

Ground-based Tracking of Geosynchronous Space Objects with a GM-CPHD Filter

Brandon A. Jones, Noble Hatten, Nicholas Ravago, and Ryan P. Russell

The University of Texas at Austin

ABSTRACT

This paper presents a multi-target tracker for space objects near geosynchronous orbit using the Gaussian Mixture Cardinalized Probability Hypothesis Density (CPHD) filter. Given limited sensor coverage and more than 1,000 objects near geosynchronous orbit, long times between measurement updates for a single object can yield propagated uncertainties sufficiently large to create ambiguities in observation-to-track association. Recent research considers various methods for tracking space objects via Bayesian multi-target filters, with the CPHD being one such example. The implementation of the CPHD filter presented in this paper includes models consistent with the space-object tracking problem to form a new space-object tracker. This tracker combines parallelization with efficient models and integrators to reduce the run time of Gaussian-component propagation. To allow for instantiating new objects, the proposed filter uses a variation of the probabilistic admissible region that adheres to assumptions in the derivation of the CPHD filter. Finally, to reduce computation time while mitigating the so-called “spooky action at a distance” phenomenon in the CPHD filter, we propose splitting the multi-target state into distinct, non-interacting populations based on the sensor’s field of view. In a scenario with 700 near-geosynchronous objects observed via three ground stations, the tracker maintains custody of initially known objects and instantiates tracks for newly detected ones. The mean filter estimation after a 48 hour observation campaign is comparable to the measurement error statistics.

1. INTRODUCTION

Traditional single-target tracking uses a set of measurements to estimate the state and uncertainty of one object. This paradigm is complicated if tracking of more than one object is required. In this case, a separate algorithm must be used to correlate each measurement to a specific object, and object birth and death must be considered. This added complexity is particularly relevant to Space Situational Awareness (SSA), in which the estimation and prediction of an ever-growing population of Space Objects (SOs) is necessary to retain custody of vital Earth-orbiting assets [1]. For example, in the GEosynchronous Orbit (GEO) regime, ambiguous observation-to-track associations are possible due to the potential for long times between measurement updates for an individual SO and resulting increases in state uncertainty.

One alternative to a bank of distinct, single-target filters with a hypothesis table is to combine the processes of data association, cardinality estimation, and state estimation into a single probabilistic framework (e.g., see [2]). Multi-target filters attempt to simultaneously estimate (i) the number of objects in a given space and (ii) a Probability Density Function (PDF) of the state of each object. In the context of SSA, improved multi-target tracking can reduce the number of Un-Correlated Tracks (UCTs) and improve state estimates [3–7]. This information, in turn, can benefit a wide array of applications, including conjunction assessment, SO characterization, and sensor tasking. Recent research in this area applies formulations or approximations of the Bayes multi-target filter to SO tracking [5–8].

One such multi-target filter is the Cardinalized Probability Hypothesis Density (CPHD) filter [2]. Based on FInite Set STatistics (FISST) and an approximation to the Bayes multi-target filter, the CPHD filter accomplishes the goals of a multi-target filter by producing (i) a cardinality distribution that approximates the number of objects and (ii) an intensity function that describes the density of objects in the single-target space [9]. In addition to handling ambiguities in data association, the CPHD provides a mathematical framework to include models for object birth, probability of detection, and clutter measurements.

This paper presents a multi-target tracker for SOs near GEO based on a Gaussian Mixture (GM) CPHD filter [10]. New objects are instantiated using a partially uniform birth (PUB) model [11] combined with the probabilistic admissible region (PAR) [12]. Additionally, this GM-CPHD implementation splits the multi-target state into distinct, non-interacting populations based on the sensor’s field of view (FOV). This modification decreases both complexity and computation time while simultaneously mitigating the “spooky action at a distance” phenomenon [13]. A method

for generating the parameterization of the multi-target state for each population as a function of the CPHD multi-target state estimate and sensor FOV is also presented.

The work presented in this paper represents a confluence of independently developed tools for SO state propagation, uncertainty propagation, and multi-target filtering. Further, the tools have been produced so as to facilitate the use of results in higher-level applications that rely on, for example, accurate observation-to-track associations (e.g., SO characterization).

The paper commences with a description of the GM-CPHD filter with the PUB model. Individual components of the SO tracking algorithm are then discussed, with some emphasis placed on state and uncertainty propagation due to the significant effect of system dynamics on overall filter run time. Finally, filter results are shown for a randomly generated, near-GEO catalog of 700 SOs that is representative of the publicly available Two-Line Element (TLE) catalog. The capabilities of the proposed birth model are demonstrated through instantiation and solution convergence for newly observed, randomly generated SOs.

2. THE GM-CPHD FILTER WITH A PARTIALLY-UNIFORM BIRTH MODEL

Bayes multi-target filtering provides an analog for the Bayes single-target filter that yields a probabilistic representation of a multi-target state. Prediction of the multi-target state accounts for variations in the number of targets, i.e., spontaneous birth, spawning, and object death, while the update accounts for detection probability, clutter, and the possibility of more than one observation of a given object. Let the single-target state and measurement spaces be denoted by \mathbb{X} and \mathbb{Z} , respectively. In the context of Random Finite Set (RFS) based multi-target filters, we represent the multi-target state and observation by

$$X_k = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{N_k}\} \subset \mathbb{X}, \quad (1)$$

$$Z_k = \{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_{M_k}\} \subset \mathbb{Z}, \quad (2)$$

where $\mathbf{x}_i \in \mathbb{X}$ is a realized random vector for a single-target state, $\mathbf{z}_i \in \mathbb{Z}$ is a realized measurement, and the sets X_k and Z_k have cardinality N_k and M_k , respectively, at time t_k . For this work, we assume that errors in \mathbf{z} are Gaussian distributed with covariance matrix \mathbf{R}_k and a known single-target observation model

$$\mathbf{z} = h(\mathbf{x}; \mathbf{o}, t), \quad (3)$$

where \mathbf{o} are parameters associated with the observer (e.g., ground-station location) and are omitted without loss of clarity. When the probability of detection $p_{D,k} < 1$ or clutter measurements may be produced by a given sensor, then N_k does not strictly equal M_k . The probability density function (PDF) of X_k is denoted by $\pi_k(X)$.

Like the single-target Bayes filter, the multi-target filter requires models for prediction and update of the PDF. In the context of an RFS-based approach, prediction of π_k from t_{k-1} to t_k is described by the Chapman-Komolgorov equation [2]

$$\pi_{k|k-1}(X_{k|k-1}) = \int f(X_{k|k-1}|X) \pi(X) \delta(X), \quad (4)$$

which leverages the set integral

$$\int f(X) \delta X = \sum_{i=0}^{\infty} \frac{1}{i!} \int f(\{x_1, x_2, \dots, x_i\}) d(x_1, x_2, \dots, x_i), \quad (5)$$

and temporal variations in the multi-target state are described by the Markov transition kernel $f(X_{k|k+1}|X)$. This transition kernel describes changes in the multi-target state due to state propagation and variation in the number of targets (e.g., birth and death). Given a realized RFS Z_k , the Bayes measurement update is

$$\pi_k(X_k|Z_k) = \frac{g(Z_k|X_{k|k-1}) \pi_{k|k-1}(X_{k|k-1})}{\int g(Z_k|X) \pi_{k|k-1}(X) \delta X}, \quad (6)$$

where $g(Z|X)$ is the multi-target likelihood function for a given sensor. The function $g(Z|X)$ accounts for $p_{D,k}$ and clutter.

The following sections describe the specific multi-target tracker used in this work. While the closed-form RFS-based filter presented in [2] is often intractable, several simplifications exist that allow for practical applications.

This includes the Probability Hypothesis Density (PHD) filter [14, 15], R-FISST [16], and the various labeled multi-Bernoulli filters [17–20]. The current work uses the Cardinalized Probability Hypothesis (CPHD) filter, which is a partial second-order form of the Bayes multi-target filter [2]. The remainder of this section describes the form of the CPHD used in the proposed space-object tracker.

