

Comparing Performance Bounds for Chi-square Monitors with Parameter Uncertainty

Jason H. Rife
Tufts University

>> Accepted Article <<

CITATION:

J. H. Rife, "Comparing performance bounds for chi-square monitors with parameter uncertainty," in *IEEE Transactions on Aerospace and Electronic Systems*, vol. 51, no. 3, pp. 2379-2389, July 2015.

doi: 10.1109/TAES.2015.140638

CORRESPONDING AUTHOR:

Jason Rife
jason.rife@tufts.edu

COPYRIGHT:

© 2015 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

FINANCIAL SUPPORT:

Federal Aviation Administration GBAS Program (Grant FAA-10-G-006)

Comparing Performance Bounds for Chi-Square Monitors with Parameter Uncertainty

Jason H. Rife, *Member, IEEE*

Abstract— This paper compares methods for evaluating the performance of chi-square monitors while conservatively accounting for parameter uncertainty. Chi-square monitors, like the Signal Deformation Monitors (SDM) used in GPS augmentation, detect failures that threaten safety-critical navigation. A chi-square monitor creates a quadratic test statistic from a random vector (nominally zero-mean, Gaussian-distributed). Gaussian model parameters, which may be poorly characterized for a real system, strongly influence chi-square monitor performance. Through a combination of theory and simulation, it is established that tight yet conservative modeling of parameter uncertainty is possible with a generalized chi-square bound for false-alarm risk and with an ellipsoid bound for missed-detection risk.

Index Terms—Chi-Square, Integrity, SBAS, GBAS, RAIM

I. INTRODUCTION

SAFE navigation systems rely on continuous monitoring to detect anomalous sensor measurements. This paper introduces new methods to ensure tight, conservative performance bounds for chi-square integrity monitors.

Integrity monitoring algorithms have played an important role in safety-critical applications of satellite navigation systems like GPS, GLOSNASS, BeiDou, and Galileo [1]. For example, integrity monitors are found in Satellite-Based Augmentation Systems (SBAS), Ground-Based Augmentation Systems (GBAS), and Receiver Autonomous Integrity Monitoring (RAIM) [2]. These systems all perform continuous testing of satellite navigation signals, to ensure signals meet safety requirements for civil aviation.

Certifying that an integrity monitor meets safety-of-life requirements necessitates that monitor performance be evaluated in a conservative manner, accounting for reasonable variations in the nominal operating environment. The term *overbounding* is often used to label such analyses, in which a conservative probability estimate is obtained from the assumption that the actual input noise distribution is not known precisely but belongs to a well-defined set of possible distributions. Seminal research on navigation overbounding

was conducted by DeCleene [3]. This work was later extended to cover more general classes of error distribution [4]-[7].

To date, most overbounding research has focused on Gaussian bounding. Gaussian bounding applies to scalar random variables such as vertical-positioning error [8],[9] or excessive-acceleration, a statistic used in monitoring for satellite-clock failure [10],[11]. Recent work has extended bounding concepts to vector signals, such as in bounding 3D positioning error [12],[13] or monitoring integrity for a vector measurement [14],[15].

Chi-square integrity monitors are a particularly important class of vector monitor, one which quadratically combines a vector of input signals to produce a positive, scalar monitor statistic [16],[17]. This class of monitor is named for its noise distribution, which is nominally chi-square. When an anomaly is present, the noisy monitor statistic tends to increase, and so large anomalies become observable by continuously comparing the monitor statistic to a threshold. Examples of chi-square monitors in satellite navigation include Signal Deformation Monitoring (SDM), Ionosphere Gradient Monitoring (IGM), and least-squares RAIM [18]-[23].

In particular, the methods presented in this paper have been developed to support verification of the GBAS SDM, which monitors for anomalous GPS waveforms that might otherwise threaten aircraft during approach and landing. Signal deformation events distort correlator peaks for GPS code tracking and can result in position errors of tens or hundreds of meters [19]. To mitigate associated integrity risk, GBAS SDM constructs a monitor statistic from eight correlator taps distributed around the correlation peak [20]. This statistic is nominally small, but grows quickly in the presence of signal deformation. The distribution of the monitor statistic is approximately chi-square; however, associated covariance parameters are difficult to model precisely, on account of nonstationary effects (caused by satellite elevation changes) and sensitivity to interference (generated, for example, by “personal privacy devices” [24],[25]).

For the most part, existing overbounds cannot be applied to assess chi-square monitor performance when the covariance model is itself uncertain. Hence, new overbounding methods are necessary. Several relevant chi-square monitor models have recently been proposed by the author, including three of particular interest: a generalized chi-square bound [26], a determinant bound [27], and an ellipsoid bound [28]. This paper compares these bounding methods and establishes which are best suited to different performance analyses.

Manuscript received August XX, 2014. The authors gratefully acknowledge the Federal Aviation Administration GBAS Program (Grant FAA-10-G-006) for supporting this research. The opinions discussed here are those of the authors and do not necessarily represent those of the FAA or other affiliated agencies.

J. Rife is with Tufts University, Medford, MA 02155 USA (e-mail: Jason.Rife@tufts.edu).

The paper’s contributions are four-fold. First, the paper frames conservative analyses of false-alarm and missed-detection risk as optimization problems. Second, the paper shows that an analytical solution is possible for the false-alarm optimization problem. Third, the paper provides a quantitative example to motivate why an exact analytical solution is difficult to obtain for the missed-detection optimization problem. Fourth, the paper applies simulation to compare two analytical solutions that conservatively approximate the missed-detection optimization; simulation show that an ellipsoid bound [27] can be much tighter than a determinant-based bound [28], by an order of magnitude in simulation.

II. BACKGROUND

This section defines optimization problems to compute two important performance metrics for chi-square monitors: the worst-case false-alarm and missed-detection probabilities.

A. Chi-Square Integrity Monitors

Most integrity monitors function by comparing a monitor statistic m to a threshold T . If the statistic ever exceeds the threshold, the monitor issues an alert. The alert indicates that an anomalous signal has been detected.

$$m > T \rightarrow \text{alert} \quad (1)$$

The integrity monitor is said to be *chi-square* when it generates m quadratically from the measurement vector \mathbf{x} , nominally a zero-mean, Gaussian-distributed random vector.

$$m = \mathbf{x}^T \hat{\mathbf{P}}^{-1} \mathbf{x} \quad (2)$$

This quadratic operation includes normalization by the matrix $\hat{\mathbf{P}}$, which is a model of the actual covariance matrix \mathbf{P} for the random vector signal \mathbf{x} .

$$\mathbf{P} = E[\mathbf{x}\mathbf{x}^T] \quad (3)$$

Normalizing by estimated covariance enhances monitor performance and results in a monitor statistic m with noise that is approximately chi-square distributed. Noise is exactly chi-square distributed if the model is exact ($\hat{\mathbf{P}} = \mathbf{P}$).

In practical applications, the estimated covariance is never known exactly, which is the natural consequence of estimating a statistical quantity from finite data. Estimation is particularly difficult for satellite navigation applications, as statistics are non-stationary, meaning the true covariance matrix \mathbf{P} changes continually. Changes in \mathbf{P} are amplified when environmental conditions are unfavorable, in SDM for example, when elevated radio-frequency interference is present [32].

