probability is negligible even after the convolution of several error sources. For example, for a target of 10^{-7}, making sure that we maximize everywhere $F(x) \leq 10^{-10}$ would be sufficient.

For some parameter combinations, the distribution F(x) is bimodal (in the sense that there are two maxima in the PDF). If this is the case, an intermediate unimodal distribution can be generated (if a rigorous Gaussian overbound is desired, as in the sense of Blanch 2018.) before computing the Gaussian overbound.

LIMITATIONS

The proposed range error model is designed such that the user can compute position error bounds using the Gaussian overbound (or the bimodal distribution). This will usually require that the error distributions for different error sources are independent of each other. In our case, this will mean that we assume that both the prior distributions and the monitor measurement noise are independent across satellites. This will mostly be the case for ground monitors, but it will not necessarily be the case for applications like autonomous integrity monitoring, where the distribution of ε_t for the different range measurements will be very dependent.

EXAMPLE APPLICATION: DUAL FREQUENCY SBAS

As an example, we examine the case of dual frequency SBAS with the clock and ephemeris algorithm as described in Blanch 2011, (in the threshold version). In this algorithm, there is an effective monitor threshold for each user range error, and an effective measurement noise. The range error bound is computed based on the integrity allocation and is derived to protect the range error only. Here we compare the baseline performance between the error bound as derived in Blanch 2011, and the one proposed here.

The design of the monitor is the same in both methods. The only change is in the derivation of the overbounding standard deviation. For the proposed method, we assumed that the prior distribution of the correction errors had a core sigma of 0.50 m and a prior probability of fault α of 10^{-4}. The bounding distributions did require a small bias (as pointed out above), but because it was below 10% of the sigma overbound, it was deemed negligible (at least for this analysis). Figures 4 and 5 show the 99.9% percentile of the resulting VPL over a day computed using MAAST (with a 2 by 2 user grid, sampled every 5 minutes) for the baseline method and the proposed method. A 24 satellite GPS constellation was assumed. The proposed method shows error bounds that are reduced by almost 30%.

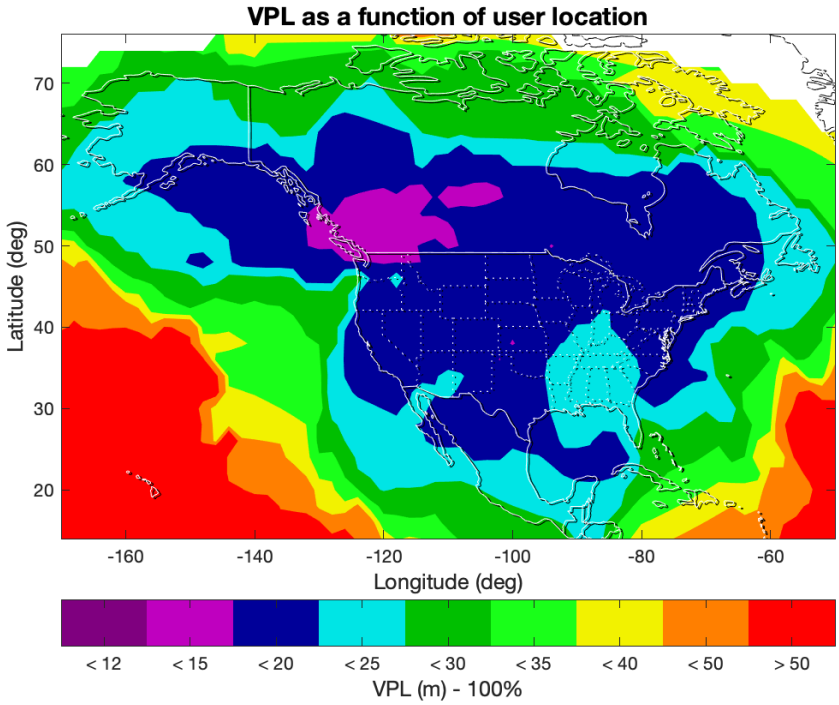

Figure 4. Vertical Protection Levels in dual frequency SBAS for a baseline clock and ephemeris bounding algorithm

