

A Proposed Integrity Equation for WAAS MOPS

Todd Walter, Per Enge and Andrew Hansen

*Dept. of Aeronautics and Astronautics
Stanford University*

ABSTRACT

In recent years there has been widespread growth in the number of proposed differential augmentation systems for GPS. Such systems are being developed by both civil authorities and commercial interests. The most stringent application of these systems is precision approach for aviation, where the system provides vertical and horizontal guidance. Precision approach has very strict requirements for accuracy, integrity, continuity and availability, and these become more stringent as the decision height decreases. To date, it appears that the Wide Area Augmentation System (WAAS) will be able to meet the accuracy requirements all the way down to a 200 foot decision height. However, the requirements on integrity and availability will require special attention. In order to guarantee integrity, confidence bounds must be sufficiently large so as to always cover the error of the broadcast correction. However, to maintain availability, the confidence bounds cannot be unduly conservative. As such, careful choice of the integrity equation is paramount.

This paper derives an integrity equation for the position solution starting from the expected distribution of errors on the corrected pseudorange measurements. We show that the details of the integrity equation depend on these expected error distributions. As a further complication, the integrity requirement applies individually to each aircraft conducting precision approach and not to an ensemble average over the airspace as a whole. Thus, there is some question as to how to calculate and apply confidence bounds. An additional constraint is the bandwidth limitation on the corrections. There are a limited number of bits with which to send corrections and confidences to the user. An integrity equation which significantly increases the amount of information broadcast to the user is not practical. The proper integrity equation must take all of these constraints into consideration.

We propose an equation which satisfies the integrity requirement while maximizing availability. Additionally this equation is fully compatible with the WAAS Minimum Operational Performance Standards (MOPS) and requires no additional bits beyond the current specification. The bits used for integrity also serve to weight the measurements for the position solution. The proposed integrity equation has been verified through Monte Carlo simulation. In addition, this equation is applied to real data from the National Satellite Test Bed (NSTB) network.

INTRODUCTION

The most important and most challenging problem when using GPS for aviation positioning is guaranteeing the integrity of the final solution. The goal is to bound the error in the position estimate. A navigation solution is termed Hazardously Misleading Information (HMI) if the error is greater than a specified bound (Vertical Alarm Limit or VAL) and not flagged to the pilot. For precision approach to within 350 to 200 feet above the ground, the Probability of HMI (Pr_{HMI}) must be less than 10^{-7} per approach for each aircraft conducting this phase of flight [1]. This problem is particularly challenging because the availability of precision approach must be kept very high ($> 99.9\%$). Thus, our methods cannot be overly conservative at every step, or we will have a system which is rarely, if ever, available. On the other hand, the bounds cannot be under conservative, as precision approach is a safety-of-life operation and the Pr_{HMI} constraint must be met at any cost. Therefore, we must describe and trace the errors in the system as accurately as possible, so that the end result is a probability density function that bounds the actual errors yet comes as close to the true distribution as is practical.

The current proposal for operating procedures [2] specifies that the master station communicate both wide area differential GPS corrections and confidences to the user. This information is combined into corrections and confidence bounds for the individual pseudorange measurements. However, all of the navigation requirements are in the user's position domain. In order to translate the pseudorange errors into position errors we need to know two things: the Probability Density Functions (PDFs) for the measurement errors, and how those PDFs combine to form the PDF of positioning error. Of course, we don't really know the PDFs for the pseudorange errors, but we can attempt to bound them. It is extremely important that the master station be able to bound the true distribution of errors. The more information we have about the measurements and how they are processed, the more accurately we can predict the distribution of errors. Even with this information, assumptions about the tail behavior of the errors will be necessary due to the finite number of observations available.

It is important to note that this paper discusses nominal operation only. It is possible, through algorithm design flaw, interference, receiver failure, etc., for the measurement error to exceed its specified bound. However, this situation is not a nominal, "fault-free", mode of operation. If some error or fault has occurred, this situation must be identified through a fault detection method. Typically these methods exploit redundancy to prevent the fault from causing an integrity breach. Such techniques might employ redundant master stations to verify the corrections before they are broadcast [3], redundant receivers to isolate receiver faults [4], or redundant measurements to isolate individual measurement errors [5] [6]. Again, we are concerned with only the nominal operation. Assuming we have incomplete information about the distribution of the measurement errors, how do we best translate that information into confidence bounds on the user's position?

GPS MEASUREMENT EQUATION

If we wish to translate the bounding PDFs of the measurement errors into the bounding PDF for the positioning error we must determine how the measurements and the user position are connected. The pseudorange measurements can be related to the true four dimensional position vector (north, east, up and clock), \mathbf{x} , through the linearized equation [7]

$$\mathbf{y} = \mathbf{S} + \mathbf{G} \cdot \mathbf{x} + \boldsymbol{\eta} \quad (1)$$

where \mathbf{y} is an N dimensional vector containing the raw or corrected pseudorange measurements, \mathbf{G} is the observation matrix, \mathbf{S} is an N dimensional vector containing information about the position and clock state of each satellite, and $\boldsymbol{\eta}$ is an N dimensional vector containing the errors on the measurements. The observation matrix consists of N rows of line of sight vectors from \mathbf{x} to each satellite, augmented by a -1 for the clock. Thus the i^{th} row corresponds to the i^{th} satellite in view and can be written in terms of the elevation angle El_i and the azimuth angle Az_i

$$\mathbf{G}_i = -[\cos El_i \cos Az_i \quad \cos El_i \sin Az_i \quad \sin El_i \quad -1] \quad (2)$$

The elements in the vector \mathbf{S} are given by

$$S_i = -\mathbf{G}_i \cdot \mathbf{X}_i \quad (3)$$

where \mathbf{X}_i is the four dimensional position vector for the i^{th} satellite in the same coordinate frame as \mathbf{x} .

