LAAS Integrity Monitoring Methodology: Least Squares versus Mid-Value Selection

P. Bastiaan Ober, *Delft University of Technology - Telecommunications and Traffic Control Systems Group*

BIOGRAPHY

Bastiaan Ober has 5 years of experience in radio navigation. He received his M.Sc. in Electrical Engineering from Delft University of Technology in 1993. Areas of experience include the influence of multipath on GPS positioning, carrier phase differential GPS, ambiguity resolution and integrity monitoring. He is currently working as a Ph.D. student at Delft University of Technology, doing research on integrity monitoring algorithms for integrated navigation systems.

ABSTRACT

This paper provides a description of the general concepts behind the integrity providence of locally augmented GPS. It shows why the exact integrity is hard to evaluate, even when simple error models are used, and explains the problems that will have to be resolved before a full performance assessment can be obtained. Integrity is mainly provided by using redundant receivers. Therefore, the paper pays special attention on two ways to combine the outputs of multiple receivers. It discusses the merits of both the traditional error detection and identification scheme and a simple robust scheme as mid-value selection, which proves much easier to evaluate.

1. Introduction

This paper is written to clarify the general methodology of LAAS integrity monitoring. This does not imply that an attempt was made to discuss all possible ways to do such monitoring; this would probably only obscure the common principles all methods basically exploit. On the other hand, the concept of LAAS integrity monitoring has already been dealt with in several other papers such as [Murphy97]. This paper therefore aims at addressing LAAS integrity monitoring on a level somewhere in between the abstraction of the concept and the specific properties of a particular architecture. Whenever appropriate the LAAS architecture that is currently being developed by the FAA LAAS Team in conjunction with the RTCA [Liu97] is used as an illustration of a practical implementation. However, the focus remains on the general approach rather than on implementation details.

The paper discusses the performance of Autonomous Integrity Monitoring (AIM) in Local Area Augmentation Systems (LAAS) that augment GPS for use in high performance landing systems. AIM exploits signal redundancy to extract information on error characteristics and removes suspicious measurements before they can effect system operation adversely. For a differential system such as LAAS AIM differs from the ‘usual’ en route AIM in three ways. First of all a cascade of two error detection schemes is involved: the first one on the ground, the second in the air. Secondly, the focus in a differential system lies on receiver errors rather than satellite errors as the latter usually cancel through the differential correction. Therefore, integrity is provided by redundant receivers. Finally, the differential corrections reduce the noise levels in the final positioning algorithm. This ensures a higher nominal accuracy and integrity, but also makes the results of performance computations more sensitive to modeling errors.

In parallel with AIM, a particular robust estimation technique is considered. In the so-called mid-value selection scheme, the available redundancy is used to reduce the effects of errors rather than for detection. This results in a completely different approach to integrity of which the particular advantages and disadvantages will be reviewed.

After a short introduction on the used notational conventions, the LAAS measurement and error models will be specified in section 3. Section 4 reviews the computations on the ground and in the air and introduces the two possible ways to combine the outputs of the redundant receivers: least squares with error detection, and mid-value selection. More detail on the error detection is supplied in section 5, after which section 6 discusses the performance of both alternatives. Section 7 provides a short overview of the main differences between the two methods and gives a limited amount of simulation results. Conclusions and recommendations conclude the paper.

2. Notational conventions and symbols

In this paper, the following notations will be used for the statistical distributions:

$x \sim X$: x has probability density function (pdf) X
$x \tilde{<} X$: x has a pdf bounded by X
$pdf(x)$: the probability density function of x
$X * Y$: convolution of pdfs X and Y
$\mathsf{N}(\vec{\mu}, Q)$: multivariate normal distribution with mean $\vec{\mu}$ and covariance matrix Q
$\mathsf{U}(\vec{a}, \vec{b})$: multivariate uniform distribution; the ith element is uniformly distributed on the interval $[\vec{a}[i], \vec{b}[i]]$

The median of multiple vectors is used on an element by element basis:

$$median(\vec{x}_1, \vec{x}_2, \vec{x}_3) = \begin{bmatrix} median(\vec{x}_1[1], \vec{x}_2[1]\vec{x}_3[1]) \\ \vdots \\ median(\vec{x}_1[N], \vec{x}_2[N]\vec{x}_3[N]) \end{bmatrix}$$

'All zero' and 'all one' vectors will be denoted as $\vec{0}$ and $\vec{1}$ respectively.

Although the following symbols will be introduced in the text as well to ease reading the text linearly, the following list is provided as a service to those readers reading otherwise and as a quick reference:

- *Numbers and indices*

M_a: number of airborne receivers
M_g: number of ground receivers
M_x: either M_a or M_g
ar: index of airborne receiver $\in a1 \ldots aM_a$
gr: index of ground receiver $\in g1 \ldots gM_g$
xr: either *ar* or *gr*
N: number of satellites tracked by ground receivers
N_C: number of corrected pseudoranges tracked by the airborne receivers

- *Ground station related symbols*

$\vec{z}_g$: pseudorange corrections from combined output of ground receivers
$\vec{e}_g$: errors in the corrections from the ground station
$\vec{z}_{gr}$: pseudorange corrections measured by receiver *gr*
$\vec{e}_{gr}$: errors in receiver *gr*
$\vec{x}_g$: pseudorange errors in satellites

- *Airborne user related symbols*

$\vec{z}_a$: corrected pseudoranges from combined airborne receiver outputs
$\vec{z}_{ar}$: corrected pseudoranges measured by receiver *ar*
$\vec{e}_{ar}$: errors in receiver *ar*
H_a : airborne observation matrix
$\vec{x}_a$: vector with airborne position and clock bias
$\Delta\hat{v}$: vertical position error
$\vec{s}_a^{\,T}$: vector that translates errors in corrected pseudoranges to the vertical position error domain

- *Either airborne user or ground station related symbols*

$\vec{e}_x$: either $\vec{e}_a$ or $\vec{e}_g$
$\vec{e}_{xr}$: either $\vec{e}_{ar}$ or $\vec{e}_{gr}$
$\vec{v}_{xr}$: receiver noise in receiver *xr*
$\vec{\mu}_{xr}^{MP}$: multipath error in receiver *xr*
$\vec{\mu}_{xr}^{F}$: receiver failure error in receiver *xr*
β_{xr}^{i} : Boolean function that is zero when the i^{th} channel of receiver *xr* is excluded by the error detector.
B_{xr}^{i} : test statistic to test on an error in the i^{th} channel of receiver *xr*
h_{xr}^{i} : threshold for B_{xr}^{i} to decide on an error

3. The LAAS system model

A GPS based LAAS contains three major parts:

- the *GPS space segment*, which delivers the navigation signals
- the *ground station*, that computes the errors on the received pseudoranges and uplinks them as corrections
- the *airborne user*, that uses both the GPS signals and the corrections for navigation

Because errors that are common between the ground segment and the user are corrected, the system is relatively

insensible to errors of the space segment. In a differential system, the focus will therefore lie on receiver errors and local signal disturbances at the ground station and the airborne user.
Before introducing the relevant error models, the system models of the ground and airborne segment will be provided. To mitigate the consequences of failing receivers, both the ground and airborne equipment will usually consist of multiple receivers. In the following analysis, it has implicitly been assumed that the clocks of the ground receivers are perfectly synchronized. This can for example be achieved by calibrating all receivers with a GPS signal generator [RTCA96].

