An AEGIS-CPHD Filter to Maintain Custody of GEO Space Objects with Limited Tracking Data

Steven Gehly, Brandon Jones, and Penina Axelrad
University of Colorado at Boulder

ABSTRACT

The Geosynchronous orbit regime is a challenging measurement environment for estimating the states of space objects. The combination of sparse data sets, missed detections, and clutter creates opportunities for estimators to fail and lose track of known objects. One important consideration in the attempt to maintain custody of these objects is the calculation of the probability of detection, a quantity used in the measurement update of many multitarget filters. This paper considers two approaches to the problem and examines the effects of computing probability of detection in a modern multitarget filter.

1. INTRODUCTION

Tracking space objects in Geosynchronous (GEO) orbits poses a number of unique challenges, many of which are associated with the collection of measurements. GEO objects are generally tracked using ground-based optical sensors, which are limited to taking measurements when lighting conditions are acceptable and are also restricted in their field of view (FOV). Given a limited number of sensors and many objects to track, the result is often long gaps between observation sets for individual objects, during which it is possible for multitarget filters to lose custody of objects they are tracking. In addition, sensors can miss detections of objects that are in their FOV, and produce false measurements, or clutter. An effective filter for Space Situational Awareness (SSA) must be robust against these conditions, and is considered to maintain custody of objects by producing a stable estimate of the correct target number and state errors that do not grow over time.

The past two decades have seen the development of a new mathematical framework for multitarget tracking termed Finite Set Statistics (FISST) [1,2]. FISST generalizes familiar concepts from single target statistics such as probability density functions (PDFs) and statistical moments in order to derive Bayes-optimal filters for multitarget problems. The use of order-independent sets of random vectors called Random Finite Sets (RFSs) to represent the multitarget state and measurement sets allows for estimation of the most likely number of targets and their states at each time. The filters also explicitly account for clutter and missed detections, the latter by computing a probability of detection, p_D. This probability is used in the measurement update step and is instrumental in maintaining custody of known objects. Previous research has produced closed form solutions to FISST-derived filters [3, 4], but no emphasis was placed on maintaining custody of objects outside the sensor FOV. This paper considers two different methods to compute p_D, incorporating information about the predicted state and uncertainty in relation to the field of view. The objective is to determine the extent to which these calculations affect filter performance, both in terms of state accuracy and the estimated number of objects.

2. BACKGROUND

2.1. AEGIS-CPHD FILTER

The simplest FISST-derived estimator is the Probability Hypothesis Density (PHD) filter, which predicts and corrects the first moment of the multitarget PDF, known as the intensity function or PHD. As a first moment approximation, the PHD filter truncates information in the PDF related to target number, and is known to suffer from high variability in the estimated number of targets, particularly in cases where detections of objects are missed due to imperfect sensors [2]. To address this issue, the Cardinalized PHD (CPHD) filter propagates and updates a cardinality distribution in addition to the PHD function [4, 5]. The cardinality is a discrete probability distribution over the possible number of targets.

Incorporating the cardinality recursion produces a more computationally complex filter, but also provides a more stable estimate of the number of objects at each time.

The full CPHD recursion is not computationally tractable, but may be simplified by approximating the PHD using a Gaussian Mixture Model (GMM) to yield a closed-form solution. Following the development of Vo et al. [4], the PHD at time t_k is approximated as a weighted sum of J_k Gaussian PDFs,

$$\nu_k(\boldsymbol{x}) \approx \sum_{j=1}^{J_k} w_k^{(j)} p_g(\boldsymbol{x}; \boldsymbol{x}_k^{(j)}, P_k^{(j)}) \tag{1}$$

where $w_k^{(j)} > 0$ are the weights and $p_g(\boldsymbol{x}; \boldsymbol{x}_k^{(j)}, P_k^{(j)})$ is a multivariate Gaussian PDF with mean $\boldsymbol{x}_k^{(j)}$ and covariance $P_k^{(j)}$. The prediction step for the PHD and cardinality is given by

$$\nu_{k|k-1}(\boldsymbol{x}_k) = \gamma_k(\boldsymbol{x}) + p_S \sum_{j=1}^{J_{k-1}} w_{k-1}^{(j)} p_g(\boldsymbol{x}; \boldsymbol{x}_{k|k-1}^{(j)}, P_{k|k-1}^{(j)}) \tag{2}$$

$$p_{k|k-1}(N) = \sum_{j=0}^{N} p_{\Gamma,k}(N-j) \sum_{l=j}^{\infty} C_j^l p_{k-1}(l) p_S^{(j)} (1-p_S)^{(l-j)} \tag{3}$$

where $\gamma_k(\boldsymbol{x})$ represents the PHD of new, or birth, targets and can be approximated by a GMM as in Eq. (1), p_S is the probability of target survival, $p_{\Gamma,k}(\cdot)$ represents the cardinality of target birth and $C_j^l = \frac{l!}{j!(l-j)!}$ is the binomial coefficient.