Because \mathbf{G} depends on \mathbf{x} , this problem is non-linear and can be complicated to solve. However, let us assume that the linearization is accurate and we have an estimate for \mathbf{x} . We wish to express the error in this estimate in terms of the observation errors. With estimates for the user's and satellites' position and clock vectors in hand, we can generate estimates for \mathbf{G} , \mathbf{y} and \mathbf{S} . Then by subtracting the estimate for \mathbf{y} from the actual measurements and by assuming that $\hat{\mathbf{x}}$ is sufficiently accurate such that $\hat{\mathbf{G}} = \mathbf{G}$, we obtain

$$\Delta \mathbf{y} = \mathbf{G} \cdot \Delta \mathbf{x} + \boldsymbol{\varepsilon} \quad (4)$$

where $\Delta \mathbf{y} \equiv \mathbf{y} - \hat{\mathbf{y}}$, $\Delta \mathbf{x} \equiv \mathbf{x} - \hat{\mathbf{x}}$, and $\boldsymbol{\varepsilon}$ is the observation error including measurement, satellite position, and clock errors:

$$\boldsymbol{\varepsilon} \equiv \mathbf{S} - \hat{\mathbf{S}} + \boldsymbol{\eta} \quad (5)$$

Equation (4) defines the linear transformation of PDFs for $\Delta \mathbf{x}$ into PDFs for $\Delta \mathbf{y}$. We are interested in the inverse of this problem. In general the number of measurements, N , exceeds the number of unknowns, 4. Thus, the solution is overdetermined and the optimal inversion depends on the PDFs of $\boldsymbol{\varepsilon}$.

PROBABILITY DENSITY FUNCTIONS

If we were to assume that the PDFs of $\boldsymbol{\varepsilon}$ could take arbitrary form, the integrity problem would be insolvable. Fortunately, we know from past experience that the GPS measurement errors are generally well

behaved. That is, they have a greater likelihood of being small than they do of being large. The errors can be traced from the measurements made at the reference stations, through the various algorithms, all the way to the corrected pseudorange observation. Along the way many different error sources are combined and averaged, biases are calibrated or estimated and removed, and outliers are detected and removed. The central limit theorem [8] predicts that each time multiple error sources are combined to form a whole, the PDF of the resulting error will look more and more gaussian. Thus, it is not unreasonable to assume that the PDFs of ϵ should be not too far from gaussian. In fact, because the master station also employs fault detection and exclusion algorithms, it is quite likely that the tails of the actual distribution will be clipped. That is, a gaussian or near gaussian distribution will probably be overly conservative in describing the true PDFs. However, when calculating the integrity, we cannot safely take advantage of this likelihood. Instead, we must consider distributions other than gaussian so that we may verify that the integrity equation does not rely too heavily on this assumption. We desire an integrity equation which is robust against a wide class of PDFs.

The approach outlined in the above paragraph relies on covariance propagation. It makes use of the central limit theorem to predict a gaussian final result. Another possible means for containing errors is a threshold method. This type of method exploits redundant measurements which are presumed to be independent of each other. If the measurements agree with one another within a specified threshold, the data is treated as if it is accurate to a level on the same order as the threshold. If the measurements do not agree, the data is flagged as less accurate or invalid. Provided the measurements are independent and the occurrence of bad data is sufficiently infrequent, this method should contain the errors within an upper and lower bound.

Another error source worth investigating is one introduced by the constraint of our correction data link. The master station is not able to communicate all of its knowledge to the user. For example, if we had measurements which showed unquestionable variations in the ionosphere at a much finer spatial distance than 5° , we could not broadcast all of that information to the users. The grid point corrections would strive to minimize the average error over the area and the confidence value would have to reflect the known errors that the users might experience along with the model uncertainty. Part of the integrity question is how to combine these error sources.

Let us start to look at some possible PDFs for the observation errors. All of the PDFs we consider will have zero mean. This is appropriate because, while

individual measurements and corrections will have biases, if we had knowledge as to the sign of that bias we could estimate it and remove it. Thus, the remaining biases must be equally likely to be positive or negative, otherwise we have incorrectly designed our master station and user algorithms. The covariance propagation method favors a zero mean gaussian distribution given by

$$n(x) \equiv \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{x^2}{2\sigma^2}} \quad (6)$$

where σ^2 is the variance of the gaussian distribution. The threshold approach favors either a clipped gaussian or a uniformly distributed variable whose PDF is given by

$$u(x) \equiv \begin{cases} \frac{1}{2a} & |x| < a \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

where a is the resulting bound on the error. A more conservative view of the threshold approach is that since the error can take any value up to a , we must assume that the user has the worst case error. This view is particularly applicable for the errors induced by the bandwidth constraint. Here we have knowledge that certain users will in fact have errors as great as the bound. The uniform distribution more accurately describes the PDF for users scattered within the service area, but for integrity purposes we must use the more conservative distribution given by

$$d(x) \equiv \frac{1}{2} \{ \delta(x+a) + \delta(x-a) \} \quad (8)$$

where $\delta(x)$ is the dirac delta function. This density function places all of the energy right at the thresholds.

