

TWO METHODS FOR LIGHT CURVE INVERSION FOR SPACE OBJECT ATTITUDE DETERMINATION

A. Burton and C. Frueh

Purdue University, School of Aeronautics and Astronautics, 710 West Stadium Ave, West Lafayette, United States, Email: {burton30, cfrueh}@purdue.edu

ABSTRACT

Information such as the shape, attitude, and rotation of space objects is not always available through resolved observation. Two methods are outlined for estimating the attitude and rotation of an object with a known shape and orbit using only light curve measurements. In the first method, multiple attitude time-histories based on the possible orientations corresponding to each measurement in the light curve are found. Each time history is then given a weight based on how well it matches a chosen dynamics model. The second method uses a Probability Hypothesis Density filter to find the most likely object state at the end of the measurement time. Both methods are applied to a simply-rotating object, and the second is applied to an object undergoing torqueless rotation about multiple axes.

1. INTRODUCTION

Space situational awareness (SSA) involves the tracking and characterization of resident space objects (RSOs) in orbit around the Earth. Much of this work focuses on establishing the orbital characteristics of RSOs, but shape and attitude information are also required for complete identification and characterization. In many cases, it is impossible to obtain a resolved image either due to distance or the small size of the object being observed. In those situations, the light curve, or the observed brightness over time, of the object may be analyzed instead. Light curve inversion is the process of retrieving shape and attitude information from the light curve measurements.

Light curve inversion has primarily been used to estimate the shapes and rotational characteristics of natural bodies such as asteroids. The method was first discussed in 1906 [28], and by the late 20th century, various methods of estimating an asteroid's shape and rotation axis had been developed [22, 16]. However, these methods generally assumed that the asteroid's shape was smooth and approximately convex, that its reflective properties were uniform, and that it was rotating about a single axis

[22, 17, 18, 35]. Some methods to estimate highly non-convex shapes have been developed, but they have limited application due to the high phase angles required for concavities to be observable [9].

None of the simplifying assumptions hold for artificial RSOs, making the inversion problem significantly more difficult [15]. Artificial satellites do not typically have smooth shapes, and the presence of features such as antennae means that they cannot be easily approximated as convex. Furthermore, a satellite's reflective properties can vary considerably from part to part, and specular glints also need to be taken into account [15]. Finally, while asteroids typically rotate about a single axis, Earth-orbiting objects can have more complex rotational behavior due to perturbing gravitational torques and active attitude control systems [15].

Some methods have been developed for extracting shape and attitude information from the light curves of artificial RSOs [19, 20]. Because of these cases' complexity, the highly-coupled shape and attitude estimation problems are often treated separately by assuming that one of the two sets of information is already known [15, 19, 4]. This is often a reasonable assumption since, for artificial objects, partial information about either shape or rotation is frequently available *a priori*.

This paper will focus on the attitude inversion problem for an RSO whose shape, reflective properties, and orbit are already known. There are several existing methods for solving this problem; however, they often require that the object be a specific shape [32, 33] or assume that rotation is torque-free over short intervals [33, 1]. Some methods are more general but still have limitations. An Unscented Kalman filter has been used to produce accurate attitude estimates using simulated light curves, but it requires a good angular velocity guess to converge correctly [7, 31].

2. ALGORITHM OVERVIEW

The flow chart in Figure 1 outlines a proposed method to solve the general problem of light curve inversion for attitude. A set of feasible orientations for each measurement

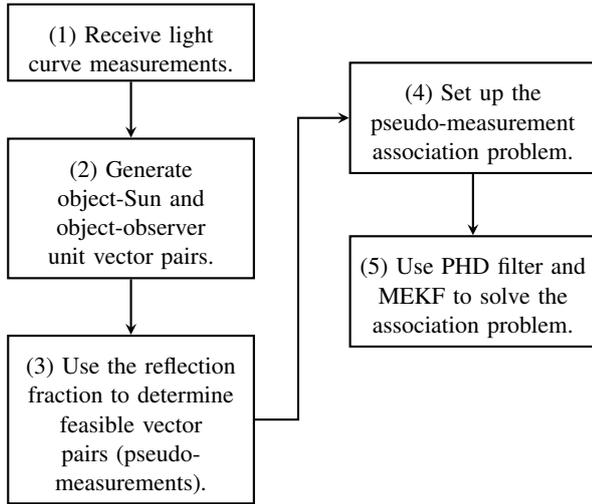


Figure 1. A flowchart for the light curve inversion method based on the chart in Burton and Frueh [3].

in the light curve are generated from the aspect graph and viewing sphere [13, 14]. These feasible orientations are also referred to as pseudo-measurements. Once the sets of pseudo-measurements at each measurement time are found, they are used to produce attitude time-histories. This is referred to as the “pseudo-measurement association problem” since it involves finding which pseudo-measurements can feasibly be associated in a sequence. In other words, the association problem is to find which of the feasible orientations can follow after each other given the rotational dynamics.

Two different methods for solving the association problem will be tested. In the first, a gating method is used to narrow down the range of possible pseudo-measurement sequences before finding the attitude time-histories. Each time-history is given a weight based on how well it fits its sequence of pseudo-measurements. In the second method, all possible sequences of pseudo-measurements are computed and weighted in parallel using a probability hypothesis density (PHD) filter. To avoid exponential growth in the number of estimates, low-weight estimates are dynamically pruned, and similar estimates are merged. The most likely estimates after the final measurement update are back-propagated to produce time-histories. A more detailed discussion of the method will elaborate on each step in the Figure 1 flowchart.

This algorithm has similarities to the multi-hypothesis tracking (MHT) method, most notably its use of multiple possible time histories with weights based on relative likelihood [27, 2]. However, there are also significant differences. The most fundamental is that each measurement is assigned to only one target in the classical MHT [8, 34]. That is a reasonable assumption for the multi-target tracking problem but does not apply to light curve inversion since there is nothing to stop multiple purely hypothetical time histories from having the same attitude at the same time. This means that many of the common measurement-assignment tools for MHT will not work

when choosing pseudo-measurements in the inversion problem. Therefore the two different methods for selecting pseudo-measurements, both the gating method and the PHD filter method, use probabilistic weights rather than hard pseudo-measurement assignments. In addition, although the PHD filter is intended for multi-target tracking, it is not a true multi-hypothesis method since it uses a probability distribution to represent possible state estimates rather than having discrete hypotheses.

