CONTROL METRIC MANEUVER DETECTION WITH GAUSSIAN MIXTURES AND REAL DATA

Andris D. Jaunzemis[*], **Midhun Mathew**[†], **Marcus J. Holzinger**[‡]

The minimum-fuel distance metric provides a natural tool with which to associate space object observation data. A trajectory optimization and anomaly hypothesis testing algorithm is developed based on the minimum-fuel distance metric to address observation correlation under the assumption of optimally maneuvering spacecraft. The algorithm is tested using inclination-change scenarios with both synthetic and real data gathered from the Wide Area Augmentation System (WAAS). Comparisons to other commonly-used association metrics such as Mahalanobis distance reveal less sensitivity in anomaly detection but improved consistency with respect to observation cadence, while providing added data through the reconstruction of the optimal maneuver. Non-Gaussian boundary conditions are also approached through an analytical approximation method, yielding significant computational complexity improvements.

INTRODUCTION & BACKGROUND

Correlating on-orbit observations and detecting space object maneuvers is a challenging endeavor in the Space Situational Awareness (SSA) field. There are currently at least 17,000 trackable on-orbit objects, 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] however, recent events highlight the mutual interest that national and private operators share for accurate object correlation and maneuver detection capability.[5]

This study examines a novel method for correlating space object tracks with known resident space objects (RSOs) and characterizing and reconstructing RSO maneuvers. The problem of associating uncorrelated tracks (UCTs) over large time periods is particularly difficult when objects 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.[6] Since both the initial and final UCTs are best estimates, with associated uncertainty distributions, the question of correlation becomes difficult to answer in operational settings, particularly in densely-populated regions of the space environment. Given a propagated best estimate of the state and its associated distribution, correlating UCTs tests whether a new observation (with its associated uncertainty) is a previously observed object, and if not, determines the discrepancy. There are many distance- or

---
[*]Graduate Research Assistant, Georgia Institute of Technology, Atlanta, GA, AIAA Student Member
[†]Graduate Research Assistant, Georgia Institute of Technology, Atlanta, GA, AIAA Student Member
[‡]Assistant Professor, Georgia Institute of Technology, Atlanta, GA, AIAA Senior Member

pseudo-distance metrics that may be used to measure the discrepancy between two state distributions (e.g. Mahalanobis distance). Problematically, none of these distance metrics directly quantify the level of propulsive effort required to cause the observed state change.

In this study, the minimum-fuel control distance metric is utilized in a hypothesis testing algorithm to approach data association and maneuver detection problems. Since on-board fuel remains a scarce commodity for operational spacecraft, operators are likely to execute optimal or near-optimal maneuvers.[6] Under the assumptions of optimal control, multiple deterministic observations 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 has shown, through the properties of strict positivity, symmetry, and triangle inequality, that control performance is a metric,[6] allowing comparisons to other commonly used distance metrics. The inputs to the maneuver detection algorithm are two UCTs, containing the estimated states and associated covariances of the trajectory boundary conditions. An optimal connecting trajectory is then computed in a minimum-quadratic-control trajectory optimization implementation. The resulting optimal trajectory is used to generate a control cost distribution associated with perturbations in the boundary conditions based on the input covariances.[6] The optimal control cost and control cost distribution are used to test the hypothesis that a maneuver occurred and complete the maneuver detection task.[6,7,8] This framework is tested using simulated and real-world operational data.

In addition to the minimum-fuel control distance maneuver detection algorithm, this study endeavors to analytically incorporate non-Gaussian boundary conditions for increased realism in representing the boundary condition UCTs. State uncertainty is a fundamental problem of space situational awareness in the cases of track association and maneuver detection.[9] The probability density function must be accurately represented to address the requirement of uncertainty consistency. Initially, the state uncertainty may be assumed to be Gaussian, which is often appropriate for resident space objects. However, the dynamics governing the orbital mechanics (e.g. gravity, drag, third body perturbations, etc.) are nonlinear. After propagation for some time, the state uncertainty may become non-Gaussian.[9] Gaussian mixtures provide a convenient method to represent non-Gaussian uncertainties and to generate a combined probability model for the total overlapping uncertainties of each state. The uncertainty in measurements may be non-Gaussian, especially after propagation of the previous best estimate, or there may be multiple sensor measurements. However, many probability density functions can be closely approximated by a Gaussian mixture.[9] Multiple state estimate and covariance pairs can be used to construct a Gaussian mixture which can more accurately describe the UCT and measurement uncertainties.

The contributions of this paper are 1) the implementation of a maneuver detection hypothesis testing algorithm using the minimum-fuel control distance metric and its evaluation using both synthetic and real-world data, 2) the comparison of the minimum-fuel control distance metric to other data association approaches (e.g. Euclidean and Mahalanobis distance), and 3) the statistical treatment of control distributions using Gaussian mixtures for non-Gaussian boundary conditions.

THEORY

The algorithm described in this section endeavors to determine a correlation between two UCTs under the hypothesis that the tracks belong to the same maneuverable space object and the object has maneuvered optimally between the observations. As pictured in Figure 1, an uncorrelated track has been generated at time t_A with best-estimate state $\mathbf{x}_A$ and covariance $\mathbf{P}_A$. This track has been prop-
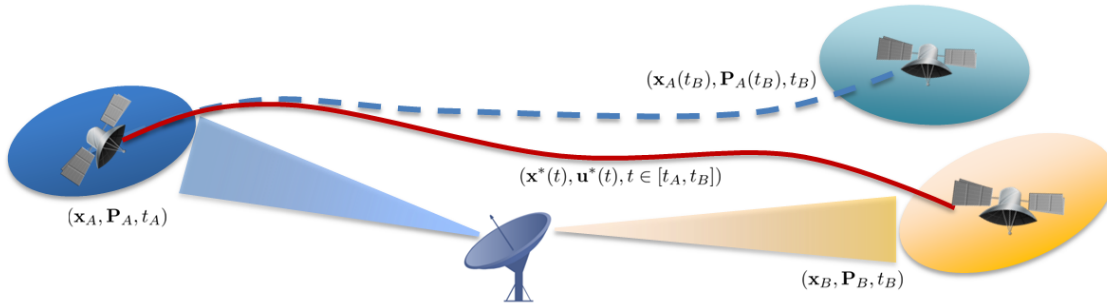
Figure 1. Maneuver detection scenario

agated forward to time t_B, yielding the propagated state $\mathbf{x}_A(t_B)$ and covariance $P_A(t_B)$. Another uncorrelated track has been generated at time t_B as well, with best-estimate state $\mathbf{x}_B$ and covariance P_B. While many algorithms function by comparing the propagated state $\mathbf{x}_A(t_B)$ to the new state $\mathbf{x}_B$, the minimum-fuel maneuver detection algorithm instead connects states $\mathbf{x}_A$ and $\mathbf{x}_B$ with an optimal trajectory and comparing the magnitude of the maneuver required to the state uncertainty. The individual components of the algorithm are addressed in the following sections.

Trajectory Optimization

The maneuver detection algorithm begins with a trajectory optimization routine, which uses the nominal states of the boundary conditions to generate the optimal connecting trajectory. 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 the chosen cost function. In this study, the trajectory is optimized by minimizing the quadratic control cost, shown in Eq. (1).[10] The quadratic control cost function is ideal for variable specific impulse (VSI) engines, often used in low-thrust applications, but the algorithm could be adjusted to use a different cost function (e.g. an impulsive cost function).

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

The deterministic two-point boundary value problem is formulated into a direct 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 (J2, J22, J3, 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, improving performance of the implementation. The output of the direct optimization step is a nominal optimal state and control trajectory connecting the UCTs. 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. Adjusting tunable convergence parameters of the optimization algorithm, such as absolute and relative tolerances, can reduce this error.

3

The optimal states $\mathbf{x}^*(t)$ and controls $\mathbf{u}^*(t)$ for $t_A \leq t \leq t_B$ are then used to compute a time-history of optimal control costates $\mathbf{p}^*(t)$. The costates are constructed by guessing a costate trajectory based on the optimal control trajectory and then modifying the costates through an iterative least-squares method until the control cost is within a user-defined convergence tolerance of the optimal trajectory. As an additional check, the new optimal costate trajectory is compared to the original optimal trajectory to ensure the endpoint error has not changed significantly.