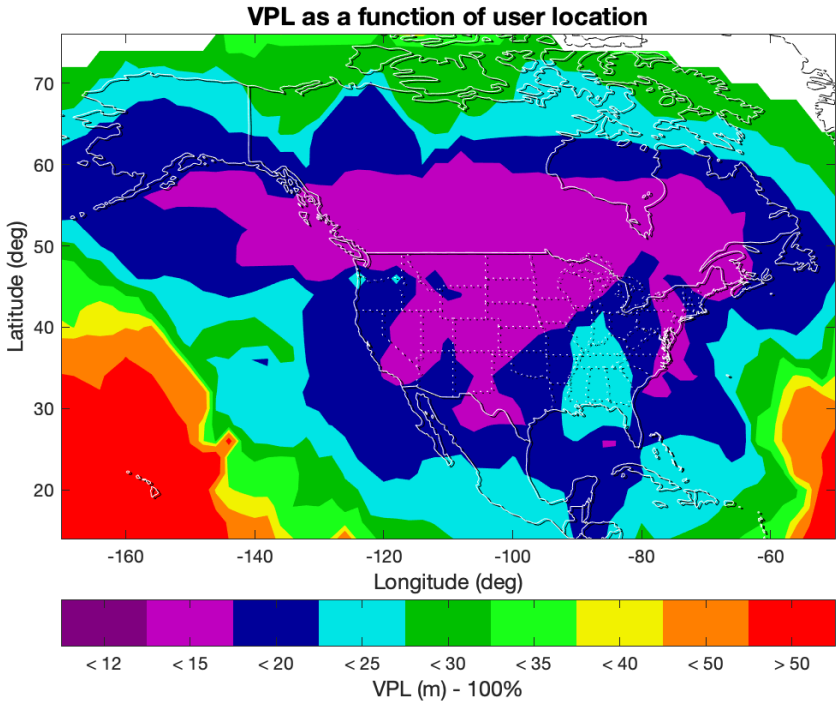

Figure 5. Vertical Protection Levels in dual frequency SBAS for a clock and ephemeris bounding algorithm using the proposed bounding method.

ACKNOWLEDGEMENTS

This work was funded by Hexagon.

REFERENCES

DeCleene, Bruce, "Defining Pseudorange Integrity - Overbounding," Proceedings of the 13th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GPS 2000), Salt Lake City, UT, September 2000, pp. 1916-1924.

Rife, J., Pullen, S., Pervan, B., and Enge, P. Paired Overbounding for Nonideal LAAS and WAAS Error Distributions. IEEE Transactions on Aerospace and Electronic Systems, 2006, 42, 4, 1386 -1395.

Blanch J., Walter, T., and Enge, P., “Gaussian Bounds of Sample Distributions for Integrity Analysis”. IEEE Transactions on Aerospace and Electronic Systems.” PP. 1-1. 10.1109/TAES.2018.2876583.

Blanch, Juan, Liu, Xinwei, Walter, Todd, "Gaussian Bounding Improvements and an Analysis of the Bias-sigma Tradeoff for GNSS Integrity," Proceedings of the 2021 International Technical Meeting of The Institute of Navigation, January 2021, pp. 703-713. https://doi.org/10.33012/2021.17861

Walter, Todd, Gunning, Kazuma, Phelts, R. Eric, Blanch, Juan, "Validation of the Unfaulted Error Bounds for ARAIM", NAVIGATION, Journal of The Institute of Navigation, Vol. 65, No. 1, Spring 2018, pp. 117-133.

Walter, T., Blanch, J., Gunning, K., Joerger, M., & Pervan, B. (2019). Determination of Fault Probabilities for ARAIM. *IEEE Transactions on Aerospace and Electronic Systems*, *55*(6), 3505-3516. [8684918]. https://doi.org/10.1109/TAES.2019.2909727

Blanch, J., Walter, T., and Enge, P., “A Clock and Ephemeris Algorithm for Dual-Frequency SBAS,” *Proceedings of the Institute of Navigation GNSS-11*, Portland, September 2011.