

Relative Multiple Space Object Tracking Using Intensity Filters

Keith A. LeGrand
Dept. of Mechanical and
Aerospace Engineering

Missouri University of Science and Technology
Rolla, Missouri 65409-0050
Email: kal7cd@mst.edu

Kyle J. DeMars
Dept. of Mechanical and
Aerospace Engineering

Missouri University of Science and Technology
Rolla, Missouri 65409-0050
Email: demarsk@mst.edu

Abstract—Multitarget intensity filters, such as the probability hypothesis density (PHD) filter and cardinalized probability hypothesis density (CPHD) filter have been recently proposed as a means to track multiple space objects from both ground-based and space-based platforms. In many applications, the CPHD is chosen over the PHD filter, as it has been claimed to offer significant improvements in the accuracy of both its cardinality estimates and state estimates. To that end, in this study, Gaussian mixture implementations of both the PHD and CPHD filters are developed to track the relative states of nearby space objects with respect to an inspector spacecraft using angles-only measurements. The performance of each solution is evaluated over several metrics, including cardinality error, optimal sub-pattern assignment distance, and execution speed.

I. INTRODUCTION

All modern approaches to space object tracking can be categorized as either space-based or ground-based. Ground-based systems utilize light arc data from optical telescopes and/or RADAR returns and benefit from the availability of high-performance computational resources. Space-based systems typically consist of one or more inspector satellites equipped with visual or infrared cameras. Space-based systems, in theory, surpass the capabilities of terrestrially-bound systems in that their imagery is not distorted by the atmosphere, and close proximities can be realized such that imagery can be resolved. Of course, the intrinsic size, weight, and power (SWAP) limitations of state-of-the-art satellite systems make the realization of coupling high-resolution imagery with high-performance computing infeasible.

Single-target space-based tracking is well-studied, with a recent research focus on non-cooperative relative navigation, that is, navigation in close proximity to an object that is not actively sharing information with the inspector [1], [2], [3], [4]. One promising approach to such proximity operations is through the use of a dual-camera system, or stereoscopic imager, as recently demonstrated with a single target by the Visual Estimation for Relative Tracking and Inspection of Generic Objects (VERTIGO) program [2]. The VERTIGO program, which was conducted in the Synchronized Position Hold, Engage, Reorient, Experimental Satellites (SPHERES) laboratory onboard the International Space Station, demon-

strated the first fully autonomous robotics vision-based navigation strategy with a non-cooperative spacecraft [2].

The proliferation of space objects in Earth orbit and, in particular, the subset of objects belonging to densely populated constellations, furthers the demand for the development and deployment of more sophisticated tracking algorithms. Examples of these dense constellations include debris clouds from satellite breakups or collisions [5], as well as controlled satellite swarms, which have been recently proposed as a means to provide broadband Internet across the globe, among countless other applications. Many of the approaches to the ground-based multitarget tracking problem involve some form of hypothesis formulation and explicit data association [6]. These techniques, such as multiple hypothesis tracking (MHT), suffer from time-exponential computational complexity increases due to the required data association, and consequently rely on heuristic reduction methods to achieve tractability.

More recent methods, based on moment approximations of the multitarget Bayes filter, have shown promise as more computationally tractable and statistically-consistent approaches to ground-based multitarget tracking [7], [8], as they mitigate the need for explicit data association. Two such methods, referred to as the probability hypothesis density (PHD) and cardinalized probability hypothesis density (CPHD) filters, approximate the full multitarget Bayes filter by propagating the posterior intensity function, which is the first-order moment of the full multitarget probability density function (pdf). The CPHD filter additionally propagates the posterior target cardinality probability mass function (pmf) [9]. Although the inclusion of cardinality statistics increases the complexity of the CPHD filter, it has been shown to dramatically improve cardinality estimates *and* state estimates [10].

In the current work, the PHD and CPHD filters are employed to serve as a means to track multiple targets in close proximity of an inspector satellite. The ability to autonomously track multiple targets in neighboring orbits using angles from video data could enable a new class of proximity operations in which cooperation between vehicles is not possible, or simply serve as a means to monitor the surrounding area in an effort to mitigate collision risks.

For successful tracking of multiple targets, systems must

be able to autonomously and statistically introduce new targets into the tracking solution using measurement data, be resilient to irregular, sparse, and spurious data, and be implementable on computationally-limited flight computers. Herein, a measurement-driven birth model is discussed which makes use of modern initial relative orbit determination methods. To account for data irregularity, target detection and clutter models are developed. The last issue of computational tractability is explored through a comparison of two modern multitarget tracking schemes, in which execution time and solution complexity are considered.

Similar to VERTIGO, this paper considers an inspector satellite system with a stereo imaging sensor that provides angular measurements of multiple nearby targets from two different optical viewpoints. These stereo measurements, when available, are processed in an intensity filter to recursively estimate the number of nearby targets, as well as the targets' relative positions and velocities with respect to the inspector satellite. The primary purpose of this study is to provide an investigation into the modeling considerations and challenges unique to angles-only space-based relative multitarget tracking. Special consideration is given to the birth model, which is responsible for incorporating previously untracked targets into the multitarget intensity when they appear. The performance of the presented framework is evaluated in the tracking of a nearby debris cloud using synthetic data.

II. RELATIVE DYNAMICS

When the orbit of an inspector spacecraft is nearly circular, the relative motion of nearby targets with respect to the inspector can be accurately modeled using a set of linearized equations of motion known as the Clohessy-Wiltshire equations. Rather than tracking the inertial states of nearby targets, tracking the relative motion of nearby targets alleviates the need for highly accurate knowledge of the inspector spacecraft's inertial state, which is a well-known principle that is exploited in most modern proximity and docking operations.