3.1 Ground station observation model

Because the position of the ground antennas is known, the ground station measures the pseudorange errors straightforwardly. Therefore, the pseudorange error measurements that are measured by the gr^{th} ground receiver can simply be modeled as:

$$\vec{z}_{gr} = \vec{x}_g + \vec{e}_{gr} \tag{3.1.1}$$

with

$\vec{z}_{gr}$: N-vector of measured pseudorange corrections
$\vec{x}_g$: N-vector of pseudorange errors
$\vec{e}_{gr}$: N-vector of receiver introduced errors

N is the number of satellites that are tracked by the ground station.

3.2 Airborne observation model

Although the exact relation is nonlinear, it is assumed that a linear regression model describes the relation between the corrected GPS pseudoranges and the user position with sufficient accuracy. The model for the ar^{th} airborne receiver can thus be written as

$$\vec{z}_{ar} = H_a \cdot \vec{x}_a + \vec{e}_{ar} + \vec{e}_g \tag{3.2.1}$$

in which

$\vec{z}_{ar}$: N_C-vector of corrected pseudoranges
H_a : $N_C \times 4$ airborne observation matrix
$\vec{x}_a$: 4-vector of unknowns (position, clock bias)
$\vec{e}_{ar}$: N_C-vector with receiver errors
$\vec{e}_g$: errors in the pseudorange corrections

N_C is the number of satellites for which the ground station broadcasts corrections. Of course, it is possible that the airborne receivers track some pseudoranges for which they do not receive corrections. This situation occurs when they track more satellites than the ground station or when the error detection circuitry on the ground prevents a certain correction to be uplinked. Although it would still be possible to use the uncorrected pseudoranges, they are not expected to contribute significantly to either the accuracy or integrity due to their relatively high noise levels. Their use will therefore not be considered in this paper.

Note that the errors $\vec{e}_{ar}$ and $\vec{e}_g$ are assumed not to contain the common satellite errors, but comprise only receiver related error components. They can therefore be considered independent.

3.3 The error models

Because GPS pseudorange errors introduced by satellite (clock) failures or ionospheric and tropospheric delays are mostly common to the ground and airborne receiver, they are not expected to dominate the quality of the final position fix. In this paper, it is assumed that the most important error sources that are present in the system are receiver related. The models for the ground and airborne receiver errors will be very similar. To avoid duplication of many of the equations, the index *xr* will frequently be used to represent both the earlier introduced *gr* and *ar* simultaneously.

The basic error model contains three supposedly independent error components:

- a noise-like component $\vec{v}_{xr}$
- a multipath bias component $\vec{\mu}_{xr}^{MP}$
- a receiver failure bias $\vec{\mu}_{xr}^{F}$

that are added to arrive at the total receiver error:

$$\vec{e}_{xr} = \vec{v}_{xr} + \vec{\mu}_{xr}^{MP} + \vec{\mu}_{xr}^{F} \tag{3.3.1}$$

The models of each of the three error sources will be the topic of the coming three paragraphs. Finally, the consequences of the choice of these ‘submodels’ for the distribution of (3.3.1) will be considered in this section’s closing paragraph.

3.3.1 Noise models

It is a usual assumption that all noisy error contributions are characterized sufficiently well when they are modeled as independent zero mean Gaussian noise. When the covariance of the noise is Q_{xr}, the probability density functions of these contributions become:

$$\vec{v}_{xr} \sim \mathsf{N}(\vec{0}, Q_{xr}) \qquad (3.3.1.1)$$

The covariance matrices have the variances of the different satellites -usually as a function of satellite elevation- on their diagonal.

3.3.2 Multipath modeling

Because of the low noise levels multipath will have a large impact on the overall system performance. It is therefore extremely important to avoid using are overly optimistic models. Unfortunately, it is hard to precisely characterize the multipath present in a certain sample either statistically or otherwise. The influence of multipath on the measured pseudoranges not only depends on the number and power of the reflections, but also on the antenna and receiver design.

Computationally, it is attractive to model multipath errors as zero mean Gaussian. This approach is followed in [Liu97] and [VanDyke97] but these papers provide no justification. Still, it might well be that this model is appropriate at least for the airborne segment because of the airplane dynamics.

Some studies such as [RTCA96] assume that with specially designed antennas, the errors caused by multipath may be limited in norm. Consequently, multipath is sometimes modeled as a uniform distribution with boundaries of 20 (or maybe 15) centimeters [Skidmore96]:

$$\vec{\mu}_{xr}^{MP} \sim \mathsf{U}(-0.2 \cdot \vec{1}, 0.2 \cdot \vec{1}) \qquad (3.3.2.1)$$

[Skidmore96] states that full verification of (3.3.2.1) is still a topic of current research.

Assuming that the multipath error is indeed bounded, it would also be possible to model the multipath as 'worst case', effectively concentration the probability density function in one single point. This makes the model deterministic and safe, as it is always conservative. Other advantages are the independence of the model of the exact pdf, and its computational simplicity. Obvious disadvantage is that worst case modeling might be *overly* conservative and limit system availability.

3.3.3 Receiver failures

When a receiver functions normally, its measurements should be free of significant biases (other than multipath biases), and therefore $\vec{\mu}_{xr}^{F} = \vec{0}$. In the case of a hard- or software failure of a receiver, the receiver will most likely have some undefined state. This makes it impossible to say anything on the resulting pseudorange errors of the receiver output. Therefore, for this type of failures the use of a 'worst case' assumption seems to be the only procedure that always stays 'on the safe side' [Ober97].

The interpretation and determination of the worst case is not at all trivial. As it will also depend on the algorithms to compute the system's unknowns and detect errors, the discussion will have to be postponed until section 6. This section finishes with the overall error model that results from the three submodels of all error contributions.

3.3.4 The overall error model

Within the LAAS model, failures from several different sources accumulate. When the disturbances are additive and independent, the probability density functions of their sum equals the convolution of the probability density functions of all the contributions. The pdf of the total error of (3.3.1) therefore reads:

$$\mathrm{pdf}(\vec{e}_{xr}) = \mathrm{pdf}(\vec{v}_{xr}) * \mathrm{pdf}(\vec{\mu}_{xr}^{MP}) * \mathrm{pdf}(\vec{\mu}_{xr}^{F}) \qquad (3.3.4.1)$$

Computation of such a convolution is generally hard. Some distributions like the Gaussian and chi-squared distributions, are reproductive under convolution, but this is not true for most other distributions. This is one of the reasons why the uniform distribution model for the multipath is computationally unattractive. One of the possible remedies to ease computations is to bound the exact overall error distribution by a bounding Gaussian distribution, as has been suggested in [VanGraas97]. Similar 'pdf bounding' techniques will be applied to the mid-value selection performance computations in section 6, but have to be used with great care.

