Issues and Results Concerning the LAAS σ_{pr_gnd} Overbound

Boris Pervan and Irfan Sayim

Illinois Institute of Technology, Chicago, IL

Abstract – 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 paper 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.

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 limits 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 assume a normally distributed fault-free error model for the broadcast pseudorange corrections. Such a model is consistent with reference receiver ranging errors due to thermal noise and diffuse multipath. (Remaining errors, such as ground reflection multipath or systematic reference receiver/antenna errors are bounded in size and can in principle be separately accounted for by including a non-zero mean for the correction error in the computation of VPL.) The standard deviation of correction error is presumed by the aircraft to be equal to the broadcast value of 'σ_{pr_gnd}' for each satellite. 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 this regard, the finite 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 σ_{pr_gnd}.

This paper investigates the sensitivity of integrity risk to statistical uncertainty in the knowledge of 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. Specific approaches for the airborne processing may be found in [2], [3] and [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 VPL_{H0} and VPL_{H1} are computed at the aircraft based on values of broadcast correction error standard deviation (σ_{pr_gnd}) for each satellite also broadcast by the reference station. In addition, the prescribed computation of VPL_{H1} requires that the ground broadcast differences between the pseudorange 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_{md_ff} and k_{md}, which are defined below.

The general approach taken in this analysis is to first quantify 'true' missed detection probability given that the actual value of reference receiver error standard deviation (σ) deviates form the broadcast value σ_{pr_gnd}. 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 σ_{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 (VPL_{H1}) is given by the following expression derived in [2], [3]:

$$VPL_{H1} = \left| \sum_{n=1}^{N} S_{zn} B_n \right| + k_{md} \sqrt{\sum_{n=1}^{N} S_{zn}^2 \left[\frac{\sigma^2_{pr\_gnd_1}(n)}{M(n)-1} + \sigma^2_{pr\_air}(n) + \sigma^2_{pr\_res}(n) \right]} \quad (1)$$

where:

- n is the satellite index;
- S_{zn} is n-th 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 use to generate the broadcast correction;
- N is the number of available satellites;
- k_{md} is a multiplier used to set the desired level of missed detection probability assuming gaussian errors. k_{md} holds a value of 2.898 for a Category 1 approach with three reference receivers;
- σ_{pr_air} is the airborne measurement error standard deviation;
- σ_{pr_res} is the standard deviation of residual errors not directly attributable to ground of airborne error (such as tropospheric decorrelation);
- $\sigma_{pr\_gnd_1} \equiv \sqrt{M}\,\sigma_{pr\_gnd}$.

The maximum acceptable values of the standard deviations σ_{pr_gnd}, σ_{pr_air}, σ_{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 will explicitly consider only variations in σ_{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 standard deviation (σ) differs from the nominal value ($\sigma_{pr\_gnd_1}$) used to generate VPL_{H1}, the effective missed detection multiplier for the computed value of VPL_{H1} is given by

$$k_{md_e} \equiv k_{md} \frac{\sqrt{\sum_{n=1}^{N} S_{zn}^2 \left[\frac{\sigma^2_{pr\_gnd_1}(n)}{M(n)-1} + \sigma^2_{pr\_air}(n) + \sigma^2_{pr\_res}(n) \right]}}{\sqrt{\sum_{n=1}^{N} S_{zn}^2 \left[\frac{\sigma^2(n)}{M(n)-1} + \sigma^2_{pr\_air}(n) + \sigma^2_{pr\_res}(n) \right]}}. \quad (2)$$

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

$$P_{H1}(MD) = Q(k_{md_e}), \quad (3)$$

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

Clearly $P_{H1}(MD)$ 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.
- Elevation Mask: 5 deg
- Simulated Duration: 24 Hours
- 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 $VPL_{H0} < VAL$ using nominal value of σ_{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 from nominal on all visible SVs. The resulting integrity risk is shown in Figure 1 as a function of $\sigma/\sigma_{pr\_gnd_1}$. The discrete distribution of data points along the horizontal ($\sigma/\sigma_{pr\_gnd_1}$) direction in the figure corresponds to the discrete values of $\sigma/\sigma_{pr\_gnd_1}$ simulated. The vertical distribution of data points at each value of $\sigma/\sigma_{pr\_gnd_1}$ is due to the varying geometries accumulated over a 24 hour period. The upper bound integrity risk curve (solid) represents the highest level of integrity risk over the 24 hour duration. Note that when $\sigma/\sigma_{pr\_gnd_1} = 1$, the missed detection probability attains a nominal value of $Q(k_{md}) = 0.0019$.

Given that all 24 satellites are available, the results in Figure 1 are undoubtedly conservative since it is unlikely that broadcast σ_{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 Figure 2. When compared with the results of Figure 1, integrity risk is reduced for values of $\sigma/\sigma_{pr\_gnd_1} > 1$ (as expected) but increased for values of $\sigma/\sigma_{pr\_gnd_1} < 1$. The latter increase is due to the fact that σ is reduced on only one SV (in contrast with Figure 1 where σ was reduced on all SVs). Figure 3, which superposes the upper bound curves from Figures 1 and 2, clearly shows the difference in integrity risk sensitivity under the two sets of assumptions.

