Control Cost and Mahalanobis Distance Binary Hypothesis Testing for Spacecraft Maneuver Detection

AE8900 MS Special Problems Report
Space Systems Design Lab (SSDL)
Guggenheim School of Aerospace Engineering
Georgia Institute of Technology
Atlanta, GA

Author:
Andris D. Jaunzemis

Advisor:
Marcus J. Holzinger

December 11, 2015

Control Cost and Mahalanobis Distance Binary Hypothesis Testing for Spacecraft Maneuver Detection

Andris D. Jaunzemis*, Midhun V. Mathew†, Marcus J. Holzinger‡
Georgia Institute of Technology, Atlanta, GA 30332

An anomaly hypothesis testing technique using the minimum-fuel control distance metric is extended to incorporate non-Gaussian boundary condition uncertainties and employ binary hypothesis testing. The adjusted control distance metric utilizes Gaussian mixtures to model non-Gaussian boundary conditions, and binary hypothesis testing allows inclusion of anomaly detection thresholds and allowable error rates. An analogous framework accommodating Gaussian mixtures and binary hypothesis testing is developed. Both algorithms are compared using simulated and empirical satellite maneuver data. The North-South station-keeping scenario shows control distance to be less sensitive with increased uncertainty than Mahalanobis distance but more consistent with respect to observation gap duration, a trend which is corroborated using available real-world data. The same consistency with respect to observation gap is observed in East-West station-keeping while also showing control distance metric to be more sensitive for shorter observation gaps. In the non-Gaussian boundary condition case, control distance outperforms Mahalanobis distance in both detection and computational complexity.

I. Introduction

Correlating on-orbit observations and detecting space object maneuvers is a challenging endeavor in Space Situational Awareness (SSA). There are currently at least 20,000 trackable on-orbit objects, over 1,000 of which are active [1, 2], and these numbers are expected to grow significantly due to improved tracking capabilities, new launches, and continued debris generation [3]. Predicting conjunction events is a difficult task [4]. Recent events, such as the Iridium-Cosmos on-orbit collision, highlight the mutual interest that national and private operators share for accurate object correlation and maneuver detection capability [5]. Detecting maneuvers is particularly important when tracking active objects for which no operational information is available, as detecting maneuvers in real-time is required to adequately react to anomalies or possible conjunctions [6, 7]. Timely detection of maneuvers allows for responsiveness in follow-up tracking, which is crucial for post-maneuver orbit characterization [6]. This task is especially difficult in low-thrust maneuvers, where state change is more subtle [8]. Hoots et al. provide an overview of tracking techniques used in the space surveillance system, noting that the operational centers still depend largely on orbit models and applications from the 1950s and 1960s [9]. These techniques often rely on input from human analysts who excel at evaluating information and drawing conclusions based on incomplete data, whereas automated tools often fall short in these situations [3]. Pending improvements in space surveillance sensors are predicted to increase the number of cataloged objects by an order of magnitude and the frequency of measurement updates by a factor of 4 or more [3]. The increase in data volume and complexity of decision will render human analyst processes impractical, motivating a need for automated decision processes.

A fundamental problem in SSA is state uncertainty, particularly in track association and maneuver detection [10]. Initially, state uncertainty may be assumed to be Gaussian, which is often appropriate for resident space objects; however, the orbital dynamics (e.g. gravity, drag, third body perturbations) are nonlinear, and after propagation the state uncertainty may become non-Gaussian [10]. One typical solution is to model a non-Gaussian distribution using a weighted sum of Gaussian distributions, termed a Gaussian mixture model. Many arbitrary probability density functions can be closely approximated by a Gaussian mixture [10].

Data association algorithms for object correlation and maneuver detection have been well explored in literature. The data association task hypothesizes an association and attempts to compute a measure of the probability that the hypothesis is true. Methods often focus on admissible regions or probabilistic approaches to compare uncorrelated tracks and detect maneuvers. For instance, Tommei et al. address object correlation and orbit determination with admissible region-based methods and a virtual debris algorithm that were applied to optical observations and radar observations [11]. Maruskin et al. also use admissible regions for object correlation by mapping admissible regions to Delaunay or Hamiltonian orbit elements and by comparing an observation with an earlier estimate propagated forward in time [12]. Fujimoto et al. employ highly constrained probability distributions in Poincare orbit element space, where distributions are defined by admissible-region maps such that the intersection between

admissible regions, or lack thereof, reflects the correlation between respective observations [13].

DeMars et al. approach the SSA data association problem using finite-set statistics to create a Bayesian approach for multi-object estimation. The adaptive entropy-based Gaussian-mixture information synthesis (AEGIS) approach is used to track objects while utilizing finite-set statistics (FISST) to account for uncertainty, false alarms and misdetections [14]. Kelecy and Jah apply batch least-squares and extended Kalman filter based strategies to detect and reconstruct low thrust finite maneuvers [8]. Huang et al. relate uncorrelated tracks (UCTs) using a nonlinear least squares iterative process to optimally estimate maneuvers following a maximum a posteriori criterion for object correlation and maneuver detection [15].

Likewise, there are many distance or pseudo-distance metrics that may be used to measure the discrepancy between two state distributions (e.g. Mahalanobis distance) [16]. Problematically, existing metrics do not directly quantify the level of propulsive effort required to cause the observed state change. The problem of associating UCTs over large time periods is particularly difficult when resident space objects (RSOs) maneuver during observation gaps. Even relatively small station-keeping maneuvers at geostationary Earth orbit (GEO) can result in position discrepancies of many kilometers after an observation gap. UCT correlation is further confounded by state estimate uncertainties [17]. Since both the initial and final UCTs are best estimates, with associated probability density functions (PDFs), correlation is difficult particularly in densely-populated regions of the space environment. Given a propagated estimate of the state PDF, correlating UCTs tests whether a new observation (with its associated uncertainty) is a previously observed object, and if not, determines the discrepancy.

Holzinger et al. propose a minimum-fuel control distance metric to approach data association and maneuver detection while considering propulsive effort and reconstructing maneuvers [18]. Since on-board fuel remains a scarce commodity for operational spacecraft, operators are likely to execute fuel-optimal, or near-optimal, maneuvers [17]. Under the assumption of optimal control, multiple deterministic UCTs can be related by computing the control effort required for a spacecraft trajectory to meet those boundary conditions. This approach necessitates the reconstruction of a minimum-fuel trajectory consistent with the a priori information and new observations. Holzinger et al. have shown, through the properties of strict positivity, symmetry, and triangle inequality, that control performance is a metric [17], allowing objective comparisons to other commonly used distance metrics. The control cost distributions required can be computed along each relevant trajectory by considering boundary condition uncertainty [18]. In the previous work, a single cost distribution was developed using the maneuvered trajectory, testing the anomaly hypothesis by comparing the size of the observed maneuver to the amount of uncertainty in control cost due to uncertainty in the boundary conditions [18]. This straight-forward approach encounters problems when attempting to draw conclusions regarding error rates, making comparisons to other metrics incomplete.

A commonly-used statistical approach in anomaly detection that incorporates error rates is binary hypothesis testing, wherein integration over a pair of PDFs allows determination of false alarm and missed detection probabilities. Binary hypothesis testing has been implemented for anomaly detection in various fields, such as signal processing [19]. One variant of binary hypothesis testing, the Neyman-Pearson approach, devises the most powerful likelihood-ratio test for a given significance level and threshold [20]. In typical Neyman-Pearson detector implementation, these thresholds are selected through analysis of a number of observations with associated PDFs; however, in applications involving only one observation, this method reduces to a more basic form of binary hypothesis testing, which is the case for UCT association where only one PDF is available at each time epoch.

Evaluating metrics for anomaly detection requires selection of a representative subsection of the infinite continuum of possible maneuvers. Since optical observations are primarily useful for space objects at high altitudes, such as GEO, maneuvers relevant to GEO spacecraft are particularly interesting. Spacecraft in GEO are assigned to specific longitude slots for their operational lifetime, but are also subject to a number of perturbations that must be rejected using station-keeping maneuvers. For instance, North-South station-keeping maneuvers adjust inclination, which is primarily perturbed by third-body gravitational effects, to maintain an equatorial orbit. East-West station-keeping maneuvers adjust true longitude (or phase), which is primarily perturbed by Earth oblateness sectoral harmonics (i.e. J_{22}), to maintain the spacecraft's GEO slot. These two primary station-keeping maneuvers form a representative subset of maneuvers that could potentially be encountered in operation.

This study modifies the approach to control distance anomaly detection by introducing a binary hypothesis structure, using control cost distributions from both the maneuvered and quiescent (non-maneuvered) trajectories. For this study, the hypothesis integration threshold is selected to match a desired rate of false alarms using the pair of control cost PDFs. Additionally, the control distance metric is extended for application to non-Gaussian boundary conditions using Gaussian mixtures. Alternate formulations of the Mahalanobis distance have already been formulated using Gaussian mixtures [21], so this study develops a similar formulation for the control distance metric. The primary contributions of this paper are: 1) a computationally tractable Gaussian mixture approximation to computing control cost distributions for non-Gaussian boundary condition probability densities,

2) a binary hypothesis testing framework for anomaly detection using the control distance metric that permits specification of false alarm rates, 3) an analogous binary hypothesis testing framework for anomaly detection using Mahalanobis distance, and 4) the implementation of these two methods with quantitative performance comparisons drawn between the control distance metric and Mahalanobis distance using both synthetic and real-world data. These metrics are compared using a GEO spacecraft in both North-South and East-West station-keeping test cases.

The remainder of the paper begins with a review of the relevant background theory, which is presented in Section II, allowing discussion of the novel theoretical contributions of this paper in Section III. Section IV contains a discussion of the simulation implementation of the control distance and Mahalanobis distance detection methods, followed by simulation results using synthetic data for a GEO spacecraft performing both North-South and East-West station-keeping maneuvers. Corresponding results using real data are presented in Section V. A subset of the data for these simulations is presented in each section for discussion, but Appendix A contains a more complete selection of the simulation data for reference. Finally, the results from the synthetic and real data simulations are distilled into potential operational applications in Section VI, followed by concluding remarks in Section VII.

II. Background

This section provides background information required for novel theory development in Section III, including description of the anomaly detection scenario and background on Gaussian mixture models and Mahalanobis distance.

A. Scenario

The following notional scenario is relevant to the task of associating a pair of UCTs to detect maneuvers. As pictured in Fig. 1, an uncorrelated track UCT_A at time t_0 is used to generate a PDF, $f_A(\mathbf{x}(t_0))$. Using the mean state of the PDF, $\mathbf{x}_\mu(t_0) = \mathbb{E}[\mathbf{x}(t_0)]$, the track is propagated to time t_f under assumed quiescent dynamics, $\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t) = \mathbf{0}, t)$, yielding the propagated mean state $\mathbf{x}_\mu(t_f)$ and its associated PDF $f_A(\mathbf{x}(t_f))$. Another uncorrelated track at time t_f, UCT_B, yields its own PDF $f_B(\mathbf{x}^*(t_f))$ and mean state $\mathbf{x}^*_\mu(t_f) = \mathbb{E}\left[\mathbf{x}^*(t_f)\right]$. An optimal maneuvered trajectory is generated connecting the mean states of the UCTs, $\mathbf{x}_\mu(t_0)$ and $\mathbf{x}^*_\mu(t_f)$, under the same dynamics, $\dot{\mathbf{x}}^*(t) = \mathbf{f}(\mathbf{x}^*(t), \mathbf{u}^*(t), t)$, by minimizing some objective function J. This yields the maneuvered trajectory $\mathbf{x}^*_\mu(t)$ and associated control $\mathbf{u}^*_\mu(t)$.

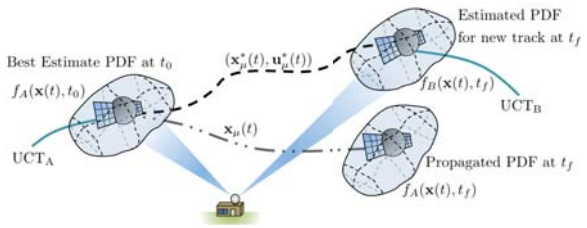

Figure 1. Maneuver detection scenario

B. Optimal Connecting Trajectories

The anomaly detection algorithm begins with a trajectory optimization routine, which uses mean states of the boundary conditions to generate boundary time-fixed optimal connecting trajectories. The UCT pair is considered a two-point boundary value problem (TPBVP), and a trajectory is computed to optimally connect the two UCTs based on a chosen cost function, J. In this study, the trajectory is optimized by minimizing the quadratic control cost, shown in Eq. (1) [22]. This particular cost function satisfies the metric properties of non-negativity, coincidence, symmetry, and triangle inequality.

$$J(\mathbf{u}(\mathbf{t}); t_0, t_f) = \frac{1}{2}\int_{t_0}^{t_f} \mathbf{u}(\tau)^\text{T}\mathbf{u}(\tau)\,\text{d}\tau. \tag{1}$$

The quadratic control cost function is ideal for variable specific impulse (VSI) engines, often used in low-thrust applications, which have become more prevalent in recent years. Alternate cost functions could also be implemented, such as an impulsive cost function, as long as they also satisfy metric properties.

C. Gaussian Mixture Models

The existing method to generate control cost PDFs assumes normally distributed boundary conditions for simplicity [17]; however, spacecraft state estimates are not always well-modeled as Gaussian distributions. This motivates the use of a more general formulation able to handle non-Gaussian boundary conditions using Gaussian mixture models. Previous uses of Gaussian mixtures for space applications include AEGIS-FISST, which leverages an adaptive algorithm to spawn or destroy Gaussian components of the Gaussian mixture to increase its fidelity while maintaining computational tractability [14].