3. LIGHT CURVE SIMULATION

The reflection model and method of generating simulated light curves are discussed before the inversion problem proper is addressed. For simulated light curves, this corresponds to Block (1) in Figure 1. This material was previously presented at the 2020 Astrodynamics Specialist Conference [3].

The object of interest is modeled as a set of flat surfaces or faces [10]. For now, it is assumed that only Lambertian reflections contribute to the light curve measurements. The intensity of the observed light from this sort of reflection is taken from Frueh [11]:

$$I_L = \frac{I_0}{\pi r_{\text{obs}}^2} \sum_{i=1}^N B_i(\hat{u}, \hat{s}) A_i C_{d,i}$$

$$B_i = \begin{cases} (\hat{u} \cdot \hat{n}_i)(\hat{s} \cdot \hat{n}_i) & \text{if } (\hat{u} \cdot \hat{n}_i) > 0 \text{ and } (\hat{s} \cdot \hat{n}_i) > 0 \\ 0 & \text{otherwise} \end{cases}, \quad (1)$$

where I_0 is the steradian integrated irradiation of the light from the Sun at the object, approximated as the solar constant; N is the total number of object faces; \hat{u} and \hat{s} are the object-observer and object-Sun direction unit vectors; A_i and $C_{d,i}$ are the area and reflection coefficient of the i^{th} face; and \hat{n}_i is the face’s normal vector. If the observer or the Sun is below any face, that face does not contribute to I_L . This effect is accounted for by the B_i function.

To find feasible orientations for each measurement, the reflection fraction γ will be used to represent the effect of a particular orientation on the observed brightness. The reflection fraction is defined as:

$$\gamma = \sum_{i=1}^N B_i(\hat{u}, \hat{s}) A_i C_{d,i}, \quad (2)$$

which only depends on the shape and reflective properties of the object of interest, and on the object-observer and object-Sun unit vectors. These vectors are already known in the space-fixed frame, so if they are known in an object-fixed frame, the object’s orientation can be easily found. Therefore, since each $\hat{u} - \hat{s}$ vector pair in the

body frame corresponds to a unique orientation, the reflection fraction γ only depends on the object's properties and its orientation. The reflection fraction required for a particular measurement, γ_M , can be found by solving Equation (1) for the reflection fraction:

$$\gamma_M = \frac{I_L \pi r_{\text{obs}}^2}{I_0}. \quad (3)$$

The reflection fraction makes finding orientations that could feasibly correspond to a given observation more efficient by providing a measure of how orientation affects the observed brightness that does not need to be recomputed for each new observation.

4. DETERMINING FEASIBLE PSEUDO-MEASUREMENTS

Once the light curve measurements, whether real or simulated, have been received, the sets of orientations that could feasibly correspond to each measurement are found. This corresponds to Block (2) through Block (3) in Figure 1. Much of the material in this section was previously presented at the 2020 Astrodynamics Specialist Conference [3].

4.1. Viewing Sphere

One of the primary difficulties with the light curve inversion problem is its inherent ambiguity. There are many different orientations for any diffuse measurement that could all produce that exact measurement, and there are multiple attitude-time histories that could have produced any individual light curve. Before determining the feasible orientations at each measurement time, candidate orientations are generated along with their corresponding reflection fractions using the viewing sphere. These γ values will be compared to the measurement γ_M 's to find the set of feasible orientations.

The viewing sphere is a conceptual construct that helps determine which faces of the object are visible to the observer and the Sun. It is a unit sphere with the object of interest at the origin. Since the Sun and observer will be very far away compared to the size of most RSOs, the object of interest can be approximated as a point [13]. The observer and the Sun are on the viewing sphere at the locations specified by \hat{u} and \hat{s} , respectively. A viewing event occurs whenever one face of the object goes from being visible to hidden for either the observer or the Sun, or vice versa. A face is said to be visible from one of these points of view if:

$$\begin{aligned} \hat{u} \cdot \hat{n}_i &> 0 \text{ or} \\ \hat{s} \cdot \hat{n}_i &> 0, \end{aligned} \quad (4)$$

where \hat{n}_i is the normal vector of the i^{th} face. For convex objects, viewing events always occur at a point on the viewing sphere perpendicular to an \hat{n}_i , i.e., on a great circle.

The great circles corresponding to each of the object faces divide the viewing sphere into three types of regions: vertices, where two or more circles intersect; edges, arcs of the circles between vertices; and areas bounded on all sides by edges. Because no viewing event occurs inside one of these divisions, the same set of sides is visible from every point in a region. A face must be visible to both the observer and the Sun to contribute to the measured intensity I_L . A combination of faces that could be both visible to the observer and illuminated by the Sun for a given phase angle is called a visibility group. The viewing sphere will be used to numerically generate $\hat{u} - \hat{s}$ pairs in such a way that every visibility group is represented by at least one pair. Skipping a visibility group could mean missing a unique grouping of orientations. Without a tool like the viewing sphere, a search would possibly miss some visibility groups, and it would be impossible to tell whether or not they had all been found.

Orientations, represented by $\hat{u} - \hat{s}$ vector pairs, are generated as follows. First, at least one \hat{u} vector is placed in every region on the viewing sphere. A vector is placed directly on each vertex, and they are evenly distributed over each edge and area. Then, a circle with angular radius δ is drawn around each \hat{u} vector, and a paired \hat{s} vector is placed at each point where the drawn circle intersects with a viewing event. This guarantees that one \hat{s} vector is placed in each vertex and edge that is δ away from the original \hat{u} vector. Then an extra \hat{s} vector is placed in the area between each intersection so that every visibility group is accounted for. Finally, additional \hat{s} vectors are created a distance δ away from \hat{u} in order to have more data points. The only way a visibility group might be missed by this method is if two area regions are separated by slightly less than δ at their closest point, which is not the case for any of the objects used for the results of this paper.

Note that since Equation (1) behaves identically if \hat{u} and \hat{s} are switched, each stored $\hat{u} - \hat{s}$ vector pair actually represents two different orientations. Generating the unit vector pairs finishes the process corresponding to Block (2) in the Figure 1 flowchart.

