Paired Overbounding and Application to GPS Augmentation

Jason Rife*, Sam Pullen*, Boris Pervan† and Per Enge*
*Stanford University; Palo Alto, CA 94305
†Illinois Institute of Technology; Chicago, IL 60616

Abstract—The relationship between range-domain and position-domain errors remains an open issue for GPS augmentation programs, such as the Federal Aviation Administration's Local Area Augmentation System (LAAS). This paper introduces a theorem that guarantees a conservative error bound (overbound) in the position domain given similarly conservative overbounds for broadcast pseudorange statistics. This paired overbound theorem requires that a cumulative distribution function (CDF) be constructed to bound both sides of the range-domain error distribution. The paired overbound theorem holds for arbitrary error distributions, even those that are non-zero mean, asymmetric or multimodal. Two applications of the paired overbound theorem to GPS augmentation are also discussed. First, the theorem is employed to construct an inflation factor for a non-zero mean Gaussian distribution; in the context of a simulation of worst-case satellite geometries for 10 locations in the United States and Europe, the required inflation factor for broadcast sigma is only 1.18, even for biases as large as 10 cm for each satellite. Second, the theorem is applied to bound a bimodal multipath model tightly; the result shaves more than 40% off the previously established inflation factor derived through a more overly conservative analysis.

I. INTRODUCTION

The FAA has sponsored two programs to supplement the Global Positioning System (GPS) to provide increased accuracy and integrity for precision landing applications. The Wide Area Augmentation System (WAAS) collects data across the coterminous United States and broadcasts corrections from satellites in geosynchronous orbit. The Local Area Augmentation System (LAAS) employs local arrays of GPS receivers located near individual airports to broadcast short-range differential corrections over a VHF radio link. In general, LAAS corrections provide higher accuracy and integrity than WAAS corrections, but apply to a much smaller geographic area.

In both classes of differential GPS augmentation, the broadcast correction messages incorporate diagnostic information that permits the user to assess system availability on-the-fly. An important component of this broadcast signal is an estimate of pseudorange error, σ_i, for each satellite, i. Given the local satellite geometry, a user may transform these range-domain errors into the position domain, where they define an error bound, or protection level, around the aircraft.

To maintain the validity of these protection limits, the pseudorange error statistics broadcast in the DGPS message must be conservative. This paper describes a method for generating conservative (or overbounding) distributions in the range domain that remain conservative when transformed into the position domain. The technique applies generally to any error distribution, be it non-zero mean, asymmetric, or multimodal.

A. Background: Overbounding Distributions

Protection limits are defined based on FAA specifications for various categories of risk. Generally, these risk requirements are most stringent in the vertical direction. Equation (1) expresses a general form for the VPL, or vertical protection level.

$$VPL = K_H \sigma_p + f \tag{1}$$

For each fault-mode hypothesis, the protection level depends on a hypothesis-specific factor, K_H, and also on σ_p, the instantaneous vertical position error. Some fault hypotheses may also introduce an offset, f. The LAAS one-satellite-out hypothesis (H1), for instance, sets f_j for each receiver, j, based on the outputs, $B_{i,j}$, of a multiple-reference consistency check (MRCC).

$$f_j = \left| \sum_i S_{v,i} B_{i,j} \right| \tag{2}$$

In (1), the position-domain error is composed from the range-domain errors, weighted according to satellite geometry.

$$\sigma_p = \sqrt{\sum_i S_{v,i}^2 \sigma_i^2} \tag{3}$$

Here the range error for each individual satellite, σ_i, is weighted by a coefficient of the vertical-projection row of the geometry matrix, S.

The specifications for integrity risk enter (1) through K_H. This factor is set such that the user vertical position error exceeds the VPL only rarely, as required by the integrity allotment, I_H, for a particular fault hypothesis, H. Assuming the position-domain error has a Gaussian distribution, the K_H factors have been established according to the following equation:

$$K_H = \sqrt{2}\, \mathrm{erf}^{-1}\left(1 - I_H\right). \tag{4}$$

The actual error distribution is not Gaussian. However, the error distribution may be modeled conservatively by a Gaussian distribution, as long as that Gaussian distribution contains more probability mass in its tails beyond the vertical alert limit (VAL) than the actual distribution. Such a conservative distribution is called a tail-overbounding distribution. In tail overbounding, the overbounding cumulative distribution function (CDF), G_o, obeys the following relationship with respect to the actual CDF, G_a. Here x indicates a vertical error value.

$$G_o\left(x=-VAL\right)\geq G_a\left(x=-VAL\right)\ \& \\ \left(1-G_o\left(x=VAL\right)\right)\geq\left(1-G_a\left(x=VAL\right)\right) \quad (5)$$

The tail-overbounding concept provides a powerful tool to abstract the position-domain error in order to simplify the broadcast signal. The broadcast signal, however, carries range-domain error statistics and not position-domain statistics. The user derives position-domain error, according to (3), using range statistics and local satellite geometry. Thus, to be useful for DGPS, an overbounding distribution must remain overbounding when transformed through (3). Unfortunately, tail overbounding in the range domain does not guarantee tail overbounding in the position domain.