In the CPHD, the multi-target state X_k is parameterized by a discrete cardinality distribution $\rho_k(n)$ describing the number of targets and a PHD $v_k(\mathbf{x})$ representing the density of targets in \mathbb{X} . This PHD is also called the single-target *intensity*. This work uses the Gaussian mixture CPHD filter proposed in [10] combined with a partially-uniform birth model to instantiate newly discovered targets [11]. The intensity is approximated by a Gaussian mixture, i.e.,

$$v_k(\mathbf{x}) = \sum_{i=1}^{J_k} w_k^{(i)} \mathcal{N}(\mathbf{x}; \bar{\mathbf{x}}_k^{(i)}, \mathbf{P}_k^{(i)}), \quad (7)$$

where each of the J components has a weight $w^{(i)} > 0$, a mean $\bar{\mathbf{x}}^{(i)} \in \mathbb{X}$, and a corresponding covariance matrix $\mathbf{P}^{(i)}$. The Partially-Uniform Birth (PUB) model presented in [11] assumes a newborn target may appear anywhere in the sensor FOV with uniform probability, which shares similarities with SO tracking. The PUB model for the newborn target intensity is

$$v_{B,k}(\mathbf{x}) = w_{B,k} \mathcal{U}(\boldsymbol{\theta}(\mathbf{x}); \mathcal{B}) \sum_{i=1}^{J_B} w_{b,k}^{(i)} \mathcal{N}(\boldsymbol{\phi}(\mathbf{x}); \bar{\boldsymbol{\phi}}_k^{(i)}, \mathbf{P}_{\phi,k}^{(i)}), \quad (8)$$

where the subscript B denotes spontaneous birth, $w_{B,k}$ is the mean number of new targets appearing at t_k , the intensity in the FOV is uniform with bounds \mathcal{B} , $\boldsymbol{\theta}$ is the observable part of the single-target state, and $\boldsymbol{\phi}$ is the unobservable portion with the dependence on \mathbf{x} removed for the sake of notational clarity in future discussions. When using optical observations for SO tracking, $\boldsymbol{\theta}$ could be topocentric right ascension and declination while $\boldsymbol{\phi}$ includes the range, range-rate, and angle rates. The model in Eq. (8) assumes a Gaussian mixture in the unobservable direction where $w_{b,k}^{(i)} > 0$, $\sum w_{b,k}^{(i)} = 1$, and the given mean and covariance are $\bar{\boldsymbol{\phi}}^{(i)}$ and $\mathbf{P}_{\phi,k}^{(i)}$, respectively. Like the presentation of the CPHD filter in [21], the following describes the prediction and update of the surviving target intensity in terms of \mathbf{x} and the birth model as $[\boldsymbol{\theta}^T \ \boldsymbol{\phi}^T]^T$.

To distinguish between intensities for surviving and newborn targets, the PUB-CPHD filter augments the state vector with a binary variable β . The case $\beta = 1$ denotes a component of the birth intensity and $\beta = 0$ corresponds to a surviving target. Given the previous cardinality estimate $\rho_{k-1}(n)$,

$$\rho_{k|k-1}(n) = \sum_{i=0}^n \rho_{B,k}(n-i) \sum_{j=i}^{n_{\max}} C_i^j \rho_{k-1}(j) p_S^i (1-p_S)^{j-i}, \quad n = 0, \dots, n_{\max}, \quad (9)$$

where $\rho_{B,k}$ is the cardinality distribution for new targets with binomial coefficient $C_i^j = j!/(i!(j-i)!)$. Prediction of the surviving target intensity is

$$v_{k|k-1}(\mathbf{x}, 0) = \sum_{i=1}^{J_{k-1}} w_{k|k-1}^{(i)} \mathcal{N}(\mathbf{x}; \bar{\mathbf{x}}_{k|k-1}^{(i)}, \mathbf{P}_{k|k-1}^{(i)}) \quad (10)$$

with

$$w_{k|k-1}^{(i)} = p_S w_{k-1}^{(i)}, \quad (11)$$

$J_{k|k-1} = J_{k-1}$, and the mean $\bar{\mathbf{x}}_{k|k-1}^{(i)}$ and covariance $\mathbf{P}_{k|k-1}^{(i)}$ of the predicted Gaussian component are produced via the unscented transform and a given single-target dynamics model. This work employs the square-root unscented transform presented in [22] with constants $\alpha = 1$ and $\beta = 2$ (here using the notation in the provided reference). The predicted birth intensity is then

$$v_{k|k-1}(\mathbf{x}, 1) = v_{B,k}(\mathbf{x}). \quad (12)$$

From t_k to t_{k+1} , newborn targets become surviving targets and $\beta = 1$ for all surviving targets. The CPHD measurement update when using the PUB model is

$$\rho_k(n) = \frac{\Psi_0 [Z_k, v_{k|k-1}] (n)}{\langle \Psi_0 [Z_k, v_{k|k-1}], \rho_{k|k-1} \rangle} \rho_{k|k-1}(n), \quad (13)$$

$$\begin{aligned} v_k(\mathbf{x}, 0) &\approx \sum_{i=1}^{J_{k|k-1}} w_{m,k}^{(i)} \mathcal{N}(\mathbf{x}; \bar{\mathbf{x}}_{k|k-1}^{(i)}, \mathbf{P}_{k|k-1}^{(i)}) \\ &\quad + \sum_{\mathbf{z} \in Z} \sum_{i=1}^{J_{k|k-1}} w_{d,k}^{(i)}(\mathbf{z}) \mathcal{N}(\mathbf{x}; \bar{\mathbf{x}}_k^{(i)}, \mathbf{P}_k^{(i)}), \end{aligned} \quad (14)$$

$$v_k(\mathbf{x}, 1) \approx \sum_{\mathbf{z} \in Z} w_{B,k}(\mathbf{z}) \mathcal{N}(\boldsymbol{\theta}(\mathbf{x}); \mathbf{z}, \mathbf{R}) \times \sum_{i=1}^{J_B} w_{b,k}^{(i)} \mathcal{N}(\phi(\mathbf{x}); \bar{\phi}_k^{(i)}, \mathbf{P}_{\phi,k}^{(i)}), \quad (15)$$

where

$$\chi = \frac{\langle \Psi_1 [Z_k, v_{k|k-1}], \rho_{k|k-1} \rangle}{\langle \Psi_0 [Z_k, v_{k|k-1}], \rho_{k|k-1} \rangle}, \quad (16)$$

$$\chi(\mathbf{z}) = \frac{\langle \Psi_1 [Z_k - \{\mathbf{z}\}, v_{k|k-1}], \rho_{k|k-1} \rangle}{\langle \Psi_0 [Z_k, v_{k|k-1}], \rho_{k|k-1} \rangle}, \quad (17)$$

$$\Psi_u [Z, v] (n) = \sum_{j=0}^{\min(|Z|, n)} (|Z| - j)! \rho_\kappa (|Z| - j) P_{j+u}^n \times \frac{((1 - p_D) \langle 1, v \rangle)^{n-(j+u)}}{(\langle 1, v \rangle + w_{B,k})^n} e_j (\Xi(Z, v)), \quad (18)$$

$$\Xi(Z, v) = \left\{ \frac{\langle 1, \kappa_k \rangle}{\kappa_k(\mathbf{z})} \left(\frac{w_{B,k}}{V_B} + p_D \sum_{i=1}^{J_{k|k-1}} w_{k|k-1}^{(i)} q_k^{(i)}(\mathbf{z}) \right) : \mathbf{z} \in Z \right\}, \quad (19)$$

$$q_k^{(i)}(\mathbf{z}) = \mathcal{N}(\mathbf{z}; \bar{\mathbf{z}}_k^{(i)}, \mathbf{P}_{\mathbf{z}\mathbf{z}}^{(i)}), \quad (20)$$

and the updated weights are

$$w_{m,k}^{(i)} = (1 - p_D) \chi w_{k|k-1}^{(i)}, \quad (21)$$

$$w_{d,k}^{(i)}(\mathbf{z}) = p_D q_k^{(i)}(\mathbf{z}) \frac{\langle 1, \kappa_k \rangle}{\kappa_k(\mathbf{z})} \chi(\mathbf{z}) w_{k|k-1}^{(i)}, \quad (22)$$

$$w_{B,k}(\mathbf{z}) = \frac{\langle 1, \kappa_k \rangle}{\kappa_k(\mathbf{z})} \chi(\mathbf{z}) \frac{w_{B,k}}{V_B}. \quad (23)$$

The update of $v_k(\mathbf{x}, \beta)$ in Eqs. (14) and (15) is an approximation since their derivation assumes that

$$\mathcal{N}(\boldsymbol{\theta}; \mathbf{z}, \mathbf{R}) 1_B(\boldsymbol{\theta}) \approx \mathcal{N}(\boldsymbol{\theta}; \mathbf{z}, \mathbf{R}), \quad (24)$$

where

$$1_B(\mathbf{z}) = \begin{cases} 1 & \mathbf{z} \in \mathcal{B}, \\ 0 & \mathbf{z} \notin \mathcal{B}. \end{cases} \quad (25)$$

For the cases where the uncertainty is small compared to the FOV \mathcal{B} , the simplification in Eq. (24) holds. A newborn target becomes a surviving target after first appearing, hence, before the CPHD time update

$$v_k(\mathbf{x}, 0) = v_k(\mathbf{x}, 0) + v_k(\mathbf{x}, 1). \quad (26)$$

This sum of intensities is then used in Eq. (10) for the surviving target prediction. To maintain tractability due to an increase in J_k , GM components are merged and pruned as outlined in [10] with additional component management methods discussed in later sections.

Upon updating the cardinality and intensity using the CPHD, target states may be extracted to form X_k . Using the cardinality distribution

$$N_k = \arg \max_n \rho_k(n), \quad (27)$$

the estimated multi-target state is

$$\hat{X}_k = \left\{ \bar{\mathbf{x}}_k^{(i)} \right\}_{i=1}^{N_k}, \quad (28)$$

where $\bar{\mathbf{x}}_k^{(i)}$ are the mean vectors of the N_k components in $v_k(\mathbf{x}, 0)$ with the largest weights.

