Core Overbounding and its Implications for LAAS Integrity

Jason Rife, Sam Pullen, *Stanford University*
Boris Pervan, *Illinois Institute of Technology*

BIOGRAPHY

Jason Rife is a Research Associate studying the Local Area Augmentation System (LAAS) at Stanford University. After receiving his B.S. in mechanical and aerospace engineering from Cornell University (1996), he spent one year working in the turbine aerodynamics group of the commercial engine division of Pratt & Whitney. He resumed his studies at Stanford and earned M.S. (1999) and Ph.D. (2004) degrees in mechanical engineering. His thesis work focused on positioning and navigation technologies for underwater robots.

Sam Pullen is a Senior Research Engineer for GNSS research at Stanford University. He has supported the FAA in developing LAAS architectures, requirements, and integrity algorithms since receiving his Ph.D. from Stanford in 1996. He has also developed performance assessment and optimization methods for LAAS, WAAS, GPS III, and the Stanford Gravity Probe B Relativity Mission. He was awarded the ION Early Achievement Award in 1999.

Boris Pervan is an Assistant Professor of Mechanical and Aerospace Engineering at the Illinois Institute of Technology in Chicago. He 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 LAAS research and development. He was 1996 recipient of the RTCA William E. Jackson Award and the 1999 M. Barry Carlton Award from the IEEE Aerospace and Electronic Systems Society.

ABSTRACT

Empirical studies have observed non-Gaussian behavior in pseudorange correction errors for the Local Area Augmentation System (LAAS). This paper introduces an overbounding technique, called core overbounding, designed to validate integrity for error distributions with heavier than Gaussian tails. The core overbounding process generates a conservative representation of the actual error distribution by decomposing the bound into two parts: an explicit function describing the distribution core and an implicit function describing the far tails. For LAAS applications, it is convenient to express the explicit core bound as a Gaussian Core (GC) or as a Gaussian Core with Gaussian Sidelobes (GCGS). These forms balance bound sharpness, important for heavy-tail mitigation, with operational simplicity, necessary for bound broadcast and protection level establishment in a real-time implementation. Analysis using the GC and GCGS bounds indicates that heavy tails have a relatively weak impact on availability for Category I LAAS but a significant impact for Category III.

INTRODUCTION

In a Ground-Based Augmentation System (GBAS), the ground station measures GPS signals and transmits differential pseduorange corrections. The ground station also broadcasts error statistics to enable airborne users to determine bounds for navigation accuracy and integrity. Current implementations of GBAS, such as the Local Area Augmentation System (LAAS) have assumed these error statistics follow a Gaussian distribution.

In fact, LAAS experiments have repeatedly measured error distributions with heavier-than-Gaussian tails [1]. Heavy tails occur even in tests designed to minimize mixing, a phenomenon which may produce apparent heavy tails when data are pooled from disparate Gaussian distributions. For instance, the Federal Aviation Administration Technical Center (FAATC) has observed that heavy-tailed statistics describe multiple passes by a single satellite observed at a fixed azimuth and elevation.

In broadcasting error statistics to users, the ground station must account for heavy tails and other irregularities in the error distribution. These irregularities may vary somewhat with satellite azimuth, satellite elevation, receiver hardware, terrain, and season. To keep system

analysis manageable, the ground station employs a conservative representation of the error distribution, called an overbound, that represents the worst possible error distribution in the absence of a hardware fault. In the design process, engineers use the overbound to construct monitoring thresholds and to validate allocations for continuity and availability. In real-time operations, users employ the overbound to construct a conservative protection level (PL) that is used to detect integrity hazards which occur when the PL has magnitude greater than a predefined alert limit (AL). The PL and AL are illustrated in Figure 1.

System integrity defines the primary constraint on overbounding. To protect integrity, the overbound tail probability, beyond the PL, must exceed the actual distribution's tail probability. This condition implies that the overbound must be conservative in both the range domain and the position domain, as the PL, which describes a position-error envelope, is derived from overbounds on line-of-site measurements, which describe range-domain errors.

In addition to the integrity constraint, the overbound must also satisfy secondary constraints on availability, continuity and real-time operations. Availability and continuity are reduced if an integrity overbound is overly conservative. Thus availability and continuity constraints favor a sharp overbound, a bound that is conservative but that closely resembles the shape of the actual error distribution. At the same time, the shape of the overbound must allow for transmission to airborne users who compute their PL "on the fly." A simple, parameterized form of the overbound is required to support limited-bandwidth communication and rapid PL computation. Because of the difficulty in achieving a simply parameterized yet sharp bound, separate overbounds may be employed for off-line design and for real-time operations, as depicted in Figure 2. Ultimately, all four factors—integrity, availability, continuity, and real-time operations—act in opposition to influence the selection of a heavy-tail GBAS overbound.

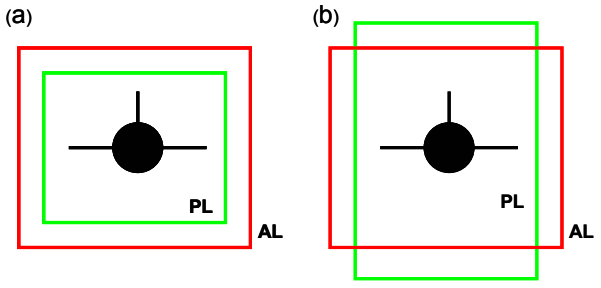

Figure 1. Defining Integrity Risk: Protection Level (PL), represents an instantaneous error envelope, evaluated on the fly. (a) The integrity specification is met if PL remains smaller than a predefined Alert Limit (AL). (b) Integrity is not satisfied if the PL extends beyond the AL.

Core bounding proves a useful tool to develop heavy-tailed overbounds that balance conservatism, sharpness and operational functionality. This paper introduces the core bounding concept, which is characterized by the decomposition of the error distribution into a well-defined core region and a weakly-defined tail region. The first section of this paper derives the mathematical foundations for the core bound and shows that core bounds guarantee conservatism in both the range and position domains. A second section develops two specific types of core bounds: the Gaussian Core (GC) and the Gaussian Core with Gaussian Sidelobes (GCGS). Whereas the GC bound is consistent with the requirements for real-time LAAS operations, the GCGS bound, with its extra parameters, provides a sharper fit that enables more precise error analysis during off-line design. A final section of the paper explores applications of the GC and GCGS bounds for LAAS. Specifically, two applications are considered: (1) formal validation of Gaussian bounds for heavy-tail error distributions and (2) mitigation of the impact of heavy tails on system availability. Mitigation is achieved by combining the ground station error distribution with other system errors, a process which attenuates heavy tails and drives the combined error distribution closer to Gaussian.

