Sigma Estimation, Inflation, and Monitoring In the LAAS Ground System

Boris Pervan[*], Sam Pullen[†], and Irfan Sayim[*]

[*]*Illinois Institute of Technology*
[†]*Stanford University*

BIOGRAPHY

Boris Pervan received a B.S. from the University of Notre Dame (1986), M.S from the California Institute of Technology (1987), and Ph.D. from Stanford University (1996), all in Aerospace Engineering. From 1987 to 1990, he was a Systems Engineer at Hughes Space and Communications Group. Dr. Pervan was a Research Associate at Stanford from 1996 to 1998, serving as project leader for GPS Local Area Augmentation System (LAAS) research and development. He was the 1996 recipient of the RTCA William E. Jackson Award and the 1999 M. Barry Carlton Award from the IEEE Aerospace and Electronic Systems Society. Currently, Dr. Pervan is Assistant Professor of Mechanical and Aerospace Engineering at the Illinois Institute of Technology in Chicago.

Sam Pullen received two S.B. degrees (in Aeronautics/ Astronautics and History) from the Massachusetts Institute of Technology (1989) and received M.S. (1990) and Ph.D. (1996) degrees from Stanford University in Aerospace Engineering. Since graduating, Dr. Pullen has served as Research Associate and as Technical Manager at Stanford, where he has supported GPS Local Area Augmentation System (LAAS) and Wide Area Augmentation System (WAAS) research and development and now serves as the project leader for LAAS. His work in these fields and his support of the Johns Hopkins University Applied Physics Laboratory (JHU/APL) GPS Risk Assessment earned him the ION Early Achievement Award in 1999.

Irfan Sayim received a B.S. degree from Marmara University (1990), Istanbul, Turkey, Certificate of Business and Administration from University of Istanbul (1991), Istanbul, Turkey. M.E. from Illinois Institute of Technology (1996), in Mechanical and Aerospace Engineering. Currently, he is a Ph.D. candidate at Illinois Institute of Technology and working on GPS Local Area Augmentation System (LAAS) integrity for aircraft precision landing.

ABSTRACT

The prescribed Local Area Augmentation System (LAAS) algorithms for the generation of vertical and horizontal protection levels implicitly assume a zero-mean, normally distributed fault-free error model for the broadcast pseudorange corrections. While such a model appears to be consistent with reference receiver ranging errors due to thermal noise and diffuse multipath, remaining errors such as ground reflection multipath and systematic reference receiver/antenna errors may be non-gaussian, non-zero-mean, or both. To ensure navigation integrity, special care must therefore be taken on the ground in the definition of the broadcast correction error standard deviation (σ_{pr_gnd}). In this regard, this paper defines a new approach toward the estimation and inflation of the broadcast value of σ_{pr_gnd} to account for both ground reflection multipath errors as well as gaussian errors due to thermal noise and diffuse multipath. A sufficient condition for the acceptability of non-zero means is also presented.

An additional requirement concerning the broadcast σ_{pr_gnd} is that it be possible to monitor for unexpected violations (failure events where, during operation, the true sigma exceeds σ_{pr_gnd}) with acceptable residual integrity risk. The use of a 'real-time' sigma estimate as a monitor statistic is undesirable because the real-time estimate is slow to respond to actual sigma violations due

to the averaging required to reduce the estimate error. A better approach based on CUSUM statistical process control is described in this paper. Because the CUSUM method is quicker to detect sigma increases from nominal to levels above the inflated value of σ_{pr_gnd}, it significantly reduces the residual risk posed by potential σ_{pr_gnd} violations.

1.0 INTRODUCTION

In the Local Area Augmentation System (LAAS), the final quantitative assessment of navigation integrity is realized through the computation of vertical and horizontal protection levels at the aircraft (termed '*VPL*' and '*HPL*,' respectively). In principle, these limits are the position bounds that can be ensured with an acceptable level of integrity risk. For example, for a Category 1 approach, the maximum permissible integrity risk is of the order of 10^{-8} with respect to a vertical alert limit (*VAL*) of 10 m. [5] The prescribed algorithms for the generation of the protection limits assume a zero-mean normally distributed fault-free error model for the broadcast pseudorange corrections. The standard deviation of correction error is presumed by the aircraft to be equal to the broadcast value of 'σ_{pr_gnd}' for each satellite.

This paper first briefly summarizes relevant prior work concerning the sensitivity of integrity risk to statistical uncertainty in the knowledge of correction error standard deviation (σ_{pr_gnd}). The effect of a non-zero mean (μ_{pr_gnd}) is then considered, and sufficient conditions under which μ_{pr_gnd} can be neglected are derived. Next, ground reflection multipath is considered in some detail. Because the environmental factors that affect ground multipath change very slowly with time, it is impractical to rely on experimental data alone. Therefore, theoretical approaches are emphasized in this paper. Finally, two sigma-monitoring algorithms are also examined for the detection of anomalies that cause the true correction error standard deviation to exceed the broadcast σ_{pr_gnd} during LAAS operations.

1.1 Prior Results for Gaussian Errors

The sensitivity of LAAS integrity risk was investigated and quantified with respect to the statistical uncertainty in the knowledge of reference receiver error standard deviation and correlation between multiple reference receivers. A detailed methodology was presented in reference [1] to define the minimum acceptable buffer parameters for the value of σ_{pr_gnd} broadcast to the aircraft. This work implicitly addressed only the gaussian error structures associated with receiver thermal noise and diffuse multipath. In this context, it is shown in [1] that any value of σ_{pr_gnd} may be broadcast provided that the following inequality is satisfied

$$\sigma_{pr\_gnd} > \sigma^* \sqrt{1+(M-1)\rho^*} / \sqrt{M}$$

Where $\sigma^* \equiv a(n_s)$, $\rho^* \equiv r + b(n_r)$, s and r are the maximum values of standard deviation and correlation coefficient for any reference receiver and reference receiver pair, respectively, and n_s and n_r are the number of independent samples used to obtain s and r. For the Category 1 case where M = 3, Figures 1A and 1B show the values of a(·) and b(·) plotted against n_s and n_r for a number of discrete values (contours) of integrity risk tolerance (the amount of additional integrity risk, beyond nominal, that is tolerable). In each figure, the contours from top to bottom correspond to integrity risk tolerances of 1*%,* 5*%,* 10*%,* and 30*%,* respectively. As one would naturally expect, Figure 1A illustrates that for any given value of n_s, the buffer factor decreases as the integrity risk tolerance is relaxed. Analogous behavior is exhibited for the required correlation buffer parameter in Figure 1B. It is also clear that only marginal reductions in buffer parameters will be realized for sample sets larger than 200 points. However, it is equally evident that sample sets smaller than 100 points will typically require rather large buffer parameters.