4.2. Finding Feasible Pseudo-Measurements

Once the set of orientations has been generated, it is necessary to find which orientations could feasibly correspond to particular measurements. This is done by comparing the reflection fraction, γ , for the generated vector pairs to the γ_M of the intensity measurements, as mentioned in Block (3) of Figure 1. Since the vector pairs are generated numerically, there is no guarantee that any γ is going to match γ_M exactly. To account for this, any orientation with a γ value within some range Γ of the

measurement's reflection coefficient is counted as feasible.

These feasible orientations are called ‘‘pseudo-measurements’’ since they are treated as measurements by the solver. One pseudo-measurement must be chosen at each time step to estimate an attitude time-history using a state estimator. The problem of finding feasible time histories using sequences of associated pseudo-measurements is referred to as the ‘‘pseudo-measurement association problem.’’

5. MULTIPLICATIVE EXTENDED KALMAN FILTER

Before discussing two methods of solving the pseudo-measurement association problem, the Multiplicative Extended Kalman filter (MEKF) will be discussed. This filter is used by both methods to estimate rotational states from the pseudo-measurements [6, 26].

5.1. State Estimation

The estimated states at each time step are the orientation error, represented by three rotation angles $\delta\alpha_k$, and the angular velocity, $\hat{\omega}$:

$$\delta\hat{x}_k = [\delta\hat{\alpha}_k^T \quad \hat{\omega}^T]^T = [\delta\alpha_x \quad \delta\alpha_y \quad \delta\alpha_z \quad \omega_x \quad \omega_y \quad \omega_z]^T. \quad (5)$$

A quaternion \hat{q} is used to represent the estimated or ‘‘reference’’ attitude of the object. At each measurement update the error in \hat{q} is estimated in terms of three small rotation angles $\delta\hat{\alpha}$. During the update step, the quaternion only stays at unit length to the first order, so it must be normalized after each update [25, 6].

The angular velocity estimate $\hat{\omega}$ also needs to be found. This is done using a modified form of the MEKF described by Crassidis and Junkins [6]. Their MEKF uses vector measurements to estimate orientation and the bias in gyroscope angular velocity measurements. Estimating angular velocity instead results in slightly different covariance propagation equations:

$$\dot{P} = FP + PF^T + GQG^T \quad (6)$$

$$F = \begin{bmatrix} -[\hat{\omega} \times] & \mathbf{I}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} \end{bmatrix} \quad G = \begin{bmatrix} -\mathbf{I}_{3 \times 3} & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{I}_{3 \times 3} \end{bmatrix}, \quad (7)$$

where P is the state covariance and Q is the process noise matrix. The filter takes the pseudo-measurements from

each sequence as its measurements. The \hat{u} and \hat{s} vectors in the pseudo-measurements are assumed to be generated by multiplying the space-fixed object-observer and object-Sun unit vectors by a rotation matrix:

$$\mathbf{y}_k = \begin{bmatrix} \hat{u}_k^M \\ \hat{s}_k^M \end{bmatrix} = \begin{bmatrix} A(\mathbf{q}_k)^T \hat{u}_k^I \\ A(\mathbf{q}_k)^T \hat{s}_k^I \end{bmatrix}, \quad (8)$$

where \mathbf{y}_k is the measurement vector for time t_k , \hat{u}_k^I and \hat{s}_k^I are the object-observer and object-Sun directions in the space-fixed frame, \hat{u}_k^M and \hat{s}_k^M are the unit vector from the pseudo-measurement, and $A(\mathbf{q}_k)$ is the rotation matrix corresponding to the true orientation \mathbf{q}_k . The space-fixed vectors are known from the problem geometry. The measurement matrix before an update is then:

$$H_k = \begin{bmatrix} [A(\hat{\mathbf{q}}_k^-)^T \hat{u}_k^I \times] \\ [A(\hat{\mathbf{q}}_k^-)^T \hat{s}_k^I \times] \end{bmatrix}, \quad (9)$$

where $\hat{\mathbf{q}}_k^-$ is the *a priori* orientation estimate. The rest of the filter is as in Crassidis and Junkins [6]. The two methods of solving the association problem both have different ways of generating initial state guesses. As such, the discussion of filter initialization is saved until after each method has been outlined.

5.2. Pseudo-Measurement Ambiguity

Pseudo-measurements store information about the object's attitude in the form of $\hat{u} - \hat{s}$ vector pairs. Due to the nature of the reflection model, each pseudo-measurement corresponds to two different orientations: one where the object-observer and object-Sun vectors are assigned as listed, and one where their directions are swapped. This introduces additional ambiguity into the state estimation problem, which the MEKF is not capable of addressing as-is.

The ambiguity is addressed by comparing the MEKF's propagated orientation to the unit vectors in the pseudo-measurement using a distance measure:

$$D = 2 - (\hat{u}^P \cdot \hat{u}^M) - (\hat{s}^P \cdot \hat{s}^M), \quad (10)$$

where the superscript ^P indicates the unit vector for the propagated orientation, and the superscript ^M indicates the pseudo-measurement unit vectors. The MEKF is then updated using whichever orientation is closest to the propagated orientation.

6. PSEUDO-MEASUREMENT ASSOCIATION PROBLEM

It is practically impossible to try every possible combination of pseudo-measurements for the different time steps;

even for a light curve with less than two dozen measurements, the number of possible combinations can be on the order of 10^{52} . Therefore, there needs to be some method of limiting the number of combinations to be tried. Two methods for limiting the number of combinations are discussed. This corresponds to Block (4) and Block (5) in Figure 1.

6.1. Method 1: Finding Feasible Sequences

The first method is to use a simple gating test to rule out unlikely sequences of pseudo-measurements. Assuming that there is no sudden change to the object's dynamics during the period of observation, then the pseudo-measurement chosen for one time step should be consistent with the pseudo-measurements that came before it.