Both the 'worst case' deterministic and the Gaussian multipath model have the advantage of keeping all pdfs in (3.3.4.1) normally distributed. In the deterministic case, the mean of this distribution will be nonzero; when " ' " indicates worst case values, the overall pdf becomes nothing but the following shifted version of the noise pdf:

$$\vec{e}_{xr} \sim \mathsf{N}(\vec{\mu}_{xr}'^{MP} + \vec{\mu}_{xr}'^{F}, Q_{xr}) \qquad (3.3.4.2)$$

When the multipath is assumed normally distributed, it can be incorporated in an increase of Q_{xr}, while $\vec{\mu}_{xr}'^{MP}$ remains zero, and thus:

$$\vec{e}_{xr} \sim \mathsf{N}(\vec{\mu}_{xr}'^{F}, Q_{xr}) \qquad (3.3.4.3)$$

4. Computing the unknowns

This section will cover the computation of the unknowns in a LAAS system. In the airplane, this is done to compute the position, on the ground to compute the pseudorange errors corrections that should be uplinked.

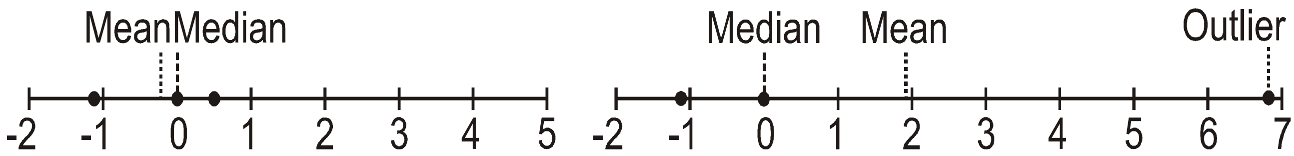

Figure 1. The mean and median in cases of zero and one outlying measurements. The mean is pulled towards the outlying measurement, while the median is not.

Special attention will be paid to the way the outputs from redundant receivers are combined. One of the possibilities is to use least squares estimation. When all receivers are identical, least squares estimation results in taking the mean of the outputs. The advantage of using least squares is that receiver noise is averaged out optimally. Because of its sensitivity to errors, integrity will have to be provided by a separate error detection and isolation scheme that works in conjunction with the least squares estimation.

An alternative approach is the use of an error resistant (robust) computation scheme. One of the most straightforward and robust schemes uses the mid-value (median) rather than the mean. This scheme will be discussed as well. Because of its relative insensitivity to errors, a separate error detection and identification scheme is not needed although it still can be used.

4.1 Combining the output of multiple receivers

When combining M_x equally accurate measurements of essentially the same quantity, the most intuitive way to combine the different outcomes is to take their mean value as the best representative:

$$\vec{z}_x = \frac{1}{M_x}\sum_{r=1}^{M_x}\vec{z}_{xr} \tag{4.1.1}$$

The mean value is exactly the value that minimizes the sum of the squared estimation residuals. In the remainder of the text the term 'least squares estimation' will therefore be used frequently to refer to (4.1.1).

Using (4.1.1), the error in the combined output simply becomes the average of the errors in the different receivers:

$$\vec{e}_x = \frac{1}{M_x}\sum_{r=1}^{M_x}\vec{e}_{xr} \tag{4.1.2}$$

A well known property of the mean is its sensitivity to errors: although decreased by a factor $1/M_x$, an error in one of the receivers can still cause arbitrarily high errors in the combined output. To obtain integrity, the least squares technique will therefore *always* have to be combined with error detection algorithms. One possible error detection algorithm will be explained in more detail in section 5.

An alternative to least squares estimation is the use of mid-value selection. Mid-value selection is a simple robust estimation technique, which makes the final estimate insensitive to large errors as long as less than half of the receivers are malfunctioning. Using the notation from section 2, the combined output and its errors can be written as

$$\vec{z}_x = median(\vec{z}_{x1},...,\vec{z}_{xM_x}) \tag{4.1.5}$$

$$\vec{e}_x = median(\vec{e}_{x1},...,\vec{e}_{xM_x}) \tag{4.1.6}$$

The median offers 'build in' integrity because large receiver errors will have a very limited influence on the final position solution. Additional integrity can still be obtained by using error detection and identification schemes, but is beyond the scope of this paper.

Figure 1 shows why mid-value estimation is error resistant. When there are three receivers of which one is in error, the mean moves strongly towards the erroneous outlying measurement. The median is relatively insensitive: the outlier can become infinitely large while the median still remains close to the correct value (0).

A third possible way to use multiple receivers is to use only part of the receivers for computing the unknowns and the remaining receivers exclusively for error detection [Kovach97]. In this setup, the unknowns are in fact computed twice by two subsets of the receivers, while error detection is performed by checks on the mutual consistency of the subsets' outputs. Compared to the integrated solution the accuracy will be less as the noise is averaged out over fewer receivers. On the other hand, the complete independence of the monitor might offer improved integrity. Although the final performance of this setup might differ essentially from the performance of integrated architectures, the methodology of operation is not, and this paper will therefore not provide explicit equations.

4.2 Ground pseudorange correction computation

There is not much to add to the discussion on combining receiver outputs when considering the computation of the pseudorange corrections on the ground. Because the pseudorange corrections are measured straightforwardly by the ground receivers, the final correction to be uplinked is simply the combined output $\vec{z}_g$ from (4.1.1).

4.3 Airborne position computation

In the aircraft, the position of the GPS antenna is computed. The accuracy and integrity of this computation determine the actual LAAS performance. Due to the properties of GPS in combination with the primary use of a LAAS as a landing aid, it is usually the vertical position that is critical in the performance evaluations. Because it simplifies the problem without effecting the treated methodology, this paper will therefore focus on the vertical position only. Suffice it to say that a similar approach can be used to assess the horizontal performance as well.

The position can be computed from the combined outputs of the airborne receivers by using least squares or robust estimation techniques. Because receiver errors are already dealt with, least squares might be the preferred method. In that case, the position estimation becomes

$$\hat{\vec{x}}_a = S_a \cdot \vec{z}_a \tag{4.3.1}$$

with $S_a=(H_a^TH_a)^{-1}H_a^T$. When $\vec{s}_a^{\,T}$ is the row from S_a that corresponds to the vertical position, the error in the vertical component can be written as

$$\Delta\hat{v} = \vec{s}_a^{\,T} \cdot (\vec{e}_a + \vec{e}_g) \tag{4.3.2}$$

This equation shows that the vertical position error is a linear combination of the sum of the ground and airborne receiver errors. Although this would be the appropriate place to provide an equation for the pdf of the vertical position, this pdf depends on the pdfs of the combined receiver outputs. These pdfs on their turn depend on the implementation of the error detection and identification scheme. Discussion of this scheme is thus a necessary first step towards assessing the final position error pdf.

5. Error detection

In general, one could say that in order to detect an error, the mutual consistency of multiple signals that represent the same quantity is used. As soon as the signals disagree sufficiently, an error is assumed. To identify the source of an error as well, a related but slightly different approach is appropriate. The suspected erroneous source is isolated and compared to all other signals. When there is disagreement, the suspicion is assumed correct, and the signal is removed from the solution.

In the remainder of this section, the LAAS architecture that is currently being developed and investigated by Ohio University is taken as a starting point. Using similar methods, other architectures would be possible as well, but discussing all of them is beyond the scope of this paper.

