A Multiple Hypothesis Approach to Satellite Navigation Integrity

BORIS S. PERVAN, SAMUEL P. PULLEN, and JOCK R. CHRISTIE

Stanford University, Stanford, California
Received December 1997
Revised March 1998

ABSTRACT

A general methodology has been developed for failure mitigation in satellite-based navigation systems. This new approach is founded on the direct evaluation of integrity risk under the unified consideration of all single-element failure hypotheses and the no-failure hypothesis. It is applicable to fault-tolerant estimation and integrity monitoring. In particular, the algorithm has been investigated for application as an airborne element of the Local Area Augmentation System (LAAS) integrity monitoring architecture. In this system, GPS ranging measurements from multiple ground-based reference receivers are to be used to mitigate the effects of reference receiver failures. The multiple hypothesis algorithm has been evaluated through analysis and simulation and has been shown to provide the tightest realizable protection limits for LAAS. The operational performance of the new algorithm has been compared with that of the approach currently under consideration by RTCA Special Committee-159. The performance of the multiple hypothesis algorithm has been demonstrated using experimental data collected during a series of LAAS prototype flight tests in a Beechcraft King Air.

NAVIGATION: *Journal of The Institute of Navigation*
Vol. 45, No. 1, Spring 1998
Printed in U.S.A.

INTRODUCTION

The management of measurement redundancy in satellite navigation systems has been a subject of great interest in recent years [1–3]. This has been especially true for aircraft precision approach and landing navigation [4–7], where stringent integrity and continuity requirements have demanded the highest level of navigation system performance. Under Category III visibility conditions, it is necessary that the requirements for integrity risk (10^{-9} per approach) and continuity risk (2×10^{-6} for any 15 s of the approach) be satisfied with respect to a vertical alert limit of approximately 5 m and a time-to-alarm of 1–2 s [8–10]. Compliance with the time-to-alarm specification in particular constrains the nature of fault management to a 'snapshot' approach, since the potential fault-detection benefits associated with post-failure noise averaging are precluded.

In this context, a new position-domain methodology has been developed for failure mitigation in satellite navigation systems. The approach is based on the direct computation of integrity risk under the combined consideration of all single-element failure hypotheses and the no-failure hypothesis. (Integrity risk, in this regard, is defined as the likelihood that the position estimate error exceeds a prespecified alarm limit.) The new algorithm provides greater flexibility and generality than existing fault management approaches and, significantly, introduces a new capability for fault-tolerant estimation.

This paper is organized as follows. First, the mathematical development of the generalized algorithms for fault-tolerant estimation and integrity monitoring is presented. This discussion is followed by a detailed application of the methodology to integrity monitoring for the Federal Aviation Administration (FAA) Local Area Augmentation System (LAAS) differential GPS architecture (for application to aircraft precision approach and landing navigation). The achievable vertical protection limits are analytically derived, and the resulting LAAS operational availability is quantified through simulation. These results are compared with the performance of the LAAS monitor algorithm presently under consideration by RTCA, Inc. Finally, experimental data are presented from LAAS prototype flight trials performed in a National Aeronautics and Space Administration (NASA) Beechcraft King Air at Moffett Federal Airfield in California.

THEORY

Consider the generalized linear observation equation

$$\mathbf{z} = \mathbf{C}\mathbf{x} + \boldsymbol{\nu} \tag{1}$$

where $\mathbf{x}$ is the $m \times 1$ vector of parameters to be estimated (including three position elements), $\mathbf{z}$ is the $n \times 1$ ($n > m$) measurement vector, and $\mathbf{C}$ is the $n \times m$ observation matrix. The $n \times 1$ vector $\boldsymbol{\nu}$ represents the unknown error in the measurement vector, which under fault-free conditions is normally distributed with zero mean and covariance $\mathbf{V}$ ($n \times n$).

In this paper, we consider the vertical component of position error exclusively (as it will generally have the smallest alert limit); the horizontal components may be treated in an identical manner. In this case, for an *arbitrary* estimate of vertical position ($\hat{x}_v$), we define the vertical integrity risk as

$$I_v|\mathbf{z} \equiv P(|x_v - \hat{x}_v| > VAL|\mathbf{z}) \tag{2}$$

where x_v is the true vertical position, and VAL is the prespecified vertical alert limit. Expanding the righthand side of equation (2), we can write

$$I_v|\mathbf{z} \equiv P(x_v - \hat{x}_v > VAL|\mathbf{z}) + P(x_v - \hat{x}_v < -VAL|\mathbf{z}) \tag{3}$$

If $f_{x_v}(\tilde{x}_v|\mathbf{z})$ is the probability density function of the true vertical position given the measurement set $\mathbf{z}$, then equation (3) may be written as

$$I_v|\mathbf{z} \equiv \int_{\hat{x}_v+VAL}^{\infty} f_{x_v}(\tilde{x}_v|\mathbf{z})d\tilde{x}_v + \int_{-\infty}^{\hat{x}_v-VAL} f_{x_v}(\tilde{x}_v|\mathbf{z})d\tilde{x}_v \tag{4}$$

We now define H_0 as the no-failure hypothesis and the failure hypotheses as H_i ($i = 1, \cdots, n$), corresponding, respectively, to failure in each of the n measurement elements in the system defined by equation (1). It will be assumed that the prior probability of simultaneous failures on multiple elements is much smaller than the integrity risk requirement under consideration. Therefore, the probability density function of the state vector $\mathbf{x}$ given the measurements $\mathbf{z}$ may be expressed as

$$f_{x_v}(\tilde{x}_v|\mathbf{z}) = \sum_{i=0}^{n} f_{x_v}(\tilde{x}_v|H_i, \mathbf{z})\, P(H_i) \tag{5}$$

where $P(H_i)$ is the prior probability of failure of element i.

In general, no prior knowledge on the distribution of failure magnitudes is available. In this case, the conditional probability density of the true vertical position, given the measurement set $\mathbf{z}$ and that element i (only) has failed, is not a function of measurement i. For example, under the assumption that the measurement errors are normally distributed, the conditional probability density functions on the righthand side of equation (5) are given by

$$f_{x_v}(\tilde{x}_v|H_i, \mathbf{z}) = N_{x_v}[\hat{x}_{v_i}, \sigma_{v_i}^2] \equiv \frac{1}{\sqrt{2\pi\sigma_{v_i}^2}} \exp\left[-\frac{(\tilde{x}_v - \hat{x}_{v_i})^2}{2\sigma_{v_i}^2}\right] \tag{6}$$

where $\hat{x}_{v_i}$ is the maximum-likelihood vertical position estimate, and $\sigma_{v_i}^2$ is its associated vertical estimate error variance that results when the element corresponding to failure hypothesis H_i is removed from the observation equation (1). Note that for H_0, the full measurement set is applicable.