Consider two candidate pseudo-measurements for time t_i and time t_{i+1} . The object's motion from the orientation corresponding to a particular t_i pseudo-measurement to the orientation described by a t_{i+1} pseudo-measurement is referred to as a transition. The transition can be found using a quaternion representation of orientation. A quaternion is defined in terms of rotation as:

$$\mathbf{q} = \begin{bmatrix} \sin \frac{\theta}{2} \hat{e} \\ \cos \frac{\theta}{2} \end{bmatrix} = \begin{bmatrix} \mathbf{q} \\ q_4 \end{bmatrix}, \quad (11)$$

where θ is the rotation angle and \hat{e} is the rotation axis. The quaternions representing the candidate pseudo-measurements are $\tilde{\mathbf{q}}_i$ and $\tilde{\mathbf{q}}_{i+1}$. The tilde indicates that the quaternion corresponds to a pseudo-measurement. The rotation for the transition from $\tilde{\mathbf{q}}_i$ to $\tilde{\mathbf{q}}_{i+1}$ is:

$$\Delta \tilde{\mathbf{q}}_{i,i+1} = \tilde{\mathbf{q}}_i \otimes \tilde{\mathbf{q}}_{i+1}^*, \quad (12)$$

where \otimes represents quaternion multiplication and $\tilde{\mathbf{q}}_{i+1}^*$ is the conjugate of $\tilde{\mathbf{q}}_{i+1}$. The rotation $\Delta \tilde{\mathbf{q}}$ is compared to a reference rotation from t_i to t_{i+1} to determine whether the transition from $\tilde{\mathbf{q}}_i$ to $\tilde{\mathbf{q}}_{i+1}$ is feasible. For example, under constant, single-axis rotation, the reference rotation is the first rotation in the sequence from the t_1 pseudo-measurement to the t_2 pseudo-measurement. Since the rotation is constant, every other transition in the sequence must be similar to the first. Similarity is measured by finding the rotation between the transition under consideration and the first transition:

$$\Delta \mathbf{q}_{t,i} = \Delta \tilde{\mathbf{q}}_{1,2} \otimes \Delta \tilde{\mathbf{q}}_{i,i+1}, \quad (13)$$

where $\Delta \tilde{\mathbf{q}}_{1,2}$ is the rotation for the first transition in the sequence. The rotation angle θ for $\Delta \mathbf{q}_t$ can be found from (11):

$$\theta_{t,i} = 2 \cos^{-1} (\Delta q_{t,i,4}). \quad (14)$$

If $\theta_{t,i}$ is above some upper bound, then the transition is counted as infeasible. A transition may also be counted as infeasible if its absolute rotation angle is greater than some set maximum. Depending on the exact gating condition, this selection process can have a bias towards selecting high-energy rotations. Setting a limit to the rotation angle between pseudo-measurements is a way of countering that bias and restricting the search to lower-energy rotations. The set of all feasible transitions for a given reference transition is referred to as T . However, just because a particular transition is feasible does not mean that it can be reached using only other feasible transitions. To eliminate those cases, all transitions that cannot be reached without using some transition outside of T are discarded. "Dead end" transitions that could not reach the final measurement time without leaving T are also eliminated. The process is repeated using each unique transition from t_1 to t_2 in turn as the reference. This leaves a smaller set of feasible sequences of pseudo-measurements and reduces the number of attitude-time histories that need to be generated.

The above discussion assumed that each pseudo-measurement corresponds to a single orientation. In reality, each pseudo-measurement could correspond to two different orientations, so that four different transitions need to be considered. Two of those transitions will be duplicates in practice, so only two unique transitions need to be considered for each pair. These additional possibilities are dealt with in two ways. First, each $t_1 - t_2$ set of pseudo-measurements will be used to find feasible sequences twice, once for each possible transition. Second, each subsequent transition is considered feasible if at least one of its possible rotations is a close match for the initial transition according to the gating criteria for $\theta_{t,i}$.

6.2. Initial State Estimate for Method 1

Since the MEKF is a nonlinear state estimator, it needs a guess for the initial orientation and angular velocity states. The initial guess for the orientation is taken from the first pseudo-measurement in the sequence that the MEKF is processing. As each pseudo-measurement corresponds to two different orientations, each sequence will need to be initialized twice. The two filters will then be run in parallel on the pseudo-measurements in the sequence. At each step, the filters will pick between the orientations represented by the pseudo-measurement as discussed in Section 5.2. While both filters could, in theory, choose to use the same orientation for a particular pseudo-measurement, they always choose differently in practice. The paired MEKFs produce two separate attitude time-histories based on the same sequence of pseudo-measurements.

The initial angular velocity estimate can be found multiple ways. The simplest method is to use the first transition in the pseudo-measurement sequence:

$$\hat{\omega}_1 = \frac{\theta_{1,2}}{\Delta t} \hat{e}_{1,2}, \quad (15)$$

where $\theta_{1,2}$ and $\hat{e}_{1,2}$ are the rotation angle and axis for the first transition in the sequence, and Δt is the time between measurements.

Other estimates are possible, however. If angular velocity is either constant or only varying slowly for the first few time steps, then a linear least-squares fit can be used to find the angular velocity from the first several pseudo-measurements [5, 21]. For a pseudo-measurement sequence, the second-order central limit estimate is:

$$\hat{\omega} = -\frac{1}{\Delta t} \left([\hat{u}_k^M \times]^T [\hat{u}_k^M \times] + [\hat{s}_k^M \times]^T [\hat{s}_k^M \times] \right)^{-1} \left([\hat{u}_k^M \times]^T [\hat{u}_{k+1}^M - \hat{u}_{k-1}^M] + [\hat{s}_k^M \times]^T [\hat{s}_{k+1}^M - \hat{s}_{k-1}^M] \right) \quad (16)$$

where \hat{u}_k^M and \hat{s}_k^M are the pseudo-measured vectors at time t_k and Δt is the time between measurements. Since Equation (16) assumes that angular velocity is approximately constant between measurements, the central limit method can be used to estimate the initial angular velocity even though it relies on measurements from before time t_k .