5.1 Testing between receivers

As indicated before, the use of a least squares algorithm to combine the outputs of different receivers needs consistency checks to provide integrity. One of the important characteristics of the Ohio University architecture is that it uses a set of test statistics B_{xr}^i for each channel of each receiver to perform 'one step' error detection and identification. Failures in a specific receiver are detected by comparing their individual outputs to those of the other receivers. The test statistics used to check the output on the i^{th} channel of receiver xr can be written as:

$$B_{xr}^i = \tfrac{1}{M_x}\Big(\vec{z}_{xr}[i] - \tfrac{1}{M_x-1}\sum_{j=1,\,j\neq r}^{M_x} \vec{z}_{xj}[i]\Big) \tag{5.1.1}$$

and can be proven to be nothing but an estimate of the output error of receiver xr on channel i. An error is detected as soon as the estimated output error exceeds a certain limit h_{xr}^i:

$$\left|B_{xr}^i\right| > h_{xr}^i \Rightarrow error\ in\ receiver\, xr \tag{5.1.2}$$

When the multipath is modeled as either worst case deterministic or Gaussian, the test statistic B_{xr}^i will be distributed according to

$$B_{xr}^i \sim N(\mu_{xr}^i, \tfrac{1}{M_x^2 - M_x} Q_x[i,i]) \tag{5.1.3}$$

When the multipath is Gaussian, the multipath error is accounted for in Q_x and

$$\mu_{xr}^i = \tfrac{1}{M_x}\vec{\mu}_{xr}'^F - \tfrac{1}{M_x-1}\sum_{j=1,\,j\neq r}^{M_x} \vec{\mu}_{xj}'^F \tag{5.1.4}$$

while for worst case multipath modeling it becomes

$$\mu_{xr}^i = \tfrac{1}{M_x}(\vec{\mu}_{xr}'^{MP} + \vec{\mu}_{xr}'^F) - \tfrac{1}{M_x-1}\sum_{j=1,\,j\neq r}^{M_x} (\vec{\mu}_{xj}'^{MP} + \vec{\mu}_{xj}'^F) \tag{5.1.5}$$

Note that the given scheme identifies erroneous receivers on a 'per channel' basis. Additional support on the identification of erroneous receivers could be obtained by a check over all channels, for example by looking at the norm of the vector $\vec{z}_{xr-}\vec{z}_x$. When this norm is high for a certain receiver and remains high over a sufficiently long period, the xr^{th} receiver will probably have to be replaced. Of course, such 'replacement' decisions should also be made when one of the receiver channels is failing repeatedly.

When mid-value selection is used, error detection is performed implicitly. To identify a failing receiver, the same residual techniques that were discussed above can be used. Identification can also be done by monitoring whether the output of a certain receiver is selected approximately $100/M_x$ percent of the time. If one of the receivers is used significantly less often, this would indicate that it produces too many outlying values and should be replaced.

5.2 Additional tests

Although the tests in section 5.1 will form the 'core' of the LAAS integrity providence, it is possible to do some extra checks. These can deal with some of the failure scenarios that would not be covered by only comparing the outputs of different receivers.

In the air, the same AIM algorithms that are used to provide protection against failing satellites en route can be used in a LAAS setting as well. These algorithms test the mutual consistency of the corrected pseudoranges used in the position computation, and provide extra protection, for example in case of:

- Multiple receiver failures
- Datalink failures (corrupting the corrections)
- Common pseudorange errors, for example caused by radio frequency interference or excessive multipath

On the ground, there is no mutual consistency as all corrections are independent. However, the corrections can be checked 'as such'. Extremely large corrections still indicate a problem with the used models. Although the error *might* be common to the airborne user and the reference station, it might still be wise to send no correction. Highly dynamic behavior of the corrections can be another sign of potential problems. Even if the corrections would still be correct, high dynamics lead to a fast decorrelation of the corrections in time, which makes the final position solution sensitive to timing delays and errors.

6. Computing LAAS performance

The final decision whether it is safe to land using a LAAS system should be taken in the air. This decision will be based on the probability that the system delivers vertical position errors that are larger than the Vertical Alarm Limit (VAL) requirement. Only when this probability is extremely small, the system will be considered sufficiently safe. Usually, the computation is approached from the other side. Instead of computing the probability that the VAL is exceeded, the Vertical Protection Level (VPL) is computed. The VPL is defined as the position error that bounds the actual position with a given (high) probability. Only when the VPL is smaller than the VAL the system is safe to use.

In either case, the probability density function of the vertical position error is needed. This pdf depends on the pdfs of the different error sources *and* on the error detection and identification algorithms that are operational. Recall that the vertical position error is a linear combination of the pseudorange correction errors from the ground and the pseudorange errors onboard the airplane:

$$\Delta\hat{v} = \vec{s}_a^{\,T} \cdot (\vec{e}_a + \vec{e}_g) \tag{6.1}$$

where $\vec{e}_a$ and $\vec{e}_g$ are the output errors of a censored combination of multiple receiver outputs. It is exactly the censoring which makes it hard to calculate the pdfs of $\vec{e}_a$ and $\vec{e}_g$.

Suppose that a combination of least squares and error detection and identification is used. The removal of suspected erroneous measurements is a non-linear operation. Despite the attempts to describe the output errors of the individual receivers in terms of Gaussian distributions, the resulting pdf after error detection will still not be Gaussian. Mid-value selection is a nonlinear operation as well. The next paragraphs will illustrate the effects of the non-linearities and indicate which principles should be applied to assess the guaranteed system performance. In particular, this involves the determination of 'worst case' multipath and receiver failure biases. The worst case position error pdf will then have to be used to assess whether the probability that the VAL is exceeded is sufficiently small to continue using the system.

6.1 The pdf of the combined receiver outputs

This paragraph investigates what can be said about the pdfs of the combined output of multiple receivers using either least squares with error detection or mid-value selection. The first step is the determination of worst case multipath, the second to find the worst case receiver failure. The latter will problem will prove much harder, especially for the least squares case.

6.1.1 Worst case multipath

When looking at system integrity, worst case multipath is multipath that is common to all receivers. This multipath can not be detected, as it gives zero mean test statistics, and has the largest possible influence on the combined receiver output. This means that the effects of multipath are most destructive when all receivers suffer from the maximum amount of multipath. When this amount equals 20 centimeters, the combined receiver output suffers 20 centimeter errors in as well[2]. Mid-value selection does not protect against common failures either, and the effect on its output will be 20 centimeters as well.

6.1.2 Least squares and receiver failures

Consider the case in which the outputs if multiple receivers are combined using least squares with error detection and identification. The set of test statistics that looks for receiver errors in the i^{th} channel is $B^i_{x1}...B^i_{xM_x}$.