Using equations (5) and (6), equation (4) can now be written as

$$I_v|\mathbf{z} = \sum_{i=0}^{n} P(H_i)\left\{\int_{\hat{x}_v+VAL}^{\infty} N_{x_v}[\hat{x}_{v_i}, \sigma_{v_i}^2]d\tilde{x}_v + \int_{-\infty}^{\hat{x}_v-VAL} N_{x_v}[\hat{x}_{v_i}, \sigma_{v_i}^2]d\tilde{x}_v\right\} \tag{7}$$

We note that given $\mathbf{z}$, $\mathbf{C}$ and $\mathbf{V}$, and a set of known prior state probabilities, the only unknown on the righthand side of equation (7) is the (as yet arbitrary) estimate of vertical position $\hat{x}_v$. While it is clear that $\hat{x}_{v_0}$ represents the maximum likelihood estimate of the vertical position under the no-failure hypothesis, it is not necessary that we choose $\hat{x}_v = \hat{x}_{v_0}$. Instead, we make the following definitions:

$$\Delta_v \equiv \hat{x}_v - \hat{x}_{v_0}$$
$$\Delta_{v_i} \equiv \hat{x}_{v_i} - \hat{x}_{v_0} \quad (\Delta_{v_0} \equiv 0) \tag{8}$$

Note that Δ_v represents the deviation in the arbitrary vertical position estimate from the no-fault estimate. Substituting these definitions into equation (7), we obtain (see the appendix)

$$I_v|\mathbf{z} = \frac{1}{2}\sum_{i=0}^{n} P(H_i)\left\{\text{erfc}\left[\frac{VAL - (\Delta_{v_i} - \Delta_v)}{\sqrt{2}\sigma_{v_i}}\right] + \text{erfc}\left[\frac{VAL + (\Delta_{v_i} - \Delta_v)}{\sqrt{2}\sigma_{v_i}}\right]\right\} \tag{9}$$

which is the expression defining integrity risk as a function of Δ_v. We will see shortly that selecting $\Delta_v = 0$ does not generally minimize integrity risk.

Equation (9) represents the weighted sum of $n + 1$ normal density function 'tail' areas as indicated graphically in Figure 1. This figure illustrates qualitatively a simple hypothetical scenario in which $n = 3$. The four normal density functions shown correspond to the three failure hypotheses (H_1, H_2, H_3) and the no-failure

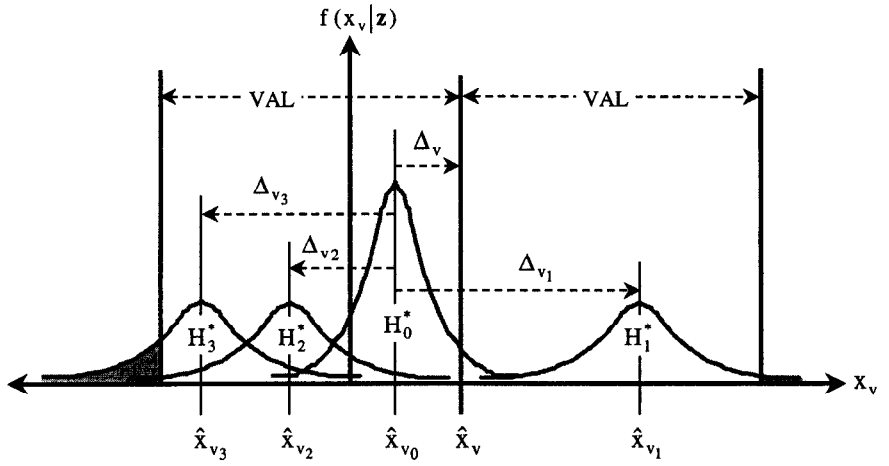

Fig. 1 – Multiple Hypothesis Concept

hypothesis (H_0). The prominence of the H_0 density function relative to the three failure density functions is intended to signify the relative prior probabilities of the associated hypotheses; in general, the prior probability of failure will be low.

Fault-Tolerant Estimation

As suggested above, it is possible to choose a nonzero value of Δ_v that lowers integrity risk. In fact, integrity risk may be minimized by the appropriate selection of Δ_v. Differentiating equation (9) with respect to Δ_v and setting the result equal to zero yields (see the appendix)

$$\sum_{i=0}^{n} \frac{P(H_i)}{\sigma_{v_i}} \left\{ \exp\left[-\left(\frac{VAL + \Delta_{v_i} - \Delta_v}{\sqrt{2}\sigma_{v_i}} \right)^2 \right] - \exp\left[-\left(\frac{-VAL + \Delta_{v_i} - \Delta_v}{\sqrt{2}\sigma_{v_i}} \right)^2 \right] \right\} = 0 \quad (10)$$

Equation (10) may be solved iteratively for Δ_v, which, in turn, may be substituted into equation (8) to obtain the minimum-integrity-risk estimate of vertical position:

$$\hat{x}_v = \hat{x}_{v_0} + \Delta_v \quad (11)$$

It is significant that this result is a *fault-tolerant* estimate, as it is the value of vertical position that minimizes integrity risk under the unified consideration of all single-element failure hypotheses and the no-fault hypothesis.

Integrity Monitor

Given a vertical position estimate (equation (11)), obtained through the solution of equation (10) (or, alternatively, by choosing the no-fault estimate where $\Delta_v = 0$), vertical integrity risk may be directly computed using equation (9). A hazardous condition is declared when

$$I_v|\mathbf{z} > I_v^{req} \quad (12)$$

where I_v^{req} is the predefined vertical integrity risk requirement. In qualitative terms, we can say that in this case, the use of the estimate $\hat{x}_v$ is inconsistent with the measurement set $\mathbf{z}$ by an amount greater than that permitted by the integrity specification. As an alternative, it is also possible to (iteratively) compute the *minimum* VAL, given the measurement set $\mathbf{z}$, such that equality exists in equation (12). We define this minimum VAL as the *statistical* vertical protection limit (VPL_s). Further discussion on VPL_s and treatment of the fault-free alarm rate resulting from equation (12) is deferred to the next section.

APPLICATION

The multiple hypothesis approach described above is general in the sense that it can be applied for snapshot redundancy management in a wide variety of satellite navigation architectures (including receiver autonomous integrity monitoring [RAIM] applications). Below we discuss in detail an example of how this methodology can be applied as an airborne element of a LAAS integrity monitor.