6.3. Weighing History

For the attitude time-histories to be useful, there needs to be some way of sorting the “better” or “more likely” time histories from those that are less likely to correspond to the true motion of the object. This sorting is done by assigning each of the time histories a weight based on how the filter converges. Similar material was previously presented at the 2020 Astrodynamics Specialist Conference [3]; only the measurement matrix definition is different. As is sometimes the case in MHT applications, the weight will be probabilistic [27]. To compute the weight, the variance of the distance measure D between the estimated orientation and the pseudo-measured orientation at each time is found. This variance is equal to the *a posteriori* state covariance P_k^+ for the estimate multiplied before and after by the measurement matrix:

$$H_{D,k} = \begin{bmatrix} \frac{\partial D_k}{\partial \delta \alpha_x} & \frac{\partial D_k}{\partial \delta \alpha_y} & \frac{\partial D_k}{\partial \delta \alpha_z} & 0 & 0 & 0 \end{bmatrix}, \quad (17)$$

where the measurement time is t_k . Note that angular velocity has no effect on the variance since it does not appear in (10). The variance of D at t_k is then:

$$\sigma_{D,k}^2 = H_{D,k} P_k^+ H_{D,k}^T, \quad (18)$$

The sequence of distance measures D_k are then assembled into a multivariate normal distribution with mean μ_D and covariance P_D :

$$\mu_D = [D_1 \quad D_2 \quad \dots]^T \quad (19)$$

$$P_D = \begin{bmatrix} \sigma_{D,1}^2 & 0 & 0 & \dots \\ 0 & \sigma_{D,2}^2 & 0 & \dots \\ \vdots & & \ddots & \end{bmatrix}. \quad (20)$$

The weight assigned to the path is equal to the value of the probability density function when all D_k are equal to zero. The weight is relative and used only to compare different paths; there is no absolute cutoff for a path to be considered “likely” or “unlikely”. Once the weight has been computed, the path and its weight are stored, and the next sequence of pseudo-measurements is processed.

6.4. Method 2: PHD Filter

The method of pre-selecting feasible sequences of pseudo-measurements was tested for an object under constant rotation, as discussed in Section 7.1. This gave promising results, but it was difficult to tune the upper bound on $\theta_{t,i}$ even for simple dynamics. To extend the solver to more complex dynamics, a new method based on a probability hypothesis density (PHD) filter was developed. A PHD filter is a multi-target tracking filter that estimates the states of objects using a probability function over the state space [29, 23, 30]. The filter used for this method approximates the distribution as a sum of Gaussian distributions, referred to as Gaussian components [30].

At each update step, every component is updated with a Kalman filter. For the orientation estimation problem, the standard Extended Kalman filter is replaced with an MEKF. As discussed in Section 5.2, the pseudo-measured orientation closest to the propagated estimate is used in the MEKF update step. If there are multiple measurements at a particular time step, then each propagated component is updated using each measurement individually, producing a separate *a posteriori* component for each one. For example, if there are two Gaussian components prior to an update step with three measurements, the probability distribution will have six components after the update is done.

Pruning and merging are used at each time step to keep the number of Gaussian components from growing exponentially. Each Gaussian component has an assigned weight to represent the relative likelihood of its mean representing the true state. When using the PHD filter for multi-target tracking, weight calculations take into account the probability of detection for each target, the likelihood of a new target appearing or an existing target disappearing (i.e., births and deaths), and the presence of

clutter. However, for the light curve inversion problem there is known to be only one object underlying the light curve. As such, probability of detection is 1, there are no births or deaths, and there are no clutter measurements. This makes the weight calculations at time t_k as follows [30, 12, 24]:

$$w_{ij,k} = \frac{w_{i,k} \mathcal{N}(z_j - \hat{z}_i | H_{i,k}^- P_{i,k}^- (H_{i,k}^-)^T + R)}{\sum_{l=1}^{n_k} w_{l,k} \mathcal{N}(z_j - \hat{z}_l | H_{l,k}^- P_{l,k}^- (H_{l,k}^-)^T + R)} \quad (21)$$

where $w_{ij,k}$ is the weight computed using the *a priori* weight corresponding to the i^{th} component, $w_{i,k}$, and the j^{th} measurement z_j ; \hat{z}_i is the expected measurement from the i^{th} component; $H_{i,k}^-$ is the measurement matrix from the propagated component mean; $P_{i,k}^-$ is the propagated component covariance; R is the measurement covariance; and n_k is the number of components at t_k .

To keep down the number of components, any Gaussian with a weight less than some threshold is discarded. Similarly, any components which are close according to Mahalanobis distance are merged together. The weight, state, and covariance of a merged component are [12, 24]:

$$w_{il,k} = w_{i,k} + w_{l,k} \quad (22)$$

$$\hat{x}_{il,k} = \frac{w_{i,k} \hat{x}_{i,k} + w_{l,k} \hat{x}_{l,k}}{w_{i,k} + w_{l,k}} \quad (23)$$

$$P_{il,k} = \frac{w_{i,k} P_{i,k} + w_{l,k} P_{l,k}}{w_{i,k} + w_{l,k}} + (\delta \hat{x}_{il,k})(\delta \hat{x}_{il,k})^T \quad (24)$$

where $\delta \hat{x}_{il,k}$ is the length six vector as in Equation (5) giving the difference between the two components' state estimates.

Once all of the measurement time steps have been processed, the means of the highest-weighted components are used as probable estimates of the object's orientation and angular velocity at the final time step. The PHD filter, as implemented, does not associate components at one time step with those at previous times, so the attitude time-histories are found by back-propagating from the final time step.

6.5. Initial State Estimate for Method 2

The PHD filter is initialized in a similar way to the various sequences found for the first method. First, a number of orientations corresponding to the pseudo-measurements at the first time step are generated. Since each pseudo-measurement corresponds to two different

orientations, there will be two orientations per pseudo-measurement at the initial time. These orientations are all equally weighted, and any that are sufficiently close to each other are merged per Equation (22) through Equation (24). Second, each initial orientation's possible angular velocities are found by comparing the orientations to the pseudo-measurements at the second time step. The angular velocities are calculated using Equation (15), and a separate component is generated for each angular velocity at its corresponding initial orientation. Finally, any components nearby to each other are merged. This process gives the set of Gaussian components at the initial time step.

7. RESULTS

7.1. Simply-Rotating Results

The light curve inversion method is applied to the case of a regular tetrahedron rotating at a constant rate $\hat{\omega}_{\text{ex}}$:

$$\hat{\omega}_{\text{ex}} = [1 \quad 1 \quad 0]^T \text{ rad/s.} \quad (25)$$

A measurement is taken every 0.25 seconds for five seconds, producing the light curve from Figure 2. This is a difficult case since the object is highly symmetrical, each side reflects light in the same manner, and there are no specular glints. The possible orientations for each visibility group were found using the viewing sphere, as shown in Figure 3. After searching the viewing sphere, between 106 and 1241 possible pseudo-measurements were found for each time step. A total of 17,280 feasible sequences were found out of $1.8 \cdot 10^{52}$ possible sequences of these pseudo-measurements. The feasible sequences were run through the MEKF using the angular velocity from Equation (15) as the initial guess.