To model the difference between the estimate $\hat{\mathbf{P}}$ and the true covariance \mathbf{P} , it is useful to define an intermediate vector \mathbf{y} , which is a vector of unit-variance, independent Gaussian random variables, related to \mathbf{x} as follows.

$$\mathbf{x} = \mathbf{P}^{1/2} \mathbf{y} \quad (4)$$

Here $\mathbf{P}^{1/2}$ is the matrix square root, a symmetric matrix that results, for example, from an *LDL^T* Decomposition [33]. Rewriting (2) in terms of (4), the monitor statistic is

$$m = \mathbf{y}^T \mathbf{Q}^{-1} \mathbf{y}, \quad (5)$$

where \mathbf{Q} is a residual covariance matrix,

$$\mathbf{Q} = \mathbf{P}^{1/2} (\hat{\mathbf{P}}^{-1}) \mathbf{P}^{1/2}. \quad (6)$$

The covariance \mathbf{Q} , like \mathbf{P} , is unknown; however, if the domain of \mathbf{Q} can be bounded, then it is possible to assess the validity of chi-square error modeling. The matrix \mathbf{Q} becomes the identity \mathbf{I} only if the model $\hat{\mathbf{P}}$ is a perfect estimate of the true \mathbf{P} . Otherwise m is not the sum-square of independent, unit-variance Gaussian variables and, hence, not chi-square distributed. Instead, the noise distribution for m is said to be generalized chi-square [29]. The cumulative distribution function P_{gx2} for a generalized chi-square distribution, inside a threshold T with covariance \mathbf{Q} and mean $\boldsymbol{\mu}$, is

$$P_{gx2}(T; \mathbf{Q}, \boldsymbol{\mu}) = \int_{\mathbf{y}^T \mathbf{y} \leq T} p_N(\mathbf{y}; \mathbf{Q}, \boldsymbol{\mu}) dV. \quad (7)$$

Here p_N is a multivariate normal distribution with the form

$$p_N(\mathbf{y}; \mathbf{Q}, \boldsymbol{\mu}) = |2\pi\mathbf{Q}|^{-\frac{1}{2}} e^{-\frac{1}{2}(\mathbf{y}-\boldsymbol{\mu})^T \mathbf{Q}^{-1}(\mathbf{y}-\boldsymbol{\mu})}. \quad (8)$$

In assessing monitor performance, two important metrics are the false-alarm and missed-detection risks. A false alarm occurs when an alert triggers though no anomaly is present. False alarms are also called *type I errors* (in statistics [30]) or *continuity risk* (in navigation [31]). When the true covariance matrix is uncertain, the false-alarm probability is

$$P_{fa} = 1 - P_{gx2}(T; \mathbf{Q}, \mathbf{0}). \quad (9)$$

This equation computes the probability that the random signal lies outside the threshold T , even though the signal is unbiased (i.e., has mean $\mathbf{0}$). By comparison, a missed detection occurs when an alert is not triggered though an anomaly is present. Missed detections are also called *type II errors* (in statistics [30]) or *integrity risk* (in navigation [31]). When the true covariance matrix is uncertain, missed-detection probability is

$$P_{md} = P_{gx2}(T; \mathbf{Q}, \boldsymbol{\mu}). \quad (10)$$

This equation computes the probability that random noise results in a monitor statistic inside the threshold, even though the anomaly introduces a nonzero mean $\boldsymbol{\mu}$.

B. Performance Bounding

Although the precise value of the covariance \mathbf{P} and mean $\boldsymbol{\mu}$ are uncertain, it may still be possible to constrain these parameters to well-characterized sets, $\mathbf{Q} \in \Theta$ and $\boldsymbol{\mu} \in \Psi$. For safety-of-life navigation, it is critical to identify the worst possible performance metrics (P_{md} and P_{fa}) over these sets.

First, consider the worst-case false-alarm probability P_{fa}^* . This probability is defined by an optimization that maximizes P_{fa} over all allowable residual covariance matrices \mathbf{Q} .

$$P_{fa}^* = \max_{\mathbf{Q} \in \Theta} (1 - P_{gx2}(T; \mathbf{Q}, \mathbf{0})) \quad (11)$$

Several approaches for defining the set Θ are possible [27],[32]. Consider covariance matrices $\mathbf{Q} \in \mathbb{R}^{N \times N}$ with eigenvalues λ_n for axes $n = \{1, \dots, N\}$. In this paper, Θ is defined to contain all \mathbf{Q} with eigenvalues λ_n between positive upper and lower limits for the n -th axis: $\lambda_{max,n}$ and $\lambda_{min,n}$, respectively.

$$\Theta = \left\{ \mathbf{Q} \mid \lambda_{min,n} < \lambda_n(\mathbf{Q}) < \lambda_{max,n} \right\} \quad (12)$$

For SDM, these limits correspond to RFI-free conditions ($\lambda_{min,n}$) and to worst-RFI conditions ($\lambda_{max,n}$). Geometrically, these limits can be visualized as constraints on the lengths of the principal axes of \mathbf{Q} . To be precise, each eigenvalue is the square of a principal axis of the ellipsoid defined by (5) for any particular value of m .

Second, consider the worst-case missed-detection probability P_{md}^* . This probability is defined by an optimization that maximizes P_{md} over all allowable values of \mathbf{Q} and $\boldsymbol{\mu}$.

$$P_{md}^* = \max_{\mathbf{Q} \in \Theta, \boldsymbol{\mu} \in \Psi} (P_{gx2}(T; \mathbf{Q}, \boldsymbol{\mu})) \quad (13)$$

Again, for this paper, the set of allowable \mathbf{Q} will be specified by (12). The set of allowable mean vectors Ψ is defined

$$\Psi = \left\{ \boldsymbol{\mu} \mid \|\boldsymbol{\mu}\| = b \right\}. \quad (14)$$

This set includes all mean vectors with the same bias magnitude b , where b is a positive constant. Geometrically, this set represents the idea that the orientation of the mean vector is arbitrary.

In subsequent analysis, it is convenient to introduce a change of variables that has no impact on the Θ or Ψ constraints. Specifically, the new variable is $\mathbf{z} = \mathbf{X}^T \mathbf{y}$, a transform that rotates the coordinate axes to align with the principal axes of \mathbf{Q} . The rotation matrix \mathbf{X} is a unitary eigenvector matrix that complements the diagonal eigenvalue matrix $\boldsymbol{\Lambda}$ in the eigen factorization

$$\mathbf{Q} = \mathbf{X} \boldsymbol{\Lambda} \mathbf{X}^T. \quad (15)$$

After transformation, the new random vector \mathbf{z} has covariance $\boldsymbol{\Lambda}$ and mean $\boldsymbol{\mu}_z = \mathbf{X}^T \boldsymbol{\mu}$. Hence, the generalized chi-square distribution in (7) can be re-written substituting

$$P_{gx2}(T; \mathbf{Q}, \boldsymbol{\mu}) = P_{gx2}(T; \boldsymbol{\Lambda}, \boldsymbol{\mu}_z). \quad (16)$$

The covariance constraint (12) is directly expressed in terms of eigenvalues, and is thus unaffected by the change of variables. Similarly, the bias-magnitude constraint (14) also applies to $\boldsymbol{\mu}_z$, since for a unitary matrix \mathbf{X} , $\|\boldsymbol{\mu}_z\| = \|\mathbf{X}^T \boldsymbol{\mu}\| = \|\boldsymbol{\mu}\|$.

III. FALSE-ALARM RISK

This section derives an analytic result for the worst-case false-alarm risk P_{fa}^* , which is the solution to optimization problem (11). The analytic solution is obtained by showing that P_{fa} always increases as the eigenvalues of \mathbf{Q} increase. As such, the maximum P_{fa} occurs for the matrix \mathbf{Q} with all eigenvalues at the upper limits defined by (12).

$$P_{fa}^* = 1 - P_{gx2}(T; \boldsymbol{\Lambda}_{max}, \mathbf{0}). \quad (17)$$

Here the matrix $\boldsymbol{\Lambda}_{max}$ is the diagonal matrix with each diagonal element λ_n set equal to the upper limit (12).