A Gaussian mixture is created by overlapping multiple Gaussian distributions with different weights. For reference, the probability density of a state $\mathbf{x}$ in a multivariate normal distribution of n-dimensions with mean $\mathbf{x}_{\mu,i}$ and covariance $\mathbf{P_i}$ is defined using Eq. (2).

$$f(\mathbf{x}; \mathbf{x}_{\mu,i}, \mathbf{P_i}) = \frac{1}{(2\pi)^{k/2}\det[\mathbf{P_i}]^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{x}-\mathbf{x}_{\mu,i})^\text{T}\mathbf{P_i}^{-1}(\mathbf{x}-\mathbf{x}_{\mu,i})\right). \tag{2}$$

Similarly, the probability density of a state $\mathbf{x}$ in a Gaussian mixture model is defined using the weighted sum of

a set of k multivariate normal density functions, as in Eq. (3). A multivariate random variable $\boldsymbol{\zeta}$ sampled from this mixture is defined in Eq. (4):

$$f(\mathbf{x}; \mathbf{x}_{\mu,1}, \mathbf{P}_1, \ldots, \mathbf{x}_{\mu,k}, \mathbf{P}_k) = \sum_{i=1}^{k} w_i f(\mathbf{x}; \mathbf{x}_{\mu,i}, \mathbf{P}_i) \tag{3}$$

$$\boldsymbol{\zeta} \sim f(\mathbf{x}_{\mu,1}, \mathbf{P}_1, \ldots, \mathbf{x}_{\mu,k}, \mathbf{P}_k) = \sum_{i=1}^{k} w_i \mathcal{N}(\mathbf{x}_{\mu,i}, \mathbf{P}_i) \tag{4}$$

where $w_1, \ldots, w_k$ are the weightings such that $w_i \geq 0$ and $\sum_{i=1}^{k} w_i = 1$.

D. Mahalanobis Distance with Gaussian Mixtures

The standard Mahalanobis distance and its Gaussian mixture analog can be derived from the general definition of the inner-product distance between two vectors $\mathbf{q}$ and $\mathbf{r}$ in k-dimensions:

$$d(\mathbf{q}, \mathbf{r})^2 = \|\mathbf{q} - \mathbf{r}\|^2 \tag{5}$$

$$= (\mathbf{q} - \mathbf{r})^{\mathrm{T}} \mathbf{A}^{-1} (\mathbf{q} - \mathbf{r}) \quad , \quad \mathbf{A} \in \mathbb{S}_{+}^{n \times n} \tag{6}$$

where the distance matrix $\mathbf{A}$ is positive-definite and symmetric. From observation, using the $n \times n$ identity matrix $\mathbf{I}_n$ as the distance matrix will yield the Euclidean distance. Likewise, using the covariance matrix $\mathbf{P}_i$ of a Gaussian distribution will yield the Mahalanobis distance. Li et al. define a new distance matrix to be used in the case of a Gaussian mixture using Kullback-Leiber divergence [21].

The Mahalanobis distance is a measure of the distance between a two points within a Gaussian distribution, scaled by the covariance of the distribution. The Mahalanobis distance thus calculates an n-dimensional standard deviation between the points. Often, one of the points is defined as the mean of the distribution, in which case Mahalanobis distance measures how many standard deviations a query is from the mean. The Mahalanobis distance of vector $\mathbf{x}$ from a distribution with mean $\mathbf{x}_\mu$ and covariance $\mathbf{P}$ is defined by Eq. (7).

$$d_M\left(\mathbf{x}, \mathbf{x}_\mu, \mathbf{P}\right) = \sqrt{\left(\mathbf{x} - \mathbf{x}_\mu\right)^{\mathrm{T}} \mathbf{P}^{-1} \left(\mathbf{x} - \mathbf{x}_\mu\right)} \tag{7}$$

A metric similar to Mahalanobis distance has been proposed by Hill et al. to identify outliers when comparing UCTs [23]. The particular metric shown in Eq. (8) compares the new PDF at time t_f, $f_B(\mathbf{x}) = \mathcal{N}(\mathbf{x}_{\mu,B}, \mathbf{P}_B)$ with a previous estimate propagated to time t_f, $f_A(\mathbf{x}) = \mathcal{N}(\mathbf{x}_{\mu,A}(t_f), \mathbf{P}_A(t_f))$. In the case of maneuver detection, $\mathbf{P}_A(t_f) + \mathbf{P}_B$ is the distance matrix used to account for the combined uncertainty of both distributions [23].

$$d_M\left(\mathbf{x}_{\mu,A}, \mathbf{x}_{\mu,B}, \mathbf{P}_A(t_f) + \mathbf{P}_B\right) = \sqrt{\left(\mathbf{x}_{\mu,A} - \mathbf{x}_{\mu,B}\right)^{\mathrm{T}} \left(\mathbf{P}_A(t_f) + \mathbf{P}_B\right)^{-1} \left(\mathbf{x}_{\mu,A} - \mathbf{x}_{\mu,B}\right)} \tag{8}$$

For the general case using Gaussian mixture boundary conditions, the Mahalanobis distance equation must be modified to account for multiple normal distributions. Minimizing the KL divergence for a Gaussian mixture yields the following distance matrix in Eq. (9) [21].

$$\mathbf{P}_\mu(\mathbf{x}) = \left[\frac{\sum_{i=1}^{n} w_i f(\mathbf{x}; \mathbf{x}_{\mu,i}, \mathbf{P}_i) \mathbf{P}_i^{-1}}{\sum_{j=1}^{n} w_j f(\mathbf{x}; \mathbf{x}_{\mu,j}, \mathbf{P}_j)} \right]^{-1} \tag{9}$$

From this result, $\mathbf{P}_\mu$ is a weighted reciprocal sum of the component covariances, where the weight is the probability of $\mathbf{x}$ occurring in distribution i compared to the entire the Gaussian mixture distribution.

In this work, the Gaussian mixture Mahalanobis distance will use Eq. (9) as the distance matrix. For a Gaussian mixture composed of k Gaussian distributions, the expected value, or centroid state, of the mixture is shown in Eq. (10).

$$\mathbf{x}_\mu = \mathbb{E}\left[\boldsymbol{\zeta}\right] = \sum_{i=1}^{k} w_i \mathbf{x}_{\mu,i} \tag{10}$$

The Mahalanobis distance analog for Gaussian mixtures is therefore defined in Eq. (11).

$$d_{GM}(\mathbf{x}, \mathbf{x}_\mu; \mathbf{P}_\mu^{-1}) = \sqrt{\left(\mathbf{x} - \mathbf{x}_\mu\right)^{\mathrm{T}} \mathbf{P}_\mu^{-1} \left(\mathbf{x} - \mathbf{x}_\mu\right)} \tag{11}$$

In the case when the Gaussian mixture only has a single component, the mixture is in fact a normal distribution, so Eqs. (9-10) reduce to the covariance and mean of the individual distribution respectively, and the standard Mahalanobis distance is recovered.

E. Binary Hypothesis Testing

A binary hypothesis structure uses separate PDFs for the null and alternative hypotheses to allow prescription of allowable error rates and detection thresholds. In order to apply binary hypothesis testing, a pair of mutually exclusive hypotheses, a null hypothesis $\mathcal{H}_0$ and an alternative hypothesis $\mathcal{H}_1$, must be developed along with corresponding probability density functions, $f_0(x)$ and $f_1(x)$. Given PDFs for the binary hypotheses, the following probabilities can be computed:

$$P_{FA} = \int_{\nu}^{\infty} f_0(x) \mathrm{d}x \tag{12}$$

$$P_{FN} = \int_{-\infty}^{\nu} f_1(x) \mathrm{d}x \tag{13}$$

$$P_D = \int_{\nu}^{\infty} f_1(x) \mathrm{d}x \tag{14}$$

where ν is a selected threshold, P_{FA} is the probability of false alarm (Type I error), P_{FN} is the probability of false negative or missed detection (Type II error), and P_D is the probability of detection. Noting that each integral shares

the same integration threshold ν, the probabilities in Eqs. (13-14) can be related by $P_{FN} + P_D = 1$. Also note that, assuming both PDFs have been normalized, Eq. (13) is equivalently defined as the cumulative distribution function (CDF) for the $\mathcal{H}_1$ hypothesis, while Eqs. (12) and (14) are the complementary CDFs, or tail distributions. Figure 2 notionally depicts the computation of these probabilities from the $\mathcal{H}_0$ and $\mathcal{H}_1$ PDFs.

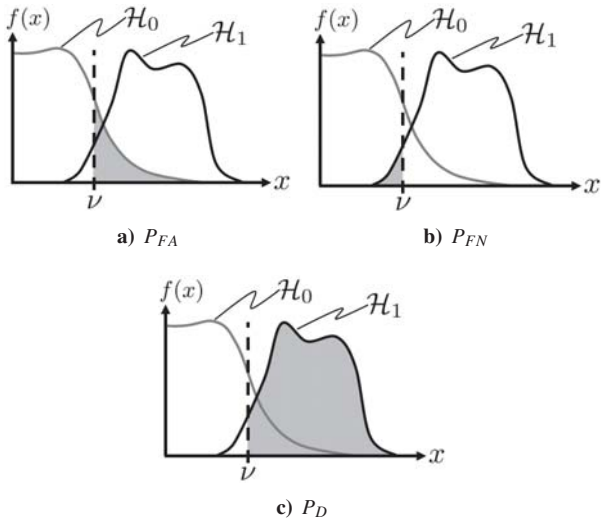

a) P_{FA} **b)** P_{FN}

c) P_D

Figure 2. One-sided Binary Hypothesis Testing Illustration

The threshold ν is typically selected in order to match a maximum allowable rate of false alarms, P_{FA}, using the null hypothesis PDF and Eq. (12).

III. Theory

The goal of this paper is to develop an algorithm using distance metrics for spacecraft anomaly detection. The control distance metric provides a method for computing the distance between two state distributions, generated by different spacecraft tracks. This is accomplished by linearizing about a nominal optimal trajectory between connecting the mean states of the PDFs. Due to the nonlinear nature of the general orbit determination problem, these PDFs can be non-Gaussian, particularly after propagation for some time. Therefore, the previous control distance metric work is extended using a Gaussian-mixture approximation for application to non-Gaussian boundary conditions. Using the extended control distance metric approach, control cost distributions are computed for the binary hypothesis pair of a non-maneuvered and a maneuvered trajectory. The control cost PDFs are used in a binary hypothesis test, subject to a desired false alarm rate, to determine the probability that an anomaly has occurred between the state estimates.

This study develops a framework to implement the control cost distribution method into an anomaly detection algorithm. The cost distribution approximation method by Holzinger et al. is extended for application to non-Gaussian boundary conditions in Section III.A, and binary hypothesis testing techniques are applied to the control cost PDFs in Section III.B. Similar treatments for Mahalanobis distance are developed and presented in Section III.C.

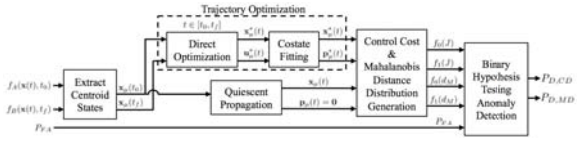

Figure 3. Minimum-fuel control distance anomaly detection and characterization framework

The anomaly detection algorithm is assembled as shown in Fig. 3. The inputs to the algorithm are a pair of UCTs, represented by PDFs at times t_0 and t_f. The output of the algorithm is the probability that the anomaly hypothesis is true , P_D, indicating whether an anomaly has been detected. Stated differently, this is the probability that something outside the modeled quiescent dynamics (e.g. a maneuver) has occurred, under the preliminary assumption that the two UCTs are generated by the same object. Using theory developed in the following sections, anomaly probability is calculated using two different methods, the control distance metric and Mahalanobis distance, yielding two anomaly probabilities for this study. This allows the analogous methods to be compared in both anomaly detection sensitivity and error rates.

A. Gaussian Mixture Approximation for Control Distance Metric

The derivation for the Gaussian mixture approximation for the control distance metric follows a similar derivation from Holzinger et al. that was restricted to Gaussian boundary conditions [17]. Using the optimal state and costate trajectory ($\mathbf{x}^*_\mu(t)$, $\mathbf{p}^*_\mu(t)$), state uncertainties are incorporated to generate an approximate probability distribution of the control cost associated with propagation between the boundary conditions. This is accomplished by linearizing about the nominal optimal trajectory and applying perturbations within the boundary condition distributions. Performing this many times generates a distribution of trajectories, and the control cost associated with the perturbed trajectories can be computed. In order to improve accuracy of boundary condition representation, the following method adapts the previous work on control distance metrics [17] for use with Gaussian mixtures. To approximate the quadratic control cost connecting Gaussian mixture boundary conditions, the trajectory connecting the centroid states of the initial and final boundary conditions is used. The centroid state, or expectation, of the initial boundary condition is computed as the weighted average of the mean states of each Gaussian component,

as shown in Eq. (15):

$$\mathbf{x}_\mu(t_0) = \mathbb{E}\left[\mathbf{x}(t_0)\right] = \mathbb{E}\left[\sum_{i=1}^{n} w_i \mathbf{x}_i(t_0)\right] = \sum_{i=1}^{n} w_i \mathbf{x}_{\mu,i}(t_0) \quad (15)$$

The centroid state of the final boundary condition can be computed similarly. The centroid states enable the trajectory to be reduced to a two-point boundary value problem, similar to the previous technique [18]. The control cost distribution is constructed by linearizing about this mean connecting trajectory and sampling initial and final states from the non-Gaussian boundary conditions, which introduces new perturbing terms. Figure 4 shows a notional depiction of the key variables introduced in the Gaussian mixture approximation. The optimal trajectory

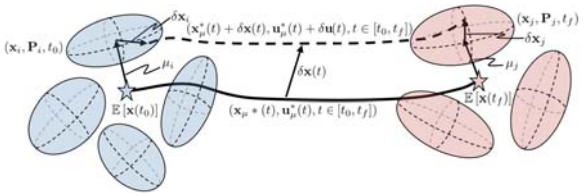

Figure 4. Gaussian mixture scenario geometry

between the centroid states is used as the nominal optimal trajectory $(\mathbf{x}^*_\mu(t), \mathbf{u}^*_\mu(t))$, computed using the trajectory optimization routine. The deviation of the mean state of an individual Gaussian component from the centroid state of its mixture is defined as $\boldsymbol{\mu}_i$ and $\boldsymbol{\mu}_j$ for the initial and final boundary conditions, respectively. This deviation term adds an extra perturbation to the control cost distribution.