There are many other approaches that could be taken in describing the PDFs of the measurement errors. However, we feel that these approaches can be combined to create a reasonable bounding PDF for any fault free measurement errors. For integrity purposes, we consider each element of ϵ to be composed of the sum of two random variables, one with the conservative distribution (8) and one of normal distribution. This new random variable has the following PDF:

$$f_1(x) \equiv \frac{1}{2\sigma\sqrt{2\pi}} \left\{ e^{-\frac{(x+a)^2}{2\sigma^2}} + e^{-\frac{(x-a)^2}{2\sigma^2}} \right\} \quad (9)$$

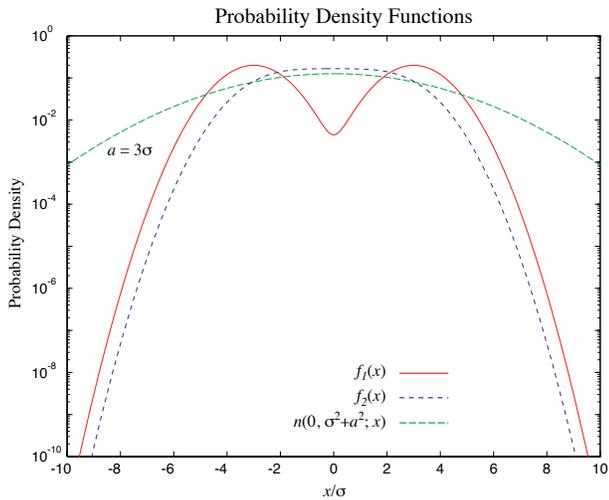


Figure 1. Three Probability Density Functions plotted as a function of x/σ .

The expected distribution for the ensemble airspace would be the sum of the gaussian with the uniform in place of the conservative distribution. Its PDF is given by

$$f_2(x) \equiv \frac{1}{4a} \left\{ \operatorname{erf}\left(\frac{x+a}{\sigma\sqrt{2}}\right) - \operatorname{erf}\left(\frac{x-a}{\sigma\sqrt{2}}\right) \right\} \quad (10)$$

The three distributions considered for this paper are shown for similar conditions in Figure 1.

Knowing the full PDF for the errors allows us to specify containment bounds for arbitrary probabilities. The bounds are found by integrating the PDFs from $-b$ to b . The value of b which integrates to 1 - Probability is the corresponding bound. Table 1 lists some confidence bounds of x/σ for various probabilities. For the distributions given by f_1 and f_2 , we have set $a = \sigma$. Of course, we have less confidence in our knowledge of the true PDF for lower probabilities. However, we can remain hopeful that the fault detection and isolation algorithms do remove true outliers. This would result in a true PDF which is bounded by (8) or (9). Notice that $f_1(x)$ is the most conservative distribution for large x . Despite this fact, the zero mean constraint creates smaller confidence bounds than would be expected from a biased gaussian. The latter case results in bounds given by $b = a + \kappa(Pr) \sigma$, where $\kappa(Pr)$ corresponds to the values listed in

Table 1 for the gaussian distribution.

POSITION SOLUTION

The next step towards finding the confidence bounds in the position domain is to determine the best method to estimate the position. If the errors were gaussian in distribution, the maximum likelihood estimator would be weighted least squares. Given (9) and (10) it is possible to devise a maximum likelihood estimator specific to each distribution. However, provided a is not much greater than σ , neither of these distributions is too far from gaussian. An estimator optimized to one of these bounding functions would not yield a large improvement over weighted least squares, even if the errors truly had that distribution. An advantage of the weighted least squares method is that it is easily solved. Other estimators require non-linear techniques which are not necessarily guaranteed to converge. Weighted least squares is simple, fast and sufficiently accurate. The weighted least squares solution requires only the variances for measurement errors. The variances are defined as:

$$\begin{aligned} \sigma_{\varepsilon_i}^2 &\equiv \langle (\varepsilon_i - \langle \varepsilon_i \rangle)^2 \rangle \\ \sigma_{\varepsilon_{ij}}^2 &\equiv \langle (\varepsilon_i - \langle \varepsilon_i \rangle) \cdot (\varepsilon_j - \langle \varepsilon_j \rangle) \rangle \end{aligned} \quad (11)$$

where the brackets, $\langle \rangle$, about an object denote its expected value. All of the PDFs we have considered are zero mean. Again this is a necessary result of well designed correction algorithms and receivers. If we had knowledge that the measurement errors had an expected value other than zero, we could remove this expected bias from the measurements and produce a better position solution. The variances for our conservative and expected distribution functions are given by:

$$\begin{aligned} \sigma_{f_1}^2 &= \sigma^2 + a^2 \\ \sigma_{f_2}^2 &= \sigma^2 + \frac{a^2}{3} \end{aligned} \quad (12)$$

The weighted least squares estimator minimizes the cost function given by

<i>Dist \ Pr</i>	10^{-2}	10^{-3}	10^{-4}	10^{-5}	10^{-6}	10^{-7}	10^{-8}	10^{-9}
$n(x)$	2.576	3.291	3.891	4.417	4.892	5.327	5.731	6.109
$f_1(x)$	3.327	4.090	4.719	5.265	5.753	6.199	6.612	6.998
$f_2(x)$	2.938	3.718	4.363	4.924	5.425	5.882	6.305	6.699

Table 1. Confidence bounds for varying Probabilities for the four Probability Density Functions. The bounds are on x/σ and we have set $a = \sigma$.

$$\Delta \hat{\mathbf{y}}^T \cdot \mathbf{W} \cdot \Delta \hat{\mathbf{y}} \quad (13)$$

where the weighting matrix, \mathbf{W} , is the inverse of the measurement covariance matrix

$$\mathbf{W}^{-1} = \langle \boldsymbol{\varepsilon} \cdot \boldsymbol{\varepsilon}^T \rangle = \begin{bmatrix} \sigma_{\varepsilon_1}^2 & \sigma_{\varepsilon_{1,2}}^2 & \text{L} & \sigma_{\varepsilon_{1,N}}^2 \\ \sigma_{\varepsilon_{1,2}}^2 & \sigma_{\varepsilon_2}^2 & \text{L} & \sigma_{\varepsilon_{2,N}}^2 \\ \text{M} & \text{M} & \text{O} & \text{M} \\ \sigma_{\varepsilon_{1,N}}^2 & \sigma_{\varepsilon_{2,N}}^2 & \text{L} & \sigma_{\varepsilon_N}^2 \end{bmatrix} \quad (14)$$

If we had no knowledge about the correlation between measurement errors, then the off-diagonal elements of the covariance matrix should be set to zero in order to form the position solution. However, knowledge of these cross-correlation terms, not merely bounds, could be included to improve the position solution. It is unlikely that this information would be known to the aircraft's receiver. Therefore, the cross-correlation terms will be set to zero, simplifying the procedure for finding the position estimate. If there were no reason to assume that any one measurement were better than any other, the problem could be further simplified by setting the covariance matrix to a constant times the identity matrix. This situation leads to the traditional non-weighted position solution [7].