The solutions found using the MEKF have weights up to $3.6 \cdot 10^{24}$. The estimated body-frame \hat{u} and \hat{s} vectors for the estimated path are compared to the true body-frame unit vectors in Figure 4. The paths of the true body

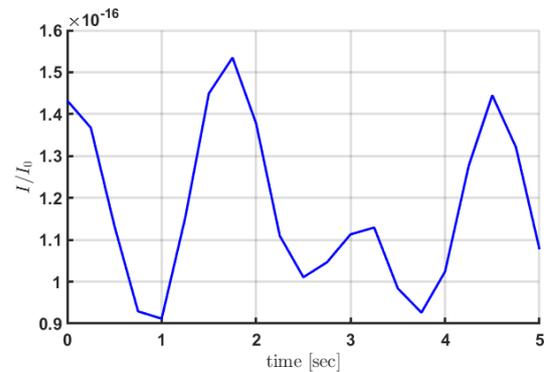


Figure 2. The measured light curve for the simply-rotating tetrahedron.

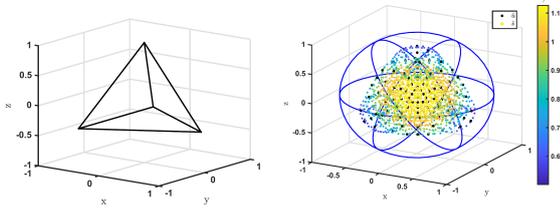


Figure 3. These images show the regular tetrahedron (left) and the \hat{u} and \hat{s} locations on the viewing sphere for one visibility group (right). The blue lines indicate viewing events, and the color of the \hat{s} vectors indicate the reflection fraction of each vector pair. See Burton and Frueh [3].

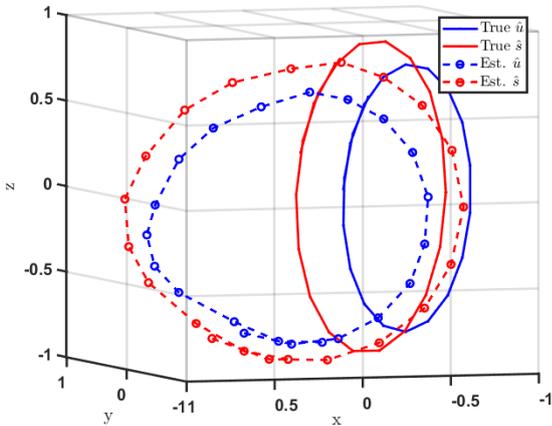


Figure 4. A plot comparing the true body-frame \hat{u} and \hat{s} directions (solid) with the highest-weighted estimated vectors (dashed) over the course of the measurement period.

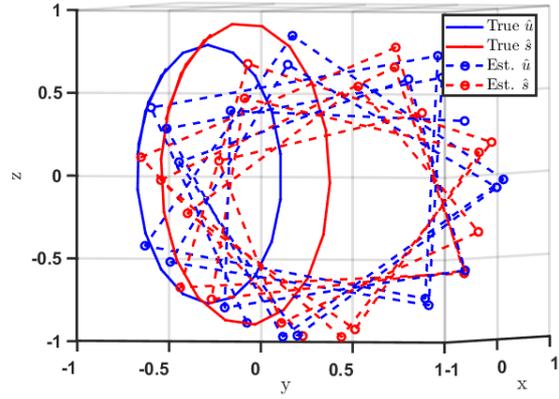


Figure 5. This plot compares the true body-frame \hat{u} and \hat{s} directions (solid) with the lowest-weighted attitude time-history (dashed).

frame vectors form circles with a common axis parallel to $\hat{\omega}_{ex}$. The estimated path's average angular velocity has a magnitude of 1.29 rad/s, compared to the 1.41 rad/s true angular velocity. However, the direction of the rotation is different, most likely because of the highly symmetric nature of the regular tetrahedron.

The lowest-weight path had a weight of $1.5 \cdot 10^{18}$. In previous work [3], weights as low as $3.3 \cdot 10^{-21}$ have been found for certain time histories, so this is still relatively high thanks to the pseudo-measurement selection process. The large differences in weights are because a path's weight increases exponentially as the time-history better fits the pseudo-measurements. The average angular velocity for the low-weight case has a magnitude of 7.09 rad/s, compared to the true magnitude of 1.41 rad/s and 1.3 rad/s magnitude for the high-weight estimate. The larger mismatch between the true and estimated angular velocity magnitudes helps explain the lower weight for this estimate compared to the high-weight path. The paths of the estimated unit vectors for this case are compared to the truth in Figure 5.

The PHD filter was also used to analyze the light curve in Figure 2. At the final time step, it had 4684 Gaussian components with weights ranging from 22.6 to 0.01. Note that the PHD filter weights are calculated differently than the weights for the feasible paths method, so these are not directly comparable to the weights discussed above. Since the second-highest weight is 21.4, compared to 6.5 for the third highest, the top two components will be discussed. Back-propagating from the final time step gives time-histories for the body-frame \hat{u} and \hat{s} over time. These are plotted against the true body-frame vectors in Figure 6. As in the first method, the angular velocity estimates are in different directions than the truth but have similar magnitude: 1.9 rad/s for the highest-weight case and 1.6 rad/s for the second-highest weight. Compare to the true angular velocity magnitude of 1.4 rad/s.