Only those measurements for which the test statistic is below a certain threshold h^i_{xr} will be used in the averaging process that generates the final output. Using a Boolean function $\beta^i_{xr} = \beta(B^i_{xr} \le h^i_{xr})$, the error in the combined output of the receiver can be written as[1]:

$$\vec{e}_x = \left(\sum_{r=1}^{M_x} \beta^i_{xr}\right)^{-1} \sum_{r=1}^{M_x} \beta^i_{xr}\vec{e}_{xr} \qquad (6.1.2.2)$$

Assume for simplicity that there are three receivers. Furthermore, assume that all receivers operate to specifications and thus $\vec{\mu}^F_{xr} = \vec{0}$. Usually, the detection threshold h^i_{xr} will be chosen such that the probability of detecting errors (false alarm) is small. Therefore, most of the time the combined receiver output will equal the mean of the three receiver outputs. [Liu97] shows that the pdf of the output in that case simply equals:

$$\text{No failure detected: } \vec{e}_x[i] \sim \mathrm{N}(\mu_x, \tfrac{1}{3}\mathrm{Q}_x[\mathrm{i,i}]) \qquad (6.1.1.3)$$

When the multipath is modeled as Gaussian it is incorporated in Q_x and thus $\mu_x = 0$, while for worst case multipath modeling $\mu_x = 0.2$.

When an error is unjustly detected and a measurement is removed, the situation becomes more complicated. [Liu97] provides expressions for the pdf of this case. In particular, it is shown that the pdf after removal of one measurement is no longer Gaussian. This is readily explained. Assume that the third measurement has been excluded. The probability that this happens will be higher the more the values of the other two measurements differ from the third, while the mean of those same two measurements becomes the combined output. This biases the output in the direction *opposite* to the excluded measurement.

When there are failures in the system, the situation becomes even more complicated. Assume that the third measurement comes from a failing receiver and is biased. As long as the bias remains *undetected,* the output becomes biased in the *same* direction as the bias in the third receiver. However, the larger the bias becomes, the likelier that it will be detected but the worse its effects when it remains undetected. Moreover, the detection is likelier when the other two measurements differ more from the biased one, just as in the no failure case. This will bias the pdf *after detection* in the direction *opposite* to the bias in the third receiver, but this final output bias will be smaller the larger the bias in the third receiver is.

To visualize what has been explained above, Figure 2 shows the behavior of the pdfs for the cases discussed. These results come from a simulation of 10^6 samples, in which the detection algorithm was tuned to a false alarm probability of 0.25. This value was deliberately chosen this high to ensure getting a significant amount of samples in all of the plots. All receiver errors were normally distributed with variance σ^2, while one of the receivers has been given a bias ranging between 0 and 6σ.

As discussed in section 4, there is usually no information on the bias in a failing receiver. The only way to overcome this lack of knowledge safely is to assume a 'worst case bias' every time a receiver fails. The question of the 'worst case' receiver failure still remains to be answered. 'Worse case' eventually means: 'affecting system integrity in the most adverse manner'.

What really makes the determination of the worst case bias tough is the non-Gaussian character of the errors. As the form of the pdf changes with the bias size, it is difficult to find an unambiguous way to rank the different pdfs from 'optimal' to 'worst case'. One of the possible ways of ordering could be to look at the pdf with the heaviest tail. Another approach could be to bound the pdf with a function that is easier to handle. The different shapes of the pdfs for different bias sizes suggest that it would really be hard to find an appropriate function though.

Obviously, this can hardly be called a satisfactory answer to the 'worst case failure' question. It is clear that further research is necessary to find ways to overcome the indicated problem.