Prior researchers have proposed several alternative strategies to tail overbounding in an attempt to establish a bridge between the range and position domains [1]. Notably, these strategies include probability density function (PDF) overbounding and cumulative distribution function (CDF) overbounding. Both these definitions are more restrictive than tail overbounding. A PDF overbound is defined such that the overbounding distribution exceeds the actual distribution for every point outside the VAL:

$$g_o\left(x\right)\geq g_a\left(x\right),\ \forall\left|x\right|>VAL\ . \quad (6)$$

A CDF overbound is defined such that the cumulative distribution function of the overbound, G_o, is always shifted toward its tails relative to the actual cumulative distribution function, G_a, according to (7).

$$G_o\left(x\right)\geq G_a\left(x\right),\ \forall G_a<\tfrac{1}{2} \\ G_o\left(x\right)\leq G_a\left(x\right),\ \forall G_a\geq\tfrac{1}{2} \quad (7)$$

In (7), as elsewhere in this paper, capital G is used to refer to a CDFs, while lowercase g refers to a PDF.

Although the PDF-based strategy appears to offer little advantage over tail overbounding, the CDF-based strategy offers an effective way to link range and position-domain overbounding, at least for certain distributions. Specifically, DeCleene established this link for symmetric, zero-mean, unimodal distributions [2]. DeCleene's proof relies on the fact that the position-domain measurement is the weighted summation of a set of range-domain measurements, according to (3). As the distribution function for a summation of independent random variables is found through convolution, the key step in linking the position and range domains was proving that the convolution operation preserved the CDF-overbounding property. As shown by DeCleene, the convolution of two CDF-overbounding distributions overbounds the convolution of two actual distributions, as long as both the actual and overbounding distributions are zero-mean, symmetric and unimodal.

B. Motivation

In general, convolution does not preserve CDF overbounds for distributions that are asymmetric, shifted-median or multimodal. This section provides examples to demonstrate the collapse of CDF overbounding in the cases of median shifting and asymmetry. A multimodal example is also presented to illustrate the limitations of the Gaussian CDF overbound used in the GPS augmentation application.

Shifted-Median Distributions

The median of a distribution, $g(x)$, is the x value which evenly splits its probability mass (that is the argument for which the cumulative distribution function equals one-half). According to the definition of CDF overbounding by (7), the medians of two distribution functions must be equal for one to provide CDF overbounding of the other. This fact is illustrated by Fig. 1. Thus if two distributions are biased with respect to their medians, the two distributions both provide partial CDF overbounding of the other. After repeated self-convolution, this partial overbounding becomes more pronounced. The following section provides an example to illustrate this condition.

For a symmetric distribution, the mean and median are equivalent. When considering distributions symmetric about zero, the shifted-median condition reduces to the zero-mean condition discussed in [2].

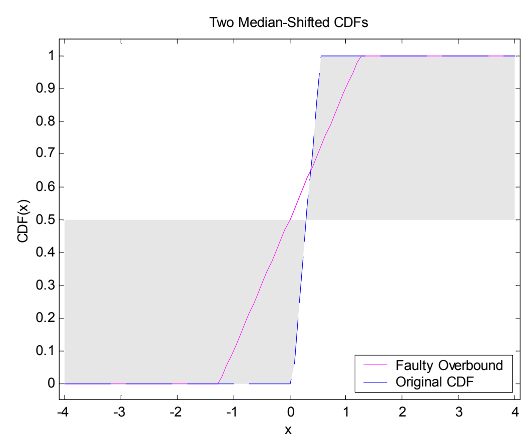

Fig. 1. CDF overbounding is invalid for median-shifted distributions

Asymmetric Distributions

The CDF overbounding definition, (7), holds as long as distribution medians are coincident, even if one or both of the actual and overbounding distributions is asymmetric. Convolution, however, may shift the medians of the actual and overbounding distribution relative to one another in the case of asymmetry. As an example, consider the following two distributions, which satisfy (7):

$$g_o(x) = \tfrac{1}{2} \quad |x| < 1$$
$$g_a(x) = \begin{cases} \tfrac{1}{2} & -1 < x \le 0 \\ 1 & 0 < x \le \tfrac{1}{2} \end{cases} \quad (8)$$

Fig. 2 illustrates these two PDFs, a simple uniform overbounding distribution, g_o, and a quasi-uniform actual distribution, g_a. The CDFs for these distributions are plotted in Fig. 3.