Both methods gave comparable results for the single-axis rotation problem. However, the number of feasible se-

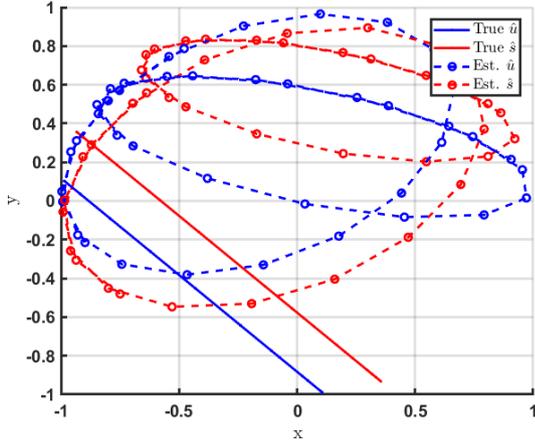


Figure 6. This plot shows the two highest-weighted body-frame \hat{u} and \hat{s} directions (dashed) for the PHD filter method along with the true body-frame vectors (solid).

quences found by the first method is extremely sensitive to the gating criteria for $\theta_{t,i}$. By contrast, the second method, using the PHD filter, required significantly less tuning and is more suitable for use on objects with more complex rotation.

7.2. Non-Constant Rotation

The PHD filter was tested on a second light curve produced by an object undergoing non-constant torqueless rotation. Once again a regular tetrahedron was used, but to reduce the symmetry of the problem each side was set to have a different reflection coefficient. The object began rotating with angular velocity $\hat{\omega}_{\text{ex}}$:

$$\hat{\omega}_{\text{ex}} = [1 \quad 1 \quad 0]^T \text{ rad/s} \quad (26)$$

and was allowed to tumble freely. A total of 21 light curve measurements were taken at a rate of one per second, plotted as the solid line in Figure 7. Between one and seventy-three pseudo-measurements were found for each measurement time. After running the filter, the component with the highest weight was back-propagated to find an attitude time-history. Since the tumbling motion makes visually comparing the true and estimated body-frame unit vectors difficult, the resulting light curve is instead compared to the actual measurements in Figure 7. From that plot, it can be seen that the light curve for the estimated time-history follows similar trends to the true light curve while being phase-shifted to the left by two to three seconds.

The merging and pruning cutoffs for the PHD filter were not carefully tuned for this test run, so it is likely that better results could be obtained by altering the filter parameters. Furthermore, the set of pseudo-measurements was extremely small; recall that one update step had only a

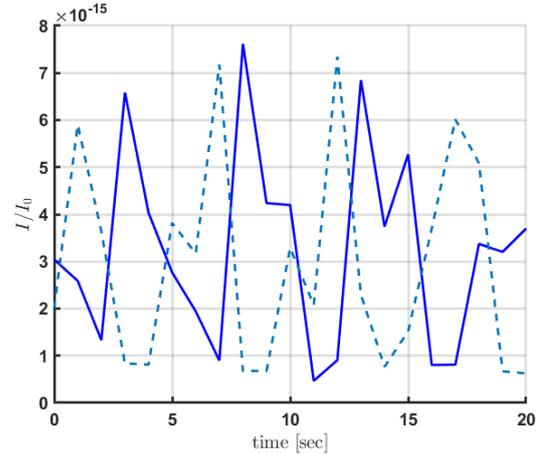


Figure 7. The measured light curve (solid) for the free-tumbling regular tetrahedron is compared to the light curve for the highest-weighted estimate (dashed).

single pseudo-measurement. Loosening the requirements for an orientation to be considered a feasible pseudo-measurement would give the filter a larger set of data and likely lead to a better estimate.

7.3. Extended Data Set

To confirm that the PHD filter with the MEKF would produce an accurate estimate when given good pseudo-measurement data, the tumbling object case was run with a slightly modified set of pseudo-measurements: in addition to those generated using the viewing sphere, the true body-frame \hat{u} and \hat{s} unit vectors were added to the pseudo-measurement list at each time step. The light curve for the highest-weighted estimate from this run is shown in Figure 8.

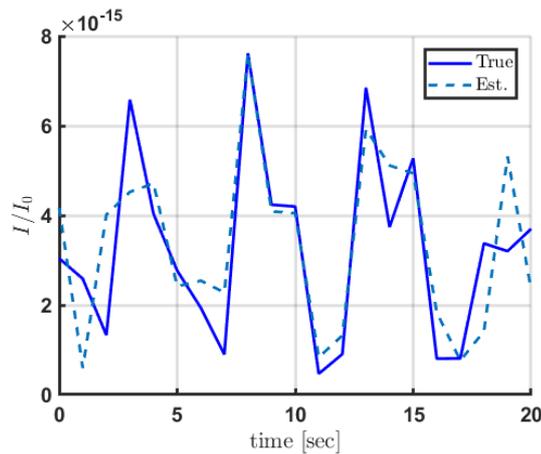


Figure 8. The measured light curve (solid) for the tumbling tetrahedron is compared to the light curve for the highest weighted estimate (dashed) when the pseudo-measurement sets are supplemented with the true body-frame \hat{u} and \hat{s} vectors.

8. CONCLUSIONS

When no resolved image of a space object can be obtained, light curve inversion is a useful tool for estimating its attitude and rotational characteristics. To avoid making assumptions about the object's dynamics, orientations of the object are generated with knowledge of its geometry and reflective properties alone using the viewing sphere. The orientations are compared to the light curve measurements to find feasible orientations, or pseudo-measurements, at each time step. The number of possible time-histories based on these pseudo-measurements is large, so two methods to process them efficiently were tested. In the first method, the set of possible pseudo-measurement sequences is narrowed down using a gating criteria. The time-histories for each sequence are given weights based on how well the estimated orientations fit the pseudo-measurements. The number of feasible sequences was found to be extremely sensitive to the gating criteria chosen.

The second method uses a probability hypothesis density filter to create a Gaussian-mixture density function of the object's state at each measurement time. To avoid an exponential growth in the number of Gaussian components, components are pruned and merged after each update step. This method was applied successfully to both a simply-rotating object and an object undergoing torqueless multi-axis rotation. Tuning the filter should further improve the results.

