Sigma Inflation for the Local Area Augmentation of GPS

BORIS PERVAN
IRFAN SAYIM
Illinois Institute of Technology

The Local Area Augmentation System (LAAS) is the differential satellite navigation architectural standard for civil aircraft precision approach and landing. While the system promises great practical benefit, a number of key technical challenges have been encountered in the definition of the architecture. Perhaps chief among these has been the need to ensure compliance with stringent requirements for navigation integrity. In this context, this research investigates the sensitivity of integrity risk to statistical uncertainties in the knowledge of reference receiver error standard deviation (σ_{pr_gnd}) and error correlation across the multiple reference receivers to be used in the LAAS ground segment. A new, detailed approach toward mitigating the integrity risk due to parameter statistical uncertainty is presented.

Manuscript received May 30, 2000; revised March 30 and June 7, 2001; released for publication July 2, 2001.

IEEE Log No. T-AES/37/4/10993.

Refereeing of this contribution was handled by J. L. Leva.

This work was supported by the Federal Aviation Administration.

Authors' address: Department of Mechanical and Aerospace Engineering, Illinois Institute of Technology, Engineering 1 Building, 10 West 32nd Street, Chicago, IL 60616-3793.

INTRODUCTION

In the Local Area Augmentation System (LAAS), the final quantitative assessment of navigation integrity is realized through the computation of vertical and horizontal protection levels at the aircraft (termed VPL and HPL, respectively). [1] In principle, these limits are the position bounds that can be ensured with an acceptable level of integrity risk. For example, for a Category 1 approach, the maximum permissible integrity risk is of the order of 10^{-8} with respect to a vertical alert limit (VAL) of 10 m [2]. The prescribed algorithms for the generation of the protection limits presume a normally distributed fault-free error model for the broadcast pseudo-range corrections. The standard deviation of correction error is further assumed by the aircraft to be equal to the broadcast value of σ_{pr_gnd} for each satellite. It is clear that to ensure navigation integrity, special care must be taken on the ground in the definition of broadcast σ_{pr_gnd}. In particular, the finite test sample sizes generally available to compute error standard deviation and the correlation of errors between multiple reference receivers (whose measurements are averaged to generate the broadcast correction) must be accounted for in the definition of σ_{pr_gnd}. The zero-mean Gaussian model is consistent with reference receiver ranging errors due to thermal noise and diffuse multipath. Other errors, such as ground reflection multipath, can slowly vary with environmental conditions making the underlying distribution difficult to characterize by experimental means alone. A number of theoretical approaches addressing such errors are covered in a separate paper [3]. For the purposes of the present work, errors are assumed to be zero-mean and Gaussian.

This paper investigates the sensitivity of integrity risk to statistical uncertainty in the knowledge of the correction error standard deviation (σ_{pr_gnd}) and error correlation between multiple reference receivers. A general approach toward mitigating the integrity risk due to parameter statistical uncertainty is presented.

LAAS INTEGRITY

The basic function of the LAAS ground facility (LGF) integrity monitoring system is the detection and removal of anomalies present in the LAAS signal-in-space (SIS) that would otherwise result in an unacceptable integrity risk to an aircraft on final approach. The notion of SIS is introduced primarily to distribute accountability between the ground and airborne navigation subsystems. In general, the aircraft is responsible for the proper functionality of the airborne equipment (which would typically include the implementation of redundant sensor tracks to provide the means for detection and removal of airborne equipment failures), while the LGF is

responsible for the detection of anomalies in both the received satellite signals and the LAAS reference data broadcast to the aircraft. The satellite signals and broadcast reference data collectively define the LAAS SIS.

As currently envisioned, LAAS SIS integrity monitoring is comprised of both ground and airborne elements. The need for an airborne processing component, even for SIS monitoring, is motivated by the fact that the integrity specifications are expressed in the *position* (rather than *range*) domain. Because the LGF is generally unaware of the specific satellites being tracked by the airborne receiver at any given time, an airborne processing component is implemented specifically to convert ground-broadcast range domain statistics to position domain protection levels. More detail on the airborne processing may be found in [2, 4].

SIGMA SENSITIVITY ANALYSIS

In the LAAS architecture, and in this analysis, integrity risk under the hypotheses of fault-free conditions (H0) and integrity risk in the event of a single reference receiver failure (H1) are considered separately. (The likelihood of simultaneous failures on multiple reference receivers is required to be negligibly small by design specification.) Nominally, the vertical protection limits $\mathrm{VPL}_{\mathrm{H0}}$ and $\mathrm{VPL}_{\mathrm{H1}}$ are computed at the aircraft based on values of broadcast correction error standard deviation ($\sigma_{\mathrm{pr\_gnd}}$) for each satellite also broadcast by the reference station. In addition, the prescribed computation of $\mathrm{VPL}_{\mathrm{H1}}$ requires that the ground broadcast differences between the pseudo-range corrections derived from various subsets of the multiple (typically 3 or 4) LGF reference receivers. The precise mathematical structure of these differences, termed B-values, is defined in the LGF System Specification [5]. (In contrast, the nominal correction broadcast for each satellite, which is used for positioning but not in the computation of protection limits, is based on an average across *all* reference receivers.) The prescribed missed detection (MD) probabilities for H0 and H1 are specified, respectively, in terms of Gaussian multipliers $k_{\mathrm{md\_ff}}$ and k_{md}, which are defined below.

The general approach taken in this analysis is to first quantify 'true' MD probability given that the actual value of reference receiver error standard deviation (σ) deviates from $\sqrt{M}$ times the broadcast value $\sigma_{\mathrm{pr\_gnd}}$. (M is the number of reference receivers used to generate the broadcast correction.) Since it is recognized that any realizable estimate of standard deviation will be based on a finite number of error samples, it is then also necessary to ensure that the broadcast value of $\sigma_{\mathrm{pr\_gnd}}$ accounts for any statistical uncertainty that may lead to increased integrity risk.

H1 Case

The vertical protection limit under the hypothesis of a failure on any given reference receiver ($\mathrm{VPL}_{\mathrm{H1}}$) is given by the following expression derived in [2, 4]:

$$\mathrm{VPL}_{\mathrm{H1}} = \left|\sum_{n=1}^{N} S_{zn} B_n\right| + k_{md}\sqrt{\sum_{n=1}^{N} S_{zn}^2\left[\frac{\sigma^2_{\mathrm{pr\_gnd}_1}(n)}{M(n)-1} + \sigma^2_{\mathrm{pr\_air}}(n) + \sigma^2_{\mathrm{pr\_res}}(n)\right]} \tag{1}$$

where

n is the satellite index;

S_{zn} is nth element of the third row (representing the vertical component) of the weighted geometry projection matrix used to generate the position estimate;

B_n is the broadcast B-value for satellite n associated with the given reference receiver;

M is the number of reference receivers used to generate the broadcast correction;

N is the number of available satellites;

k_{md} is a multiplier used to set the desired level of MD probability assuming Gaussian errors. k_{md} holds a value of 2.898 for a Category 1 approach with three reference receivers [2];

$\sigma_{\mathrm{pr\_air}}$ is the airborne measurement error standard deviation;

$\sigma_{\mathrm{pr\_res}}$ is the standard deviation of residual errors not directly attributable to ground or airborne error (such as ionospheric and tropospheric decorrelation);

$\sigma_{\mathrm{pr\_gnd}_1} \equiv \sqrt{M}\sigma_{\mathrm{pr\_gnd}}$.

The maximum acceptable values of the standard deviations $\sigma_{\mathrm{pr\_gnd}}$, $\sigma_{\mathrm{pr\_air}}$, $\sigma_{\mathrm{pr\_res}}$, are functions of satellite elevation given in [2]. In this analysis, we first assume a Category 1 system with a class B3 ground facility ($M = 3$) and B class airborne equipment as defined in [2]. Although we explicitly consider only variations in $\sigma_{\mathrm{pr\_gnd}}$, it should be noted that the method of analysis described below is in principle applicable to airborne and residual errors as well.

When the actual ground error standard deviation (σ) differs from the nominal value ($\sigma_{\mathrm{pr\_gnd}_1}$) used to generate $\mathrm{VPL}_{\mathrm{H1}}$, the effective MD multiplier for the computed value of $\mathrm{VPL}_{\mathrm{H1}}$ is given by

$$k_{md_e} \equiv k_{md}\frac{\sqrt{\sum_{n=1}^{N} S_{zn}^2\left[\frac{\sigma^2_{\mathrm{pr\_gnd}_1}(n)}{M(n)-1} + \sigma^2_{\mathrm{pr\_air}}(n) + \sigma^2_{\mathrm{pr\_res}}(n)\right]}}{\sqrt{\sum_{n=1}^{N} S_{zn}^2\left[\frac{\sigma^2(n)}{M(n)-1} + \sigma^2_{\mathrm{pr\_air}}(n) + \sigma^2_{\mathrm{pr\_res}}(n)\right]}}. \tag{2}$$

Note that the B-value term is not present since it is invariant with respect to changes in σ. The associated MD probability is then

$$P_{H1}\{\text{MD} \mid \sigma(1)\ldots\sigma(N)\} = Q(k_{md_e}) \qquad (3)$$

where the function $Q(x)$ is defined as the area to the right of x under a standard normal density function (i.e., the tail probability).

Clearly $P_{H1}\{\text{MD} \mid \sigma(1)\ldots\sigma(N)\}$ will in general be a strong function of satellite geometry through the projection matrix S. In this regard, a GPS constellation simulation was executed to establish this sensitivity. The details of the geometry simulation follow.

Constellation: Nominal 24 satellite (SV) constellation defined in RTCA DO-229A [6].

Elevation Mask: 5 deg.

Simulated Duration: 24 h.

LGF Location: Chicago O'Hare International Airport.

SV Outage Conditions: Both the complete 24 SV constellation and worst case (most sensitive) 22 SV constellation subsets were simulated.

Geometries not meeting $\text{VPL}_{H0} < \text{VAL}$ using the nominal value of $\sigma_{\text{pr\_gnd}}$ were excluded since approaches would not be conducted in these cases.

In the first set of simulations, all 24 SVs were assumed to be usable, and the true standard deviation σ was varied in (2) on all visible SVs by simultaneously scaling the nominal value of $\sigma_{\text{pr\_gnd}_1}$ for each satellite by the same factor. The scale factor $(\sigma/\sigma_{\text{pr\_gnd}_1})$ was varied from 0.5 to 2, and the resulting integrity risk was computed from (3) for each value of $\sigma/\sigma_{\text{pr\_gnd}_1}$. The integrity risk results are shown in Fig. 1 as a function of $\sigma/\sigma_{\text{pr\_gnd}_1}$. The discrete distribution of data points along the horizontal $(\sigma/\sigma_{\text{pr\_gnd}_1})$ direction in the figure corresponds to the discrete values of $\sigma/\sigma_{\text{pr\_gnd}_1}$ simulated. The vertical distribution of data points at each value of $\sigma/\sigma_{\text{pr\_gnd}_1}$ is due to the varying geometries accumulated over a 24 h period. The upper bound integrity risk curve (solid) represents the highest level of integrity risk over the 24 h duration. Note that when $\sigma/\sigma_{\text{pr\_gnd}_1} = 1$, the MD probability attains a nominal value of $Q(k_{md}) = 0.0019$.