Fig. 3 also illustrates the CDFs of each distribution when convolved with itself 1, 2, or 5 times. The convolution of a distribution with itself *N*-1 times describes the result of the summation of *N* random variables with identical, independent distributions (IIDs). The CDF overbound theorem guarantees, for symmetric g_a and g_o with a common median, that the convolution of the overbounding distribution overbounds the convolution of the actual distribution. For the asymmetric case of (8), however, overbounding breaks down after the 1st convolution (the 2 IID instance). The unbounded probability first appears near the medians of the two distributions as shown in Fig. 3. The mass of unbounded probability increases with each additional convolution, progressing farther and farther toward the tails. This behavior is a clear violation of the conservatism required for precision GPS navigation.

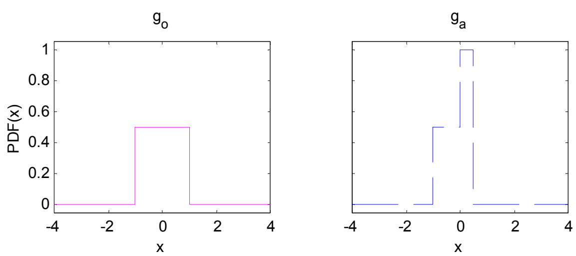

Fig. 2. Symmetric and Asymmetric PDFs

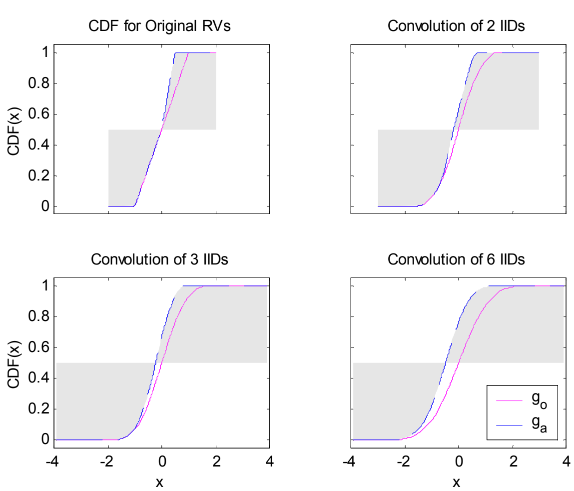

Fig. 3. Breakdown of CDF overbounding for asymmetric PDF

Multimodal Distributions

The condition of unimodality is not strictly necessary for preservation of CDF overbounding through convolution. However, a unimodal-overbound distribution, like the Gaussian assumed for GPS augmentation, does not satisfy the CDF overbounding requirement of (7) for the case of an arbitrary multimodal error distribution. As an example consider the paired delta-function distribution:

$$g_a(x) = \tfrac{1}{2}\delta(x-1) + \tfrac{1}{2}\delta(x+1)\,. \quad (9)$$

The CDF for the paired delta-function is illustrated in Fig. 4. As shown by the figure, this distribution does not have a unique point at which the median exists. Rather, the CDF takes a value of one-half over an extended region of *x* values between 1 and -1. To apply (7) to this distribution requires that the overbounding distribution have a zero probability mass located between *x* values of 1 and -1. Clearly, a unimodal distribution, such as a Gaussian distribution, cannot overbound the paired delta-function distribution.

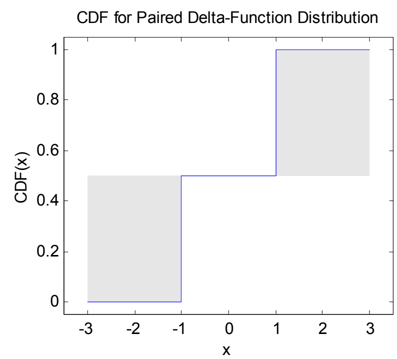

Fig. 4. Paired delta CDF

II. PAIRED OVERBOUND THEOREM

This section introduces a new overbounding strategy called paired overbounding. Paired overbounds are preserved through convolution for arbitrary distributions. Thus the paired overbounding concept effectively relates range-domain and position-domain overbounding even for distributions that have shifted medians, asymmetry and/or multiple modes.

The paired overbound is, in fact, a set of two CDF bounds from which an overbounding CDF may be derived. This set of bounds consists of a left bound, G_L, and a right bound, G_R, defined relative to the actual CDF, G_a, as follows:

$$\begin{aligned} G_L(x) &\ge G_a(x), \ \forall x \\ G_R(x) &\le G_a(x), \ \forall x \end{aligned} \,. \quad (10)$$

In paired overbounding, the overbounding CDF is constructed from the left and right bounds.

$$G_o = \begin{cases} G_L(x) & \forall G_L < \tfrac{1}{2} \\ \tfrac{1}{2} & \text{otherwise} \\ G_R(x) & \forall G_R > \tfrac{1}{2} \end{cases} \quad (11)$$