The update equations for the cardinality distribution and PHD are coupled, and include the previously mentioned probability of detection term in several places.

$$\nu_k(\boldsymbol{x}) = \frac{\langle \Psi_k^1[w_{k|k-1}, Z_k], p_{k|k-1} \rangle}{\langle \Psi_k^0[w_{k|k-1}, Z_k], p_{k|k-1} \rangle} (1-p_D) \nu_{k|k-1}(\boldsymbol{x}) + \sum_{z \in Z_k} \sum_{j=1}^{J_{k|k-1}} w_k^{(j)}(z) p_g(\boldsymbol{x}; \boldsymbol{x}_k^{(j)}, P_k^{(j)}) \tag{4}$$

$$p_k(N) = \frac{\Psi_k^0[w_{k|k-1}, Z_k](N) p_{k|k-1}(N)}{\langle \Psi_k^0[w_{k|k-1}, Z_k], p_{k|k-1} \rangle} \tag{5}$$

Here, Z_k is the measurement set, and the following definitions are used

$$\Psi_k^u[w, Z](N) = \sum_{j=0}^{\min(m_k, N)} (m_k - j)! p_\kappa(m_k - j) P_{j+u}^n \frac{\langle 1 - p_D, \nu \rangle^{N-(j+u)}}{\langle 1, \nu \rangle^N} \sigma_j(\Lambda_k(w, Z)) \tag{6}$$

$$\Lambda_k(w, Z) = \left\{ \frac{\langle 1, \kappa_k \rangle}{\kappa_k(z)} p_D \sum_{j=1}^{J_{k|k-1}} w_{k|k-1}^{(j)} p_g(\boldsymbol{z}; \boldsymbol{z}_k^{(j)}, P_{zz}^{(j)}) : \boldsymbol{z} \in Z \right\} \tag{7}$$

$$w_k^{(j)}(z) = p_D w_{k|k-1}^{(j)} p_g(\boldsymbol{z}; \boldsymbol{z}_k^{(j)}, P_{zz}^{(j)}) \frac{\langle \Psi_k^1[w_{k|k-1}, Z_k \backslash \{z\}], p_{k|k-1} \rangle}{\langle \Psi_k^0[w_{k|k-1}, Z_k], p_{k|k-1} \rangle} \frac{\langle 1, \kappa_k \rangle}{\kappa_k(z)} \tag{8}$$

where m_k is the number of measurements, p_κ is the cardinality of clutter, $\kappa_k(z)$ is the PHD of clutter, and $\sigma_j(\cdot)$ are the elementary symmetric functions, which can be computed using a formulation from Mahler [2]. At time t_k, the number of targets can be estimated using the PHD or the estimated cardinality distribution.

$$N_k = \int \nu_k(\boldsymbol{x}) \approx \sum_{j=1}^{J_k} w_k^{(j)} \tag{9}$$

$$N_k = \sum_{N=1}^{\infty} N \cdot p_k(N) \tag{10}$$

The PHD update Eq. (4) contains two terms, the first of which accounts for the possibility of missed detections. In this case, the filter retains the *a priori* PHD function. In the GMM approximation, the *a priori* component means and covariances are kept and are not updated by any measurement, and the updated component weights are computed from

$$w_k^{(j)} = \frac{\langle \Psi_k^1[w_{k|k-1}, Z_k], p_{k|k-1} \rangle}{\langle \Psi_k^0[w_{k|k-1}, Z_k], p_{k|k-1} \rangle}(1 - p_D) w_{k|k-1}^{(j)} \tag{11}$$

This term produces $J_{k|k-1}$ GMM components whose weights are scaled by a factor including $(1 - p_D)$. If the probability of detection is high, the resulting weight will be low, and the component will not contribute much to the updated PHD. If a detection for an object is missed, a high p_D may cause the filter to lose track of the object, therefore calculation of p_D is significant to maintaining custody of objects.

The second term of Eq. (4) accounts for detected objects, and computes a measurement update for each GMM component using each measurement, thereby producing $J_{k|k-1} \cdot m_k$ components. The updated weights for these components are scaled by a factor including p_D and the individual likelihood of measurement to component association $p_g(\boldsymbol{z}; \boldsymbol{z}_k^{(j)}, P_{zz}^{(j)})$. Therefore, if a measurement is far from the predicted measurement of a component, the weight will be small and the component ignored. Additionally, if a low p_D is incorrectly selected for an object that is detected, it may cause the filter to lose custody of the object.