Given that all 24 satellites are available, the results in Fig. 1 are undoubtedly conservative since it is unlikely that broadcast $\sigma_{\text{pr\_gnd}}$ would underestimate the true σ for *all* visible satellites. In this regard, a second simulation was performed varying σ on only one (the most integrity-risk-sensitive) satellite for each geometry. The results are shown in Fig. 2. When compared with the results of Fig. 1, integrity risk is reduced for values of $\sigma/\sigma_{\text{pr\_gnd}_1} > 1$ (as expected) but increased for values of $\sigma/\sigma_{\text{pr\_gnd}_1} < 1$. The latter increase is due to the fact that σ is reduced on only one SV (in contrast with Fig. 1 where σ was reduced on all SVs). Fig. 3, which superposes the upper bound curves from Figs. 1 and 2, clearly shows the difference in integrity risk sensitivity under the two sets of assumptions.

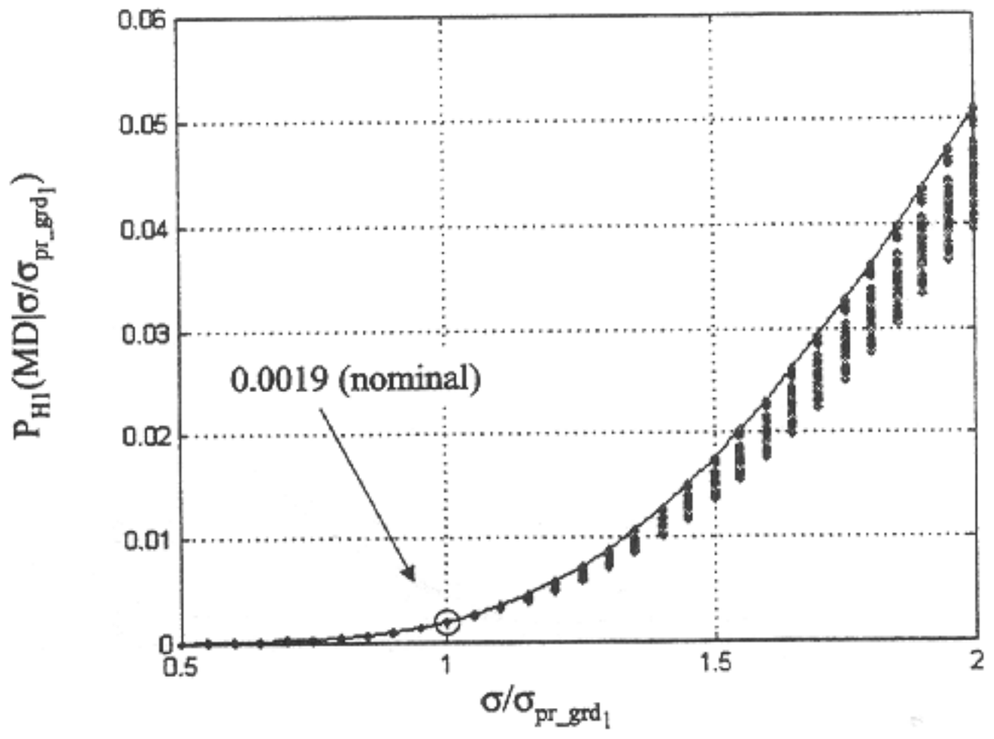

Fig. 1. H1 integrity risk sensitivity to σ-variations, 24 SV case. σ varied on all SVs in view.

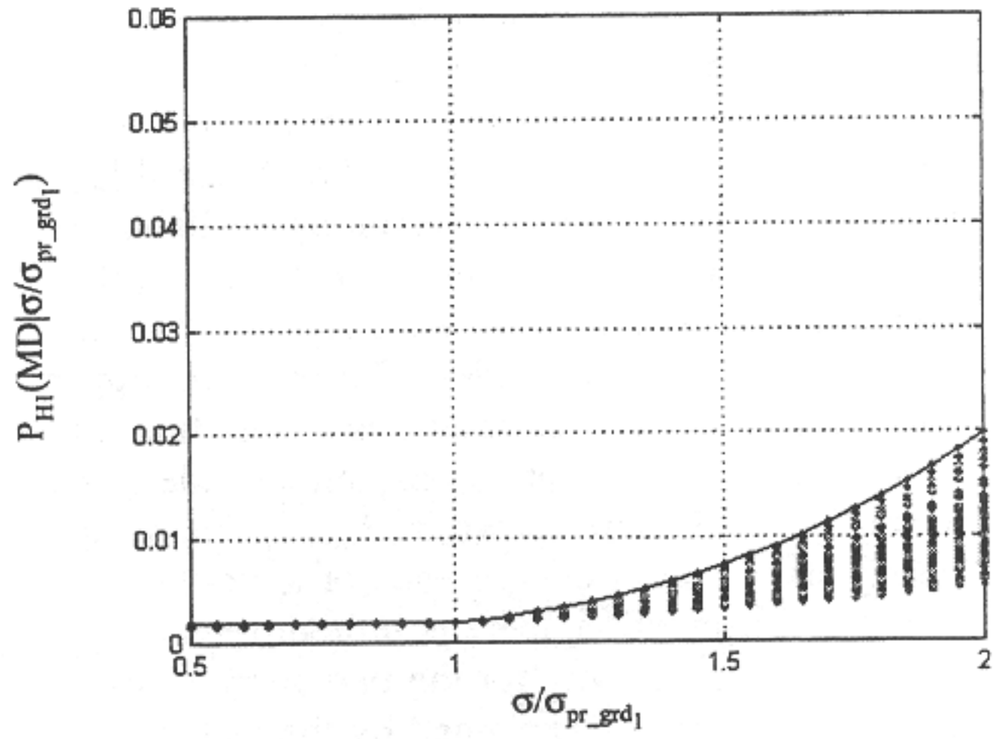

Fig. 2. H1 integrity risk sensitivity to σ-variations, 24 SV case. σ varied on worst case SV only.

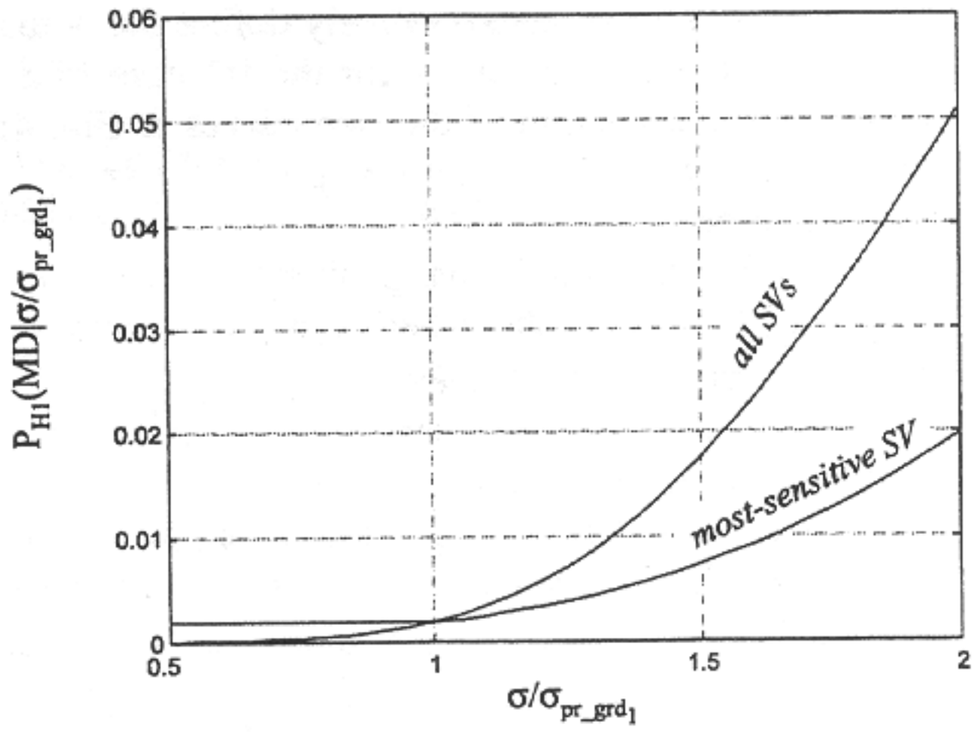

Fig. 3. H1 integrity risk sensitivity to σ-variations, 24 SV case. Upper bound curves from Figs. 1 and 2.

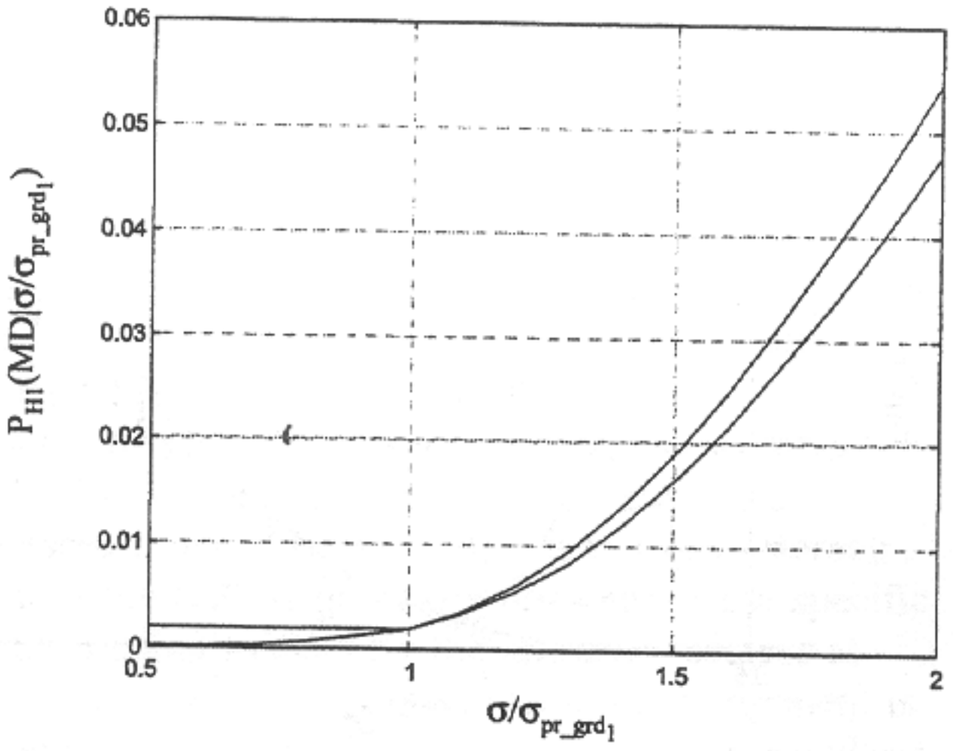

Fig. 4. H1 integrity risk sensitivity to σ-variations, 22 SV case. Upper bound curves.