To obtain this result, combine (11) with (7) and (16) to give:

$$P_{fa}^* = \max_{\mathbf{Q} \in \Theta} \left(1 - \int_{\mathbf{z}^T \boldsymbol{\Lambda} \mathbf{z} \geq T} p_N(\mathbf{z}; \boldsymbol{\Lambda}, \mathbf{0}) dV \right). \quad (18)$$

Now integrate the normal distribution p_N in a single direction, along the i^{th} coordinate axis. The result is

$$\int_{T \geq \mathbf{z}^T \boldsymbol{\Lambda} \mathbf{z}} p_N(\mathbf{z}; \boldsymbol{\Lambda}, \mathbf{0}) dV = \int_{\tilde{T} \geq 0} P_N(\sqrt{\tilde{T}}; \lambda_i, 0) \frac{e^{-\left(\frac{\sum_{k \neq i} z_k^2}{2\lambda_k}\right)}}{\prod_{k \neq i} (2\pi\lambda_k)^{\frac{1}{2}}} \prod_{k \neq i} dz_k. \quad (19)$$

Here $\tilde{T} = T - \sum_{k \neq i} z_k^2$, and P_N is the normal cumulative distribution between symmetric upper and lower limits:

$$P_N(\tau; \sigma^2, 0) = \int_{-\tau}^{\tau} \frac{1}{\sqrt{2\pi\sigma}} e^{-\frac{z^2}{2\sigma^2}} dz. \quad (20)$$

The slope of (19) with respect to any one eigenvalue λ_i is

$$\frac{\partial}{\partial \lambda_i} \int_{\mathbf{z}^T \leq T} p_N(\mathbf{z}; \boldsymbol{\Lambda}, \mathbf{0}) dV = \int_{\tilde{T} \geq 0} \left(-\frac{\sqrt{\tilde{T}} e^{-\frac{\tilde{T}}{2\lambda_i}}}{\lambda_i (2\pi\lambda_i)^{\frac{1}{2}}} \right) \frac{e^{-\left(\sum_{k \neq i} \frac{z_k^2}{2\lambda_k}\right)}}{\prod_{k \neq i} (2\pi\lambda_k)^{\frac{1}{2}}} \prod_{k \neq i} dz_k. \quad (21)$$

The above expression is nonpositive, since the leading negative sign multiplies a series of nonnegative terms (noting the λ_n , \tilde{T} , and the exponentials are nonnegative). Thus,

$$\frac{\partial}{\partial \lambda_i} \left(\int_{\mathbf{z}^T \leq T} p_N(\mathbf{z}; \boldsymbol{\Lambda}, \mathbf{0}) dV \right) \leq 0 \quad \forall i. \quad (22)$$

Returning to optimization problem (18), the above result can be used to show that

$$\frac{\partial}{\partial \lambda_i} P_{fa}(T; \boldsymbol{\Lambda}, \mathbf{0}) \geq 0 \quad \forall i. \quad (23)$$

The above relationship implies that P_{fa} always increases as eigenvalues increase, and so false-alarm risk is maximized when the eigenvalues on each principal axis reach their maximum values, defined by constraint (12). The resulting false-alarm risk is the expression for P_{fa}^* given by (17).

IV. MISSED-DETECTION RISK

This section considers analytical solutions for bounding missed-detection risk. First, the section introduces examples to explain why an exact, analytical solution to the worst-case missed-detection risk optimization problem is difficult to obtain. Second, the section describes how relaxing the original optimization problem, described by (13), makes it possible to obtain an analytical, approximate solution. The approximate solution \bar{P}_{md} is a conservative model of the exact solution P_{md}^* in the sense that

$$\bar{P}_{md} \geq P_{md}^*. \quad (24)$$

Some conservatism is critical for a safety-of-life system; however, excess conservatism may mean that the system cannot perform its required function. (In the extreme, a probability \bar{P}_{md} going to its maximum value of one is equivalent to the assumption that the monitor always fails to detect faults, a result which would render the monitor useless.)

To ensure good monitor performance without sacrificing safety, it is desirable to ensure the approximation is tight, meaning that the slack in inequality (24) is small. The tightness of the ellipsoid approximation proposed in this paper is assessed at the end of this section, where a simulation-based analysis compares its performance to that of the determinant bound, previously introduced in [27].

A. Obstacles to Obtaining an Exact Analytical Solution

This section provides a rationale describing why it is difficult to obtain an exact, analytical solution to problem (13), the problem of identifying worst-case missed-detection risk over the space of unknown parameters. The basic idea is that an analytical solution would be easy to obtain if the worst-case parameters fell on the boundaries of the allowed parameter space ($\mathbf{Q} \in \Theta$ and $\boldsymbol{\mu} \in \Psi$). This was the case for false-alarm risk, where the worst-case risk occurred at the edge of the domain (with the principal axes of \mathbf{Q} reaching their maximum possible values). By comparison, the worst-case parameters for the missed-detection risk optimization problem can occur in the middle of the allowed domain, and not only at the edges. Unfortunately, no analytical formula has been derived to identify the worst-case parameters when they lie interior to the domain. (Formulae obtained from setting derivatives to zero result in difficult-to-solve integral equations).

Remark 1: Worst Bias May Not Align with Major Axis. First consider the question of which bias direction maximizes missed-detection risk, noting that the bias magnitude is fixed by constraint (14). The bias direction may be envisioned as a series of rotations between the direction of the vector $\boldsymbol{\mu}$ and the directions of the largest principal axis of the covariance matrix $\boldsymbol{\Lambda}$. A zero angle represents one limiting case, where the mean vector and ellipsoid major axis align, as illustrated in Fig. 1(a). It is intuitive to imagine that the missed-detection risk will be largest for the zero-angle case, as this configuration would seem to project probability directly from the distribution center toward the threshold. To illustrate this case, the figure shows the threshold as a sphere (or really a circle in 2D) with radius \sqrt{T} . The Gaussian distribution is illustrated as the combination of a mean vector (arrow tip) and a covariance matrix (ellipse). The quantity P_{md} is the integral of the Gaussian distribution inside the threshold.

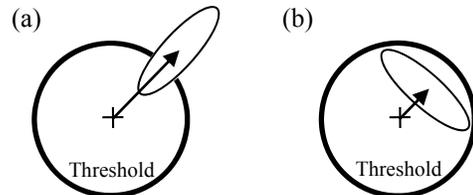


Fig. 1. Limiting cases for worst-case bias direction: (a) bias aligned with largest principal axis of covariance ellipse, for b somewhat greater than \sqrt{T} , and (b) bias aligned with shortest principal axis of covariance ellipse, for b somewhat less than \sqrt{T} .

Computations confirm that, when the bias magnitude b is somewhat larger than the threshold \sqrt{T} , the worst-case P_{md} does indeed occur for the zero-angle configuration shown in Fig. 1(a). More generally, though, if the bias shrinks, the worst-case may occur when the bias aligns with the minor axis or at some intermediate angle. That the worst angle varies over this range is shown in Fig. 2. The figure plots P_{md} as a function of the angle θ between the bias and the major axis of the covariance ellipse. Results are plotted for a 2D case with

$\sqrt{T} = 7$. Contours are shown for several values of bias magnitude b in the range $b \in [6, 8]$, a range chosen to include biases both above and below the threshold at $\sqrt{T} = 7$.