The quadratic control cost function, in Eq. (16), can be expanded by decomposing the control effort $\mathbf{u}(t)$ into three components as shown in Eq. (17):

$$J = \frac{1}{2}\int \mathbf{u}(\tau)^{\mathrm{T}}\mathbf{u}(\tau)\mathrm{d}\tau \quad (16)$$

$$\mathbf{u}(t) = \mathbf{u}^*_\mu(t) + \delta\mathbf{u}_{ij}(t) + \delta\mathbf{u}(t) \quad (17)$$

where $\mathbf{u}^*_\mu(t)$ represents the optimal control associated with the mean trajectory $\mathbf{x}^*_\mu(t)$, $\delta\mathbf{u}_{ij}(t)$ represents the control perturbation due to the $\boldsymbol{\mu}_i$ and $\boldsymbol{\mu}_j$ variations in mean state of Gaussian initial component i and Gaussian final component j from the centroid states, and $\delta\mathbf{u}(t)$ represents the control perturbation due to $\delta\mathbf{x}_0$ and $\delta\mathbf{x}_f$, variations in the state sampled from boundary condition uncertainty. Similarly, from Holzinger et al. [17], the optimal control effort can be written as a function of the costates using Eq. (18):

$$\mathbf{u}(t) = -\frac{\partial \mathbf{f}}{\partial \mathbf{u}}^{\mathrm{T}}\left(\mathbf{p}^*_\mu(t) + \delta\mathbf{p}_{ij}(t) + \delta\mathbf{p}(t)\right) \quad (18)$$

where $\mathbf{p}^*_\mu(t)$ represents the optimal costate associated with the mean trajectory $\mathbf{x}^*_\mu(t)$, $\delta\mathbf{p}_{ij}(t)$ represents the costate perturbation due to $\boldsymbol{\mu}_i$ and $\boldsymbol{\mu}_j$, and $\delta\mathbf{p}(t)$ represents the costate perturbation due to $\delta\mathbf{x}_0$ and $\delta\mathbf{x}_f$. There exists a function $\boldsymbol{\Lambda}(t, t_0) \in \mathbb{R}^{n\times 2n}$, shown in Eq. (19), that maps variations in initial and final states to variations in the costate at time t, $\delta\mathbf{p}(t)$.

$$\begin{aligned}\boldsymbol{\Lambda}(t, t_0) = [&\boldsymbol{\Phi}_{px}(t, t_0) \\ &- \boldsymbol{\Phi}_{pp}(t, t_0)\boldsymbol{\Phi}_{xp}(t_f, t_0)^{\dagger}\boldsymbol{\Phi}_{xx}(t_f, t_0)\boldsymbol{\Phi}_{pp}(t, t_0)\boldsymbol{\Phi}_{xp}(t_f, t_0)^{\dagger}]\end{aligned} \quad (19)$$

Note that this function is composed of portions of the state transition matrix partitioned as

$$\begin{bmatrix}\delta\mathbf{x}(t)\\ \delta\mathbf{p}(t)\end{bmatrix} = \boldsymbol{\Phi}(t, t_0)\begin{bmatrix}\delta\mathbf{x}(t_0)\\ \delta\mathbf{p}(t_0)\end{bmatrix} \quad (20)$$

$$= \begin{bmatrix}\boldsymbol{\Phi}_{xx}(t, t_0) & \boldsymbol{\Phi}_{xp}(t, t_0)\\ \boldsymbol{\Phi}_{px}(t, t_0) & \boldsymbol{\Phi}_{pp}(t, t_0)\end{bmatrix}\begin{bmatrix}\delta\mathbf{x}(t_0)\\ \delta\mathbf{p}(t_0)\end{bmatrix} \quad (21)$$

where $\boldsymbol{\Phi}(t, t_0)$ is the state transition matrix mapping variations $\delta\mathbf{x}$ and $\delta\mathbf{p}$ to time t about the optimal trajectory. Also note that, while the pseudoinverse term $\boldsymbol{\Phi}_{xp}(t_f, t_0)^{\dagger}$ is not guaranteed to exist for arbitrary systems, its existence implies controllability through the optimal control problem. This portion of the state transition matrix determines how variations in the costates affect the state, or in other words whether the state is controllable. In this paper, controllability is assumed, so for present purposes the pseudoinverse is also assumed to exist [17].

Using $\boldsymbol{\Lambda}(t, t_0)$, the components $\delta\mathbf{u}_{ij}(t)$ and $\delta\mathbf{u}(t)$ can be computed as shown in Eqns. (22) and (23).

$$\delta\mathbf{u}_{ij}(t) = -\frac{\partial \mathbf{f}}{\partial \mathbf{u}}^{\mathrm{T}}\boldsymbol{\Lambda}(t, t_0)\begin{bmatrix}\boldsymbol{\mu}_i\\ \boldsymbol{\mu}_j\end{bmatrix} \quad (22)$$

$$\delta\mathbf{u}(t) = -\frac{\partial \mathbf{f}}{\partial \mathbf{u}}^{\mathrm{T}}\boldsymbol{\Lambda}(t, t_0)\begin{bmatrix}\delta\mathbf{x}_i\\ \delta\mathbf{x}_j\end{bmatrix} \quad (23)$$

The following terms are defined for ease of notation:

$$\boldsymbol{\mu}_{ij} = \begin{bmatrix}\boldsymbol{\mu}_i\\ \boldsymbol{\mu}_j\end{bmatrix} \quad (24)$$

$$\mathbf{P}_{ij} = \begin{bmatrix}\mathbf{P}_i & \mathbf{0}\\ \mathbf{0} & \mathbf{P}_j\end{bmatrix} \quad (25)$$

$$\delta\mathbf{z}_{ij} = \begin{bmatrix}\delta\mathbf{x}_i\\ \delta\mathbf{x}_j\end{bmatrix} \quad (26)$$

Note that $\boldsymbol{\mu}_{ij} \in \mathbb{R}^{12\times 1}$ is a constant vector for each (i, j) boundary condition pair. Similarly, $\mathbf{P}_{ij} \in \mathbb{R}^{12\times 12}$ is a constant matrix for each (i, j) boundary condition pair. The zero-mean random variable $\delta\mathbf{z}_{ij}$ is sampled from the i and j boundary condition uncertainties such that $\delta\mathbf{x}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{P}_i)$ and $\delta\mathbf{x}_j \sim \mathcal{N}(\mathbf{0}, \mathbf{P}_j)$.

Since $\delta\mathbf{z}_{ij}$ is independent of time τ, the approximate quadratic control cost for a single term of the Gaussian mixture connecting initial distribution i to final distribution j can be expressed as seen in Eq. (27) by substituting the definitions in Eqs. (28-30) into Eq. (16):

$$\begin{aligned}J_{ij} \approx J^* &+ \boldsymbol{\omega}(t_f, t_0)^{\mathrm{T}}\delta\mathbf{z}_{ij} + 2\boldsymbol{\mu}_{ij}^{\mathrm{T}}\boldsymbol{\Omega}(t_f, t_0)\delta\mathbf{z}_{ij} + \boldsymbol{\omega}(t_f, t_0)^{\mathrm{T}}\boldsymbol{\mu}_{ij}\\ &+ \boldsymbol{\mu}_{ij}^{\mathrm{T}}\boldsymbol{\Omega}(t_f, t_0)\boldsymbol{\mu}_{ij} + \delta z^{\mathrm{T}}\boldsymbol{\Omega}(t_f, t_0)\delta\mathbf{z}_{ij}\end{aligned} \quad (27)$$

$$J^* = \frac{1}{2}\int_{t_0}^{t_f} \mathbf{u}^*(\tau)^{\mathrm{T}}\mathbf{u}^*(\tau)\mathrm{d}\tau \tag{28}$$

$$\boldsymbol{\omega}(t_f, t_0) = \int_{t_0}^{t_f} \boldsymbol{\Lambda}(\tau, 0)^{\mathrm{T}} \frac{\partial \mathbf{f}}{\partial \mathbf{u}} \mathbf{u}^*(\tau)\mathrm{d}\tau \tag{29}$$

$$\boldsymbol{\Omega}(t_f, t_0) = \frac{1}{2}\int_{t_0}^{t_f} \boldsymbol{\Lambda}(\tau, 0)^{\mathrm{T}} \frac{\partial \mathbf{f}}{\partial \mathbf{u}} \frac{\partial \mathbf{f}}{\partial \mathbf{u}}^{\mathrm{T}} \boldsymbol{\Lambda}(\tau, 0)\mathrm{d}\tau \tag{30}$$

J^* is the quadratic control cost of the optimal trajectory $(\mathbf{x}^*_\mu(t), \mathbf{p}^*_\mu(t))$ without boundary-condition variations. The terms $\boldsymbol{\omega}(t, t_0)$ and $\boldsymbol{\Omega}(t, t_0)$ are defined relative to the optimal trajectory $(\mathbf{x}^*_\mu(t), \mathbf{u}^*_\mu(t))$ connecting the centroid states. Combining terms, the cost J_{ij} in Eq. (27) can be re-written in a format similar to the strictly Gaussian result from Holzinger et al. [17], as shown in Eq. (31).

$$\begin{aligned} J_{ij} &= J^* + \delta J_{ij} \qquad (31)\\ &\approx J^* + \boldsymbol{\omega}(t_f, t_0)^{\mathrm{T}}(\boldsymbol{\mu}_{ij} + \delta\mathbf{z}_{ij}) + \\ &\quad (\boldsymbol{\mu}_{ij} + \delta\mathbf{z}_{ij})^{\mathrm{T}}\boldsymbol{\Omega}(t_f, t_0)(\boldsymbol{\mu}_{ij} + \delta\mathbf{z}_{ij}) \qquad (32)\end{aligned}$$

where δJ_{ij} is the variational control cost due to uncertainties in the i^{th} initial and j^{th} final terms of the boundary conditions. From Holzinger et al. [17] Appendix B, the analytic first and second moments of the variational control cost, δJ_{ij}, are:

$$\mathbb{E}\left[\delta J_{ij}\right] = \mu_{J,ij} \approx \boldsymbol{\omega}^{\mathrm{T}}\boldsymbol{\mu}_{ij} + \boldsymbol{\mu}_{ij}^{\mathrm{T}}\boldsymbol{\Omega}_{ij}\boldsymbol{\mu}_{ij} + \mathrm{Tr}\left[\boldsymbol{\Omega}_{ij}\mathbf{P}_{ij}\right] \tag{33}$$

$$\begin{aligned}\mathbb{E}\left[\delta J_{ij}^2\right] = \sigma^2_{J,ij} &\approx \left(\boldsymbol{\omega}^{\mathrm{T}} + 2\boldsymbol{\mu}_{ij}^{\mathrm{T}}\boldsymbol{\Omega}\right)^{\mathrm{T}}\mathbf{P}_{ij}\left(\boldsymbol{\omega}^{\mathrm{T}} + 2\boldsymbol{\mu}_{ij}^{\mathrm{T}}\boldsymbol{\Omega}\right)\\ &\quad + 2\mathrm{Tr}\left[\boldsymbol{\Omega}\mathbf{P}_{ij}\boldsymbol{\Omega}\mathbf{P}_{ij}\right] \qquad (34)\end{aligned}$$

The control cost considering all i initial boundary conditions and j final boundary conditions is then the weighted sum of the individual costs between each i and j:

$$J \approx J^* + \sum_{i=1}^{N_0}\sum_{j=1}^{N_f} w_i w_j \delta J_{ij} \tag{35}$$

where w_i is the weight of the i^{th} initial boundary condition and w_j is the weight of the j^{th} final boundary condition. Thus, the analytic expected value of the initial and final Gaussian sum boundary condition may be written as:

$$\begin{aligned}\mathbb{E}\left[J\right] &\approx \mathbb{E}\left[J^* + \sum_{i}^{N_0}\sum_{j}^{N_f} w_i w_j \delta J_{ij}\right]\\ &\approx J^* + \sum_{i}^{N_0}\sum_{j}^{N_f} w_i w_j \mathbb{E}\left[\delta J_{ij}\right]\\ &\approx J^* + \sum_{i}^{N_0}\sum_{j}^{N_f} w_i w_j \mu_{J,ij} \qquad (36)\end{aligned}$$

which leads to the expected value of the total control cost distribution in Eq. (37):

$$\begin{aligned}\mu_J &\approx J^*\\ &\quad + \sum_{i}^{N_0}\sum_{j}^{N_f} w_i w_j \left(\boldsymbol{\omega}^{\mathrm{T}}\boldsymbol{\mu}_{ij} + \boldsymbol{\mu}_{ij}^{\mathrm{T}}\boldsymbol{\Omega}\boldsymbol{\mu}_{ij} + \mathrm{Tr}\left[\boldsymbol{\Omega}\mathbf{P}_{ij}\right]\right) \qquad (37)\end{aligned}$$

This yields the approximate expected value of the control cost distribution connecting two Gaussian-mixture boundary conditions, and therefore can be used to approximate control cost distributions for non-Gaussian boundary conditions. Note that, in the case where the boundary conditions each consist of a single Gaussian component, the summation and weighting terms drop out and $\boldsymbol{\mu}_{ij} = 0$, recovering the Gaussian expressions for control cost distributions from Holzinger et al. [17].

Note that the approximations shown here assume that the Gaussian components are within the region of convergence of the map $\boldsymbol{\Lambda}(t_0, t_f)$. This assumption is more likely to be valid when the Gaussian components of the mixture are more tightly packed, as is the case after propagation of an initially Gaussian state estimate. However, if the Gaussian mixture components are more separate, this assumption may no longer be valid. Results in Section III.B show a decidedly non-Gaussian case with accurate control cost reproduction, but this is not guaranteed. Preliminary work [24] used an intermediate method where optimal trajectories were computed between each of the Gaussian components, instead of linearizing about the single optimal trajectory and applying the perturbing cost due to $\boldsymbol{\mu}_\mathbf{i}$ and $\boldsymbol{\mu}_\mathbf{j}$. This introduces significant computational complexity over the approximations in this paper due to the combinatorial nature of this approach, requiring $N_0 \times N_f$ optimal trajectories, but may be preferred based on the boundary conditions provided. Cost distributions developed using a combinatorial method are still valid for use in the remaining theory.