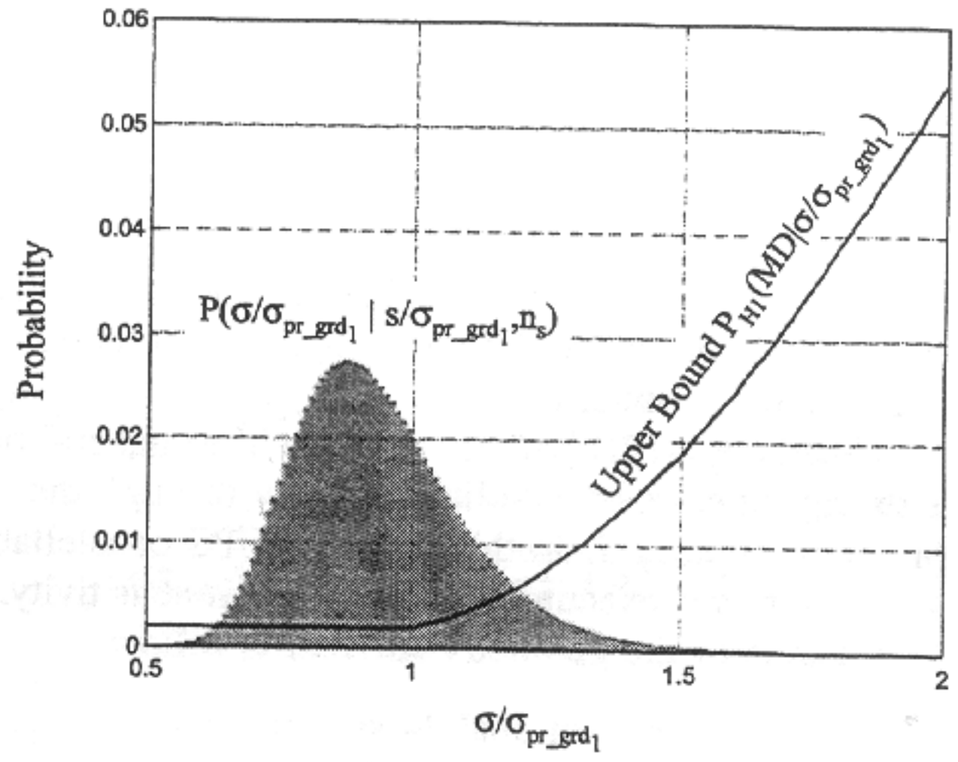

Fig. 5. Probability distribution for $\sigma/\sigma_{pr\_grd_1} = 0.9$, $n_s = 20$.

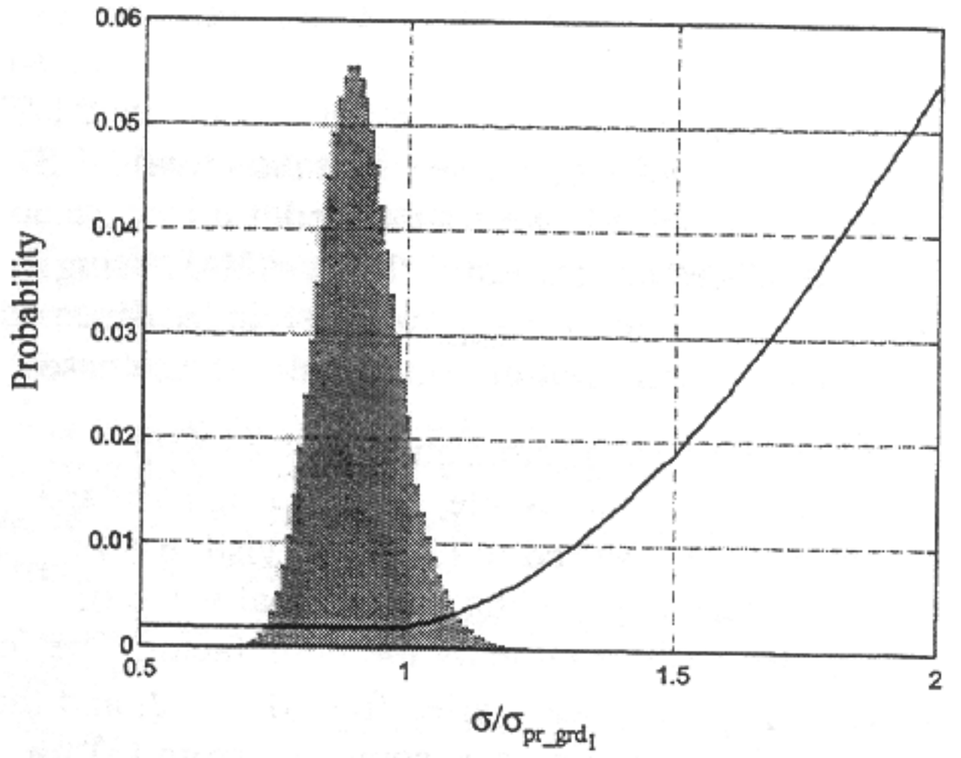

Fig. 6. Probability distribution for $\sigma/\sigma_{pr\_grd_1}$ for $s/\sigma_{pr\_grd_1} = 0.9$, $n_s = 80$.

In general, however, it cannot be assumed that all 24 satellites will always be available for use. For example, existing simulation results of LAAS operational availability in the LAAS standard [2] are based on worst case (lowest resulting availability) 22 satellite subset geometries. In this context, the simulations executed above were repeated for all 22-satellite subset geometries. The resulting upper bound sensitivity curves assuming σ variation on all satellites and σ variation only on a single (most sensitive) satellite are shown in Fig. 4. The results clearly show that in the presence of a modestly depleted constellation, there is little difference in integrity risk sensitivity for the two approaches. This result is readily explained by the fact that when fewer satellites are available the effect in the position domain of error variations on individual geometry-critical satellites is more pronounced. For our analysis, we can conservatively define the actual integrity risk sensitivity curve for the H1 case as a piecewise superposition of the two curves in Fig. 4; for any value of $\sigma/\sigma_{pr\_gnd_1}$, the upper of the two curves is used.

Given that the conditional probability $P_{H1}(MD \mid \sigma/\sigma_{pr\_gnd_1})$ has been established, it is still necessary to define a probability density for $\sigma/\sigma_{pr\_gnd_1}$ so that the overall risk probability can be quantified. In this regard, given n_s independent measurements derived from a Gaussian distribution with a sample variance s^2, Box and Tiao [7] show that the probability density function for σ is

$$p(\sigma \mid s, n_s) = \left[\frac{1}{2}\Gamma\left(\frac{n_s}{2}\right)\right]^{-1} \left(\frac{n_s s}{2}\right)^{n_s/2} \sigma^{-(n_s+1)} \exp\left(\frac{n_s s^2}{2\sigma^2}\right). \tag{4}$$

The derivation of (4) uses Bayes' Theorem with a noninformative prior distribution for $p(\sigma)$. (It is noted that this derivation and other principles of Bayesian statistical inference have not gained universal acceptance among statisticians.) For a given computed value of s, the probability that σ lies in any specified interval may be computed by the integration of (4) over the interval. For example, Fig. 5 shows the resulting probability mass function $P(\sigma/\sigma_{pr\_gnd_1} \mid s/\sigma_{pr\_gnd_1}, n_s)$ for $\sigma/\sigma_{pr\_gnd_1}$ intervals of width 0.01 and $s/\sigma_{pr\_gnd_1} = 0.9$ and $n_s = 20$. The result is plotted together with the conditional $P_{H1}(MD \mid \sigma/\sigma_{pr\_gnd_1})$ curve already established. Clearly, despite the fact that s is lower than $\sigma_{pr\_gnd_1}$ for this case, the likelihood that the actual value of σ exceeds $\sigma_{pr\_gnd_1}$ is nonnegligible. As the number of available samples n_s is increased, however, the likelihood that σ exceeds σ_{pr_gnd} decreases. Fig. 6 illustrates the case where $n_s = 80$. Similarly, if the computed value of s is lower, the likelihood that σ exceeds $\sigma_{pr\_gnd_1}$ is also lower. Fig. 7 shows the case where $n_s = 20$ and $s/\sigma_{pr\_gnd_1} = 0.7$.

A parametric analysis was performed in which n_s was varied with discrete values 20, 50, 100, and 200, and $s/\sigma_{pr\_gnd_1}$ between 0.7 and 1.3 (in increments of 0.01). The overall H1 MD probability given $s/\sigma_{pr\_gnd_1}$

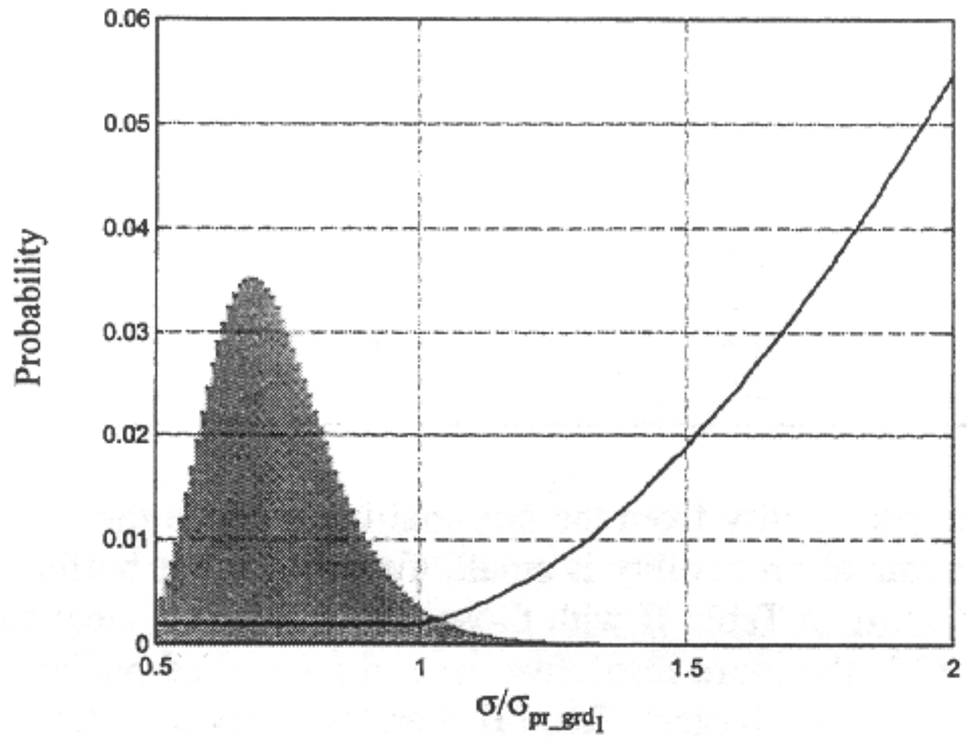

Fig. 7. Probability distribution for $\sigma/\sigma_{\text{pr\_grd}_1}$ for $s/\sigma_{\text{pr\_grd}_1} = 0.7$, $n_s = 20$.