For additional clarification, Fig. 3 plots as a function of bias magnitude b the angle θ maximizing P_{md} for each contour. As such, Fig. 3 indicates that the worst-case bias direction varies continuously, transitioning from alignment with the largest principal axis ($\theta = 0^\circ$) when b is somewhat larger than \sqrt{T} to alignment with the shortest principal axis ($\theta = 90^\circ$) when b is somewhat smaller than \sqrt{T} . Intermediate values between these two limits are relatively sensitive to problem parameters, including the vector dimension and the values of T and \mathbf{Q} .

At present, no analytical formulae have been identified to compute the width of the transition region or the worst-case P_{md} values in the transition region; thus, somewhat computationally expensive numerical methods were used to find P_{md}^* in this example described by Fig. 1-Fig. 3.

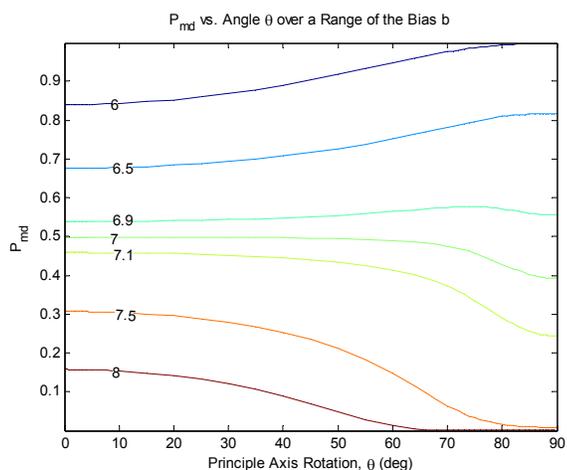


Fig. 2. P_{md} can increase or decrease as the rotation angle θ between bias and largest principal axis increases.

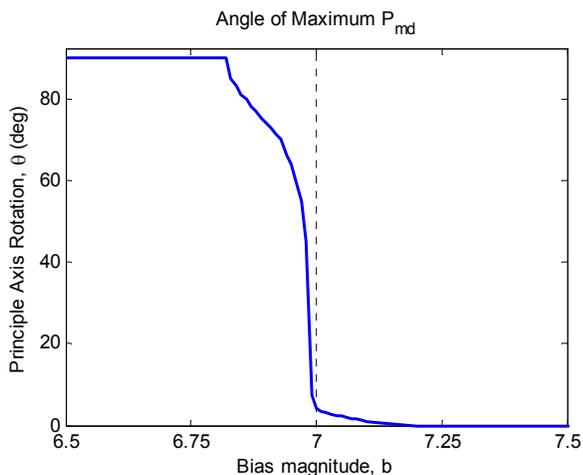


Fig. 3. As bias magnitude increases, the angle of the worst-case P_{md} decreases, indicating that the worst bias direction rotates toward the covariance matrix's largest principal axis.

Remark 2: Sigma Inflation is Not Always Conservative. In concept, bias-direction sensitivity might be resolved by inflating the covariance matrix. In other words, one might imagine widening the covariance ellipsoid shown in Fig. 4(a) into a larger covariance sphere, with equal variance in all directions. By widening the ellipsoid into a sphere, the problem would become rotationally invariant, and bias direction would not matter. It is even intuitive that this idea might be conservative, since widening a distribution (a process sometimes called “sigma inflation” in analysis of 1D distributions) is generally thought to increase P_{md} .

Although it is intuitively appealing to “inflate sigma” to remove bias-direction sensitivity, it cannot be assumed that covariance inflation is truly conservative. In fact, P_{md} may either increase or decrease when the covariance matrix is inflated, depending on the orientation of the covariance ellipse relative to the mean vector $\boldsymbol{\mu}$. A simple illustration of a case in which inflation is not conservative is shown in Fig. 4. In Fig. 4(a), an elongated covariance ellipse is shown, whose major axis is aligned with the bias direction. If the covariance ellipse is inflated such that the ellipse's minor axis length matches that of the major axis, then the covariance ellipse becomes a circle, as shown in Fig. 4(b). In the process of inflation, a portion of the probability distribution (shaded gray) that was originally inside the threshold (bold circle) moves outside the threshold. None of the remaining probability distribution (white area) moves inside the threshold, and so the total probability inside the threshold decreases. The net result is that covariance inflation underestimates risk in this case. This type of inflation is clearly not conservative if our goal is to obtain an estimate of P_{md} that is at least as large, or larger, than the actual risk.

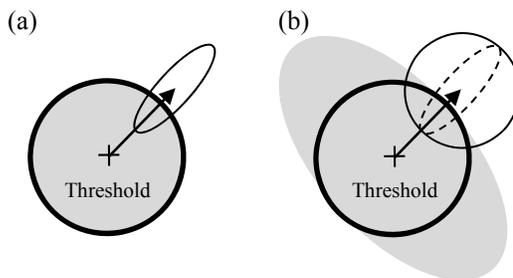


Fig. 4. Sigma inflation does not always result in higher P_{md} . Note the P_{md} integral sums all probability within the circular threshold. (a) Probability mass that contributes to the original P_{md} integral, before inflation, is shaded gray. (b) Inflation elongates the gray region, moving some gray-shaded probability mass outside the threshold and therefore reducing P_{md} .

This result that inflation sometimes underestimates risk is further supported by computation. Again consider the 2D case where the threshold is $\sqrt{T} = 7$. Also, consider bias magnitudes b varying between 6 and 8. Assume that the maximum principal axis of the covariance ellipse has unit length and aligns with the bias vector $\boldsymbol{\mu}$. The results of computing P_{md} versus minor axis length for contours of fixed b is shown in Fig. 5. It is evident that the computed risk P_{md} decreases as λ_{min} (the minor axis length squared) increases.

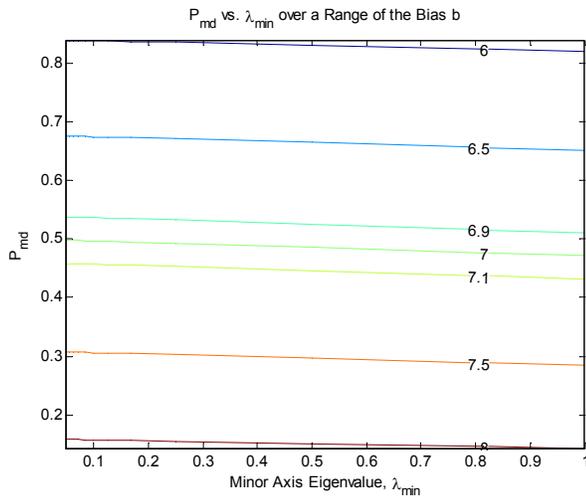


Fig. 5. Missed-detection risk P_{md} increases with minor axis length.

B. Relaxed Optimization Problem

In typical applications, missed-detection risk must be evaluated repeatedly, for a wide range of conditions, and so it is useful to have an algebraic (as opposed to a numerical) solution for P_{md} . The analysis above shows that intuitive bounding methods do not necessarily yield conservative P_{md} estimates if the bound is assumed to lie at an edge of the allowed parameter space. This section discusses an alternative approach, one that relaxes the optimization to obtain an analytical solution. The approach, called ellipsoid bounding, was first introduced in [28]. Ellipsoid bounding is guaranteed to be conservative.