In addition to the assumption of a circular inspector orbit, some non-restrictive assumptions are made for convenience and notational simplicity. First, only planar relative motion is considered; that is, target motion in the out-of-plane direction is assumed to be zero. The target dynamics are modeled in the rotating Hill frame (depicted in Fig. 1), which for circular orbits rotates at a rate equal to the reference orbit's mean motion. Second, for simplicity, the origin of the Hill frame is taken to be the center of mass of the inspector spacecraft. With this, a target's position and velocity relative to the inspector spacecraft evolve over time as

$$\mathbf{x}_k = \mathbf{F}_{k-1} \mathbf{x}_{k-1}, \quad (1)$$

where the state vector \mathbf{x}_k consists of the relative position and velocity in the form $\mathbf{x}_k = [\mathbf{r}_k^T \quad \mathbf{v}_k^T]^T$, and

$$\mathbf{F}_{k-1} = \begin{bmatrix} \Phi_{rr} & \Phi_{rv} \\ \Phi_{vr} & \Phi_{vv} \end{bmatrix},$$

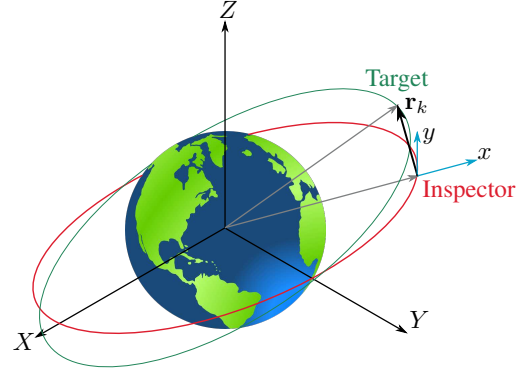


Fig. 1: Inspector/Target orbit geometry.

where

$$\begin{aligned} \Phi_{rr} &= \begin{bmatrix} 4 - 3 \cos \psi & 0 \\ 6(\sin \psi - \psi) & 1 \end{bmatrix}, \\ \Phi_{rv} &= \begin{bmatrix} \frac{1}{n_c} \sin \psi & \frac{2}{n_c} (1 - \cos \psi) \\ \frac{2}{n_c} (\cos \psi - 1) & \frac{4}{n_c} \sin \psi - \frac{3}{n_c} \psi \end{bmatrix}, \\ \Phi_{vr} &= \begin{bmatrix} 3n_c \sin \psi & 0 \\ 6n_c (\cos \psi - 1) & 0 \end{bmatrix}, \\ \Phi_{vv} &= \begin{bmatrix} \cos \psi & 2 \sin \psi \\ -2 \sin \psi & -3 + 4 \cos \psi \end{bmatrix}, \end{aligned}$$

and $\psi = n_c(t_k - t_{k-1})$, where n_c is the mean motion of the inspector satellite. Note that the Clohessy-Wiltshire (CW) model is only one of many available relative motion models. Other options include those presented in [11] and [12]. The CW model is chosen here because of its linear nature, which allows for simple target propagation and a straightforward solution to the relative Lambert problem.

The relative Lambert problem, whose namesake is borrowed from the classic Lambert problem of two-body orbital mechanics, consists of finding the a target's relative velocity at t_{k-1} that satisfies a known set of relative positions at t_{k-1} and t_k . Using the CW model, if \mathbf{r}_{k-1} , \mathbf{r}_k , and $(t_k - t_{k-1})$ are known, it follows that

$$\mathbf{v}_{k-1} = \Phi_{rv}^{-1}(\mathbf{r}_k - \Phi_{rr} \mathbf{r}_{k-1}). \quad (2)$$

The relative Lambert solution of (2) is used extensively in the creation of birth intensities for new targets, as discussed in Section V. Specifically, the relative Lambert problem is used to form a pdf in both relative position *and* velocity when only two positional pdfs and their temporal separations are known.

III. INTENSITY FILTERS

At time t_k , there are $N(k)$ targets with states $\mathbf{x}_{k,1}, \mathbf{x}_{k,2}, \dots, \mathbf{x}_{k,N(k)}$. The multitarget state $\mathbf{X}_k \in \mathcal{X}$ is composed of all of the target states, in no specified order, such that

$$\mathbf{X}_k = \{\mathbf{x}_{k,1}, \mathbf{x}_{k,2}, \dots, \mathbf{x}_{k,N(k)}\}.$$

Propagation of the multitarget pdf, which is expressed in the multitarget space \mathcal{X} , is generally intractable [13]. Tractable

recursions can be achieved by propagating the first-moment approximation of the multitarget pdf, which is known as the intensity or PHD and exists in the single-target space. In some situations, such as when the number of targets is large, the PHD filter can produce poor estimates of target cardinality. This is due to the fact that the PHD filter estimates the cardinality using only one parameter (the mean), thus effectively approximating the cardinality with a Poisson distribution, in which the mean and variance are equal [10]. In such situations, the CPHD filter often offers improvements in performance over the PHD filter by jointly propagating the target intensity $v_{k|k}$ and target cardinality distribution $p_{k|k}$.