REFERENCES

1. BENSON, C. J., SCHEERES, D. J., RYAN, W. H., AND RYAN, E. V. Cyclic complex spin state evolution of defunct geo satellites. In *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference, Maui, HI* (2018).
2. BLACKMAN, S. S. Multiple hypothesis tracking for multiple target tracking. *IEEE Aerospace and Electronic Systems Magazine* 19, 1 (2004), 5–18.
3. BURTON, A., AND FRUEH, C. Light curve attitude estimation using the viewing sphere. In *Astrodynamics Specialist Conference 2020* (2020).
4. CALEF, B., AFRICANO, J., BIRGE, B., HALL, D., AND KERVIN, P. Photometric signature inversion. In *Unconventional Imaging II* (2006), vol. 6307, International Society for Optics and Photonics, p. 63070E.
5. CRASSIDIS, J. L. Angular velocity determination directly from star tracker measurements. *Journal of guidance, control, and dynamics* 25, 6 (2002), 1165–1168.
6. CRASSIDIS, J. L., AND JUNKINS, J. L. *Optimal Estimation of Dynamic Systems*. CRC Press, 2012.
7. CRASSIDIS, J. L., AND MARKLEY, F. L. Unscented filtering for spacecraft attitude estimation. *Journal of guidance, control, and dynamics* 26, 4 (2003), 536–542.
8. DAVEY, S. J. *Extensions to the probabilistic multi-hypothesis tracker for improved data association*. Citeseer, 2003.
9. ĀURECH, J., AND KAASALAINEN, M. Photometric signatures of highly nonconvex and binary asteroids. *Astronomy & Astrophysics* 404, 2 (2003), 709–714.
10. FAN, S. *The Light Curve Simulation and Its Inversion Problem for Human-Made Space Objects*. PhD thesis, Purdue University, August 2020.
11. FRUEH, C. Space traffic management. Lecture Notes, 2019.
12. FRUEH, C. Notes on r. mahler: Statistical multisource-multitarget information fusion, artech house, 2007. Purdue University, 2021.
13. GIGUS, Z., AND MALIK, J. Computing the aspect graph for line drawings of polyhedral objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12, 2 (1990), 113–122.
14. GUALTIERI, J. A., BAUGHER, S., AND WERMAN, M. The visual potential: One convex polygon. *Computer Vision, Graphics, and Image Processing* 46, 1 (1989), 96–130.
15. HALL, D., CALEF, B., KNOX, K., BOLDEN, M., AND KERVIN, P. Separating attitude and shape effects for non-resolved objects. In *The 2007 AMOS Technical Conference Proceedings* (2007), Maui Economic Development Board, Inc. Kihei, Maui, HI, pp. 464–475.
16. KAASALAINEN, M., LAMBERG, L., LUMME, K., AND BOWELL, E. Interpretation of lightcurves of atmosphereless bodies. i-general theory and new inversion schemes. *Astronomy and Astrophysics* 259 (1992), 318–332.

17. KAASALAINEN, M., AND TORPPA, J. Optimization methods for asteroid lightcurve inversion: I. shape determination. *Icarus* 153, 1 (2001), 24–36.
18. KAASALAINEN, M., TORPPA, J., AND MUINONEN, K. Optimization methods for asteroid lightcurve inversion: II. the complete inverse problem. *Icarus* 153, 1 (2001), 37–51.
19. LINARES, R., AND CRASSIDIS, J. L. Space-object shape inversion via adaptive hamiltonian markov chain monte carlo. *Journal of Guidance, Control, and Dynamics* 41, 1 (2018), 47–58.
20. LINARES, R., JAH, M. K., CRASSIDIS, J. L., AND NEBELECKY, C. K. Space object shape characterization and tracking using light curve and angles data. *Journal of Guidance, Control, and Dynamics* 37, 1 (2014), 13–25.
21. LIU, H.-B., YANG, J.-C., YI, W.-J., WANG, J.-Q., YANG, J.-K., LI, X.-J., AND TAN, J.-C. Angular velocity estimation from measurement vectors of star tracker. *Applied optics* 51, 16 (2012), 3590–3598.
22. MAGNUSON, P., BARUCCI, M. A., DRUMMOND, J. D., LUMME, K., OSTRO, S. J., SURDEJ, J., TAYLOR, R. C., AND ZAPPALA, V. Determination of pole orientations and shapes of asteroids. In *Asteroids II*, R. P. Binzel, T. Gehrels, and M. S. Matthews, Eds. The University of Arizona Press, 1989, ch. 5, pp. 66–97.
23. MAHLER, R. P. A theoretical foundation for the stein-winter probability hypothesis density (phd) multitarget tracking approach. Tech. rep., ARMY RESEARCH OFFICE ALEXANDRIA VA, 2000.
24. MAHLER, R. P. *Statistical Multisource-Multitarget Information Fusion*. Artech House, 2007.
25. MARKLEY, F. L. Attitude error representations for kalman filtering. *Journal of guidance, control, and dynamics* 26, 2 (2003), 311–317.
26. MARKLEY, F. L. Attitude estimation or quaternion estimation? *Journal of the Astronautical Sciences* 52, 1-2 (2004), 221–238.
27. REID, D. An algorithm for tracking multiple targets. *IEEE transactions on Automatic Control* 24, 6 (1979), 843–854.
28. RUSSELL, H. N. On the light-variations of asteroids and satellites. *The Astrophysical Journal* 24, 1 (1906), 1–18.
29. STREIT, R. L., AND LUGINBUHL, T. E. Maximum likelihood method for probabilistic multihypothesis tracking. In *Signal and Data Processing of Small Targets 1994* (1994), vol. 2235, International Society for Optics and Photonics, pp. 394–405.
30. VO, B.-N., AND MA, W.-K. The gaussian mixture probability hypothesis density filter. *IEEE Transactions on signal processing* 54, 11 (2006), 4091–4104.
31. WETTERER, C. J., AND JAH, M. Attitude determination from light curves. *Journal of Guidance, Control, and Dynamics* 32, 5 (2009), 1648–1651.
32. WILLIAMS, J. The determination of the orientation of a tumbling cylinder from the shape of the light curve. *OFFICE OF AEROSPACE RESEARCH* (1969), 31.
33. WILLIAMS, V. Location of the rotation axis of a tumbling cylindrical earth satellite by using visual observations: Part i: Theory. *Planetary and Space Science* 27, 6 (1979), 885–890.
34. YILMAZ, A., JAVED, O., AND SHAH, M. Object tracking: A survey. *Acm computing surveys (CSUR)* 38, 4 (2006), 13–es.
35. YU, L., JI, J., AND WANG, S. Shape, thermal and surface properties determination of a candidate spacecraft target asteroid (175706) 1996 fg3. *Monthly Notices of the Royal Astronomical Society* 439, 4 (2014), 3357–3370.