The “spooky action” effect in the CPHD results from Eqs. (13)-(23). The additional factor added by the cardinality through the measurement set likelihoods Ψ_u enforces that $\sum w^{(i)} = \mathbb{E}[\rho]$. However, when a missed detection occurs, the corresponding component weight decreases by the factor $1 - p_D$ in Eq. (21). If, for example, the cardinality indicates a high probability of two targets, but one is missed, then the weight on the other target increases to compensate. Hence, targets separated in space such that they are statistically independent may influence each other in the CPHD measurement update. This is known as “spooky action at a distance” [2] and may be mitigated through the designation of distinct, statistically non-interacting populations. Our method for mitigating this is discussed in the following sections.

3. SPACE OBJECT TRACKER

The described CPHD filter provides a mathematical framework for the proposed space-object tracker, which utilizes several tools to reduce computation time, improve robustness, and increase portability. Additionally, they address elements specific to space-object tracking when compared to the classical CPHD literature (e.g., see [10, 11]). The following sections begin with specific elements required for the tracker, and conclude with a description of the integrated toolset. Except where noted, software is implemented in MATLAB with some computationally expensive operations written in C/C++ and Fortran.

3.1. COMPUTATIONALLY EFFICIENT ORBIT STATE AND UNCERTAINTY PROPAGATION

The predicted mean and covariance of each GM component – required for the computation of $v_{k|k-1}$ – are calculated using the adaptive entropy-based Gaussian-mixture information synthesis (AEGIS) method [23]. In the AEGIS formulation, the multi-target state is represented using a GM, the number of components of which is allowed to increase over time based on the nonlinearity of the system dynamics. The state and uncertainty of each component are propagated in time using the unscented transform [22]. The orbital equations of motion are formulated using Cowell’s method, and all sigma points for a given component are placed in a single state vector and propagated simultaneously. The linearly propagated differential entropy, propagation of which is required by the AEGIS method, is also placed in the state vector. For the case of \mathbf{x} being position and velocity, this results in a 79-element state vector for each GM component: 13 sigma points of six states each, joined by the differential entropy. While this strategy precludes the possibility of parallelizing the propagation of each sigma point for a given component, the more attractive parallelization at the component level remains viable.

The AEGIS method detects nonlinearity by comparing the propagated differential entropy with a differential entropy value calculated at the end of each propagation step based on the propagated covariance via the sigma points. If the difference between the two differential entropy values is greater than a user-defined threshold, then the GM component splits into multiple components, each of which is then propagated individually. Though an event-finding algorithm may be used to more precisely determine the time at which the differential entropy tolerance is violated, no such technique is implemented in the current work due to the required small propagation step sizes. (See further discussion of step sizes below.)

It is noted that the linearized time derivative of the differential entropy for a GMM component is equal to the trace of the Jacobian of the mean orbital state of that component. In general, this relation necessitates the calculation of the full Jacobian. However, the trace of the Jacobian of a conservative force (e.g., gravity) is zero, so only the Jacobians of nonconservative contributions to the dynamics model need be evaluated [24].

State and uncertainty propagation dominate overall computation time of the GM-CPHD filter because of the large number of GMM components and non-trivial orbital dynamics. Therefore, several measures are taken to prioritize efficiency. First, while the majority of filter operations are performed using MATLAB, propagation code is written in Fortran to take advantage of the speed of the compiled language. The Fortran routines are compiled into a shared library, which are loaded into the MATLAB workspace. From there, the propagation routines are then invoked directly from MATLAB, and output variables are returned to the MATLAB workspace. Using this paradigm, multi-CPU parallelization at the GM component propagation level is accomplished via the MATLAB `parfor` construct.

For many astrodynamics applications, high-order, variable-step-size linear multistep or explicit Runge-Kutta methods are typically used to propagate states and uncertainties [25–27]. However, for the current application, the intervals between measurements – and the corresponding propagation durations – are short compared to the relatively slowly and uniformly varying dynamics of the near-GEO environment. As a result, the step sizes for typical high-order methods would be artificially truncated from their optimal values to accommodate measurement updates, leading to inefficiency. A fourth-order, fixed-step, explicit Runge-Kutta method is therefore used instead [28]. To prevent information loss due to frequent conversions of a covariance matrix to or from sigma points, the propagated state vector is stored and only regenerated after a measurement update of that component.

The efficiencies of the individual elements of the single-state dynamics model are addressed, as well. Ephemeris calculations for the positions of the Sun and Moon and for the orientation of the Earth are performed using cubic-spline interpolations of publicly available NAIF SPICE ephemeris data [29–31]. Coefficients are precalculated for the timespan of interest and saved in binary files. During filter initialization, the coefficients are loaded into RAM and stored for the duration of the filter runtime. This strategy has been shown to improve ephemeris retrieval speeds by one or more orders of magnitude compared to the common technique of directly calling SPICE routines, and provides adequate accuracy for the current application. Additionally, because ephemeris calculations are functions of time only, the simultaneous propagation of all sigma points for a given component allows for a single ephemeris calculation to be reused in the dynamics models of all 13 sigma points.

The spherical harmonics coefficients used to calculate the geopotential are also loaded into RAM a single time during filter initialization to avoid unnecessary repetitive operations [32]. It is noted that interpolation-based geopotential calculation methods are capable of significantly outperforming spherical harmonics implementations for high-degree/order fields [33, 34]. However, the spherical harmonics formulation is used in the current work due to the relatively low-degree/order field required for acceptable propagation of the near-GEO catalog under consideration.

The final element of the single-state dynamics model is solar radiation pressure (SRP). The SRP acceleration is calculated using a standard cannonball SO model with constant area-to-mass-ratio and coefficient of reflectivity [35]. A conical Earth shadow model is used. Unlike gravitational forces, SRP is nonconservative; however, the trace of the SRP Jacobian in the current model is zero, so the calculation is avoided for the propagation of differential entropy.

The single-state dynamics model is summarized in Table 1.

Table 1. Single-state dynamics model.

Force	Implementation	Notes
Geopotential	10×10 spherical harmonics potential	Coefficients loaded into RAM at initialization; cubic spline Earth orientation calculation
Sun and Moon gravity	Point-mass potential	Cubic spline ephemerides
SRP	Cannonball SO with conical Earth shadow	Jacobian not required

3.2. MODELING BIRTH WITH THE PROBABILISTIC ADMISSIBLE REGION

Methods based on the admissible region provide an initial probabilistic description of a newly-detected space object [12,36]. The admissible region approach describes the initial state using a measurement z with physics-based constraints to limit the possible solutions in the unobserved directions (e.g., see [37,38]). In the case of four-dimensional (4D) optical observations, i.e., observations of angles and angle-rates, the constraints reduce the space of admissible range and range-rate solutions. Typical constraints include bounds on the semimajor axis and/or eccentricity. The Constrained Admissible Region (CAR) approach assumes a uniform distribution of possible range and range-rate

pairs in the admissible space, which is approximated using a GM model to yield a PDF of the initial satellite state [23]. Alternatively, the Probabilistic Admissible Region (PAR) uses sampling-based methods to yield a more accurate approximation of the admissible region as a function of random inputs. For example, the PAR approach may be used to account for the measurement PDF and constraints (e.g., a uniform distribution of possible semimajor axis values). Generated samples are then employed to produce a GM model for the admissible region for use in a filter. See the provided references for details on the CAR and PAR and Fig. 1 for a comparison of the two resulting PDFs for a case outlined in [21].




Fig. 1. Comparison of the CAR (left) and PAR (right) for the same optical observation and semimajor axis/eccentricity constraints [21]

In [21], we demonstrate that the PAR may be combined with the PUB-CPHD filter while not violating any assumptions in the filter derivation. In this context, θ is the four-dimensional angle and angle-rate observation and the PAR produces a GM for the range and range-rate vector ϕ . The CAR cannot be used in this filter without violating assumptions in the birth model. Specifically, the birth model in Eq. (8) assumes independence between the realized θ and the Gaussian mixture for ϕ , but the CAR produces the range/range-rate GM as a function of the observation. However, the PAR may be defined as a function of the FOV and generated using the observation distribution $\mathcal{U}(\mathcal{B})$, which is consistent with the independence assumption. Hence, a PAR-based representation of the range/range-rate PDF may be used in the PUB-CPHD filter. Generation of the samples uses a C++ tool that interfaces with MATLAB via `mex`. Other software required to generate the PAR-based GM uses MATLAB. This process requires approximately 7 seconds per PAR, which can increase filter runtime when executed with each measurement scan.

To reduce computation time, this tracker only generates a PAR when the CPHD indicates a high likelihood of a new object. For a given measurement z , Eq. (23) produces the intensity of a new object via $w_{B,k}(z)$. We only create a PAR when this weight is greater than the GM component pruning threshold in the CPHD filter. Since any component with a weight below this value will be removed after the measurement update, the software skips the generation of the PAR for such cases. A low weight occurs when $\chi(z)$ is small, which corresponds to cases when the filter determines that the measurement is likely clutter or is produced by a surviving object.