The PHD and CPHD are the same in that they propagate the posterior intensity function at time t_{k-1} to time t_k as [10], [14]

$$v_{k|k-1}(\mathbf{x}) = \int p_{S,k}(\boldsymbol{\xi}) f_{k|k-1}(\mathbf{x}|\boldsymbol{\xi}) v_{k-1}(\boldsymbol{\xi}) d\boldsymbol{\xi} + \gamma_k(\mathbf{x}),$$

where $f_{k|k-1}(\cdot|\boldsymbol{\xi})$ is the single-target transition density at time t_k conditioned on previous state $\boldsymbol{\xi}$, $p_{S,k}(\boldsymbol{\xi})$ is the state-dependent probability of target survival, and $\gamma_k(\mathbf{x})$ is the intensity of spontaneous births at time t_k . The difference between the PHD and the CPHD filter is that the CPHD filter additionally propagates the discrete cardinality distribution as [10]

$$p_{k|k-1}(n) = \sum_{j=0}^n p_{\Gamma,k}(n-j) \Pi_{k|k-1}[v_{k-1}, p_{k-1}](j),$$

where

$$\Pi_{k|k-1}[v, p](j) = \sum_{\ell=j}^{\infty} C_j^\ell \frac{\langle p_{S,k}, v \rangle^j \langle 1 - p_{S,k}, v \rangle^{\ell-j}}{\langle 1, v \rangle^\ell}, \quad (3)$$

with $p_{\Gamma,k}(\cdot)$ representing the cardinality distribution of target births and C_j^ℓ denoting the binomial coefficient, which is defined as

$$C_j^\ell = \frac{\ell!}{j!(\ell-j)!}.$$

The operator $\langle \cdot, \cdot \rangle$ denotes the inner product, which is defined between α and β as

$$\langle \alpha, \beta \rangle = \int \alpha(x) \beta(x) dx \quad \text{and} \quad \langle \alpha, \beta \rangle = \sum_{\ell=0}^{\infty} \alpha(\ell) \beta(\ell)$$

in the cases that α and β are scalar real-valued functions or real-valued sequences, respectively. The cardinality distribution is represented by a pmf over the variable n . The variable n represents a discrete cardinality hypothesis; that is, $p_{k|k-1}(n) = \Pr(N(k) = n)$. The cardinality pmf is propagated according to the target birth cardinality pmf $p_{\Gamma,k}$ and existing target survival probability $\Pi_{k|k-1}$. Note that because the multitarget state has no specific ordering, the ℓ choose j different ways in which targets can survive or die is accounted for by the binomial coefficient that appears in (3).

At time t_k , a new measurement set $\tilde{\mathbf{Z}}_k$ is made available. The measurement set is an unordered set which contains both

false alarms (or clutter) and target-originated measurements. Measurements used for new target births (as described in Section V) are processed separately, and the set of all non-birth target measurements is denoted by $\mathbf{Z}_k = \tilde{\mathbf{Z}}_k \setminus \{\mathbf{Z}_{\text{birth},k}\}$. When the measurement set \mathbf{Z}_k is received, the CPHD filter updates the cardinality distribution and intensity function according to [15]

$$p_k(n) = \frac{\Upsilon_k^0[v_{k|k-1}, \mathbf{Z}_k](n) p_{k|k-1}(n)}{\langle \Upsilon_k^0[v_{k|k-1}, \mathbf{Z}_k], p_{k|k-1} \rangle} \quad \text{and}$$

$$v_k(\mathbf{x}) = \frac{\langle \Upsilon_k^1[v_{k|k-1}, \mathbf{Z}_k], p_{k|k-1} \rangle}{\langle \Upsilon_k^0[v_{k|k-1}, \mathbf{Z}_k], p_{k|k-1} \rangle} [1 - p_{D,k}(\mathbf{x})] v_{k|k-1}(\mathbf{x})$$

$$+ \sum_{\mathbf{z} \in \mathbf{Z}_k} \frac{\langle \Upsilon_k^1[v_{k|k-1}, \mathbf{Z}_k \setminus \{\mathbf{z}\}], p_{k|k-1} \rangle}{\langle \Upsilon_k^0[v_{k|k-1}, \mathbf{Z}_k], p_{k|k-1} \rangle} \psi_{k,z}(\mathbf{x}) v_{k|k-1}(\mathbf{x}),$$

where

$$\Upsilon_k^u[v, \mathbf{Z}](n) = \sum_{j=0}^{\min(|\mathbf{Z}|, n)} (|\mathbf{Z}| - j)! p_{K,k}(|\mathbf{Z}| - j) P_{j+u}^n$$

$$\times \frac{\langle 1 - p_{D,k}, v \rangle^{n-(j+u)}}{\langle 1, v \rangle^n} e_j(\Xi(v, \mathbf{Z})), \quad (4)$$

$$\psi_{k,z}(\mathbf{x}) = \frac{\langle 1, \kappa_k \rangle}{\kappa_k(\mathbf{z})} g_k(\mathbf{z}|\mathbf{x}) p_{D,k}(\mathbf{x}),$$

$$\Xi(v, \mathbf{Z}) = \langle v, \psi_{k,z} \rangle : \mathbf{z} \in \mathbf{Z},$$

$g_k(\cdot|\mathbf{x})$ is the single-target measurement likelihood at time t_k given current state \mathbf{x} , $p_{D,k}(\mathbf{x})$ is the state-dependent probability of target detection at time t_k given current state \mathbf{x} , $\kappa_k(\cdot)$ is the intensity of clutter measurements at time t_k , $p_{K,k}(\cdot)$ is the cardinality distribution of clutter at time t_k , and $e_j(\cdot)$ is the elementary symmetric function of order j , which is defined as

$$e_j(\mathbf{Z}) = \sum_{\mathbf{S} \subseteq \mathbf{Z}, |\mathbf{S}|=j} \left(\prod_{\zeta \in \mathbf{S}} \zeta \right) \quad (5)$$