B. Anomaly Detection using Binary Hypothesis Testing

The existing implementation of control cost maneuver detection forms the anomaly hypothesis using only the control cost distribution from the maneuvered trajectory, comparing the deterministic optimal control cost to the uncertainty in the maneuvered control cost distribution. This approach computes a probability that the optimal control cost was detectable over the noise in the measurements, or the probability of a maneuver [25]. However, this method does not allow specification of acceptable error rates, making comparisons to existing distance metrics, such as Mahalanobis distance, incomplete. While the cost distribution for the maneuvering (alternative) hypothesis is well defined from the previous method, an opposing distribution for the null hypothesis was not developed.

This study modifies the anomaly detection method by applying binary hypothesis testing. For the anomaly detection problem, the binary hypotheses are formulated as follows:

$\mathcal{H}_0$ (Null Hypothesis): Observed trajectory adequately explained by quiescent state propagation with boundary condition uncertainty.

$\mathcal{H}_1$ (Alternative Hypothesis): Observed trajectory not adequately explained by quiescent state propagation with boundary condition uncertainty.

Recalling the scenario from Fig. 1, the binary distributions for the null and alternate hypotheses are drawn from the maneuvering and non-maneuvering trajectories, respectively.

In the event that the null hypothesis is accepted, the change in state between observations is small enough that it more likely explained solely by uncertainty in the boundary conditions. Therefore, the associated control cost PDF $f_0(J)$ is derived from a quiescent propagated trajectory:

$$f_0(J) \approx \sum_{i=1}^{N_0} \sum_{j=1}^{N_0} w_i w_j \mathcal{N}\left(\mu_{J,ij}, \sigma^2_{J,ij}\right) \tag{38}$$

The Gaussian mixture approximation for control cost distribution is applied along the quiescent trajectory connecting the initial UCT boundary condition and its propagated counterpart at the later time, such that each initial Gaussian component $i = 1, ..., N_0$ has a corresponding final Gaussian component $j = 1, ..., N_0$.

Alternately, if the null hypothesis is rejected, then the change in state between observations is too large to be adequately explained solely by uncertainty in the boundary conditions. The associated control cost PDF $f_1(J)$ is derived from the maneuvering trajectory:

$$f_1(J) \approx \sum_{i=1}^{N_0} \sum_{j=1}^{N_f} w_i w_j \mathcal{N}\left(\mu_{J,ij}, \sigma^2_{J,ij}\right) \tag{39}$$

Note that here the Gaussian mixture approximation for control cost distribution is applied along the optimal trajectory connecting the two input UCT boundary conditions, so the final condition Gaussian components j have changed between Eqns. (38) and (39).

This formulation was selected by analyzing a number of different theoretical scenarios to ensure the entire space of possible binary hypotheses revealed the desired behavior, as illustrated in Fig. 5. For instance, in the case of a small maneuver, the null and alternative hypothesis PDFs will mostly overlap. Since the allowable rate of false alarm, P_{FA}, is likely to be small, the corresponding threshold ν will cause the anomaly probability P_D from Eq. (14) to be small as well. However, in the event that a large maneuver has occurred, the alternative hypothesis PDF will be shifted far to the right of the null hypothesis PDF. Using the same P_{FA} and corresponding ν, the anomaly probability P_D is large, indicating that a maneuver has likely occurred.

The following approach applies binary hypothesis testing to spacecraft anomaly detection using control distance PDFs and prescribed acceptable false alarm and missed detection rates:

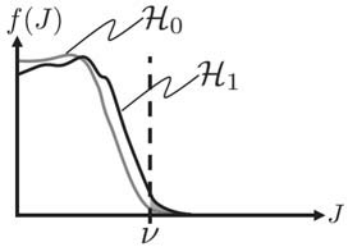

a) Significant overlap, low P_D

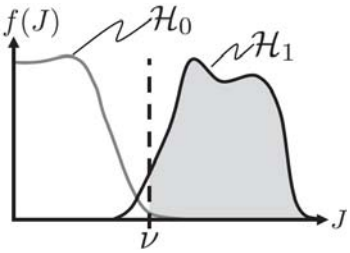

b) Small overlap, high P_D

Figure 5. Theoretical scenarios for control cost binary hypotheses

1) Construct control cost PDFs for binary hypotheses, $f_0(J)$ and $f_1(J)$, using Eqs. (38) and (39).

2) Using allowable false alarm rate P_{FA}, compute integration threshold cost ν using Eq. (12).

3) Compute anomaly probability P_D using Eq. (14) using $f_1(J)$ and ν from previous step.

Once P_D is computed, a final thresholding process can be used to determine whether to flag as an anomaly.

4a) If threshold detection probability for anomaly is directly prescribed, $P_{D,thresh}$ is given.

4b) If allowable false negative (missed detection) rate prescribed, $P_{D,thresh} = 1 - P_{FN}$.

5) If $P_D >= P_{D,thresh}$, flag as anomaly.

C. Binary Hypotheses for Mahalanobis Distance

Since Mahalanobis distance has been proposed as a potential data association and anomaly detection metric, an analogous formulation for binary hypothesis testing is also developed using Mahalanobis distance. The binary hypotheses are reinterpreted using the Mahalanobis distance formulation from Eq. (11).

The null hypothesis considers the state distribution $\mathbf{f}_A(\mathbf{x}(t_f))$, the initial distribution propagated quiescently to time t_f. Given a point sampled from distribution $\mathbf{f}_A$ at t_f, namely $\mathbf{x}_A(t_f) \sim \mathbf{f}_A(\mathbf{x}(t_f))$, Eq. (40) computes the Mahalanobis distance from this point to the quiescent distribu-

tion $\mathbf{f}_A$.

$$d_{M|\mathcal{H}_0} = \sqrt{\left(\mathbf{x}_A(t_f) - \mathbb{E}\left[\mathbf{x}_A(t_f)\right]\right)^{\mathrm{T}} \mathbf{P}_A^{-1} \left(\mathbf{x}_A(t_f) - \mathbb{E}\left[\mathbf{x}_A(t_f)\right]\right)} \tag{40}$$

The alternative hypothesis considers the state distribution $\mathbf{f}_B(\mathbf{x}(t_f))$ from the new observation at time t_f. Given a point sampled from distribution $\mathbf{f}_B$, namely $\mathbf{x}_B(t_f) \sim \mathbf{f}_B(\mathbf{x}(t_f))$, Eq. (41) computes the Mahalanobis distance from this point to the quiescent distribution $\mathbf{f}_A$.

$$d_{M|H_1} = \sqrt{\left(\mathbf{x}_B(t_f) - \mathbb{E}\left[\mathbf{x}_A(t_f)\right]\right)^{\mathrm{T}} \left(\mathbf{P}_A(t_f) + \mathbf{P}_B\right)^{-1} \left(\mathbf{x}_B(t_f) - \mathbb{E}\left[\mathbf{x}_A(t_f)\right]\right)} \tag{41}$$

The use of $(\mathbf{P}_A(t_f) + \mathbf{P}_B)$ as the distance matrix follows the convention set by Hill et al. per Eq. (8) to describe the distance between a maneuvered distribution and the quiescent propagated distribution [23].

Similar to the control cost PDFs, this interpretation using Mahalanobis distance is justified by considering the theoretical cases of quiescent propagated state distribution $\mathbf{f}_A(\mathbf{x}(t_f))$ and new state distribution $\mathbf{f}_B(\mathbf{x}(t_f))$ as shown in Fig. 6. Cases I and III refer to a situation where uncertainty in the initial spacecraft state estimate is high, but uncertainty in the new estimate is reduced. Cases II and IV refer to a situation where uncertainty in the initial estimate is low, but uncertainty in the new estimate is significantly larger. In case I, the propagated distribution envelops the new state distribution, so the Mahalanobis distance PDFs show significant overlap, yielding a low probability of anomaly. In case II, the new state distribution envelops the propagated distribution, so there is once again significant overlap. The mean of the distributions are offset and uncertainty in the new observation is larger, leading to a non-negligible anomaly probability and indicating that an anomaly likely occurred between the observations to cause the change. In both cases III and IV, the state distributions no longer overlap, causing the Mahalanobis distance distributions to be further separate as well, leading to high anomaly probabilities.

In each case, this binary hypothesis Mahalanobis distance formulation effectively determines the anomaly probability as designed, providing an analogous formulation for comparison to the control distance metric.

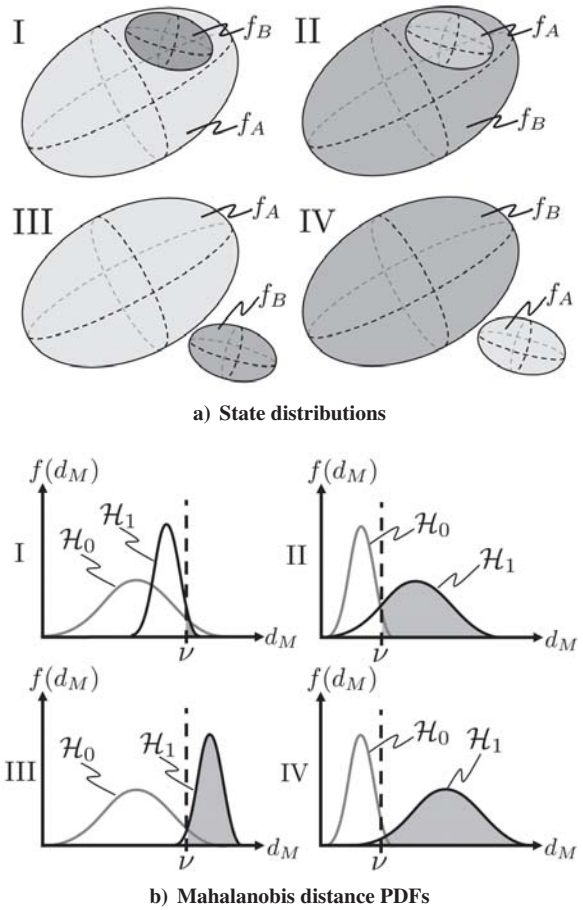

a) State distributions

b) Mahalanobis distance PDFs

Figure 6. Theoretical scenarios for Mahalanobis distance binary hypotheses

IV. Simulation Results

This section begins with a discussion of the algorithm implementation, followed by a demonstration and analysis of the performance of the algorithm. Sensitivity studies are also performed to assess the applicability of the methods to typical spacecraft maneuver classes.

A. Implementation Details

The algorithm as described in Fig. 3 is implemented in MATLAB to evaluate its performance and effectiveness.

Specifics for each portion of this block diagram are explained below.

1. Trajectory Optimization

Using the input UCT state PDFs, the deterministic two-point boundary value problem between the expected-value states is formulated into an optimization problem, discretizing the simulation into a user-defined number of time-steps, which is solved using the constrained minimization function in MATLAB, *fmincon()*. The decision variable for this minimization is a stacked vector of the thrust accelerations at each discrete time instant, and the thrust accelerations are held constant for each discrete time step. Keplerian dynamics, along with a number of user-selectable perturbation accelerations (J_2, J_{22}, J_3, lunar gravitational, and solar gravitational perturbations), are enforced between steps of the trajectory as equality constraints to ensure the generated trajectory dynamics are accurate. Since the partial derivatives of the dynamics with respect to the decision variables (thrust accelerations) are well known, the gradient of the constraint is supplied to the optimization function to improve convergence. The output of the direct optimization step is a nominal optimal direct trajectory of states and controls connecting the UCTs, which is refit to an optimal indirect trajectory of states and costates. The generated optimal trajectory is validated using the nonlinear dynamics to numerically integrate the proposed control vector and quantify the error between the integrated final condition and the specified final UCT boundary condition. In addition, the quiescent trajectory from the initial condition is computed under the same dynamics with the assumption that no control input is used.

2. PDF Generation

To accurately construct the control cost PDFs for Gaussian-mixture boundary conditions, a localized Monte-Carlo-like method is employed to sample from the boundary conditions and apply the Gaussian-mixture control cost approximation. This is done by selecting one of the boundary condition Gaussian components from the mixture randomly with a probability based on the weighting for that particular Gaussian component. Once the Gaussian components are selected for the initial and final states, the deviation $\boldsymbol{\mu}_{ij}$ and covariance $\mathbf{P}_{ij}$ terms are known. The zero-mean random variable $\delta\mathbf{z}_{ij}$ is selected using the covariance information for the chosen Gaussian components. With all the required terms gathered, the quadratic cost is computed using Eq. (16). When generating the numerical PDFs, if a negative-valued cost is generated, that sample is discarded and another sample is drawn. The number of samples required to construct this distribution depends on the scenario. In this study, the control cost distribution is sampled $n_{samp} = 10000$ times, an amount shown in previous studies to be able to reconstruct the non-linearized control cost distributions in a similar scenario [24], and the samples are used to construct the approximate control cost PDF. This process is performed using both the quiescent and maneuvering trajectories to form both null hypothesis $f_0(J)$ and alternative hypothesis $f_1(J)$ PDFs. This process could equivalently be performed analytically by forming a Gaussian-mixture using $\mu_{J,ij}$, and $\sigma_{J,ij}$ from Eqs. (10) and (34) along with w_i and w_j for each boundary condition. The localized Monte-Carlo-like approach is chosen for simplicity of implementation. Mahalanobis distance PDFs are constructed in a similar manner. Samples from each boundary condition are used in conjunction with Eqs. (40-41) to develop PDFs for each hypothesis.

3. Anomaly Detection

Once PDFs are obtained for each hypothesis, anomaly detection is performed the same for both control distance and Mahalanobis distance. CDFs for both null and alternative hypotheses are computed, noting that the lower limit of integration becomes 0 since the distance cannot be negative.

$$F_0(x) = \int_0^x f_0(y)\mathrm{d}y \tag{42}$$

$$F_1(x) = \int_0^x f_1(y)\mathrm{d}y \tag{43}$$

Applying Eq. (12), the input allowable false alarm rate P_{FA} is used to compute the threshold ν by interpolating on the null hypothesis CDF, F_0. Using the threshold ν and applying Eq. (14), the probability of anomaly is computed by interpolating on the alternative hypothesis CDF, F_1.