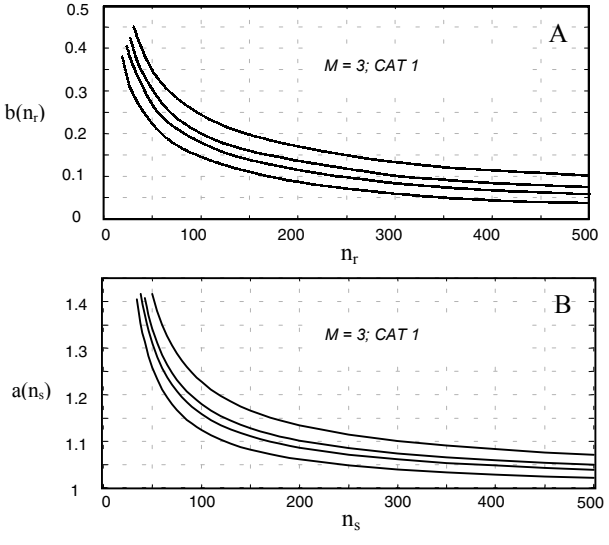

Figure 1A-1B. Sigma/Correlation Buffer Parameters vs. Number of Samples

2.0 THE EFFECT OF A NON-ZERO MEAN

The existence of a non-zero mean in the correction error can result unacceptable integrity risk. Therefore, either the mean must be specifically accounted in LAAS integrity risk or it must be shown to have a negligible effect on integrity. Since the LAAS protection level equations have no explicit parameters with which to account for non-zero means, the second alternative is pursued here.

The computed LAAS *VPL* under the fault-free hypothesis (*H0*) is defined in [5] as

$$VPL_{HO} = k_{md\_ff}\sqrt{\sum_{i=1}^{N} S_{zi}^2[\sigma_{pr\_gnd}^2(i) + \sigma_{pr\_air}^2(i) + \sigma_{pr\_ion}^2(i) + \sigma_{pr\_trop}^2(i)]}$$

where

- i is the satellite index;
- S_{zi} is i-th element of the third row (representing the vertical component) of the weighted geometry projection matrix used to generate the position estimate;
- N is the number of available satellites;
- k_{md_ff} is a multiplier (approximately 5.8 for Category I) used to set the desired level of missed detection probability assuming gaussian errors;
- σ_{pr_air} is the airborne measurement error standard deviation;
- σ_{pr_ion} and σ_{pr_trop} are the standard deviations of residual errors due to ionospheric and tropospheric decorrelation);

When a mean error (μ_{pr_gnd}) exists on the broadcast correction value, then the actual *VPL* bound is given by

$$VPL_{H0act} \le VPL_{H0} + \sum_{i=1}^{N} |\, S_{zi}\mu_{pr\_gnd}(i)\,| \,.$$

To illustrate the sensitivity of *VPL* to nonzero means, simulation results are presented in Figure 2. To generate this figure a depleted (22 out of 24) GPS satellite constellation) was used and a mean magnitude of 0.2 m was assumed for all satellites. An approximate linear relationship between the computed *VPL* (using the nominal equation) and the actual value is exhibited. The residual integrity risk incurred by the existence of the mean is evident in the slope of this linear characteristic, which is larger than unity. In the most stringent interpretation of LAAS and associated Ground-Based Augmentation System (GBAS) requirements, a slope greater than unity is unacceptable. Unfortunately this situation is unavoidable since the LAAS *VPL* equation does not explicitly account for non-zero means. In a more practical interpretation of requirements, however, one may note that for a Category I approach using a Space Based Augmentation System (SBAS), the required *VAL* is 12 m – in contrast to 10 m for GBAS/LAAS. The implication is that while a conservative 'threshold' *VAL* of 10 m is used for GBAS/LAAS, an 'operational' value of 12 m is actually sufficient to ensure integrity. In this regard, it is only necessary to ensure that the slope of the linear characteristic is 12/10 = 1.2 or lower.

Using the equations above, a geometry-independent upper bound for the actual *VPL* may be defined:

$$VPL_{act} \le \left(1 + \frac{\beta\sqrt{N}}{k_{md}}\right) VPL_{comp},$$

where

$$\beta = \max_i \left|\frac{\mu_{pr\_gnd}(i)}{\sigma_{pr\_gnd}(i)}\right| .$$

This bounding result also applies for the $H1$ case (failed reference receiver) if '$H0$' is replaced with '$H1$' and k_{md_ff} is replaced with k_{md} (which holds a minimum value of 2.878 for Category I).

To ensure a limit slope of less than 1.2, it is sufficient to limit the acceptable value of β (which is simply the maximum magnitude of the mean relative to sigma in the range domain). For the worst case where $N_{\max} = 12$, the maximum acceptable values of β are 0.335 and 0.167 for the $H0$ and $H1$ cases, respectively. Clearly, for both conditions to be met: $\beta < 0.167$. Note that this is a *sufficient condition* to ensure that the mean is negligible. It may be achieved either by verifying that the range-domain mean is small enough in comparison to the standard deviation or by inflating the broadcast standard deviation as necessary to ensure that the sufficient condition is satisfied.