CORE OVERBOUNDING

A core bound is a conservative representation of an underlying error distribution. As its distinguishing feature, the core bound decomposes the error distribution into two regions, a core and a tail. Each region is bounded separately. This decomposition offers an advantage in the development of an overbound, especially

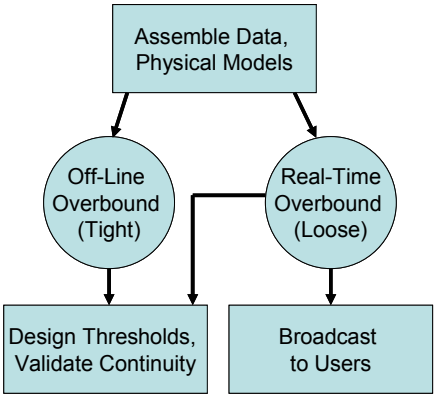

Figure 2. Using More than One Overbound: Real-time operations require an overbound with a simple parameterization. Off-line design, however, benefits from a tighter overbound with a more complex form.

if the far-tail region is strongly non-Gaussian or unknown. In these cases, the distribution extremes may be bounded *implicitly* using a general bound that places few restrictions on the shape of the tails. After removal of the extremes, the core region is bounded *explicitly* using a distribution that approximates the shape of the actual error distribution. In effect, the implicit tails truncate the actual error distribution to permit sharper bounding and improved availability.

Core and Tail Regions

A core bound, G_{cb}, is a cumulative distribution function (CDF) consisting of two fractional CDFs: an explicit core, $\widehat{G}_{o,ex}$, and an implicit tail, $\widehat{G}_{o,im}$.

$$G_{cb}(x) = \widehat{G}_{o,ex}(x) + \widehat{G}_{o,im}(x) \tag{1}$$

The fractional CDFs, indicated by the hat notation, are CDFs multiplied by a scalar between zero and one. Each fractional distribution may be viewed as a conditional CDF weighted by the prior probability of that condition.

The explicit and implicit fractional bounds are associated with the core and the tail, respectively, of the actual error CDF, G_a.

$$G_a(x) = \widehat{G}_{a,core}(x) + \widehat{G}_{a,tail}(x) \tag{2}$$

The fractional CDFs for the actual error distribution are defined in terms of their derivatives, which are discontinuous at a core-tail transition point, T:

$$\widehat{g}_{a,core}(x) = \begin{cases} g_a(x) & |x| \le T \\ 0 & |x| > T \end{cases} \tag{3}$$

and

$$\widehat{g}_{a,tail}(x) = \begin{cases} 0 & |x| \le T \\ g_a(x) & |x| > T \end{cases}. \tag{4}$$

Lower case variable names are used here to indicate probability density functions (PDFs), which are the derivatives of cumulative distribution functions (CDFs) identified by capital variable names.

The remainder of this section provide a mathematical basis for the core bound theorem, which holds that G_{cb} is a conservative bound for G_a if the fractional CDFs, $\widehat{G}_{o,ex}$ and $\widehat{G}_{o,im}$, are conservative representations of $\widehat{G}_{o,core}$ and $\widehat{G}_{o,tail}$.

Generalized CDF Bounding

An important property of an overbound is conservatism in both the range and position domains. Core bounds inherit this conservative property because they belong to the set of overbounds known as CDF overbounds. In order to develop the core bound, it is first necessary to summarize CDF overbounding, a concept first introduced by DeCleene [2] and later extended by Rife *et al.* [3].

DeCleene proved that a set of range-domain overbounds could be combined into a position-domain bound without loss of integrity. This derivation employed a set of "central overbounds," one for each ranging source. Each central overbound was defined as a CDF, G_o, with more cumulative mass in its tails than the actual error CDF, G_a. Mathematically, the central bound is described:

$$\begin{aligned} G_o(x) &\ge G_a(x), \ \forall x < 0 \\ G_o(x) &\le G_a(x), \ \forall x \ge 0 \end{aligned}. \tag{5}$$

The position-domain error that results from a weighted summation of independent range-domain errors is bounded by the convolution of central CDF overbounds for each ranging source. The resulting position-domain bound is also a CDF bound, described by (5), and is thus conservative. This conservatism applies, however, only if the actual error distributions and their overbounds are symmetric, unimodal, and zero mean.

To remove these restrictions on the central CDF bound, Rife *et al.* proposed a modified overbounding approach based on a pair of CDF bounds, G_L and G_R, one to the left and one to the right of the actual error distribution.

$$\begin{aligned} G_L(x) &\ge G_a(x), \ \forall x \\ G_R(x) &\le G_a(x), \ \forall x \end{aligned} \tag{6}$$

Called paired bounding, this approach remains valid even when the overbound or the actual error CDF have asymmetry, multiple modes, or a non-zero mean.

The central, left, and right bounds may be compactly described in a generalized form. This general form replaces the inequalities of the central bound, (5), and the left-right bounds, (6), with a slack (or margin) variable, M. The sign of the margin variable depends on the type of bound.

$$\operatorname{sign}(M) \begin{cases} \in \left\{0, \operatorname{sign}\left(\frac{1}{2} - G_a\right)\right\} & \text{Central Bound} \\ \le 0 & \text{Right Bound} \\ \ge 0 & \text{Left Bound} \end{cases} \tag{7}$$

With this definition of the margin variable, the generalized form of the CDF bound is

$$G_o(x) = G_a(x) + M(x). \tag{8}$$

The following section uses this generalized form of CDF overbounding to validate bound decomposition. It is significant to note that the general overbound definition applies both to standard CDFs and to fractional CDFs which have total mass less than one.

Bound Decomposition

The central notion of core bounding is the decomposition of an error distribution into two parts – a core and a tail. This section establishes that a generalized CDF overbound, as described by (8), results from the summation of conservative fractional bounds.

Any actual error CDF, G_a, may be decomposed into fractional distributions, $\hat{G}_{a,i}$.

$$G_a(x) = \sum_i \hat{G}_{a,i}(x) \tag{9}$$

Applying the definition of the generalized overbound, (8), a conservative bound, $\hat{G}_{o,i}$, may be defined for each $\hat{G}_{a,i}$, with equivalent total mass.

$$\hat{G}_{o,i}(x) = \hat{G}_{a,i}(x) + \hat{M}_i(x) \tag{10}$$

Together, these fractional overbounds sum to form a CDF overbound with a total probability mass equal to one.

$$G_o(x) = \sum_i \hat{G}_{o,i}(x) \tag{11}$$

If the margin variables for each fractional bound are consistent—that is they are all either central, left, or right bounded according to (7)—then the sum of the margin variables has the same sign as each individual M_i.

$$\operatorname{sign}\left(\sum_i \hat{M}_i\right) = \operatorname{sign}(M_i) \quad \forall i \tag{12}$$

Consistency of the margin variables, according to (12), implies that the overbound recovered by summing the fractional bounds is a generalized CDF overbound according to (8).

$$G_o(x) = G_a(x) + \sum_i \hat{M}_i(x) \tag{13}$$

Thus the core bound, (1), is a conservative overbound in the range and position domains if it is based on conservative fractional CDF bounds.