Control Cost Distribution Generation

Using the optimal state and costate trajectory ($\mathbf{x}^*(t)$ and $\mathbf{p}^*(t)$), uncertainties in position and velocity are included to generate an approximate probability distribution of the control cost. This is accomplished by linearizing about the nominal optimal trajectory and applying perturbations in the initial and final conditions based on the covariance of each associated UCT. Performing this many times generates a distribution of optimal or sub-optimal trajectories, and the control cost associated with these perturbed trajectories can be computed using the results from Holzinger et al.[6] Key results used in generating and analyzing the control cost distributions for this work are reiterated here. In general, the cost distribution can be written succinctly as shown in Eq. (2):

$$J = J^* + \delta J \tag{2}$$

where δJ is a perturbation to the optimal cost, J^*, caused by variations in the initial and final boundary conditions. This cost function is utilized by taking the Taylor series expansion up to the first-order terms to approximate a value for the perturbed cost. From Holzinger et al.,[6] there exists a function $\mathbf{\Lambda}(t, t_0)$, shown in Eq. (3), that maps variations in initial and final states to variations in the costate $\delta \mathbf{p}(t)$ at time t.

$$\mathbf{\Lambda}(t, t_0) = \mathbf{\Phi}_{px}(t, t_0) - \mathbf{\Phi}_{pp}(t, t_0)\mathbf{\Phi}_{xp}(t_f, t_0)^\dagger \mathbf{\Phi}_{xx}(t_f, t_0)\mathbf{\Phi}_{pp}(t, t_0)\mathbf{\Phi}_{xp}(t_f, t_0)^\dagger \tag{3}$$

Note that this function is composed of portions of the state transition matrix. The pseudoinverse term $\mathbf{\Phi}_{xp}(t_f, t_0)^\dagger$ is not guaranteed to exist for arbitrary systems, but for present purposes it is assumed to exist.[6] The cost function in Eq. (1) is expanded by substituting $\mathbf{u}(\tau) = \mathbf{u}^*(\tau) + \delta\mathbf{u}(\tau)$:

$$J \approx \frac{1}{2}\int_{t_0}^{t_f} \mathbf{u}^*(\tau)^{\mathrm{T}}\mathbf{u}^*(\tau)\mathrm{d}\tau + \int_{t_0}^{t_f} \mathbf{u}^*(\tau)^{\mathrm{T}}\frac{\partial \mathbf{f}}{\partial \mathbf{u}}^{\mathrm{T}}\mathbf{\Lambda}(\tau, 0)\delta\mathbf{z}\mathrm{d}\tau \tag{4}$$

$$+ \frac{1}{2}\int_{t_0}^{t_f} \delta\mathbf{z}^{\mathrm{T}}\mathbf{\Lambda}(\tau, 0)^{\mathrm{T}}\frac{\partial \mathbf{f}}{\partial \mathbf{u}}\frac{\partial \mathbf{f}}{\partial \mathbf{u}}^{\mathrm{T}}\mathbf{\Lambda}(\tau, 0)\delta\mathbf{z}\mathrm{d}\tau$$

where $\delta\mathbf{z}^{\mathrm{T}} = \begin{bmatrix} \delta\mathbf{x}_0^{\mathrm{T}} & \delta\mathbf{x}_f^{\mathrm{T}} \end{bmatrix} \in \mathbb{R}^{2n}$. Since $\delta\mathbf{z}$ is independent of τ, the cost function can be rewritten using the following definitions:

$$J^* = \frac{1}{2}\int_{t_0}^{t_f} \mathbf{u}^*(\tau)^{\mathrm{T}}\mathbf{u}^*(\tau)\mathrm{d}\tau \tag{5}$$

$$\boldsymbol{\omega}(t_f, t_0) = \int_{t_0}^{t_f} \mathbf{\Lambda}(\tau, 0)^{\mathrm{T}}\frac{\partial \mathbf{f}}{\partial \mathbf{u}}\mathbf{u}^*(\tau)\mathrm{d}\tau \tag{6}$$

$$\boldsymbol{\Omega}(t_f, t_0) = \frac{1}{2}\int_{t_0}^{t_f} \mathbf{\Lambda}(\tau, 0)^{\mathrm{T}}\frac{\partial \mathbf{f}}{\partial \mathbf{u}}\frac{\partial \mathbf{f}}{\partial \mathbf{u}}^{\mathrm{T}}\mathbf{\Lambda}(\tau, 0)\mathrm{d}\tau \tag{7}$$

Note that J^* is the optimal cost of the optimal trajectory $(\mathbf{x}^*(t), \mathbf{p}^*(t))$ without boundary-condition variations. The cost function in Eq. (4) is rewritten in a more compact form:

$$J \approx J^* + \boldsymbol{\omega}(t_f, t_0)^\mathrm{T}\delta\mathbf{z} + \delta\mathbf{z}^\mathrm{T}\boldsymbol{\Omega}(t_f, t_0)\delta\mathbf{z} \tag{8}$$

The process for computing the control cost distribution is simplified using statistical treatments for the mean (μ_J) and variance (σ_J^2) to sample and generate a cost distribution. The process for calculating the mean and standard deviation from the optimal trajectory and covariance is based on Eqns. (9) and (10) respectively.[6]

$$\mu_J = J^* + \mathrm{Tr}\left[\boldsymbol{\Omega}\mathbf{P_z}\right] \tag{9}$$

$$\sigma_J^2 = \boldsymbol{\omega}^\mathrm{T}\mathbf{P_z}\boldsymbol{\omega} + 2\mathrm{Tr}\left[\boldsymbol{\Omega}\mathbf{P_z}\boldsymbol{\Omega}\mathbf{P_z}\right] \tag{10}$$

If it is assumed that the boundary conditions are independent of each other and do not share information, $\mathbf{P}_z \in \mathbb{R}^{12\times12}$ is composed of $\mathbf{P}_A$ in the top-left 6×6 block and $\mathbf{P}_B$ in the bottom-right 6×6 block. Using the analytical values for the first and second moments of the distribution, a Gaussian approximation of the cost distribution is generated. The distribution of control costs associated with state uncertainty, found by subtracting the optimal control cost J^* from cost distribution J, is termed the uncertainty cost distribution. Therefore, from Eq. (2), the normally distributed perturbation in the cost function is $\delta J \sim \mathcal{N}(\mathrm{Tr}\left[\boldsymbol{\Omega}\mathbf{P_z}\right], \sigma_J^2)$. In this study, the expectation of this control cost perturbation is referred to as the minimum detectable control cost, defined in Eq. (11):

$$J_{det} = \mathrm{Tr}\left[\boldsymbol{\Omega}\mathbf{P_z}\right] \tag{11}$$

so-called because it is a threshold used in the anomaly detection algorithm to determine if the state uncertainty accounts for the state change.

In order to validate the uncertainty cost distribution, a more exact distribution was constructed by sampling locally around the optimal trajectory with a random stacked vector of initial and final state perturbations, $\delta\mathbf{z} \sim \mathcal{N}(0, \mathbf{P_z})$. The quadratic cost due to this state uncertainty can then be approximated as shown in Eq. (12):[6]

$$\delta J \approx \boldsymbol{\omega}^\mathrm{T}\delta\mathbf{z} + \delta\mathbf{z}^\mathrm{T}\boldsymbol{\Omega}\delta\mathbf{z} \tag{12}$$

In this implementation, the analytical mean and standard deviation are used to size the range for the histogram used in generating the normalized PDF. Furthermore, the minimum detectable cost for this distribution can then found using the cumulative distribution function (CDF). The minimum detectable control cost (J_{det}) is defined, similar to the Gaussian approximation, as the expectation of the uncertainty control cost distribution. In testing, the Gaussian approximation agreed well with the locally sampled distribution, producing nearly identical cost distributions, thereby validating the use of the Gaussian approximation in this study.