An efficient method for the calculation of (5) is described in [10]. The permutation coefficient P_j^n is defined as

$$P_j^n = \begin{cases} \frac{n!}{(n-j)!} & , n \geq j \\ 0 & , n < j \end{cases}.$$

The PHD filter update is much simpler. By approximating the target cardinality as Poisson-distributed, the PHD filter updates the intensity according to

$$v_k(\mathbf{x}) = [1 - p_{D,k}(\mathbf{x})] v_{k|k-1}(\mathbf{x})$$

$$+ \sum_{\mathbf{z} \in \mathbf{Z}_k} \frac{p_{D,k}(\mathbf{x}) g_k(\mathbf{z}|\mathbf{x}) v_{k|k-1}(\mathbf{x})}{\kappa_k(\mathbf{z}) + \int p_{D,k}(\boldsymbol{\xi}) g_k(\mathbf{z}|\boldsymbol{\xi}) v_{k|k-1}(\boldsymbol{\xi}) d\boldsymbol{\xi}}$$

The PHD/CPHD recursions admit both sequential Monte Carlo (SMC) [16]/[13] and Gaussian mixture (GM) [14]/[10] implementations. Because particle methods such as the SMC are, in general, extremely computationally expensive, they are ill-suited for real-time satellite navigation systems, in

which state vectors with 30+ states are common. Although the practicality of GM implementations in the context of the current state-of-the-art flight processor limitations is still under investigation, the GM approach is often the most feasible for computationally-limited systems. To that end, GM implementations of the PHD and CPHD filters are considered in this work.

In the case of linear Gaussian dynamics and measurements, the PHD and CPHD recursions admit a closed-form solution when the posterior intensity is modeled by a GM. The closed-form recursion for linear Gaussian models can be extended to accommodate nonlinear dynamics and measurements via linearization of the transition density and measurement likelihood functions in the same fashion as the single-target extended Kalman filter [10], [14]. Because only linear dynamics are considered in this work, local linearizations are only required in the measurement update.

IV. MEASUREMENT MODEL

Angular measurements are taken from two cameras which are attached to the inspector spacecraft. The use of two cameras separated at a known baseline (as opposed to a single-camera system) allows the range ambiguity associated with an angular measurement to be resolved, thus mitigating the need for prior knowledge of the target's geometry. The concurrent use of two distributed cameras to achieve depth perception is commonly known as stereoscopic vision and is an active topic of research across multiple disciplines [2].

A. Target-Originated Measurements

Robust image processing algorithms are available [17], [18] that identify unique features on an imaged target and match the features between a stereoscopic image pair, such as in Figure 2. Similar to the case of Figure 2, multiple features are identified and matched for a single target. Although the angles associated with these matched features will not, in general, correspond to the target's exact center of mass (CM), a simple averaging of the features' angles will result in a good approximation of the target's geometric center, which for geometrically small targets, will be close to the CM. Further improvements in accuracy can be made with prior knowledge of the target's geometry and mass distribution, as might be the case with cooperative swarms.




Fig. 2: Example of image processing output for a single target. The lines demonstrate matched features between the two images of a stereoscopic image pair.

One of the primary benefits of the feature matching algorithm is that, because the features are matched between

images, angles can be processed in pairs as opposed to individually. To that end, a single measurement is modeled as containing both angles from the stereo imager with additive Gaussian white noise as

$$\mathbf{z}_k = [\theta_{C1,k} \quad \theta_{C2,k}]^T + \boldsymbol{\nu}_k,$$

where θ_{C_n} is the azimuth angle measured from the y axis of the C_n camera frame, $E[\boldsymbol{\nu}_k] = 0$, and $E[\boldsymbol{\nu}_k \boldsymbol{\nu}_k^T] = \mathbf{R}_k = \text{diag}\{\sigma_{\theta_{C1,k}}^2, \sigma_{\theta_{C2,k}}^2\}$. For convenience, the camera frames are assumed to be aligned with the Hill frame. With this assumption, a stereoscopic measurement is related to a target's state by the measurement function

$$\mathbf{h}(\mathbf{x}_k) = \begin{bmatrix} \tan^{-1} \left(\frac{u_{C1,k,x}}{u_{C1,k,y}} \right) \\ \tan^{-1} \left(\frac{u_{C2,k,x}}{u_{C2,k,y}} \right) \end{bmatrix}, \quad \mathbf{u}_{C_n} = \frac{\mathbf{r}_k - \mathbf{d}_{k,C_n}}{\|\mathbf{r}_k - \mathbf{d}_{k,C_n}\|},$$

where \mathbf{r}_k is the relative position vector portion of the state \mathbf{x}_k and \mathbf{d}_{k,C_n} is the position of the C_n camera frame origin with respect to the Hill frame origin.