Implicit and Explicit Bounding

In ensuring that $\hat{G}_{o,ex}$ and $\hat{G}_{o,im}$ are conservative, the explicit core bound and the implicit tail bound are treated in different fashions. It is assumed that empirical data, physical models, or statistical assumptions are available to define the form of the error distribution precisely in the core. The overbound for the core region may thus be defined as an explicit function bounding the worst-case CDF for the actual error distribution.

$$\hat{G}_{o,ex}(x) = \int_{-\infty}^{x} \hat{g}_{a,core}(x) + \hat{M}_{ex}(x). \tag{14}$$

By contrast, no assumptions are made about the shape of the tail region. Thus, the actual tail error, $\hat{g}_{a,tail}$, is left as an arbitrary function and is bounded implicitly:

$$\hat{G}_{o,im}(x) = \int_{-\infty}^{x} \hat{g}_{a,tail}(x) + \hat{M}_{im}(x). \tag{15}$$

Only the total probability of a tail error, P_t, is defined:

$$P_t = \int_{-\infty}^{\infty} \hat{g}_{a,tail}\, dx. \tag{16}$$

This total probability defines the worst case arbitrary tail. For central bounding, the tail probability could, in an extreme case, extend uniformly toward infinity. In the right and left bounding cases, which are not constrained by unimodality, the worst case tails are spikes (delta functions) located at +∞ (right bounding) or -∞ (left bound). The resulting worst-case values of $\hat{G}_{o,im}$ are summarized by (17).

$$\hat{G}_{o,im,worst}(x) = \begin{cases} \frac{1}{2}P_t & \text{Central Bound} \\ 0 & \text{Right Bound} \\ P_t & \text{Left Bound} \end{cases} \tag{17}$$

The core bounding theorem can now be stated formally. The theorem holds that a core bound, (1), is conservative in both the range and position domains if the explicit and implicit overbounds satisfy condition (14) and (15) with consistent slack margins, as defined by (12).

Range Domain to Position Domain Conversion

The core bound theorem guarantees conservatism for a position-domain bound synthesized from range-domain core bounds. A corollary result, proved in the Appendix and summarized here, states that the synthesized position-domain bound is, in fact, also a core bound. The position-domain core bound, H_p, consists of an explicit part and and implicit part.

$$H_p(z) = \hat{H}_{p,ex}(z) + \hat{H}_{p,im}(z). \quad (18)$$

The explicit part of the position-domain bound is derived by convolving the explicit CDF core bound for one error source with the weighted PDF core bounds for all other independent error sources. The weighting factors are labeled S_i.

$$\hat{H}_{ex}(z) = \hat{G}_{core,1}(S_1 x_1) * \hat{g}_{core,2}(S_2 x_1) * \ldots * \hat{g}_{core,3}(S_N x_N), \quad (19)$$

The worst case bound for the implicit term is similar in form to (17).

$$\hat{H}_{p,im,worst}(x) = \begin{cases} \frac{1}{2} P_{im} & \text{Central Bound} \\ 0 & \text{Right Bound} \\ P_{im} & \text{Left Bound} \end{cases}. \quad (20)$$

The magnitude of the total implicit probability, P_{im}, is bounded by the sum of the implicit tail probabilities, P_t, for each of the independent error sources.

$$P_{im} \le \sum_{n=1}^{N} P_{t,n} \quad (21)$$

CORE BOUNDING IMPLEMENTATIONS

This section proposes specific implementations of core bounding that are based on the general mathematical framework of the previous section. Two types of core bound are considered: the Gaussian Core (GC) bound and the Gaussian Core with Gaussian Sidelobes (GCGS) bound. Both variations incorporate the Gaussian distribution, since this function behaves in a favorable manner during convolution. (i.e., The output of the convolution of two Gaussian functions is Gaussian.)

The Gaussian Core (GC) Bound

The GC bound is among the simplest possible implementations of core bounding. The GC bound is advantageous for LAAS, furthermore, because no changes to the VHF broadcast message are required to implement GC bounding. In GC bounding, the explicit core distribution is Gaussian.

$$\hat{G}_{GC,core}(x) = (1 - P_t)\psi(x;\sigma_n) \quad (22)$$

Here ψ represents the CDF of the Gaussian distribution:

$$\psi(x;\sigma) = \int_{-\infty}^{x/\sigma} \frac{1}{\sqrt{2\pi}} \exp(-x^2/2)dx. \quad (23)$$

The GC bound is characterized by two parameters: a width parameter, σ_n, and the implicit tail probability, P_t Figure 3 illustrates how the two parameters determine the shape of the GC distribution in relation to a conventional Gaussian. Increasing σ_n widens the bound; increasing P_t compresses the bound in the vertical direction.

The position-domain formulation for the central GC bound after N convolutions is found by combining (18) with (22):

$$H_{GC}(z) = (1 - P_t)^N \psi(z;\sigma_p) + \tfrac{1}{2} P_{im}. \quad (24)$$

Here the position domain width, σ_p, is the root of the sum of the squares of the weighted range-domain values, σ_n:

$$\sigma_p = \sqrt{\sum_n S_n^2 \sigma_n^2}. \quad (25)$$

Assuming the same value of implicit tail probability, P_t, for all range-domain distributions, (21) implies

$$P_{im} \le N P_t. \quad (26)$$

The Gaussian Core Gaussian Sidelobes (GCGS) Bound

A second useful class of core bound is the Gaussian Core with Gaussian Sidelobes (GCGS) bound. Sidelobes augment the explicit core bound with heavier tails. In this fashion, sidelobes permit tighter bounding for off-line analysis of LAAS error distributions. GCGS bounds may employ one or more sidelobes in each tail. Figure 4 illustrates the simplest case of GCGS bounding, using only one sidelobe per tail. In the figure, GCGS bounds, represented as PDFs, are compared with a reference Gaussian distribution (dashed line).

The set of parameters that define a GCGS distribution are a superset of the GC parameters. GCGS parameters include distribution width, σ_n, tail probability, P_t, the sidelobe shift parameter, u, and the sidelobe scaling