Given an initial estimate for position and an estimate of the weighting matrix, we can perform the weighted least squares inverse of (4) to improve our position estimate. The solution is given by

$$\Delta \hat{\mathbf{x}} = (\mathbf{G}^T \cdot \mathbf{W} \cdot \mathbf{G})^{-1} \cdot \mathbf{G}^T \cdot \mathbf{W} \cdot \Delta \hat{\mathbf{y}} \equiv \mathbf{K} \cdot \Delta \hat{\mathbf{y}} \quad (15)$$

where the definition has been made

$$\mathbf{K} \equiv (\mathbf{G}^T \cdot \mathbf{W} \cdot \mathbf{G})^{-1} \cdot \mathbf{G}^T \cdot \mathbf{W} \quad (16)$$

As we mentioned before, the position solution may need to be iterated, updating $\hat{\mathbf{x}}$, $\hat{\mathbf{G}}$, $\hat{\mathbf{S}}$, and $\hat{\mathbf{y}}$ at each step. We obtain our best estimate for position, $\hat{\mathbf{x}}$, when we have driven $\Delta \hat{\mathbf{x}}$ to zero and the cost function (13) is at a minimum.

Notice that the definition for \mathbf{K} has the weighting in both its denominator and its numerator. If all variance estimates were uniformly scaled up or down, it would have no impact on the position solution. The best estimate of \mathbf{x} depends only on the relative confidence between the measurements and not on their absolute values. The same is not true for the confidence bounds on

the position estimate. In fact, we will show that the bounds are very much dependent on the overall scaling of the measurement variances.

CONFIDENCE BOUNDS

Finding the confidence bound is not as simple as finding the position estimate. Again there are at least two approaches one can take: error thresholding and covariance propagation. The threshold approach assumes that the user has N measurements each with some error, ε_i , bounded by b_i . These measurement errors are related to the position errors by the \mathbf{K} matrix. Thus, the Vertical Protection Limit (VPL), assuming a single error, would be given by $K_{3,i} \cdot b_i$. However, each individual measurement has an associated error and confidence bound. The final vertical position bound depends on the relationship of these errors and their PDFs. The most conservative bound, which would work for arbitrary PDF, is given by:

$$\text{VPL}_{\text{absolute}} \equiv \sum_{i=1}^N |K_{3,i} \cdot b_i| \quad (17)$$

This bound is overly conservative in most situations. This protects against the worst case errors combining in the worst possible way. While this method has high integrity, the availability would almost certainly be very low. Instead we would like to find a less conservative method that provides full integrity while yielding better availability.

The covariance propagation method's estimate of the final PDF of position error would be given by the convolution of each element $K_{3,i} \cdot \varepsilon_i$. The resulting distribution approaches a gaussian as N becomes large. The mean and variance of the positioning distribution are given respectively by the sums of the means and variances of the original variables, $K_{3,i} \cdot \varepsilon_i$. Thus, the second moment of vertical estimate can be written:

$$\sigma_{\hat{x}_3} = \sqrt{\sum_{i=1}^N (K_{3,i} \cdot \sigma_{\varepsilon_i})^2} \quad (18)$$

We can define another Vertical Protection Limit using (18). The assumption that the position error is bounded by a gaussian with variance (18), and confidence bounds from Table 1, results in the less conservative VPL:

$$\text{VPL}_{\sigma} \equiv \kappa(Pr) \cdot \sigma_{\hat{x}_3} \quad (19)$$

If each individual error ε_i has a gaussian distribution, then κ may be pulled inside the radical and applied to each element. Here, the final bound is expressed as the square root of the sum of the squares of the measurement confidence bounds

$$\text{VPL}_{\text{sum of squares}} \equiv \sqrt{\sum_{i=1}^N (K_{3,i} \cdot b_i(Pr))^2} \quad (20)$$

where $b_i(Pr)$ is the confidence bound on ε_i for a given probability, Pr . However, one must be extremely cautious when using this as a protection limit. It is only valid if each individual measurement error distribution is gaussian. Otherwise (20) will likely underestimate the bound on position creating an integrity risk. As we will later see, even for leptokurtic distributions such as (9) and (10) the predicted limit in (20) does not bound the actual error.