Anomaly Detection

Using the optimal control cost and uncertainty control cost distribution generated in the previous steps, the algorithm can now perform the anomaly hypothesis testing. In this algorithm, hypothesis testing involves the computation of the cumulative distribution function, as shown in Eq. (13):[8]

$$p_A = \int_{-\infty}^{J^*} f\left(J\right)\mathrm{d}J \tag{13}$$

The hypothesis being tested is that the state difference observed between two UCTs is attributable to an anomaly (e.g. controlled maneuver, unmodeled perturbations, etc.) as opposed to quiescent propagation and measurement uncertainty. This hypothesis is tested by comparing the magnitude of the optimal control cost to the uncertainty control cost distribution to determine whether the control required to connect the UCTs is buried in the noise of the uncertainty. This comparison can be easily seen in the CDF generated from the uncertainty cost distribution, and the probability that the hypothesis is true, termed the probability of anomaly, is computed by finding the value of the CDF at a control cost value equal to the optimal control cost, as in Eq. (13). A trajectory with a nominal control cost (J^*) greater than the minimum detectable cost (J_{det}) has a higher probability of being the result of an anomaly and not the result of propagated boundary condition uncertainty.

A sample result to illustrate this method is shown in Figure 2. The figure shows the PDF and CDF for the quadratic control cost as a function of control cost. The minimum detectable control cost can be seen as the control cost where the CDF equals 0.5. The vertical dashed line represents the nominal optimal control cost for that maneuver. On the CDF plot, it can be seen that the intersection of the CDF and nominal quadratic cost yields a probability of anomaly of approximately 0.75, or 75%. The uncertainty control cost CDF is the key to computing the probability that a maneuver has occurred for a given optimal trajectory and state uncertainty.
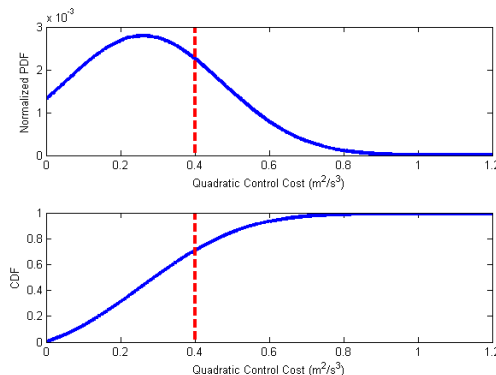

Figure 2. Notional control cost distribution (normalized PDF and CDF)

Since the control cost and cost distribution magnitudes vary widely with both simulation time horizon and UCT uncertainty, it is useful to define an anomaly detectability ratio (ADR) as the ratio of the nominal optimal control cost to the minimum detectable cost, as shown in Eq. (14):

$$\text{ADR} = \frac{J^*}{J_{det}} = \frac{J^*}{\text{Tr}\left[\mathbf{\Omega P_z}\right]} \tag{14}$$

If the nominal optimal control cost is equal to the minimum detectable quadratic cost for a given state uncertainty, then the detectability ratio is 1 and the probability of anomaly is equal to 0.5, or 50%, as seen in Figure 2. That is to say, there is an equal probability that the observed state change is caused by an anomaly (such as a maneuver) or explained by uncertainty in the initial and final states. As the nominal control cost increases with respect to the minimum detectable cost, the detectability ratio increases and the probability that the state change is caused by an anomaly increases. Eventually, the nominal cost becomes substantially larger than the cost due to uncertainty, so the probability that state uncertainty could cause the state change is approximately zero and the probability of anomaly approaches 1.0, or 100%.

6

Maneuver Detection Algorithm Framework

The entire maneuver detection algorithm is assembled as shown in Figure 3. The inputs to the algorithm are a pair of UCTs, containing the state, covariance, and time data for each track. The trajectory optimization is first performed on the best-estimate states to extract an optimal direct trajectory, which is then refit to calculate the adjoints. Then the cost distribution is generated to test the anomaly hypothesis and determine the probability that an anomaly has occurred. The output of the algorithm is the probability that the anomaly hypothesis was true, P_A. Stated differently, this is the probability that a maneuver has occurred between the UCTs.
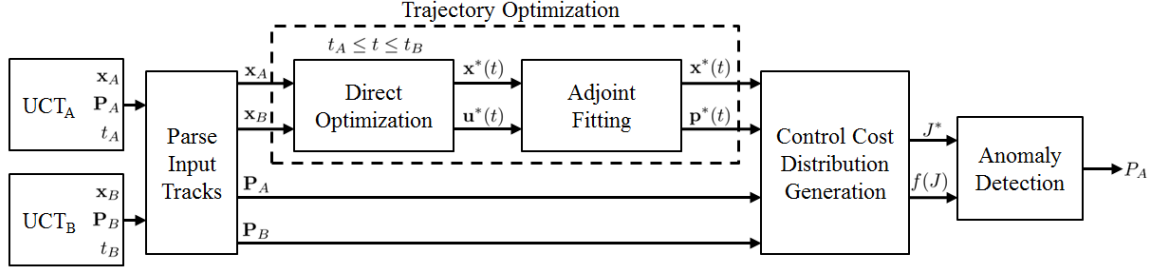

Figure 3. Minimum-fuel metric maneuver detection algorithm

Given an optimal trajectory and uncertainty in the boundary conditions, this method can compute the probability that an observed state change is due to quiescent propagation or active maneuvering, thus addressing the maneuver detection goal. Additionally, if it is determined that boundary condition uncertainty cannot account for the state change, and a maneuver or unmodeled perturbation must have occurred, the minimum magnitude and time history of the potential maneuver can be determined from the computed optimal trajectory, characterizing any such maneuver. Furthermore, if the user has knowledge of the maneuvering capabilities of the RSO and the maneuver is determined to be impossible for that specific RSO, this method can be used to associate (or disassociate) UCTs based on the control authority required.

Quadratic Control Distance Distributions using Gaussian Mixtures

By employing Gaussian mixtures to model uncertainties, the accuracy of the initial and final conditions can be improved. For boundary conditions described by Gaussian distributions, the TPBVP uses the trajectory connecting the best-estimate, or mean, state and incorporates the corresponding covariance. However, a single Gaussian distribution might not accurately describe the state or properly incorporate the uncertainties. To properly address non-Gaussian uncertainty, potential TPBVPs between separate regions of uncertainties must be considered. Using Gaussian mixtures creates a combined probability model for the non-Gaussian distributions.

Gaussian mixture models are defined as the weighted sum of a set of multivariate normal density functions. In this application, each normal distribution represents a component of the boundary condition defined by the Gaussian mixture. In general, a multivariate Gaussian mixture may be represented by the density function in Eq. (15).

$$f(x) = \sum_{i=1}^{n} w_i f_i(x) \qquad (15)$$

7

where $f_i(x) = \mathcal{N}(\mu_i, \sigma_i^2)$ is a Gaussian density function with mean μ_i and variance σ_i^2, and $w_1, \ldots, w_n$ is a set of weightings such that $w_i \geq 0$ and $\sum_i^n w_i = 1$. A random variable (ζ) sampled from this mixture may be defined as seen in Eq. (16).

$$\zeta \sim \sum_{i=1}^n w_i \mathcal{N}(\mu_i, \sigma_i^2) \tag{16}$$

In order to improve accuracy in the modeling of initial and final condition state uncertainties, the normal distribution assumption can be relaxed using Gaussian mixtures as described below. When considering boundary conditions defined by Gaussian mixtures, a combined quadratic control cost for the entire mixture must be defined. Instead of calculating the quadratic cost as a function of the cost from each trajectory connecting an initial component Gaussian distribution to each final Gaussian distribution, a trajectory connecting the mean states of the initial and final boundary conditions is used. The mean state of a boundary condition is computed as the weighted average of the best-estimate states of each component Gaussian distribution. The mean states enable the trajectory to be reduced to a single two-point boundary value problem, constructing the cost distribution by linearizing about this mean connecting trajectory and sampling initial and final states from the non-Gaussian boundary conditions.
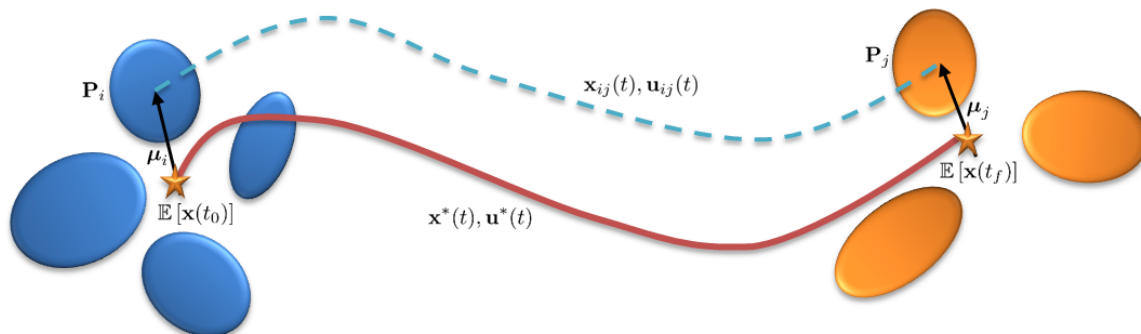