The net result of the measurement update step is the creation of $J_{k|k-1} \cdot (m_k + 1)$ GMM components, many of which will have small weights and not contribute to the PHD. To keep the problem computationally tractable, a strategy is needed to remove GMM components that are not significant. In this paper, the Adaptive Entropy-based Gaussian Information Synthesis (AEGIS) filter [6, 7] is used as the underlying estimator to predict and correct individual GMM component means and covariances and manage the overall number of GMM components.

AEGIS is an adaptive GMM filter, and can model non-Gaussian uncertainty by using multiple GMM components to represent each object if a sufficient amount of nonlinearity is detected during the prediction step. This is important in SSA applications, as the orbital dynamics are nonlinear and propagated uncertainties become non-Gaussian during long data gaps. In addition, AEGIS prescribes methods to remove or merge GMM components if the weights become small or the components are close together, typically as determined from the Mahalanobis distance. The underlying equations for the prediction and correction steps come from the Unscented Kalman Filter (UKF), therefore, the individual component means, covariances, and predicted measurements in Eqs. (2)-(8) are computed using the UKF equations [8].

2.2. CLUTTER MODEL

The CPHD filter requires an assumed distribution for both the cardinality and PHD of clutter. In this paper, the cardinality of clutter is assumed to be Poisson-distributed.

$$p_\kappa(N) = \frac{(\lambda \cdot V_s)^N}{N!} \exp(-\lambda \cdot V_s) \tag{12}$$

The Poisson mean rate is therefore $\lambda \cdot V_s$, where λ is a mean number of clutter returns per unit volume in measurement space, and V_s is the sensor volume. The clutter PHD is assumed to be uniformly distributed in the sensor FOV, and can be defined as

$$\kappa(z) = \lambda \cdot V_s \cdot \mathcal{U}(z) \tag{13}$$

$$\mathcal{U}(z) = \begin{cases} 1/V_s & \text{if } z \in \text{FOV} \\ 0 & \text{if } z \notin \text{FOV} \end{cases} \tag{14}$$

Note that this simplifies some of the terms in the measurement update,

$$\frac{\langle 1, \kappa_k \rangle}{\kappa_k(z)} = \frac{\lambda \cdot V_s \int \mathcal{U}(z) dz}{\lambda} = V_s \tag{15}$$

because $\kappa_k(z) = \lambda \cdot V_s \cdot 1/V_s = \lambda$ in the FOV, and $\int \mathcal{U}(z) dz = 1$.

2.3. PROBABILITY OF DETECTION CALCULATION

As discussed above, computing the correct probability of detection plays an important role in maintaining custody of objects in the filter. The previous section applies a constant p_D for all components, but this can be problematic for objects that are outside the sensor FOV at measurement time. In the case that objects are known to be outside the FOV and cannot be detected, assigning the same p_D as for objects in the FOV will quickly cause GMM components to be downweighted and knowledge of the object to be lost. Different values of p_D need to be assigned for each object and it is essential to account for known objects' positions in regard to the FOV.

A simple approach to determining $p_D^{(j)}$ for an individual component is to use the predicted measurement and an indicator function [9]. The probability of detection can be modeled as a product of two terms, a constant $p_{D,\text{sensor}}$ that accounts for the sensor's imperfect ability to detect objects in the FOV, and a state-dependent term

$$p_{D,\text{FOV}}^{(j)} = \begin{cases} 1 & \text{if } \boldsymbol{z}^{(j)} \in \text{FOV} \\ 0 & \text{if } \boldsymbol{z}^{(j)} \notin \text{FOV} \end{cases} \tag{16}$$

Assuming the two processes are independent, the overall $p_D^{(j)} = p_{D,\text{sensor}} \cdot p_{D,\text{FOV}}^{(j)}$.

This approach addresses the issue, but may not be sufficient in cases where the object is near the edge of the field of view. Consider the case illustrated in Fig. 1 in which a predicted measurement lies in the FOV, but the actual object is not visible. In this case, applying a probability of detection based only on the component's mean state can lead to issues. In the limiting case that $p_{D,\text{sensor}} = 1$, the missed detection component weight will be zero and the filter will lose custody of the object. Even in the case $p_{D,\text{sensor}} < 1$, if the state-dependent $p_D^{(j)}$ is used and the object remains at the edge of the FOV for several measurement epochs, the component will be downweighted significantly and can be lost.