B. False Alarms

Occasionally, the image processing algorithms will return measurements that do not correspond to actual targets, as illustrated in Figure 3. Such false alarms can be the result of




Fig. 3: Example of false matches.

sensor malfunctions, glint events [19], or other space-lighting phenomena. These false alarms are accounted for by the clutter intensity $\kappa_k(\mathbf{z})$ and cardinality distribution $p_{K,k}$. The clutter intensity is assumed to be uniform over the valid stereoscopic measurement domain D_s such that

$$\kappa_k(\mathbf{z}) = \lambda V \mathcal{U}(\mathbf{z}), \quad \mathcal{U}(\mathbf{z}) = \begin{cases} 1/V & , \quad \mathbf{z} \in D_s \\ 0 & , \quad \text{otherwise} \end{cases},$$

where λ is the average number of clutter returns per unit "volume" and V is the sensor volume. Because false alarms exhibit the same properties as target-originated measurements, namely, that they correspond to lines of sight (LOSs) that intersect in the composite field of view (CFOV), the volume V must represent the space over which stereoscopic measurements are valid. Valid stereo measurements are those for which the angular measurements belong to both of the cameras' fields of view (FOVs) and the LOSs they describe intersect in the positive- y half of the plane (in "front" of the cameras). For aligned cameras, to guarantee that the LOSs intersect, the azimuth angle from the leftmost sensor must always be greater

than the azimuth angle from the rightmost sensor. With this, the sensor volume is obtained from the volume integral

$$V = \int_{D_s} 1dz, \quad D_s = \{\text{FOV}_{C1} \cap \text{FOV}_{C2} | \theta_{C1} > \theta_{C2}\}. \quad (6)$$

When the FOV angle is identical for both cameras; that is, when $\text{FOV}_{C1} = \text{FOV}_{C2} = \text{FOV}$, the result of (6) is simply $V = \frac{1}{2}\text{FOV}^2$. The frequency of false alarms is dependent on a plethora of parameters related to the sensor, the image processing system, lighting conditions, etc. The likelihood of receiving a given number of false alarms at time t_k is modeled by the clutter cardinality distribution $p_{K,k}(n)$. In this study, $p_{K,k}(n)$ is taken as Poisson, such that

$$p_{K,k}(n) = \frac{(\lambda V)^n}{n!} e^{-\lambda V}. \quad (7)$$

For real-world operations, $p_{K,k}(n)$ should be fitted to ground-testing data before spaceflight.

C. Probability of Detection

The probability of detection $p_{D,k}(\mathbf{x}_k)$ represents the probability that a target with state \mathbf{x}_k will be measured by the stereoscopic imager at t_k . As mentioned in the previous section, targets will, in general, have multiple features that are matched, such that when the target is visible, there is a strong likelihood of detection. To that end, detection is taken to be guaranteed if the target exists in the domain of valid stereoscopic measurements and within the maximum ranges of the sensors, such that

$$p_{D,k}(\mathbf{x}_k) = \begin{cases} 1, & \mathbf{h}(\mathbf{x}_k) \in D_s \text{ and} \\ & \|\mathbf{r}_k - \mathbf{d}_{k,Cn}\| \leq \rho_{\max,Cn} \\ 0, & \text{otherwise} \end{cases}, \quad (8)$$

where $\rho_{\max,Cn}$ is the maximum range of sensor C_n .

The probability of detection model of (8) is, in fact, a strong simplification of the complex metrological considerations involved in satellite vision-based navigation. Such considerations include, but are not limited to, the solar illumination geometry, sensor solar keep-out zones, and measurement occlusion caused by other targets. These complex modeling considerations are beyond the scope of this paper, as they are very sensor- and mission-dependent.

V. TARGET BIRTH

When new targets enter the inspector's stereoscopic FOV, their intensities must be added to the intensity of existing targets, and the cardinality pmf must be updated accordingly. It is assumed that measurements associated with new targets are identified. The assumption that the measurements of new targets are explicitly identified is certainly one of interest, but the development of such an identification mechanism is beyond the scope of the current work. Such mechanisms are not infeasible, however. For instance, one potential approach consists of applying a parallel process that analyzes image feature data to detect new targets and informs the birth process when a new target is detected. Additionally, this parallel

process, when applied to the data, would identify the angular measurements in closest likelihood agreement with the new detection to be used by the previously described birth process. With this, the birth intensity can be computed from a set of stereoscopic measurements using the initial relative orbit determination (IROD) framework developed in [4]. In this framework, two stereoscopic angle measurements are collected at times t_{k-2} and t_{k-1} . At each time step, the angle measurement and corresponding uncertainty from Camera 2 are used to bound possible relative range values along Camera 1's LOS. The relative range uncertainty over these bounded regions is taken to be uniform and is approximated by a GM, as illustrated in Figure 4. The resulting pdfs are expressed in




Fig. 4: GM representing uniform relative range uncertainty over $\rho_{\min} \leq \rho \leq \rho_{\max}$.

polar coordinates in the relative position space. Because the intent is to obtain a pdf that can be expressed in the single-target space, the set of positional polar coordinate pdfs are transformed into a full-state Cartesian pdf using a Lambert-linkage process.

In the Lambert-linkage, components from the relative position mixture at t_{k-2} are linked to components from the relative position mixture at time t_{k-1} by computing the relative velocities that satisfy a given pair of relative positions and the time elapsed between them. Statistics are collected during this process such that the linkage of every component pair combination between t_{k-2} and t_{k-1} produces a unique birth component in the single-target space with appropriate weight, mean, and covariance.