Figure 4. Gaussian Mixture scenario geometry

The method for constructing this cost distribution is based off a similar derivation for Gaussian initial and final distributions,[6] and is described below. Figure 4 shows a notional depiction of the key variables introduced in the Gaussian mixture approximation. The expected value of the states at each boundary condition Gaussian mixture is computed, and the optimal trajectory between these points is computed to find the nominal optimal trajectory ($\mathbf{x}^*(t), \mathbf{u}^*(t)$). Additionally, the deviation of the mean of each individual Gaussian from the expected value of each boundary condition state is computed as $\boldsymbol{\mu}_i$ and $\boldsymbol{\mu}_j$ for the initial and final conditions, respectively.

The quadratic control cost function, in Eq. (17), can be expanded by decomposing the control effort $\mathbf{u}(t)$ into three components as shown in Eq. (18).

$$J = \frac{1}{2} \int \mathbf{u}(\tau)^{\mathrm{T}} \mathbf{u}(\tau) \mathrm{d}\tau \tag{17}$$

$$\mathbf{u}(t) = \mathbf{u}^*(t) + \mathbf{u}_{ij} + \delta\mathbf{u}(t) \tag{18}$$

$$= -\frac{\partial \mathbf{f}}{\partial \mathbf{u}}^{\mathrm{T}} \left(\mathbf{p}^*(t) + \mathbf{p}_{ij}(t) + \delta\mathbf{p}(t) \right) \tag{19}$$

8

In Eq. (18), the term $\mathbf{u}_{ij}(t)$ represents the control cost attributable to the $\boldsymbol{\mu}_i$ and $\boldsymbol{\mu}_j$, the variations in mean state of Gaussian initial component i and final component j from the expected value of the Gaussian mixture boundary conditions. Likewise, $\delta\mathbf{u}(t)$ is the control cost attributable to $\delta\mathbf{x}_0$ and $\delta\mathbf{x}_f$, the variations in sampled state from boundary condition uncertainty. These components can be computed as shown in Eqns. (20) and (21) using $\boldsymbol{\Lambda}(t, t_0)$ from Eq. (3).

$$\mathbf{u}_{ij}(t) = -\frac{\partial \mathbf{f}}{\partial \mathbf{u}}^{\mathrm{T}} \boldsymbol{\Lambda}(t, t_0) \begin{bmatrix} \boldsymbol{\mu}_i \\ \boldsymbol{\mu}_j \end{bmatrix} \tag{20}$$

$$\delta\mathbf{u}(t) = -\frac{\partial \mathbf{f}}{\partial \mathbf{u}}^{\mathrm{T}} \boldsymbol{\Lambda}(t, t_0) \begin{bmatrix} \delta\mathbf{x}_0 \\ \delta\mathbf{x}_f \end{bmatrix} \tag{21}$$

The following vectors are defined for ease of notation:

$$\boldsymbol{\mu}_{ij} = \begin{bmatrix} \boldsymbol{\mu}_i \\ \boldsymbol{\mu}_j \end{bmatrix} \tag{22}$$

$$\delta\mathbf{z} = \begin{bmatrix} \delta\mathbf{x}_0 \\ \delta\mathbf{x}_f \end{bmatrix} \tag{23}$$

where $\boldsymbol{\mu}_{ij}$ is a constant vector for each i,j boundary condition pair and $\delta\mathbf{z}$ is sampled from the boundary condition uncertainties such that $\delta\mathbf{x}_0 \sim \sum_{i=1}^{n} w_i \mathcal{N}(0, \mathbf{P}_i)$ and $\delta\mathbf{x}_f \sim \sum_{j=1}^{m} w_j \mathcal{N}(0, \mathbf{P}_j)$. Substituting the above relationships into Eq. (17), expanding, and introducing the functions $\boldsymbol{\omega}(t, t_0)$ and $\boldsymbol{\Omega}(t, t_0)$ from Eqns. (6) and (7), the quadratic control cost for a single term of the Gaussian mixture connecting Gaussian initial distribution i to Gaussian final distribution j is:

$$J_{ij} = J^* + \boldsymbol{\omega}(t_f, t_0)^{\mathrm{T}} \delta\mathbf{z} + 2\boldsymbol{\mu}_{ij}^{\mathrm{T}} \boldsymbol{\Omega}(t_f, t_0)\delta\mathbf{z} + \boldsymbol{\omega}(t_f, t_0)^{\mathrm{T}} \boldsymbol{\mu}_{ij} \tag{24}$$
$$+ \boldsymbol{\mu}_{ij}\boldsymbol{\Omega}(t_f, t_0)\boldsymbol{\mu}_{ij} + \delta z^{\mathrm{T}} \boldsymbol{\Omega}(t_f, t_0)\delta\mathbf{z}$$

The quadratic cost is dependent upon J^*, $\boldsymbol{\Omega}$, and $\boldsymbol{\omega}$ of the nominal optimal trajectory connecting the expected value initial and final states. Combining terms, the cost J_{ij} can be re-written in a format similar to the strictly Gaussian result from Holzinger et al.,[6] as shown in Eq. (25).

$$J_{ij} = J^* + \boldsymbol{\omega}(t_f, t_0)^{\mathrm{T}}(\boldsymbol{\mu}_{ij} + \delta z) + (\boldsymbol{\mu}_{ij} + \delta\mathbf{z})^{\mathrm{T}} \boldsymbol{\Omega}(t_f, t_0)(\boldsymbol{\mu}_{ij} + \delta\mathbf{z}) \tag{25}$$

From Holzinger et al.[6] Appendix B, the analytic first and second moments of J_{ij} are:

$$\mathbb{E}\left[J_{ij}\right] = \mu_{J,ij} = J^* + \boldsymbol{\omega}^{\mathrm{T}} \boldsymbol{\mu}_{ij} + \boldsymbol{\mu}_{ij}^{\mathrm{T}} \boldsymbol{\Omega}_{ij} \boldsymbol{\mu}_{ij} + \mathrm{Tr}\left[\boldsymbol{\Omega}_{ij} \mathbf{P}_{z,ij}\right] \tag{26}$$

$$\mathbb{E}\left[J_{ij}^2\right] = \sigma_{J,ij}^2 = \left(\boldsymbol{\omega}^{\mathrm{T}} + 2\boldsymbol{\mu}_{ij}^{\mathrm{T}}\boldsymbol{\Omega}\right)^{\mathrm{T}} \mathbf{P}_{z,ij} \left(\boldsymbol{\omega}^{\mathrm{T}} + 2\boldsymbol{\mu}_{ij}^{\mathrm{T}}\boldsymbol{\Omega}\right) + 2\mathrm{Tr}\left[\boldsymbol{\Omega}\mathbf{P}_{z,ij}\boldsymbol{\Omega}\mathbf{P}_{z,ij}\right] \tag{27}$$

Then, the total cost of all i initial boundary conditions and j final boundary conditions is the weighted sum of the individual costs between each i and j:

$$J = \sum_i \sum_j w_i w_j J_{ij} \tag{28}$$

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 bound-

ary condition may be written as:

$$\mathbb{E}[J] = \mathbb{E}\left[\sum_i \sum_j w_i w_j J_{ij}\right]$$

$$= \sum_i \sum_j w_i w_j \mathbb{E}[J_{ij}]$$

$$= \sum_i \sum_j w_i w_j \mu_{J,ij} \tag{29}$$

which leads to the final result:

$$\mu_J = \sum_i \sum_j w_i w_j \left(J^* + \boldsymbol{\omega}^{\mathrm{T}} \boldsymbol{\mu}_{ij} + \boldsymbol{\mu}_{ij}^{\mathrm{T}} \boldsymbol{\Omega} \boldsymbol{\mu}_{ij} + \mathrm{Tr}\left[\boldsymbol{\Omega} \mathbf{P}_{z,ij}\right]\right) \tag{30}$$

RESULTS

To validate the algorithm implementation, test cases are used to identify the limits and performance of the algorithm. For each case tested, a nominal UCT boundary condition uncertainty (covariance matrix) is parameterized using a scalar multiplier to examine the effects of uncertainty on anomaly detection. Similarly, to examine the effects of duration between UCTs, each scenario simulates various time intervals between UCTs of up to 48 hours. The uncertainty control cost distributions at the boundary conditions are computed and compared to the deterministic optimal maneuver control cost to test the maneuver (or anomaly) hypothesis. The results include a synthetic (simulated) scenario and real-world data. Since real-world operational data is available, the synthetic scenario is constructed to closely resemble that scenario for comparison purposes.