3.3. 4D OBSERVATION GENERATION

The use of the PAR assumes a 4D observation of right ascension, declination, and their rates. These measurements are not provided using optical sensors, but may be generated via, for example, a least-squares fit of a time series of angle measurements over a short time duration [39]. Such methods assume that observations are grouped using a form of observation-to-observation association. As described in [40], linear models may be employed to link observations from a series of optical images. This process also reduces the number of clutter observations to a 0.01 probability (in

the case of [40]). We use a similar approach, but in the context of the GM-PHD filter to account for missed detections and clutter. The GM-PHD is similar to the previously presented GM-CPHD filter, but without the estimation of a cardinality distribution. More details of the GM form of the PHD filter are found in [15].

Given a collection of observations over a prescribed time period from a single sensor, the 4D observation generation tool produces angle and angle-rate measurements. The GM-PHD filter used for grouping observations detects objects and estimates their right-ascension, declination, angle rates, and accelerations state. Each optical image produces one or more right-ascension/declination observations that we group into an RFS Z_k . Given the sets $Z_k, Z_{k+1}, \dots, Z_{k+K}$, each are used in the GM-PHD filter with a linear dynamics model. Upon processing all available observation sets, the GM-PHD filter produces a collection of likely objects in the measurement space and the measurements associated with those objects. The least-squares algorithm for generating 4D measurements and the observation-error covariance matrix R_k [39] uses the observations associated with likely tracks. The time duration of the series and the number of frames determines the accuracy of the measurements.

This work also considered the use of a PHD smoother [41] to reduce false detections. However, testing of the tool demonstrated a lack of robustness in cases where the last measurement scan in the time series produces a missed detection for an object. When this occurs, the smoother reduces the likelihood that a track exists, thereby increasing the number of missed detections. Instead, testing demonstrated that inspecting the last two GM-PHD solutions for tracks yielded sufficient robustness for generating 4D measurements.

3.4. APPROXIMATE MULTI-SENSOR CPHD FILTER

A multi-sensor implementation of the CPHD filter allows for processing angles-only and 4D observations in the same filter. While the generation of 4D observations allows for using the PAR-based birth model, not all objects observed produce the minimum of three measurements required for the least-square regression. For example, a space object that exits the field of view after one angle-only measurement scan fails to yield sufficient information to approximate the angle-rate. However, that single data point may be used to update the multi-target state, and the tracker uses the 4D measurements and any unused 2D observations. Processing both 4D and 2D measurements from the same sensor at the same time with the same field of view requires the multi-sensor CPHD measurement update to reduce filter error.

This work uses the approximate multi-sensor CPHD filter proposed in [42]. The iterated-corrector approach, i.e., successive measurement updates for different sensors, allows for fusing measurements from different sensors in the CPHD. However, this approach can bias the filter when, for example, the sensors have different detection probabilities [43]. While the multi-sensor CPHD filter is computationally intractable [14], [42] proposed a principled approximation to the true solution. To illustrate the method, the CPHD intensity update Eqs. (14) and (15) may be combined and written as [44]

$$v_k(\mathbf{x}, \beta) = L_k^{(j)} \left(Z_k^{(j)} | \mathbf{x}, \beta \right) v_{k|k-1}(\mathbf{x}, \beta) \quad (29)$$

where L_k is an update operator and the j superscript denotes an update via the j th sensor. The iterated-corrector update may be expressed as [43]

$$v_k(\mathbf{x}, \beta) = L_k^{(1)} \left(Z_k^{(1)} | \mathbf{x}, \beta \right) L_k^{(2)} \left(Z_k^{(2)} | \mathbf{x}, \beta \right) \cdots L_k^{(s)} \left(Z_k^{(s)} | \mathbf{x}, \beta \right) v_{k|k-1}(\mathbf{x}, \beta) \quad (30)$$

for sensors $1, \dots, s$. In this case, the updated solution varies with the order of the $L_k^{(i)}$ terms. The approximate multi-sensor CPHD filter removes dependence on the order of update through a normalization term:

$$v_k(\mathbf{x}, \beta) = N_k \frac{L_k^{(1)} \left(Z_k^{(1)} | \mathbf{x}, \beta \right) L_k^{(2)} \left(Z_k^{(2)} | \mathbf{x}, \beta \right) \cdots L_k^{(s)} \left(Z_k^{(s)} | \mathbf{x}, \beta \right) v_{k|k-1}(\mathbf{x}, \beta)}{\sum_{\beta=0,1} \int L_k^{(1)} \left(Z_k^{(1)} | \mathbf{x}, \beta \right) L_k^{(2)} \left(Z_k^{(2)} | \mathbf{x}, \beta \right) \cdots L_k^{(s)} \left(Z_k^{(s)} | \mathbf{x}, \beta \right) S_{k|k-1}(\mathbf{x}, \beta) d\mathbf{x}} \quad (31)$$

where N_k is the updated number of objects and

$$S_{k|k-1}(\mathbf{x}, \beta) = \frac{v_{k|k-1}(\mathbf{x}, \beta)}{N_{k|k-1}}. \quad (32)$$

Full details on the method may be found in [42], and [45] demonstrates improved performance when using the approximate multi-sensor PHD filter instead of an iterated-corrector. While not demonstrated in the current paper, comparisons of the iterated-corrector and approximate CPHD solution are consistent with the results of [45].

3.5. CPHD FILTER WITH TWO POPULATIONS

The current implementation of the GM-CPHD filter leverages a dual-population model to: (i) mitigate the “spooky action at a distance” discussed previously, (ii) account for only a subset of objects in the FOV, and (iii) reduce computation time. The cardinality distribution $\rho(n)$ and the PHD $v(x)$ are split to yield two multi-target states. One state represents objects inside of the FOV while the second describes those believed to be outside of the sensor’s view. This requires a splitting of both the cardinality and the intensity based on the FOV boundaries. Upon splitting the state, we perform a measurement update using the CPHD parameters describing the population inside of the FOV. Hence, the “spooky effect” is isolated to components in a given population. This also reduces the measurement update computation time by using smaller values of n_{\max} , N , etc. Upon performing the measurement update, the two populations are recombined to form a single multi-target state for the objects in the surveillance region. The remainder of this section outlines the splitting and combination process.

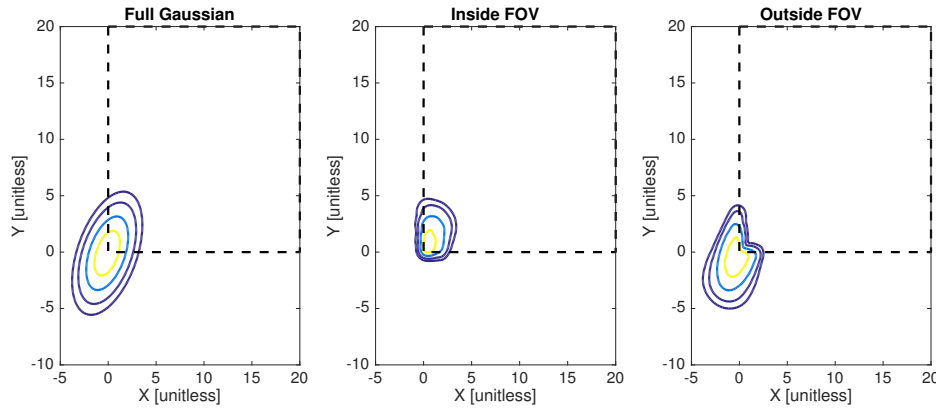


Fig. 2. Illustration of splitting a single Gaussian component to account for the FOV.

Simple splitting of the intensity based on the FOV compares the mean of the predicted measurement Gaussian with the sensor limits, i.e., if $\theta^{(i)} \in \mathcal{B}$ for the i th component, then the component is designated as in the FOV. However, for example in the case where the mean is near the FOV edge, there can be a significant portion of the density function both inside and outside of the FOV. Fig. 2 illustrates this in the leftmost image. (Note that this is an example for the sake of illustration and is not based on an SO simulation.) As seen in the figure, while the mean of the component at $(0, -0.1)$ is outside of the FOV, there is a non-zero density within the field of view that indicates the object may be detected. However, this component would be included as part of the population outside of the FOV using the initial definition. For tests with the GM-CPHD filter for space objects, the filter would assume a missed detection for any objects with true states just outside of the FOV. This result causes the filter to reduce the weight of the component, and thereby confidence in its existence, when there is no observation but a high probability of detection. Early testing demonstrated a loss of known objects due to such a phenomenon. Alternatively, the filter may generate a possible new track based on the realized measurement, thereby adding a filter bias. This work outlines a new method for splitting the PHD, i.e., the object density function, to better account for objects near the FOV boundary.