Since the condition defined above is sufficient but not strictly necessary (in particular, it is conservative for good satellite geometries), alternatives are possible. For example, we may consider at the reference station (at any given time) all subset geometries that the aircraft may use to compute *VPL*. First assume the largest specification-compliant values of standard deviation for *air*, *iono*, and *tropo* errors, and for all candidate subset geometries compute both VPL_{H0} and VPL_{H1}. The LAAS ground station verifies that in all cases where these *VPL*s are less than 10 m, that the addition of the unaccounted for position domain mean term does not cause the actual

bounding *VPL* to exceed 12 m. Note that by itself, this is not truly a sufficient condition because the *air*, *iono*, and *tropo* standard deviation values used at the aircraft may be smaller than the maximum values assumed on the ground. If they are smaller at the aircraft, it is possible that geometries exist for which both *VPL*s computed on the ground exceed 10 m (and thus are not verified) but at the aircraft are slightly less than 10 m. To protect against such occurrences, the ground station would also have to consider a *range* of standard deviation values for *air*, *iono*, and *tropo* below the largest values assumed above. Such an approach would probably be computationally intensive, but is possible in principle.

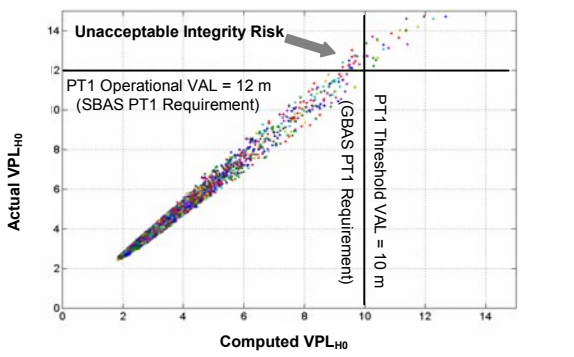

Figure 2. Actual VPL_{H0} vs. computed VPL_{H0} with depleted constellation (22out of 24 SVs)

3.0 GROUND REFLECTION MULTIPATH

As noted earlier, due to the slowly varying environmental factors involved, it is unlikely that the effect of ground multipath can be quantified by experimental means alone. In this section, therefore, we define a number of candidate theoretical approaches for the establishment and inflation of correction error standard deviation to account for the effect of ground reflection multipath. In this analysis, we assume a Multipath Limiting Antenna (MLA) implementation at the LGF. Under this assumption, the amplitude of ground-reflected signal relative to direct signal is less than -30 dB (0.032). The intention of making such an assumption is not necessarily to restrict LAAS ground implementations to the use of the MLA, but instead to make the task of demonstrating at least one acceptable means for sigma establishment more easily tractable. In the analysis that follows, a number of variables and parameters are used:

d is half correlator spacing (e.g, 0.05 chip or 15 m)
δ is multipath delay
α is amplitude of reflected signal relative to direct
θ is phase of reflected signal relative to direct
e is ranging error due to ground reflection multipath
h is antenna height
E is satellite elevation

Following the multipath development in [2] and neglecting second order and higher terms in α (since $\alpha^2 << 1$), the resulting multipath delay is:

$$e = \min[\delta, d]\alpha\cos\theta ,$$

where $\delta = 2h\sin E$.

Now define a constant c for a given antenna height and satellite elevation angle: $c = \min[2h\sin E, d]$. The *normalized* MP error is the defined as:

$$\hat{e} = e/c = \alpha\cos\theta .$$

The variation of multipath error vs. delay is plotted in Figure 3.

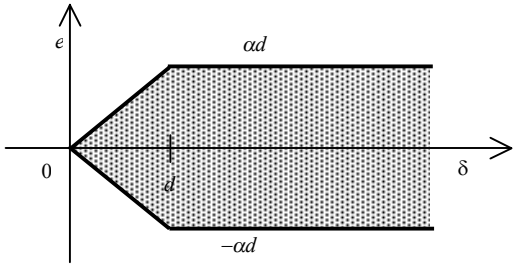

Figure 3. MP error vs. MP delay

For a given satellite location, θ will vary over widely spaced days between 0 and 2π due to small changes in reflection surface height and reflectivity. Therefore, θ is typically treated as a uniformly distributed random variable: $\theta \sim U[0,2\pi]$. The corresponding probability density function (pdf) is shown in Figure 4.

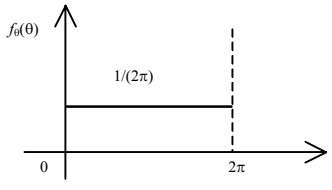

Figure 4. Uniformly distributed RV θ

Because θ is passed through a cosine function in our application, we define a new random variable (RV) z is appropriate as it follows:

$$z = \cos\theta$$

The corresponding pdf for z is

$$f_z(z) = \frac{1}{\pi\sqrt{1-z^2}}, \qquad \forall \qquad |z| < 1.$$

This function is plotted in Figure 5.

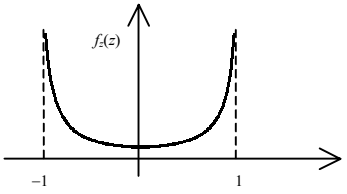

Figure 5. Pdf for RV z.

In the case of reflection amplitude α, we will consider three different models: Uniformly distributed, Rayliegh distributed, and constant value (a deterministic model). These models are expressed mathematically below:

Model 1: Uniform α [Figure(6)].

$$f_\alpha(x) = \begin{cases} 1/b & 0 \le x \le b \\ 0 & otherwise \end{cases}$$

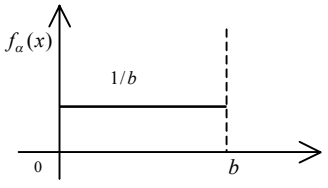

Figure 6. Pdf for uniformly distributed α

Model 2: Rayleigh α [Figure(7)].

$$f_\alpha(x) = \frac{x}{a^2}\exp\left[-\frac{x^2}{2a^2}\right], \qquad \forall \qquad x > 1$$

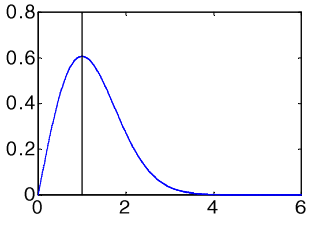

Figure 7. Pdf for Rayleigh distributed α

Model 3: Constant α.