LAAS Integrity Monitoring Background

The FAA LAAS is a differential GPS architecture that will provide the means for aircraft precision approach and landing using satellite navigation. In general, the goal of the LAAS integrity monitoring subsystem 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 allocate accountability between the ground and airborne navigation subsystems. In general, the aircraft is responsible for the proper functioning of the airborne equipment (including the implementation of redundant sensor tracks to provide the means for detection and removal of airborne equipment failures), while the LAAS service provider 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. In the most stringent case (Category III), it is necessary that the LAAS integrity monitoring architecture meet the following SIS requirements [8–10]:

- *Integrity risk* of 10^{-9} per approach with respect to a vertical alert limit of approximately 5 m and a time-to-alarm of 1–2 s.
- *Continuity risk* (alarm rate) of 2×10^{-6} for any 15 s of the approach.

The application under consideration in the remainder of this paper is the detection and removal of failures in the LAAS ground segment. Such events are ultimately the source of anomalous data broadcast to the aircraft. To this end, multiple reference receivers are to be implemented in the baseline LAAS architecture. Specifically, at least three receivers are required (for Category III) to provide the desired capability for both detection and removal of a failure in a single reference receiver.

As currently envisioned, LAAS integrity monitoring comprises both ground and airborne elements. On the ground, a screening function will be implemented to isolate and remove reference receiver ranging failures by direct range measurement comparison between reference receivers (see, for example, [4]).

The protection limits resulting from known thresholds in the ground-based screening function can be computed in the aircraft. However, if the aircraft does not have access to the actual deviations between the multiple reference measurements, it must always be assumed that all measurement differences are at the maximum permitted by the ground threshold. (This must be assumed since such a condition is, in fact, possible.) As a result, the computed protection limits at the aircraft are generally quite conservative and thus reduce LAAS availability. The conservatism of this approach is alleviated by sending to the aircraft the actual differences between the multiple reference range measurements in addition to the nominal differential range correction (which is based on the average of all valid ground receiver measurements). In fact, it is for this reason that the currently proposed LAAS range measurement message [11] has been designed to include this additional data. With this data, it is possible to reconstruct at the aircraft the measurements obtained at all reference receivers. In this context the remainder of this paper focuses on the application of the multiple hypothesis approach to the airborne function of the LAAS integrity monitor.

Multiple Hypothesis Application Fundamentals

Our development considers the case of three sets of reference receiver measurements that, when processed separately at the aircraft, generate three vertical position estimates defined by z_1, z_2, and z_3, respectively. If these individual vertical position estimates are interpreted as individual 'measurements' of vertical position, their relation to the true vertical position (x_v) is given by

$$\begin{bmatrix} z_1 \\ z_2 \\ z_3 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} x_v + \begin{bmatrix} \nu_{g_1} + \nu_a \\ \nu_{g_2} + \nu_a \\ \nu_{g_3} + \nu_a \end{bmatrix}, \tag{13}$$

where

$$\nu_a \sim N(0, \sigma_{\nu_a}^2) \tag{14}$$

is the component of vertical position error due to measurement errors at the aircraft receiver, and ν_{g_i} is the component of vertical position error due to measurement errors at reference receiver i. Under normal error conditions (i.e., no reference receiver failure)

$$\nu_{g_i} \sim N(0, \sigma_{\nu_g}^2) \tag{15}$$

It is assumed that ν_{g_1}, ν_{g_2}, and ν_{g_3} are statistically independent (approximately true if the reference stations are spatially separated sufficiently to decorrelate ground multipath) and identically distributed under fault-free conditions. However, note that the vertical position 'measurements' z_1, z_2, and z_3 clearly are not independent because they all include the common airborne error term ν_a.

Under the hypothesis of no reference receiver failure (H_0), the maximum-likelihood estimate of vertical position is simply

$$\hat{x}_{v_0} = \tfrac{1}{3}(z_1 + z_2 + z_3) \tag{16}$$

and its associated error variance is

$$\sigma_{v_0}^2 = \tfrac{1}{3}\sigma_{\nu_g}^2 + \sigma_{\nu_a}^2 \tag{17}$$

Similarly, the maximum-likelihood vertical

position estimate under the hypothesis that reference receiver i has failed (H_i) is

$$\hat{x}_{v_i} = \frac{1}{2}\sum_{\substack{j=1 \\ j\neq i}}^{3} z_j, \tag{18}$$

with the associated error variance

$$\sigma^2_{v_i} = \frac{1}{2}\sigma^2_{v_g} + \sigma^2_{v_a} \tag{19}$$

Note that the estimate error variance (equation (19)) is identical for all i, so $\sigma^2_{v_1} = \sigma^2_{v_i}$.

Substituting the system of equations (13) into equations (16) and (18) and differencing to construct the discriminators defined in equation (8) results in

$$\begin{bmatrix}\Delta_{v_1}\\ \Delta_{v_2}\\ \Delta_{v_3}\end{bmatrix} = \frac{1}{6}\begin{bmatrix}-2 & 1 & 1\\ 1 & -2 & 1\\ 1 & 1 & -2\end{bmatrix}\begin{bmatrix}v_{g_1}\\ v_{g_2}\\ v_{g_3}\end{bmatrix} \tag{20}$$

Note that the discriminators are only a function of the reference station error components since v_a has been removed by differencing. In addition, because

$$\Delta_{v_3} = -(\Delta_{v_1} + \Delta_{v_2}) \tag{21}$$

only two of the discriminators are actually independent. (It is also noteworthy that the Δ_{v_i} discriminators [equation (20)] may be interpreted as the RTCA-defined 'B-values' [12] expressed in the vertical position domain.)

Fault-Free Conditions

Under the normal error conditions described by equation (15), it can easily be shown using equation (20) that the discriminator means are zero, and the discriminator covariance matrix is

$$V_\Delta \equiv \begin{bmatrix}E(\Delta^2_{v_1}) & E(\Delta_{v_1}\Delta_{v_2})\\ E(\Delta_{v_1}\Delta_{v_2}) & E(\Delta^2_{v_2})\end{bmatrix} = \frac{1}{6}\begin{bmatrix}1 & -\frac{1}{2}\\ -\frac{1}{2} & 1\end{bmatrix}\sigma^2_{v_g} \tag{22}$$

Solving equations (17) and (19) for $\sigma^2_{v_g}$ and substituting into equation (22) results in

$$V_\Delta = \begin{bmatrix}1 & -\frac{1}{2}\\ -\frac{1}{2} & 1\end{bmatrix}\left[\left(\frac{\sigma_{v_1}}{\sigma_{v_0}}\right)^2 - 1\right]\sigma^2_{v_0} \tag{23}$$

where

$$\left(\frac{\sigma_{v_1}}{\sigma_{v_0}}\right)^2 = \left[\frac{1}{2} + \left(\frac{\sigma_{v_a}}{\sigma_{v_g}}\right)^2\right] \Big/ \left[\frac{1}{3} + \left(\frac{\sigma_{v_a}}{\sigma_{v_g}}\right)^2\right] \tag{24}$$