To reduce filter biases when the mean of a component is near the FOV boundary, we split a component in right-ascension/declination space to generate GM approximations of the portion inside and outside of the FOV. In [46], the authors split a component along the major principle axis based on the three-component library in [23]. Instead, the proposed method splits a given component along the principle axis closest to the shortest line between the mean and the FOV boundary. Splitting is done recursively until the probability of a realized measurement for each component is less than 0.05 or greater than 0.95. Fig. 1 also illustrates the resulting split of the GM components to isolate contributions inside and outside of the a given region. A case near the corner of the FOV is selected to illustrate that the solution for the two densities captures the boundary in the description of the new intensities. While not illustrated, the leftmost image is recovered (qualitatively) when merging all components output via the previously described procedure.

The cardinality distribution for a given population is computed using the weights of GM components in the corresponding intensity function. Assuming that weights approximate the probability that an object exists in the i th

component and $w^{(i)} \leq 1$ for all i , then the cardinality distribution generated via a Bernoulli approximation is [47]

$$\rho(n) = \left(\prod_{j=1}^{J_{FOV}} (1 - w^{(j)}) \right) \cdot e_j \left(\left\{ \frac{w^{(1)}}{1 - w^{(1)}}, \dots, \frac{w^{(J)}}{1 - w^{(J)}} \right\} \right) \quad (33)$$

where $e_j(\cdot)$ are the elementary symmetric functions. Computational complexity of these elementary symmetric functions is $\mathcal{O}(J^2)$, and generation of these functions uses a C tool to reduce computation time. This efficiency gain is critical for cases in the numeric results section where $J = 700$.

The combination of the updated CPHD state for objects inside the FOV with the cardinality and intensity for the outside objects is straightforward. For the intensity

$$v(\mathbf{x}) = v^{(I)}(\mathbf{x}) + v^{(O)}(\mathbf{x}), \quad (34)$$

which, in the case of the GM-CPHD, is a union of the weights, means, and covariance matrices of the two intensities. The cardinality distribution of the full multi-target state is the convolution of the probability mass functions, i.e.,

$$\rho(n) = \sum_{k=0}^{\infty} \rho^{(I)}(k) \rho^{(O)}(n - k), \quad n = 0, \dots, n_{\max}. \quad (35)$$

Upon combining the two populations, the GM-CPHD time update is performed to the next time of interest.

3.6. INTEGRATED SPACE OBJECT TRACKER

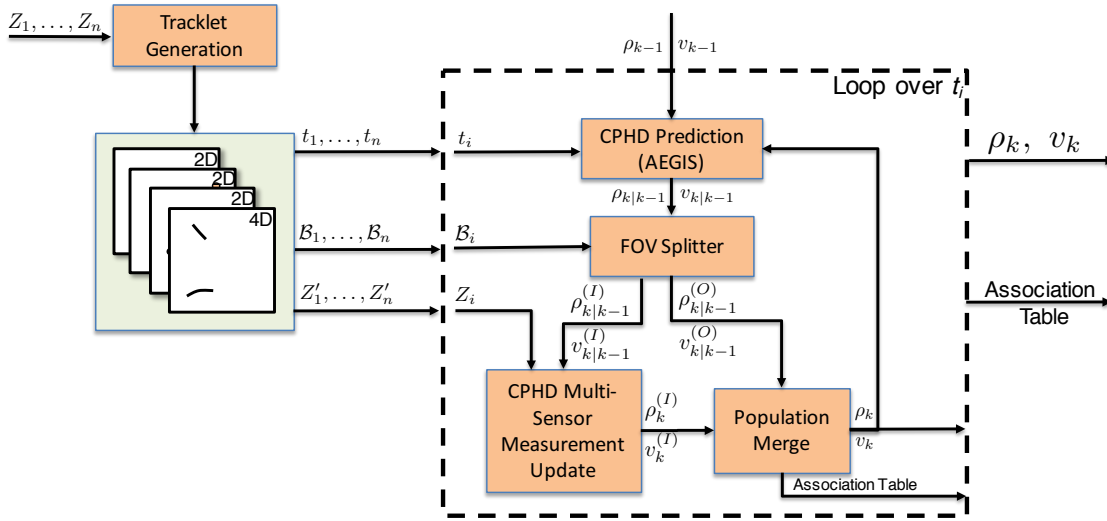


Fig. 3. Block diagram for the space-object tracker

Fig. 3 illustrates the combination of these tools to produce the proposed space-object tracker. The process begins with observation RFSs Z_i to produce the 4D observations (denoted as tracklets in the image) and the set of unused 2D observations. This produces the time t_i of each new measurement set Z'_1 with FOV boundaries \mathcal{B}_i . The filter then loops over time to process observations at each time t_i . Although not explicitly stated, observation sets at the same time use the previously described multi-sensor measurement update. The output of the tracker includes the multi-target state parameterized by ρ_k and v_k along with a table of measurements used in the update of each GM component in v_k .

4. NUMERICAL RESULTS

This section describes the performance of the proposed space-object tracker for maintaining custody of a catalog of space objects while establishing custody of newly detected objects. The section begins with a description of the general test scenario, followed by presentation and discussion of results.

4.1. SCENARIO DEFINITION

The near-GEO catalog used for this simulation is representative of the public TLE catalog available online¹. We extract all TLEs from the public catalog for Feb. 27, 2013 with mean motions between 0.9 and 1.1 revolutions per sidereal day, which yields 1,111 samples. These limits on mean motion correspond to semimajor axis limits of [39376, 45100] km. We then form a 20-component Gaussian mixture PDF in the six-dimensional orbit element space describing the density of objects near GEO using the EM algorithm [48]. Finally, we sample the PDF to generate a 700-object catalog for use in these simulations. The selection of the simulated catalog is constrained to only include objects within the field of regard (FOR) for the three simulated optical sensors.

Table 2. Upper and Lower Limits for Initial Covariance Standard Deviations

Orbital Element	Lower Bound	Upper Bound
Semimajor Axis (km)	0.1	1.0
Eccentricity	10^{-5}	10^{-4}
Inclination ($^{\circ}$)	10^{-3}	10^{-2}
Right Ascension ($^{\circ}$)	10^{-3}	10^{-2}
Argument of Perigee ($^{\circ}$)	10^{-3}	10^{-2}
True Anomaly ($^{\circ}$)	10^{-3}	10^{-2}




Fig. 4. Statistical consistency (as Mahalanobis distance) and absolute position RSS errors for the synthetic space-object catalog at epoch.

The initial filter estimate of the space-object catalog uses the 700 generated orbital element states and a randomly generated covariance matrix. For each object, the standard deviation for uncertainty for each element is sampled from a uniform distribution with the boundaries given in Table 2. The mean state in orbital elements (from the randomly generated TLE) and the diagonal covariance are used with the unscented transform and converted to Cartesian position and velocity. Sigma points for the transformation are selected such that no invalid values are produced, e.g., a negative eccentricity. The true object state is produced by sampling this covariance matrix, which yields an initial space-object catalog statistically consistent with the true catalog. Fig. 4 provides the distribution of initial Root-Sum-Square (RSS) of the position state errors and the distribution of Mahalanobis distances at the initial epoch time. Both the true state and the multi-target state are then propagated forward for 24 hours before simulating measurements. For this reason, presented results begin at hour 24. This yields an approximate Root-Mean-Square (RMS) position error of 8 km when measurements are first available, which is illustrated in the results presented below. The Gaussian components are propagated using the AEGIS method for approximating a non-Gaussian posterior PDF. This GM representation of

¹ Available at <https://celestrak.com> or <https://www.space-track.org/>

the space-object catalog is provided to the filter as $v_0(\mathbf{x})$. The initial cardinality distribution is initialized such that $\rho(N) = 1$ where N is the number of initially known objects for a given case.