Ellipsoid bounding relaxes the missed-detection risk maximization problem (13) by replacing the original volume of integration $F = \{\mathbf{y} | \mathbf{y}^T \mathbf{y} \leq T\}$ with a larger, inclusive volume of integration G . The worst-case missed-detection risk computed for the relaxation is \bar{P}_{md} , where

$$\bar{P}_{md} = \max_{\mathbf{Q} \in \Theta, \boldsymbol{\mu} \in \Psi} \left(\int_G p_N(\mathbf{y}; \mathbf{Q}, \boldsymbol{\mu}) dV \right). \quad (25)$$

It can be shown that \bar{P}_{md} is a conservative estimate of the actual worst case risk P_{md}^* , so long as the new volume of integration G contains the original volume of integration F . This result follows from the fact that the probability distribution is everywhere positive, so the integral over G must be larger than the integral over F if $G \supset F$.

$$P_{gx2}(T; \mathbf{Q}, \boldsymbol{\mu}) = \int_F p_N(\mathbf{y}; \mathbf{Q}, \boldsymbol{\mu}) dV \leq \int_G p_N(\mathbf{y}; \mathbf{Q}, \boldsymbol{\mu}) dV \quad (26)$$

For an appropriate choice of the volume G , the solution to the relaxed problem (25) is the following expression.

$$\bar{P}_{md} = P_{ncx} \left(T^*; DoF, b^{*2} \right) \quad (27)$$

Here P_{ncx} indicates a noncentral chi-square distribution. The apparent threshold and bias for this distribution are T^* and b^* .

$$T^* = \frac{\lambda_{\max}}{\lambda_{\min}^2} T \quad (28)$$

$$b^* = \left(b - \sqrt{T} \right) \sqrt{\lambda_{\max}^{-1}} + \sqrt{T^*}. \quad (29)$$

In the above expression, the values λ_{\min} and λ_{\max} are the smallest and largest possible eigenvalues over all axes, according to (12). In other words, $\lambda_{\min} \leq \lambda_n$ and $\lambda_{\max} \geq \lambda_n$ for all eigenvalues λ_n of covariance matrix \mathbf{Q} .

An outline of the proof that (27) is the solution to the relaxed optimization problem (25) is provided below; however, some details of the proof are left to the earlier reference [28].

In constructing the proof, the most important step is the definition of the volume of integration G . In order to simplify analysis, the relaxed boundary of integration is chosen to be an ellipsoid of the same shape as the covariance ellipsoid \mathbf{Q} . In the following analysis, the covariance ellipsoid is rotated into \mathbf{z} coordinates, as defined above; consequently the covariance matrix \mathbf{Q} is replaced with the diagonal eigenvalue matrix $\boldsymbol{\Lambda}$. In these coordinates, the boundary ∂G of the relaxed volume of integration G is

$$\partial G = \left\{ \mathbf{z} \mid \left(\mathbf{z} - \boldsymbol{\mu}_g \right)^T \boldsymbol{\Lambda}^{-1} \left(\mathbf{z} - \boldsymbol{\mu}_g \right) = T_g \right\}. \quad (30)$$

The parameter T_g is a scaling parameter that controls the size of the ellipse. The ellipse center is shifted by the vector $\boldsymbol{\mu}_g$. In order to ensure a tight bound, the offset $\boldsymbol{\mu}_g$ is selected such that the volumes F and G are tangent at a contact point \mathbf{z}_c that lies along the direction of the distribution mean $\boldsymbol{\mu}_c$.

$$\mathbf{z}_c = \frac{\sqrt{T}}{b} \boldsymbol{\mu}_c \quad (31)$$

Making the bound tangent to the original threshold at \mathbf{z}_c ensures the new boundary of integration is a close approximation to the original boundary precisely where probability is highest (near the mean $\boldsymbol{\mu}_c$). As such, this choice adds only a small degree of overconservatism when G replaces F . An example of an ellipsoid bound that is tangent to a spherical threshold is illustrated in 2D in Fig. 6.

As proved in [28], the two volumes are tangent if $\boldsymbol{\mu}_g$ is

$$\boldsymbol{\mu}_g = \left(\mathbf{I} - \frac{\sqrt{T_g}}{\sqrt{\mathbf{z}_c^T \boldsymbol{\Lambda} \mathbf{z}_c}} \boldsymbol{\Lambda} \right) \mathbf{z}_c. \quad (32)$$

It is shown in [28] that the ellipse G is large enough to entirely contain F as long as

$$T_g \geq \frac{\mathbf{z}_c^T \boldsymbol{\Lambda} \mathbf{z}_c}{\min_n (\lambda_n^2)}. \quad (33)$$

This condition depends on uncertain parameters: the covariance matrix $\boldsymbol{\Lambda}$, which is the eigenvalue matrix for the original covariance \mathbf{Q} , and the bias vector $\boldsymbol{\mu}_z$, which sets the contact point \mathbf{z}_c according to (31). To ensure that all possible parameter values are covered on the range $\mathbf{Q} \in \Theta$ and $\boldsymbol{\mu} \in \Psi$, it is sufficient to set $T_g = T^*$, where T^* is defined by (28). This result can be shown by noting in (33) that the largest value of the numerator, over $\boldsymbol{\mu} \in \Psi$, is $\lambda_{max} T$; likewise, the smallest value of the denominator, over $\mathbf{Q} \in \Theta$, is λ_{min}^2 .

Having set $T_g = T^*$, the ellipse G is fully specified. The next step is to show that expression (27) is a solution to the relaxed optimization problem (25) for G described by (30)-(32). To do this, introduce another coordinate transformation, substituting \mathbf{w} for \mathbf{z} where

$$\mathbf{w} = \boldsymbol{\Lambda}^{-1/2} (\mathbf{z} - \boldsymbol{\mu}_g). \quad (34)$$

This substitution transforms G from an ellipsoid into a sphere. The substitution does the same for the Gaussian probability distribution, mapping its contours of constant probability from ellipsoids into spheres. The resulting covariance is the identity matrix \mathbf{I} .

$$\bar{P}_{md} = \max_{\mathbf{Q} \in \Theta, \boldsymbol{\mu} \in \Psi} \left(\int_{\mathbf{w}^T \mathbf{w} \leq T_g} P_N(\mathbf{w}; \mathbf{I}, \boldsymbol{\mu}_{tot}) dV_w \right) \quad (35)$$

The transformed bias is

$$\boldsymbol{\mu}_{tot} = \boldsymbol{\Lambda}^{-1/2} (\boldsymbol{\mu}_z - \boldsymbol{\mu}_g). \quad (36)$$

In \mathbf{w} coordinates, the integral has now become a noncentral chi-square distribution. The noncentral chi-square distribution P_{ncx} is a special case of the generalized chi-square distribution, one with an identity covariance matrix. Because the identity covariance matrix is symmetric under rotation, the direction of the bias $\boldsymbol{\mu}_{tot}$ does not matter. Only its magnitude b_{tot} matters.

The magnitude-squared $b_{tot}^2 = \|\boldsymbol{\mu}_{tot}\|^2$ is sometimes called the *noncentrality parameter*, and the rank of the identity matrix \mathbf{I} is sometimes labeled degrees-of-freedom *DOF*. Using this notation, the following formula can be written to relate the noncentral and generalized chi-square distributions.

$$P_{ncx}(T^*; DoF, b_{tot}^2) = \int_{\mathbf{w}^T \mathbf{w} \leq T^*} P_N(\mathbf{w}; \mathbf{I}, \boldsymbol{\mu}_{tot}) dV \quad (37)$$

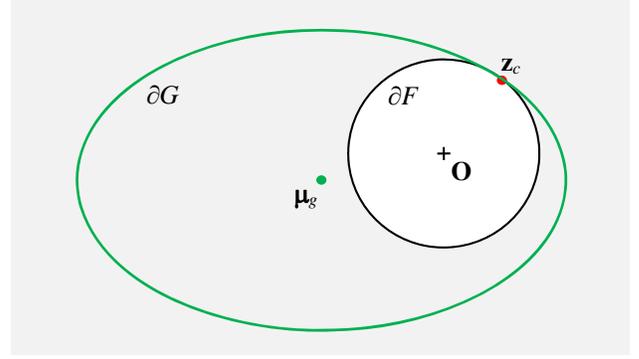


Fig. 6. Bounding Concept – An ellipsoid ∂G with fixed shape must contain a spherical volume ∂F when both surfaces are tangent at \mathbf{z}_c .