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 is shown in Figure 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 are 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 Figure 4; for any value of $\sigma/\sigma_{pr\_gnd_1}$, the upper of the two curves is used.

Given that the conditional probability $P_{H1}(MD|\sigma/\sigma_{pr\_gnd_1})$ has been established, it is still necessary to define a distribution for $\sigma/\sigma_{pr\_gnd_1}$ so that the overall risk probability can be quantified. In this regard, it is well known that the sample variance s^2 of n_s independent measurements derived from a gaussian distribution is Chi Square distributed:

$$(n_s - 1)\frac{s^2}{\sigma^2} \sim \chi^2_{n_s-1} \qquad (4)$$

Thus for a given computed value of s, the probability that σ lies in any specified range is easily computed. For example, Figure 5 shows the resulting probability mass function $P(\sigma/\sigma_{pr\_grd_1} | s/\sigma_{pr\_gnd_1}, n_s)$ for the case where $s/\sigma_{pr\_gnd_1} = 0.9$ and $n_s = 20$. The result is plotted together with the conditional $P_{H1}(MD|\sigma/\sigma_{pr\_gnd_1})$ 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 non-negligible. As the number of available samples n_s is increased, however, the likelihood that σ exceeds σ_{pr_gnd} decreases. Figure 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. Figure 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 missed detection probability given $s/\sigma_{pr\_gnd_1}$ and n_s was then computed numerically via

$$P_{H1}\left(MD \mid s/\sigma_{pr\_gnd_1}, n_s\right) = \sum_{\sigma} P_{H1}\left(MD \mid \sigma/\sigma_{pr\_gnd_1}\right) \times P\left(\sigma/\sigma_{pr\_gnd_1} \mid s/\sigma_{pr\_gnd_1}, n_s\right) \quad (5)$$

The results are plotted in Figure 8, which shows quantitatively how the H1 missed detection probability increases as $s/\sigma_{pr\_gnd_1}$ increases and as n_s decreases. The results are plotted in terms of percentage error (above the nominal value of 0.0019) in Figure 9.

To ensure integrity in an absolute sense, $P_{H1}\left(MD \mid s/\sigma_{pr\_gnd_1}, n_s\right)$ should not exceed the nominal specified value of 0.0019. However, the results in Figure 9 show that this criterion cannot be realistically attained because an infinitely large sample set is required. Nevertheless, it can be ensured that the missed detection probability does not differ from the nominal value by a significant amount. For example, Table 2 summarizes the results obtained from Figure 9 assuming that a 5% tolerance is acceptable. Under this assumption, the table quantifies the minimum value of $\sigma_{pr\_gnd_1}$ that may be broadcast given any value of s obtained from n_s samples. Clearly, the broadcast $\sigma_{pr\_gnd_1}$ must in general be larger than s. As expected, however, the buffer factor (the amount by which s must be scaled to define $\sigma_{pr\_gnd_1}$) approaches 1 as n_s grows large. (It should be noted that the quantitative results in Table 1 apply for the H1 case only.)

n_s	Minimum Value of $\sigma_{pr\_gnd_1}$
50	1.16×s
100	1.09×s
200	1.05×s
500	1.02×s

Table 1: Sigma Buffer Factors for H1

H0 Case

For the fault–free hypothesis, the vertical protection limit is given by the following expression [2], [3]:

$$VPL_{H0} = k_{md\_ff}\sqrt{\sum_{n=1}^{N} S_{zn}^2\left[\frac{\sigma_{pr\_gnd_1}^2(n)}{M(n)} + \sigma_{pr\_air}^2(n) + \sigma_{pr\_res}^2(n)\right]}. \quad (6)$$

When the actual ground standard deviation (σ) differs from the nominal value ($\sigma_{pr\_gnd_1}$) used to generate VPL_{H0}, the effective missed detection multiplier for the computed value of $VPL_{H0}$is given by

$$k_{md\_ff_e} \equiv k_{md\_ff}\frac{\sqrt{\sum_{n=1}^{N} S_{zn}^2\left[\frac{\sigma_{pr\_gnd_1}^2(n)}{M(n)} + \sigma_{pr\_air}^2(n) + \sigma_{pr\_res}^2(n)\right]}}{\sqrt{\sum_{n=1}^{N} S_{zn}^2\left[\frac{\sigma^2(n)}{M(n)} + \sigma_{pr\_air}^2(n) + \sigma_{pr\_res}^2(n)\right]}} \quad (7)$$

where k_{md_ff} is 5.810 for Category 1 approach with three reference receivers. The associated MD probability is

$$P_{H0}(MD) = 2Q(k_{md\_ff_e}). \quad (8)$$

The sensitivity analysis executed for the H1 case was repeated for H0. The resulting sensitivity of missed detection probability for the H0 case is quantified in Figure 10. As with the H1 case, $P_{H0}(MD|\ s/\sigma_{pr\_gnd}, n_s)$ increases as s/σ_{pr_gnd} increases and as n_s decreases. In Figure 11, the results of Figure 10 are plotted in terms of percentage error (relative to the nominal value of $2Q(k_{md\_ff}) = 6.2\text{x}10^{-9}$).