In the limiting case where the ground error component is much larger than the airborne error component, $(\sigma_{v_1}/\sigma_{v_0})^2 \to \frac{3}{2}$. Alternatively, when the airborne error component is much larger than the ground component, $(\sigma_{v_1}/\sigma_{v_0})^2 \to 1$. In general, then, it is true that

$$1 < \frac{\sigma_{v_1}}{\sigma_{v_0}} < \sqrt{\frac{3}{2}} \tag{25}$$

In practical LAAS applications $\sigma_{v_1}/\sigma_{v_0} \approx 1.2$, although the actual value, along with the value of σ_{v_0}, is easily computed directly at the aircraft based on the assumed ground and airborne noise statistics and the satellite geometry at hand. Figure 2 shows the 10^{-7} normal error probability ellipse associated with equation (23) for the example case of $\sigma_{v_1}/\sigma_{v_0} = 1.2$.

Prior Probabilities

To compute integrity risk (or VPL_s) at the aircraft under the consideration of reference station failure hypotheses, it is first necessary to define the prior probabilities of failure. In this regard, it is implicitly assumed that the probability of multiple simultaneous reference station failures is much smaller than the integrity risk specification of 10^{-9}. Therefore, it is necessary that $P(H_1) = P(H_2) = P(H_3) = 10^{-5}$ or lower for independent reference stations, and for the no-failure hypothesis that one have approximately $P(H_0) = 1 - 3 \times 10^{-5}$.

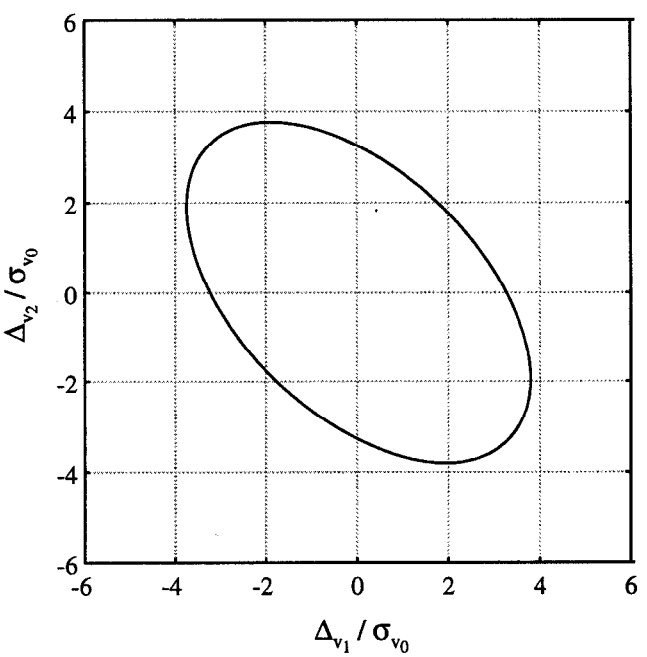

Fig. 2—Normal Error Probability Ellipse

Multiple Hypothesis Integrity Monitor with $\Delta_v = 0$

In this section, it is assumed that the nominal no-fault estimate of vertical position ($\Delta_v = 0$) is used for navigation. In this case, given the prior probabilities defined above and $\sigma_{v_1}/\sigma_{v_0}$, it is possible to express integrity risk as defined in equation (9) as a function of only $\Delta_{v_1}/\sigma_{v_0}$, $\Delta_{v_2}/\sigma_{v_0}$, and VAL/σ_{v_0}. As an example, Figure 3 shows a set of contours of constant 10^{-9} integrity risk for $\sigma_{v_1}/\sigma_{v_0} = 1.2$ and $VAL/\sigma_{v_0} = 7$, 8, and 9. It is noteworthy that these contours may also be interpreted as detection boundaries on the discriminators, since a set of discriminator values within a particular VAL contour has an associated integrity risk of less than 10^{-9} for that VAL, whereas a set of discriminator values outside the contour has an associated integrity risk of more than 10^{-9}. In this interpretation, it is easy to see that as the VAL is reduced, the detection boundary becomes more compact. Clearly, this will also result in a higher fault-free alarm rate. In this context, the predictive vertical protection limit (VPL_p) is defined as the minimum supportable VAL that can achieve a prescribed fault-free alarm rate requirement for given values of σ_{v_0} and σ_{v_1}. For the $\sigma_{v_1}/\sigma_{v_0} = 1.2$ example, Figure 4 shows the detection boundary corresponding to $VPL_p/\sigma_{v_0} = 8.127$ for a fault-free alarm probability of 10^{-7}. (The alarm rate was computed by direct integration of the Gaussian probability density function defined by equation (23)). The associated 10^{-7} normal error ellipse (dashed) is also shown in the figure for comparison. In this respect, it may alternatively be said that at any arbitrary measurement epoch, the probability that $VPL_s/\sigma_{v_0} > VPL_p/\sigma_{v_0}$ under fault-free conditions is 10^{-7}. The variation in VPL_p/σ_{v_0} across the entire realizable range of $\sigma_{v_1}/\sigma_{v_0}$ is shown in Figure 5. Clearly, VPL_p/σ_{v_0} is also a function of the prior probability of failure (which has been assumed to be 10^{-5}). For example, in the $\sigma_{v_1}/\sigma_{v_0} = 1.2$ case, the value of VPL_p/σ_{v_0} ranges between 7.37 for a prior probability of 10^{-6} and 8.78 for a prior probability of 10^{-4} (neglecting simultaneous multiple receiver failures).