Thus the integral in (35) may be rewritten as follows.

$$\bar{P}_{md} = \max_{\mathbf{Q} \in \Theta, \boldsymbol{\mu} \in \Psi} \left(P_{ncx}(T^*; DoF, b_{tot}^2) \right) \quad (38)$$

In this expression, the parameters T^* and DoF are fixed constants. The only remaining variable is the mean vector $\boldsymbol{\mu}_z$, whose direction influences the magnitude of b_{tot} . Importantly, the dependence of the distribution on b_{tot} is a simple relationship; the noncentral chi-square probability always decreases as the noncentrality parameter grows:

$$\frac{\partial}{\partial b_{tot}} P_{ncx}(T^*; DoF, b_{tot}^2) < 0. \quad (39)$$

Thus, (38) is maximized when the bias magnitude b_{tot} takes its minimum value. Label this value b^* ,

$$b^* = \min_{\mathbf{Q} \in \Theta, \boldsymbol{\mu} \in \Psi} (b_{tot}). \quad (40)$$

The original problem of maximizing an integral has been transformed into the much simpler problem of minimizing a scalar expression. It is convenient to square the scalar b_{tot} before finding this minimum, as the squaring operation does not affect the parameter values resulting in the minimum, given that b_{tot} is positive. An expression for b_{tot} squared is

$$b_{tot}^2 = \|\boldsymbol{\mu}_{tot}\|^2 = \left(\frac{b}{\sqrt{T}} - 1 \right)^2 \mathbf{z}_c^T \boldsymbol{\Lambda}^{-1} \mathbf{z}_c + 2 \left(\frac{b}{\sqrt{T}} - 1 \right) \frac{T \sqrt{T_g}}{\sqrt{\mathbf{z}_c^T \boldsymbol{\Lambda} \mathbf{z}_c}} + T_g. \quad (41)$$

This expression is obtained by combining (31), (32), and (36). Further simplification is possible by noting the contact point \mathbf{z}_c lies on the sphere of radius \sqrt{T} in an arbitrary direction, which can be described by an arbitrary unit vector $\hat{\mathbf{u}}$. Substituting $\mathbf{z}_c = \sqrt{T} \hat{\mathbf{u}}$ gives

$$b_{tot}^2 = (b - \sqrt{T})^2 \hat{\mathbf{u}}^T \mathbf{\Lambda}^{-1} \hat{\mathbf{u}} + 2(b - \sqrt{T}) \frac{\sqrt{T_g}}{\sqrt{\hat{\mathbf{u}}^T \mathbf{\Lambda} \hat{\mathbf{u}}}} + T_g. \quad (42)$$

This expression is minimized when the vector $\hat{\mathbf{u}}$ aligns with the maximum principal axis of the eigenvalue matrix $\mathbf{\Lambda}$.

$$(b^*)^2 = (b - \sqrt{T})^2 \frac{1}{\max_n(\lambda_n)} + 2(b - \sqrt{T}) \frac{\sqrt{T_g}}{\sqrt{\max_n(\lambda_n)}} + T_g \quad (43)$$

Noting the maximum eigenvalue of $\mathbf{\Lambda}$ over the set Θ has previously been labeled λ_{max} , complete the squares.

$$b^* = (b - \sqrt{T}) (\sqrt{\lambda_{max}})^{-1} + \sqrt{T_g} \quad (44)$$

This step recovers equation (29) and completes the proof that (27) is an analytical solution that optimizes (25).

C. Quantifying Conservatism in Relaxed Bound

The bound \bar{P}_{md} derived in the previous section is guaranteed to be conservative, such that $\bar{P}_{md} \geq P_{md}^*$. It has been hypothesized that the ellipsoid bound is tight, because the ellipsoid threshold closely approximates the original spherical threshold at the location where probability is highest as shown in Fig. 6; however, this claim has not previously been tested. The goal of this section is to apply simulation to assess overconservatism in two ways: by comparison to cases from the original design space ($\boldsymbol{\mu} \in \Psi$ and $\mathbf{Q} \in \Theta$) and also by comparison to an alternative bounding strategy called determinant bounding [27].

Simulations considered in this section are representative of a particular class of integrity monitor used in safety-critical satellite navigation: Signal Deformation Monitoring (SDM). SDM is a key monitor for two civil aviation systems, the Satellite-Based Augmentation System (SBAS) and Ground-Based Augmentation System (GBAS) [19]-[21]. To model the GBAS implementation of SDM, a 7 degree-of-freedom system is considered. The \mathbf{Q} matrix is assumed to be uncertain but constrained according to (12), with eigenvalues between upper and lower limits. Upper eigenvalue limits are unity on all axes. Lower eigenvalue limits $\lambda_{min,n}$ are drawn from [27]. The set of lower eigenvalue limits is the following.

$$\lambda_{min,n} = \{0.42, 0.45, 0.56, 0.59, 0.64, 0.75, 0.93\} \quad (45)$$

Per constraint (14), bias direction is arbitrary. The monitor threshold is set to a value of $T = 45.7$ (representing a nominal chi-square false-alarm probability of 10^{-7}). Bias magnitude values b ranging from 0 to 10 are considered, so as to fall on either side of \sqrt{T} (i.e., on either side of 6.76).

For these parameters, four different estimates of P_{md} are

evaluated. These include (1) a baseline model, (2) an actual case, (3) the ellipsoid bound, and (4) the determinant bound. Each of these estimates is described in more detail, below.

The first P_{md} estimate is the baseline, which assumes the covariance matrix has the widest principal axes possible ($\mathbf{Q} = \mathbf{I}$). The baseline is not guaranteed conservative, as P_{md} risk can shrink for a wider covariance matrix, as seen in Fig. 4. The baseline model P_{base} is evaluated

$$P_{base} = P_{ncx}(T; DoF, b^2). \quad (46)$$

The second P_{md} estimate is an actual case from the design space ($\mathbf{Q} \in \Theta$). Specifically, eigenvalues are set to their minimum values for all principle axes except the largest principle axis, for which the eigenvalue is set to unity:

$$\lambda_n = \{0.42, 0.45, 0.56, 0.59, 0.64, 0.75, 1.0\}. \quad (47)$$

The bias vector $\boldsymbol{\mu}$ is chosen to be aligned with the largest principle axis of \mathbf{Q} . This choice appears to be the set of parameters resulting in the worst P_{md} when the bias magnitude b is large, as illustrated in Fig. 1 through Fig. 3. To control numerical accuracy in integrating the generalized chi-square distribution function of (10) for this actual case, a Monte Carlo simulation with 10^6 trials is used.

The third P_{md} estimate is the ellipsoid bound \bar{P}_{md} , which is computed by evaluating equations (27) through (29). Because \bar{P}_{md} has been shown to be a conservative bound, \bar{P}_{md} is expected to exceed the actual P_{md} values for all parameter sets in the allowed parameter space (including the ‘‘actual case’’ simulated with Monte Carlo methods, as described above).