The results of the H0 σ-sensitivity analysis are summarized in Table 2, where a 5% missed detection 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 in probability from the nominal value when the nominal probability is small. Comparing the buffer factors in Table 2 with those in Table 1, it is clear that H0 is the more restrictive case (the required buffer factors are larger). Table 2 therefore serves to define the minimum σ buffer factor.

n_s	Minimum Value of $\sigma_{pr\_gnd_1}$
50	1.29×s
100	1.16×s
200	1.09×s
500	1.04×s

Table 2: Sigma Buffer Factors for H0

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 $\sigma_{pr\_gnd_1}$ 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 $\sigma_{pr\_gnd_1}$. 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:

$$\sigma = \sigma_{pr\_gnd_1}\sqrt{1+(M^*-1)\rho}\,. \qquad (9)$$

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:

$$\rho = \frac{(\sigma/\sigma_{pr\_gnd_1})^2 - 1}{M^* - 1}\,. \qquad (10)$$

H1 Case

Using the transformation above, the horizontal axis of Figure 3 may be rescaled in terms of ρ. The resulting upper bound curve for the conditional probability $P_{H1}(MD|\rho)$ is shown in Figure 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 [6] which is approximately gaussian:

$$Z \equiv \frac{1}{2}\ln\left(\frac{1+r}{1-r}\right) \sim N_Z\left[\frac{1}{2}\ln\left(\frac{1+\rho}{1-\rho}\right), (n_r-3)^{-\frac{1}{2}}\right] \qquad (11)$$

Figure 13 illustrates an example probability mass function $P(\rho|r,n_r)$ for r = 0.3, n_r = 20 and finely spaced intervals of ρ. (The $P_{H1}(MD|\rho)$ curve from Figure 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 missed detection probability given ρ and n_s was then computed numerically using a summation procedure equivalent to the sigma sensitivity case. The results are plotted in Figure 14, which shows quantitatively how the H1 missed detection probability increases as r increases and as n_s decreases.

When compared to the nominal missed detection 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 effect of positive correlation, and furthermore, integrity risk is magnified by uncertainty in correlation coefficient. Note, however, that in principle correlation can be accounted for by simply increasing the value of $\sigma_{pr\_gnd_1}$. For example for precisely known values of error standard deviation (σ) and correlation coefficient (ρ) we may use any value of $\sigma_{pr\_gnd_1}$ such that

$$\sigma_{pr\_gnd_1} > \sigma\sqrt{1+(M^*-1)\rho}\,. \qquad (12)$$

However, to account for the fact that σ is not known precisely (only the estimate s based on n_s samples is available), s in the equation above must be replaced with $a(n_s)s$, where $a(n_s)$ is the scale factor defined in Table 2. 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 $P(MD|\rho)$ was recomputed assuming that the value of $\sigma_{pr\_gnd_1}$ has

already been buffered using the above equation for selected values of ρ^* between 0 and 0.5. The resulting curves are given in Figure 15. Note that the curve corresponding to $\rho^* = 0$ (which represents the case where there is no buffering on $\sigma_{pr\_gnd_1}$) is identical to the $P(MD|\rho)$ curve in Figure 12. As expected, the influence of non-zero correlation coefficient on integrity risk is decreased as ρ^* is increased (i.e., as the buffer on $\sigma_{pr\_gnd_1}$ is increased).

For each of the ρ^* curves in Figure 15, it is possible to compute $P_{H1}(MD|r,n_r)$ as was done for the $\rho^* = 0$ case in Figure 13. For example, the results for $n_r = 20$ are plotted in Figure 16. Note that the uppermost curve, which corresponds to the case $\rho^* = 0$, is identical to the $n_r = 20$ curve on Figure 14. For a 5% acceptable tolerance on integrity risk relative to the nominal value of 0.0019, it is possible to obtain from Figure 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 compute correlation coefficient estimate r. Figure 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 Figure 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 thus be approximately defined by the simple linear functions in Table 3. Note that ρ^* must always be larger than r in order to account for the statistical uncertainty due to finite number of samples. As n_r becomes large, however, the minimum acceptable value of ρ^* asymptotically approaches r.

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

Table 3: Correlation Buffer Parameters for H1

H0 Case

The correlation sensitivity analysis executed for the H1 case was repeated for H0. Figure 18 shows the resulting minimum values of ρ^* given a computed correlation coefficient estimate r. Comparison of this figure with Figure 17 shows that, as with the sigma sensitivity analysis, the H0 case is more restrictive because 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 Figure 18 are given in Table 4.

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

Table 4: Correlation Buffer Parameters for H0

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, Figure 19 shows the upper bound curves for $P_{H0}(MD|\sigma/\sigma_{pr\_gnd_1})$ 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 equation (7) reduces to

$$k_{md\_ff_e} \equiv \frac{k_{md\_ff}}{\sigma/\sigma_{pr\_gnd_1}}. \qquad (13)$$

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 5 for Category 1 approach and M = 3. Note, as expected, that the buffer parameters are slightly higher than those in Tables 2 and 4 (which were derived using the maximum permissible values for σ_{pr_air} and σ_{pr_res}).

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

Table 5: Sigma/Correlation Buffer Parameters (Results for Category 1 and M = 3)

INTEGRITY RISK TOLERANCE

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:

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

where $\sigma^* = a(n_s)s$, $\rho^* = r + b(n_r)$, and s and r are the maximum values of standard deviation and correlation coefficient for any receiver and reference receiver pair, respectively. For the Category 1 case where M = 3 and a 5% relative missed detection (integrity risk) tolerance, the values of $a(\cdot)$ and $b(\cdot)$ are given for a number of discrete values of n_s and n_r in Table 5.