Suppose that at times t_{k-2} and t_{k-1} , the measurements $\mathbf{z}_{k-2,\ell} \in \bar{\mathbf{Z}}_{k-2}$, $\mathbf{z}_{k-1,j} \in \bar{\mathbf{Z}}_{k-1}$ are identified as measurements that originate from a new target; that is, $\mathbf{z}_{k-2,\ell} \in \mathbf{Z}_{\text{birth},k-2}$ and $\mathbf{z}_{k-1,j} \in \mathbf{Z}_{\text{birth},k-1}$. For notational simplicity, the ℓ and j subscripts are dropped. Each measurement contains angles from both cameras; that is

$$\mathbf{z}_{k-2} = [\theta_{C1,k-2} \quad \theta_{C2,k-2}]^T \quad \text{and} \\ \mathbf{z}_{k-1} = [\theta_{C1,k-1} \quad \theta_{C2,k-1}]^T.$$

At t_{k-2} and t_{k-1} , the LOS of Camera 2 is rotated by a predetermined deviation in all directions, and the intersections between the rotated lines and LOS of Camera 1 are used to determine minimum and maximum relative range bounds ρ_{\min} and ρ_{\max} . In this study, the rotational deviation is chosen as three times the angular standard deviation of Camera 2. The points ρ_{\min} and ρ_{\max} represent the minimum and maximum relative range bounds (with respect to the camera frame) within which the position of the birth target is hypothesized and thus serve as the bounds in the construction of the relative range pdf. Note that ρ_{\max} is not the same as the sensor range limitation $\rho_{\max,Cn}$. However, in the event that the bounding

process generates a relative range ρ_{\max} that is larger than the known sensor range limitation, ρ_{\max, C_n} is appropriately used in the place of ρ_{\max} as the upper bound on the relative range. Equal uncertainty in relative range is assumed over these bounds, such that the relative range pdf is

$$p_{\rho,k}(\rho) = \begin{cases} \frac{1}{\rho_{\max} - \rho_{\min}} & , \quad \rho_{\min} \leq \rho \leq \rho_{\max} \\ 0 & , \quad \text{otherwise} \end{cases} \quad (9)$$

The relative range pdf of (9) is modeled by a GM of L_k components; that is

$$p_{\rho,k}(\rho) = \sum_{\ell=1}^{L_k} w_{\rho,k}^{(\ell)} \mathcal{N}(\rho; m_{\rho,k}^{(\ell)}, \sigma_{\rho,k}^2). \quad (10)$$

A more detailed discussion on the construction of this model is provided in [4]. Note that while the angular measurement from Camera 2 and its associated uncertainty provide mixture bounds over the LOS of Camera 1, the direction of the LOS of Camera 1 is intrinsically noisy. Thus, Camera 1's measurement $\theta_{C1,k}$ is treated as the mean, and its uncertainty is accounted for in the positional pdf by augmenting the components' covariance with Camera 1's measurement noise covariance $\sigma_{\theta_{C1,k}}^2$. With this, the relative position pdf is formed over the polar coordinate position \mathbf{y} as

$$p_{y,k}(\mathbf{y}) = \sum_{\ell=1}^{L_k} w_{y,k}^{(\ell)} \mathcal{N}(\mathbf{y}; \mathbf{m}_{y,k}^{(\ell)}, \mathbf{P}_{y,k}^{(\ell)}), \quad (11)$$

where

$$w_{y,k}^{(\ell)} = w_{\rho,k}^{(\ell)}, \quad \mathbf{m}_{y,k}^{(\ell)} = \begin{bmatrix} m_{\rho,k}^{(\ell)} & \theta_{C1,k} \end{bmatrix}^T, \\ \text{and} \quad \mathbf{P}_{y,k}^{(\ell)} = \text{diag}\{\sigma_{\rho,k}^2, \sigma_{\theta_{C1,k}}^2\}.$$

The single-target pdf is fully described in the relative position space by (11), but it does not contain relative velocity information. However, if two of these relative position pdfs are formed at two unique times, the pdf in relative velocity can be formed through a Lambert-linkage process. Suppose that at time t_k , two measurement sets at t_{k-2} and t_{k-1} have been identified as belonging to a new target. GMs describing the single-target probability density using (11) are formed as

$$p_{y,k-2}(\mathbf{y}) = \sum_{\ell=1}^{L_{k-2}} w_{y,k-2}^{(\ell)} \mathcal{N}(\mathbf{y}; \mathbf{m}_{y,k-2}^{(\ell)}, \mathbf{P}_{y,k-2}^{(\ell)}) \\ \text{and} \quad p_{y,k-1}(\mathbf{y}) = \sum_{\ell=1}^{L_{k-1}} w_{y,k-1}^{(\ell)} \mathcal{N}(\mathbf{y}; \mathbf{m}_{y,k-1}^{(\ell)}, \mathbf{P}_{y,k-1}^{(\ell)}).$$

The Lambert-linkage consists of "linking" these two pdfs together to obtain a new pdf that includes relative velocity information. Every component pair combination must undergo a conversion from polar coordinates to Cartesian coordinates and a transformation from the camera frame to the Hill frame, after which a relative Lambert solution is found to provide the relative velocity that satisfies the two relative position components and their temporal separation. This linkage process is nonlinear, and thus to properly capture the statistics of this

process, an unscented transform [20] is employed. For brevity, this process is not described here, but a thorough discussion is found in [4]. The result of the Lambert-linkage process is an $L_{k-2} \cdot L_{k-1}$ component GM that approximates the single-target pdf in relative position and velocity and is expressed in Cartesian coordinates as

$$p_{x,k-2|k-1}(\mathbf{x}) = \sum_{\ell=1}^{L_{k-2}L_{k-1}} w_{x,k-2}^{(\ell)} \mathcal{N}(\mathbf{x}; \mathbf{m}_{x,k-2}^{(\ell)}, \mathbf{P}_{x,k-2}^{(\ell)}).$$