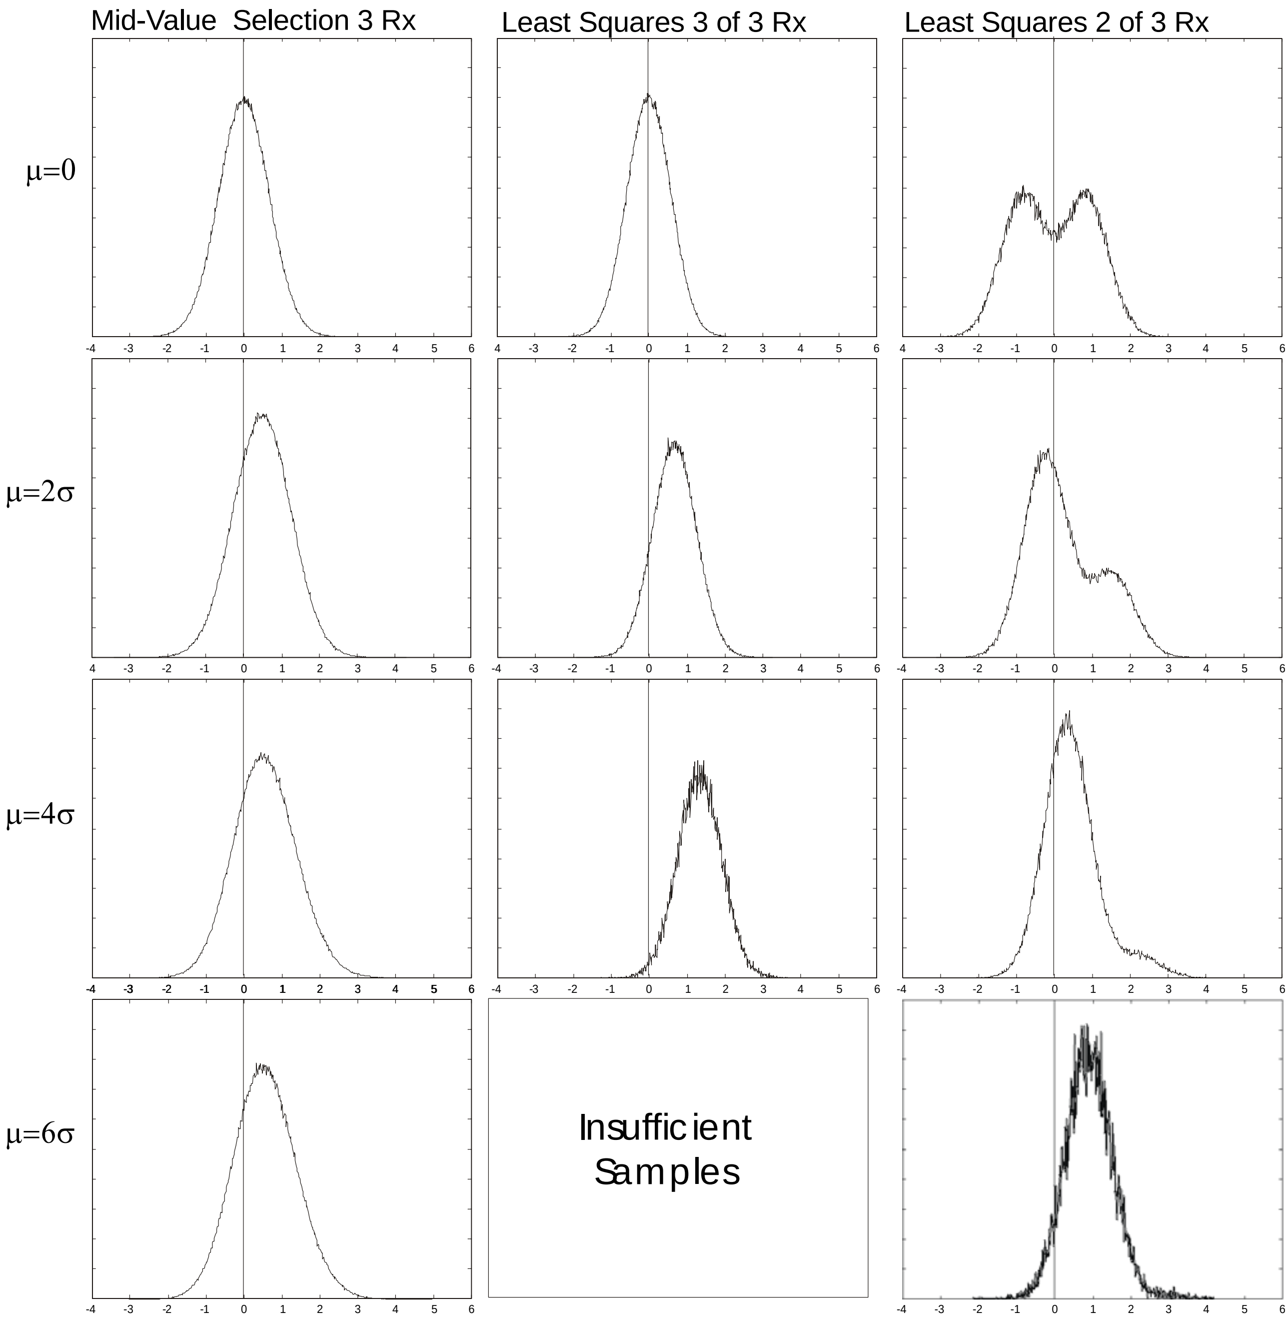

Figure 2. The probability density functions of combined receiver outputs. All receiver errors are normally distributed with variance σ^2. One of the receivers has a bias with values ranging from 0 to 6σ. The three columns correspond to mid-value selection, least squares with no error detected, and least squares with one error detected respectively. In the last case, the erroneous receiver has been removed from the solution.

6.1.3 Mid-value selection and receiver failures

Mid-value selection has characteristics that are quite different when compared to least squares. It really shines in case there are errors, but performs suboptimally when all errors are zero mean Gaussian.

First, consider the probability that the output error of the combined receivers exceed a certain number e_0 when all receivers function well. Of the three outputs, the one with the value that lies in between the two others is selected. This implies that *at least two* of the three receiver outputs have to be in error by more than e_0. Using the equal

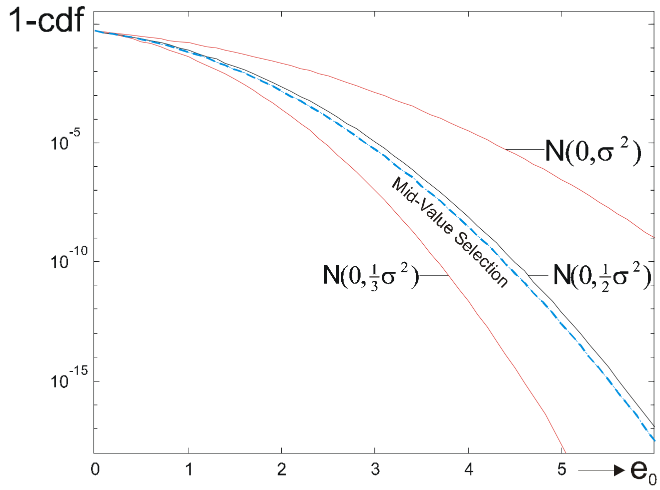

Figure 3. One minus the cumulative distribution function of the combined output of three correctly operating receivers using mid-value selection, compared to those of three different Gaussian distributions.

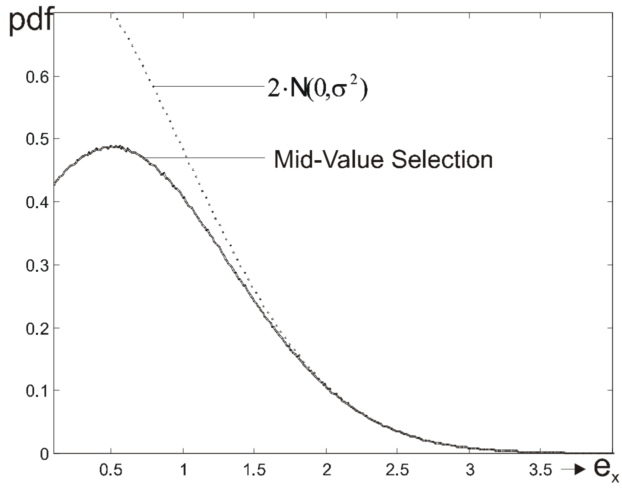

Figure 4. The probability density function of the combined output of three receivers of which one failing using mid-value selection, compared to twice the Gaussian distribution.

statistical properties of all three receivers, the following equation is obtained [David70]:

$$P_0(\vec{e}_x[i] > e_0) = 3P(\vec{e}_{x1}[i] > e_0)^2 - 2P(\vec{e}_{x1}[i] > e_0)^3 \tag{6.1.3.1}$$

Figure 3 shows a plot of the values of (6.1.3.1) as compared to corresponding values of three different Gaussian distributions. The plot suggests that

$$no\ \ failures:\ \vec{e}_x[i] \,\tilde{<}\, \mathrm{N}(0, \tfrac{1}{2}Q_{xr}[i,i]) \tag{6.1.3.2}$$

is a reasonable upperbound on the probability density function of $\vec{e}_x[i]$. Note that the variance is one and a half times larger than in the least squares case when no detection occurs. This is the price paid for the advantage of error insensitivity and the impossibility of false detections.

Now consider the probability that the output error of the combined receivers exceeds a certain number e_0 *while one receiver is in error*. This probability is largest when the erroneous output always exceeds e_0. In that case, the combined output will also exceed e_0 when *at least one* of the remaining two receivers is in error by more than e_0. The following result is therefore obtained:

$$P_1(\vec{e}_x[i] > e_0) = 2P(\vec{e}_{x1}[i] > e_0) - P(\vec{e}_{x1}[i] > e_0)^2 \tag{6.1.3.3}$$

For sufficiently large e_0 the last term becomes negligible, which leads to the following upperbound on the pdf of the combined output:

$$one\ \ failure:\ \vec{e}_x[i] \,\tilde{<}\, 2\mathrm{N}(0, Q_{xr}[i,i]) \tag{6.1.3.4}$$

Note that for small values of e_0 (6.1.3.4) is conservative by at most a factor 2, while the bound becomes tighter when e_0 grows. Values of the upperbound from (6.1.3.4) are plotted in Figure 4 together with the probability density function of $\vec{e}_x[i]$ obtained from a large simulation (10^6 samples).