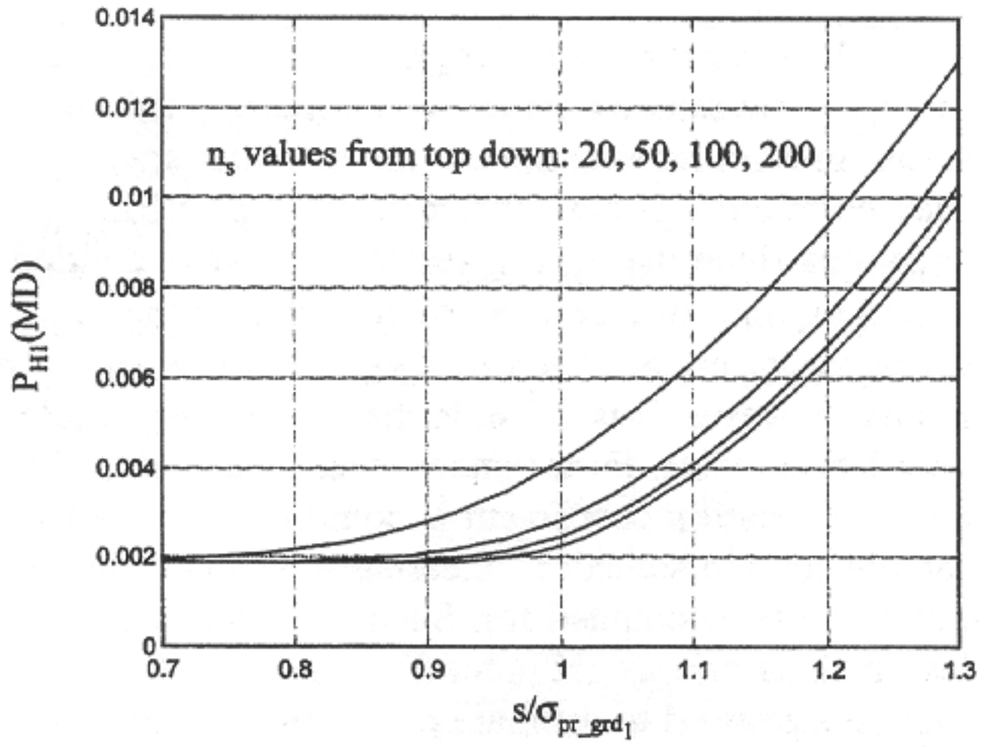

Fig. 8. H1 integrity risk versus $s/\sigma_{\text{pr\_grd}_1}$.

and n_s was then computed numerically via

$$P_{\text{H1}}(\text{MD} \mid s/\sigma_{\text{pr\_gnd}_1}, n_s) = \sum_{\sigma} P_{\text{H1}}(\text{MD} \mid \sigma/\sigma_{\text{pr\_gnd}_1}) \times P(\sigma/\sigma_{\text{pr\_gnd}_1} \mid s/\sigma_{\text{pr\_gnd}_1}, n_s). \quad (5)$$

The results are plotted in Fig. 8, which shows quantitatively how the H1 MD probability increases as $s/\sigma_{\text{pr\_gnd}_1}$ increases and as n_s decreases. The results are plotted in terms of percentage error (above the nominal value of 0.0019) in Fig. 9.

To ensure integrity in an absolute sense $P_{\text{H1}}(\text{MD} \mid s/\sigma_{\text{pr\_gnd}_1}, n_s)$ should not exceed the nominal specified value of 0.0019. However, the results in Fig. 9 show that this criterion cannot be realistically attained because an infinitely large sample set is required. Nevertheless, it can be ensured that the MD probability does not differ from the nominal value by a significant amount. For example, Table I summarizes the results obtained from Fig. 9 assuming that a 5% tolerance is acceptable. Under this assumption, the

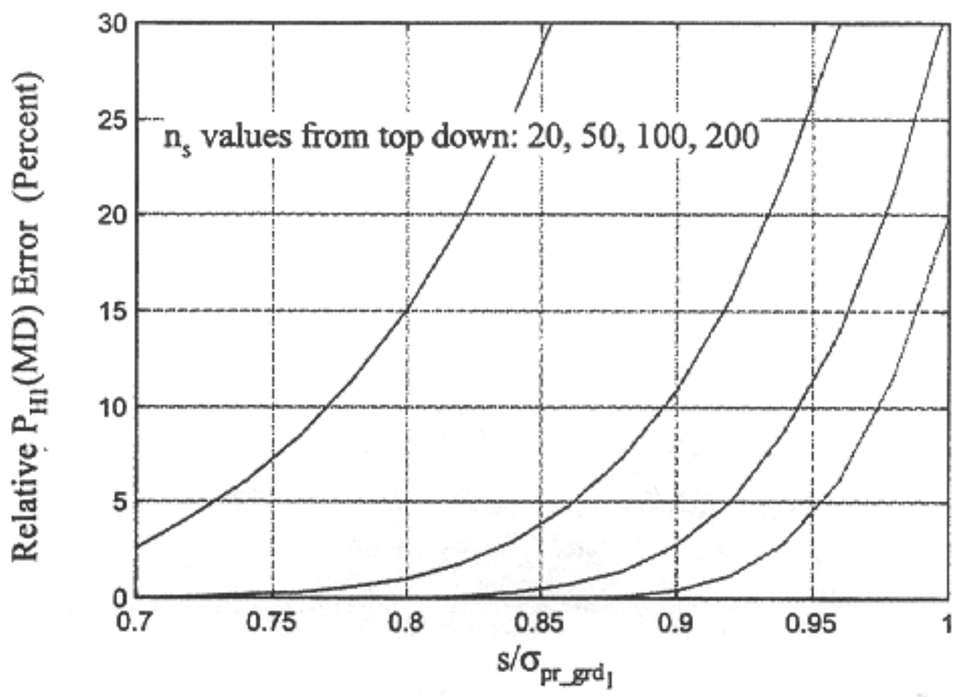

Fig. 9. H1 integrity risk versus $s/\sigma_{\text{pr\_grd}_1}$ (relative error).

TABLE I
Sigma Buffer Factor for H1

n_s	Minimum Value of $\sigma_{\text{pr\_gnd}_1}$
50	$1.16 \times s$
100	$1.09 \times s$
200	$1.05 \times s$
500	$1.02 \times s$

table quantifies the minimum value of $\sigma_{\text{pr\_gnd}_1}$ that may be broadcast given any value of s obtained from n_s samples. Clearly, the broadcast $\sigma_{\text{pr\_gnd}_1}$ must in general be larger than s. As expected, however, the buffer factor (the amount by which s must be scaled to define $\sigma_{\text{pr\_gnd}_1}$) approaches 1 as n_s grows large. (It should be noted that the quantitative results in Table I apply for the H1 case only.)

H0 Case

For the fault-free hypothesis, the vertical protection limit is given by the following expression [2, 4]:

$$\text{VPL}_{\text{H0}} = k_{\text{ff\_md}} \sqrt{\sum_{n=1}^{N} S_{zn}^2 \left[\frac{\sigma_{\text{pr\_gnd}_1}^2(n)}{M(n)} + \sigma_{\text{pr\_air}}^2(n) + \sigma_{\text{pr\_res}}^2(n) \right]}. \quad (6)$$

When the actual ground error standard deviation (σ) differs from the nominal value ($\sigma_{\text{pr\_gnd}_1}$) used to generate VPL_{H0}, the effective MD multiplier for the computed value of VPL_{H0} is given by

$$k_{\text{ff\_md}_e} \equiv k_{\text{ff\_md}} \frac{\sqrt{\sum_{n=1}^{N} S_{zn}^2 \left[\frac{\sigma_{\text{pr\_gnd}_1}^2(n)}{M(n)} + \sigma_{\text{pr\_air}}^2(n) + \sigma_{\text{pr\_res}}^2(n) \right]}}{\sqrt{\sum_{n=1}^{N} S_{zn}^2 \left[\frac{\sigma^2(n)}{M(n)} + \sigma_{\text{pr\_air}}^2(n) + \sigma_{\text{pr\_res}}^2(n) \right]}} \quad (7)$$

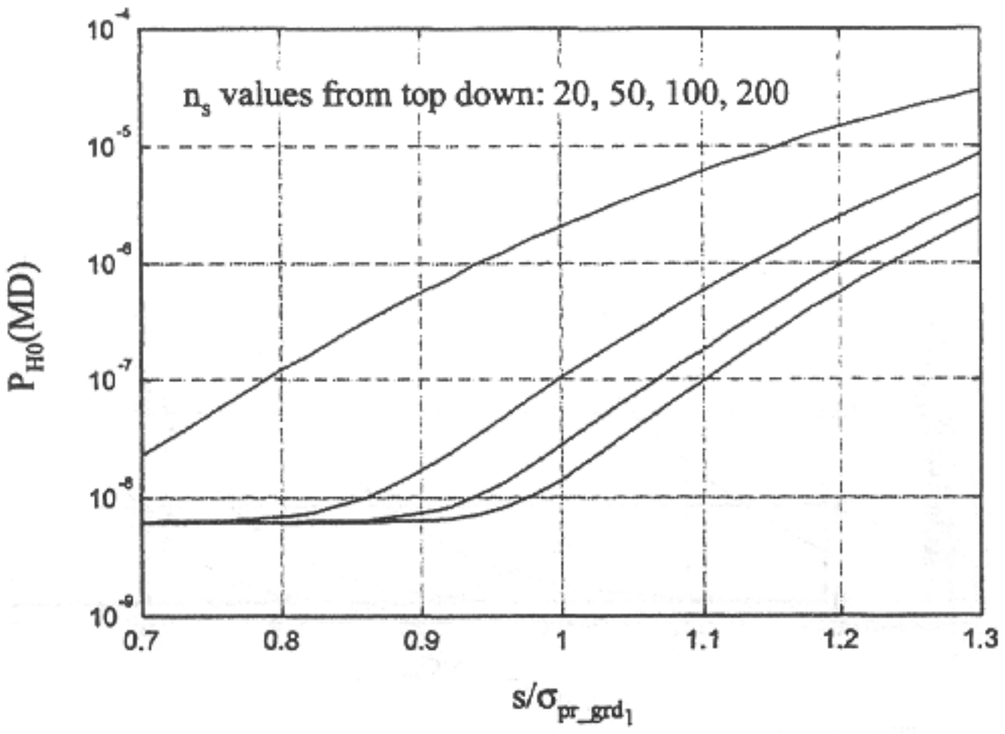

Fig. 10. H0 integrity risk versus $s/\sigma_{\text{pr\_grd}_1}$.