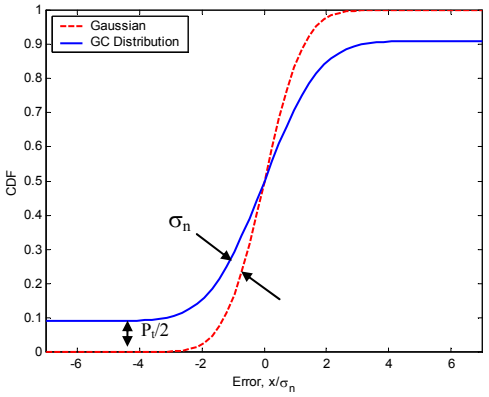

Figure 3. Cumulative Distribution of GC Bound

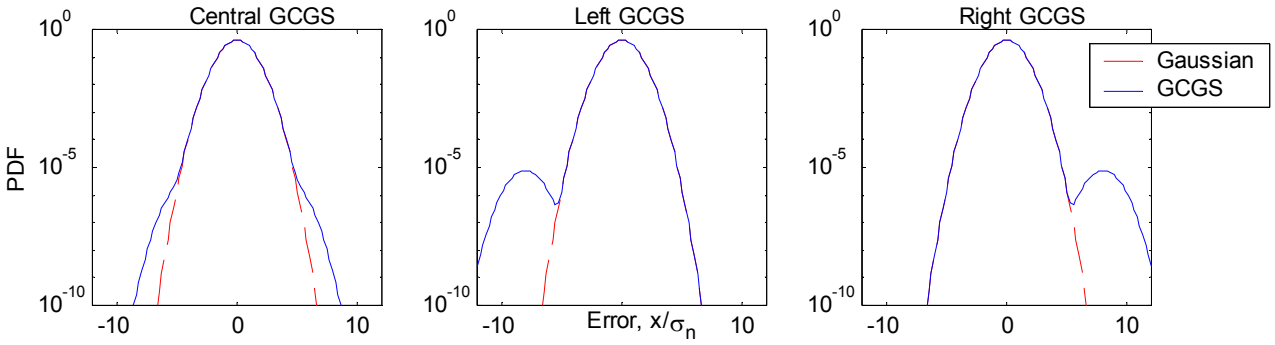

Figure 4. PDF for Explicit Core of GCGS Bound

parameter, ε. The sidelobe shift parameter describes the distance, normalized by σ_n, between the center of the core Gaussian and its sidelobes. The sidelobe scaling parameter indicates the conditional probability associated with each sidelobe. For the case of a GCGS distribution with multiple sidelobes, the parameters u and ε are vectors.

All lobes of the GCGS share the same width parameter, σ_n. The GCGS distribution can thus be expressed as the convolution of a single Gaussian distribution with a set of delta functions:

$$\hat{G}_{GCGS,core}(x) = \psi(x;\sigma_n) * d(x), \tag{27}$$

where

$$d(x) = \alpha\delta(x) + \sum_i \varepsilon_i \delta(x + \sigma_n u_i). \tag{28}$$

The sidelobe shift parameter, u, has all elements positive for a left bound, negative for a right bound, and symmetrically distributed for a central bound. The main lobe weighting parameter, α, is defined by the CDF total mass constraint:

$$\alpha = 1 - \sum_i \varepsilon_i - P_t. \tag{29}$$

The central GCGS may have only one maximum, as required by the unimodality constraint of [2]. This constraint imposes a limit on sidelobe placement. The constraint on u may be derived as a function of the relative weights of the sidelobes and the main lobe. Figure 5 plots the relationship between α/ε and u for the one-lobe case. With left and right GCGS bounds, in pair bounding, there is no unimodality constraint. Thus sidelobe locations are unconstrained and multiple peaks are allowed.

The convolution of multiple range-domain GCGS bounds results in a GCGS bound in the position-domain. By invoking the commutative property of convolved functions, it is straightforward to show that the convolution of GCGS core distributions, described by (27), results in an explicit position-domain bound that consists of a Gaussian distribution convolved with a delta function set.

$$\hat{H}_{GCGS,ex}(z) = \psi(z;\sigma_p) * {}^N d(z) \tag{30}$$

In the position-domain, the delta function set, ${}^N d$, is a convolution of the delta groupings, d, associated with each of N weighted errors.

$${}^N d(z) = d_1(S_1 x_1) * d_2(S_2 x_2) * \ldots * d_3(S_3 x_N) \tag{31}$$

The complete form for the position-domain GCGS bound also includes an implicit term, which is described by (20) and (21).

LAAS APPLICATIONS OF CORE BOUNDING

Heavy-tailed error distributions motivate the application of core bounding concepts to LAAS. Two significant challenges associated with the bounding of heavy-tail errors are (1) formal overbound validation and (2) inflation-factor mitigation. The first challenge occurs because desirable overbounding functions, like the traditional Gaussian distribution, cannot formally overbound the extremes of a heavy-tailed error distribution. The implicit tails of a core bound address this problem, as they truncate the most extreme errors from the explicitly bounded core region. Although

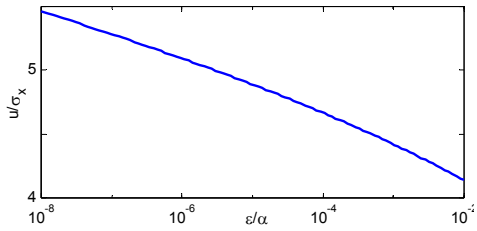

Figure 5. Central GCGS Unimodality Limit

truncation does not remove the non-Gaussian behavior from the core distribution, remaining heavy tails shrink when the core distribution is combined with other system errors. Since distributions tend toward Gaussian during convolution, the magnitude of the σ inflation-factor for Gaussian bounding is reduced in this process. Harnessing this potential for inflation-factor mitigation presents a second challenge for overbounding a heavy-tailed distribution. The sidelobes of the GCGS bound address this challenge by providing a sharp, non-Gaussian core bound without sacrificing the analytical tractability of the Gaussian CDF. Thus the GCGS provides a means to establish heavy-tail mitigation in a rigorous analytical fashion, without resorting to Monte Carlo simulation.

Validation of an Overbound with Slowly Decaying Tails

The first overbounding challenge for LAAS focuses on validation of overbounds for heavy-tail distributions. The current specifications for LAAS assume central Gaussian bounds for all errors. Gaussian overbounds are not trivial to establish, however, for the heavy-tailed case.

A central overbound must have more tail mass than the actual error distribution, everywhere, as expressed by (5). This formal condition extends to infinity. If the heavy-tails of the actual error distribution extend to infinity, it may be mathematically impossible to define an inflated-sigma Gaussian that satisfies condition (5). Thus, Gaussian overbounds can only be defined for heavy-tail error distributions that are in some way truncated.

The Normal Inverse Gaussian (NIG) represents one example of probability distribution for which no Gaussian overbound exists. The NIG may, however, approximate the heavy tails of the ground-station error distribution for LAAS, as discussed by Braff [4]. The following equation, with $M = 1$, describes the unit-variance NIG:

$$f_{NIG}(x) = \frac{M\delta_0^2 \exp(M\delta_0^2)}{\pi\sqrt{x^2+\delta_0^2}} K_1\left(M\delta_0\sqrt{x^2+\delta_0^2}\right). \qquad (32)$$

For $M > 1$, $f_{NIG}(x)$ gives the average of M identical NIG distributions. K_1 is a modified Bessel function of the second kind, degree one. The shape parameter, δ_0, determines the weight of the NIG tail. Braff recommends setting $\delta_0 = 0.65$ to conservatively represent the tails of the LAAS ground station error distribution.