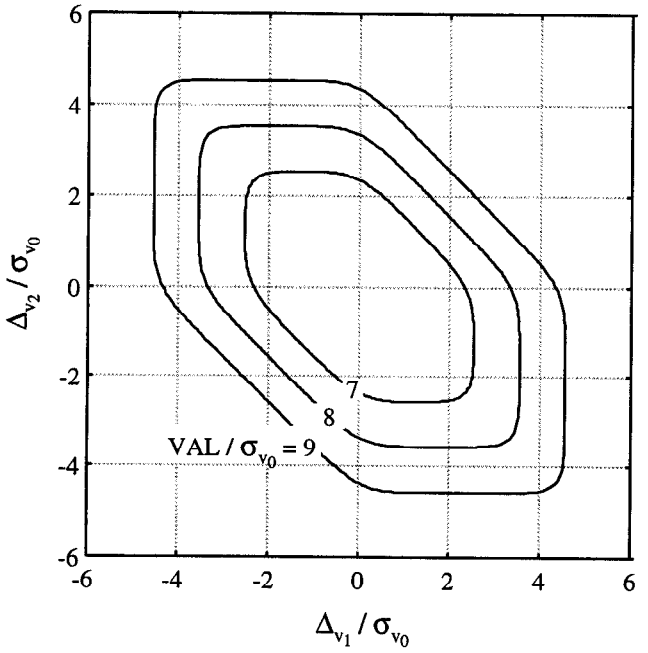

Fig. 3—Constant Integrity Risk Contours ($\Delta_v = 0$)

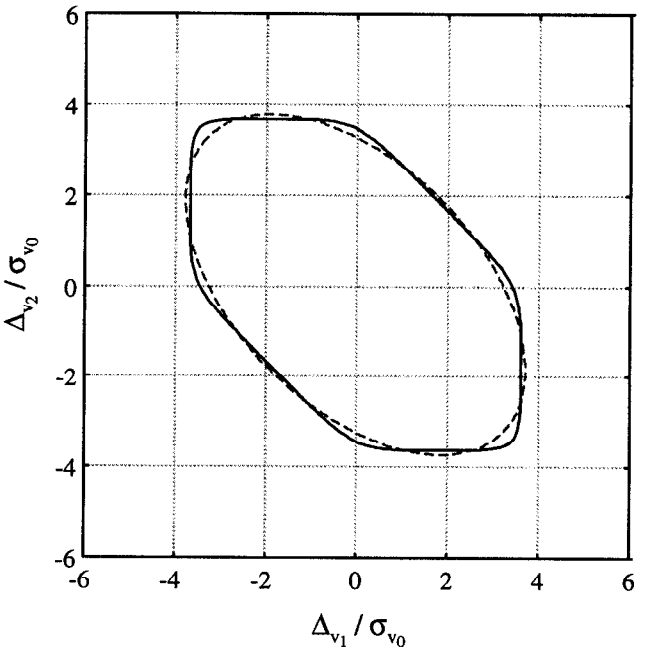

Fig. 4—Detection Boundary ($\Delta_v = 0$)

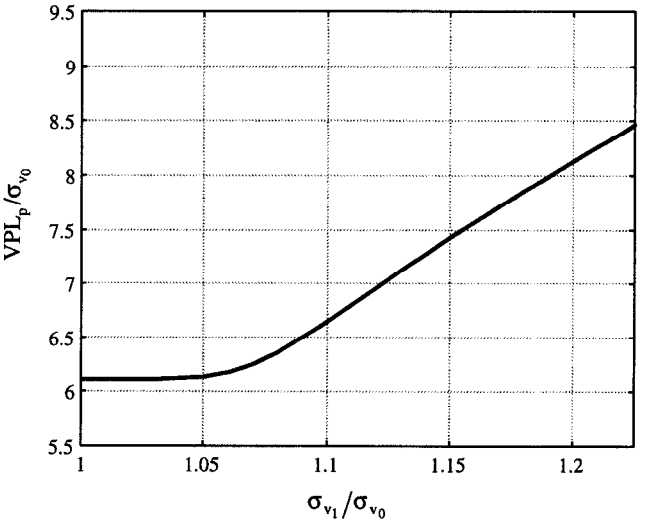

Fig. 5—Multiple Hypothesis ($\Delta_v = 0$) VPL_p Performance

Operationally, at the beginning of an approach, σ_{v_0} and σ_{v_1} are computed based on the known satellite geometry and fault-free ground and airborne ranging error statistics. These are used to find $\sigma_{v_1}/\sigma_{v_0}$ and then VPL_p/σ_{v_0} by table lookup of the data plotted in Figure 5. Using the known σ_{v_0}, VPL_p is computed. If the resulting VPL_p is larger than the VAL prescribed for the category of approach, the LAAS navigation function is declared unavailable because of deficient ranging geometry. If VPL_p is smaller than VAL, the LAAS navigation function is declared available, and the approach may be initiated.

At each measurement epoch during the approach, integrity risk is computed using equations (8), (9), (16), and (18); the values of σ_{v_0} and σ_{v_1}; and the prior probabilities assumed above. If the resulting integrity risk exceeds the 10^{-9} specification, a hazardous navigation condition is detected. As noted earlier, an alternative, but equivalent, means of detection may be implemented (at the expense of computation time) by directly comparing VPL_s with the VAL specification.

Multiple Hypothesis Integrity Monitor with Fault-Tolerant Estimation

In general, a smaller VPL_p (and therefore, higher navigation system availability) is achievable when the fault-tolerant estimate of vertical position is used in place of the nominal no-fault estimate. In this case, the deviation (Δ_v) in vertical position from the nominal no-fault estimate is given by the solution of equation (10). Using equation (10) in place of $\Delta_v = 0$ assumed above, it is again possible to express integrity risk as defined in equation (9) as a function only of $\Delta_{v_1}/\sigma_{v_0}$, $\Delta_{v_2}/\sigma_{v_0}$, and VAL/σ_{v_0} (given the prior probabilities and $\sigma_{v_1}/\sigma_{v_0}$). For the $\sigma_{v_1}/\sigma_{v_0} = 1.2$ example, Figure 6 shows the detection boundary corresponding to $VPL_p/\sigma_{v_0} = 7.839$ for a fault-free alarm probability of 10^{-7}. (Again, the associated 10^{-7} normal error ellipse is also shown in the figure for comparison.) It is clear that the minimum supportable VAL consistent with a prescribed fault-free alarm rate requirement (i.e., VPL_p) is lower using the fault-tolerant estimator than the nominal case in which $\Delta_v = 0$ (where $VPL_p/\sigma_{v_0} = 8.127$). The variation in VPL_p/σ_{v_0} across the entire realizable range of $\sigma_{v_1}/\sigma_{v_0}$ is shown in Figure 7 for both the fault-tolerant estimator (labeled 'FTE') and the nominal case of $\Delta_v = 0$.