If condition (10) holds, the paired overbound (11) is, in fact, a CDF overbound in the sense of (7). Strictly speaking, the overbound properties of (11) are not preserved through convolution. However, the properties of the left and right overbounds are preserved through convolution. This is to say if left and right bounds are defined for two arbitrary CDF functions, G_{a1} and G_{a2}, according to (12), then the convolution of the left and right distributions will still bound the convolution of the actual distribution on the left and right according to (13).

$$\text{If}\quad \begin{aligned} &G_{L1}(x) \geq G_{a1}(x),\ \forall x \\ &G_{R1}(x) \leq G_{a1}(x),\ \forall x \\ &G_{L2}(y) \geq G_{a2}(y),\ \forall y \\ &G_{R2}(y) \leq G_{a2}(y),\ \forall y \end{aligned} \quad \text{and} \tag{12}$$

$$\text{then}\quad \begin{aligned} &G_{L1+L2}(z) \geq G_{a1+a2}(z),\ \forall z \\ &G_{R1+R2}(z) \leq G_{a1+a2}(z),\ \forall z \end{aligned}\ . \tag{13}$$

In other words, the distribution of the sum of two random variables, $z=x+y$, is bounded on the left and right by the convolution of the individual bounds for the x and y distributions. A new overbounding distribution for the summation, z, may be constructed from the new left and right bounds, (13), in the form of the overbound (11).

Together, (12) and (13) comprise the paired overbound theorem. The proof of (13), given the assumption (12), follows.

According to standard probability theory, (14) describes the cumulative distribution function for the random variable z, where $z=x+y$ and where x and y are random variables with PDFs $g_{a1}(x)$ and $g_{a2}(y)$, respectively.

$$\begin{aligned} G_{a1+a2}(z) &= \int_{-\infty}^{\infty}\int_{-\infty}^{z-x} g_{a1}(x)\, g_{a2}(y)\, dy dx \\ &= \int_{-\infty}^{z-y}\int_{-\infty}^{\infty} g_{a1}(x)\, g_{a2}(y)\, dy dx \end{aligned} \tag{14}$$

Rearranging the integrals and introducing the appropriate cumulative distribution function yields:

$$\begin{aligned} G_{a1+a2}(z) &= \int_{-\infty}^{\infty} g_{a1}(x)\, G_{a2}(z-x)\, dx \\ &= \int_{-\infty}^{\infty} G_{a1}(z-y)\, g_{a2}(y)\, dy \end{aligned}\ . \tag{15}$$

Applying (15) to the summation of a variable drawn from the *a1* distribution and a second from the *L2* distribution results in:

$$\begin{aligned} G_{a1+L2}(z) &= \int_{-\infty}^{\infty} g_{a1}(x)\, G_{L2}(z-x)\, dx \\ &= \int_{-\infty}^{\infty} G_{a1}(z-y)\, g_{L2}(y)\, dy \end{aligned}\ . \tag{16}$$

Combining this equation with assumption (12),

$$\begin{aligned} \int_{-\infty}^{\infty} g_{a1}(x)\, G_{L2}(z-x)\, dx &\geq \int_{-\infty}^{\infty} g_{a1}(x)\, G_{a2}(z-x)\, dx \\ G_{a1+L2}(z) &\geq G_{a1+a2}(z) \end{aligned}\ . \tag{17}$$

Similarly, because

$$G_{L1+L2}(z) = \int_{-\infty}^{\infty} G_{L1}(z-y)\, g_{L2}(y)\, dy\ , \tag{18}$$

and because of assumption (12), it follows that

$$\begin{aligned} \int_{-\infty}^{\infty} G_{L1}(z-y)\, g_{L2}(y)\, dy &\geq \int_{-\infty}^{\infty} G_{a1}(z-y)\, g_{L2}(y)\, dy \\ G_{L1+L2}(z) &\geq G_{a1+L2}(z) \end{aligned}\ . \tag{19}$$

Thus the left bound property is invariant to convolution:

$$G_{L1+L2}(z) \geq G_{a1+a2}(z)\ . \tag{20}$$

The proof of the invariance of the right bound property has the same form as the development (14) through (20), with the right bound substituted for the left and with less than or equal signs substituted for greater than or equal signs. Given that the left and right bounds are preserved after convolution, the CDF overbound, constructed in the manner of (11), bounds the convolution of the actual distributions a1 and a2.

Thus this proof establishes that the left and right bounding properties are preserved through convolution and, hence, that an overbound based on the convolved left and right distributions overbounds the convolution of the actual distributions. By extension, paired overbounding in the range domain guarantees paired overbounding in the position domain, given an arbitrary error distribution, even one that is not zero-mean, symmetric and unimodal.