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, Figure 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 Figure 21. In both Figure 20 and 21, it is also clear that only marginal reductions in buffer parameters will be realized for sample sets larger than 200 points. However, it is equally 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 implicitly addressed only 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 or systematic reference receiver/antenna errors will be necessary. If these errors are bounded in size they can in principle be separately accounted for by introducing a non-zero mean for the correction error in the computation of VPL. Continuing work in this regard is underway.

ACKNOWLEDGEMENTS

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

REFERENCES

[1] Braff, R., "Description of the FAA's Local Area Augmentation System (LAAS)," *Navigation*, Vol. 44, No. 4, Winter 1997-98.

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

[3] Liu, F., Murphy, T., and Skidmore, T., "LAAS Signal-in-Space Integrity Monitoring Description and Verification Plan," Proceedings of ION GPS-97, Kansas City, MO, September 1997.

[4] Pervan, B., Pullen, S. and Christie, J., "A Multiple Hypothesis Approach to Satellite Navigation Integrity," *Navigation*, Vol. 45, No. 1, Spring 1998.

[5] Federal Aviation Administration, *Local Area Augmentation System Ground Facility Specification*, FAA-E-2937, Washington, D.C., September 21, 1999.

[6] Spiegel, M. and Stephens, L., *Theory and Problems of Statistics*, McGraw-Hill Schaum's Outline Series, 3rd Edition, 1999.

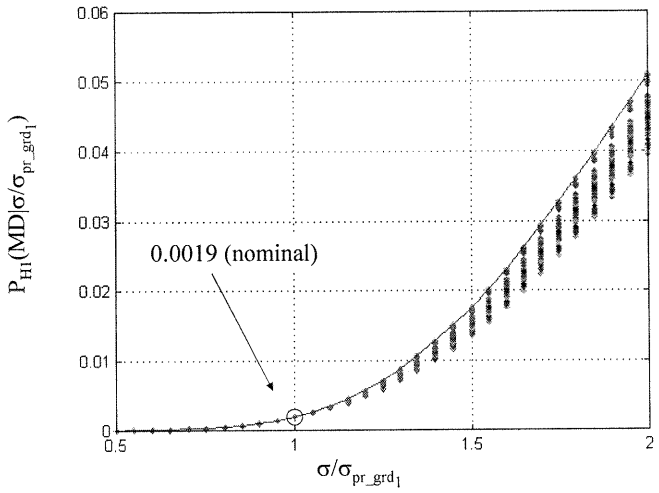

Figure 1: H1 Integrity Risk Sensitivity to σ-Variations -- 24 SV Case
(σ Varied on All SVs In View)

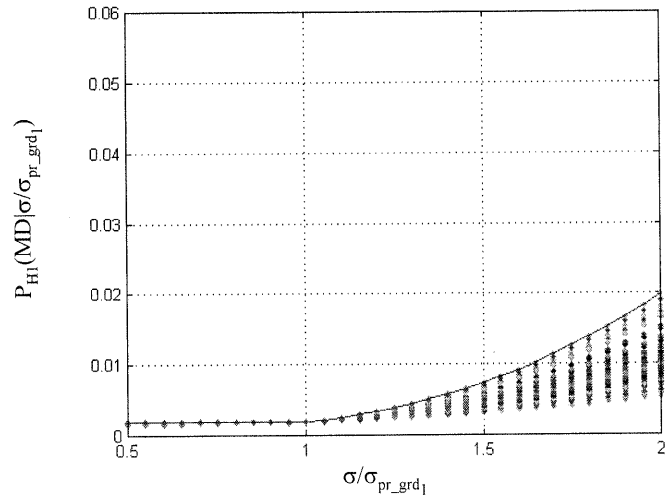

Figure 2: H1 Integrity Risk Sensitivity to σ-Variations -- 24 SV Case
(σ Varied on Worst-Case SV Only)

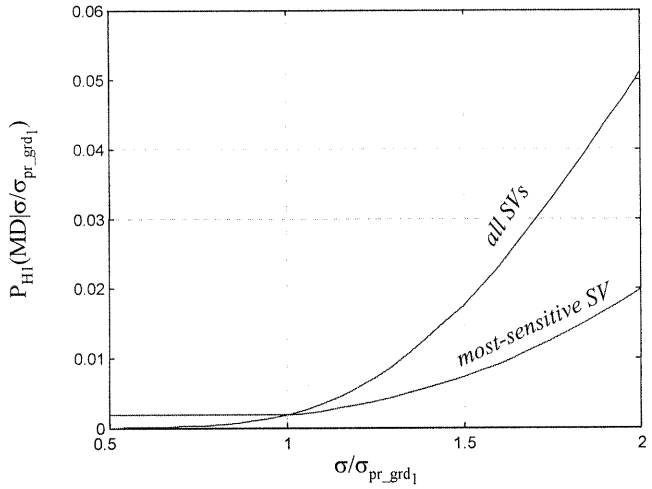

Figure 3: H1 Integrity Risk Sensitivity to σ-Variations -- 24 SV Case
(Upper Bound Curves from Figures 1 and 2)