In contrast with finite-σ Gaussian distributions, which cannot bound the NIG to infinity, a core bound can provide a formal overbound for the NIG. To illustrate this application of the core bound, Figure 6 plots the CDFs for the GC distribution and the NIG distribution on Gaussian axes. (On these axes, a Gaussian CDF appears as a straight line with slope inversely proportional to σ.)

The arcing NIG tails undermine conventional Gaussian overbounding, but they fall cleanly inside the implicit tails of the GC overbound. In effect, the implicit tail clips the actual error distribution, allowing for the establishment of a formal Gaussian bound for the truncated core distribution.

LAAS specifications limit the maximum size of the implicit tail. The most stringent constraint on the tail probability, P_t, is imposed by the integrity allotment for fault-free missed detections, P_{ffmd}. This P_{ffmd} allocation depends on the number of reference receivers (M) and the landing system category [5], as summarized by the following table.

Table 1. Probability of Fault Free Missed Detection

	P_{ffmd} (M=3)	P_{ffmd} (M=4)	Exposure Time
Category I	$6.25\text{x}10^{-9}$	$5.0\text{x}10^{-9}$	150 s
Category III	$3.13\text{x}10^{-11}$	$2.5\text{x}10^{-11}$	15 s

Fault free missed detections occur when the actual position-domain error exceeds the PL. This situation may occur for all errors in the implicit tail. A fraction, ζ, of P_{ffmd} must therefore be diverted to absorb the implicit tail probability. This fraction must be relatively small to ensure minimal impact on the PL equations, which are applied to the explicit core distribution. This probability fraction, moreover, must be split among the various ranging sources that sum to form the implicit position-domain distribution, as described by (21). For all ranging sources, the allowable probability mass in each implicit tails, ½P_t, is thus determined by ζ and the total number of ranging sources, N:

$$\frac{P_t}{2} = \frac{\zeta P_{ffmd}}{2N}. \qquad (33)$$

For this paper, the quantity ζ/2N was set to 1/100.

The probability mass allocated to each implicit tail is a very small number. Even this small implicit probability, however, permits formal validation of Gaussian Core bounding for a heavy-tailed distribution.

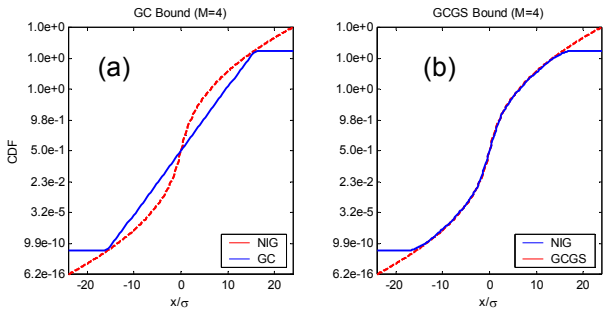

Figure 6. Bounding a Normal Inverse Gaussian Error Distribution with (a) Gaussian Core and (b) Gaussian Core with Gaussian Sidelobe distributions

Sharpening Overbounds to Improve Availability

The GC bound features two characteristics that are important for LAAS. First, the integrity of the GC bound can be formally validated for heavy-tailed error distributions. Second, implementation of the GC bound, which requires only the broadcast of a single Gaussian σ parameter, is fully compatible with the current LAAS architecture. The principal disadvantage of the GC bound is overconservatism. As Figure 6(a) clearly shows, the GC bound is not a sharp bound for a heavy tail error distribution such as the NIG. As a consequence, the GC bound may incur a significant system availability penalty.

The high inflation factors associated with heavy tails may be mitigated by considering convolution smoothing, which results when multiple error sources are added together. By the Central Limit Theorem, sums of random variables tend toward a Gaussian distribution. As a corollary, the convolution of non-Gaussian distributions tends to mitigate heavy tails. Taking advantage of convolution mitigation requires accurate representation of heavy tails through the use of a sharp, non-Gaussian overbound, such as the raw NIG distribution or a GCGS like that shown in Figure 6(b).

This section performs analysis using a GCGS bound, rather than the underlying NIG error representation, because of the convenient analytical properties of the GCGS. No closed-form analytic expression describes the convolution of an NIG with a Guassian or the convolution of multiple NIG distributions with differing sigmas. By contrast, the convolution of multiple GCGS results in a GCGS output, as described by (30). The convolution of a GCGS with a Gaussian also results in a GCGS output. Specifically, if a conventional Gaussian of standard deviation σ_g is convolved with a GCGS of the form (27), with width parameter σ_n, the resulting distribution is

$$\hat{G}_{GCGS,core}\left(\frac{x}{\sigma_{tot}}\right) = \psi\left(\frac{x}{\sigma_{tot}};1\right) * d\left(\frac{x}{\sigma_{tot}}\right) \tag{34}$$

with

$$d\left(\frac{x}{\sigma_{tot}}\right) = \alpha\delta\left(\frac{x}{\sigma_{tot}}\right) + \sum_i \varepsilon_i \delta\left(\frac{x}{\sigma_{tot}} + \frac{u_i}{\phi}\right). \tag{35}$$

The output of the Gauss-GCGS convolution is thus a GCGS distribution identical to the original distribution, but with the shift parameter, u, scaled by the inverse of the Gauss Ratio, ϕ.

$$\phi = \frac{\sigma_{tot}}{\sigma_n} = \frac{\sqrt{\sigma_n^2 + \sigma_g^2}}{\sigma_n}. \tag{36}$$

The Gauss ratio describes the relative magnitude of the final and initial sigmas. This scaling property of ϕ in (35) greatly simplifies the analysis of the convolution of a GCGS with other Gaussian error sources. Plotting the output GCGS as a function of the Gauss ratio, ϕ, indicates the degree of inflation-factor mitigation achieved for a particular Gauss ratio. Figure 7 illustrates this sensitivity. The initial GCGS distribution ($\phi = 1$) is a close fit to the NIG distribution with $\delta_0 = 0.65$ (dashed red line). Contours for higher values of ϕ are also shown. Convergence of the GCGS distribution toward the Gaussian limit (cyan line) is rapid as ϕ increases. Thus the inflation factor decreases toward one as ϕ increases. This effect is also clear in Figure 8, which plots inflation factor against Gauss ratio.