III. APPLICATION TO DGPS

The paired overbounding strategy offers two advantages over standard CDF overbounding for GPS augmentation applications. First, the paired overbound's additional degrees of freedom enable construction of tighter error bounds, particularly for the case of non-zero mean error distributions. Second, the generality of the paired overbound permits bounding of arbitrary multipath distributions, including those with more than one mode. This section develops a pair of examples to illustrate tight bounding both of a biased Gaussian error distribution and of a bimodal multipath model. In both cases, the paired overbound theorem reduces extra conservatism and thereby improves integrity without sacrificing availability.

A. Reducing Overbounding Conservatism

For WAAS and LAAS applications, the user assumes that broadcast σ values describe a range-domain error distribution that is zero-mean Gaussian. Because these distributions are not actually zero-mean and Gaussian, the differential GPS reference station must transmit, instead, a Gaussian σ that overbounds the actual error distribution in a conservative fashion. As compared with single-CDF overbounding, paired overbounding introduces additional parameters that may be used to bound more tightly an actual error distribution and thereby to mitigate unnecessary conservatism in the bounding Gaussian.

In this context, a paired Gaussian overbound (with each Gaussian shifted symmetrically from zero), proves useful. The left and right CDFs from which this paired overbound is constructed take the form:

$$\begin{aligned} G_L(x) &= \int_{-\infty}^{x} \mathcal{N}(-b_o, \sigma_o)\, dx \\ G_R(x) &= \int_{-\infty}^{x} \mathcal{N}(b_o, \sigma_o)\, dx \end{aligned} \quad . \qquad (21)$$

Here $\mathcal{N}(b_o,\sigma_o)$ refers to the Gaussian density function with mean b_o and standard deviation σ_o. The subscript o indicates that the mean and deviation describe the overbounding distribution, rather than the actual distribution. Fig. 5 illustrates the left and right bounds and the associated paired overbound.

These paired Gaussians bound a region that is symmetric around zero. The effective VPL for either overbound is shifted by the weighted sum of the Gaussian offsets, $b_{o,i}$, for each of the N satellites in view.

$$VPL_o = K_H \sqrt{\sum_{i=1}^{N} S_{v,i}^2 \sigma_{o,i}^2} + \sum_{i=1}^{N} \left| S_{v,i} \right| b_{o,i} + f_H \qquad (22)$$

In effect, the overbounding biases shrink the alert limit, VAL, the upper allowable bound of the VPL.

Because neither the VAL nor the $b_{o,i}$ parameters are broadcast, the paired overbound cannot be applied directly to shrink the VAL on a user-by-user basis in real time. The $b_{o,i}$ can, however be incorporated in a scaling factor on $\sigma_{o,i}$. This scaling factor, β, remaps the VPL so that both (22) and (23) reach the VAL for the same values of sigma.

$$\begin{aligned} VPL_{o,effective} &= \beta \left(K_H \sqrt{\sum_{i=1}^{N} S_{v,i}^2 \sigma_{o,i}^2} + f_H \right) \\ \beta &= \frac{VAL}{VAL - \sum_{i=1}^{N} \left| S_{v,i} \right| b_{o,i}} \end{aligned} \qquad (23)$$

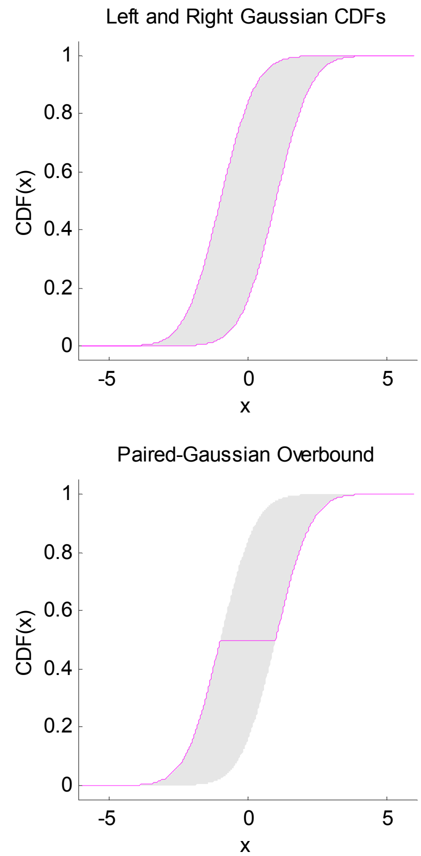

Fig. 5. Paired Gaussian CDFs

The scaling factor multiplies both the range-domain errors and, in the case of the H1 hypothesis, the MRCC B-values of (2). The range-domain error itself consists of four terms, including errors for the ground receiver, the airborne receiver, the ionosphere and the troposphere.

$$\sigma_{o,i}^2 \geq \sigma_i^2 = \sigma_{gr,i}^2 + \sigma_{air,i}^2 + \sigma_{iono,i}^2 + \sigma_{tropo,i}^2 \qquad (24)$$

All of the range-domain error terms are broadcast to the user, except for the airborne value, σ_{air}. Because the broadcast signal cannot dynamically inflate σ_{air}, it is necessary that this error always be inflated to cover a worst-case scenario.