Table 3. Ground Station Coordinates [49]

Station	Latitude (deg)	Longitude (deg)	Height (m)
Socorro, NM	33.82	-106.66	1510.2
Maui, HI	20.71	-156.25	3058.6
Diego Garcia	-7.41	72.45	-61.2

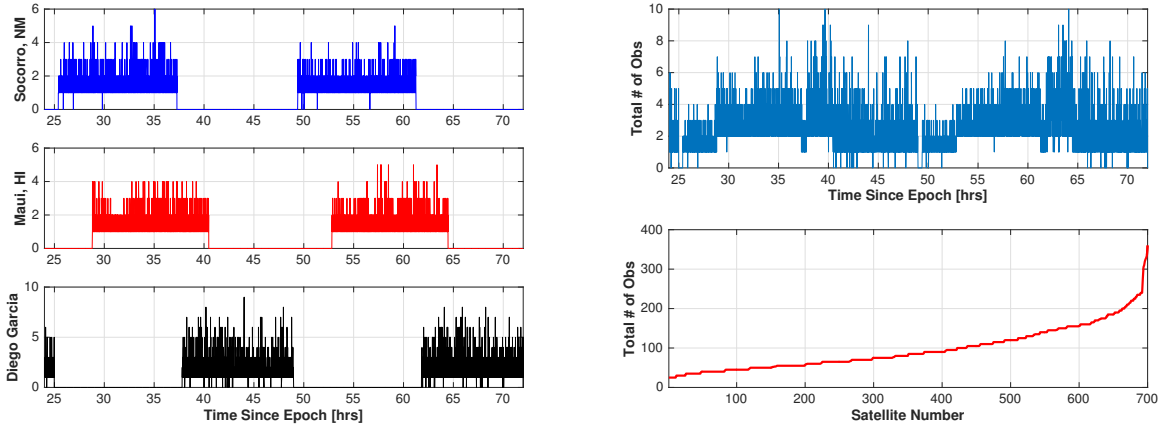


Fig. 5. Number of observations per measurement scan (left) and summary statistics over the scenario and for each satellite (right)

Table 3 describes the positions of the ground stations used for generating simulated angles-only measurements. Coordinate system reduction uses the IAU2000/2006 standard [50]. All objects within the $2^\circ \times 2^\circ$ field of view (FOV) may produce a measurement in a single scan. Observations are generated every six seconds with $p_D = 0.9$ and the sensor changing its pointing direction every 30 seconds. This yields five measurements from which the 4D measurements may be generated. As mentioned previously, the tools assume a minimum of three single-epoch measurements to form a 4D observation vector. The assumption of independence between measurement scans yields $p_D \approx 0.9914$ for the 4D measurements (the probability of at least three successes given five independent trials with a 0.9 probability of success). The remaining angle-only observations then have $p_D \approx 7.74 \times 10^{-3}$, which equals the probability that an object in the FOV is detected and no 4D measurement is generated. This disparity in the p_D values for the 4D and 2D measurements necessitates the use of the multi-sensor CPHD filter. Measurement noise for single-epoch angles measurements is 1 arcsec, which is reduced through the generation of the 4D measurements. No clutter measurements are included in the presented results and such tests are designated as future work. Fig. 5 illustrates the number of observations for each sensor at each time and the number of observations per space object over the two-day simulation.

As described previously, we use the PAR in combination with a PUB model to instantiate new tracks in the CPHD filter in a manner similar to [21]. The PAR itself is generated with random inputs defined uniformly over the instantaneous FOV, the semimajor axis limits [39000, 45100] (commensurate with, but slightly larger than the catalog limits), and a maximum orbit eccentricity of 0.1. The Gaussian mixture is formed using 20,000 samples in the range and range-rate direction. Since some components produced by the procedure have large axes, e.g., 1,000 km, these components are further split to have a maximum range uncertainty of 150 km to be more manageable for propagation. Each PAR requires 1000-5000 of these smaller GM components depending on the case. Most of these are quickly merged and/or pruned by the CPHD filter.

The CPHD filter requires additional parameters describing object survival and GM component management. This filter assumes $p_S = 0.99$. This assumption is conservative since space objects near GEO seldom depart that region of

space, but death can occur through, for example, a satellite break-up event. To prevent an exponential increase in the number of GM components in the filter, components within a Mahalanobis distance of 3 are merged and any with a weight less than 10^{-5} are pruned. See [15] for more details on the merging and pruning algorithm. Additionally, since some GM components are added to the filter and not observed again, we also prune components based on a maximum size. We quantify size using differential entropy

$$T = \frac{1}{2} \log |2\pi e \mathbf{P}| \quad (36)$$

where any components with $T > 10$ are removed from the multi-target state. Differential entropy is one method for quantifying the “diversity” of a PDF, and larger values reflect a larger covariance matrix in the case of a Gaussian distribution. Operationally, sensor tasking (e.g., see [46]) may be employed to limit growth in the size of the covariance ellipsoid and avoid the need for this requirement. Future work will explore other methods for putting a maximum on the size of a covariance before an object is considered “lost” or nonexistent.

One metric for describing the filter accuracy is the Optimal Sub-Pattern Assignment (OSPA) [51]

$$d_p^{(c)}(X, Y) \equiv \left(\frac{1}{n} \left[\min_{\pi \in \Pi_n} \sum_{i=1}^m d^{(c)}(x_i, y_{\pi(i)})^p + c^p(n - m) \right] \right)^{1/p} \quad (37)$$

where X and Y are RFSs, i.e., the estimated and true RFSs, Π is an assignment set relating an element $x \in X$ to $y \in Y$, $d^{(c)}$ is the minimum of c and the distance between the input arguments, and c and p are the cut-off and order parameters, respectively. The OSPA metric quantifies the minimum RMS distance between elements of the estimated and true RFSs and includes a penalty c for errors in cardinality. If the penalty term is ignored (the second term inside the brackets), then it describes the RMS error of the $\min[m, n]$ elements in the sets. The assignment Π is determined using minimization via auction. For the results present in this paper, only position is considered for OSPA with $p = 2$, $c = 1000$ km, and the penalty term ignored. The optimal value of c of the calculation of OSPA varies with time for this application, and we simplify the presentation of accuracy by selecting a large value and ignoring the penalty.

The following sections details results for the proposed space-object tracker for maintaining a catalog. The first section describes performance with all objects initially known to the filter. The next section then considers the case with a subset of the 700 objects unknown to the filter and custody is established through the track instantiation procedure.

4.2. KNOWN OBJECT CATALOG MAINTENANCE

This section describes the ability of the filter to maintain custody of the 700 objects in the catalog when all are initially known. The initial cardinality estimate for the case is $\rho_0(700) = 1$ and $\rho(n) = 0$ for $n \neq 700$. The AEGIS-propagated GM for the 700 object catalog yields $v_0(x)$. While theoretically not required since all objects are known to the filter, the birth model is not disabled for this test. The space-object tracker requires ~ 110 minutes using six cores to complete the 48-hour scenario on a Linux desktop computer with a dual 2.6 GHz Intel Xeon E5-2690 v4 running Red Hat Enterprise release 7.2. As outlined previously, the majority of the software is written in MATLAB with some routines in C/C++ using the gcc v.4.8.5 compiler. The propagator is written in Fortran and compiled using the gfortran version 4.8.5 compiler.

Fig. 6 describes the accuracy of the estimated number of objects, i.e., cardinality, and a description of the number of GM components required for the filter estimate. While not exhibited in the figure, the probability of fewer objects, e.g., $\rho(699)$, fluctuates slightly due to missed detections but the filter compensates with follow-up observations. The cardinality estimate remains constant at 700 objects and requires between 700 and 720 GM components. For this test, increases in the number of GM components result from the possible combination of hypotheses when using the multi-sensor CPHD approximation and not a hypothesized new object.

Fig. 7 provides the filter position OSPA accuracy. The filter converges on a solution accurate to approximately 200 m, which matches the observation noise of 1 arcsec for near-GEO objects. While not illustrated here, the mean number of measurements per object determines the convergence rate when not employing sensor tasking [46]. The reduction in OSPA accuracy shortly before 60 hours after the epoch results from an observation of a single object not recently observed at that point in the simulation.

4.3. SIMULTANEOUS KNOWN AND NEWBORN OBJECT TRACKING

The test case presented in this section duplicates that of the previous one, but with 15 satellites selected as initially unknown to the space-object tracker. Newborn objects are selected by the simulation framework such that: (i) they

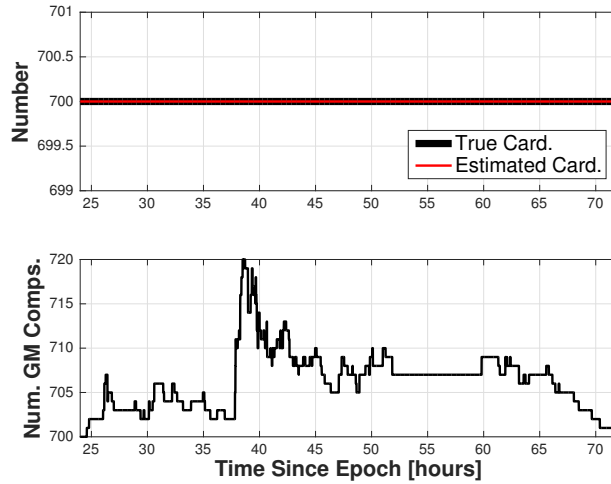


Fig. 6. Estimated cardinality and number of GM components for the case with known objects

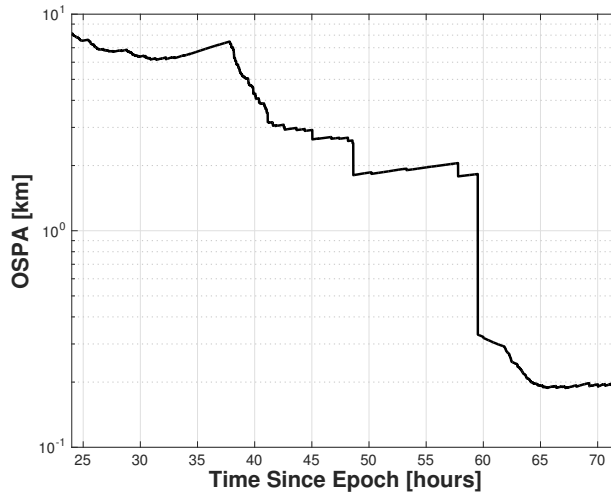


Fig. 7. Position OSPA accuracy for the case with known objects

are initially detected in the first ten hours of the simulation, and (ii) there are at least 100 angles-only observations over the two-day timespan to allow for possible solution convergence. Assuming no missed detections and perfect 4D observation generation, this yields a minimum of 20 angle/angle-rate observations over the 48-hour observation campaign. Future work will characterize tracker performance with fewer observations of newborn targets. The initial CPHD state for this case assumes $\rho(650) = 1$ and includes the AEGIS-propagated components of the GM for the known objects. When initially observed such that a 4D observation is generated, the filter determines the weight of the birth components and, when deemed necessary, uses the PAR to instantiate an initial intensity for the newborn objects. For this case, the space-object tracker requires approximately 121 minutes using six cores on the same Linux desktop computer.