Synthetic Scenario: Inclination Correction

The first test case studied is a simulated inclination correction. For this purpose, a geostationary satellite inclination change scenario was constructed, similar to the real-world data available. In the scenario, UCTs are initialized with identical orbit elements except for the inclination and true anomaly. The initial inclination is set to 0.02 degrees, while the final inclination is set to 0 degrees, values chosen to reflect the available real data. The change in true anomaly is related to the duration of the simulation, which is another variable in the simulation. This set of boundary conditions allows for the simulation of an inclination-change station-keeping maneuver for varying observation gaps to assess the detectability of this maneuver.

A sensitivity study is conducted to generate confidence curves of probability of anomaly based on a) maneuver detectability ratio and b) UCT time separation. The boundary condition uncertainty for this particular scenario is assumed Gaussian and is initialized at 1 meter in position and 0.01 meters-per-second in velocity, as shown in Eq. (31):

$$\mathbf{P}_A = \mathbf{P}_B = \mathrm{diag}\left(\alpha^2, \quad \alpha^2, \quad \alpha^2, \quad (0.01\alpha)^2, \quad (0.01\alpha)^2, \quad (0.01\alpha)^2\right) \tag{31}$$

where α is the scalar that is varied between scenarios to change the uncertainty (and modulate the detectability ratio). The scalar is allowed to vary logarithmically from 1 to 10^3, a range determined by trial and error to cover a wide range of detectability ratio values (from near-zero to well above 9). This corresponds to maximum boundary condition position and velocity uncertainties of 10 kilometers and 100 meters-per-second, respectively.

Preliminary results from this scenario were used for validation of the algorithm framework, using a nominal scenario with the initial condition 30 minutes before the ascending node passage and the final condition 30 minutes after the ascending node passage. The probability of anomaly for this scenario using different detectability ratios is shown in Figure 5. As desired, the probability of anomaly at a detectability ratio of 1 is exactly 0.5, so there is an equal likelihood that the observed state change is caused by a control maneuver or uncertainty at this test case. As the detectability ratio increases above 1, the probability that a maneuver occurred increases, and opposite below a detectability ratio of 1.
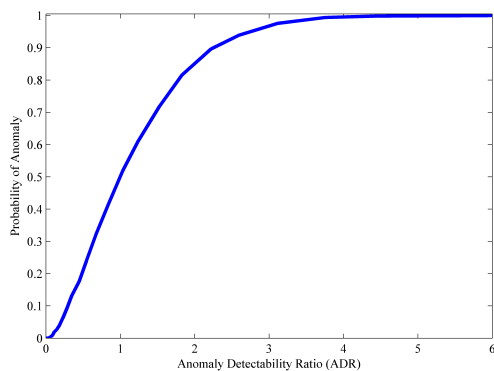
Figure 5. Probability of Anomaly vs Detectability Ratio, Synthetic Scenario (1 hour simulation duration)

To test the result validity, an upper-bound cost can be applied based on the quadratic cost from the maneuver optimization.[11] Additionally, the continuously-thrusting maneuver cost should be lower-bounded by the impulsive maneuver cost for an inclination change of 0.02 degrees. These sanity checks are shown in Table 1, and the trajectory (actual) cost is within the acceptable range.

Table 1. Simulated Maneuver Control Cost and Comparisons

Parameter	Value	Units
Quadratic Control Cost	$1.609022285 \times 10^{-4}$	$\frac{m^2}{s^3}$
Impulsive ΔV	1.0732610	$\frac{m}{s}$
Actual ΔV	1.0763310	$\frac{m}{s}$
Upper-bound ΔV	1.0763345	$\frac{m}{s}$

Having analyzed the performance of the algorithm using a nominal scenario, the next step is to examine the sensitivity of the algorithm to the time between UCT observations. To perform this sensitivity study, the initial and final conditions are varied to change simulation duration, keeping the ascending node passage at the midpoint of the simulation duration. The initial and final times are varied from 5 minutes before and after the ascending node (10 minute simulation duration) to 24 hours before and after the ascending node (48 hour simulation duration), and the initial and final states are adjusted accordingly. The trajectory optimizer is not confined to a specific time within the simulation time window to perform the correction maneuver, and results show that the optimal trajectories selected use the maximum thrust acceleration at the node passages. This is consistent with intuition since an impulsive inclination-change maneuver would be performed entirely at the
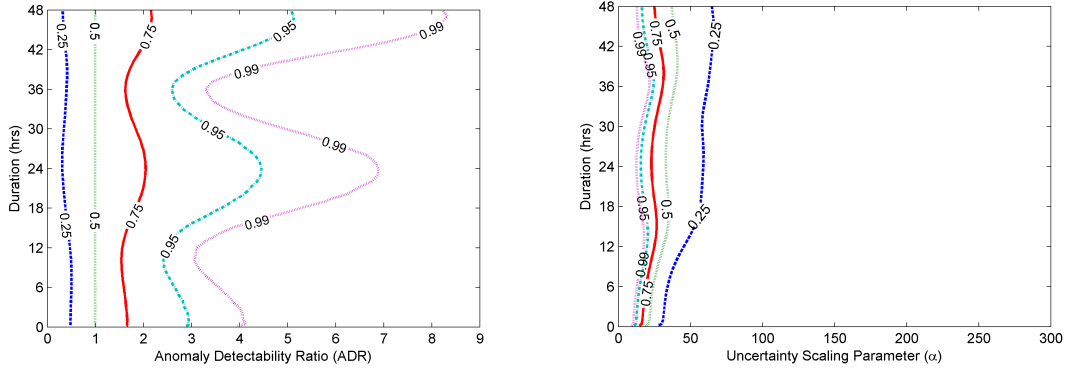
11

Figure 6. Anomaly detection probability vs Anomaly Detectability Ratio and Uncertainty Scaling Parameter, Synthetic Scenario

node passages, so an optimal low-thrust maneuver should maximize its thrusting at the nodes.

The results of this sensitivity study are distilled in the contour plot in Figure 6. On the left-hand side plot, detectability ratio and duration are varied on the x- and y-axes respectively, and the contours show lines of constant probability of anomaly, increasing from 0.25% to 0.99% from left to right. The contours shown represent different confidence levels (or probabilities of anomaly) that the state change is attributable to an anomaly such as a maneuver given the observation duration and uncertainties. The 0.5, or 50%, confidence level is a straight line at a detectability ratio of 1, as prescribed by the anomaly detection method. Otherwise, even moderate confidence levels are relatively consistent, though a periodic behavior can be seen. The improvement in maneuver detection at 12 and 36 hours coincides with the highest out-of-plane state difference between the final observation and propagated state estimate. Conversely, at durations of an orbit period, maneuver detection requires lower state uncertainty because the boundary condition states are more similar.

The same sensitivity study data was plotted using α, the uncertainty scaling parameter, as the x-axis variable, shown in the right-hand side of Figure 6. Viewed as a function of uncertainty scaling parameter, the control distance metric better shows its consistency with respect to duration, especially in comparison to the Mahalanobis distance results presented later. The difference in the two contours of Figure 6 reveals an issue in viewing the results as a function of ADR, since ADR is a function of both minimum detectable control distance and optimal control distance. The plot using uncertainty scaling parameter avoids conflating these variables and emphasizes the duration-independence of the control distance metric.