$\alpha = b$

In each case, the multipath (MP) error is the product of α and z.

3.1 Model 1: Uniformly Distributed α

In this case, the normalized MP (NMP) error is the product of RV z and the uniformly distributed RV α. The resulting pdf of NMP error can be shown to be

$$f_{\hat{e}}(x) = \frac{x}{a^2}\ln\left[\frac{1+\sqrt{1-(x/b)^2}}{|x|/b}\right], \qquad \forall \qquad |x| > 1.$$

This function is plotted in Figure 8 for an example value of $b = 1$.

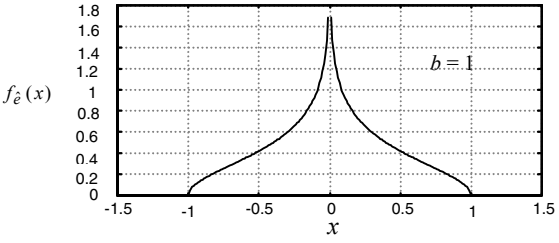

Figure 8. Pdf of MP error due to Model 1.

It is clear that the pdf is symmetric, unimodal, and truncated. The associated cumulative distribution function (cdf) of NMP error is:

$$F_{\hat{e}}(x) = 1 + \frac{1}{\pi}\frac{\ln\left[|b/x| + \sqrt{(b/x)^2 - 1}\right]}{b/x} - \frac{1}{\pi}\sec^{-1}(b/x)$$

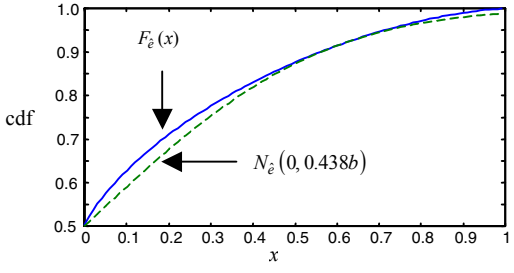

Figure 9. Gaussian overbounding of Model 1

Because the pdf is truncated it is possible to overbound the cdf of NMP error with a gaussian cdf for all values of x. In Figure 9, it is shown that a gaussian cdf with $\sigma \ge 0.438b$ is sufficient in this regard. The significance

of this result is made evident by the use of a theorem derived in reference [4], which states that if a RV has the following properties:

a. Cdf is overbounded by gaussian for all x,
b. Pdf is symmetric, and
c. Pdf is strictly unimodal,

then the cdf resulting from an arbitrary linear combination of such random variables is also overbounded by the cdf of the linear combination the corresponding gaussian RVs for all x. Because Model 1 satisfies these three conditions, range-domain overbounding implies position-domain overbounding.

3.2 Model 2: Rayleigh Distributed α

In this case, the normalized MP (NMP) error is the product of RV z and the Rayleigh distributed RV α. The resulting pdf of NMP error can be shown to be [6]

$$f_{\hat{e}}(x) = \frac{x}{a\sqrt{2\pi}} \exp\left[\frac{x^2}{2a^2}\right] \equiv N_{\hat{e}}(0, a)$$

Since the result is gaussian, the issue of overbounding need not be considered. However, an appropriate value of the parameter a must be selected. In this regard, we consider two properties of Rayleigh Distribution (as shown in Figure 10):

A- Maximum-likelihood value (MLV) of $\hat{e} = a$
B- Expected value (EV) of $\hat{e} = a\sqrt{\pi/2}$

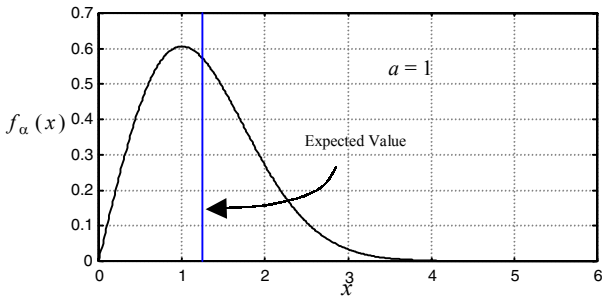

Figure 10. Rayleigh pdf

Therefore, if we select $a \equiv b$ (b=–30dB is the MLV of α), any gaussian with $\sigma \geq b$ may be used. Alternatively, if we choose $a \equiv b/\sqrt{\pi/2}$ (b=–30dB is the EV of α), we may use any gaussian such that $\sigma \geq 0.798b$. Note that while some flexibility clearly exists in the precise implementation of the Rayleigh model, the fact that the magnitude of α will always be unbounded suggests that the results derived from the model may be conservative. Conversely, the relative advantage of the Rayleigh model is that it can accommodate uncertainty (should there be any) in the knowledge of the value of the maximum reflection strength.

3.3 Model 3: Constant value of α

With this model, NMP error will simply the product of a scalar constant (α=b=–30dB) and the RV z. The resulting pdf of NMP error is therefore

$$f_{\hat{e}}(x) = \frac{1}{\pi b\sqrt{1-(x/b)^2}}, \qquad \forall \qquad |x| < b,$$

which is plotted in Figure 11 for b = 1.

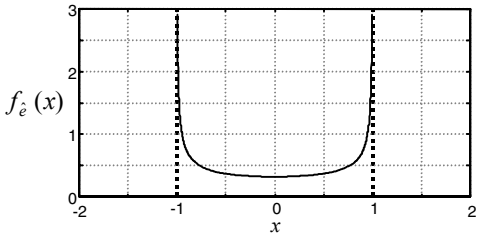

Figure 11. Pdf for MP error due to Model 3.