The results (6.1.3.3) and (6.1.3.4) can be combined to get a *fully worst case* upperbound on the vertical position error pdf. As discussed earlier, the vertical position is a weighted sum of $2N_C$ channel output errors:

$$\Delta\hat{v} = \vec{s}_a^{\,T} \cdot \vec{e}_a + \vec{s}_a^{\,T} \cdot \vec{e}_g \tag{6.1.3.5}$$

When all receivers – both on the ground and in the air – are operating correctly, (6.1.3.2) can be applied to all errors. In that case, an upperbound on the vertical position error is obtained as

$$\Delta\hat{v} \,\tilde{<}\, \mathrm{N}\left(0, \tfrac{1}{2}\vec{s}_a[i]^2 \sum_{i=1}^{N_C} \left(Q_{ar}[i,i] + Q_{gr}[i,i]\right)\right) \tag{6.1.3.6}$$

When one of the ground receivers is in error, (6.1.3.4) has to be used for the ground segment and (6.1.3.2) for the airborne part:

$$\Delta\hat{v} \tilde{<} 2^{N_C} \mathrm{N}\left(0, 2^{N_C - 1} \vec{s}_a[i]^2 \sum_{i=1}^{N_C} \left(Q_{ar}[i,i] + 2Q_{gr}[i,i]\right)\right) \quad (6.1.3.7)$$

The two factors 2^{N_C} originate in the convolution of the N_C error pdf bounding functions (6.1.3.4), see Appendix A. In a similar fashion, the following expression for the case in which both a ground and an airborne receiver failing is obtained:

$$\Delta\hat{v} \tilde{<} 2^{2N_C} \mathrm{N}\left(0, 2^{2N_C} \vec{s}_a[i]^2 \sum_{i=1}^{N_C} \left(Q_{ar}[i,i]^2 + Q_{gr}[i,i]\right)\right) \quad (6.1.3.8)$$

The probability that the vertical position error exceeds the VAL can be found by using a weighted sum of the given three different pdfs. The weights should equal the probability that a certain pdf is valid at a certain moment in time, see for example [Ober97].

7. Least squares and mid-value selection compared

Having discussed both least squares and mid-value selection it might be useful to summarize the main differences between the methods shortly. The operational differences have been summarized in Table 1. Most, but not all of these differences have been addressed in the text already. Table 2 gives some indication of the performance differences.

Because a rigorous method to evaluate the worst case bias performance of the least squares with error detection has not been found yet, it is currently impossible to assess the operational performance differences. However, current methods exist to do a performance assessment. These simulation figures to be presented here compare the mid-value selection vertical protection limits to those computed with algorithms developed by the RTCA. These algorithms use least squares and error detection as described in this paper, but include no 'rigorous' worst case bias principles. The main reason for presenting the simulation result is therefore to provide a limited indication of the performance that can be expected from mid-value selection as compared to a well-studied and available method.

The RTCA algorithms and their parameters are described in [VanDyke97] that gives VPL expressions for the fault free and single failure case. The latter uses the actual values of the test statistics (5.1.3) (see also the comment in the first note). Substitution of the stochastic model for the test statistics in the error-free case gives a so-called 'predicted' VPL: the VPL that is expected to be computed in the airplane when the system is operating well. Therefore, the predicted VPL determines the continuity of service that can be expected from a correctly functioning system. The airborne segment in the expression in [VanDyke97] is modeled as if it were operating a 'single receiver' only. Therefore, the expressions were slightly adapted to incorporate three correctly working airborne receivers instead of just one. All parameters for three receivers and Cat II conditions have been taken from [VanDyke97]. About 400,000 samples (equally spread over 24 hours and the globe) have been computed. Multipath has been modeled as a Gaussian disturbance.

Least Squares	Mid-Value Selection
Separate error detection algorithm	No error detection required
Optimal accuracy	Reduced accuracy
Hard to assess 'worst case' biases	Readily computable 'worst case' biases
High influence of undetected errors	Small influence of undetected errors
Position jump after detections	Graceful degradation

Table 1. Differences between least squares and mid-value selection

VPL Method	Availability	Mean VPL
No failures – LS (RTCA)	0.9990	1.31
Predicted (RTCA)	0.9985	2.16
No failures – Mid Value Sel. (6.1.3.6)	0.9986	2.06
1 Ground Rx failure - Mid Value Sel. (6.1.3.7)	0	38
1 Ground 1 Air Rx failure - Mid Value Sel. (6.1.3.8)	0	596

Table 2. RTCA Algorithms compared to Mid-Value selection

As expected, the no failure VPL computation of the RTCA gives the best VPL values because of the reduced amount of noise as compared to mid-value selection. The ‘no failure’ VPL using mid-value selection is comparable with the RTCA’s ‘predicted VPL’ values. When the probability of a receiver failure is small, this value is a good indication of the overall system performance as well, due to the fact that the VPL is not that much larger than the ‘no failure’ VPL.

8. Concluding remarks

This paper has described two different ways to exploit signal redundancy to provide system integrity for locally augmented GPS. Most attention has been paid to the way the output of multiple receivers can be combined into a single value that should be as error-free as possible.

Because of the extremely high requirements, the modeling of all signal disturbances must be done with great care. Unfortunately, available multipath models do not seem to be fully validated yet. Therefore, the discussion in this paper considered three possible models when discussing the LAAS integrity methodology.

The least squares scheme has to be combined with error detection and removal algorithms in order to provide integrity. One of the main difficulties is that the combined output signal after censoring by the error detector is hard to describe. This makes the determination of ‘worst case’ biases difficult, while this is the only way to get a tight upperbound on the overall performance. Further research will be required to resolve this issue.

The mid-value selection scheme has proven to be much easier to describe and bounds on its performance have been given. This algorithm provides less accuracy than the least squares solution. On the other hand, its sensitivity to errors is reduced, and no separate error detection algorithm is necessary.

NOTES

[1] One could argue whether one should employ the actual values of the B^i_{xr} or just the fact that they are smaller or larger than h^i_{xr} . Only exploiting the values of β^i_{xr}, as has been done here, has the advantage that many computations can be done off-line. However, the B^i_{xr} values provide more information on the actual system performance.

[2] System *continuity* will be affected most when the multipath has opposite values on the different receivers. This will cause a maximum value of the test statistic while having no influence on the combined output at all.

ACKNOWLEDGMENT

Tim Murphy is gratefully acknowledged for his comments.

REFERENCES

[David70] David, H.A.; “Order Statistics”, John Wiley and Sons, 1970

[DeCleene97] DeCleene, B., T. Murphy and T. Skidmore; “Specifics of the LAAS Context Model G(x)”, Proceedings of the 10th International Technical Meeting of the Satellite Division of the Institute of Navigation, September 1997, pp. 1583-1592

[Jaynes94] Jaynes, E.T.; “Probability Theory: The Logic of Science”, Fragmentary Edition of June 1994, http://omega.albany.edu:8008/JaynesBook.html

[Kovach97] Kovach, K. and L. Huffman; “SCAT-I Integrity Method: Detection or Estimation”?, Proceedings of the 53rd Annual Meeting of the Institute Of Navigation, June 1997, pp. 103-120

[Liu97] Liu, F., T. Murphy and T. Skidmore; “LAAS Signal-in-Space Integrity Monitoring Description and Verification Plan”, Proceedings of the 10th International Technical Meeting of the Satellite Division of the Institute of Navigation, September 1997, pp. 485-498

[Murphy97] Murphy, T., B. DeCleene and T. Skidmore; “G(x)”, Proceedings of the National Technical Meeting of the Institute of Navigation, January 1997, pp. 475-484