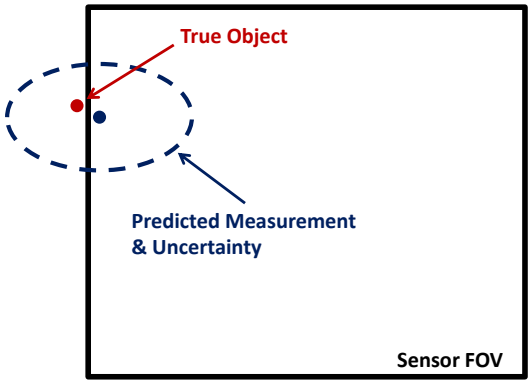

Fig. 1. Illustration of Sensor Field of View

An alternative approach is to use the predicted uncertainty to compute the probability that an object is in the field of view. In this case,

$$p_{D,\text{FOV}}^{(j)} = \int_{\text{FOV}} p_g(\boldsymbol{z}; \boldsymbol{z}_k^{(j)}, P_{zz}^{(j)}) d\boldsymbol{z} \tag{17}$$

which will produce $0 \leq p_{D,\text{FOV}}^{(j)} \leq 1$, and importantly, reduces the overall $p_D^{(j)}$ in cases where an object is near the edge of the FOV.

Either approach requires minor changes in the measurement update equations above. In Eqs. (4) and (8), the appropriate $p_D^{(j)}$ must be used for each component, more easily seen in Eq. (11) for the missed detection case. In Eq. (7), the $p_D^{(j)}$ must be moved inside the summation, and finally the modification to Eq (6) is most easily simplified by writing out the inner product for the GMM approximation.

$$\langle 1 - p_D^{(j)}, \nu \rangle = \sum_{j=1}^{J} (1 - p_D^{(j)}) w^{(j)} \tag{18}$$

3. NUMERICAL SIMULATION

To assess the effects of using different probability of detection calculations in the CPHD filter, a numerical simulation is included. Two test cases are analyzed, each estimating the state of 4 objects in GEO orbit with no target birth or death considered. The system dynamics are assumed to be given by the two-body problem,

$$\ddot{\mathbf{r}} = \frac{-\mu}{r^3}\mathbf{r} \tag{19}$$

where $\mathbf{r}$ is the position vector and μ is the gravitational parameter. The state vector consists of the position and velocity in the ECI frame.

$$\boldsymbol{X} = \begin{bmatrix} x & y & z & \dot{x} & \dot{y} & \dot{z} \end{bmatrix}^T \tag{20}$$

Both test cases involve short observation windows in which not all objects are detected, and both satellite and clutter measurements are received, followed by gaps in measurement data. Measurements used are topocentric right ascension and declination taken from a single ground station located in Maui at the coordinates given in Table 1. The sensor is assumed to have a 2 degree field of view in both axes, and white measurement noise is added with standard deviations provided in Table 3.

Table 1. Station Location in ECEF

Station	x_s (km)	y_s (km)	z_s (km)
Maui	-5465.210	-2403.610	2242.120

The measurement equations are given by

$$\alpha = \tan^{-1}\left(\frac{y - y_{si}}{x - x_{si}}\right) \qquad \delta = \sin^{-1}\left(\frac{z - z_{si}}{\rho}\right) \tag{21}$$

where the subscript si denotes a ground station coordinate in ECI and $\rho = \sqrt{(x - x_{si})^2 + (y - y_{si})^2 + (z - z_{si})^2}$ is the range.

To generate the initial state of each object, Keplerian elements are used. In both test cases the orbital parameters for the EchoStar1 satellite define the the first object, with starting positions for other objects randomly initialized by varying the eccentricity, inclination, and mean anomaly from the EchoStar1 elements.[1] At each time step, the sensor is pointed at EchoStar1, and the variations in initial elements cause the other objects to drift in and out of the field of view. The initial state and standard deviations used to generate objects are provided in Table 2.

Table 2. Initial State Parameters

Parameter	Value
a	42164.573 (km)
e	0.0002878
i	0.006 (deg)
Ω	278.657 (deg)
ω	139.8697 (deg)
M	181.4332 (deg)
σ_i	1.0 (deg)
σ_M	0.25 (deg)
σ_e	0.006

[1] TLE obtained from www.space-track.org

To initialize the filter covariance, uncertainties in the element space are used and rotated to the ECI frame using an unscented transform, values are provided in Table 3. The initial cardinality distribution sets an equal probability for each number of objects from 1-15, and the initial PHD function is approximated with 4 GMM components randomly perturbed from the true object states using the standard deviations in Table 3. Finally, a process noise covariance is added at the end of each filter prediction step with values defined in the RIC frame and rotated to ECI at each time, see Table 4. The values listed are multiplied by the time interval between observations during measurement arcs, but not following data gaps.