The fourth P_{md} estimate is the determinant bound P_{det} . This bound was proposed in [27] as a conservative overestimate of the generalized chi-square distribution ($P_{det} > P_{md}^*$). The form of the determinant bound is the following:

$$P_{det} = \min \left[1, \prod_n \frac{1}{\lambda_{min,n}^{1/2}} P_{ncx}(T; DoF, b^2) \right]. \quad (48)$$

The determinant bound is considered here for the sake of comparison, as a means of evaluating the level of overconservatism for the ellipsoid bound.

The four models are plotted together on axes of P_{md} versus bias magnitude b in Fig. 7. The baseline curve is shown as a dash-dotted green line, the Monte Carlo trials are plotted as red circles, the ellipsoid bound \bar{P}_{md} is plotted as a solid blue line, and the determinant bound as a dotted mauve line. On the plot, the region where the bias falls inside the threshold has a gray shaded background, and the region where the bias falls outside the threshold has a white background.

A first observation is that the baseline estimate does not conservatively bound the actual (Monte Carlo) case; rather, the

baseline case is everywhere lower. Hence, the baseline estimate should not be used to compute the missed-detection risk for a safety-critical system if input noise parameters are uncertain.

A second observation is that both the ellipsoid and determinant bounds are confirmed to be conservative, in that they both somewhat overestimate risk for the actual case simulated with Monte Carlo methods. It is readily apparent that, at least for this analysis, the ellipsoid bound is significantly tighter than the determinant bound. That the ellipsoid bound is a much tighter estimate than the determinant bound is confirmed in Fig. 8, which plots the data from Fig. 7 after normalizing each P_{md} value by the baseline estimate (46) for the same bias magnitude b . Normalization removes the trend in Fig. 7, making it clear that the determinant bound (dotted curve) lies much farther than the ellipsoid bound (solid curve) from the Monte Carlo case (circles). Even on the right side of the plot, where the ellipsoid bounds starts to deviate slightly from Monte Carlo simulation (30% inflation), the determinant bound is still substantially more conservative (approximately 250% inflation).

A minor note regarding Fig. 8 is that the ellipsoid bound (solid curve) appears to approach the determinant bound (dotted curve) as bias increases. The determinant bound reaches a plateau (a constant multiplier times the baseline model) once the overbound value drops below unity, which happens when the bias becomes somewhat larger than the threshold. It is a reasonable question to ask whether the ellipsoid and determinant bounds ever cross, such that the determinant bound becomes tighter than the ellipsoid bound. The answer is that a crossing does eventually occur, as illustrated by Fig. 9, which plots the normalized P_{md} bounds from Fig. 8 as a function of the normalizing value P_{opt} , from (46). The crossing occurs at a P_{opt} of 10^{-80} , which is much, much smaller the P_{md} requirements for aviation safety systems (which typically put requirements on P_{md} risk at a level of 10^{-9} or larger). In other words, for practical safety-of-life designs, it is unlikely that the determinant bound will ever outperform the ellipsoid bound.

Full details of applying the ellipsoid overbound to integrity verification are beyond the scope of this paper; however, it is worth briefly noting the relevance of P_{md} overbounding for SDM verification in GBAS. Monitor verification requires that, for a particular value of user-ranging error E , no threat has P_{md} greater than a specified limit. A P_{md} versus b curve (as shown in Fig. 8) can easily be converted into the required P_{md} versus E curve, by using existing threat models that relate bias magnitude b to a worst associated error E [19]-[21]. In this sense, the ellipsoid overbounding model presented here can serve a key role in SDM verification.

V. SUMMARY

This paper introduced tools for tight but conservative performance characterization of chi-square integrity monitors for which the input noise is characterized by uncertain parameters. A first contribution was to frame as optimization problems the determination of worst-case parameter configurations that maximize false-alarm and missed-detection

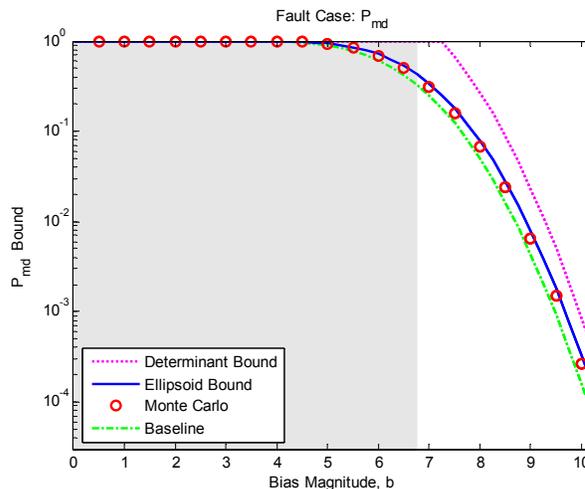


Fig. 7. Comparison of models of missed-detection probability as a function of bias magnitude b

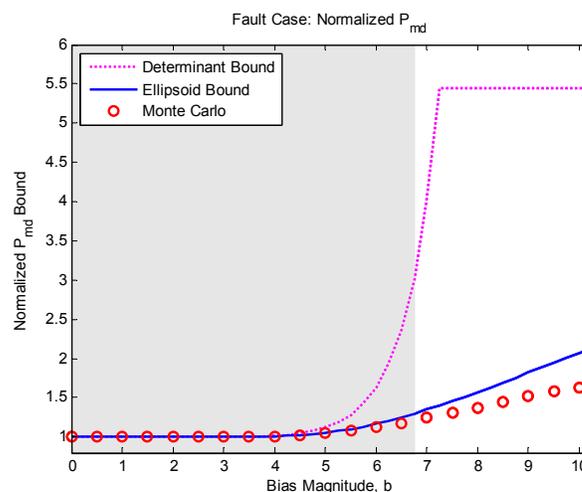


Fig. 8. Comparison of models of missed-detection probability, after normalization by a noncentral chi-square cumulative distribution

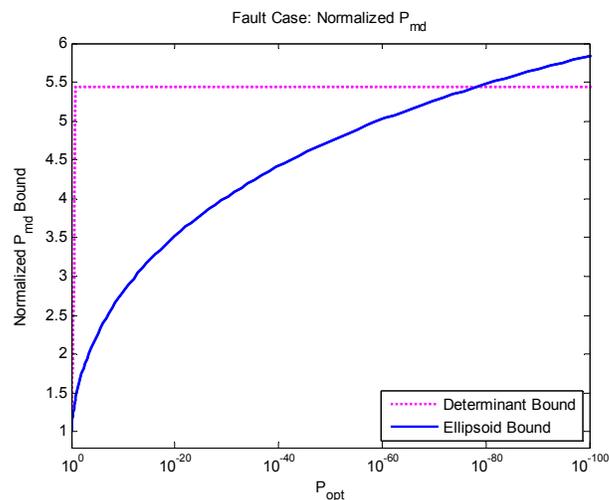


Fig. 9. Cross-over point where ellipsoid bound meets determinant bound

risk. A second contribution was to prove that an exact, analytical solution is possible for the problem of computing false-alarm risk. A third contribution was to provide an illustrative example to explain why a similar exact, analytical solution has not been found for the missed-detection risk optimization problem. A final contribution was to demonstrate, through simulation, that an approximate analytical solution called the ellipsoid bound can provide a very tight approximation for missed-detection risk.