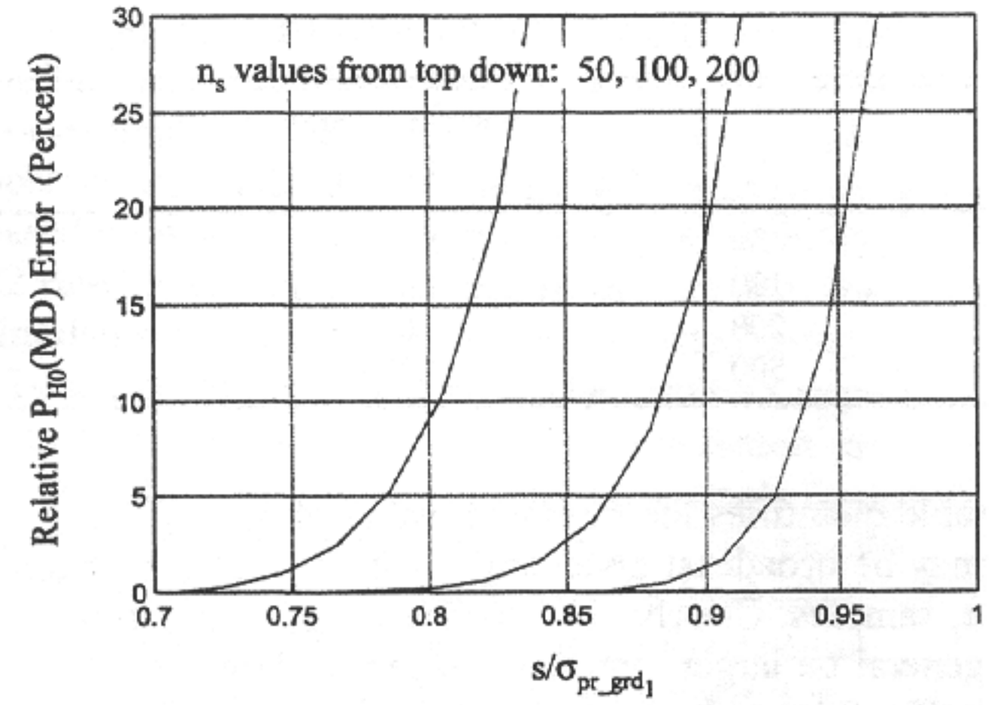

Fig. 11. H0 integrity risk versus $s/\sigma_{\text{pr\_grd}_1}$ (relative error).

where $k_{\text{ff\_md}}$ is 5.810 for Category 1 approach with three reference receivers [2]. The associated MD probability is

$$P_{\text{H0}}\{\text{MD} \mid \sigma(1)\ldots\sigma(N)\} = 2Q(k_{\text{ff\_md}_e}). \tag{8}$$

The factor of 2 in (8) is present for the H0 case because the underlying fault-free distribution has a zero mean, and therefore both positive and negative errors are of interest. The sensitivity analysis executed for the H1 case was repeated for H0. The resulting sensitivity of MD probability for the H0 case is quantified in Fig. 10. As with the H1 case, $P_{\text{H0}}(\text{MD} \mid s/\sigma_{\text{pr\_gnd}_1}, n_s)$ increases as $s/\sigma_{\text{pr\_gnd}_1}$ increases and as n_s decreases. In Fig. 11, the results of Fig. 10 are plotted in terms of percentage error (relative to the nominal value of $2Q(k_{\text{ff\_md}}) = 6.2 \times 10^{-9}$).

The results of the H0 σ-sensitivity analysis are summarized in Table II, where a 5% MD tolerance (relative to nominal) has again been used. As with the H1 case, the buffer factor approaches 1 as n_s becomes very large. Note that, in general, H0 σ-sensitivity is greater than that found for the H1 case. This is due to the fact that a small variation in σ for a Gaussian random variable will cause a larger relative deviation

TABLE II
Sigma Buffer Factor for H0

n_s	Minimum Value of $\sigma_{\text{pr\_gnd}_1}$
50	$1.29 \times s$
100	$1.16 \times s$
200	$1.09 \times s$
500	$1.04 \times s$

in probability from the nominal value when the nominal probability is small. Comparing the buffer factors in Table II with those in Table I, it is clear that H0 is the more restrictive case (the required buffer factors are larger). Table II therefore serves to define the minimum σ buffer factor.

CORRELATION SENSITIVITY ANALYSIS

In the preceding analysis, it was implicitly assumed that ranging errors were uncorrelated across ground receivers. Note, in fact, that any such correlation is not consistent with the VPL equations since the $\sigma^2_{\text{pr\_gnd}_1}$ terms are always divided by the number of receivers (to account for the averaging of uncorrelated receiver measurements). In reality, however, it is possible that some measurable correlation exists. Furthermore, even if a negligibly small correlation coefficient is computed from a finite sample set, the statistical uncertainty in the estimate must also be accounted for. Such uncertainty is lessened, as one would naturally expect, as the sample size used to estimate correlation coefficient increases.

To examine sensitivity to correlation, we assume that the ground error standard deviation for any given reference receiver is $\sigma_{\text{pr\_gnd}_1}$. The effect of positive correlation between receivers when averaging M^* errors can be (with modest conservatism) modeled as an effective increase in σ as follows:

$$\sigma = \sigma_{\text{pr\_gnd}_1}\sqrt{1 + (M^* - 1)\rho} \tag{9}$$

where ρ is the maximum correlation coefficient between any pair of receivers, $M^* = M$ for H0, and $M^* = M - 1$ for H1. The sigma sensitivity analysis results (in particular, the satellite geometry simulations) are directly applicable to correlation sensitivity as well through the following simple transformation:

$$\rho = \frac{(\sigma/\sigma_{\text{pr\_gnd}_1})^2 - 1}{M^* - 1}. \tag{10}$$

H1 Case

Using the transformation above, the horizontal axis of Fig. 4 may be rescaled in terms of ρ. The resulting upper bound curve for the conditional probability

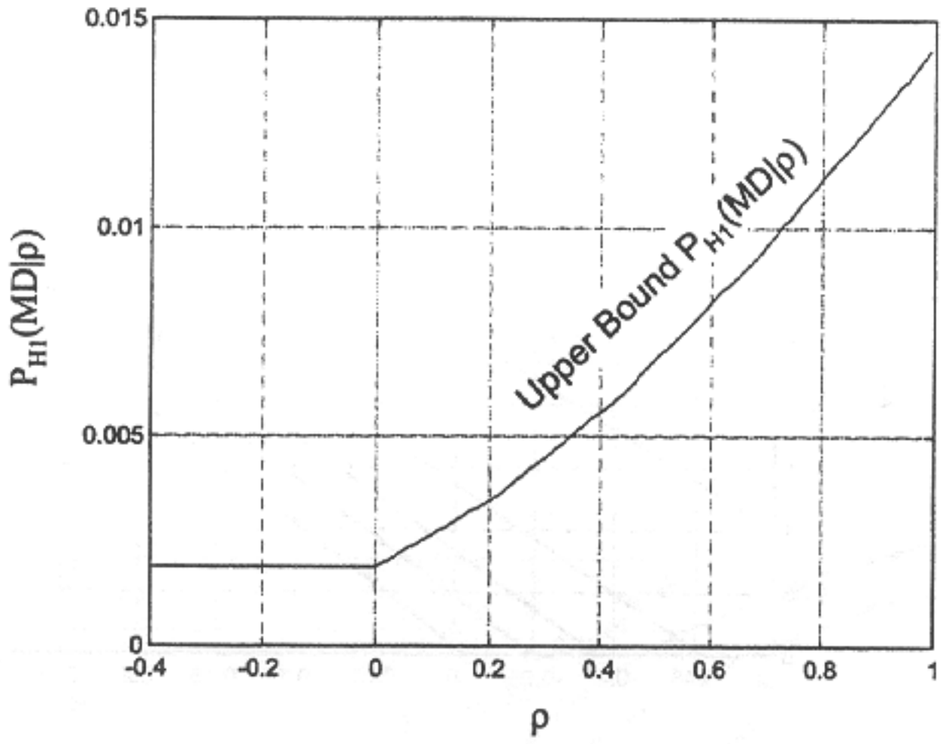

Fig. 12. H1 integrity risk sensitivity to ρ-variations, 22 SV case. Upper bound curve.

$P_{H1}(MD \mid \rho)$ is shown in Fig. 12. Now given any pair of reference receivers, each with n_r samples of measurement error, we may compute a sample correlation coefficient r. To define a distribution for ρ given r, we use a similar approach to that in the sigma analysis except that the chi-square distribution no longer applies. In this case, however, we make use of the Fisher Z statistic [8] which is approximately Gaussian:

$$Z \equiv \tfrac{1}{2}\ln\left(\frac{1+r}{1-r}\right) \sim N_z\left[\tfrac{1}{2}\ln\left(\frac{1+\rho}{1-\rho}\right),(n_r-3)^{-1/2}\right]. \tag{11}$$

Fig. 13 illustrates an example probability mass function $P(\rho \mid r, n_r)$ for $r = 0.3$, $n_r = 20$ and finely spaced intervals of ρ. (The $P_{H1}(MD \mid \rho)$ curve from Fig. 12 is overlayed.) A parametric analysis was then executed in which n_r varied with discrete values 20, 50, 100, 200, and 1000 and r varied between -0.2 and 0.3. The overall H1 MD probability given ρ and n_r was then computed numerically using a summation procedure equivalent to the sigma sensitivity case. The results are plotted in Fig. 14, which shows quantitatively how the H1 MD probability increases as r increases and as n_r decreases.