There is a more efficient method to calculate the second moment of expected positioning accuracy than (18). The covariance matrix for the position estimate is given by:

$$\begin{aligned} \langle \Delta \hat{\mathbf{x}} \cdot \Delta \hat{\mathbf{x}}^T \rangle &= \langle (\mathbf{K} \cdot \Delta \hat{\mathbf{y}}) \cdot (\mathbf{K} \cdot \Delta \hat{\mathbf{y}})^T \rangle \\ &= \mathbf{K} \cdot \langle \Delta \hat{\mathbf{y}} \cdot \Delta \hat{\mathbf{y}}^T \rangle \cdot \mathbf{K}^T \end{aligned} \quad (21)$$

Then, because our best estimate of $\langle \Delta \hat{\mathbf{y}} \cdot \Delta \hat{\mathbf{y}}^T \rangle$ is the same as our best estimate for $\langle \varepsilon \cdot \varepsilon^T \rangle$, we can use (14) to simplify this to:

$$\langle \Delta \hat{\mathbf{x}} \cdot \Delta \hat{\mathbf{x}}^T \rangle = (\mathbf{G}^T \cdot \mathbf{W} \cdot \mathbf{G})^{-1} \quad (22)$$

Thus, the full position estimate covariance matrix is available when the position solution is calculated. The variance of the vertical position estimate is given by the third diagonal element of the position estimate covariance matrix. We need not apply the \mathbf{K} matrix to the expected measurement variances. Equation (18) is replaced with the more efficient equation

$$\sigma_v \equiv \sqrt{\left[(\mathbf{G}^T \cdot \mathbf{W} \cdot \mathbf{G})^{-1} \right]_{3,3}} \quad (23)$$

where σ_v replaces the notation $\sigma_{\hat{x}_3}$, and with this minor change in notation, (19) becomes

$$\text{VPL}_{\sigma_v} \equiv \kappa(Pr) \cdot \sigma_v \quad (24)$$

For integrity purposes, \mathbf{W} could either be the same diagonal weighting matrix used to form the position estimate, or we can include bounds on cross-correlation

terms. While these bounds are unlikely to improve our position estimate, they will most likely make the expected integrity bounds more conservative, by properly accounting for correlations between measurement errors. These off diagonal elements should be applied only when we actually have knowledge of their value. For example, the error in the vertical tropospheric delay will be correlated for all of the users measurements, but the tropospheric mapping function errors may not be. Ionospheric delays will have some correlation across the user's measurements, but differentially corrected ionospheric delays will have much less correlation. Assumed cross correlations should be added with hesitancy, and only if we truly believe they are necessary to protect integrity.

We now have two candidates, (17) and (24), for providing the vertical protection level for precision approach. Equation (20) will be investigated for comparison purposes as it is only applicable for gaussian measurement errors and is not robust to different distributions.

MONTE CARLO RESULTS

The equations of the preceding sections were verified using Monte Carlo simulations. Random geometries were generated by choosing arbitrary user locations within the Conterminous United States (CONUS) and using random times to position 24 of the GPS satellites using almanac information. The user weighting is based on a fit to the curve in [6], the nominal value as a function of elevation is given by

$$\sigma_{\varepsilon_i} \approx e^{(1.4175 \sin^2 \varepsilon_i - 2.9125 \sin \varepsilon_i)} \quad (25)$$

where the amplitude of the fit, $\sigma_A = 3.45$ m, has been left off as it is the relative scaling which is most important. Once the geometry was selected and the elevation angle determined for each satellite, the values for σ and a were independently selected from a uniform distribution ranging from 0.5 to 1.5 multiplied by the value calculated in (25). This process is depicted in Figure 2. The curve in (25) and Figure 2, shows nearly a 4 to 1 ratio between the parameters very high and very low satellites. In addition, the 0.5 to 1.5 spread in the selection of σ and a implies that the ratio of a/σ for any given satellite can take on the value of 1/3 to 3.

Using a randomly selected \mathbf{G} matrix we were able to create a random \mathbf{W} matrix. We then used (17), (20) and (24) to generate predicted Vertical Protection Limits. These VPLs could be compared to actual vertical errors

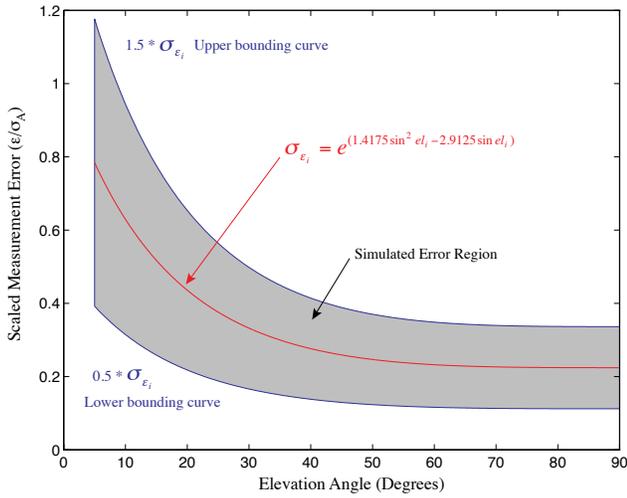


Figure 2. Curves from which values of σ and a were generated as a function of elevation angle. The center curve corresponds to (25) and the shaded region shows the range of possible values. For example, if the elevation angle for a satellite is 5° , then both a and σ are chosen from a uniform distribution between 0.4 and 1.18.

created by simulating error terms distributed according to (9). For a given probability, Pr , we used fixed \mathbf{G} and \mathbf{W} matrices and generated $30/(1 - Pr)$ position solutions to determine the true bound. For example, for a 99.9% probability we generated 30,000 position solutions with errors distributed according to (9) and calculated the “true” 99.9% position bound. The “true” bound was then compared to the candidate VPLs. The uncertainty in the “true” bound was found by using the same number of cases to estimate the bound for a gaussian. This value was then compared to the calculated values in Table 1. It was found that, for a gaussian, this numerical method of estimating the “true” bound was accurate to better than 3% more than 95% of the time.