Fig. 8 presents the performance of the space-object tracker based on the cardinality and estimated state accuracy. During the first 10 hours as newborn objects appear, the estimated cardinality increases and approximately equals the true number of objects. In some cases, a delay results from an initially low CPHD weight $w_{B,k}(z)$ for a possible new object, which requires a follow-up observation to confirm the object's existence. The filter does include some clutter tracks arising from duplicate instances of the PAR for the same object. In some instances, follow-up observations of the initially unknown object occur before there is a sufficiently high confidence in its existence to prevent another instance of the birth model. Over time, as illustrated in the plot, the filter prunes all but one of these duplicate objects.

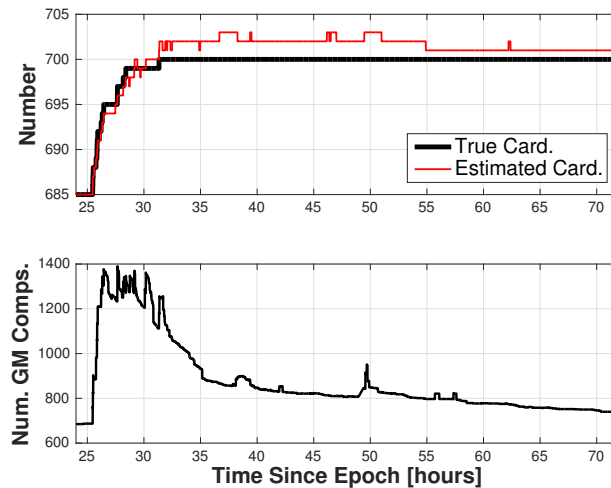


Fig. 8. Estimated cardinality and number of GM components for the case with newborn objects

The filter requires up to 1400 GM components, which the tracker prunes with follow-up observations of the newborn objects.

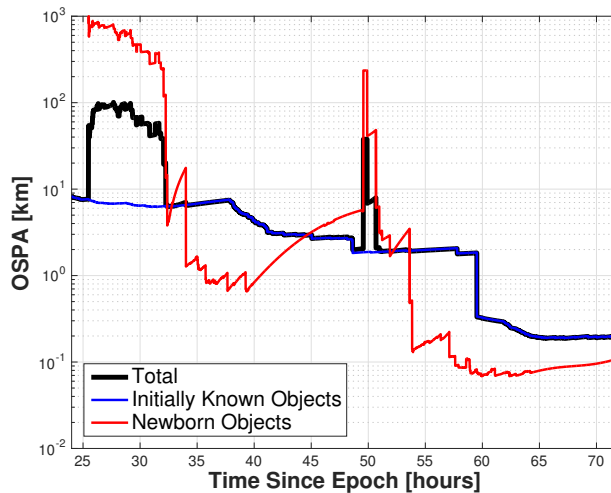


Fig. 9. Position OSPA Accuracy for the Case with Newborn Objects

Fig. 9 illustrates the accuracy of the filter-estimated state via position OSPA without the cardinality penalty. The error is further broken down into components from the initially known and newborn objects. Initially known object accuracy remains similar to the previous case in which only known objects were simulated. The OSPA accuracy for newborn objects determines the total accuracy, which is expected. The increase in error for the newborn objects at the 50-hour mark results from an instance of the birth model. Fig. 8 demonstrates a similar increase in the cardinality and number of GM components at that time. This instance of the birth model introduces an error in the estimated state, which is resolved with future observations. Upon viewing the RMS position accuracy, the filter has converged on a mean state accuracy of approximately 100-200 m over the two-night scenario. The error in the initially known objects is larger at the final time due to a nontrivial number of targets not recently observed.

5. CONCLUSIONS

The proposed space-object tracker, which is based on the GM-CPHD filter, provides a tractable means for estimating the number of objects in a given surveillance region and the single-target states. Tracking of SOs using the CPHD filter requires several tools customized for space-object tracking, (i) which reduces the computation load and compensates for limited information content when using angles-only measurements, (ii) a restricted FOV when compared to the region of interest, and (iii) the instantiation of new objects in the multi-target state. The space-object tracker leverages newly implemented orbit state and uncertainty propagation tools developed for the system that reduce computation time. Splitting of the Gaussian components near the edge of the FOV adjusts the CPHD intensity function used in the measurement update, which reduces instances of lost targets due to a believed missed detection. The space-object tracker requires a birth model to add a possible new target to the multi-target state. An object birth model based on the PAR in the CPHD requires the use of angles and angle-rates, which is generated from a time series of single-epoch observations using a GM-PHD filter in the measurement space. Combining the reduced four-dimensional observations with additional angles-only measurements requires the use of the approximate multi-sensor CPHD measurement update to prevent a filter bias. Combining these tools yields a space-object tracker capable of performing joint data association and state estimation. This includes the simultaneous maintenance of a prior space-object catalog and establishing custody of newly discovered objects with estimation accuracies at the final time comparable to the measurement errors. These tools run faster than realtime using a combination of MATLAB, C/C++, and Fortran software.

Future work includes expanded testing of the software and processing of real data. This includes testing the tools to establish custody of newborn objects after a satellite breakup event, such as a collision. Additionally, robustness will be improved by combining GM component track labels with an auction algorithm, which prevents potential swapping of GM component labels between closely spaced targets. Processing of real data will impact the measurement modeling in the filter, which will affect runtime and accuracy.

ACKNOWLEDGEMENTS

This work was funded by U.S. Air Force Small Business Innovation Research (SBIR) contract number FA9453-14-C-0295, awarded to Emergent Space Technologies, Incorporated. The authors wish to thank Ryan Weisman of the Air Force Research Laboratories, and Elfego Pinon, Sun Hur-Diaz, and Jessica Anderson of Emergent Space Technologies for their conversations and contributions to this work. The authors additionally acknowledge Ben Bradley for providing the software used to parse and propagate the TLEs used to help generate the near-GEO catalog. Additionally, Steven Gehly wrote the software for computing the OSPA accuracy and minimization via the auction algorithm.

6. REFERENCES

1. Committee for the Assessment of the U.S. Air Force's Astrodynamics Standards; Aeronautics and Space Engineering Board; Division on Engineering and Physical Sciences; National Research Council. *Continuing Kepler's Quest: Assessing Air Force Space Command's Astrodynamics Standards*. The National Academies Press, 2012.
2. Ronald P. S. Mahler. *Statistical Multisource-Multitarget Information Fusion*. Artech House, Boston, Massachusetts, 2007.
3. Navraj Singh, Joshua T. Horwood, Jeffrey M. Aristoff, Aubrey B. Poore, Carolyn Sheaff, and Moriba K. Jah. Multiple hypothesis tracking (MHT) for space surveillance: Results and simulation studies. In *Proceedings of the 2013 Advanced Maui Optical and Space Surveillance Technologies Conference*, Wailea, Maui, Hawaii, September 2013.
4. Jason Stauch, Moriba Jah, Jason Baldwin, Thomas Kelecyc, and Keric A. Hill. Mutual application of joint probabilistic data association, filtering, and smoothing techniques for robust multiple space object tracking (invited). In *AIAA/AAS Astrodynamics Specialist Conference*, San Diego, California, August 4-7 2014.
5. Islam I. Hussein, Kyle J. DeMars, Carolin Früh, Moriba K. Jah, and Richard S. Erwin. An AEGIS-FISST algorithm for multiple object tracking in space situational awareness. In *AIAA/AAS Astrodynamics Specialist Conference*, Minneapolis, Minnesota, August 13-16 2012.