According to (23), β is a function of satellite geometry. The worst case value of β may be empirically bounded by identifying the worst prospective satellite geometries, including all instances with one or two-satellites removed from the constellation. The pathological worst case for a given geometry occurs when all elements of $b_{o,i}$ are assumed to have magnitude equal to the worst individual element.

$$\sum_{i=1}^{N} \left| S_{v,i} \right| b_{o,i} \leq \left\| b_o \right\|_\infty \sum_{i=1}^{N} \left| S_{v,i} \right| = \left\| b_o \right\|_\infty \left\| S_v \right\|_1 \qquad (25)$$

In this pathological case, the geometry and bias terms are decoupled as an infinity-norm and a one-norm, respectively.

Antenna and satellite geometry models may thus be applied independently to generate the worst case bound according to (25). Based on LAAS test prototype (LTP) data, a reasonable model for antenna error, both for the dipole-array multipath limiting antenna (MLA) and the high zenith antenna (HZA), is a Gaussian distribution with an unknown bias, $b_{o,i}$, no larger than 10 cm [3]. Thus $\|b_o\|_\infty$ may be taken as 10 cm.

A practical bound on $\|S_v\|_1$ may be calculated by simulating satellite geometry, including all possible one and two satellite out combinations. To generate such a practical bound, worst-case values of $\|S_v\|_1$ were computed for 10 sites (6 in the coterminous United States and 4 in Europe) at one minute intervals over a sidereal day. Geometries with an un-inflated VPL greater than VAL were discarded. The histogram of the available $\|S_v\|_1$ values over all times and locations is illustrated by Fig. 6a. The most unfavorable one-norm for this set of worst two-satellite-out geometries was $\|S_v\|_1$=14.24. Accordingly, an empirical bound for the one-norm term of (25) is $\|S_v\|_1 \leq 15$. Inserting this one-norm bound into (23), along with an assumed 10 cm $\|b_o\|_\infty$ and a 10 m VAL, results in an inflation factor of β=1.18. This inflation factor bounds the pathological worst-case bias that might occur during a sidereal day.

Allowing the inflation factor to vary over time can reduce excess conservatism. Specifically, the ground station can use the current satellite geometry to generate the inflation factor through (25). Less conservative than the steady-valued inflation factor that covers the entire sidereal day, this instantaneous inflation factor nonetheless provides valid overbounding at a particular time. To implement this solution, the ground station would need to refer to a pre-computed table of worst case $S_v \cdot b_o$ values for all times in one sidereal day. Unfortunately, the airborne user cannot take advantage of the instantaneous inflation factor, if the factor is not broadcast explicitly, and must instead inflate σ_{air} with the more conservative steady-valued inflation factor, 1.18.

A reasonable estimate of average inflation factor for this time-varying case can be constructed assuming that, at any time, only the individual satellites with the worst geometry factor ($S_{v,i} = |S_v\|_\infty$) have large bias errors (i.e. $b_{o,i}$ = 10 cm). For the 10 site simulation, the average instantaneous worst-element was $\langle\|S_v\|_\infty\rangle$=2.67, as shown in Figure Fig. 6b. The resulting average instantaneous inflation is:

$$\beta_{instant} \approx \frac{VAL}{VAL - 2\|b_o\|_\infty \left\langle \|S_v\|_\infty \right\rangle} = 1.06 . \qquad (26)$$

The comparison of this averaged time-varying inflation factor, 1.06, to the steady-valued factor, 1.18, indicates the potential gains of overbounding using instantaneous geometry.

In both its steady and instantaneous forms, this geometry-based inflation strategy compares favorably to more conservative geometry-free inflation factors for mean bounding described in previous work [2,4,5]:

$$\beta = \left(1 + \max_i \left|\frac{\mu_i}{\sigma_i}\right| \frac{\sqrt{N}}{K_H}\right). \qquad (27)$$

Equation (27) assumes the existence of a maximum ratio of mean to standard deviation for the error distribution. For the LTP antennas, unfortunately, this ratio is greater than unity at several elevation angles [3]. Taking an optimistic mean-to-deviation ratio as unity and assuming a standard satellite constellation (N=9) with K_{H0} equal 5.81, the level of inflation recommended by (27) is β=1.52.

The comparison of the two strategies for mean bounding indicates the utility of the new approach. By taking into account a bound on the worst-case available satellite geometry, the new approach reduces conservatism relative to the older, geometry-free approach. Thus, with (23), the inflation factor required to overbound 10 cm antenna biases is as low as 1.06 for instantaneous inflation and 1.18 for worst-case static inflation.