This process was applied to 10,000 geometries for each probability level. Due to processing constraints, the probabilities we evaluated were 99.9%, 99.99% and 99.999% rather than at the full P_{HMI} . Figure 3 shows the histograms of the ratios of the “true” error to the three different vertical protection levels. Ideally the histogram would be a narrow spike right at or below unity. This would indicate that the bound exactly matched the actual error and correctly predicted its magnitude. As is evident in the figure, VPL_{absolute} (17), provides full integrity, but sacrifices quite a bit of availability by predicting an error roughly twice as large as the actual. $VPL_{\text{sum of squares}}$ (20) does not provide integrity and cannot be safely used. VPL_{σ_v} (24) has nearly the ideal distribution. It is pushed up against 1, maximizing availability, while maintaining the tightest distribution among the candidates. In addition, the long tail of the distribution is toward the left,

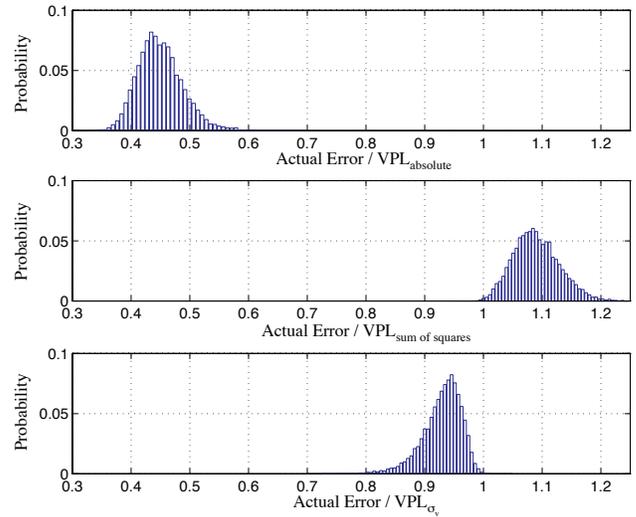


Figure 3. Histograms of the “true” 99.999% position error compared to that predicted by the various vertical protection levels (17), (20), and (24) assuming perfect knowledge of a and σ .

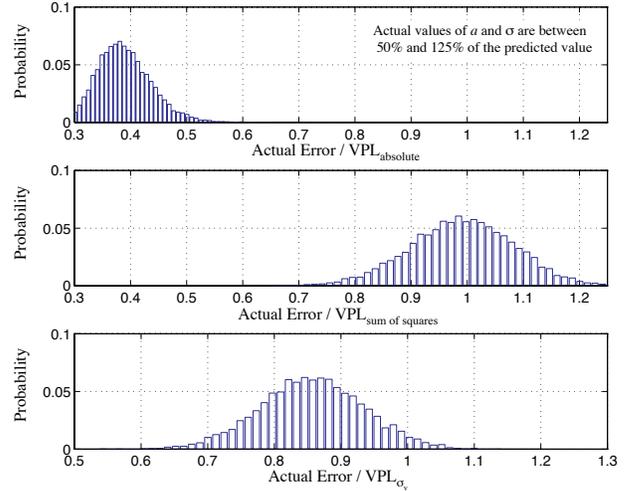


Figure 4. Histograms of the “true” 99.9% position error compared to that predicted by the various vertical protection levels (17), (20), and (24) where a and σ are imprecisely known.

meaning that when it is most in error, it predicts high. The other two protection levels have their tails running toward the right. Thus, even if we were to rescale them (divide VPL_{absolute} by 1.5 to increase availability or multiply $VPL_{\text{sum of squares}}$ by 1.25 to increase integrity) the distributions still have worse shapes. Neither one could be made to optimize both integrity and availability as effectively as VPL_{σ_v} .

In order to form the $Pr\%$ bound for (17) and (20) we used the approximation $b(Pr) = a + \kappa(Pr)\sigma$. This value is a little conservative as can be seen in Table 1. Depending on the ratio of a/σ and the value of Pr , the

bound could be nearly 10% conservative. Unfortunately, calculating the true bound, b , requires inverting the integral of (9). While this can be done numerically, it is time consuming. To see the effect this overestimation had on VPL_{absolute} and $VPL_{\text{sum of squares}}$, we ran one simulation calculating the exact 99.9% bound. While this resulted in a slightly better magnitude for VPL_{absolute} , it shifted $VPL_{\text{sum of squares}}$ more than 5% to the right. Thus using the exact bound in (20) caused it to become even worse at providing integrity. In neither case did it improve the shape of the histograms.

For Figure 3, the exact values of a and σ were known and could be used in the weighting matrix and to form the integrity bounds. This situation is not realistic, but from an integrity viewpoint is close to the worst allowable case. In reality a and σ are not known exactly, instead we must use upper bounds. If the estimated upper bounds are larger than the actual value, integrity is preserved. However if the estimates are lower than the actual values then none of these equations can guarantee integrity. Using estimates that exactly match the true values optimizes availability, but come close to sacrificing integrity. When the values are unknown, the largest possible values must be used. In this case the integrity is even better than is demonstrated in Figure 3.

Figure 4 demonstrates that estimates for a and σ not bounding the actual values will lead to integrity breaches. Here the actual values of a and σ are independently uniformly distributed between 50% and 125% of the predicted values. Because the predicted values do not necessarily bound the true values we do not expect these equations to preserve integrity despite the fact that most of the errors are likely over bounded. It is evident from Figure 4 that the overly conservative VPL_{absolute} still protects the user while the other two equations do not. The important lesson from these results, however, is that the estimated values for a and σ must bound the true ones. If there is a possibility that the true values will be 125% of the estimates then we must inflate our estimates by at least 125%. It is evident that if the x-axis of the graphs in Figure 4 are divided by 1.25 both VPL_{absolute} and VPL_{σ_v} would protect integrity while our third candidate would not. It is important to remember that for our integrity equations *we cannot use our best estimate of variance, instead we must use our bounding estimate.*

BANDWIDTH CONSIDERATIONS

The proposed integrity equation (24) makes efficient use of the correction bandwidth. Both (17) and (20) require the variance, $\sigma_i^2 = \sigma^2 + a^2$, and the bound, $b_i(Pr) = a + \kappa(Pr)\sigma$, for each satellite to be sent.