B. Gaussian Mixture Approximation Validation

The Gaussian mixture approximation promises a more computationally tractable method for addressing non-Gaussian boundary conditions; however, it still must provide an accurate reconstruction of the uncertainty cost distribution. To validate the Gaussian mixture approximation of control cost distributions, a synthetic scenario is constructed for a GEO spacecraft with a non-Gaussian boundary condition state distribution, represented using Gaussian mixtures. The orbital elements for the boundary conditions are listed in Tables 1-2. This particular scenario is selected to illustrate the ability to generate and handle non-Gaussian control cost distributions using the binary hypothesis testing approach outlined above. The initial condition occurs 30 minutes before the ascending node passage, and the 3 Gaussian terms vary only in inclination. The final condition occurs 30 minutes after the ascending node passage, with 2 Gaussian terms varying only in inclination. Note that the asymmetry in the Gaussian mixtures ($N_0 = 3, N_f = 2$) is entirely allowed by

the formulation in Section III.A. These particular boundary conditions represent a 1 hour observation gap wherein the observed spacecraft has performed a small inclination correction. The $1-\sigma$ boundary condition uncertainties are initialized at 10 meters in position and 10 centimeters-per-second in velocity. These values are selected to generate multimodal cost distributions when scaled by α to emphasize the generality of the analytical contributions.

Table 1. Boundary condition orbital elements for Gaussian mixture inclination change

Parameter	Initial Condition $i=1$	$i=2$	$i=3$
Weighting, w	0.34	0.33	0.33
Semi-major Axis, a (km)	42164	42164	42164
Eccentricity, e	0	0	0
Inclination, i (deg)	0.015	0.025	0.035
Long. of Asc. Node, Ω (deg)	0	0	0
Arg. of Periapsis, ω (deg)	0	0	0
True Anomaly, θ (deg)	352.5	352.5	352.5

Table 2. Initial boundary condition orbital elements for Gaussian mixture inclination change

Parameter	Final Condition $j=1$	$j=2$
Weighting, w	0.5	0.5
Semi-major Axis, a (km)	42164	42164
Eccentricity, e	0	0
Inclination, i (deg)	0.00	0.01
Long. of Asc. Node, Ω (deg)	0	0
Arg. of Periapsis, ω (deg)	0	0
True Anomaly, θ (deg)	7.5	7.5

For validation purposes, an alternate method is used to generate control cost distributions through direct sampling of the boundary conditions and non-linearized trajectory optimization. For each run, a state is selected at random from the initial and final Gaussian mixture boundary conditions. The optimal control between the chosen boundary conditions is computed, and the resulting control cost is collected before selecting another pair of boundary conditions. This process is repeated $n_{samp} = 5000$ times to sample the control cost space, binning the results to construct a control cost PDF. Since this method does not make the simplifying assumption of linearizing about a best-estimate trajectory, it more accurately generates the actual control cost distribution between the Gaussian mixture boundary conditions, at the expense of much longer computation times.

Figure 7 shows the normalized PDF and associated CDF resulting from both these methods. The solid lines indicate the non-linearized propagation, labeled "True," and the dashed lines indicate the control cost distribution generated using the Gaussian mixture approximation method, labeled "Aprx." Inspection of the normalized PDFs shows significantly non-Gaussian control cost distributions, as expected given the separation in the boundary condition components: transitioning from initial component $i = 1$ to final component $j = 1$ uses significantly less fuel than from initial component $i = 3$ to final component $j = 1$. Despite the differences in approach, the control cost PDFs and CDFs agree well between the non-linearized and approximated methods, demonstrating the efficacy of contribution 1: the Gaussian mixture control cost approximation.

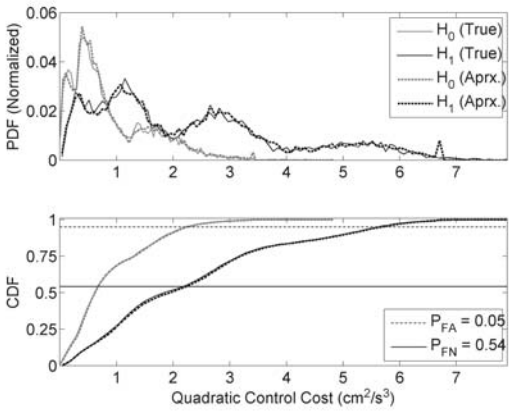

Figure 7. Comparison of control cost distributions using the non-linearized and approximate methods

Since the CDFs are nearly identical, the binary hypothesis testing algorithm yields a nearly identical probability of anomaly using either method. To further demonstrate the applicability of contribution 2, the binary hypothesis testing algorithm, a sample calculation of anomaly probability is illustrated using Fig. 7. Assuming an allowable false alarm rate of 5%, or $P_{FA} = 0.05$, the threshold ν is calculated from Eq. (12) by interpolating on the $\mathcal{H}_0$ CDF to find the cost where $F_0(J) = 1 - P_{FA}$. Given the threshold cost, the false negative probability is calculated using Eq. (13) by interpolating on the $\mathcal{H}_1$ CDF, such that here $P_{FN} = F_1(\nu) = 0.54$. Recalling that $P_D = 1 - P_{FN}$, the probability of anomaly for this scenario is therefore $P_D = 0.46$, or 46%.

Table 3. Timing comparison for validation scenario

Method	Complexity	Actual Time (s)
Non-Linearized	$O(n_{samp})$	53450
GM Approximation	$O(c)$	2.835

Given the similarity of the PDFs, the computational complexity savings of the Gaussian mixture approximation method offers a significant benefit inherent in the first contribution. Table 3 shows the time required to perform the full anomaly detection algorithm for both of these scenarios. These results are in-line with expectation, as the trajectory optimizer is expected to be the computational bottleneck. The non-linearized validation method runs

in roughly $O(n_{samp})$ time since a trajectory optimization is required for each sample (with small variations in the time required for each sample based on the boundary conditions). In comparison, the Gaussian-mixture approximation method only requires a single trajectory optimization, and therefore runs in nearly constant time, or $O(c)$. Therefore, when considering the computational complexity improvements, the Gaussian mixture approximation method significantly outperforms the non-linearized approach while maintaining accuracy in the control cost PDF and therefore performing similarly in anomaly detection.

In order to test the efficacy of the Mahalanobis distance Gaussian-mixture method proposed, the same scenario is used to develop PDFs and CDFs for the binary hypotheses using Mahalanobis distance. In comparison to the control distance method, the Mahalanobis distance PDFs from the individual Gaussian components are more distinct, so the combined PDF forms narrow peaks in an extremely non-Gaussian manner as seen in Fig. 8. In particular, the null hypothesis distribution covers a wide range of Mahalanobis distance values, since the quiescent propagated distributions are very dissimilar. The wide null hypothesis distribution causes the distance threshold for 95% anomaly confidence to be large, yielding an effectively 0% anomaly detection, similar to Case I in Fig. 6.

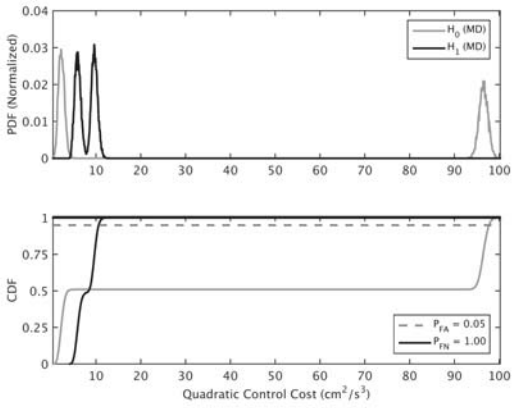

Figure 8. Mahalanobis distance distributions for non-Gaussian boundary conditions

Additionally, the Mahalanobis distance Gaussian mixture formulation is significantly more computationally expensive due to the complicated combined covariance calculation in Mahalanobis distance PDF generation. Without this calculation, in the purely Gaussian case, Mahalanobis distance is computed quicker than control distance, since Mahalanobis distance does not require trajectory optimization. Adding in the combined covariance calculation for non-Gaussian boundary conditions significantly increases computation time over the control distance method, as shown in Table 4.

Using the formulations developed in the Theory section, these results demonstrate the ability to take non-Gaussian boundary conditions and form PDFs and CDFs for use in binary hypothesis testing. The control distance method proved to be better in the detection of this particular maneuver, but the following simulation results will delve deeper into performance of each algorithm in different maneuver detection scenarios.

Table 4. Timing comparison by algorithm segment

Algorithm Segment	CD (s)	MD (s)
Trajectory Optimization	1.845	0
PDF Generation	0.762	11.347
Anomaly Testing	0.229	0.225
Total	2.835	11.572

C. Synthetic Data - Inclination Change

Having shown the ability to accurately construct control cost PDFs for non-Gaussian boundary conditions, the anomaly detection algorithm is next evaluated by parameterizing the problem to assess sensitivity. For the remaining results in this paper, the boundary conditions are simplified to single Gaussian distributions for ease of parameterization and discussion, but similar simulations could be performed using Gaussian mixtures as shown above.

A simulated scenario is constructed to emulate an inclination change performed at GEO, termed a North-South station-keeping maneuver in operations. Typically, satellites at GEO will be placed into orbit slots and given allowable deviations in the North-South and East-West directions. This particular scenario employs a 0.02 degree inclination change, similar in magnitude to that observed in the available real-world data (Section V.A). The goal of this scenario is to analyze the sensitivity of both the control distance and Mahalanobis distance binary hypothesis testing methods to changes in observation gap, uncertainty, and false alarm rate. Each of these three parameters are varied systematically in simulation for evaluation.

The observation gap is varied between 10 minutes and 48 hours. Each element of the 6×6 covariance matrix is varied by an uncertainty scaling parameter α. The $1-\sigma$ boundary condition position uncertainty is initialized at 1 meter and varied up to 250 meters using the scaling parameter α. Likewise, the $1-\sigma$ boundary condition velocity uncertainty is also varied between 1 centimeter-per-second and 250 centimeters-per-second using the scaling parameter α. The prescribed false alarm rate is varied between 0.5% and 10%, or $P_{FA} = 0.005$ to 0.10. For each combination of observation gap, boundary condition uncertainty, and prescribed false alarm rate, the anomaly probability is computed using both the control distance and Mahalanobis distance metrics. This produces a 4-dimensional dataset, $(P_D : \delta t, \alpha, P_{FA})$, which is best viewed as a set of contour plots using slices of constant P_{FA}.

Figure 9 presents a contour plot for a subset of the control distance sensitivity study data using a false alarm probability $P_{FA} = 0.05$. Note that, though the uncertainty scaling parameter α is varied from $1 \leq \alpha \leq 250$, a subset of this range ($\alpha >= 10$) has been plotted to highlight trends. Toward the left of the uncertainty axis, boundary condition uncertainty is low, so it is easier to distinguish between the non-maneuvered and maneuvered trajectories; therefore, anomaly probability is high ($P_D \approx 1$). Increasing boundary condition uncertainty causes more overlap in the control distance distributions, introducing values of $P_D < 1$ as the propagated uncertainty is large enough to account for the new observation. Additionally, a slight spike in anomaly probability can be seen around half-orbit period observation gaps, which coincides with the furthest out-of-plane difference between the quiescent and maneuvered trajectories. At these points, boundary condition uncertainty must be significantly greater to cause overlap between the trajectories, showing increased sensitivity to the maneuver at that condition.

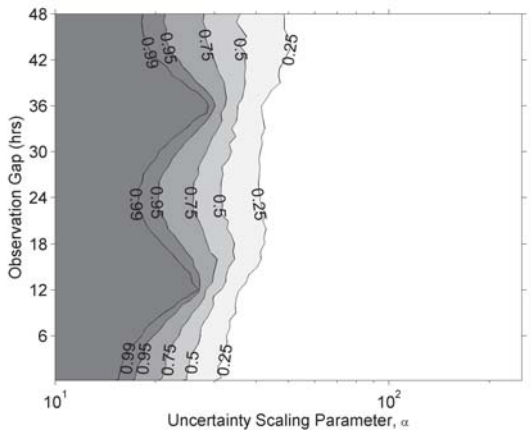

Figure 9. Control distance probability of anomaly contours vs uncertainty scaling parameter and observation gap, simulated inclination change maneuver, $P_{FA} = 0.05$.

Additional contour plots for different false alarm probabilities can be seen in Fig. 16 of Appendix A. Each subfigure shows a contour plot of anomaly probability for a prescribed false alam rate, where the darker colors indicate a higher probability of anomaly. Here, false alarm rates of 0.5%, 1%, 5%, and 10% were used. As expected, the algorithm declares higher probabilities of anomaly for a fixed observation gap and uncertainty as the allowable rate of false alarm increases because the threshold for maneuver detection is lessened. The remainder of the trends are consistent regardless of false alarm probability, so they are relegated to Appendix A for reference.

Figure 10 presents the same anomaly probability data using Mahalanobis distance distributions, with a more complete set of plots featured in Fig. 17 of Appendix A. Similar trends observed with control distance can also be noted for Mahalanobis distance: increasing the rate of false positives increases the probability of anomaly, and increasing uncertainty yields lower confidence in anomaly detection. However, Mahalanobis distance also shows significant variability, specifically with resonances near orbital half-periods (shown 12 and 36 hours for GEO). At these points, the non-maneuvered and maneuvered orbits, for $\mathcal{H}_0$ and $\mathcal{H}_1$ respectively, are at their furthest separation ($\mathcal{H}_0$ at its maximum out-of-plane distance), making the state difference large. Additionally, a smaller spike can be seen at the orbit periods (24 and 48 hours) where both orbits lie in the equatorial plane, due to significant velocity vector differences between the $\mathcal{H}_0$ and $\mathcal{H}_1$ orbits despite similar positions.