As a system, LAAS accounts for ground station errors, ionospheric errors, tropospheric errors and airborne noise. LAAS specifications dictate the relative magnitude of these error terms, which are assumed to be Gaussian overbounded [5]. Based on these specified error levels, the LAAS ground facility can estimate the Gauss ratio which applies for each individual satellite. The total error is dominated by the ground noise, airborne noise and ionosphere gradient terms. According to the specifications, the broadcast standard deviation for the ground error can be modeled by the Ground Accuracy Designator C (GAD-C) curve. An airborne noise term with σ defined by the Airborne Accuracy Designator B (AAD-B) curve can also be assumed. The ionospheric gradient standard deviation is

$$\sigma_{iono} = OF \cdot \sigma_{vig}\left(X + 2\tau V_{air}\right). \tag{37}$$

This relationship incorporates two system parameters: the vertical ionospheric gradient parameter, σ_{vig}, taken to be 4 mm/km, and the time constant, τ, of the carrier-smoothing filter, taken to be 100 s. The obliquity factor, *OF*, is a function of satellite elevation angle, known by

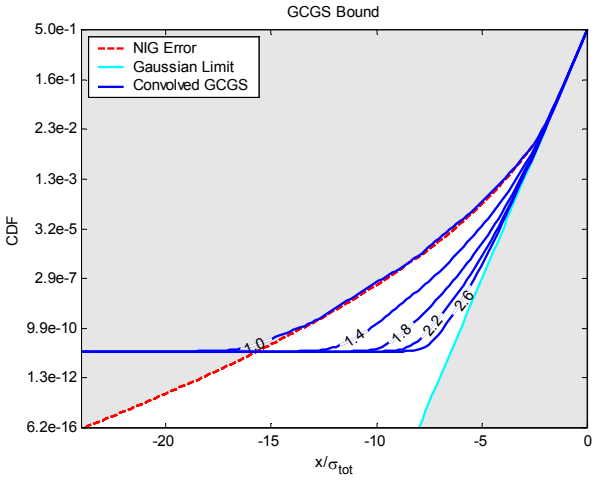

Figure 7. Heavy Tail Mitigation with Increasing Gauss Ratio

the ground station. Although the remaining parameters are not known by the ground station, the aircraft speed, V_{air}, can be assumed to be about 70 m/s. The aircraft to receiver distance, X, provides no benefit, however, as this distance goes to zero during final approach.

For the purposes of computing a Gauss ratio, ϕ, the total Category I error, σ_{tot}, is the root-sum-squared of the ground, airborne and ionospheric σ-values. The total error allowed for Category III operations is significantly smaller, however, because the vertical alert limit is reduced from 10 m to 5.3 m. To represent these stringent requirements, the total error for Category III was computed using the standard GAD-C curve for ground error along with an airborne error at half the AAD-B level and no ionospheric error. In [7], Shively showed that this level of total error achieves the Category III availability requirements. Figure 9 plots curves for the total error for both the Category I and Category III error representations.

In computing the Gauss ratio, ϕ, the σ_n for the non-Gaussian error term is modeled as an NIG with $\delta_0 = 0.65$. This is based on data from the Federal Aviation Administration Technical Center (FAATC) [1] that describes the performance of an integrated multipath limiting antenna (IMLA) [6]. The IMLA σ is only a fraction of the assumed GAD-C curve, as shown in Figure 9. The Gauss ratios obtained by comparing the FAATC data to the Category I and Category III total error curves (for either three or four operating reference receivers) is illustrated by the lower graph of Figure 9.

The following table shows the inflation factors, based on Figure 8, for the worst ϕ value at or below 35°. The ground station switches between the IMLA and the high zenith antenna (HZA) at this 35° transition point.

Table 2. PL Inflation Factor after Mitigation

	M = 4	M = 3
Category I	1.01	1.01
Category III	1.06	1.14
Category III, 30% inflation	1.17	1.33

The inflation factors in Table 2 are derived assuming the off-line GCGS bound is replaced by a real-time GC bound for broadcast to the airborne user. It is important to note that the inflation factors of the table apply to the entire error distribution (the PL), and not just to the ground error, σ.

Table 2 considers six cases, including three and four reference receiver cases for three conditions: Category I operation, Category III operation, and inflated Category III operation. The last case represents inflation that may be necessary to compensate for uncertainty in the actual error distribution. For the inflated case, all Gauss ratios from Figure 9 were reduced by an additional factor of 1.3.

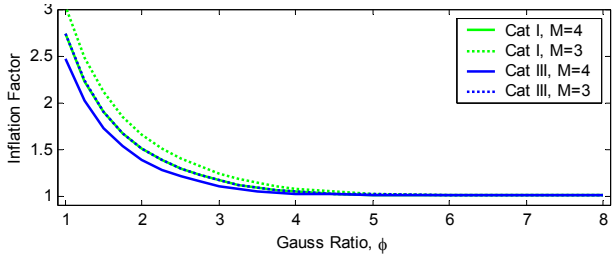

Figure 8. Heavy-Tail Inflation Factor vs. Gauss Ratio. Error distributions are first fitted with a GCGS, to assess convolution mitigation, then fitted with a GC bound, to enable broadcast.

In all cases, convolution mitigation improves the level of PL inflation as compared with the no-mitigation ($\phi = 1$) case. For comparison, Table 3, below, lists the PL inflation factors for the case of no convolution mitigation.

Table 3. PL Inflation Factor for GC Bound

	M = 4	M = 3
Category I	1.07	1.10
Category III	1.19	1.29
Category III, 30% inflation	1.31	1.46

In the Category I case, convolution mitigation makes a significant impact. Without mitigation, heavy tails inflate the PL for Category I operations by at least 7%. This penalty is high enough to reduce system availability below the desired 99.9% level. When convolution mitigation is considered, however, the PL inflation factor is 1.01. To maintain integrity given heavy-tailed errors, the ground station need merely decrease its broadcast alert limit by 1%, an action which minimally impacts system availability.

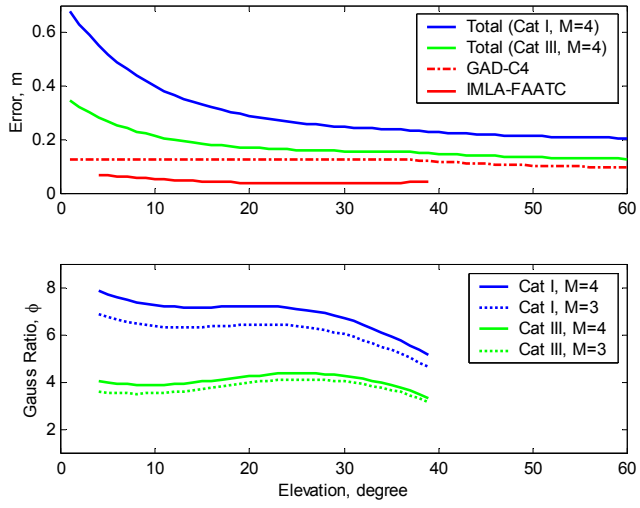

Figure 9. Deriving Gauss Ratio for LAAS: The error standard deviation, σ, and Gauss ratio, ϕ, plotted against satellite elevation angle

In contrast with the Category I case, heavy tails have a severe impact on Category III operations. Even after convolution mitigation is applied, the PL inflation factor may be as high as 1.14, or higher still if additional inflation is applied for uncertainty in the ground error distribution. The table therefore motivates further mitigation considerations for Category III operations.

Further Tail Mitigation Using Excess H1 Probability

Aside from convolution mitigation, another strategy for reducing the inflation factor for the explicit core is to transfer a greater amount of probability into the implicit tails. The fault free (H0) hypothesis, however, places a severe restriction on the implicit tail probability, P_t, as described in Table 1. The H0 limit may be circumvented if the largest ground station errors are defined as faults under the single receiver failure (H1) hypothesis. This concept of defining the extremes of the ground station error distribution as H1 faults was first introduced by Brenner, who studied the concept using Monte Carlo methods [8].