The associated cdf of normalized MP error is:

$$F_{\hat{e}}(x) = \frac{1}{2} + \frac{1}{\pi}\sin^{-1}\frac{x}{b}.$$

Note that, in this case, the NMP error pdf is symmetric but not unimodal. While we can find a gaussian distribution to overbound the cdf at all x (since the pdf is truncated), this property will not be retained after convolution. Therefore, range domain overbounding does not ensure overbounding in the position domain, and this case must be treated in a different way from the previous two models.

A first approach, which is simple and conservative, is to treat NMP error as a mean offset $\mu = b$. In this case we may apply the results for mean overbounding derived earlier in this paper: With 12 satellites, to protect an operation *VAL* of 12 m with a threshold *VAL* of 10 m, it must be that $b < 0.167\sigma$. Thus NMP error will be overbounded with a gaussian with $\sigma \geq 6.02b$. In the general case where corrections for N < 12 satellites are broadcast by the ground station $\sigma \geq (5\sqrt{N}/k_{md})b$ is sufficient.

It is important to clarify that his result is directly applicable for a *constant phase* model, but it is conservative for uniform phase model under consideration. The reason is that the NMP error actually

varies between $+b$ and $-b$, while the mean overbound assumes the worst-case magnitude and sign for each satellite. To accommodate the uniform phase model, we consider the direct convolution of Model 3 pdfs. In the analysis which follows, we rely on the fact that it is only strictly necessary to overbound above 2.878σ (set by the minimum k_{md} value for VPL_{H1}) for linear combinations of up to 12 SVs.

Direct *analytical* convolution of Model 3 pdfs yields integrals that are not tractable in closed form. The characteristic function (Fourier Transform) of the Model 3 pdf can be readily shown to be a Bessel Function of first kind. However, the inverse Fourier Transform of products (equivalent to convolution in the range domain) of Bessel Functions is not readily accessible in closed form. Furthermore, direct *numerical* convolution of Model 3 pdfs is also difficult since the pdf function has singularities at $\pm b$, requiring impractically fine discretization for accurate results. To circumvent these difficulties, we introduce a conservative approximation for the Model 3 pdf as follows:

$$f_{\hat{e}}(x) = \tfrac{1}{2}[\delta(x-b) + \delta(x+b)]$$

A graphical comparison of the actual Model 3 pdf and this conservative model is shown in Figure 12.

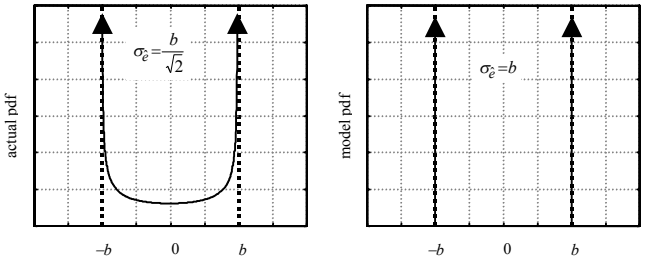

Figure 12.

We first consider the hypothetical limit case of convolution of large number of independent, identically distributed (i.i.d.) sources. In this case, the Central Limit Theorem defines necessary bounding value of sigma for actual pdf $\sigma \geq \sigma_{\hat{e}} = b/\sqrt{2}$ and for the conservative model $\sigma \geq \sigma_{\hat{e}} = b$. Unfortunately, while this result holds when N is very large, it is not sufficient for finite values of N. For the example case where N=9, the cdf plots in Figure 13 shows that using a gaussian pdf with $\sigma = \sigma_{\hat{e}} = b$ will not always overbound in the region of interest (above 2.878).

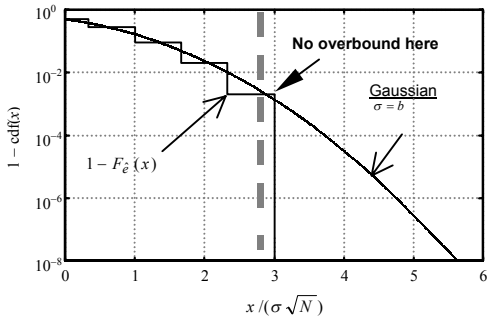

Figure 13.

From this example, it is clear that further inflation of gaussian σ will be necessary.

For a given satellite i, we assume the NMP error e_i distributed according to our model:

$$e_i \sim \tfrac{1}{2}[\delta(x-\beta_i) + \delta(x+\beta_i)] \qquad (i=1,\ldots N)$$

Since $\sigma_i = \beta_i$, the maximum possible NMP error resulting from a linear combination of the N sources is:

$$e_{\max} \equiv \max \sum_{i=1}^{N} e_i = \sum_{i=1}^{N} \beta_i = \sum_{i=1}^{N} \sigma_i \leq \sqrt{N}\sqrt{\sigma_1^2 + \sigma_2^2 + \ldots + \sigma_N^2} \equiv \sqrt{N}\sigma_{tot}$$

Equality (largest e_{max}) in the bound above occurs when all β_i are the same (i.e., N i.i.d. sources). In this case, no inflation of σ_i is required when $\sqrt{N} \leq 2.878$ ($N \leq 8$) since error will never exceed $2.878\sigma_{tot}$.

In general, an inflation/deflation factor $\sqrt{N}/2.878$ may be used in the range domain to ensure that the position domain NMP error never exceeds $2.878\sigma_{tot}$. For example for N =12, an inflation factor of 1.204 is implied. However, a zero probability of exceeding $2.878\sigma_{tot}$ is clearly not necessary. We only require that gaussian *bounds* the actual NMP error in the cdf sense above $2.878\sigma_{tot}$. A sufficient inflation factor of 1.05 is obtained from direct convolution of N i.i.d. sources. This result is sufficient for all N up to 12 (and conservative for N<9). Figure 14 illustrates the cdf overbounding results for N = 12.