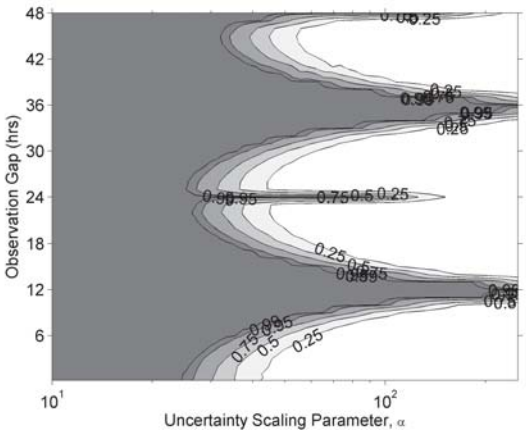

Figure 10. Mahalanobis distance probability of anomaly contours vs uncertainty scaling parameter and observation gap, simulated inclination change maneuver, $P_{FA} = 0.05$.

It can be seen that, at some observation gap durations, Mahalanobis distance is significantly more sensitive to inclination change maneuvers, as evidenced by the higher probabilities of anomaly at higher uncertainty. However, this is not always the case, as at times the control distance metric and Mahalanobis distance metric are comparable in detection probabilities. This variation with observation gap time is a dangerous aspect of using Mahalanobis distance to develop maneuver detection thresholds. Since Mahalanobis distance is very dependent upon the observation time, this requires a better understanding and more careful consideration of the dynamics of the spacecraft and what kind of maneuver it would perform at what point in the orbit. Control distance, however, is more consistent with time, allowing the development of a more general-use threshold for the North-South stationkeeping maneuver at GEO. Additionally, the reliability of uncertainty quantification is a concern when considering the uncertainty sensitivity advantage of Mahalanobis distance. In Section VI a synergistic implementation is discussed, using both control distance and Mahalanobis distance to leverage the advantages of both methods.

Alternate test cases (omitted from this paper for space and uniformity between examples) showed that both al-

gorithms were more sensitive to velocity uncertainty than position uncertainty. For instance, in the control distance test cases shown here, the algorithm struggles to detect maneuvers above $\alpha = 100$, which corresponds to 100 meters in position uncertainty and 100 centimeters-per-second in velocity uncertainty. Manually changing the boundary condition uncertainties to include cases lower velocity uncertainty showed improved anomaly detection performance even at higher position uncertainty.

D. Synthetic Data - Phasing Maneuvers

A similar sensitivity study is conducted for a phasing maneuver at GEO. This maneuver type is termed East-West station-keeping, as it refers to the satellite maintaining a specific longitude over the Earth. A 0.1 degree change in longitude is prescribed, selected to represent drifting completely across a ±0.05 degree GEO slot. The same parameters (observation gap, boundary condition uncertainty, and false alarm rate) are varied over the same ranges.

Figures 11 and 12 present the probability of anomaly contours for the orbit phasing maneuver for both control distance and Mahalanobis distance, respectively. More complete sets of data are featured in Appendix A, Fig. 18 for control distance and Fig. 19 for Mahalanobis distance. Once again, some trends hold true for both control distance and Mahalanobis distance: increasing uncertainty decreases anomaly probability while increasing allowable false alarm rate increases anomaly probability. In this case, however, both methods show significant variation with observation gap. Performing a 0.1 degree longitude change over 10 minutes requires significantly more fuel than over 24 hours, where the spacecraft can more effectively utilize natural dynamics and slightly adjust its semi-major axis to transfer to a different point in the orbit. Therefore, maneuvers of this kind over small observation gaps are much easier to detect.

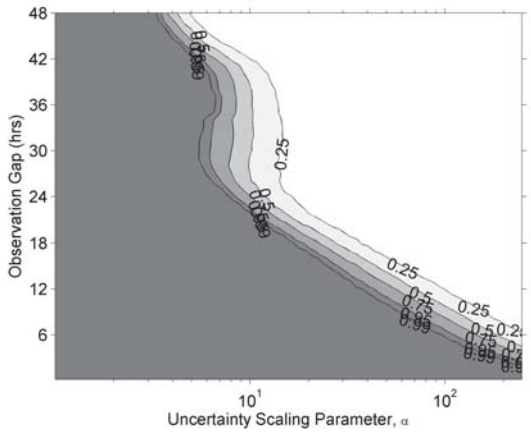

Figure 11. Control distance probability of anomaly contours vs uncertainty scaling parameter and observation gap, simulated phasing maneuver, $P_{FA} = 0.05$.

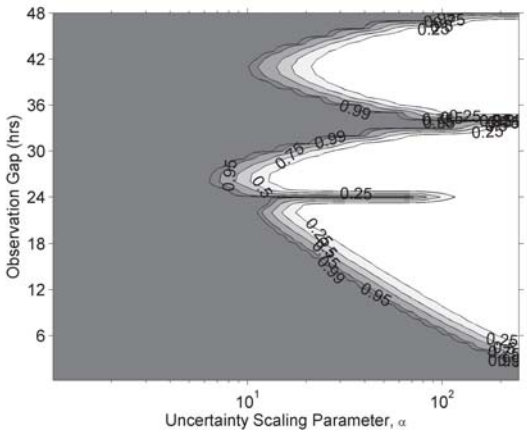

Figure 12. Mahalanobis distance probability of anomaly contours vs uncertainty scaling parameter and observation gap, simulated phasing maneuver, $P_{FA} = 0.05$.

An interesting note in this scenario is that control distance either matches or outperforms Mahalanobis distance for much of the first orbit period, predicting higher probabilities of anomaly for the same uncertainty, observation gap, false alarm rate triplet. However, as the observation gap increases, Mahalanobis distance again shows considerable variation, yielding higher anomaly probabilities than control distance at 34 hours and lower probabilities at 26 hours. Once again, the observation time can be seen to be a significant factor in anomaly detection for Mahalanobis distance, requiring knowledge of the spacecraft's position in its orbit. For control distance, while it is more sensitive to time for East-West maneuvers than North-South, it still shows improved consistency in medium-duration cases, albeit with lower anomaly sensitivity.

V. Empirical Data Results

A. Station-Keeping

To complement the simulated scenarios, the algorithm is also tested using real operational data, the availability of which drove the construction of the simulated inclination change scenario. The real data, taken from observations of the Galaxy 15 geostationary satellite by the Wide Area Augmentation System (WAAS), spans a month of operation and includes Earth-centered Earth-fixed (ECEF) position and velocity, as well as radial, in-track, and cross-track (RIC) acceleration, as seen by a rotating Hill frame attached to the spacecraft. WAAS is an extremely accurate navigation system that uses a network of ground-based reference stations to measure small variations in GPS satellite signals to develop deviation corrections (DCs). The DCs are then broadcast by GPS satellites to improve position accuracy calculations for WAAS-enabled GPS receivers.

Figure 13 shows the cross-track acceleration data for

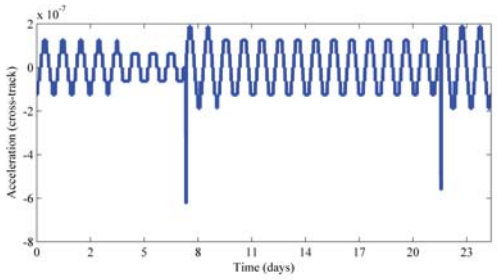

Figure 13. Cross-track acceleration for Galaxy 15 satellite, real data

the empirical dataset. Inspection of the acceleration data reveals two large anomalous cross-track acceleration events, candidates for North-South inclination station-keeping maneuvers, during days 7 and 22. The selected maneuver, the peak during day 7, resulted in a 0.03 degree inclination change. Simulation initial and final conditions are selected corresponding to the desired observation gap such that the maneuver is always in the middle of the selected time span. For instance, for a 6 hour observation gap, the initial condition is the spacecraft state 3 hours before the maneuver, and the final condition is the state 3 hours after the maneuver. The real-world data is analyzed in a similar manner to the synthetic data by varying observation gap, boundary condition uncertainty, and prescribed false alarm rate.

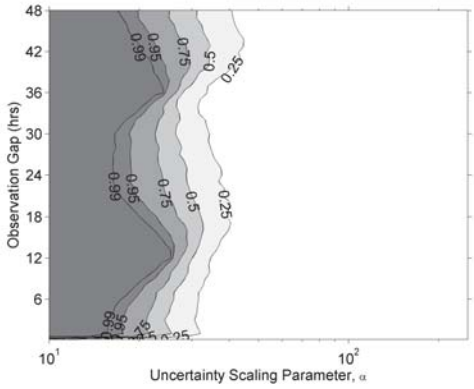

Figure 14. Control distance probability of anomaly contours vs uncertainty scaling parameter and observation gap, real-data (WAAS) inclination change maneuver, $P_{FA} = 0.05$.

The results of the real data sensitivity study are seen in Fig. 14 and 15 for control distance and Mahalanobis distance, respectively. This data is nearly identical to that of the simulated inclination change scenario, showing the same trends with respect to all varied parameters. In this scenario, Mahalanobis distance is still more sensitive in its detection of the maneuver at higher uncertainties. However, control distance remains more consistent with respect to observation gap, allowing for improved application to arbitrary space objects without requiring specific knowledge of the object's spot in its orbit. The agreement

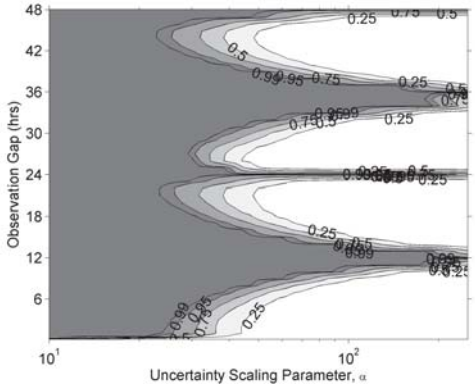

Figure 15. Mahalanobis distance probability of anomaly contours vs uncertainty scaling parameter and observation gap, real-data (WAAS) inclination change maneuver, $P_{FA} = 0.05$.

between these results and Fig. 9 and 10 lends confidence to the applicability of this approach in an operational setting. Additional results using different false alarm rates are presented in Appendix A, Fig. 20 and 21 for control distance and Mahalanobis distance, respectively. As with the synthetic scenarios, increasing P_{FA} lowers the threshold for anomaly probability calculation and thus increases anomaly probability. The remaining observation gap and uncertainty trends remain unaffected by P_{FA}.

VI. Operational Application

The simulation results presented in this paper explore a range of different maneuver and boundary condition cases; this section condenses these results and provides recommended operational use cases for each anomaly detection method. In situations where the separation between the Gaussian components of the Mahalanobis distance Gaussian mixture are greater than their covariances, combined Mahalanobis distance PDFs tend to be segmented (see Fig. 8), which significantly reduces the ability to detect anomalies using the Mahalanobis distance method. Moreover, regardless of the mixture component separations, the Mahalanobis distance method incurs a significant computation time penalty for non-Gaussian boundary conditions due to the costly requirement to recompute the relevant covariance at each intermediate propagation time-step. In contrast, the computational burden of control distance is approximately constant regardless of whether or not boundary conditions are Gaussian. In the Gaussian case, control distance takes approximately 2- to 3-times as long to compute, and Mahalanobis distance is more sensitive in maneuver detection in most but not all cases. The two methods are complementary; when one performs less effectively, the other remains sensitive to anomalies. The WAAS data supports results generated by the simulated cases, providing empirical support for these results.

Based on these observations, the control distance met-

ric is recommended in applications with non-Gaussian boundary condition uncertainties to avoid significant computation penalties, as well as poorly conditioned anomaly detection in more-segmented Gaussian mixture cases. For Gaussian boundary conditions, Mahalanobis distance is in general more sensitive to anomalies. However, since neither method always dominates in anomaly detection for Gaussian boundary conditions, the control distance method could be implemented alongside Mahalanobis distance without significant computational complexity increase. Implementing control cost in both cases gracefully handles transitions from Gaussian to non-Gaussian boundary conditions, maintaining performance without introducing significant computational complexity.

VII. Conclusions

The control distance metric provides a natural means of associating spacecraft observations by subsuming the orbital dynamics into the association metric. This effort modifies the control distance anomaly detection approach to address error-rate shortcomings in the single hypothesis method and relax Gaussian boundary condition assumptions. The inclusion of Gaussian mixture approximations and a binary hypothesis test approach allow control of allowable error rates and enable comparisons using anomalous and quiescent hypotheses. An analogous set of hypotheses is constructed for Mahalanobis distance and extended to allow for non-Gaussian uncertainties. For both Gaussian and non-Gaussian boundary conditions, simulation results show control distance is able to compute anomaly probabilities at an operationally acceptable computational cost. The Gaussian test cases show Mahalanobis distance to be generally but not uniformly more sensitive to anomalies. However, Mahalanobis distance anomaly detection is less consistent with observation gap, and the Mahalanobis distance approach is less effective in both anomaly detection and computational complexity with non-Gaussian boundary conditions. This study concludes that, the control distance method is preferred for use with non-Gaussian boundary conditions, while both Mahalanobis distance and control distance should be implemented for Gaussian boundary conditions for added robustness.

VIII. Acknowledgements

This material is based on research sponsored by the Air Force Research Laboratory under contract number FA9453-13-1-0282. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Research Laboratory or the U.S. Government.

This material is also based on work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1148903. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

The authors would also like to acknowledge Dr. Morgan Baldwin for much guidance in the development of the maneuver detection framework and Dr. Kyle DeMars for enlightening discussions regarding Gaussian mixtures.

References

[1] James, L., "Keeping the Space Environment Safe for Civil and Commercial Users," *Statement of Lieutenant General Larry James, Commander, Joint Functional Component Command for Space, Before the Subcommittee on Space and Aeronautics*, 2009.

[2] Liou, J.-C., "Modeling the Large and Small Orbital Debris Populations for Environment Remediation," Tech. rep., NASA Orbital Debris Program Office, Johnson Space Center, Houston, TX, June 2014.

[3] Cox, C., Degraaf, E. J., Wood, R. J., and Crocker, T. H., "Intelligent Data Fusion for Improved Space Situational Awareness,," *Space 2005*, 2005.