Note that, in an operational sense, the use of fault-tolerant estimation is potentially beneficial only when the VPL_p in the nominal estimator case

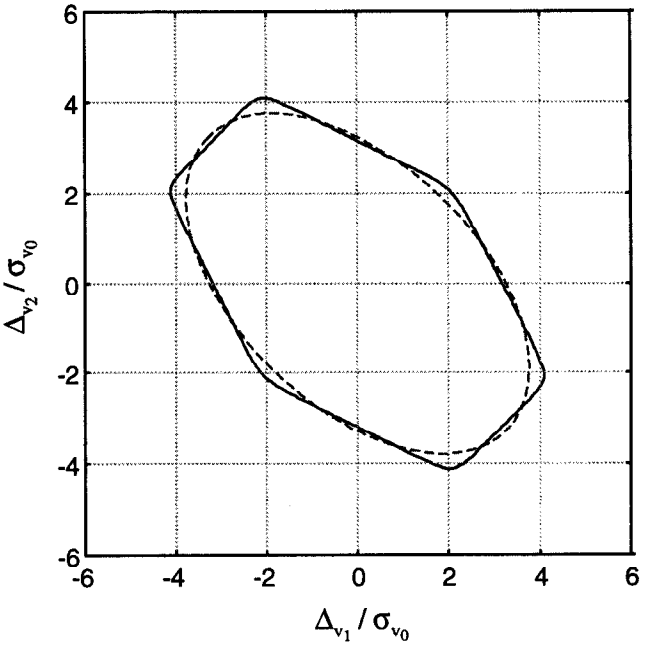

Fig. 6—Detection Boundary (FTE)

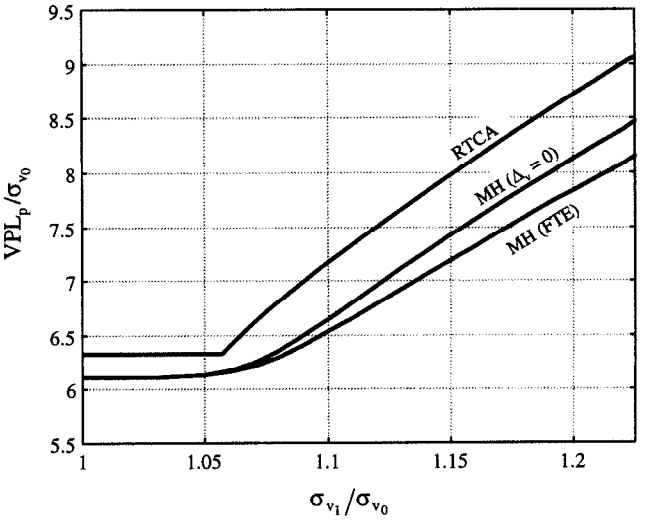

Fig. 7—Predictive VPL Performance Comparison

($\Delta_v = 0$) is larger than the specified VAL. In this event, the VPL_p may be lowered to the degree identified in Figure 8 if fault-tolerant estimation is used for navigation during the approach. All other operational integrity monitoring considerations are identical with the nominal ($\Delta_v = 0$) case.

Performance Comparison with Proposed RTCA Algorithm

The RTCA-proposed approach to the airborne element of LAAS integrity monitoring is described in some detail in [12] (with minor differences from this paper in failure hypothesis definition and general notation). Interestingly, however, the

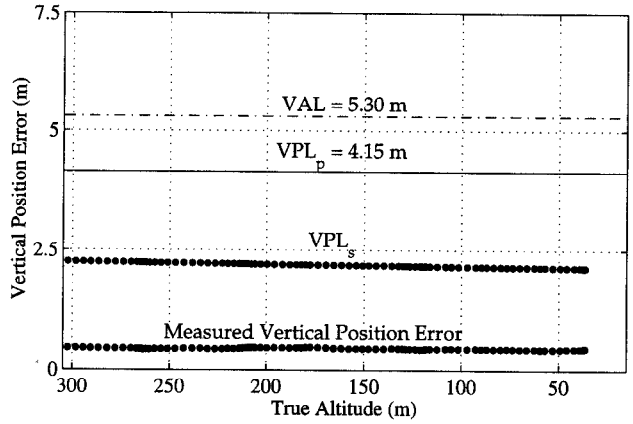

Fig. 8—Vertical Position Error and VPL for Typical Approach

RTCA approach can also be interpreted as a simplified version of the multiple hypothesis approach using the nominal no-fault estimate $\Delta_v = 0$. Specifically, in the RTCA approach, the integrity risk associated with the no-fault hypothesis and each single-reference-receiver-failure hypothesis is considered separately. The RTCA approach can be derived from the same basic principles as the full multiple hypothesis approach as follows.

First, the vertical integrity risk requirement I_v^{req} is allocated to all individual failure hypotheses. In this derivation, the example of three LAAS reference receivers is again considered. Although a specific allocation had not been prescribed by RTCA as of this writing, it is not unreasonable to choose the maximum allowable integrity risk for any of the four hypotheses (three failure hypotheses and the no-failure hypothesis) to be simply $I_v^{req}/4$. In comparison, the actual integrity risk associated with any given failure hypothesis i is simply the i-th term on the righthand side of equation (9) taking $\Delta_v = 0$. Alternatively, in terms of VPL_s, one can write

$$\tfrac{1}{4} I_v^{req} = \tfrac{1}{2} P(H_i)\left\{\operatorname{erfc}\left[\frac{VPL_{s_i} - \Delta_{v_i}}{\sqrt{2}\sigma_{v_i}}\right] + \operatorname{erfc}\left[\frac{VPL_{s_i} + \Delta_{v_i}}{\sqrt{2}\sigma_{v_i}}\right]\right\} \quad (26)$$

To solve for VPL_s explicitly, the following bound is applied:

$$\tfrac{1}{4} I_v^{req} \le P(H_i) \operatorname{erfc}\left[\frac{VPL_{s_i} - |\Delta_{v_i}|}{\sqrt{2}\sigma_{v_i}}\right] \quad (27)$$

In this case, if one conservatively assumes equality in equation (27),

$$VPL_{s_i} = |\Delta_{v_i}| + k_{md_i}\sigma_{v_i} \quad (28)$$

where the factor k_{md_i} corresponds to an integrated Gaussian tail probability of $\frac{1}{4} I_v^{req}/P(H_i)$. For example, if $I_v^{req} = 10^{-9}$ and $P(H_0) \approx 1$, then $k_{md_0} = 6.327$. Similarly, if all failure hypotheses have prior probability $P(H_1) = 10^{-5}$, then $k_{md_1} = k_{md_2} = k_{md_3} = 4.215$.