The birth intensity at time t_k is then obtained by simply propagating $p_{x,k-2|k-1}(\mathbf{x})$ to time t_k , such that

$$\gamma_k(\mathbf{x}) = \int p_{S,k}(\boldsymbol{\xi}) f_{k|k-2}(\mathbf{x}|\boldsymbol{\xi}) p_{x,k|k-2}(\mathbf{x}) d\boldsymbol{\xi}.$$

In the event that multiple births occur simultaneously, the birth intensity is the summation of the new single-target pdfs given by

$$\gamma_k(\mathbf{x}) = \sum_{\mathbf{Z}_{\text{birth},k-2,k-1}} \int p_{S,k}(\boldsymbol{\xi}) f_{k|k-2}(\mathbf{x}|\boldsymbol{\xi}) p_{x,k|k-2}(\mathbf{x}) d\boldsymbol{\xi}.$$

where $\mathbf{Z}_{\text{birth},k-2,k-1}$ denotes the set of temporally-paired (between t_{k-2} and t_{k-1}) birth target measurements. The cardinality of births is taken as Poisson, yielding

$$\Gamma_k(n) = \frac{|\mathbf{Z}_{\text{birth},k-2,k-1}|^n}{n!} e^{-|\mathbf{Z}_{\text{birth},k-2,k-1}|},$$

so that higher target number birth events are treated with more uncertainty than lower target number birth events. Because the





Fig. 5: Block diagram of space-based relative multitarget tracking framework.

birth intensity is formed from measurement data at t_{k-2} and t_{k-1} , the birth intensity is not introduced into the intensity model until another measurement is available, such that the birth intensity is not doubly conditioned on the same measurement data. A high-level block diagram of the complete tracking framework can be found in Figure 5.

Figure 6 illustrates an example of a new birth target with two existing targets. In Figure 6a, the distribution on the left corresponds to a birth intensity at t_k . Its elongated nature is due to the large uncertainty in relative range and the uncertainty growth from propagation from t_{k-2} to t_k . Because of this, some birth components, when added to the intensity of surviving targets, reside outside of the detectable region, which is represented by the darker gray intersection of the cameras' FOVs. These components, if left unhandled, are treated as misdetections and do not undergo a measurement



(a) Untruncated birth intensity. (b) Truncated birth intensity.

Fig. 6: Example of one birth with two existing targets.

update. In this specific example, these components will likely enter the detectable region at the following measurement update and thus undergo an update at t_{k+1} . In other cases, however, components may reside outside of detection region due to their relative range and, furthermore, may never reenter the region. In applications where *all* target tracks are maintained indefinitely (i.e. $p_{S,k} = 1.0$), the persistence of such components can complicate numerical aspects of the filter. Instead, if a measurement is known to be associated with the birth target at t_k , its existence implies target detectability, and thus, all components that fall outside the detectable region can be safely truncated, as illustrated in Figure 6b. The truncated distribution is then renormalized and treated as the birth intensity.

The necessity of this truncation is more apparent when measurements are processed at high frequency, and more specifically, when little time is elapsed between Lambert-linkage measurements. As discussed in [4], pdfs formed from small linkage times generally lead to a higher uncertainty in the target relative velocity as compared to pdfs formed from longer linkage times. In some cases, for exceptionally small linkage times, the large uncertainty in relative velocity will lead to propagated components which fall in the negative- y half plane (“behind” the camera).

Another approach to handling these “runaway” components is to implement a non-unity probability of survival. The effect of this is that components are downweighted at every step, which over time, significantly reduces the contribution of components that are undetectable. For many terrestrial applications, specifically applications when target motion is secular, values between 0.95 and 0.99 for $p_{S,k}$ are common [21], [22]. Conversely, for ground-based tracking of geostationary satellites, where a future measurement is almost certainly guaranteed due to the satellites’ intrinsic periodic motion, $p_{S,k} = 1$ is appropriate. Satellite relative motion is both secular and periodic, and as a result, target survival is not as straightforward. Potential approaches could involve computing the probability of re-detection based on states extracted from the intensity function. The states of targets with low probabilities of re-detection may be converted to inertial states for long-

term propagation while their corresponding components are downweighted. Currently, more sophisticated target survival schemes are being investigated and will be discussed in future work.

VI. RESULTS

The performance of the PHD and CPHD filters is investigated in the context of a simulated debris cloud tracking problem. The inspector satellite occupies a 400 km altitude circular Earth orbit. For simplicity, the satellite is modeled to rotate with the Hill frame, such that the inspector’s cameras always point along the positive y axis of the Hill frame, and perfect attitude knowledge is assumed.

Two cameras are fixed to the inspector spacecraft, and due to the chosen rotation of the inspector spacecraft, their locations within the Hill frame are constant and are described by

$$\mathbf{d}_{C1} = [-2 \ 0]^T \text{ [m]} \quad \text{and} \quad \mathbf{d}_{C2} = [2 \ 0]^T \text{ [m]}.$$

Measurements from the stereo imager are received every 60 seconds in the form of bearings angles, which are corrupted by zero-mean Gaussian noise with a standard deviation of 750 arcseconds, which corresponds to a five pixel deviation on a 1080×1080 pixel focal plane array with a 45 deg FOV. Targets are treated as undetectable if they fall outside of either camera’s FOV or are more than 150 m away from either camera. False alarms are generated in accordance with a Poisson cardinality distribution (7) with mean $\lambda V = 2$ returns per collection.