[4] Oltrogge, D., Alfana, S., and Gist, R., "Satellite Mission Operations Improvements Through Covariance Based Methods," *AIAA 2002-1814, SatMax 2002: Satellite Performance Workshop*, April 2002.

[5] Kelso, T. S., "Analysis of the Iridium 33 - Cosmos 2251 Collision," *Advanced Maui Optical and Space Surveillance Technologies Conference*, September 2009.

[6] Kelecy, T., Hall, D., Hamada, K., and Stocker, D., "Satellite Maneuver Detection Using Two-line Element (TLE) Data," *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*, 2007.

[7] Ray, A., "Symbolic Dynamic Analysis of Complex Systems for Anomaly Detection," *Mechanical Systems and Signal Processing*, Vol. 21, 2007.

[8] Kelecy, T. and Jah, M., "Detection and Orbit Determination of a Satellite Executing Low Thrust Maneuvers," *Acta Astronautica*, Vol. 66, No. 5, 2010, pp. 798–809.

[9] Hoots, F. R. and Schumacher, P. W., "History of Analytical Orbit Modeling in the U.S. Space Surveillance System," *Journal of Guidance, Control, and Dynamics*, Vol. 27, 2004, pp. 174–185.

[10] Horwood, J. T., Aragon, N. D., and Poore, A. B., "A Gaussian Sum Filter Framework for Space Surveillance," *Signal and Data Processing of Small Targets*, Vol. 8137, San Diego, CA, 2011.

[11] Tommei, G., Milani, A., and Rossi, A., "Orbit determination of space debris: admissible regions," *Celestial Mechanics and Dynamical Astronomy*, Vol. 97, No. 4, 2007, pp. 289–304.

[12] Maruskin, J. M., Scheeres, D. J., and Alfriend, K. T., "Correlation of Optical Observations of Objects in Earth Orbit," *Journal of Guidance, Control, and Dynamics*, Vol. 32, No. 1, Jan. 2009, pp. 194–209.

[13] Fujimoto, K. and Scheeres, D., "Correlation of Optical Observations of Earth-Orbiting Objects and Initial Orbit Determination," *Journal of Guidance, Control, and Dynamics*, Vol. 35, No. 1, 2012, pp. 208–221.

[14] DeMars, K. J., Hussein, I. I., Frueh, C., Jah, M. K., and Scott Erwin, R., "Multiple-Object Space Surveillance Tracking Using Finite-Set Statistics," *Journal of Guidance, Control, and Dynamics*, March 2015, pp. 1–16.

[15] Huang, J., Hu, W.-D., Xin, Q., and Du, X.-Y., "An Object Correlation and Maneuver Detection Approach for Space Surveillance," *Research in Astonomy and Astrophysics*, Vol. 12, 2012, pp. 1402–1416.

[16] Gallego, G., Cuevas, C., Mohedano, R., and Garcia, N., "On the Mahalanobis Distance Classification Criterion for Multidimensional Normal Distributions," *IEEE Transactions on Signal Processing*, Vol. 61, No. 17, June 2013, pp. 4387–4396.

[17] Holzinger, M. J., Scheeres, D. J., and Alfriend, K. T., "Object Correlation, Maneuver Detection, and Characterization Using Control-Distance Metrics," *Journal of Guidance, Control, and Dynamics*, Vol. 35, No. 4, July 2012, pp. 1312–1325.

[18] Holzinger, M. J., "Delta-V Metrics for Object Correlation, Maneuver Detection, and Maneuver Characterization," *AIAA Guidance, Navigation and Control Conference*, Portland, OR, August 2011.

[19] Yan, Q., "Distributed Signal Detection under the Neyman-Pearson Criterion," *IEEE Transactions on Information Theory*, Vol. 47, 2001, pp. 1368–1377.

[20] Neyman, J. and Pearson, E. S., "On the Problem of the Most Efficient Tests of Statistical Hypotheses," *Philosophical Transactions of the Royal Society A: Mathetmatical, Physical and Engineering Sciences*, Vol. 231, 1933, pp. 239–337.

[21] Li, X. Q. and King, I., "Gaussian Mixture Distance for Information Retrieval," *IEEE International Joint Conference on Neural Netrowks*, Vol. 4, 1999.

[22] Prussing, J. E. and Conway, B. A., *Orbital Mechanics*, Oxford University Press, 2012.

[23] Hill, K., Alfriend, K. T., and Sabol, C., "Covariance-based Uncorrelated Track Association," *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, 2008.

[24] Jaunzemis, A. D., Mathew, M. V., and Holzinger, M. J., "Maneuver Detection using Gaussian Mixtures and Real Data," *25th AAS/AIAA Space Flight Mechanics Meeting*, Williamsburg, VA, January 2015.

[25] Holzinger, M. J., Scheeres, D. J., and Alfriend, K. T., "Delva-V Distance Object Correlation and Maneuver Detection with Dynamis Parameter Uncertainty and Generalized Constraints," *Proceedings of the 219th American Astronomical Society Meeting*, Austin, TX, January 2012.

IX. Appendix A - Sensitivity Study Results

This appendix contains a larger selection of data from the hypothesis testing sensitivity studies. Trends using this data are highlighted in the Simulation Results and Empirical Results sections.

A. Synthetic Data: North-South Station-Keeping

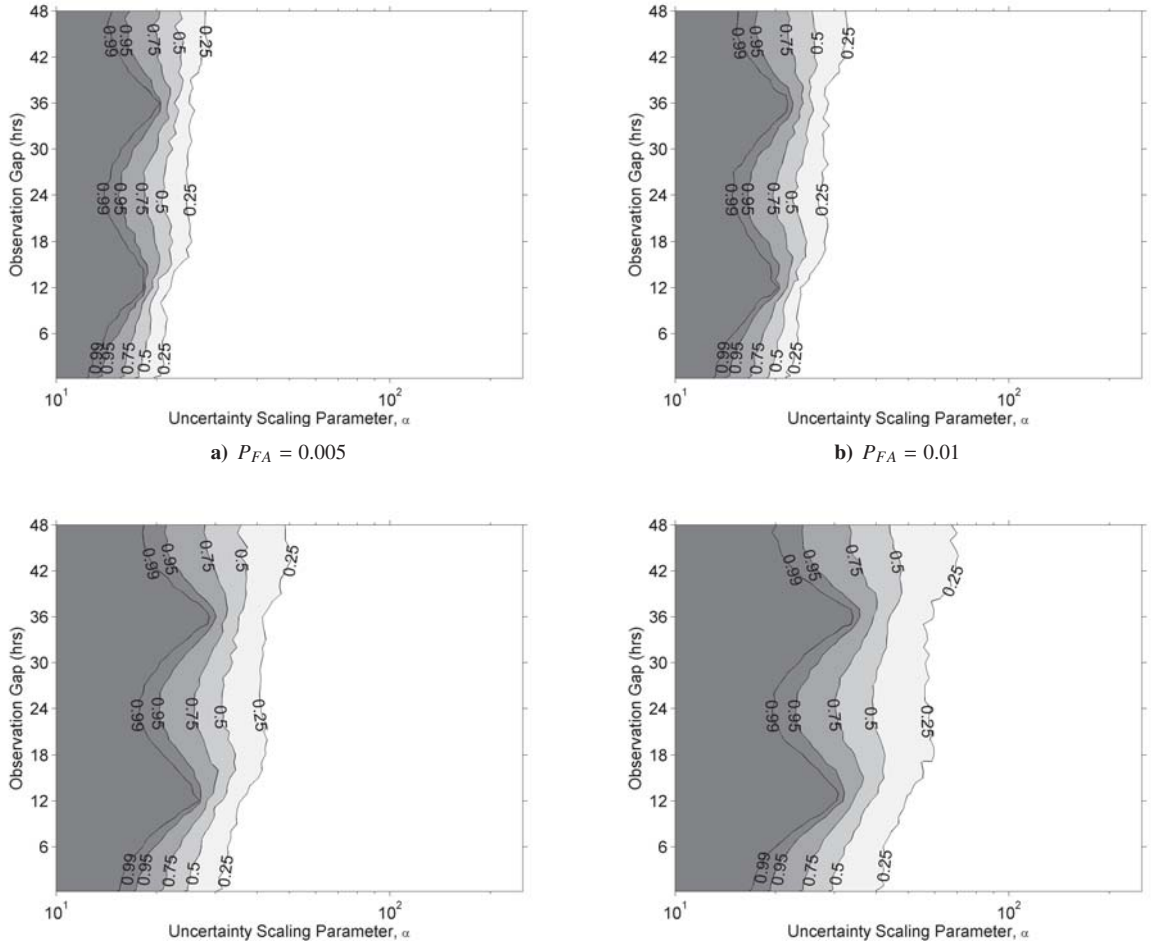

a) $P_{FA} = 0.005$ **b)** $P_{FA} = 0.01$

c) $P_{FA} = 0.05$ **d)** $P_{FA} = 0.1$

Figure 16. Control distance probability of anomaly contours vs uncertainty scaling parameter and observation gap, simulated inclination change maneuver.

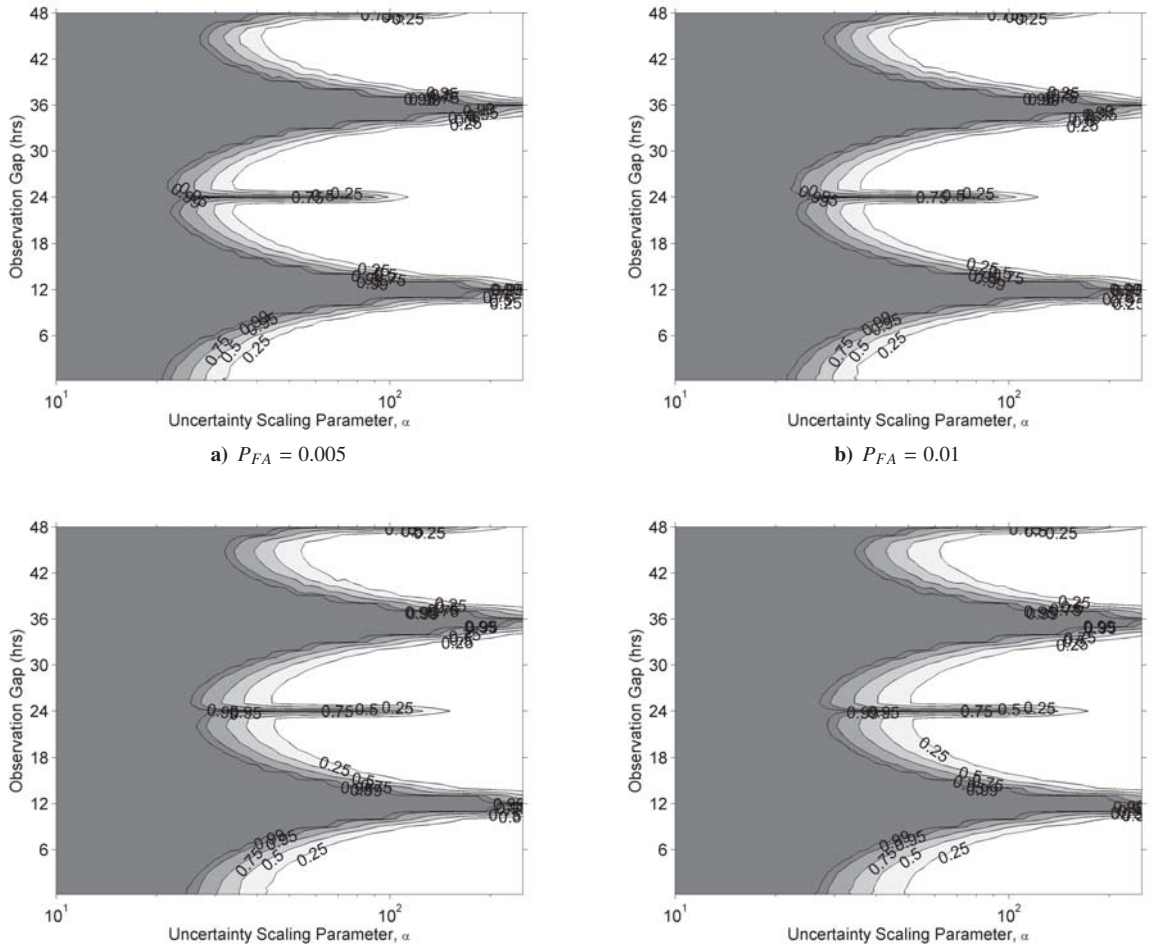

a) $P_{FA} = 0.005$ **b)** $P_{FA} = 0.01$

c) $P_{FA} = 0.05$ **d)** $P_{FA} = 0.1$

Figure 17. Mahalanobis distance probability of anomaly contours vs uncertainty scaling parameter and observation gap, simulated inclination change maneuver.