When compared with the nominal MD probability for H1 (0.0019), substantial increases in relative integrity risk are clearly evident in the results. However, this is a not unexpected result because the VPL equations have no direct means to accommodate the effect of positive correlation, and furthermore, integrity risk is magnified by uncertainty in the correlation coefficient. Note, however, that in principle correlation can be accounted for by simply increasing the value of $\sigma_{\text{pr\_gnd}_1}$. For example for precisely known values of error standard deviation (σ) and correlation coefficient (ρ) we may use any value of $\sigma_{\text{pr\_gnd}_1}$ such that

$$\sigma_{\text{pr\_gnd}_1} > \sigma\sqrt{1+(M^*-1)\rho}. \tag{12}$$

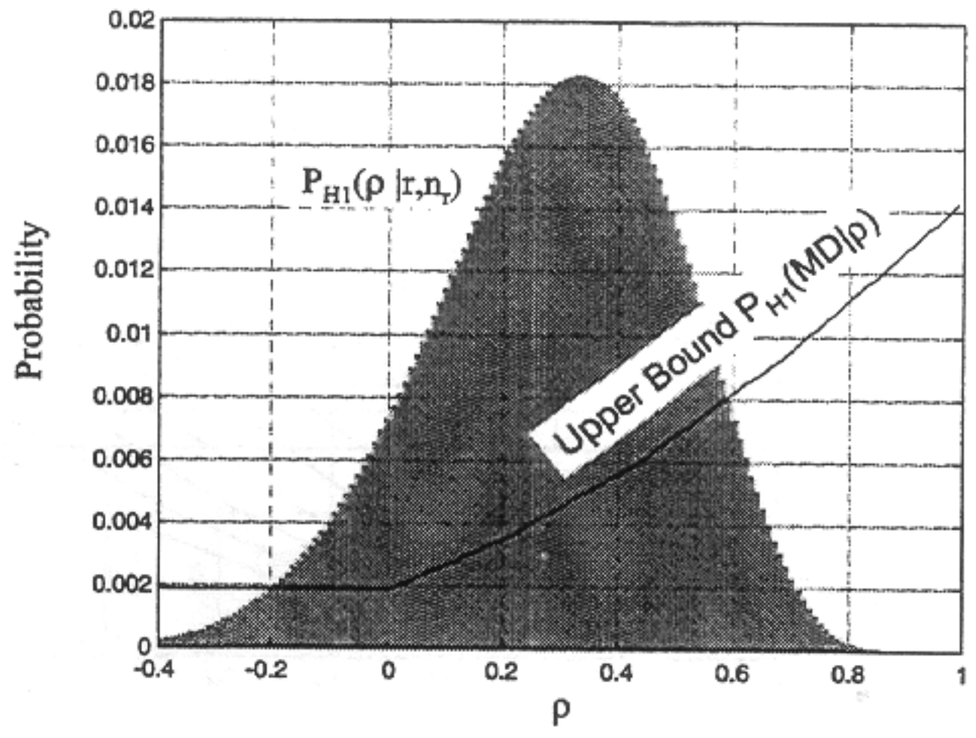

Fig. 13. Probability distribution for ρ given $r = 0.3$, $n_r = 20$.

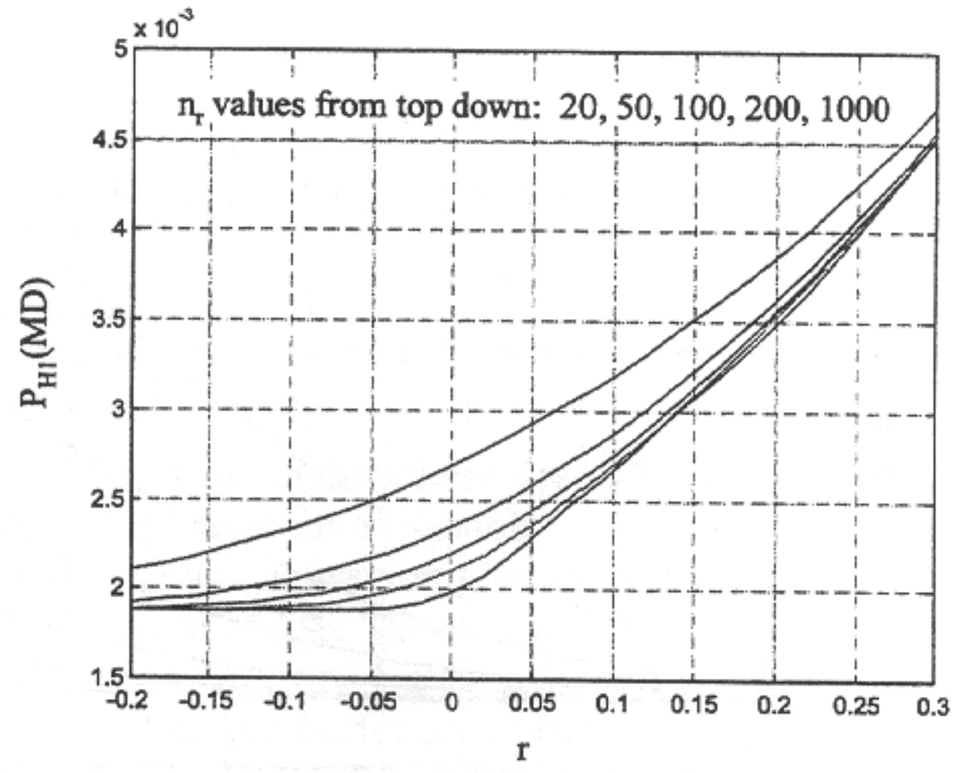

Fig. 14. H1 integrity risk versus r.

However, to account for the fact that σ is not known precisely (only the estimate s based on n_s samples is available), σ in the equation above must be replaced with $a(n_s)s$, where $a(n_s)$ is the scale factor defined in Table II. Similarly, because ρ cannot be known precisely (only the correlation coefficient estimate r based on n_r samples is available), ρ must be replaced in the equation above with a buffered value ρ^*, where ρ^* is a function of r and n_r that is yet to be defined.

To determine the required value of ρ^* for a given number or samples n_r, the conditional probability $P(MD \mid \rho)$ was recomputed assuming that the value of $\sigma_{\text{pr\_gnd}_1}$ has already been buffered using the above equation for selected values of ρ^* between 0 and 0.5. The resulting curves are given in Fig. 15. Note that the curve corresponding to $\rho^* = 0$ (which represents the case where there is no buffering on $\sigma_{\text{pr\_gnd}_1}$) is identical to the $P(MD \mid \rho)$ curve in Fig. 12. As expected, the influence of non-zero correlation coefficient on integrity risk is decreased as $\rho*$ is increased (i.e., as the buffer on $\sigma_{\text{pr\_gnd}_1}$ is increased).

For each of the ρ^* curves in Fig. 15, it is possible to compute $P_{H1}(MD \mid r, n_r)$ as was done for the $\rho^* = 0$

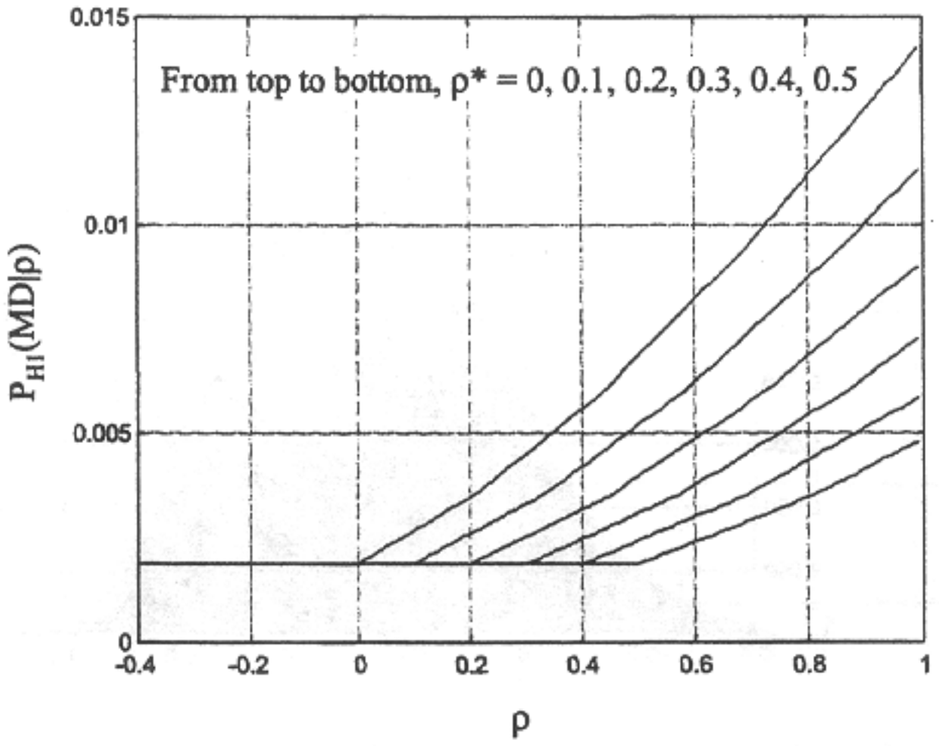

Fig. 15. H1 integrity risk sensitivity to ρ-variations, 22 SV case (with correlation buffering).

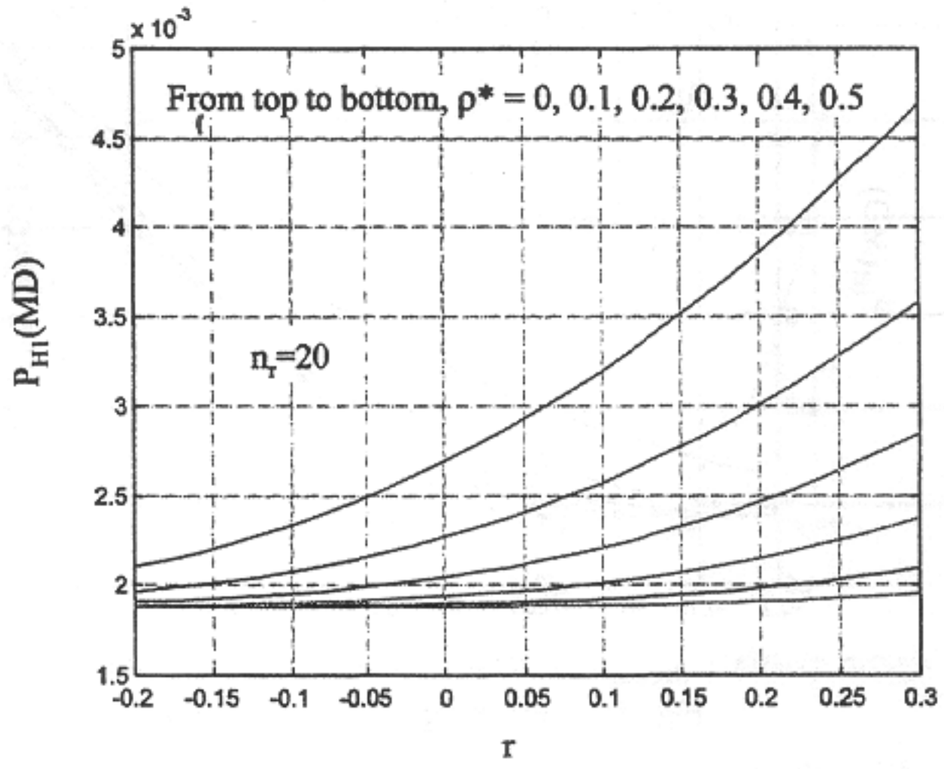

Fig. 16. H1 integrity risk versus r (with correlation buffering).