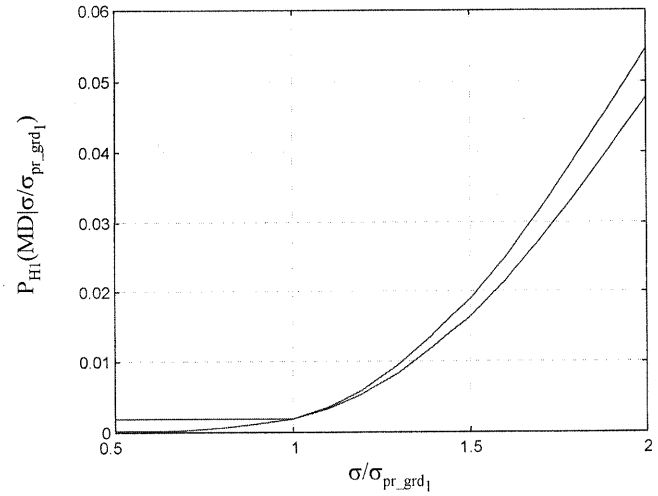

Figure 4: H1 Integrity Risk Sensitivity to σ-Variations -- 22 SV Case
(Upper Bound Curves)

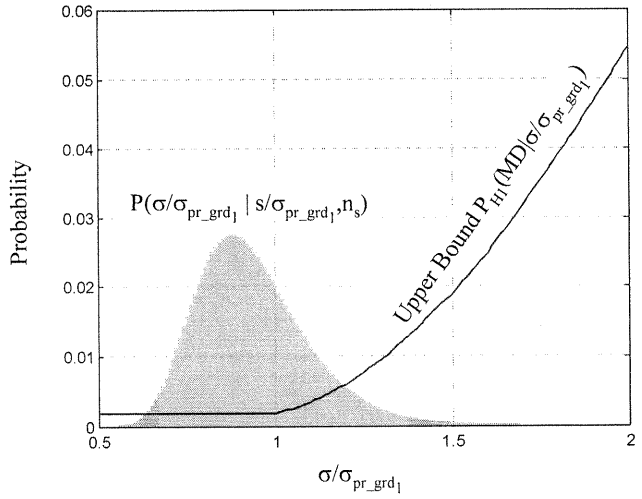

Figure 5: Probability Distribution for $\sigma/\sigma_{pr\_grd_1}$ for $s/\sigma_{pr\_grd_1} = 0.9$, $n_s = 20$

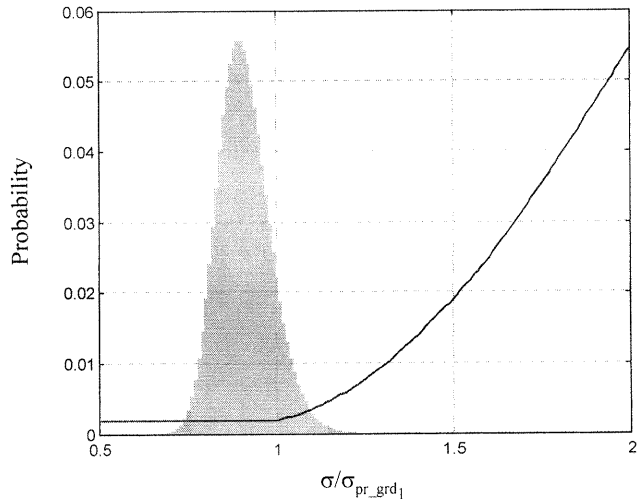

Figure 6: Probability Distribution for $\sigma/\sigma_{pr\_grd_1}$ for $s/\sigma_{pr\_grd_1} = 0.9$, $n_s = 80$

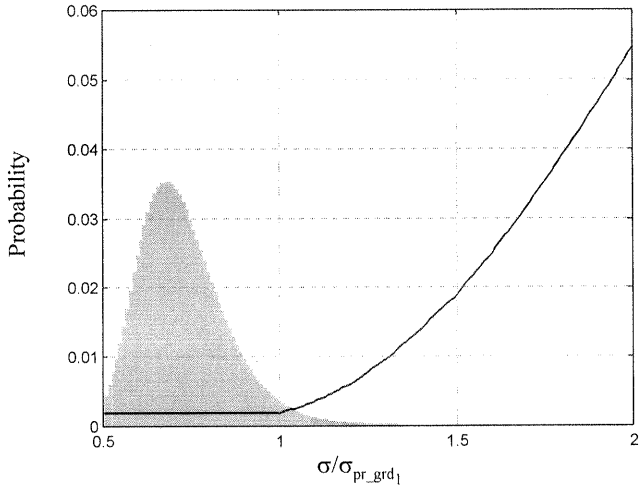

Figure 7: Probability Distribution for $\sigma/\sigma_{pr\_grd_1}$ for $s/\sigma_{pr\_grd_1} = 0.7$, $n_s = 20$