However, (24) only requires the variance. Thus we need not separate the bias and noise terms and send them separately. This integrity equation conforms most closely with the current MOPS specifications [2]. According to the MOPS, the master station must send to the user estimates of the 99.9% bound of the satellite correction error (User Differential Range Error or UDRE) and the 99.9% bound of the ionosphere correction error (Grid Ionosphere Vertical Error or GIVE). In order to take advantage of our proposed integrity equation the UDRE and GIVE values would merely have to be altered to equal the bounding variance values instead of the currently specified 99.9% bound. The number of bits required and all other formatting remains unchanged.

The master station would combine its estimate of the noise terms and estimates of bias terms as though they were distributed according to (9). The bias terms would include the maximum deviations of errors which are observed by the master station but cannot be transmitted due to bandwidth constraints. It would then generate the final bounding variance estimates as in the upper part of (12) and transmit them to the user via the UDREs and GIVES. The users would take these values and combine them with their local receiver bounding variances to determine the weighting matrix \mathbf{W} according to

$$\sigma_i^2 = \sigma_{UDRE_i}^2 + F^2(El_i)\sigma_{UIVE_i}^2 + \sigma_{SNR_i}^2 + \frac{\sigma_{m45}^2}{\tan^2 El_i} + \frac{\sigma_{trv}^2}{\sin^2 El_i} \quad (26)$$

where $F(El_i)$ is the ionospheric obliquity factor, $\sigma_{SNR_i}^2$ contains the user receiver variance, σ_{m45}^2 is the multipath variance at 45°, and σ_{trv}^2 is the variance of the vertical tropospheric delay estimate. This final variance is used to weight the position solution in (15) and to verify the integrity of the user geometry and weights according to (24).

NSTB VERIFICATION

In order to test the integrity equation using real data we used the National Satellite Test Bed (NSTB) network together with the Stanford master station algorithms. The NSTB consists of 22 reference stations scattered across the CONUS and Canada (Figure 5). The data from these reference stations is sent to Stanford in real time. With our master station software we are able to designate which stations are used to generate the differential corrections. Other stations can be made passive. The passive stations are used to observe the quality of the corrections by comparing the NSTB

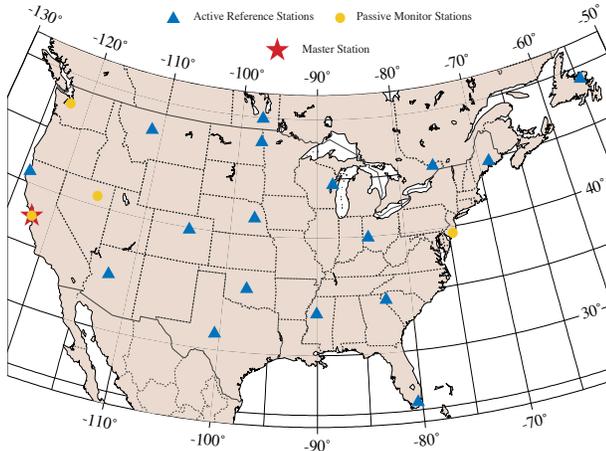


Figure 5. The location of NSTB differential reference stations, passive monitoring stations and the master station at Stanford University are shown.

calculated position against the surveyed location of the antenna.

The verification with real data tests two distinct but related processes: the ability of the master station algorithms to generate bounding UDREs and GIVEs, and the validity of the proposed integrity equation in the face of errors with unknown probability density functions. The analysis and simulations of previous sections exclusively tested the integrity equation. The real data primarily tests the ability of our master station to bound the errors with sufficient integrity.

Because the geometry and noise levels are constantly changing in the real data, we cannot generate the true $Pr\%$ confidence bounds that we used with the simulated data. Instead we divide the actual error by the square root of the variance bound (σ_v) at every epoch. The histogram of these data points is shown in Figure 6. The data is combined from four different reference stations, the FAA Technical Center in Atlantic City NJ, Elko NV, Seattle WA, and Stanford University in Stanford CA, and was collected over several days. If the UDREs, GIVEs and integrity equation are valid then the tail of the histogram will be bounded by a gaussian with unit variance (solid line shown for reference). As is evident in the figure the vertical error is well described by and protected by the bounding variance. In all cases the ratio of the error to the sigma was below 4.5. For more than 10^6 points a value closer to 4.9 would not be unreasonable. No error came close to exceeding the 5.33 σ_v integrity bound specified for precision approach

It should be noted that there were other monitor stations, which are still in the process of being evaluated, that had position errors not covered by the integrity equation. Upon investigation, we discovered that these

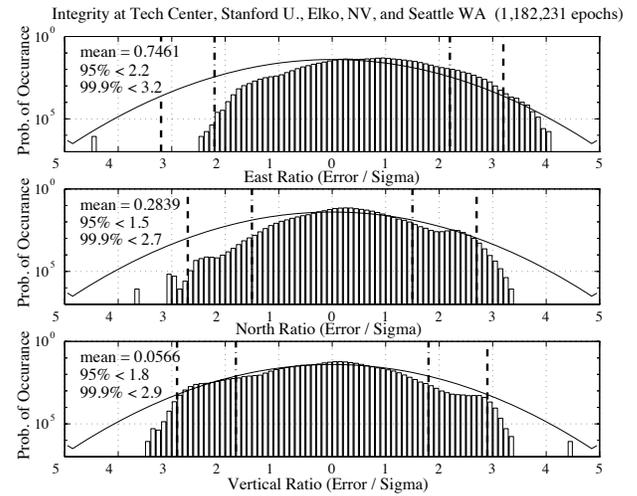


Figure 6. Histograms are shown for the actual error divided by the predicted error sigma for the East, North, and Vertical directions. The reference curve is for a zero mean gaussian with unit variance. It is evident that the tails of the actual distribution are better behaved than the expected distribution.