At a given measurement epoch, the statistical VPL is defined as

$$VPL_s = \max_i VPL_{s_i} \quad (29)$$

To compute the predictive VPL for an aircraft approach, one must consider the fault-free variation in Δ_{v_i}. In this regard, note that the fault-free distribution of the discriminators is given in equation (23). The variance on each individual discriminator is given by

$$\sigma^2_{\Delta_{v1}} = \left[\left(\frac{\sigma_{v_1}}{\sigma_{v_0}}\right)^2 - 1\right]\sigma^2_{v_0} \quad (30)$$

Although it is clear that the three discriminators are correlated random variables, for simplicity the RTCA approach has been to assume their independence in the computation of fault-free alarm rates. In this conservative approximation, the likelihood that any of $|\Delta_{v_i}|$ (i = 1, 2, 3) exceed $k_{ffa}\sigma_{\Delta_{v1}}$ is equivalent to an allocated fault-free alarm probability of 10^{-7} when $k_{ffa} = 5.523$. Using this result together with equations (28) and (30), one can write

$$VPL_{p_1} = k_{ffa}\left[\left(\frac{\sigma_{v_1}}{\sigma_{v_0}}\right)^2 - 1\right]^{\frac{1}{2}}\sigma_{v_0} + k_{md_1}\sigma_{v_1} \quad (31)$$

With regard to the no-failure hypothesis, it is noteworthy that no explicit definition of VPL_{p_0} is given in the RTCA algorithm as currently proposed. However, since $\Delta_{v_0} \equiv 0$, VPL_{s_0} is not actually a random variable and may therefore equivalently be used as the definition for VPL_{p_0}. Since it is possible that VPL_{p_0} is larger than VPL_{p_1} when $\sigma_{v_1}/\sigma_{v_0}$ approaches unity, the resulting predictive VPL is defined (for completeness in the RTCA approach) as

$$VPL_p = \max[VPL_{p_0}, VPL_{p_1}] \quad (32)$$

The variation in VPL_p/σ_{v_0} for the RTCA approach is shown in Figure 7 across the entire realizable range of $\sigma_{v_1}/\sigma_{v_0}$. It is clear that the VPL_p for the RTCA approach is always higher than the multiple hypothesis result (also shown in the figure). Fundamentally, this is the case because the RTCA approach conservatively requires that the

component of integrity risk due to each individual failure hypothesis be below an allocated value. In contrast, the multiple hypothesis approach requires only that the total risk summed over all hypotheses be lower than the overall integrity risk requirement. Because of this, a higher operational availability of LAAS can be expected using the multiple hypothesis approach. This improvement in availability is quantified in the following section.

Availability Considerations

Because VPL_p can be assessed in advance for a given ranging error model and given satellite and user locations, it is possible to estimate system availability, with respect to a prespecified VAL, at any given time, as well as the duration of periods in which $VPL_p > VAL$. To do this, a GPS satellite simulation has been conducted for LAAS users located at three major airports in the continental United States with Category III capability: Chicago-O'Hare (ORD), New York-Kennedy (JFK), and Los Angeles (LAX). In these simulations, all 276 possible cases in which 22 of the 24 satellites of the idealized GPS constellation [13] are healthy are sampled. For each outage case, the satellites are moved through their respective orbits over 1 (repeatable) day at 5 min intervals. At each update, the satellite geometry (using a 5 deg elevation mask angle), vertical position accuracy, and resulting VPL_p are calculated for all three airports. Position accuracy calculations are based on the Boeing/Collins LAAS ranging error model [14]. Records of each VPL result are maintained, and periods in which the VPL exceeds a vertical alert limit of 5.3 m (the proposed limit for Category III operations) [10] are tracked separately for each airport to measure LAAS outage durations.

Table 1 shows the predicted Category III availability results for each airport. Service availability, or the long-term likelihood of availability at any given time, is above 99 percent for all cases. It varies as much as 0.24 percent among locations for a given algorithm and shows an improvement of about 0.14 percent when the multiple hypothesis approach is used in place of the RTCA algorithm. This improvement is quite small given the larger uncertainties of number of healthy satellites and ranging accuracy. The further improvement obtained from the fault-tolerant algorithm is even smaller. The maximum outage duration is improved by no more than 5 min over that of the RTCA algorithm in any of the cases studied.

While the improvement in availability will be more pronounced at lower values of VAL, it is also true that the overall service availability at these VALs is very likely to be unacceptably low using either algorithm. It is evident, therefore, that while the reduction in VPL provided by the multiple hypothesis approach (indicated in Figure 7) is substantial, the resulting LAAS service availability improvement is quite small.

EXPERIMENTS

To demonstrate the operational performance of the multiple hypothesis approach, the algorithm was incorporated into a prototype LAAS architecture installed at Moffett Federal Airfield in California for testing in conjunction with flight trials performed in a NASA Ames King Air. (A detailed description of the Moffett/King Air LAAS prototype design, implementation, and flight test results may be found in [15].) Three reference receivers were implemented in the ground segment with a minimum antenna separation of 100 m to provide statistically independent multipath errors at each antenna. Raw code and carrier phase measurements from each receiver were transmitted to the aircraft through a VHF data link for processing. At the aircraft, code measurements were smoothed using the carrier, and nominal least-squares positioning ($\Delta_v = 0$) was performed using the measurements from all three reference receivers. The multiple hypothesis integrity monitor was then implemented to ensure position estimate integrity with respect to reference receiver failures.

Figure 8 shows the vertical position error profile for an arbitrary approach. The 'truth' trajectory from which this error profile was obtained was generated using kinematic (carrier phase) GPS, with the cycle ambiguities resolved during a preflight static survey. The bias-like position error structure exhibited is due to the effect of low-

Table 1—Conditional Availability of Category III Operations Given Two Satellite Failures

	RTCA Algorithm		MH ($\Delta_v = 0$)		MH (FTE)	
Airport	Avail.	Max. Outage (min)	Avail.	Max. Outage (min)	Avail.	Max. Outage (min)
LAX	0.9936	65	0.9944	60	0.9945	60
ORD	0.9916	55	0.9939	50	0.9943	50
JFK	0.9940	45	0.9952	45	0.9957	45
Overall	0.9931	65	0.9945	60	0.9948	60

frequency multipath at the reference receivers. Also shown in the figure are the VPL_s containment boundary generated during the approach and the VPL_p obtained at the beginning of the approach, both generated using the multiple hypothesis algorithm. In these results, VPL_s provides a 10^{-9} statistical (integrity) boundary for vertical position error at each epoch of the approach, while VPL_p represents the 10^{-7} statistical (continuity) boundary for VPL_s. Similar results were obtained in 30 additional approaches performed during these flight tests.

CONCLUSION

A new multiple hypothesis methodology for failure mitigation in satellite-based navigation systems has been developed, implemented, and tested. Integrity risk is directly evaluated under the unified consideration of all single-element failure hypotheses and the no-failure hypothesis. The methodology provides the basis for both fault-tolerant estimation and integrity monitoring.

The algorithm has been investigated for application as an airborne element of the LAAS integrity monitoring architecture. While the multiple hypothesis algorithm is shown to provide the tightest realizable protection limits for LAAS, it is also concluded that the currently proposed RTCA algorithm will provide almost equivalent LAAS availability. The operational performance of the new algorithm has been demonstrated using experimental data collected during a series of LAAS prototype flight tests in a NASA Beechcraft King Air at Moffett Federal Airfield.