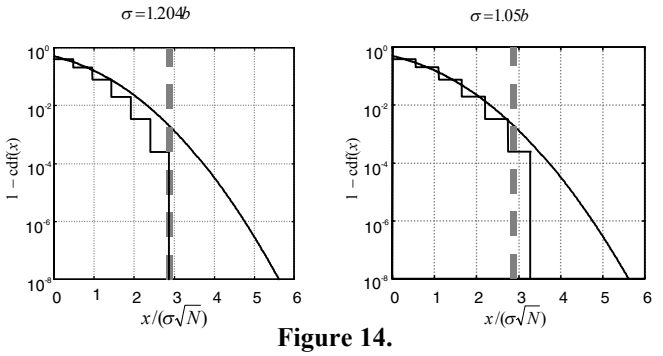

Figure 14.

Because Model 3 does not meet the strict conditions of reference [3], we must also explicitly consider the effect of the addition of other contributing gaussian sources (due to diffuse MP and receiver noise). Clearly, the addition of such errors is not an issue in the following limiting cases:

- Gaussian sources are very small: Model 3 dominates, so an inflation factor of 1.05 is sufficient.
- Gaussian sources are very large: Model 3 errors are negligible by comparison (Model 3 inflation factor is irrelevant).

In cases where gaussian and Model 3 errors are of comparable size, the results are again evaluated by direct convolution. Figure 15 shows an example result of the convolution of N=12 Model 3 sources and N=12 gaussian sources. An inflation factor of 1.05 was applied to overbound the Model 3 NMP error components. Each plot in Figure 15 corresponds to a different relative magnitude (0.01 to 10) of gaussian to ground multipath error. In all cases, the inflation factor of 1.05 is seen to be sufficient to ensure overbounding.

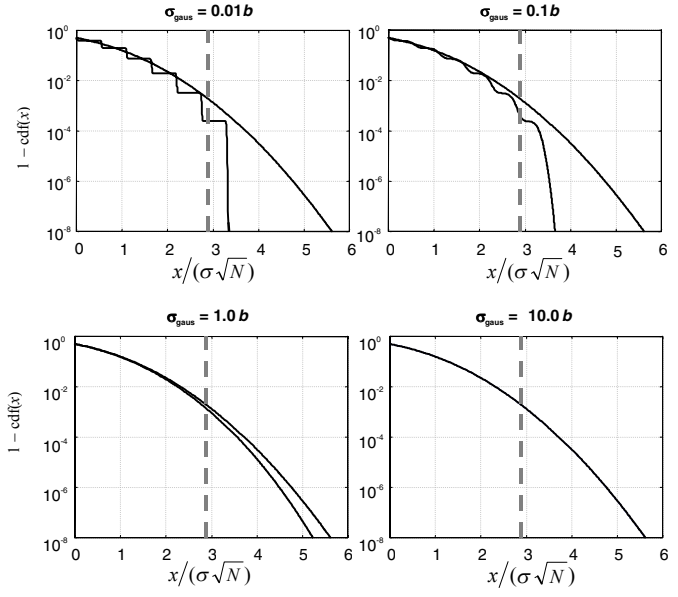

Figure 15. Convolution of N Model 3 sources with gaussian sources

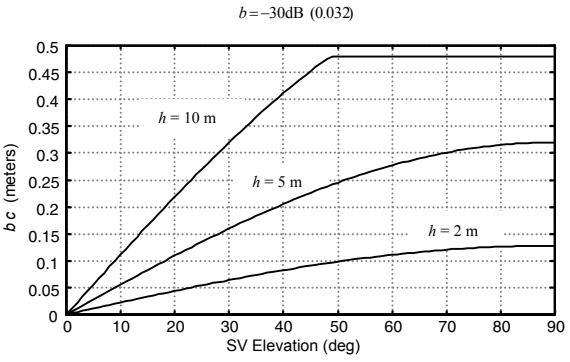

Figure 16. Size of Product $b\,c$ vs. SV Elevation

3.4 Summary of Error Models Results

Model 1: Uniform Phase, Uniform Reflection Strength

Bound Ground MP with:

$\sigma_e \geq 0.438bc$ where $c \equiv \min[2h\sin E, d]$

Model 2: Uniform Phase, Rayleigh Reflection Strength

Bound Ground MP with:

Case A: $\sigma_e \geq bc$ Case B: $\sigma_e \geq 0.798bc$

Model 3: Uniform Phase, Constant Reflection Strength

Bound Ground MP with:

$\sigma_e \geq 1.05bc$ (for worst case N=12)

Figure 16 shows the size of product bc vs. satellite elevation for three different antenna heights are plotted.

4.0 SIGMA MONITORING

In addition to selecting broadcast values σ_{pr_gnd} to overbound sigma uncertainty under nominal conditions, real-time sigma monitoring is necessary to detect anomalies that cause the true σ_{pr_gnd} to exceed the broadcast σ_{pr_gnd} during LAAS operations. Many algorithms have been proposed for monitoring sigma, and two different sigma monitor algorithms have been examined and have been implemented in the Stanford LAAS Ground Facility prototype, known as the Integrity Monitor Testbed, or IMT [8]. Details of these methods can be found in [7].

1.) Sigma Estimation: This method estimates sample standard deviations from LGF B-values normalized by their theoretical sigmas and compares results to a time-dependent threshold for each receiver channel tracking a GPS satellite. Under normal conditions, the estimated sigmas follow the chi-square distribution used in [1], which determines the

detection threshold that achieves an allowably low fault-free detection rate.

2.) *Cumulative Sum (CUSUM)*: This method maintains cumulative sums of squared and normalized B-values for each receiver channel tracking a GPS satellite and is updated every two smoothing time intervals (200 sec). A windowing parameter (k) is subtracted from the normalized, squared B-values based on target out of control sigma (σ_I), and the result is reset if the resulting sum falls below zero. If the sum is above zero at any update epoch, the CUSUM is compared to a fixed threshold (h) that does not vary with time.