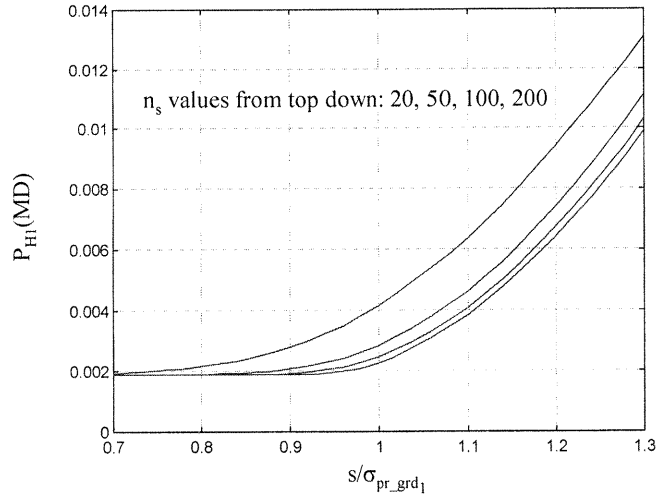

Figure 8: H1 Integrity Risk vs. $s/\sigma_{pr\_grd_1}$

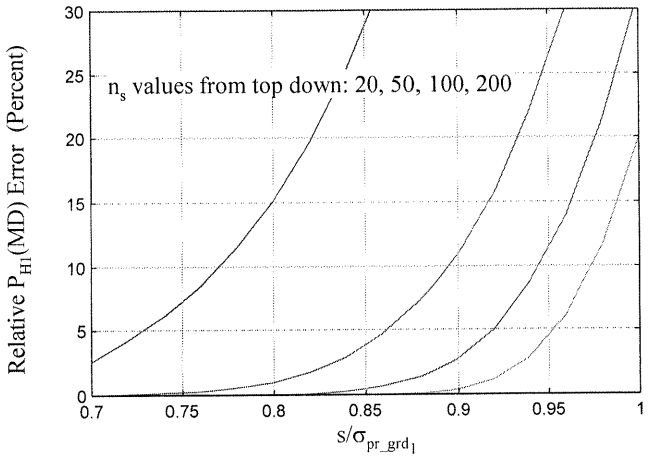

Figure 9: H1 Integrity Risk vs. $s/\sigma_{pr\_grd_1}$ (Relative Error)

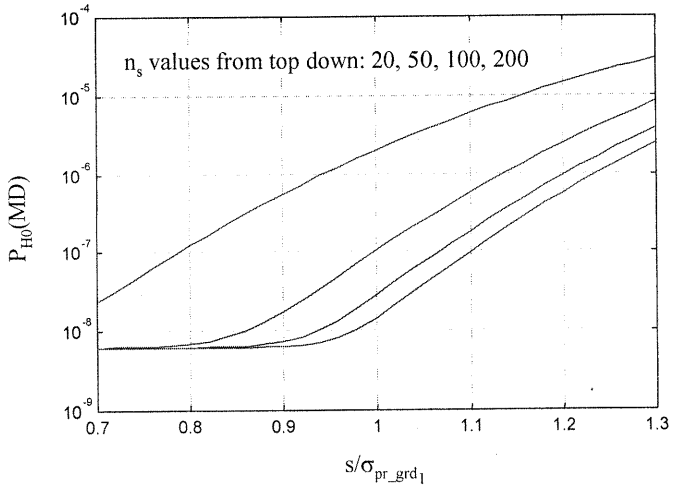

Figure 10: H0 Integrity Risk vs. $s/\sigma_{pr\_grd_1}$

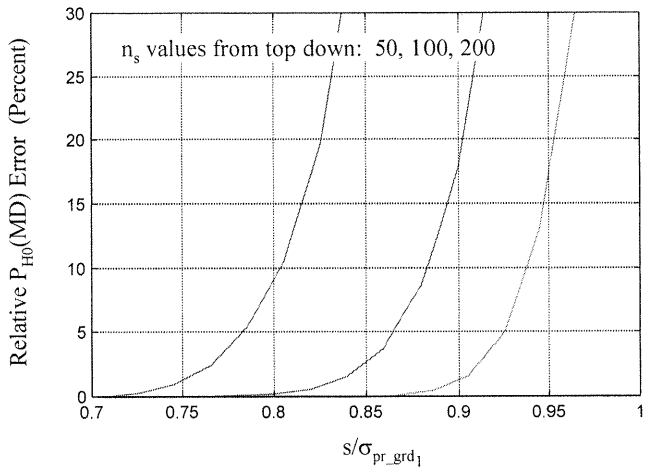

Figure 11: H0 Integrity Risk vs. $s/\sigma_{pr\_grd_1}$ (Relative Error)

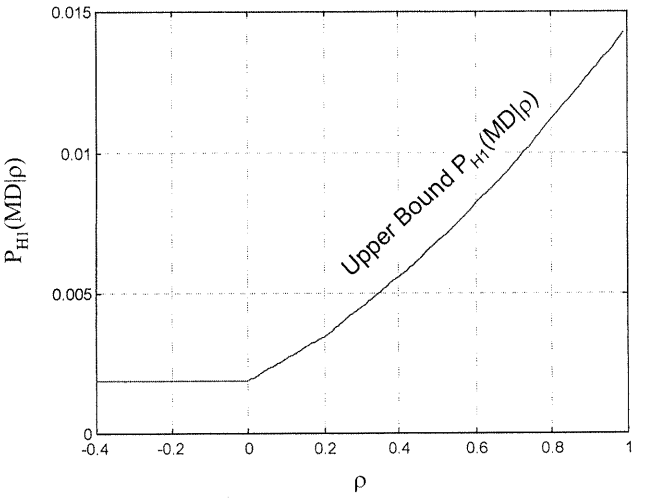

Figure 12: H1 Integrity Risk Sensitivity to ρ-Variations -- 22 SV Case
(Upper Bound Curve)