6. 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*, AIAA 2014-4311, San Diego, California, August 4-7 2014.
7. Brandon A. Jones and Ba-Ngu Vo. A labeled multi-Bernoulli filter for space object tracking. In *2015 AAS/AIAA Space Flight Mechanics Meeting*, AAS 15-413, Williamsburg, VA, January 11-15 2015.
8. Kyle J. DeMars, Islam I. Hussein, Carolin Frueh, Moriba Jah, and R. Scott Erwin. Multiple-object space surveillance tracking using finite-set statistics. *Journal of Guidance, Control, and Dynamics*, 38(9):1741–1756, 2015.
9. Ronald Mahler. A theory of PHD filters of higher order in target number. In *Proceedings of the SPIE, Signal Processing, Sensor Fusion, and Target Recognition XV*, volume 6235, pages 62350K–62350K–12, 2006.
10. 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, July 2007.
11. Micheal Beard, Ba-Tuong Vo., Ba-Ngu Vo, and Sanjeev Arulampalam. A partially uniform target birth model for Gaussian mixture PHD/CPHD filtering. *Aerospace and Electronic Systems, IEEE Transactions on*, 49(4):2835–2844, October 2013.
12. Islam I. Hussein, Christopher W. T. Roscoe, Paul W. Schumacher, Jr., and Matthew P. Wilkins. Probabilistic admissible region for short-arc angles-only observations. In *Proceedings of the 2014 Advanced Maui Optical and Space Surveillance Technologies Conference*, Wailea, Maui, Hawaii, September 9-12 2014.
13. D. Franken, M. Schmidt, and M. Ulmke. “Spooky action at a distance” in the cardinalized probability hypothesis density filter. *IEEE Transactions on Aerospace and Electronic Systems*, 45(4):1657–1664, 2009.
14. Ronald P. S. Mahler. Multitarget bayes filtering via first-order multitarget moments. *Aerospace and Electronic Systems, IEEE Transactions on*, 39(4):1152–1178, Oct 2003.
15. Ba-Ngu Vo and Wing-Kin Ma. The Gaussian mixture probability hypothesis density filter. *IEEE Transactions on Signal Processing*, 54(11):4091–4104, November 2006.
16. W. Faber, S. Chakravorty, and Islam I. Hussein. A randomized sampling based approach to multi-object tracking. In *Information Fusion (FUSION), 2015 18th International Conference on*, pages 1307–1314, Washington, DC, July 6-9 2015.
17. Ba-Tuong Vo and Ba-Ngu Vo. Labeled random finite sets and multi-object conjugate priors. *IEEE Transactions on Signal Processing*, 61(13):3460–3475, July 2013.
18. Ba-Ngu Vo, Ba-Tuong Vo, and Dinh Phung. Labeled random finite sets and the bayes multi-target tracking filter. *IEEE Transactions on Signal Processing*, 62(24):6554–6567, December 2014.
19. Ba-Ngu Vo, Ba-Tuong Vo, and Hung Gia Hoang. An efficient implementation of the generalized labeled multi-Bernoulli filter. *arXiv preprint, arXiv: 1606.08350v1*, 2016.
20. Stephan Reuter, Ba-Tuong Vo, Ba-Ngu Vo, and Klaus Dietmayer. The labeled multi-Bernoulli filter. *IEEE Transactions on Signal Processing*, 62(12):3246–3260, June 2014.
21. Brandon A. Jones. Modeling birth in a space-object CPHD filter using the probabilistic admissible region. In *Information Fusion (FUSION), 2015 19th International Conference on*, Heidelberg, Germany, July 5-8 2016.
22. Rudolph Van der Merwe and Eric A. Wan. The square-root unscented Kalman filter for state and parameter estimation. In *2001 IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 6, pages 3461–3464, Salt Lake City, Utah, May 7-11 2001.
23. 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.
24. William E. Boyce and Richard C. DiPrima. *Elementary Differential Equations and Boundary Value Problems*. John Wiley & Sons, Inc., Hoboken, NJ, eighth edition, 2005. pp. 638-639.

25. P J Prince and J. R. Dormand. High order embedded runge-kutta formulae. *J Comput Appl Math*, 7(1):67–75, 1981.
26. Krishnan Radhakrishnan and A.C. Hindmarsh. Description and use of lsode, the livermore solver for ordinary differential equations. Technical Report UCRL-ID-113855, Lawrence Livermore National Laboratory, Livermore, CA, 1993.
27. M. Berry and L. Healy. Implementation of gauss-jackson integration for orbit propagation. *The Journal of the Astronautical Sciences*, 52(3):331–357, 2004.
28. M. E. Davis. *Numerical Methods and Modeling for Chemical Engineers*. John Wiley & Sons, Inc., New York, 1984. pp. 11-21.
29. Nitin Arora and Ryan P. Russell. A fast, accurate, and smooth planetary ephemeris retrieval system. *Celest Mech Dyn Astr*, 108(2):107–124, 2010.
30. Charles H. Acton. Ancillary data services of nasa’s navigation and ancillary information facility. *Planetary and Space Science*, 44(1):65–70, 1996.
31. The navigation and ancillary information facility: Spice data (spice kernels). <http://naif.jpl.nasa.gov/naif/data.html>. Accessed: August 2016.
32. William M. Kaula. *Theory of Satellite Geodesy*. Dover Publications, Inc., Mineola, NY, 2000 dover edition, 1966. pp. 1-11.
33. Nitin Arora, Vivek Vittaldev, and Ryan P. Russell. Parallel computation of trajectories using graphics processing units and interpolated gravity models. *J Guid Control Dynam*, 38(8):1345–1355, 2015.
34. Brandon A. Jones, George H. Born, and Gregory Beylkin. Comparisons of the cubed-sphere gravity model with the spherical harmonics. *J Guid Control Dynam*, 33(2):415–425, 2010.
35. Oliver Montenbruck and Eberhard Gill. *Satellite Orbits: Models, Methods, and Applications*. Springer-Verlag, Berlin, corrected 2nd printing edition, 2001. pp. 77-82, 248.
36. Kyle J. DeMars and Moriba K. Jah. Probabilistic initial orbit determination using Gaussian mixture models. *Journal of Guidance, Control, and Dynamics*, 36(5):1324–1335, September - October 2013.
37. Kohei Fujimoto, Daniel J. Scheeres, Johannes Herzog, and Thomas Schildknecht. Association of optical tracklets from a geosynchronous belt survey via the direct Bayesian admissible region approach. *Advances in Space Research*, 53(2):295–308, 2014.
38. G. Tommei, A. Milani, and A. Rossi. Orbit determination of space debris: admissible regions. *Celestial Mechanics and Dynamical Astronomy*, 97(4):289–304, 2007.
39. Jared M. Maruskin, Daniel J. Scheeres, and Kyle T. Alfriend. Correlation of optical observations of objects in earth orbit. *Journal of Guidance, Control, and Dynamics*, 32(1):194–209, January-February 2009.
40. Carolin Früh and Thomas Schildknecht. Object image linking of earth orbiting objects in the presence of cosmoics. *Advances in Space Research*, 49(3):594–602, 2012.
41. Ba-Ngu Vo, Ba-Tuong Vo, and Ronald P. S. Mahler. Closed-form solutions to forward-backward smoothing. *IEEE Transactions on Signal Processing*, 60(1):2–17, January 2012.
42. Ronald Mahler. Approximate multisensor CPHD and PHD filters. In *Information Fusion (FUSION), 2010 13th International Conference on*, pages 1–8, July 2010.
43. Sharad Nagappa and Daniel E. Clark. On the ordering of the sensors in the iterated-corrector probability hypothesis density (phd) filter. In *Proceedings of the SPIE, Signal Processing, Sensor Fusion, and Target Recognition XX*, volume 8050, 2011.

44. Ronald P. S. Mahler. *Advances in Statistical Multisource-Multitarget Information Fusion*. Artech House, Boston, Massachusetts, 2014.
45. Sharad Nagappa, Daniel E. Clark, and Ronald Mahler. Incorporating track uncertainty into the ospa metric. In *Information Fusion (FUSION), 2011 Proceedings of the 14th International Conference on*, pages 1–8, July 2011.
46. Steven Gehly, Brandon A. Jones, and Penina Axelrad. Sensor allocation for tracking geosynchronous space objects. *Journal of Guidance, Control, and Dynamics*, accepted for publication, 2016.
47. Y. Petetin, Daniel Clark, Branko Ristic, and D. Maltese. A tracker based on a cphd filter approach for infrared applications. In *Signal Processing, Sensor Fusion, and Target Recognition XX, Proceedings of the SPIE*, volume 8050, pages 80500N–80500N–12, 2011.
48. Mario A. T. Figueiredo and A. K. Jain. Unsupervised learning of finite mixture models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(3):381–396, March 2002.
49. David A. Vallado and Wayne D. McClain. *Fundamentals of Astrodynamics and Applications*. Microcosm Press and Springer, Hawthorne, CA and New York, NY, third edition, 2007.
50. Gérard Petit and Brian Luzum. Iers conventions (2010). IERS Technical Note 36, International Earth Rotation and Reference Systems Service (IERS), Frankfurt am Main, Germany, 2010.
51. Dominic Schuhmacher, Ba-Tuong Vo, and Ba-Ngu Vo. A consistent metric for performance evaluation of multi-object filters. *IEEE Transactions on Signal Processing*, 56(8):3447–3457, August 2008.