Sigma monitoring plays an important role in ensuring that the integrity risk posed by the possibility that the true sigma exceeds the broadcast sigma is limited. Because sigma monitoring is limited by the number of independent samples that can be collected in one satellite pass, it cannot detect all cases in which the broadcast sigma is no longer an overbound. Because of this, additional inflation beyond that needed to overbound nominal sigma uncertainty may be needed to provide margin so that sigma monitoring can meet the LAAS Ground Facility integrity requirements [9].

Figure 17 compares the performance of sigma estimation to that of several varieties of CUSUM monitoring based on the time to detect, with a missed-detection probability

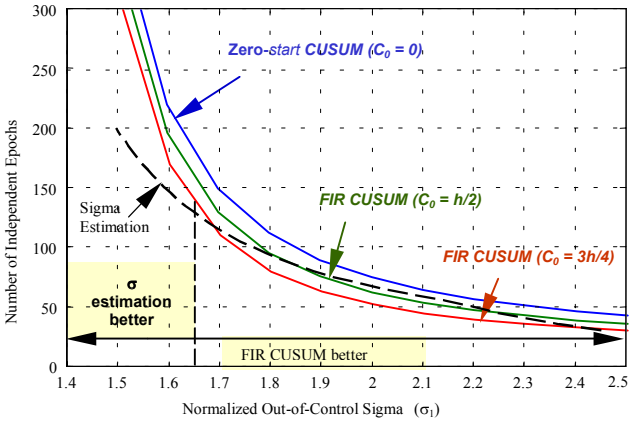

Figure 17. Sigma Estimation Monitor

of 0.001 or below, potential sigma violations as a function of the normalized out-of-control sigma, which is the ratio between the actual sigma and the theoretical sigma. Note that depending on how the overbounding inflation factor is determined, the theoretical sigma may be lower than the broadcast sigma; thus it is possible for minor out-of-control sigma cases to not exceed the broadcast sigma.

The results show that sigma estimation is superior to any CUSUM method in detecting "small" or "minimal" bounding violations, where the out-of-control sigma (normalized by the theoretical sigma) is below 1.68. Above this value, the "fast-impulse-response" (FIR) CUSUM is superior, which is important because larger sigma violations lead to larger integrity threats. The FIR CUSUM achieves faster detection by initializing the CUSUM to a non-zero value closer to the threshold every time the CUSUM resets (goes below zero). In this case, it is possible to use a "head start" of 75% of the threshold value without affecting false-alarm performance.

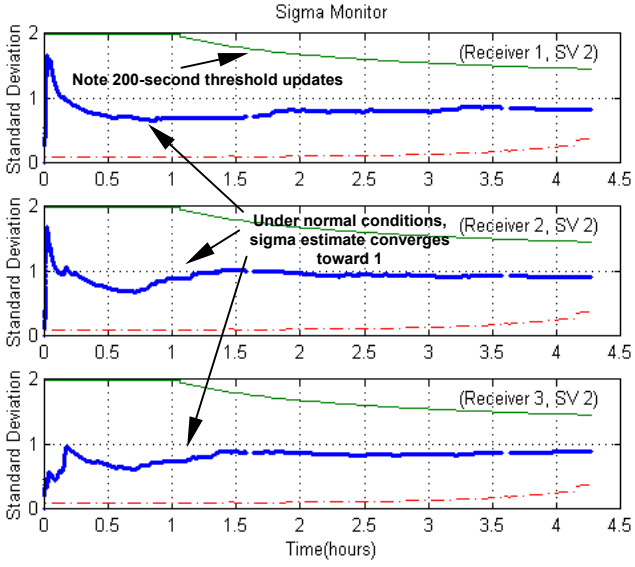

Figure 18. Sigma Estimation Results from IMT Nominal Data

Based on these analytical results, sigma estimation and CUSUMs are both implemented in the Stanford IMT and work together to detect anomalies. Figure 18 shows the results of applying the sigma estimation algorithm to the IMT under nominal conditions. A single satellite (SV 2) is shown here on all three IMT reference receivers. The dark curves show the normalized sigma estimate, and the light curves show the detection thresholds (monitoring starts after 18 independent B-values have been collected, which takes one hour). Note that the normalized sigma estimates stay well below the detection thresholds (which get smaller over time as the number of independent samples increases), and that the sigma estimates converge over time to values just below 1. Thus, at least for this satellite and this set of data, the theoretical sigma values appear slightly conservative. Very similar results have been obtained from the other satellites in this IMT dataset.

The top plot in Figure 19 shows the result of applying the 'zero-start' CUSUM variant to actual IMT data for the same satellite shown in the previous plot (SV 2) and IMT reference receiver 1, and the lower plot shows the normalized B-values that fed the CUSUM. The CUSUM in this case is targeted at an out-of-control sigma twice that of the theoretical sigma (σ_I = 2), which gives a high

windowing factor (k = 1.848). The CUSUM is updated every 200 seconds so that successive updates are statistically independent (if this were not the case, the CUSUM would continually add up similar numbers and would eventually cross the threshold). Because the windowing factor is subtracted off of each independent B^2 update, and the normalized B^2 is usually less than k at the update times, the CUSUM continually resets itself at zero except for a couple of excursions that stay well below the threshold (h) of 36. Very similar results have been obtained from the other satellites tracked by the IMT in this dataset – no flags are observed at all.