B. Synthetic Data: East-West Station-Keeping

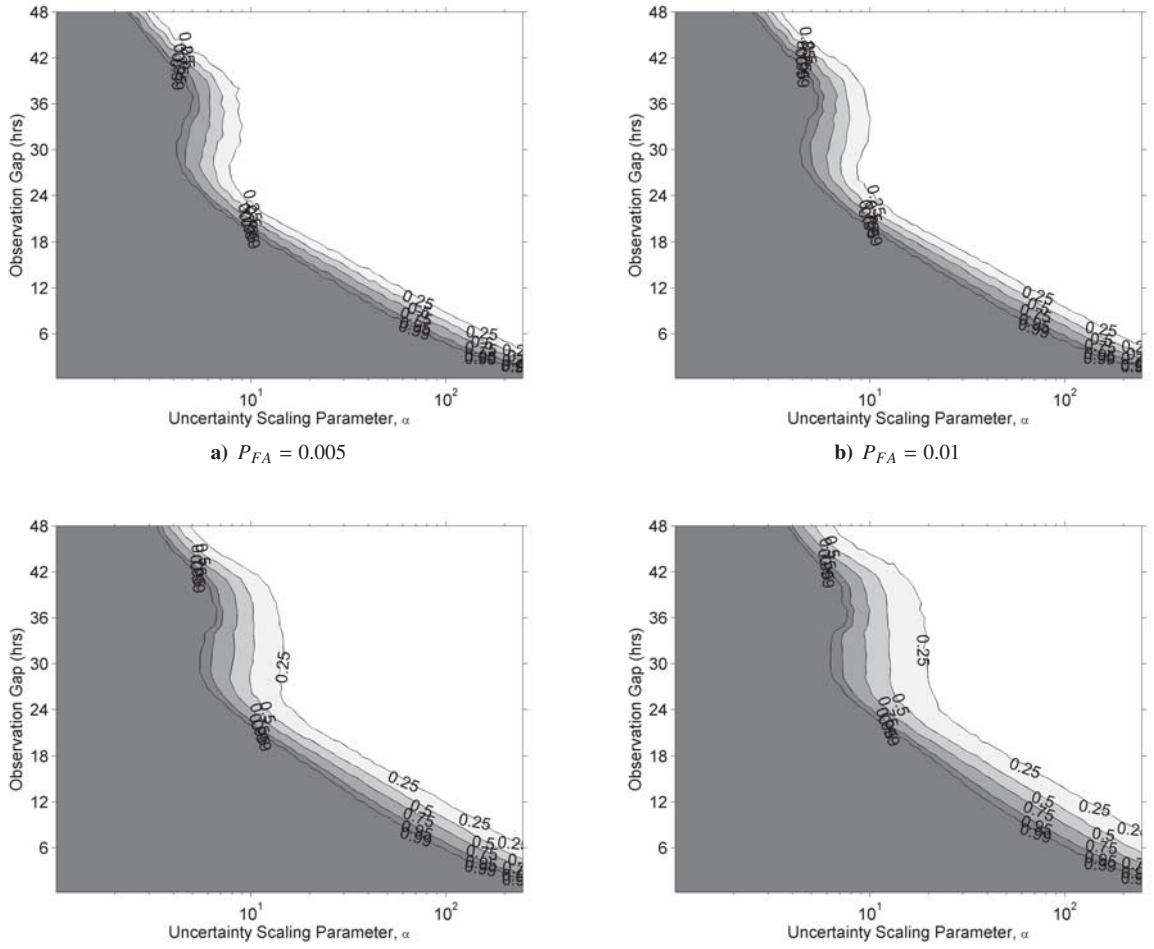

a) $P_{FA} = 0.005$ **b)** $P_{FA} = 0.01$

c) $P_{FA} = 0.05$ **d)** $P_{FA} = 0.1$

Figure 18. Control distance probability of anomaly contours vs uncertainty scaling parameter and observation gap, simulated phasing maneuver.

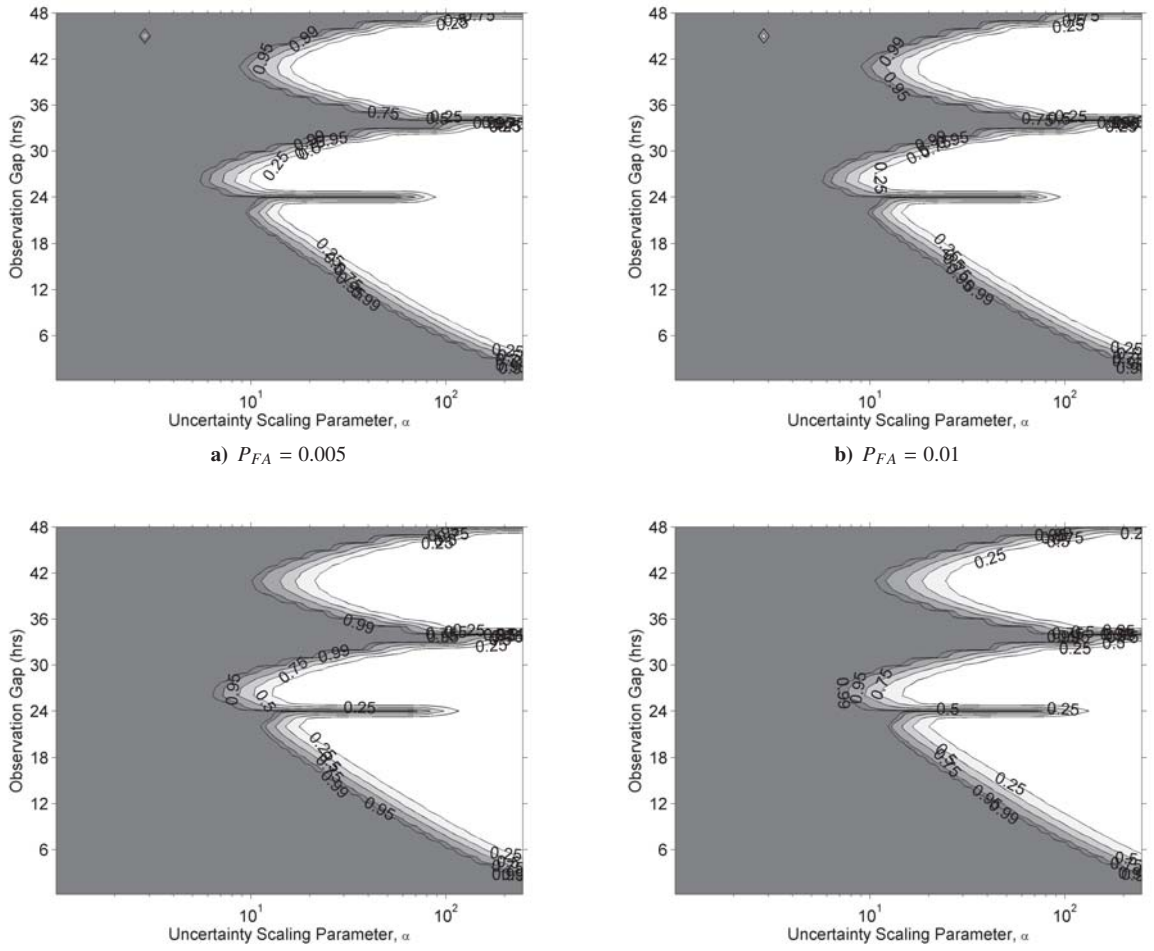

a) $P_{FA} = 0.005$ **b)** $P_{FA} = 0.01$

c) $P_{FA} = 0.05$ **d)** $P_{FA} = 0.1$

Figure 19. Mahalanobis distance probability of anomaly contours vs uncertainty scaling parameter and observation gap, simulated phasing maneuver.

C. Empirical Data: North-South Station-Keeping

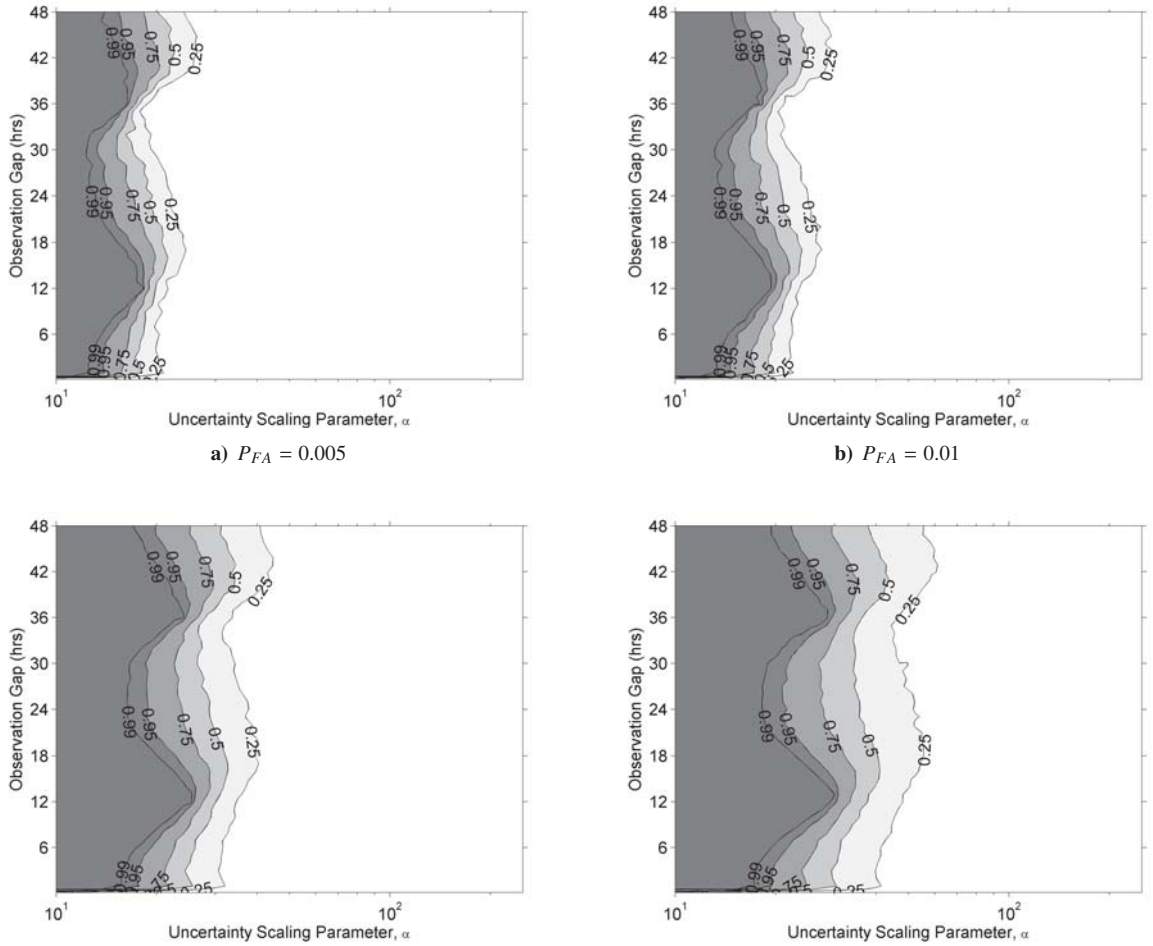

a) $P_{FA} = 0.005$ **b)** $P_{FA} = 0.01$

c) $P_{FA} = 0.05$ **d)** $P_{FA} = 0.1$

Figure 20. Control distance probability of anomaly contours vs uncertainty scaling parameter and observation gap, real-data (WAAS) inclination change maneuver.

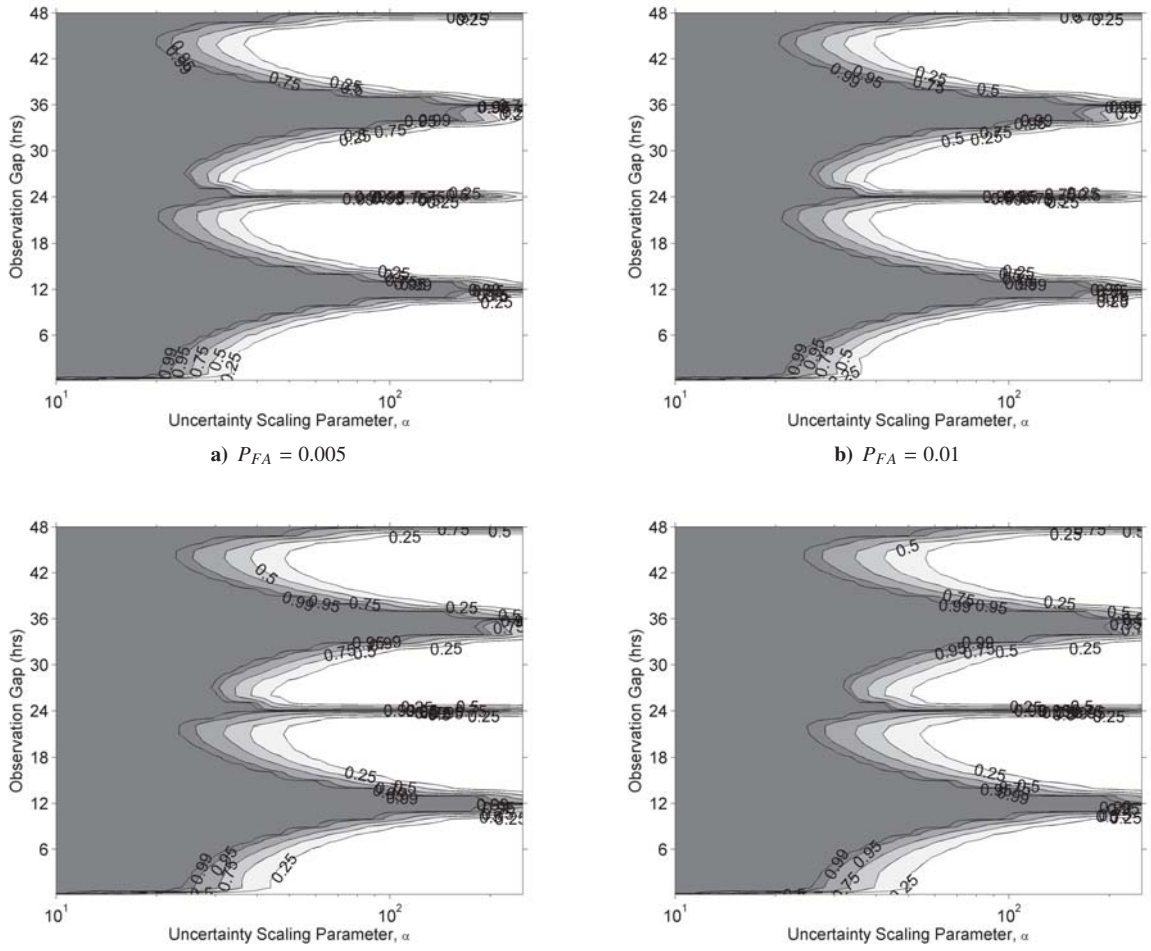

a) $P_{FA} = 0.005$ **b)** $P_{FA} = 0.01$

c) $P_{FA} = 0.05$ **d)** $P_{FA} = 0.1$

Figure 21. Mahalanobis distance probability of anomaly contours vs uncertainty scaling parameter and observation gap, real-data (WAAS) inclination change maneuver.