The notion of redefining extreme tail events as receiver faults has a direct impact on core bounding. Any excess margin in the allowed receiver fault rate can be applied to increase P_t. Excess margin may exist in the differing LAAS H1 specifications for integrity and continuity. The integrity requirement allows a higher fault rate (1E-5 / 150 s) than the continuity requirement (2.5E-6 / 150 s for Category I and 6.3E-8 / 150 s for Category III) on a per receiver basis. If the LAAS receiver hardware is built to match the tighter continuity constraint, then the margin in the H1 integrity budget is approximately 7.5E-6 per independent sample. This probability allotment must be distributed over the 2 tails of N error sources, according to (33).

The H1 hypothesis requires that excess margin be applied to the tails of the single-receiver error distribution (M =1). Figure 10 illustrates the significant increase in P_t when bounding a non-Gaussian distribution with H1 margin (blue) rather than H0 margin (green). The illustration

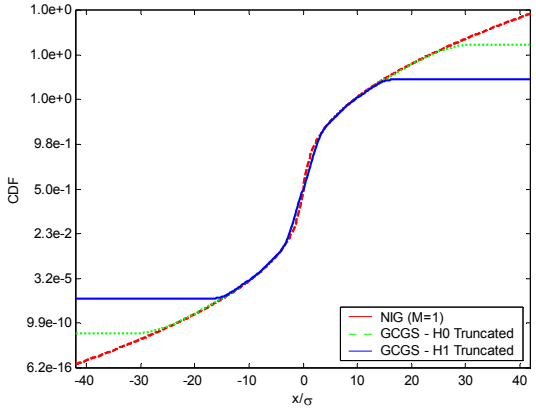

Figure 10. Using H1 Receiver Failure Budget for Implicit Tail

exaggerates the benefit somewhat, however, because the H0 margin can be applied directly to an already mitigated multi-receiver distribution (M = 3 or M = 4) to achieve improved results. The competition between these effects means that excess H1 bounding has a greater impact when the Gauss ratio, ϕ, is low. As Table 4 illustrates, the H1 bounding approach provides only modest benefits to Category III LAAS, and negligible benefits to Category I.

Table 4. PL Inflation Factor using H1 Margin

	M = 4	M = 3
Category I	1.01	1.01
Category III	1.04	1.09
Category III, 30% inflation	1.11	1.21

Broadcast of Non-Gaussian Bounds

A final mitigation strategy considers convolution of errors over multiple satellites. This strategy introduces a challenge, however, because the ground station lacks knowledge about the set of satellites in view of the airborne user. To address this lack of information, the ground station could compute statistics for all subsets of visible satellites using a position-domain monitor [9]. An alternative possibility, that of broadcasting a non-Gaussian error, is enabled by core bounding. This strategy avoids penalizing airborne users with visible satellite geometries more favorable than the worst case geometry simulated by a ground-based monitor. On the other hand, transmitting a non-Gaussian error distribution does require changing the LAAS broadcast message and PL computation.

Because these are undesirable changes from an operational viewpoint, a thorough availability analysis is not conducted in this paper. Nonetheless, it is worth noting that the broadcast of a single-lobe GCGS distribution would offer advantages in further mitigating heavy-tail effects.

The GCGS is a bound well suited to real-time operations. The single-lobe GCGS uses only two extra parameters in addition to distribution σ : a sidelobe shift, u, and a sidelobe scaling, ε, both described by (28). The GCGS, moreover, results in a straightforward computation of the PL equation based on an analytical form for the position domain error distribution. This position domain error distribution would be a GCGS distribution with multiple lobes, obtained by convolving a delta function set with a Gaussian distribution.

The GCGS broadcast would reduce inflation in two ways. An immediate benefit arises because the single-lobe GCGS bound is sharper than the no-lobe GC bound, given a fixed tail probability, P_t. The inset of Figure 11 illustrates the relative sharpness of the GC (red) and GCGS (blue) bounds in the vicinity of the P_{ffmd} limit. PL inflation factor is reduced from 1.14, as shown in Table 4,

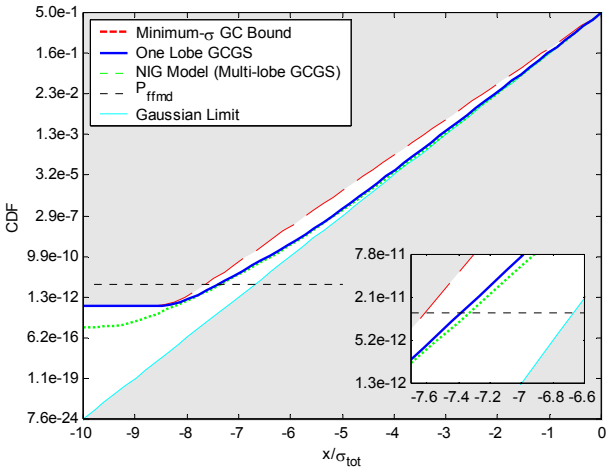

Figure 11. Impact of Broadcasting a non-Gaussian Overbound

to 1.11. The second, and more significant, benefit of non-Gaussian bounding arises from convolution mitigation. When no single satellite dominates the total error, convolution mitigation drives the error distribution tails toward the Gaussian limit (cyan). Equivalently, the PL inflation factor is pushed closer to 1.0.

Other Core Bounding Applications

This section discussed methods to validate integrity overbounds for heavy-tailed distributions and to tighten those bounds to improve availability. Other applications for core bounding techniques include (1) continuity analysis and (2) the analysis of airborne and ionosphere error distributions with heavy-tails. In the first case, core bounding methods could be applied to assess the impact of heavy tails on LAAS B-Values and to evaluate the continuity risk posed by false alarms under the H1 hypothesis. In the second case, that of heavy-tailed airborne and ionosphere noise, a core bounding analysis would indicate the extent to which heavy-tail mitigation suffers when all error distributions are non-Gaussian.

SUMMARY

A new tool, called core bounding, was introduced to establish formal overbounds for heavy-tailed error distributions. The tool decomposes the error distribution into tail and core regions to provide a sharper bound. Sharper bounds result in better availability for a safety-of-life system such as LAAS. Two types of core bound were introduced for application to LAAS: the Gaussian Core (GC) and the Gaussian Core with Gaussian Sidelobes (GCGS) distributions. The GC distribution is advantageous for real-time operations, as it enables validation of heavy-tail bounding with no changes to the LAAS broadcast message. The GCGS distribution, by contrast, proves a useful tool for off-line analysis in which heavy tails are mitigated by convolution with other error sources. A convolution mitigation analysis, using the GCGS, indicates that heavy tails have only a minor impact on Category I LAAS availability. Heavy tails have a significant impact on Category III, however. If the Category III alert limit remains at 5.3 m, availability limits may require that a non-Gaussian error distribution be broadcast to airborne users.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the Federal Aviation Administration Satellite Navigation LAAS Program Office (AND-710) for supporting this research. The opinions discussed here are those of the authors and do not necessarily represent those of the FAA or other affiliated agencies.