Simulations of P_{md} indicated that the ellipsoid bound is very tight. As compared to an actual distribution from the allowed parameter space, the ellipsoid bound was inflated by no more than 30% (out to a P_{md} risk of 10^{-5}). The implication is that ellipsoid bounding is a useful (conservative but tight) method for approximating P_{md} for chi-square monitors where parameter uncertainty exists and where an analytical solution is needed to streamline performance analysis.

REFERENCES

- [1] Hegarty, C. and Chatre, E., "Evolution of the Global Navigation Satellite System (GNSS)," *Proceedings of the IEEE*, Vol. 96, No. 12, 2008, pp. 1902-1917.
- [2] Rife, J. and Pullen, S., "Aviation applications," *GNSS Applications and Methods*, S. Gleason and D. Gebre-Egziabher, Eds., Artech House, 2009, pp. 245-268.
- [3] DeCleene, B., "Defining pseudorange integrity - overbounding," *Proceedings of the ION-GPS*, 2000, pp. 1916-1924.
- [4] Rife, J., Walter, T., and Blanch, J., "Overbounding SBAS and GBAS error distributions with excess-mass functions," *Proc. ION-GNSS*, Sydney, Australia, Dec. 6-8, 2004.
- [5] Rife, J., Pullen, S., Pervan, B., and Enge, P., "Paired overbounding for nonideal LAAS and WAAS error distributions," *IEEE Trans. Aerospace and Electronic Systems*, Vol. 42, No. 4, 2006, pp. 1386-1395.
- [6] Rife, J., Pullen, S., and Pervan, B., "Core overbounding and its implications for LAAS integrity," *Proc. ION-GNSS*, 2004, pp. 2810-2821.
- [7] Rife, J. and Pervan, B., "Overbounding revisited: discrete-error distribution modeling for safety-critical GPS navigation," *IEEE Trans. Aerospace and Electronic Systems*, Vol. 48, No. 2, 2012, pp. 1537-1551.
- [8] Liu, F., Murphy, T., Skidmore, T., "LAAS Signal-in-Space Integrity Monitoring Description and Verification Plan," *Proc. ION-GPS*, 1997, pp. 485-497.
- [9] Shively, C. and Hsiao, T., "Availability Enhancements for Cat IIIb LAAS," *NAVIGATION*, Vol. 51, No. 1, 2004, pp. 45 - 57.
- [10] Cohenour, C., van Graas, F., "GPS orbit and clock error distributions," *NAVIGATION*, Vol. 58, No. 1, 2011, pp. 17-28.
- [11] Osechas, O., Misra, P., Rife, J., "Carrier-phase acceleration RAIM for GNSS satellite clock fault detection," *NAVIGATION*, Vol. 59, No. 3, 2012, pp. 221-235.
- [12] Joerger, M. and Pervan, B., "Kalman filter-based integrity monitoring against sensor faults," *Journal of Guidance, Control, and Dynamics*, Vol. 36, No. 2, 2013, pp. 349-361.
- [13] Langel, S., Khanafseh, S., and Pervan, B., "Bounding integrity risk for sequential estimators with stochastic modeling uncertainty," *AIAA J. of Guidance, Control, and Dynamics*, Vol. 37, No. 1, 2014.
- [14] Rife, J., and Gebre-Egziabher, D., "Symmetric overbounding of correlated errors," *NAVIGATION*, Vol. 54, No. 2, 2007, pp. 109-124.
- [15] Pulford, G., "A proof of the spherically symmetric overbounding theorem for linear systems," *NAVIGATION*, Vol. 55, No. 4, 2008, pp. 283-292.
- [16] Brumback, B., Srinath, M. D., "A fault-tolerant multisensor navigation system design," *IEEE Trans. Aerospace and Electronic Systems*, Vol. 23, No.6, 1987, pp.738-756.
- [17] Da, R., "Failure detection of dynamical systems with the state chi-square test," *AIAA J. of Guidance, Control, and Dynamics*, Vol. 17, No. 2, 1994.
- [18] Brown, R.G., "Receiver autonomous integrity monitoring," In *Global Positioning System: Theory and Applications, Volume II*, B.W. Parkinson and J.J. Spilker, Eds., Washington, D.C.: AIAA, 1996.
- [19] Phelts, R.E., Walter, T., and Enge, P., "Toward real-time SQM for WAAS: improved detection techniques," *Proc. ION-GNSS*, 2003, pp. 2739-2749.
- [20] Liu, F., Brenner, M., Tang, C.Y., "Signal deformation monitoring scheme implemented in a prototype local area augmentation system ground installation," *Proc. ION-GNSS*, 2006, pp. 367-380.
- [21] Rife, J., and Phelts, R.E., "Formulation of a time-varying maximum allowable error for Ground-Based Augmentation Systems," *IEEE Trans. Aerospace and Electronic Systems*, Vol. 44, No. 2, 2008, pp. 548-560.
- [22] Khanafseh, S., Pullen, S., and Warburton, J., "Carrier Phase Ionospheric Gradient Ground Monitor for GBAS with Experimental Validation," *NAVIGATION*, Vol. 59, No. 1, 2012, pp. 51-60.
- [23] Reuter, R., Weed, D., Brenner, M., "Ionosphere gradient detection for Cat III GBAS," *Proc. ION-GNSS*, 2012, pp. 2175-2183.
- [24] Grabowski, J.C., "Field observations of personal privacy devices," *Proc. ION-ITM*, 2012, pp. 689-741.
- [25] Pullen, S., Gao, G., Tedeschi, C., Warburton, J., "The impact of uninformed RF interference on GBAS and potential mitigations," *Proc. ION-ITM*, 2012, pp. 780-789.
- [26] Rife, J.H., "Overbounding missed-detection probability for a chi-square monitor," *Proc. ION-GNSS*, 2012, pp. 1348-1368.
- [27] Rife, J. H., "The effect of uncertain covariance on a chi-square integrity monitor," *NAVIGATION*, Vol. 60, No. 4, 2013, pp. 291-303.
- [28] Rife, J.H. and Schuldt, D., "Minimum volume ellipsoid scaled to contain a tangent sphere, with application to RAIM," *Proc. IEEE PLANS*, 2014, Monterey, CA.
- [29] Joerger, M. and Pervan, B., "Kalman filter-based integrity monitoring against sensor faults," *AIAA J. Guidance, Control, and Dynamics*, Vol. 36, No. 2, 2012, pp. 349-361.
- [30] Peck, R. and Devore, J. *Statistics: The exploration & analysis of data*, Cengage Learning, 2011.
- [31] Enge, P., "Local area augmentation of GPS for the precision approach of aircraft," *Proc. of the IEEE*, Vol. 87, No. 1, 1999, pp. 111-132.
- [32] Rife, J.H., "Practical ellipsoid bounding for chi-square integrity monitors, including GNSS signal deformation monitoring," in preparation for *NAVIGATION*, 2014.
- [33] Golub, G., and Van Loan, C., *Matrix Computations*, 3rd Edition, John Hopkins University Press, 1996.



Jason H. Rife (M'01) received his B.S. in mechanical and aerospace engineering from Cornell University, Ithaca, NY, and his M.S. and Ph.D. in mechanical engineering from Stanford University, Stanford, CA.

He is currently an Associate Professor of Mechanical Engineering at Tufts University in Medford, MA. At Tufts, he directs the Automation Safety and Robotics Laboratory (ASAR), which applies theory and experiment to characterize the integrity of autonomous vehicle systems. Previously, he has worked at Pratt & Whitney, as a development engineer in the Turbine Aerodynamics group, and at the Stanford University GPS Laboratory, as a research engineer supporting FAA and Navy projects for precise landing of aircraft on land and at sea using GPS.