This sensitivity study emphasizes an advantage of the control metric approach, which is that it performs relatively consistently with respect to observation duration. Some previous knowledge of the orbit is still useful in determining the validity of the maneuver detection probability results. For instance, when attempting to detect inclination change maneuvers, observations spaced at orbit periods are not ideal because this could yield a small state difference and require lower observation uncertainty to detect maneuvers with high likelihood. Another, more obvious, result is that larger maneuvers, which yield a higher detectability ratio for the same boundary condition uncertainty, are easier to detect. This is particularly relevant for maneuvers by spacecraft with variable specific-impulse engines, where the state changes are small and spread out over a long time. Conversely, a large state-change by a nearly-impulsive maneuver could still be detected even in a short time between observations, if the control required for the maneuver significantly exceeds the state change

in the homogeneous orbit based on boundary condition uncertainty.

Real-World Data: Inclination Correction

   To complement the simulated scenario, the algorithm is tested using real operational data, the availability of which drove the construction of the simulated 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 Deveiation Corrections (DCs). The DCs are then broadcast by GPS satellites to improve position accuracy calculations for WAAS-enabled GPS receivers.[12]
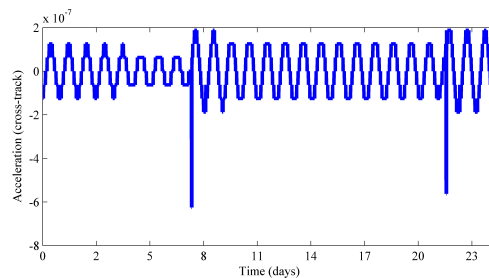

Figure 7. Cross-track Acceleration for Geosynchronous satellite, Real Data

   Figure 7 shows the cross-track acceleration data for the real dataset. Inspection of the acceleration data reveals two large anomalous cross-track acceleration events, candidates for a North-South (inclination) station-keeping maneuver, during days 7 and 22. The selected maneuver (the peak around day 7) resulted in a 0.03 degree inclination change, and simulation initial and final conditions were chosen to surround this maneuver in equal time increments around the maneuver. The real-world data is analyzed in a similar manner to the synthetic data: by varying boundary condition uncertainty and time, a sensitivity study is conducted to construct probability of anomaly confidence curves with respect to detectability ratio and simulation duration. An initial boundary condition is selected from the data given a desired time before the candidate maneuver, Therefore, the study begins with a simulation duration of 1 hour and increases to a duration of 48 hours, spaced evenly around the candidate maneuver.

   The results of this sensitivity study are shown in Figure 8, plotted against anomaly detectability ratio on the left and uncertainty scaling parameter on the right. The confidence-level contours agree closely with with Figure 6 from the synthetic scenario. In testing, when many of the perturbing accelerations were deselected, the algorithm responded by predicting higher probabilities of anomaly at lower uncertainties, since the real spacecraft deviated from the unperturbed propagation quickly. Therefore, it is important to model the relevant dynamics perturbations for a given test case. In this case, agreement between the simulated and operational data cases lends confidence that the applicable perturbations have been modeled in the trajectory optimizer, and reinforces the applicability of this algorithm to real data in operational use. Additionally, since propagated orbital mechanics uncertainties often become non-Gaussian quickly, the Gaussian mixture approximation methods presented in this paper provide further computationally feasible applicability to real data.
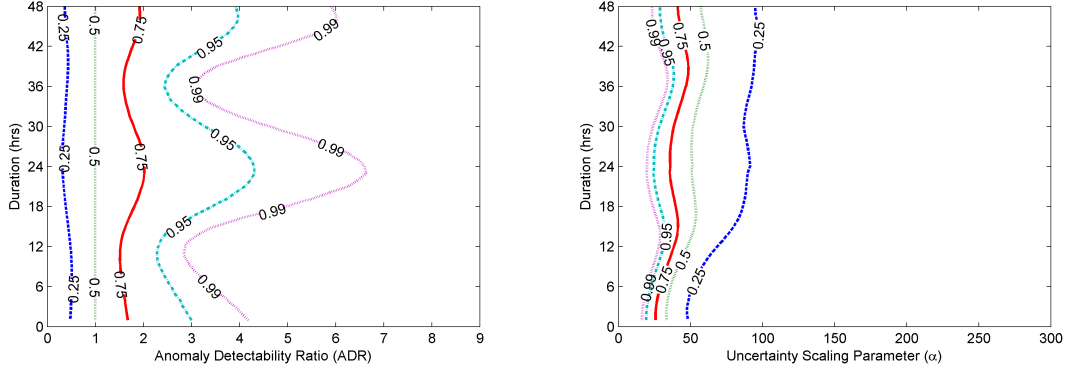
Figure 8. Anomaly detection probability vs Anomaly Detectability Ratio and Uncertainty Scaling Parameter, Real Data

Distance Metric Comparisons

The proof of control performance as a metric has been shown by proving the properties of positivity, strict positivity, symmetry, and triangle inequality.[6] In order to assess the effectiveness of the minimum-fuel control cost as a distance metric, it is important to consider comparisons to other common distance metrics. Primarily, it is important to note the differences between control cost and both geometric and multivariate distribution distance metrics, and the advantage that a control cost metric lends in orbital mechanics applications. This section and the included analysis address the second stated contribution for this paper.

The most simplistic geometric distance metric for evaluating the difference between two orbital state vectors is Euclidean distance. This considers only the position difference between two different states without consideration for velocity differences or uncertainty in the boundary conditions. The Euclidean distance, Eq. (32), is calculated by taking the 2-norm of the position difference:

$$d_E(\mathbf{x}_A, \mathbf{x}_B, t_B) = \| \begin{bmatrix} \mathbf{I}_{3\times3} & \mathbf{0}_{3\times3} \end{bmatrix} (\mathbf{x}_B - \mathbf{x}_A(t_B)) \|_2 \tag{32}$$

where $\mathbf{x}_B$ is the state vector of the final UCT and $\mathbf{x}_A(t_B)$ is the initial UCT state propagated to time t_B. The quiescent propagated state $\mathbf{x}_A(t_B)$ is calculated by taking the initial condition UCT and propagating it under the assumption that no anomaly has occurred. In Euclidean distance, only the positional difference is considered in relating the two states. For observation association in the space environment, this is not a good method since even a small control input can cause a large positional diversion over time. For instance, a spacecraft that performs a plane change maneuver at the ascending node will have a small positional difference between the perturbed and homogeneous orbits at first, but after a quarter of an orbit period the out-of-plane displacement will be significant. Furthermore, the geometric distance metrics break down as an association metric when the position difference is small because it does not account for the velocity difference. Even RSOs in very dissimilar orbits (such as a closed orbit and a fly-by trajectory) can reside in close proximity for a short time, so simply establishing a position threshold using Euclidean distance is not sufficient.

Alternate metrics exist that measure the distance between multivariate normal distributions, such as an orbital state vector, and can take into account both position and velocity information as well as uncertainty in the state vectors. This is ideal in defining a metric to discriminate between two orbital state observations. One example of such a metric is the Mahalanobis distance, which effectively applies a weight to the state difference along each axis based on the combined uncertainty along
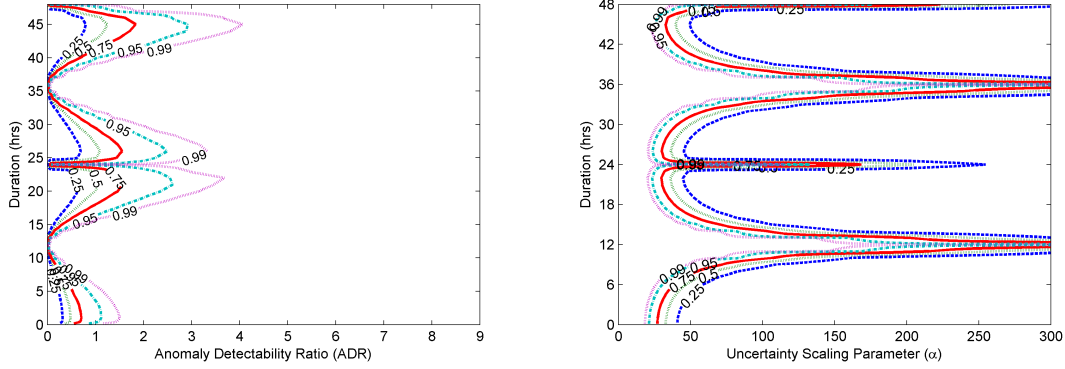
14