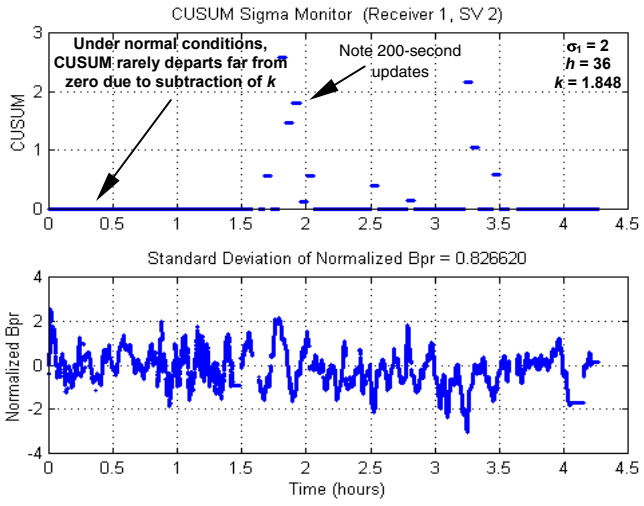

Figure 19. Zero-Start CUSUM Results from IMT Nominal Data

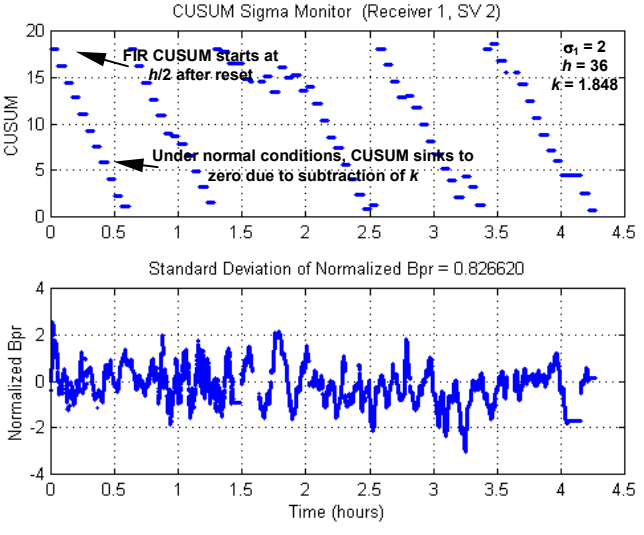

Figure 20. FIR CUSUM Results from IMT Nominal Data

The top plot in Figure 20 is similar to the CUSUM plot from Figure 19 except that an FIR CUSUM is used here on the same IMT data. In this case, the CUSUM is initialized at $h/2$ = 18 and is reset there every time the CUSUM falls below zero. As explained in the previous paragraph, under nominal conditions, B^2 is usually below k; thus the CUSUM slowly falls toward zero and is then reset. In one case (at 3.5 hours), the CUSUM begins to head higher than 18 but does not get far before resuming its expected decline toward zero, and the threshold of 36 is never threatened. Again, the other satellites included in this IMT dataset show essentially the same pattern.

5.0 SUMMARY

To accommodate all contributing, independent error sources, the suggested approach to establishing σ_{pr_gnd} is is to compute the root-sum-square (RSS) of the stadard deviations due to:

Receiver noise: Obtained from bench test data with specified LAAS interference environment, and inflated using the results of [1].

Diffuse multipath: Obtained from short duration field data and inflated using the results of [1].

Ground reflection multipath: Obtained from by theoretical means such as those described in this paper/Analysis ground multipath sigma. Use data to the extent practical to ensure validity of the underlying physical models.

Reflections from discrete objects may be mitigated by masking, using large antenna heights, or handled using the same methods as ground multipath.

In addition to ensuring that the broadcast sigmas overbound the true sigmas under nominal conditions, sigma monitoring is needed to detect violations of sigma overbounds that are due to unexpected anomalies. Both sigma estimation and CUSUM methods are useful in this respect. Sigma estimation more rapidly detects small violations, and the FIR CUSUM variant more rapidly detects significant violations that would be a larger threat to user integrity. Both methods are limited by the 200-second interval between independent B-values, as in all cases, more than one hour (18 independent samples) is needed to guarantee detection of significant violations with $P_{MD} \leq 0.001$. However, CUSUM *mean* times-to-detect are much shorter and are well under one hour for large violations [7]. Both methods work smoothly and predictably within the Stanford Integrity Monitor Testbed prototype, and failure testing with induced sigma violations will begin soon.

ACKNOWLEDGEMENTS

The constructive comments and advice regarding this work provided by R. Braff, B. DeCleene, P. Enge, D. Lamb, C. Shively and C. Varner, and J. Warburton are

greatly appreciated. The authors gratefully acknowledge the Federal Aviation Administration for supporting this research. However, the views expressed in this paper belong to the authors alone and do not necessarily represent the position of any other organization or person.

REFERENCES

[1] B.Pervan and I.Sayim, "Issues and Results Concerning The LAAS σ_{pr_gnd} Overbound," IEEE PLANS 2000, San Diego,CA.

[2] M. Braasch, "Multipath Effects, "Global Positionin System: Theory and Applications, Volume 1,AIAA Progress in Aeronautics and Astronautics Vol. 163,1996

[3] B. Pervan, "Draft SARPS Guidance Material for the GBAS Mean Requirement," paper delivered to B. DeCleene, May 19, 2000

[4] B. DeCleene, "Proof for ICAO Overbounding Requirement," Paper distributed to ICAO Overbounding Subgroup , December 2,1999

[5] RTCA (SC-159/WG-4), "Minimum Aviation System Performance Standards for the Local Area Augmentation System (LAAS)," RTCA/DO-245, RTCA Inc., Washington DC, 28 September 1998.

[6] J. Proakis, Digital Communications, McGraww-Hill, 1983.

[7] S.Pullen, *et.al*., "The Use of CUSUMs to Validate Protection Level Overbounds for Ground Based and Space Based Augmentations Systems," *Proceedings of ISPA 2000*. Munich, Germany, 18-20 July 2000.

[8] M.Luo, S.Pullen, et.al., "Development and Testing of the Stanford LAAS Ground Facility Prototype," Proceedings of ION NTM-2000. Anaheim, CA., 26-28 January 2000, pp.210-219.

[9] Specification: Performance Type One Local Area Augmentation System Ground Facility. U.S. Federal Aviation Administration, Washington, D.C., FAA-E-2937, 21 September 1999.