[Ober97] Ober, P.B.; “Autonomous Integrity Monitor Performance”, Proceedings of GNSS '97, first European Symposium on Global Navigation Satellite Systems, April 21-25, 1997, pp. 139-152. Electronic copy of this paper is obtainable at http://wwwtvs.et.tudelft.nl/people/bas/bas.htm

[RTCA96] Recommended LAAS Architecture, presented to RTCA SC-159 WG-4, Ray Swider and Bruce DeCleene, November 12, 1996

[Skidmore96] Skidmore, T., F. van Graas and F. Liu; “Ohio University LAAS Monitoring: Design and Performance”, Proceedings of the 9th International Technical Meeting of the Satellite Division of the Institute of Navigation, September 1996, pp.1615-1621

[VanDyke97] Van Dyke, K., Q. Hua, M. Sams and S. Rowson; “Precision Approach Availability and Integrity Requirements”, Proceedings of the 53rd Annual Meeting of the Institute Of Navigation, June 1997, pp. 683-690

[VanGraas97] Van Graas, F., D.W. Diggle, M. Uijt de Haag, T.A. Skidmore and M.F. DiBenedetto; “Ohio University/FAA Flight Test Demonstration Results of the Local Area Augmentation System (LAAS)”, Proceedings of the 10th International Technical Meeting of the Satellite Division of the Institute of Navigation, September 1997, pp. 1623-1630

APPENDIX A

This appendix will show how the convolution of the pdfs of all contributions to the vertical position error is used to derive (6.1.3.6) - (6.1.3.8). In particular, the zero'th, first and second moment of these pdfs are computed (all higher order moments are zero and can be left out of this discussion). After that, the behavior of these moments under convolution is exploited to find the moments of the bound on the vertical position error pdf.

When all receivers work well, the contribution of the i^{th} channel to the vertical position error from either the ground or the airborne segment has been shown to be bounded by (see section 6)

$$\vec{s}_a[i]\vec{e}_x[i] \,\tilde{<}\, f_{no\_err,x,i} \tag{A.1}$$

in which

$$f_{no\_err,x,i} = N(0, \tfrac{1}{2}\vec{s}_a[i]^2 Q_{xr}[i,i]) \tag{A.2}$$

When one receiver is in error, this contribution is bounded by

$$\vec{s}_a[i]\vec{e}_x[i] \,\tilde{<}\, f_{one\_err,x,i} \tag{A.3}$$

instead, with

$$f_{one\_err,x,i} = 2N(0, \vec{s}_a[i]^2 Q_{xr}[i,i]) \tag{A.4}$$

Using the notation of [Jaynes94] for the zero'th, first and second moment of a real function $g(x)$:

$$\begin{aligned} Z_g &= \int_{-\infty}^{\infty} g(x)dx \\ F_g &= \int_{-\infty}^{\infty} x \cdot g(x)dx \\ S_g &= \int_{-\infty}^{\infty} x^2 \cdot g(x)dx \end{aligned} \tag{A.5}$$

the moments of the pdf bounding functions can easily be derived from the moments of the normal distribution. When $g(x) = \mathrm{N}(\mu,\sigma^2)$, its zero'th, first and second moment are known to equal 1, μ, and σ^2 respectively. Note that when all higher order moments are zero, this is also true the other way around, and these values uniquely determine a normal pdf as well. The bound (A.2) found on the 'no receiver errors' behavior of the mid-value selection scheme is a zero mean normal pdf, and the following result for the 'no errors' bound on the pdf on the output of the i^{th} channel is therefore obtained immediately:

$$\begin{aligned} Z_{f_{no\_err,x,i}} &= 1 \\ F_{f_{no\_err,x,i}} &= 0 \\ S_{f_{no\_err,x,i}} &= \tfrac{1}{2}\vec{s}_a[i]^2 Q_{xr}[i,i] \end{aligned} \tag{A.6}$$

Using the fact that the bound (A.4) in the 'one failing receiver' situation is nothing but *twice* a normal pdf, the moments for this case are readily computed as:

$$\begin{aligned} Z_{f_{one\_err,x,i}} &= 2 \\ F_{f_{one\_err,x,i}} &= 0 \\ S_{f_{one\_err,x,i}} &= 2\vec{s}_a[i]^2 Q_{xr}[i,i] \end{aligned} \tag{A.7}$$

[Jaynes94, Appendix C] shows how the moments of the convolution of two functions (that are not necessarily pdfs!) can be computed from the moments of these functions. Generalizing the results to the convolution of N functions $g_i(x)$ with all first moments zero the expressions become

$$Z_{g_1(x)*g_2(x)*\ldots*g_N(x)} = \prod_{i=1}^{N} Z_{g_i} \tag{A.8}$$

$$F_{g_1(x)*g_2(x)*\ldots*g_N(x)} = 0 \tag{A.9}$$

$$S_{g_1(x)*g_2(x)*\ldots*g_N(x)} = \sum_{i=1}^{N}\left(S_{g_i(x)} \prod_{j=1,\, j\neq i}^{N} Z_{g_i(x)} \right) \tag{A.10}$$

It follows from (6.1.3.5) that in the case of no receiver errors, $pdf(\Delta\hat{v})$ is bounded by the convolution of N_C functions $f_{no\_err,a,i}$ and N_C functions $f_{no\_err,g,i}$. With (A.8)-(A.10), the first three moments of the bound can be computed to equal

$$Z_{pdf(\Delta\tilde{v})} = 1 \tag{A.11}$$

$$F_{pdf(\Delta\tilde{v})} = 0 \tag{A.12}$$

$$S_{pdf(\Delta\tilde{v})} = \sum_{i=1}^{N_C}\left(\tfrac{1}{2}s_a[i]^2 Q_{ar}[i,i] + \tfrac{1}{2}s_a[i]^2 Q_{gr}[i,i]\right) \tag{A.13}$$

When one of the ground receivers is in error, the bound on $pdf(\Delta\hat{v})$ becomes the convolution of N_C functions $f_{no\_err,a,i}$ and N_C functions $f_{one\_err,g,i}$, and will thus have moments

$$Z_{pdf(\Delta\tilde{v})} = 2^{N_C} \tag{A.14}$$

$$F_{pdf(\Delta\tilde{v})} = 0 \tag{A.15}$$

$$S_{pdf(\Delta\hat{x})} = 2^{N_C} \sum_{i=1}^{N_C} \left(\tfrac{1}{2} s_a[i]^2 Q_{ar}[i,i] + \vec{s}_a[i]^2 Q_{gr}[i,i] \right) \tag{A.16}$$

Finally, when both an airborne and a ground receiver are failing, the convolution of N_C functions $f_{err,a,i}$ and N_C functions $f_{one\_err,g,i}$ gives bounds for $pdf(\Delta\hat{v})$ with moments

$$Z_{pdf(\Delta\tilde{v})} = 2^{2N_C} \tag{A.17}$$

$$F_{pdf(\Delta\tilde{v})} = 0 \tag{A.18}$$

$$S_{pdf(\Delta\tilde{v})} = 2^{2N_C} \sum_{i=1}^{N_C} \left(s_a[i]^2 Q_{ar}[i,i] + \vec{s}_a[i]^2 Q_{gr}[i,i] \right) \tag{A.19}$$