Table 3. Initial State Covariance and Measurement Noise Standard Deviations

Parameter	Value
σ_a	10 (km)
σ_e	1e-4
σ_i	0.1 (deg)
σ_Ω	0.001 (deg)
σ_ω	0.001 (deg)
σ_M	0.001 (deg)
$\sigma_\alpha = \sigma_\delta$	1 (arcsec)

Table 4. Process Noise Standard Deviations in RIC Frame

Parameter	Value
$\sigma_r = \sigma_c$	0.0032 (km)
σ_i	0.01 (km)
$\sigma_{\dot{r}} = \sigma_{\dot{c}}$	3.2e-6 (km/s)
$\sigma_{\dot{i}}$	1e-5 (km/s)

3.1. PERFORMANCE METRIC

To evaluate filter performance, it is necessary to define a miss distance for the multitarget problem. In the context of the RFS, the multitarget miss distance must provide the difference between two sets of vectors, accounting for errors both in cardinality and the state estimates of individual objects. To address this need, the Optimal Subpattern Assignment (OSPA) metric has been developed, offering several advantages over previous multitarget error metrics [10].

OSPA is a consistent metric on the space of finite sets, and provides a meaningful physical interpretation of state and cardinality errors by assigning the largest possible subset of estimated and true states, computing the error between them, and adding a fixed error for each point that is not assigned. The computation can be summarized in three steps.

Given two random finite sets, $X = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_m\}$ and $Y = \{\boldsymbol{y}_1, \ldots, \boldsymbol{y}_n\}$, the p-th order OSPA metric with cutoff c is defined as

$$d_p^{(c)}(X,Y) = \left[\frac{1}{n}\left(\min_{\pi\in\Pi_n}\sum_{i=1}^{m} d^{(c)}(\boldsymbol{x}_i, \boldsymbol{y}_{\pi(i)})^p + c^p(n-m)\right)\right]^{1/p} \tag{22}$$

$$d^{(c)}(\boldsymbol{x}_i, \boldsymbol{y}_{\pi(i)}) = \min(c, \|\boldsymbol{x}_i - \boldsymbol{y}_{\pi(i)}\|_p) \tag{23}$$

where Π_n represents the set of permutations on $\{1, \ldots, n\}$ and it is assumed $m \leq n$. For the case $m > n$, the metric can be computed as $d_p^{(c)}(Y, X)$. The steps to compute OSPA are as follows:

1. Find the optimal subpattern assignment between X and Y that minimizes the p-norm distance between m vectors. This can be done using a 2D assignment algorithm such as auction [11].

2. For each vector $\boldsymbol{y}_i$, let $\alpha_i = c$ if no assignment is made, or the minimum of c and the p-norm distance of the assignment.

3. Compute the p-th order average $d_p^{(c)}(X,Y) = ((1/n)\sum_{i=1}^{n}\alpha_i^p)^{1/p}$.

3.2. TEST CASE 1

The first test case examines a data rich scenario, in which measurement arcs of one hour in length are used, separated by five hour data gaps. Measurements are taken at five minute intervals. The initial conditions are such that all four objects are in the field of view at the initial time, and one object crosses out of the FOV during two of the five measurement windows. The sensor probability of detection is assumed to be $p_{D,\text{sensor}} = 0.8$ for all objects within the FOV, allowing for the possibility of missed detections due to poor lighting conditions or imperfect sensor capability. Clutter is added at each measurement time, and is assumed to be Poisson-distributed in number with a mean rate of 10 returns per epoch, and uniformly distributed in the sensor FOV. Fig. 2(a) provides data for the measurement conditions, indicating which objects are detected at each time, and the number of missed detections and clutter returns throughout the simulation. Fig. 2(b) provides a visualization of the sensor FOV and all satellite-generated measurements collected during the simulation.