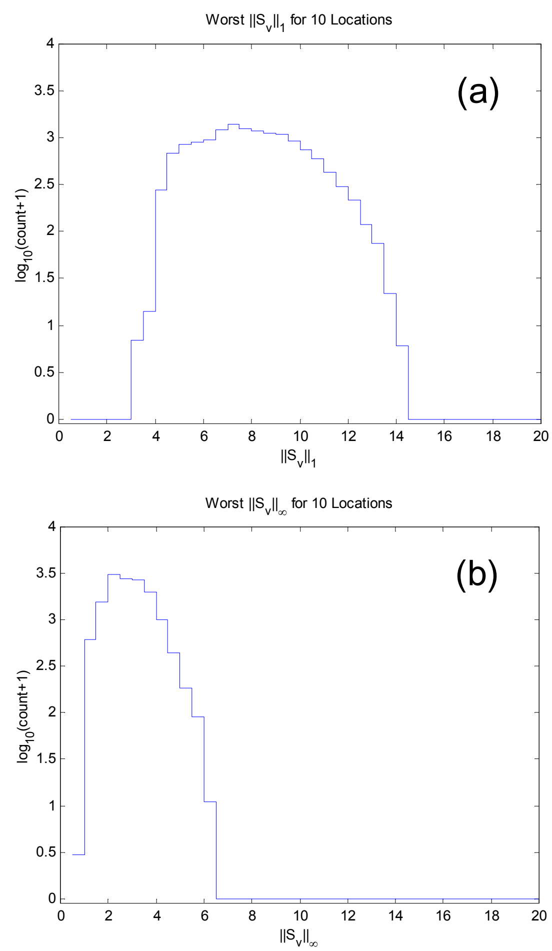

Fig. 6. Log plots of histograms for worst-case (a) $\|S_v\|_1$ and (b) $\|S_v\|_\infty$ computed at one-minute intervals at 10 locations (6 in the conterminous United States and 4 in Europe).

B. Bounding Multipath Error

In addition to handling error biases, the paired overbound theorem also permits the bounding of iregular distributions that may result from multipath modeling. In [4], for instance, Pervan develops a statistical multipath model and hypothesizes that the multipath error bound may be multimodal:

$$\begin{aligned} g_m(x) &= \frac{1}{\pi b\sqrt{1-(x/b)^2}} \quad |x| \le b \\ G_m(x) &= \tfrac{1}{2} + \tfrac{1}{\pi}\sin^{-1}(x/b) \quad |x| \le b \end{aligned} \quad . \qquad (28)$$

The distribution parameter, b, has a dimension in meters and describes the half-width of the PDF. Fig. 7 illustrates the PDF and CDF of this function.

In the absence of the paired overbound theorem, Pervan instead studied this function with tail overbounding in the position domain. To accomplish the transformation analytically, Pervan modeled (28) conservatively as a delta function pair:

$$g_o(x) = \tfrac{1}{2}\delta(x-b) + \tfrac{1}{2}\delta(x+b) \ . \qquad (29)$$

To assess the tail overbound for this paired-delta distribution, Pervan assumed identical multipath error for the range measurements for 12 satellites. The resulting inflation factor for the paired-delta functions was $k_{inf,\delta} = 1.05b/\sigma_a$.

Because this bound is highly conservative, application of the paired overbound theorem can substantially reduce the inflation factor for the multipath error distribution (28). As an example, a pair of uniform distributions could be used as the left and right bounding distributions:

$$\begin{aligned} g_L(x) &= \tfrac{1}{(1+\gamma)b} \quad |x-\gamma| < 1+\gamma \\ g_R(x) &= \tfrac{1}{(1+\gamma)b} \quad |x+\gamma| < 1+\gamma \end{aligned} \ . \qquad (30)$$

Here, the γ parameter is chosen so that the bounds lie tangent to the multipath CDF at its extreme points, as shown in Fig. 8 (γ=0.1385b). Applying these paired-uniform bounds and assuming identical multipath error for each range measurement for 12 satellites (with all $S_{v,i}$ equal their worst case value of 1), the convolved position-domain paired overbound is a twelfth order piecewise polynomial. Fig. 9 compares this polynomial to an equivalent central-limit-theorem Gaussian, $\mathcal{N}\left(N\gamma, \sqrt{N/3}(1+\gamma)\right)$.

This Gaussian, which overbounds the piecewise polynomial conservatively in the tails, can be used to compare the inflation factor derived from the paired-delta distribution to the inflation factor derived for the paired-uniform distributions. By (23) and (25), the β factor for the paired-uniform bound can be bounded assuming a worst-case value of $||S_v||_1$. In accordance with the previous section, VAL is assumed 10 m and $||S_v||_1$ is taken equal to 10.

$$\beta = VAL/(VAL - 10\gamma b) \ . \qquad (31)$$

The associated inflation factor for the paired-uniform distributions is $k_{inf,U} = 0.657b\beta/\sigma_a$. Because β is a function of b, the ratio of the inflation factors for the two distributions are also a function of b:

$$\frac{k_{\inf,\delta}}{k_{\inf,U}} = 1.598(1 - 0.1385b) \ . \qquad (32)$$

The improvement in inflation factor given by (32) is plotted in Fig. 10. Even for multipath distribution modes at ±1 m, the inflation factor is reduced by nearly 40%. Clearly, the use of the paired-overbound theorem enables a substantial reduction in the conservatism required to bound the multipath distribution (28). Still further reductions in the inflation factor could be achieved by using a tighter paired-overbound (i.e. a polynomial paired overbound rather than a uniform one) or by taking advantage of the conservatism in the overbounding Gaussian shown in Fig. 9.