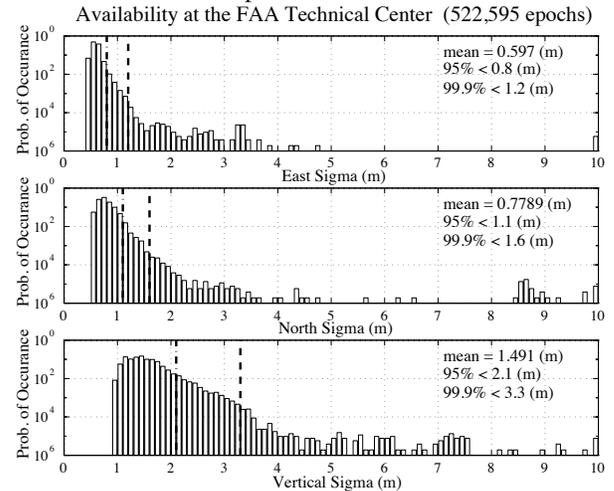


Figure 7. These histograms show the distribution of the predicted sigmas for the East, North, and Vertical directions. The system is declared available when the sigmas are below their maximum values. For 200' decision height, the maximum in the vertical direction is 3.6 meters

stations had experienced false locks in their satellite tracking. Again, it must be emphasized that neither WAAS nor this integrity equation can protect against these local fault modes. Such events must be captured through local fault detection methods independent of this integrity equation. In fact, these faults were identified by the monitor station's weighted Receiver Autonomous Integrity Monitoring (RAIM) algorithm [6].

The real data supports both the integrity equation and the ability of the master station to generate

appropriate bounding variances. We now wish to turn to the issue of availability. For a navigation sensor error bound of 19.2 meters at a probability level of better than 10^{-7} , the maximum allowable value of σ_v would be approximately 3.6 meters. If the value is larger then the system is declared unavailable. Figure 7 shows the histograms of the position sigmas for the FAA technical center. From this chart we can see that the availability of the system at this location was better than 99.9%. Because VPL_{absolute} does not fit within the MOPS message format, it was not tried. However, Figures 3 and 4 imply that had we employed that equation its availability would likely have been below 80%. Therefore we feel that our integrity equation is a practical one, capable of providing both sufficient integrity and sufficient availability.

CONCLUSIONS AND RECOMMENDATIONS

The analysis, the Monte Carlo results and the end-to-end check with real data strongly support VPL_{σ_v} as the clear choice for the integrity equation. For the observation error PDFs investigated, it maximizes availability without sacrificing integrity and is robust against non-gaussian distributions.

Our simulations only covered the tails out to 99.999%. In reality, predicting tail behavior requires extreme care. The master station and user algorithms have inexact knowledge of the actual variances of the observation errors. Thus, the values transmitted to the user and applied in the algorithms must be the bounding estimate for the variances. We must be extremely confident that the broadcast values are no smaller than the true values. Moreover, the variance must cover the position solution error to the Pr_{HMM} level. One must be certain that the tails of (9) cover the actual error to that level of probability. However, we also do not wish to grossly overestimate these values, as we wish to exploit their relative differences to weight the position solution (15).

Analysis of the integrity equation with real data provided us with more information about our UDRE and GIVE generation algorithms than about the integrity equation itself. This proof of concept is extremely important, yet there is still much work to be done towards improving the UDRE and GIVE generation algorithms. We need to improve our confidence in these algorithms to provide bounding values for future data. One possible source of concern is correlations between the error terms. However investigation of the tails of the data in Figure 6 reveals that these correlations appear to cancel one another

rather than add to each other. More research must be devoted to this area, but it is not a surprising result.

The analysis and simulation assumed a specific PDF for the observation errors in this analysis. This PDF, (9), is more pessimistic than the likely distribution of actual measurement errors. Yet (24) still provided full integrity and good availability. Further research can be done to verify that VPL_{σ_v} performs as well for other PDFs, but we already know that it performs exceedingly well against optimistic distributions (6), worst case distributions (9) and real distributions (Figure 6). The real data shows that the integrity equation works on PDFs likely to be encountered in the real world and that variance bounds can be generated despite imperfect knowledge of these PDFs. In addition this integrity equation is fully compliant with the existing WAAS MOPS.

REFERENCES

- [1] GNSSP Working Group A, *GNSS Performance Requirements*, Prepared for GNSSP meeting South Brisbane, February 20, 1997.
- [2] RTCA Special Committee 159 Working Group 2, "Wide Area Augmentation Signal Specification," September, 1995.
- [3] Walter, T., Enge, P., Van Graas, F., *Integrity for the Wide Area Augmentation System*, Proceedings of DSNS 95, Bergen, Norway, April, 1995, Vol II., pp 1-8.
- [4] FAA ANM 110 Advisory Circular, 25.1309, "System Design and Analysis," June, 1988.
- [5] Brown, R. G., *A Baseline GPS RAIM Scheme and a Note on the Equivalence of Three RAIM Methods*, **Navigation**, Vol. 39, No. 3, Fall 1992, pp.301-16.
- [6] Walter, T. and Enge P., *Weighted RAIM for Precision Approach*, proceedings of ION GPS-95, Palm Springs, CA, Sept., 1995, pp 1995-2004.
- [7] Axelrad, P. and Brown, R. G., *GPS Navigation Algorithms*, Chapter 9 in **Global Positioning System: Theory and Applications**, Washington D.C., published by AIAA, 1996.
- [8] Papoulis, A., **Probability, Random Variables, and Stochastic Processes**, Singapore, McGraw-Hill Book Co., 1984.