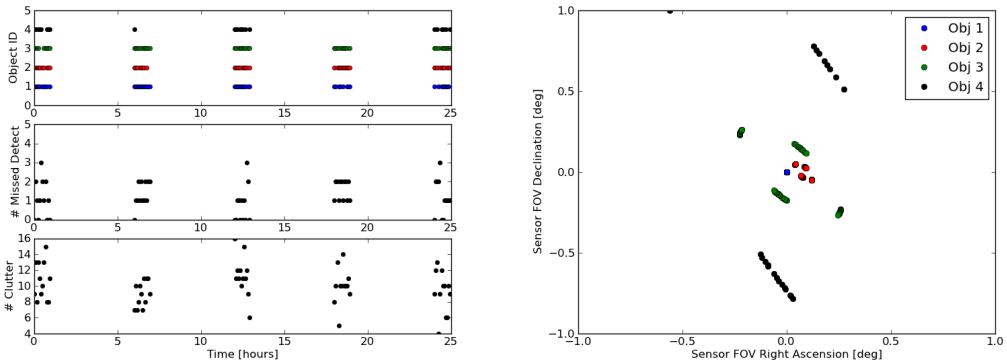

(a) Detected Objects, Missed Detections, Clutter

(b) Measurements in Sensor FOV

Fig. 2. Test Case 1 Measurement Data

From Fig. 2(a) it is apparent that most objects are in the FOV throughout the simulation. The missed detections are a result of modeling an imperfect sensor with $p_{D,\text{sensor}} = 0.8$. Object 4 crosses out of the FOV at the beginning of the second measurement arc, and only one detection of the object is registered for that pass. Fig. 2(b) shows that this measurement is at the edge of the FOV, at approximately $(-0.6°, 1.0°)$. Note that Object 1 appears stationary in the FOV because the sensor is defined to point at it for each measurement epoch.

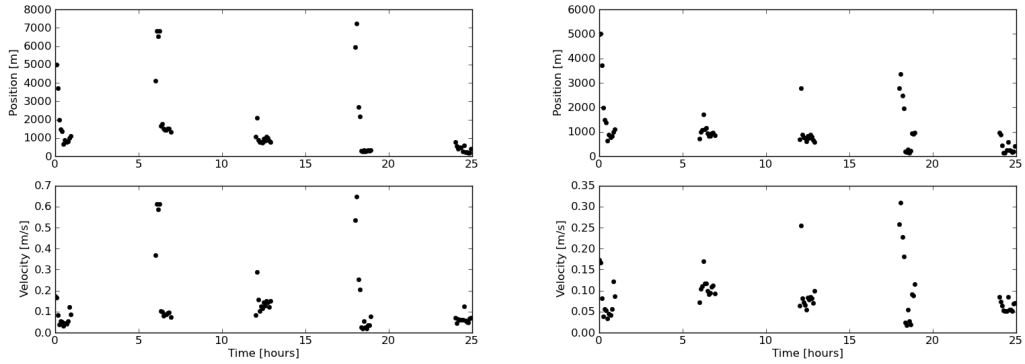

(a) State Dependent p_D Filter

(b) Covariance Dependent p_D Filter

Fig. 3. Test Case 1 Filter OSPA

Results of filter execution are provided in Figs. 3-4. The OSPA and cardinality plots demonstrate that both filters maintain custody of all four objects, producing a final estimate with errors on the order of 100 meters. The filter using

the state-dependent p_D calculation from Eq. (16) does produce larger state errors during the times when Object 4 is out of view, notably in the second measurement arc when the object is near the edge of the FOV, indicating that the covariance dependent p_D calculation may offer an advantage in this situation as expected. The effect is fairly limited, however, and by the end of the second measurement arc errors are on the order of 1 km.

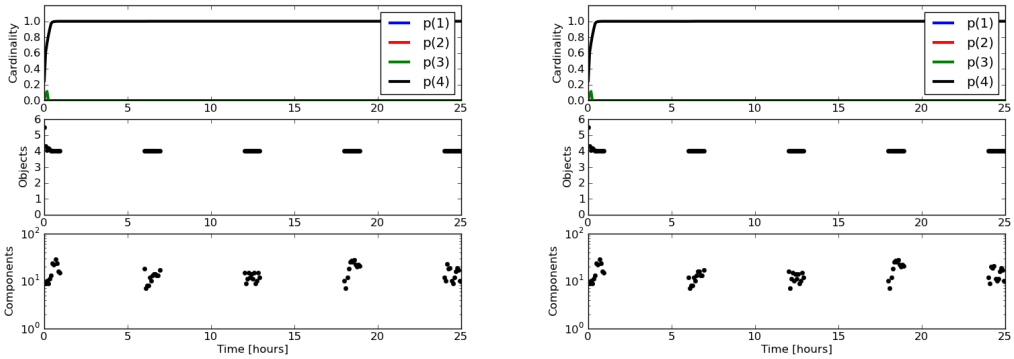

(a) State Dependent p_D Filter

(b) Covariance Dependent p_D Filter

Fig. 4. Test Case 1 Cardinality, Estimated Number of Objects, and Components

Fig. 4 provides the cardinality and estimated number of objects, computed from the PHD using Eq. (9), as well as the number of GMM components throughout the simulation. Both filters correctly converge on an estimate of 4 objects during the first measurement arc and maintain this estimate through the duration of the test case.

3.3. TEST CASE 2

The second test case employs 10 minute measurement arcs separated by 8-hour gaps, and is intended as a step toward a more realistic SSA scenario. Measurements are taken at a higher rate, once per minute, than in Case 1. Clutter measurements are added using the same parameters as Case 1. Due to the limited number of observations, it is necessary to set $p_{D,\text{sensor}} = 1$ to keep all measurements when objects are in the field of view and produce a reasonable filter estimate. Fig. 5 provides details on the measurement conditions, from which it can be seen that only one data arc contains measurements for all 4 objects, and often only one or two are in view.