Figure 9. Mahalanobis anomaly detection probability vs Anomaly Detectability Ratio and Uncertainty Scaling Parameter, Synthetic Scenario

that axis in the observations.[13] The expression for the Mahalanobis distance is given in Eq. (33):

$$d_M(\mathbf{x}_A, \mathbf{P}_A, \mathbf{x}_B, \mathbf{P}_B, t_B) = \sqrt{(\mathbf{x}_A(t_B) - \mathbf{x}_B)^{\mathrm{T}}(\mathbf{P}_A(t_B) + \mathbf{P}_B)^{-1}(\mathbf{x}_A(t_B) - \mathbf{x}_B)} \qquad (33)$$

where $\mathbf{x}_B$ and $\mathbf{P}_B$ are the full state and covariance information of the state estimate from the final UCT, and $\mathbf{x}_A(t_B)$ and $\mathbf{P}_A(t_B)$ are the state and covariance of the quiescent propagated state. Similar to Euclidean distance, the propagated homogenous state is calculated by taking the initial condition UCT and propagating it under the assumption that no maneuver has been performed, and the propagated covariance is calculated using the propagated state transition matrix to update the initial covariance. Mahalanobis distance improves upon Euclidean distance by accounting for uncertainty. Low values of the Mahalanobis distance yield a higher confidence in correlation of the full UCT, accounting for both the full state (position and velocity) and state uncertainty. Since the Mahalanobis distance metric is sensitive to deviations from expected propagation, if the dynamics are not well-modeled the method will return abnormally high numbers of uncorrelated tracks. Additionally, Mahalanobis distance does not characterize maneuvers; it simply recognizes differences in expected propagation.

In developing alternate distance metric results, the same data from the synthetic scenario sensitivity study are used to allow comparisons to the minimum-fuel control distance metric (Figure 6). In order to facilitate comparisons, the Mahalanobis distance value is mapped to the same confidence intervals. Typically, Mahalanobis distance provides representation of the dissimilarity between measurement sets and their associated covariances, analogous to standard deviations. As the dimensionality of the problem increases though, Mahalanobis distance diverges from the typical standard deviation confidence intervals for one-dimensional data (e.g. 1-sigma is 68% confidence, 2-sigma is 95% confidence, etc.). However, relationships exist to convert Mahalanobis distance to a confidence interval in n-dimensional space,[14] so the Mahalanobis distance values were converted to applicable confidence intervals to match the control distance metric analysis. The Mahalanobis distance metric plot is shown in Figure 9, presented in a similar manner as Figure 6 for the minimum-fuel metric.

At durations approaching an orbit period, the Mahalanobis distance decreases, showing that the UCTs are more closely correlated, with a higher likelihood that state uncertainty alone explains the observation difference. Consequently, a maneuver is harder to detect at these points. However, at the 24-hour duration, where the initial and final states are both at the ascending node crossing, the UCTs are no longer well correlated, likely due to the small difference in position but large difference in velocity required for the different inclination between the boundary conditions. Conversely, at the

15

12- and 36-hour durations, where the out-of-plane displacement is at its maximum, the Mahalanobis distance is considerably higher even with high uncertainties, showing that the observed state change likely requires a maneuver or other anomaly to correlate the observations. Once again, this is due to the inherent state differencing in the metric, since the propagated initial state would have the maximal out-of-plane displacement at these points. These trends are similar to those seen in Figure 6 for the control distance metric.

Since the left-hand side of Figures 6 and 9 both plot detectability ratio and duration against probabilities of anomaly, there is a one-to-one correlation between points on these plots. Initial impressions from point-to-point comparisons between the two figures show that Mahalanobis distance is able to detect the anomaly at lower detectability ratios, but is less consistent with respect to observation gap. However, recall that detectability ratio is computed using both the optimal control and minimum detectable control distance from the minimum-fuel control distance algorithm. Since neither of these parameters are independent of the maneuver detection algorithm, plotting Mahalanobis distance against detectability ratio is a convolution of the two techniques, which does not lend well to independent comparisons. Therefore, while detectability ratio is useful for the visualization of results in the minimum-fuel maneuver detection algorithm, it is more desirable to plot the minimum-fuel and Mahalanobis algorithm results against the independently varied uncertainty scalar value (α). The right-hand side of Figures 6 and 9 show the results using the uncertainty scaling parameter, providing a cleaner means of comparing the two methods. This time, as α is increasing along the x-axis, probabilities of anomaly decrease from left to right. As uncertainty increases, the probability of anomaly decreases for both methods, as expected. The trends for Mahalanobis distance are maintained in this new representation: it is more confident of an anomaly than the minimum-fuel metric at lower uncertainties, but is not as consistent with time. This is important because it makes the development of thresholds for detection much more difficult without prior knowledge of the orbit, making Mahalanobis distance more difficult to apply to un-queued UCT correlation and maneuver detection. In comparison, the minimum-fuel control distance metric is relatively consistent with respect to observation gaps, meaning thresholds can be set for maneuver detection and be used with higher confidence without prior knowledge of the orbit. The thresholds required for minimum-fuel control distance, however, are much more consistent and would allow simpler application to un-queued UCT correlation and maneuver detection.

A summary of the comparisons between Euclidean distance, Mahalanobis distance, and minimum-fuel control distance metrics can be seen in Table 2. The observed importance of observation cadence and concerns over the Mahalanobis distance's high sensitivity to outliers gave rise to questions regarding the relative rates of false-positives (type I error) and false-negatives (type II error). The concern is that flagging even 2-sigma state differences as anomalies with high-probability using Mahalanobis distance could contribute to a high rate of false-positives: reporting an anomaly when one didn't actually occur. This motivates future work into the computation of error rates for both Mahalanobis and minimum-fuel control distance. Of primary concern is minimizing the rate of false-negatives to ensure that maneuvers are not missed, which will be addressed in future work.

Gaussian Mixture Approximations

To validate the application of Gaussian mixtures to the algorithm, Gaussian mixture initial and final boundary conditions were applied to a synthetic inclination correction scenario. This synthetic scenario covers a duration of 1 hour between UCT times, centered around the ascending node. In this case, though, there are two different initial condition orbits: one with an inclination of 0.01

Table 2. Comparison of distance metric performance factors

	Euclidean	Mahalanobis	Minimum Fuel
Input	$\mathbf{r}_A, \mathbf{r}_B$	$\mathbf{x}_A, t_A, \mathbf{P}_A, \mathbf{x}_B, t_B, \mathbf{P}_B$	$\mathbf{x}_A, t_A, \mathbf{P}_A, \mathbf{x}_B, t_B, \mathbf{P}_B$
Detection	Poor	Very sensitive to deviation from predicted dynamics	Sensitive
Timing	Poor	Cadence important	Most consistent
Uncertainty	None allowed	Uncertainty incorporated	Uncertainty incorporated

degrees and another with an inclination of 0.02 degrees. The corresponding final states both have an inclination of 0 degrees, prescribing an inclination change maneuver between the UCTs. The two initial and two final states are each given an equal weighting of 0.5. Similar to the previous simulations, the position and velocity uncertainties in the initial and final boundary conditions are initialized to 1 meter in 0.01 meters-per-second, respectively, and the uncertainty is scaled by the parameter α. This particular Gaussian mixture scenario was constructed with the intention of developing a sufficiently non-Gaussian control cost distribution. Through trial and error, an uncertainty scaling parameter of $\alpha = 10$ was found to generate a bimodal uncertainty cost distribution.