case in Fig. 14. For example, the results for $n_r = 20$ are plotted in Fig. 16. Note that the uppermost curve, which corresponds to the case $\rho^* = 0$, is identical to the $n_r = 20$ curve on Fig. 14. For a 5% acceptable tolerance on integrity risk relative to the nominal value of 0.0019, it is possible to obtain from Fig. 16 the maximum value of r allowable for a given value of ρ^*. This result can also be interpreted as the minimum value of ρ^* given a computed correlation coefficient estimate r. Fig. 17 shows the results for the $n_r = 20$ case under consideration and also for values of n_r equal to 50, 100, 200, and 1000. It is clear from Fig. 17 that a nearly linear relationship (with unity slope and positive y-intercept) exists between r and the minimum acceptable value of ρ^*. Thus, the (minimum) ρ^* can be approximately defined by the simple linear functions in Table III. Note that ρ^* must always be larger than r in order to account for the statistical uncertainty due to a finite number of samples. As n_r becomes large, however, the minimum acceptable value of ρ^* asymptotically approaches r.

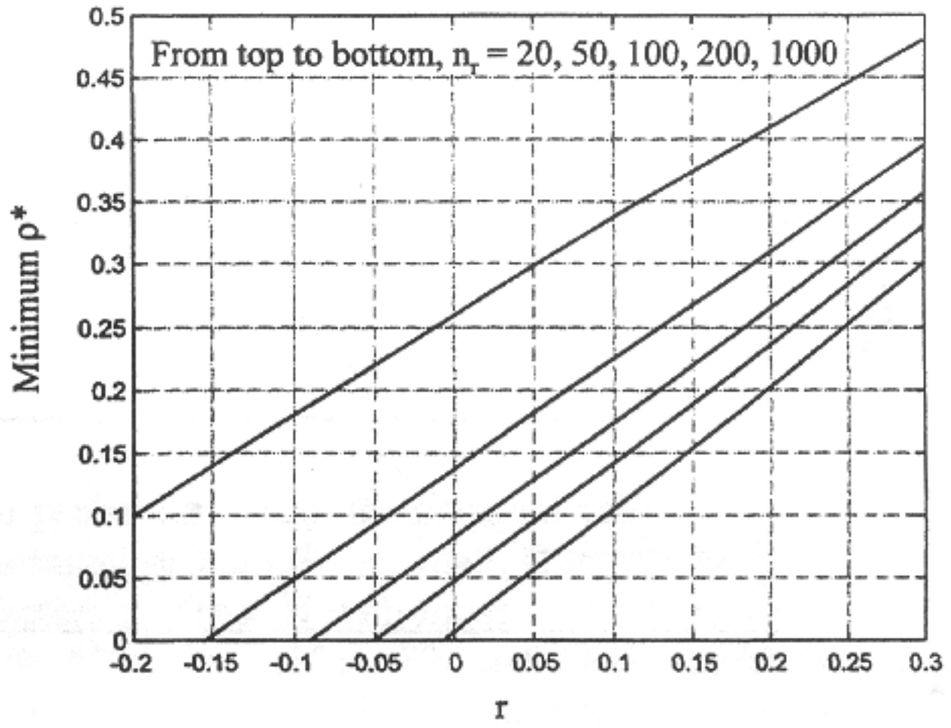

Fig. 17. Minimum ρ^* for 5% H1 integrity risk tolerance.

TABLE III
Correlation Buffer Parameters for H1

n_r	Minimum Value of ρ^*
50	$0.14 + r$
100	$0.08 + r$
200	$0.05 + r$
500	$0.02 + r$

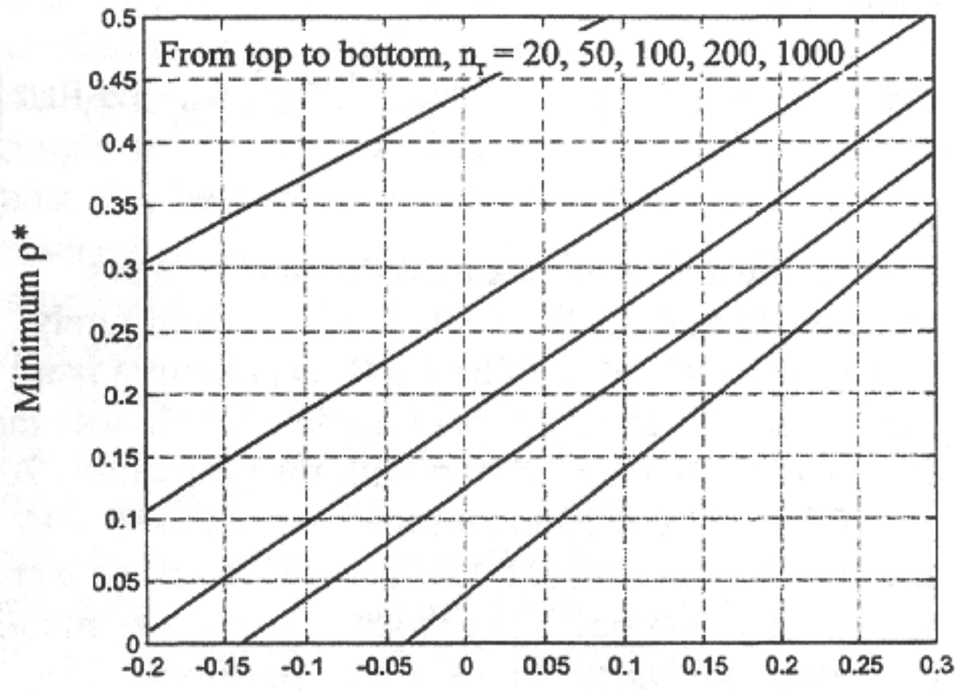

Fig. 18. Minimum ρ^* for 5% H0 integrity risk tolerance.

H0 Case

The correlation sensitivity analysis executed for the H1 case was repeated for H0. Fig. 18 shows the resulting minimum values of ρ^* given a computed correlation coefficient estimate r. Comparison of this figure with Fig. 17 shows that, as with the sigma sensitivity analysis, the H0 case is more restrictive. For a given computed value of r, the value of ρ^* required to ensure a 5% integrity risk tolerance is larger for the H0 case than the H1 case. Hence the H0 case must be the one used to define ρ^*. The approximate linear relations for H0 derived from Fig. 18 are given in Table IV.

TABLE IV
Correlation Buffer Parameters for H0

n_r	Minimum Value of ρ^*
50	$0.27 + r$
100	$0.18 + r$
200	$0.12 + r$
500	$0.06 + r$

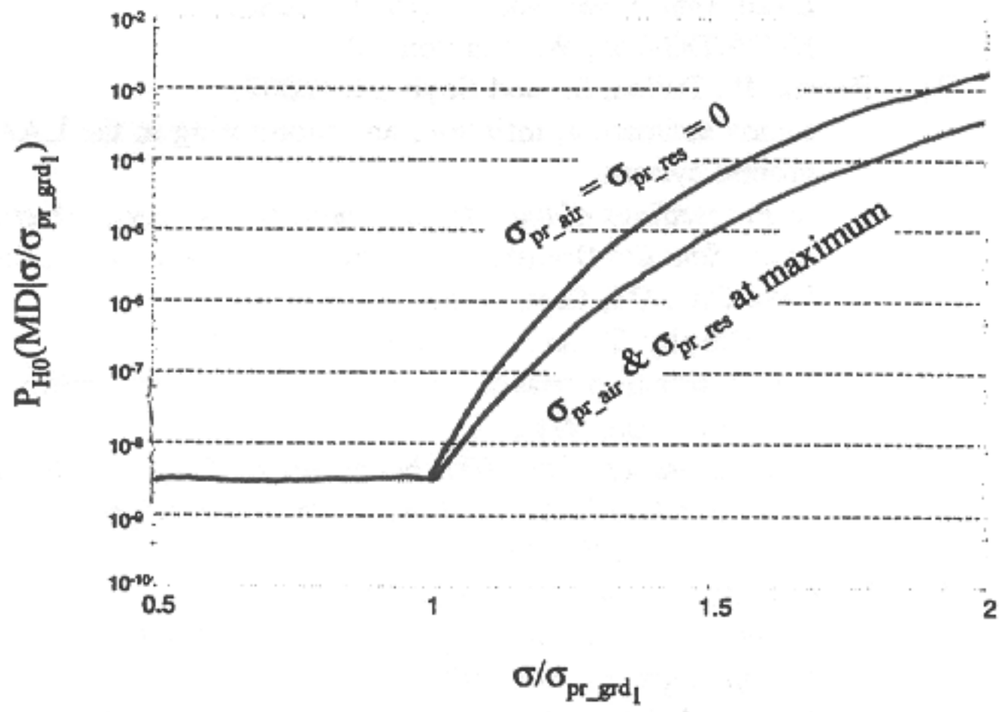

Fig. 19. H0 integrity risk sensitivity to σ-variations, 22 SV case.