REFERENCES

[1] J. Warburton, *LGF Sigma Pseudorange Ground Establishment, Overbounding and Monitoring: Support Data*. Presented by the FAA William J. Hughes Technical Center to RTCA SC-159, December 2002.

[2] B. DeCleene, *Defining Pseudorange Integrity – Overbounding*, Proceedings of the Institute of Navigation's ION-GPS 2000, pp. 1916-1924.

[3] J. Rife, S. Pullen, B. Pervan, and P. Enge, *Paired Overbounding and Application to GPS Augmentation*, IEEE Position, Location and Navigation Symposium, pp. 439-446, 2004.

[4] R. Braff and C.A. Shively, *A Method of Overbounding Ground Based Augmentation System (GBAS) Heavy Tail Error Distributions*, Manuscript accepted by The Journal of Navigation, 2004.

[5] RTCA Inc., *Minimum Aviation System Performance Standards for the Local Area Augmentation System (LAAS)*, RTCA/DO-245, 1998.

[6] D.B. Thornberg, D.S. Thornberg, M.F. DiBenedetto, M.S. Braasch, F. van Graas, and C. Bartone, *The LAAS Integrated Multipath Limiting Antenna (IMLA)*, Proceedings of the Institute of Navigation's ION-GPS 2002, pp. 1916-1924.

[7] C.A. Shively and T.T. Hsiao. *Availability Enhancements for Cat IIIB LAAS*, NAVIGATION, Journal of the Institute of Navigation, Vol. 51, No. 1, 2004.

[8] M. Brenner. *Over-Bounding*. Presented by Honeywell to the FAA, TIM#10, April 2004.

[9] J. Lee, S. Pullen, G. Xie, and P. Enge, *LAAS Position Domain Monitor Analysis and Failure-Test Verification*, AIAA 21st International Communications Satellite Systems Conference and Exhibit, 2003.

APPENDIX

Range-to-Position Domain Conversion for a Core Bound

The position-domain bound associated with a set of core-bounded ranging errors is itself a core bound. This result, previously introduced as equations (18) through (21), is derived here.

The position-domain error, z, is the sum over all satellites of error contributions, x, scaled by a sensitivity factor, S_n.

$$y_n = S_n x_n , \quad z = \sum_{n=1}^{N} y_n \tag{38}$$

The scaling factor reflects projective geometry and the relative accuracy of measurements. In LAAS, these error sources are considered to be independent. As such, the distribution of the position error, z, is a convolution of the individual ranging source error distributions, x. Similarly, a position-domain CDF overbound may be derived by convolving overbounds for each scaled ranging error [2,3].

In this paper, the notation * denotes convolution. The CDF resulting from the convolution of two PDFs, $f(y_1)$ and $g(y_2)$, is equivalent to the convolution of the CDF $F(y_1)$ with the PDF $g(y_2)$. Consistent with earlier sections of this paper, capital letters indicated CDFs and lower case letters indicate PDFs.

$$\begin{aligned} H(z) &= F(y_1) * g(y_2) = \int_{-\infty}^{\infty} F(z - y_2)\, g(y_2)\, dy_2 \\ &= \int_{-\infty}^{z} \left(f(y_1) * g(y_2) \right) dz = \int_{-\infty}^{z} h(z)\, dz \end{aligned} \tag{39}$$

As described by equations, (18)-(21), a position-domain core bound results from the convolution of a set of range-domain core bounds. This result can be derived by induction, starting with the 2-satellite case.

Consider a pair of scaled error distributions for two ranging sources, f and g. The distribution for the sum of these errors is, ${}^2H(z)$, where the leading superscript indicates the number of ranging source distributions convolved. If f and g are core bounded, ${}^2H(z)$ is:

$$\begin{aligned} {}^2H(z) = &\int_{-\infty}^{\infty} \left[\hat{F}_{core}(z - y_2) \hat{g}_{core}(y_2) \right] dy_2 \\ &+ \int_{-\infty}^{\infty} \left[\hat{F}_{core}(z - y_2) \hat{g}_{tail}(y_2) \right] dy_2 \\ &+ \int_{-\infty}^{\infty} \left[\hat{F}_{tail}(z - y_2) g(y_2) \right] dy_2 \end{aligned} \tag{40}$$

The first term of (40) is the explicit core of ${}^2H$.

$${}^2\hat{H}_{ex}(z) = \int_{-\infty}^{\infty} \hat{F}_{core}(z - y_2) \hat{g}_{core}(y_2)\, dy_2 \tag{41}$$

The remaining terms of (40) form the implicit part of ${}^2H$.

$$\begin{aligned} {}^2\hat{H}_{im}(z) = &\int_{-\infty}^{\infty} \left[\hat{F}_{core}(z - y_2) \hat{g}_{tail}(y_2) \right] dy_2 \\ &+ \int_{-\infty}^{\infty} \left[\hat{F}_{tail}(z - y_2) g(y_2) \right] dy_2 \end{aligned} \tag{42}$$

After convolution, the amount of probability in the implicit distribution is ${}^2P_{im}$. This function is related to the tail probabilities of the individual range domain bounds, as defined by (16). The implicit bounding function (42) can be evaluated at $+\infty$ to compute the total implicit probability, ${}^2P_{im}$. This approach decouples the integral of F and g as follows:

$${}^2P_{im} = \hat{F}_{core}(\infty) \int_{-\infty}^{\infty} \hat{g}_{tail}\ dy + \hat{F}_{tail}(\infty) \int_{-\infty}^{\infty} g\ dy . \tag{43}$$

Consequently,

$$\begin{aligned} {}^2P_{im} &= P_{core,f} P_{tail,g} + P_{tail,f} \\ &= (1 - P_{tail,f}) P_{tail,g} + P_{tail,f} \\ &= P_{tail,f} + P_{tail,g} - P_{tail,f} P_{tail,g} \end{aligned} \tag{44}$$

Because all probabilities must be positive,

$${}^2P_{im} < P_{tail,f} + P_{tail,g} . \tag{45}$$

In summary, for two satellites, the explicit core bound distribution is formed by convolving the core bounds for the individual error distributions, according to (41). The amount of probability in the implicit distribution is expressed by (44) and bounded by (45). The 2-satellite bound is, itself, a core bound of the form (1).

The N-satellite equations are derived from the 2-satellite equations by induction. Convolving the 2-satellite bound, ${}^2H$, with a subsequent range-domain core bound generates a core bound, ${}^3H$. Repeating this process results in the N-satellite equations, (19) and (21), which are direct analogs to the 2-satellite equations, (41) and (45).