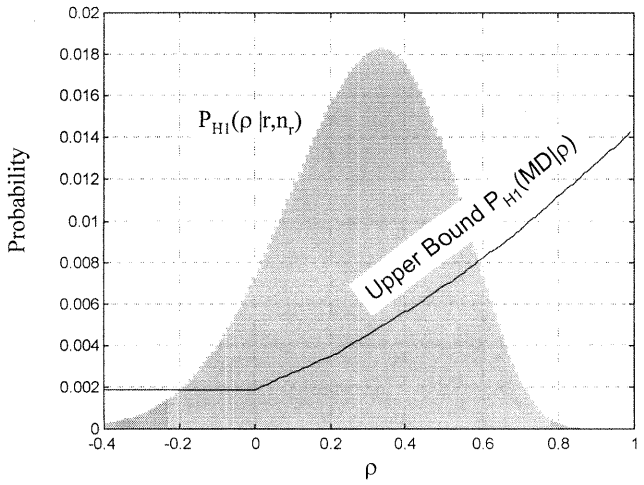

Figure 13: Probability Distribution for ρ given r = 0.3, n_r = 20

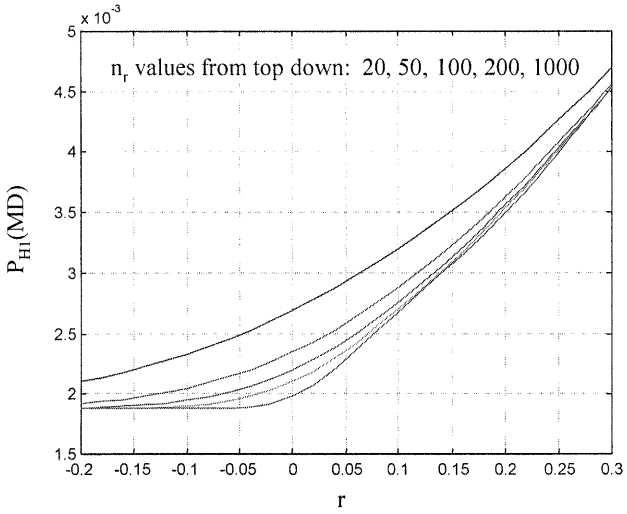

Figure 14: H1 Integrity Risk vs. r

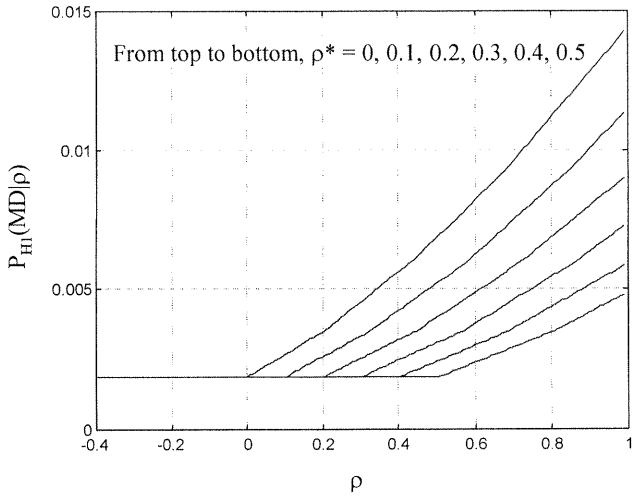

Figure 15: H1 Integrity Risk Sensitivity to ρ-Variations -- 22 SV Case
(With Correlation Buffering)

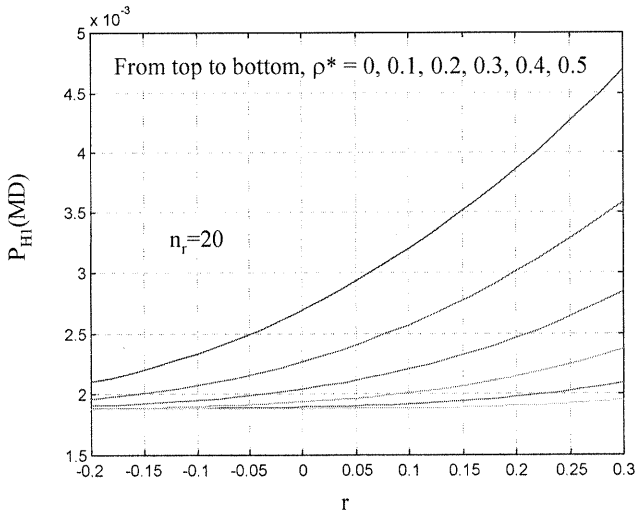

Figure 16: H1 Integrity Risk vs. r (With Correlation Buffering)