The analytial approximation promises a computationally tractable method for addressing non-Gaussian boundary conditions, but it still needs to provide an accurate reconstruction of the uncertainty cost distribution. Therefore, multiple methods for generating the uncertainty cost distribution were used for comparison. The simplest cost distribution generation method is a non-linearized direct sampling of the boundary conditions, referred to as the validation Monte-Carlo sampling. For each run, a state is selected at random from the initial and final Gaussian mixture boundary conditions. The trajectory optimizer from the algorithm is then used to generate an optimal trajectory cost between the chosen boundary conditions, and the resulting control cost is collected before selecting another pair of boundary conditions. This process is repeated a user-selected number of times to sample the control cost space and build up control cost distribution. 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. The second method takes advantage of the individual Gaussian distributions that make up the Gaussian mixture boundary conditions to perform the minimum-fuel metric algorithm on all possible combinations of the boundary conditions. This method requires performing the trajectory optimization on each combination of Gaussians, meaning in the synthetic case constructed here (2 initial condition Gaussians and 2 final condition Gaussians), 4 trajectory optimizations were required to form the 4 connecting optimal trajectories. Then, each of these connecting trajectories are sampled using the results of the minimum-fuel metric algorithm, linearizing about the trajectory and sampling from the Gaussian boundary conditions to develop control cost distributions. These individual control cost distributions are summed together to generate the final control cost distribution representing the transfer between Gaussian mixture boundary conditions. This method makes simplifying assumptions to linearize about the connecting trajectories, but makes substantial computational improvement over the non-linearized approach. The third method utilizes the full results from the Gaussian mixture approximation implementation derived in this paper. The Gaussian mixture boundary conditions are used to generate a mean initial and mean final state, which is used for the trajectory optimization. The control cost distribution is then developed by linearizing about this connecting trajectory and utilizing Eqns. (24) and (28). This

represents the most simplified version, but offers much improved computation time over the other methods since only one trajectory optimization is required.

Figure 10 shows the uncertainty control cost distributions (both PDFs and CDFs) resulting from this study, and Table 3 shows the differences in computation time, along with the optimal cost for this maneuver. The control costs themselves show high levels of agreement between the non-linearized propagation and simplified methods, with the CDFs being nearly identical. Therefore, as can be inferred from the CDFs, the probability of anomaly would be nearly identical using any of these three methods, yielding a roughly 80% probability of anomaly. The real difference in implementations is realized in the timing results. These results are in-line with expectation, as the trajectory optimizer is expected to be the computational bottleneck. For n_{samp} samples (e.g. $n_{samp} = 10,000$), the non-linearized validation method should complete in roughly $\mathcal{O}(n_{samp})$ time (small variations in the time required for each sample based on the boundary conditions). In comparison, for n initial condition Gaussians and m final condition Gaussians, the full Gaussian mixture method is expected to require $\mathcal{O}(nm)$ time. Finally, the Gaussian mixture approximation method only requires a single trajectory optimization, and therefore should run in nearly constant time, or $\mathcal{O}(c)$. From Table 3, the non-linearized method required $11689\times$ more computation time than the analytical Gaussian mixture approximation. Note that, because the non-linearized method requires significant computation time, this timing result was extrapolated from 1000 samples to match the 10,000 samples of the other methods. The full Gaussian mixture method was $3.3\times$ slower, though still runs very fast with many more samples than the non-linearized approach. However, it is apparent that the Gaussian mixture approximation is able to achieve very good accuracy while significantly improving computation time over the other methods. The computer used for this test is an i5-2500K clocked at 4.2 GHz with 16 GB RAM running 64-bit MATLAB 2013a. It is important to note that all these tests were run with MATLAB, so even more computational speed-up should be realizable through implementation in C or some other compiled code, allowing for real-time use.
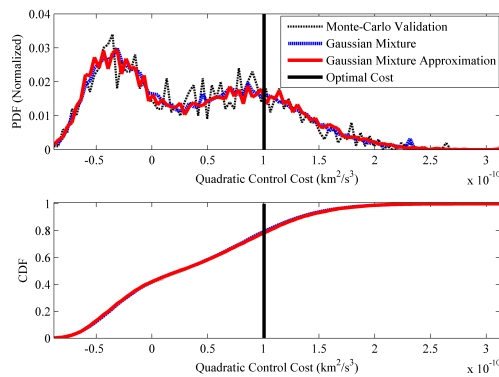

Figure 10. Comparison of cost distributions

Table 3. Time required to generate cost distributions ($n_{samp} = 10,000$, n = m = 2)

	MC Validation*	GM	GM Approximation
Computational Complexity	$\mathcal{O}(n_{samp})$	$\mathcal{O}(nm)$	$\mathcal{O}(c)$
Run Time (s)	10497.28	2.95391	0.898014

In orbital mechanics applications, Gaussian initial conditions are not typically preserved after

propagation, so representing non-Gaussian distributions consistent with orbital mechanics is an important task. Gaussian mixtures allow the representation of these non-Gaussian boundary conditions, and these results show the accuracy achievable in the simplified analytical approach to generating control cost distributions from a Gaussian mixture. When considering the computational complexity improvements, the Gaussian mixture approximation method significantly outperforms the non-linearized approach while still maintaining a similar cost distribution and therefore performing similarly in anomaly detection. However, this approximation method is expected to have limitations when dealing with vast uncertainty distributions. For instance, if the perturbation terms from the mean Gaussian mixture state ($\boldsymbol{\mu}_{ij}$) terms are large, linearizing about the mean trajectory may no longer be valid.

CONCLUSIONS

In this study, an algorithm was developed to address maneuver detection and UCT correlation using a minimum-fuel metric based on work by Holzinger et al.[6] The algorithm was evaluated using synthetic inclination-change scenarios, validating its ability to detect maneuvers, and the results were compared to real-world data. In both cases, the algorithm was able to detect the maneuvers. The probability of anomaly was compared to other distance metrics, primarily the Mahalanobis distance metric. The minimum-fuel control distance method showed less sensitivity to anomalies than Mahalanobis distance, requiring a higher detectability ratio, but is more consistent with simulation time and also has the added benefit of reconstructing the maneuver for maneuver characterization. Finally, an analytical treatment of non-Gaussian boundary conditions using Gaussian mixtures was derived and examined with a similar simulated scenario. The method was able to accurately reproduce non-Gaussian control cost distributions consistent with validation data generated through a non-linearized Monte-Carlo simulation.

The distance metric comparison motivates future work in the analysis of error rates for both control distance and Mahalanobis distance metrics. Given the high sensitivity of Mahalanobis distance to outliers, it may be causing a high rate of false-positives, preventing association of otherwise associated UCTs. The true concern though is to determine the rates of type-II errors, which are false-negatives or missed maneuvers. Future work will endeavor to understand what threshold should be set for control distance to achieve a prescribed, low error rate in false-negatives to ensure that the method can detect most maneuvers, and will compare that to Mahalanobis distance error rates. Additional future work includes testing the Gaussian mixture approximation using more non-Gaussian boundary conditions to determine its limitations.

The authors would also like to acknowledge Dr. Kyle DeMars for enlightening discussions regarding Gaussian mixtures.

REFERENCES

[1] L. James, "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] "The Space Report 2011," *Space Foundation*, 2011.

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

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

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

[6] M. J. Holzinger, D. J. Scheeres, and K. T. Alfriend, "Object Correlation, Maneuver Detection, and Characterization Using Control-Distance Metrics," *Journal of Guidance, Control, and Dynamics*, Vol. 35, July 2012, pp. 1312–1325.

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

[8] M. J. Holzinger, D. J. Scheeres, and K. T. Alfriend, "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.

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

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

[11] E. D. Gustafson and D. J. Scheeres, "Optimal Timing of Control-Law Updates for Unstable Systems with Continuous Control," *Journal of Guidance, Control, and Dynamics*, Vol. 32, May 2009, pp. 878–887.

[12] "Federal Radionavigation Plan," tech. rep., Department of Defense, Department of Homeland Security, Department of Transportation, 2005.

[13] K. T. Abou-Moustafa, F. d. l. Torre, and F. P. Ferrie, "Designing a Metric for the Difference Between Gaussian Densities," *Brain, Body and Machine: Proceedings of an International Symposium on the Occasion of the 25th Anniversary of the McGill University Centre for Intelligent Machines*, Vol. 83, 2010, pp. 57–70.

[14] G. Gallego, C. Cuevas, R. Mohedano, and N. Garcia, "On the Mahalanobis Distance Classification Criterion for Multidimensional Normal Distributions," *IEEE Transactions on Signal Processing*, Vol. 61, June 2013, pp. 4387–4396.