TABLE V
Worst Case Sigma/Correlation Buffer Parameters for H0 (Results for Category 1 and $M = 3$)

n_s or n_r	Minimum σ^*	Minimum ρ^*
50	$1.34 \times s$	$0.30 + r$
100	$1.18 \times s$	$0.20 + r$
200	$1.10 \times s$	$0.13 + r$
500	$1.05 \times s$	$0.07 + r$

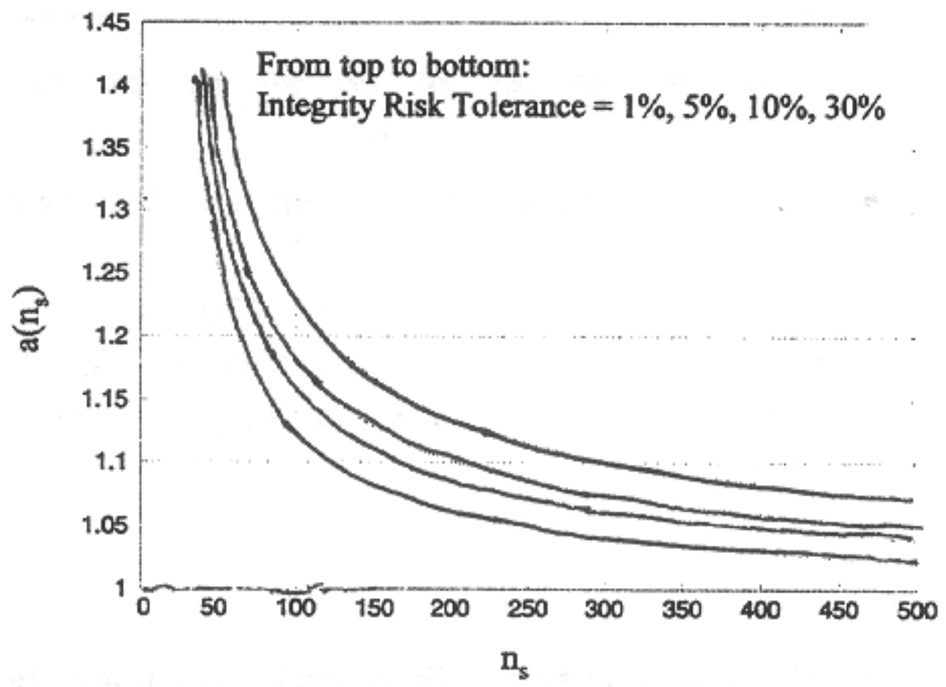

Fig. 20. Sigma buffer factor versus number of samples.

WORST CASE SENSITIVITY

The results generated thus far have been based on the nominal functions for σ_{pr_air} and σ_{pr_res} defined in [2]. Since these functions actually define the *maximum* permissible values for these parameters, in practical application it is likely that σ_{pr_air} and σ_{pr_res} will actually be smaller. In this case, it is expected that integrity risk will be more sensitive to variations in $\sigma_{pr\_gnd_1}$. It is therefore also instructive to examine the limiting scenario in which σ_{pr_air} and σ_{pr_res} are zero. We consider here only the H0 case (because it has already been shown to be more sensitive to variations in σ and ρ than H1). In this regard, Fig. 19 shows the upper bound curves for $P_{H0}(MD \mid \sigma/\sigma_{pr\_gnd_1})$ for the 22 SV constellation case. The upper (solid) curve defines integrity risk sensitivity when σ_{pr_air} and σ_{pr_res} are zero. The lower (dashed) curve, which is included only for comparison, corresponds to the case already covered where σ_{pr_air} and σ_{pr_res} hold their maximum permissible values. It is clear that integrity risk sensitivity increased for all values of $\sigma/\sigma_{pr\_gnd_1}$ greater than one. Note that in this region, where the curves are defined by the case where $\sigma/\sigma_{pr\_gnd_1}$ is varied on all satellites simultaneously, integrity risk sensitivity is invariant with respect to geometry for the case where σ_{pr_air} and σ_{pr_res} are zero. This is true since (7) reduces to

$$k_{ff\_md_e} \equiv \frac{k_{ff\_md}}{\sigma/\sigma_{pr\_gnd_1}}. \tag{13}$$

Using the same methodology described in the sections above, the minimum acceptable values for σ^* and ρ^* were computed assuming an acceptable integrity risk tolerance of 5%. These results are given in Table V for Category 1 approach and $M = 3$. Note, as expected, that the buffer parameters are slightly higher than those in Tables II and IV (which were derived using the maximum permissible values for σ_{pr_air} and σ_{pr_res}).

Taken together, the results of the sigma and correlation analyses above demonstrate that any value of σ_{pr_gnd} may be broadcast provided that the following inequality is satisfied:

$$\sigma_{pr\_gnd} > \sigma^* \sqrt{1 + (M-1)\rho^*}/\sqrt{M}. \tag{14}$$

INTEGRITY RISK TOLERANCE

Because the 5% integrity risk tolerance was arbitrarily selected, it is necessary to quantify how the buffer parameters vary with respect to integrity risk tolerance. In this regard, Fig. 20 shows the required value of the σ buffer factor as a function of the number of samples for various values of integrity risk tolerance. As one would naturally expect, the figure illustrates that for any given value of n_s, the buffer factor decreases as the integrity risk tolerance is relaxed. Analogous behavior is exhibited for the required correlation buffer parameter in Fig. 21. In both Figs. 20 and 21, it is also clear that only marginal reductions in buffer parameters will be

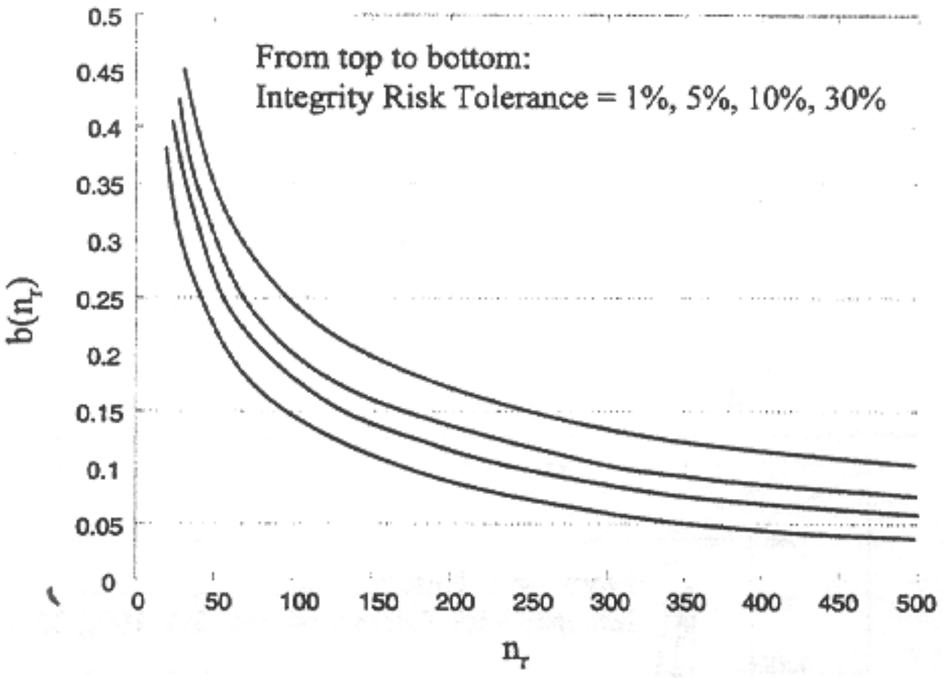

Fig. 21. Correlation buffer parameter versus number of samples.

realized for sample sets larger than 200 points. However, it is equally clear that sample sets smaller than 100 points will typically require rather large buffer parameters.

SUMMARY

In this paper, the sensitivity of LAAS integrity risk was investigated and quantified with respect to the statistical uncertainty in the knowledge of reference receiver error standard deviation and correlation between multiple reference receivers. A detailed methodology was presented to define the minimum acceptable buffer parameters for the value of σ_{pr_gnd} broadcast to the aircraft. This work addressed the Gaussian error structures associated with receiver thermal noise and diffuse multipath. It is likely that additional buffering for the effects of remaining errors, such as ground reflection multipath will be necessary. Continuing work in this regard is underway.

ACKNOWLEDGMENTS

The constructive comments and advice regarding this work provided by Steve Bellingham, Ron Braff, Barbara Clark, Bruce DeCleene, Sam Pullen, Curt Shively, and Frank van Graas are greatly appreciated.

REFERENCES

[1] Braff, R. (1997)
Description of the FAA's Local Area Augmentation System (LAAS).
Navigation, **44**, 4 (Winter 1997–1998).

[2] RTCA Inc.
Minimum Aviation System Performance Standards for the Local Area Augmentation System (LAAS).
RTCA/DO-245, Washington, DC.

[3] Pervan, B., Pullen, S., and Sayim, I. (2000)
Sigma estimation, inflation, and monitoring in the LAAS ground system.
In *Proceedings of the 13th International Technical Meeting of the Satellite Division of the Institute of Navigation*, Salt Lake City, UT, Sept. 2000.

[4] Liu, F., Murphy, T., and Skidmore, T. (1997)
LAAS signal-in-space integrity monitoring description and verification plan.
In *Proceedings of the 10th International Technical Meeting of the Satellite Division of the Institute of Navigation*, Kansas City, MO, Sept. 1997.

[5] Federal Aviation Administration (1999)
Local Area Augmentation System Ground Facility Specification, FAA-E-2937.
Federal Aviation Administration, Washington, DC, Sept. 21, 1999.

[6] RTCA Inc.
Minimum Operational Performance Standards for Global Positioning System/Wide Area Augmentation System Airborne Equipment.
RTCA/DO-229A, Washington, DC.

[7] Box, G., and Tiao, G. (1973)
Bayesian Inference in Statistical Analysis.
New York: Wiley, 1973.

[8] Spiegel, M., and Stephens, L. (1999)
Theory and Problems of Statistics (Schaum's Outline Series), (3rd ed.).
New York: McGraw-Hill, 1999.

Boris Pervan received a B.S. from the University of Notre Dame, Notre Dame, IN (1986), the M.S from the California Institute of Technology, Pasadena (1987), and Ph.D. from Stanford University, Stanford, CA (1996), all in aerospace engineering.

From 1987 to 1990, he was a Systems Engineer at Hughes Space and Communications Group involved in spacecraft mission analysis for a variety of commercial and government programs including Aussat B, the Advanced Tracking and Data Relay Satellite, the Advanced Launch System, and the Magellan mission to Venus. He was a Research Associate at Stanford from 1996 to 1998, serving as project leader for GPS Local Area Augmentation System (LAAS) research and development. Currently, he is Assistant Professor of Mechanical and Aerospace Engineering at the Illinois Institute of Technology in Chicago.

His work on the application of satellite-based navigation to aircraft precision landing earned him the 1996 RTCA William E. Jackson Award. He was also the 1999 recipient of the M. Barry Carlton Award from the IEEE Aerospace and Electronic Systems Society, the Abert J. Zahm Prize in Aeronautics at Notre Dame (1986), the Guggenheim Fellowship at Caltech (1987), and three best paper awards from the Institute of Navigation (ION) Satellite Division (1992, 1993, and 1997).

Irfan Sayim received a B.S. degree from Marmara University, Istanbul, Turkey (1990), Certificate of Business and Administration from the University of Istanbul, Istanbul, Turkey (1991), and the M.E. from Illinois Institute of Technology, Chicago (1996) in mechanical and aerospace engineering.

From 1990 to 1993, he was a Design Engineer and Technical Manager at Istanbul Fren Company. From 1993 to 1994, he was a Research Assistant at Gebze Institute of Technology, Gebze, Turkey. Currently, he is a Ph.D. candidate at Illinois Institute of Technology, Chicago, working on GPS Local Area Augmentation System (LAAS) integrity for aircraft precision landing.