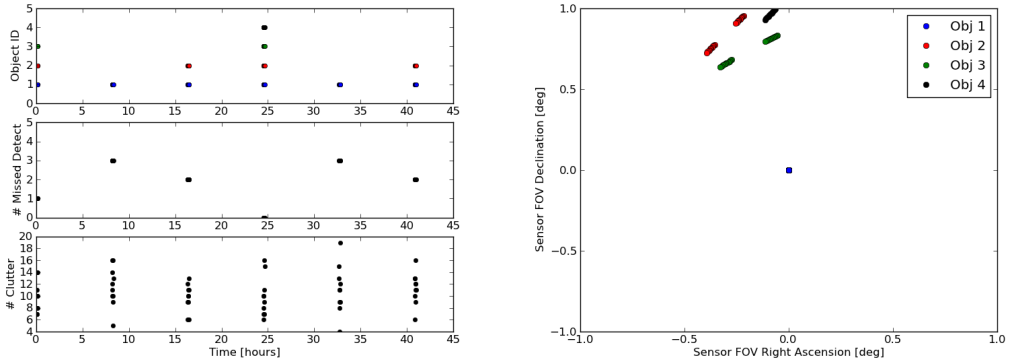

(a) Detected Objects, Missed Detections, Clutter

(b) Measurements in Sensor FOV

Fig. 5. Test Case 2 Measurement Data

OSPA results are given in Fig. 6 and both filters achieve similar results. There is an instance of higher position errors in the filter using the covariance-dependent p_D calculation, but by the end of the simulation both filters have the same errors, on the order of 5 km. While these errors are large, they do represent a reduction by an order of magnitude

from the original errors, which are on the order of 50 km RSS in position. Both filters estimate 4 objects present for most of the simulation, with an exception during the fourth measurement arc when the cardinality and estimated number of targets are momentarily reduced, see Fig. 7. It is possible this is caused by an incorrect p_D calculation when Object 4 is near the edge of the FOV, as seen in Fig. 5, though no detections are actually missed. The error is corrected within the same measurement arc and held for the remainder of the test case.

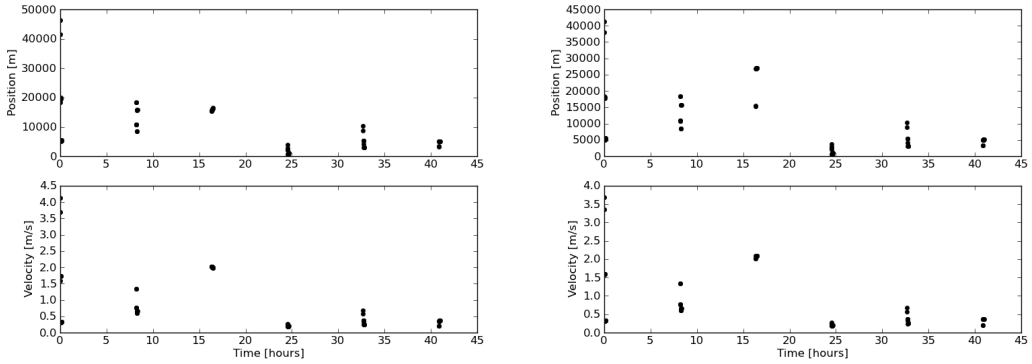

(a) State Dependent p_D Filter

(b) Covariance Dependent p_D Filter

Fig. 6. Test Case 2 Filter OSPA

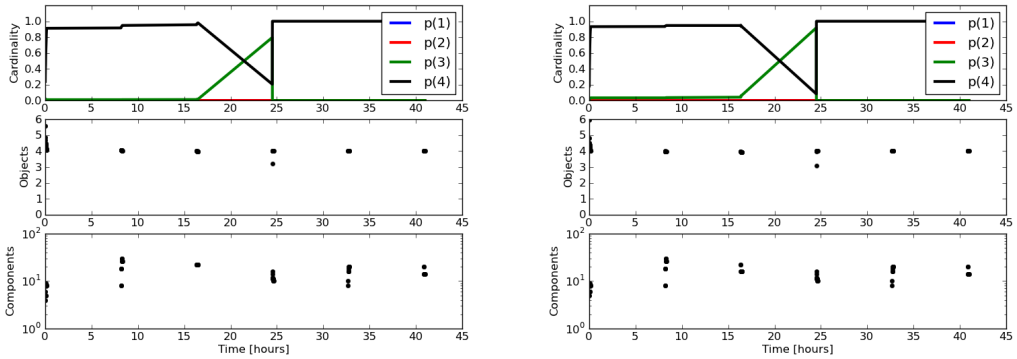

(a) State Dependent p_D Filter

(b) Covariance Dependent p_D Filter

Fig. 7. Test Case 2 Cardinality, Estimated Number of Objects, and Components

4. CONCLUSIONS AND FUTURE WORK

This paper presents background, analysis, and simulation results examining two different methods of computing the probability of detection in the CPHD filter. From the equations, it is clear that incorrectly computing p_D can have a significant impact on the filter's ability to maintain custody of objects, though two relatively simple approaches were tested and both successfully addressed the issue.