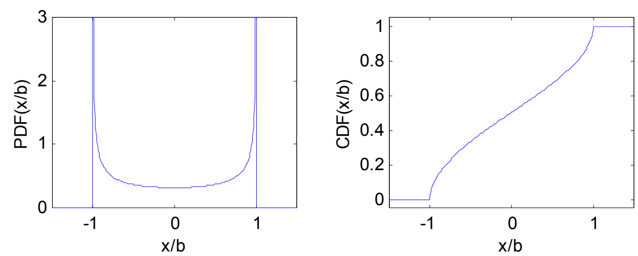

Fig. 7. Hypothesized Multipath Distribution

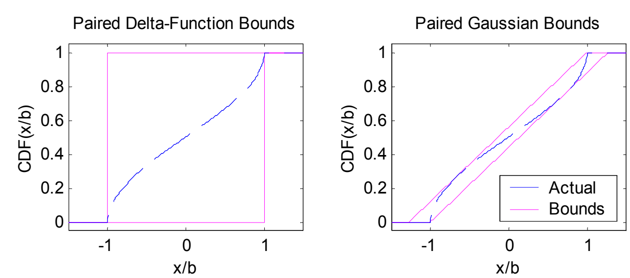

Fig. 8. Paired CDF Bounding

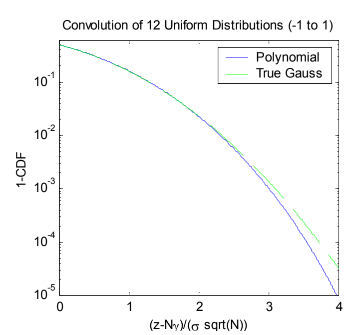

Fig. 9. Comparison of convolved Uniform Distribution CDF to Gaussian CDF

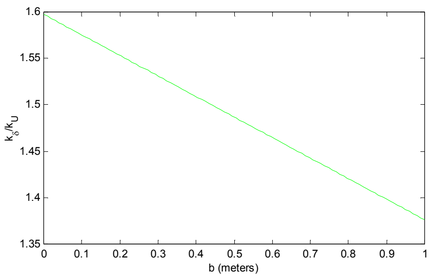

Fig. 10. Comparison of inflation factors for two bounding distributions

IV. SUMMARY

This paper introduced an overbounding theorem using paired left and right CDF overbounds. The theorem provides a means to bound error distributions in such a way as to guarantee bounding after a convolution operation. Consequently, the distribution for the sum of a set of random variables is bounded if the distributions of the individual random variables are bounded. In contrast with other overbounding methods, this paired-overbound theorem applies to arbitrarily shaped range-domain error distributions, even distributions that are non-zero mean, asymmetric, or multimodal.

In the field of differential GPS augmentation, the paired overbound theorem provides a means to guarantee bounding of the position-domain error given a bounded range-domain distribution. This property has two applications for GPS augmentation. First, the theorem provides a new tool to reduce conservatism in the broadcast σ. This process of mitigating overly conservative assumptions will prove significant in certifying GPS augmentation for increasingly stringent FAA requirements. Second, the theorem provides a new tool to model irregular multipath error distributions. No tool previously existed to tightly bound these distributions, which may, in general, be non-zero mean, asymmetric or multimodal.

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] P.B. Ober, R. Farnsworth, E. Breeuwer, and D. van Willigen. SBAS Integrity Verification. In Proceedings ION GPS 200, pp. 1805-1812, 2001.

[2] B. DeCleene. Defining Pseudorange Integrity – Overbounding. In Proceedings of ION GPS 2000, pp. 1916-1924.

[3] William J. Hughes FAA Technical Center, ACT 360. LTP data for 10/21/00, 3/26/01, 7/27/01, and 10/6/01. Unpublished presentation.

[4] B. Pervan, S. Pullen, and I. Sayim. Sigma Estimation, Inflation, and Monitoring in the LAAS Ground System. In Proceedings ION GPS 2000, pp. 1234-1244, 2000.

[5] C. Shively. A Proposed Method for Including the Mean in the LAAS Correction Error Bound by Inflating the Broadcast Sigma. Mitre Memorandum, pp.1-8, 2000.