ACKNOWLEDGMENTS

The assistance of Guttorm Opshaug and Professor Per Enge at Stanford University is greatly appreciated. The authors gratefully acknowledge the FAA 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.

APPENDIX

Beginning with equation (7), making the substitution

$$s \equiv \tilde{x}_v - \hat{x}_v \tag{A-1}$$

and using the definitions of equation (8) yields

$$I_v|\mathbf{z} = \sum_{i=0}^{n} P(H_i) \left\{ \int_{VAL}^{\infty} N_s[\Delta_{v_i} - \Delta_v, \sigma_{v_i}^2]ds + \int_{-\infty}^{-VAL} N_s[\Delta_{v_i} - \Delta_v, \sigma_{v_i}^2]ds \right\} \tag{A-2}$$

where the probability density function N_s in equation (A-2) is defined by analogy with the definition of equation (6). The definite Gaussian integrals in equation (A-2) evaluate to the well-known analytical (complementary error function) result given in equation (9).

Now making the substitution

$$u \equiv s + \Delta_v \tag{A-3}$$

into equation (A-2) results in

$$I_v|\mathbf{z} = \lim_{\mu->\infty} \sum_{i=0}^{n} P(H_i) \left\{ \int_{VAL+\Delta_v}^{\mu+\Delta_v} N_u[\Delta_{v_i}, \sigma_{v_i}^2]du + \int_{-\mu+\Delta_v}^{-VAL+\Delta_v} N_u[\Delta_{v_i}, \sigma_{v_i}^2]du \right\} \tag{A-4}$$

To obtain the minimum integrity risk value of Δ_v, equation (A-4) must be differentiated with respect to Δ_v. The form of equation (A-4) is significant in that only the limits of integration (i.e, not the integrands) are a function of the parameter Δ_v. In this case, it is convenient to use Leibnitz's theorem [16] for differentiation of an integral:

$$\frac{d}{d\Delta_v} I_v|\mathbf{z} = \lim_{\mu->\infty} \sum_{i=0}^{n} P(H_i) \left\{ N_u[\Delta_{v_i}, \sigma_{v_i}^2]_{u=\mu+\Delta_v} - N_u[\Delta_{v_i}, \sigma_{v_i}^2]_{u=VAL+\Delta_v} + N_u[\Delta_{v_i}, \sigma_{v_i}^2]_{u=-VAL+\Delta_v} - N_u[\Delta_{v_i}, \sigma_{v_i}^2]_{u=-\mu-\Delta_v} \right\} \tag{A-5}$$

Setting this derivative equal to zero and evaluating the Gaussian density functions at the limits of integration as indicated on the righthand side of equation (A-5) results in

$$\sum_{i=0}^{n} P(H_i) \left\{ 0 - \frac{1}{\sqrt{2\pi\sigma_{v_i}^2}} \exp\left[-\left(\frac{VAL + \Delta_v - \Delta_{v_i}}{\sqrt{2}\sigma_{v_i}} \right)^2 \right] + \frac{1}{\sqrt{2\pi\sigma_{v_i}^2}} \exp\left[-\left(\frac{-VAL + \Delta_v - \Delta_{v_i}}{\sqrt{2}\sigma_{v_i}} \right)^2 \right] - 0 \right\} = 0 \tag{A-6}$$

which is equivalent to equation (10).

REFERENCES

1. Parkinson, B. and Axelrad, P., *Autonomous Integrity Monitoring Using the Pseudorange Residual*, NAVIGATION, Journal of The Institute of Navigation, Vol. 35, No. 2, Summer 1988.
2. Sturza, M., *Navigation System Integrity Monitoring Using Redundant Measurements*, NAVIGATION, Journal of The Institute of Navigation, Vol. 35, No. 4, Winter 1988–89.

3. Brown, R. G., *A Baseline RAIM Scheme and a Note on the Equivalence of Three RAIM Methods*, NAVIGATION, Journal of The Institute of Navigation, Vol. 39, No. 3, Fall 1992.
4. Brenner, M., *An Integrity Monitoring Scheme for Precision Approach Applications*, Proceedings of ION GPS-96, Kansas City, MO, September 1996.
5. Shively, C. and Lee, Y., *A Position Domain Ground-Based Integrity Method for LAAS*, Proceedings of ION GPS-96, Kansas City, MO, September 1996.
6. Braff, R., *Alarm Limits for Local DGPS Integrity Monitoring*, Proceedings of The Institute of Navigation National Technical Meeting, Anaheim, CA, January 1995.
7. Markin, K. and Shively, C., *A Position-Domain Method for Ensuring Integrity of Local Area Differential GPS*, Proceedings of The Institute of Navigation 51st Annual Meeting, Colorado Springs, CO, June 1995.
8. *Requirements Document for Local Area Augmentation System*, Satellite Navigation Program Office, Federal Aviation Administration, October 1997.
9. Swider, R., Miller, D., Shakarian, A., Flannigan, S., and Fernow, J. P., *Development of Local Area Augmentation System Requirements*, Proceedings of ION GPS-96, Kansas City, MO, September 1996.
10. DeCleene, B., et al., *Rationale for (LAAS) Requirements*, RTCA SC-159, WG 4A, DCA 3-7, Washington, DC, July 15, 1997.
11. *GNSS Based Precision Approach: Local Area Augmentation System (LAAS) Signal-in-Space Interface Control Document*, RTCA SC-159, WG 4A, Draft II, September 7, 1997.
12. Liu, F., Murphy, T., and Skidmore, T., *Verification of the LAAS Integrity Monitoring Algorithm*, Proceedings of ION GPS-97, Kansas City, MO, September 1997.
13. *Minimum Operational Performance Standards for Global Positioning System / Wide Area Augmentation System Airborne Equipment*, RTCA DO-229, Washington, DC, January 16, 1996.
14. Murphy, T., *Proposed Corrected Pseudorange Error Requirements*, RTCA SC-159, WG 4A, Everett, WA, August 16, 1997.
15. Pervan, B., Lawrence, D., Gromov, G., Opshaug, G., Christie, J., Ko, P.-Y., Mitelman, A., Pullen, S., Enge, P., and Parkinson, B., *Flight Test Evaluation of an Alternative Local Area Augmentation System Architecture*, NAVIGATION, Journal of The Institute of Navigation, Vol. 45, No. 1, Spring 1998.
16. Abramowitz, M. and Stegun, I. A. (eds.), *Handbook of Mathematical Functions*, New York: Dover Publications, 1972.