From initial analysis, it was hypothesized that computation of p_D using a covariance-dependent approach would work better than a state-dependent calculation, particularly in cases where objects are near the edge of the field of view. While there is a small difference in errors computed in the first test case considered, generally the effect is not very large. In considering the size of the field of view, a 2 degree sensor FOV maps to over 1000 km in each axis at GEO, and with position errors and uncertainties down to km or 100 meter level, there is not much difference between the p_D computed from the two methods. The situation illustrated in Fig. 1 is not realistic for objects that are being

tracked with these level of estimation errors. In addition, objects near the edge of the sensor FOV are not likely to remain there for many epochs in a row, as they will drift in or out of the FOV, or the sensor may be retasked entirely, so effects on filter performance are not significant.

With this consideration, there may be other situations in which a covariance-dependent p_D calculation would be important. For newly acquired objects, initial orbit determination using a constrained admissible region can produce a large initial uncertainty which may occupy a significant portion of the sensor FOV [12]. Additionally, it is noted that the cases considered did not include data gaps long enough for uncertainties to become non-Gaussian, nor did they include mismodeled accelerations. In the end, the AEGIS propagation and uncertainty modeling was not required and the CPHD filter predicted and corrected single GMM components for each object. Longer data gaps and mismodeled accelerations could cause uncertainties to grow or become non-Gaussian, which may warrant use of more careful calculation of the expected p_D. Future work will consider these cases in an effort to more realistically model SSA measurement scenarios to determine the importance of computing probability of detection.

5. REFERENCES

1. Ronald P. S. Mahler. Multitarget bayes filtering via first-order multitarget moments. *IEEE Transactions on Aerospace and Electronic Systems*, 39(4):1152–1178, October 2003.

2. Ronald P. S. Mahler. *Statistical Multisource-Multitarget Information Fusion*. Artech House, Norwood, MA, 2007.

3. Ba-Ngu Vo and Wing-Kin Ma. The gaussian mixture probability hypothesis density filter. *IEEE Transactions on Signal Processing*, 54(11):4091–4104, 2006.

4. Ba-Tuong Vo, Ba-Ngu Vo, and Antonio Cantoni. Analytic implementations of the cardinalized probability hypothesis density filter. *IEEE Transactions on Signal Processing*, 55(7):3553–3567, 2007.

5. M. Ulmke, Ozgur Erdinc, and Peter Willett. Gaussian mixture cardianlized phd filter for ground moving targets. In *Proceedings 10th International Conference on Information Fusion*, Quebec, Canada, July 2007.

6. Kyle J. DeMars, Robert H. Bishop, and Moriba K. Jah. Entropy-based approach for uncertainty propagation of nonlinear dynamical systems. *Journal of Guidance, Control, and Dynamics*, 36(4):1047–1057, July-August 2013.

7. Kyle J. DeMars. *Nonlinear Orbit Uncertainty Prediction and Rectification for Space Situational Awareness*. PhD thesis, University of Texas at Austin, Austin, TX, 2010.

8. Simon J. Julier and Jeffrey K. Uhlmann. A new extension of the kalman filter to nonlinear systems. In *Proc. Aerosense: 11th International Symposium on Aerospace/Defense Sensing, Simulations and Controls*, pages 182–193, 1997.

9. Chee Sing Lee, Daniel E. Clark, and Joaquim Salvi. Slam with dynamic targets via single-cluster phd filtering. *IEEE Journal of Selected Topics in Signal Processing*, 7(3):543–552, June 2013.

10. Dominic Schuhmacher, Ba-Tuong Vo, and Ba-Ngu Vo. A consistent metric for performance evaluation of mulit-object filters. *IEEE Transactions on Signal Processing*, 56(8):3447–3457, August 2008.

11. Samuel Blackman and R. Popoli. *Design and Analysis of Modern Tracking Systems*. Artech House, Norwood, MA, 1999.

12. Brandon A. Jones, Steven Gehly, and Penina Axelrad. Measurement-based birth model for a space object cardinalized probability hypothesis density filter. In *AIAA/AAS Astrodynamics Specialist Conference*, San Diego, CA, August 2014.