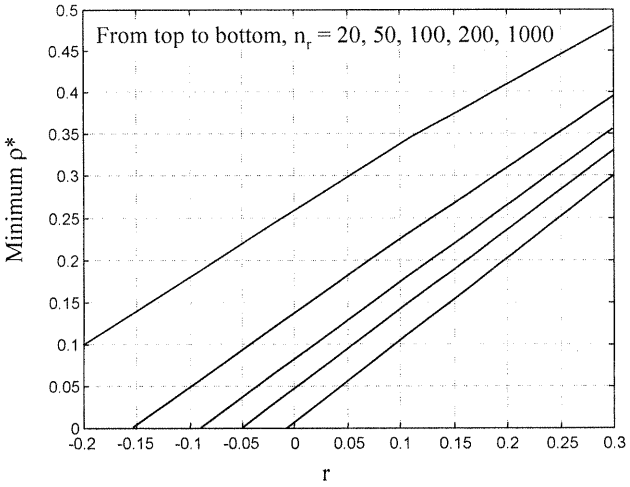

Figure 17: Minimum ρ* for 5% H1 Integrity Risk Tolerance

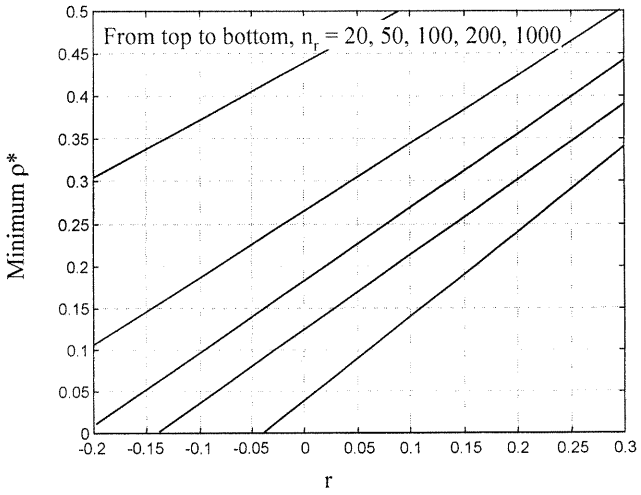

Figure 18: Minimum ρ* for 5% H0 Integrity Risk Tolerance

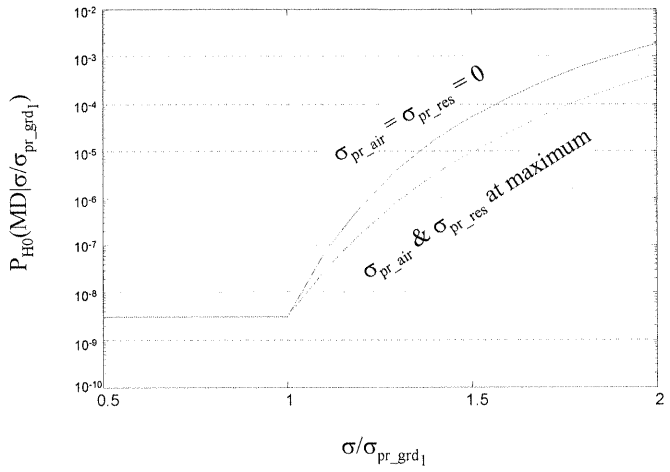

Figure 19: H0 Integrity Risk Sensitivity to σ-Variations -- 22 SV Case

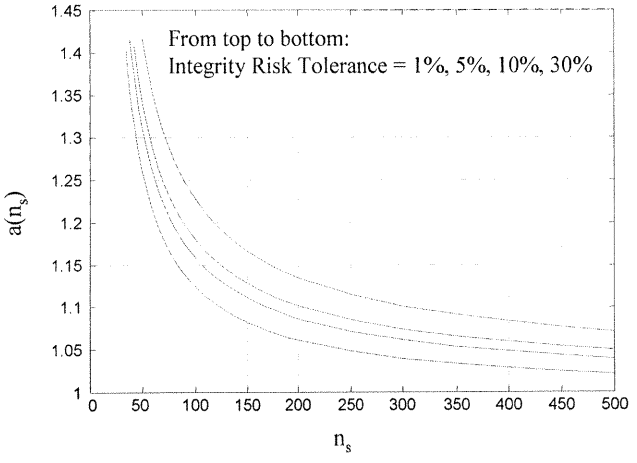

Figure 20: Sigma Buffer Factor vs. Number of Samples

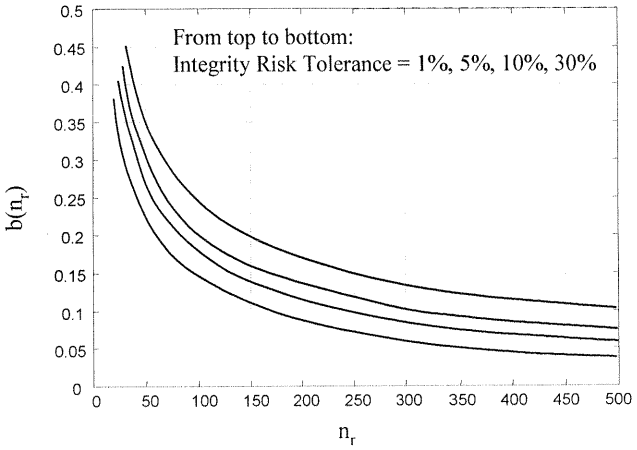

Figure 21: Correlation Buffer Parameter vs. Number of Samples