Targets are initialized in close proximity to the inspector spacecraft, and their true initial relative states are listed in Table I. The initial states are propagated using CW dynamics to produce the truth model, of which the positional data are shown in Figure 7. For convenience, the intensity function is initialized with its component means at the visible targets’ true states, with component covariances given by $\mathbf{P}_0^{(i)} = \text{diag}[2^2 \ 2^2 \ 0.01^2 \ 0.01^2]$ in m^2 and m^2/s^2 . The initial target cardinality pmf is initialized as uniform between zero and nineteen; that is, $\Pr(N(0) = n) = 0.05 \forall n \in \{0, \dots, 19\}$.

TABLE I: Initial relative states of targets.

Target	x [m]	y [m]	\dot{x} [m/s]	\dot{y} [m/s]
1	4.96026332	38.32210715	0.00441559	-0.01124242
2	5.39555258	18.63914052	-0.00715109	-0.00773215
3	-11.33700072	36.87592374	0.00532677	0.02230584
4	-8.86908651	23.98647305	-0.00436219	0.01892312
5	-14.29718807	27.12171413	-0.00070161	0.02893518
6	-3.07473223	14.79459501	-0.00459050	0.00995115
7	-0.16750985	18.95011169	-0.00591691	-0.00320291
8	-6.24943905	39.15861624	0.00750975	0.00747834
9	-16.15237908	18.14938250	-0.00111060	0.03229021

At each measurement update, the PHD and CPHD posterior intensities are reduced according to the same pruning/merging laws. Gaussian components are combined if their Mahalanobis distance falls below a predetermined merging threshold and components with weights less than a predetermined pruning threshold are removed from the mixture. Additionally, the

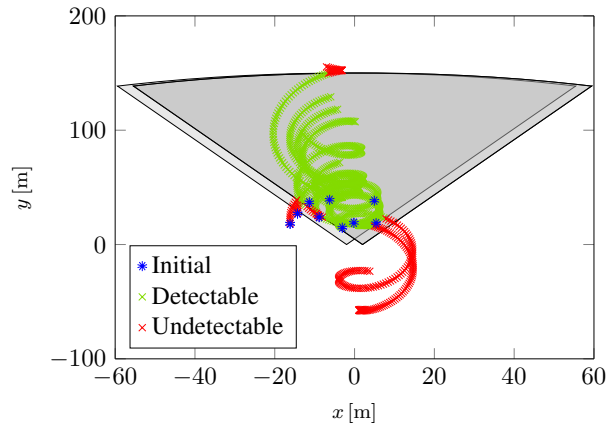


Fig. 7: Relative position histories. The gray shaded regions represent the camera FOVs.

cardinality means and standard deviations for the CPHD filter are computed as

$$\hat{N}(k) = \sum_{n=0}^{n_{\max}} np_k(n)$$

$$\text{stddev} = \sqrt{\sum_{n=0}^{n_{\max}} (n - \hat{N}(k))^2 p_k(n)}.$$

Alternatively, the cardinality estimate for both the PHD and CPHD solutions can be attained by utilizing the important relationship

$$\hat{N}(k) = \int v_k(\mathbf{x}) d\mathbf{x},$$

which can be simply computed by summing the intensity's Gaussian component weights. These cardinality estimates and the true target cardinality are directly compared in Fig. 8. As shown, the three targets that are not accounted for by




Fig. 8: PHD and CPHD cardinality estimates and true target cardinality.

the initial intensity function are introduced over the first several time steps via the birth process. The PHD filter overestimates the cardinality for two of these three births, and in general, exhibits much less accurate cardinality estimates

than its cardinalized counterpart. For a closer examination of the CPHD's cardinality estimate, the cardinality error and pmf standard deviation are provided in Fig. 9. Note that shortly after $t = 1.5$ hr, the cardinality error falls below 64-bit machine precision, thus making the log error undefined.




Fig. 9: CPHD cardinality statistics.

At each time step, the state estimates are extracted from the PHD and CPHD intensities. The “miss distance” between these extracted state estimates and the true target states is computed using the optimal subpattern assignment (OSPA) metric [23]. As shown in Fig. 10, the accuracy of the CPHD filter exceeds that of the PHD filter with the exception of only a few time steps, as is expected.




Fig. 10: First-order OSPA metric with cutoff $c = 1$.

It is often accepted that the CPHD filter requires additional execution time in its recursion in comparison to the PHD because of its cubic complexity in measurements. This generalization, however, is only valid under the assumption that the number of components between the PHD and CPHD intensity mixtures remains the same. In many cases, such as the presented example, the CPHD update and subsequent reduction step, in fact, produces a more tightly concentrated intensity function with far fewer components. To that end, the total recursion time, or time elapsed from the beginning of the prediction step to the end of the correction step, for the CPHD filter is often less than the PHD filter. This trend is exemplified in Fig. 11, which compares the non-dimensional execution times of the PHD and CPHD filters. As execution




Fig. 11: Nondimensionalized recursion execution times.

time is highly system-dependent, all of the times are divided by the maximum execution time, such that all of the execution times fall between zero and one. As shown in Fig. 11, in this example, the PHD filter is responsible for the highest single-step recursion time. In fact, the recursion time of the CPHD filter is often less than that of the PHD filter, specifically during and following time steps when new targets are initiated.

VII. CONCLUSION

Frameworks for tracking an unknown number of noncooperative targets in nearby orbits from a space-based platform using the probability hypothesis density (PHD) and cardinalized probability hypothesis density (CPHD) filters are presented. The relative motion of several targets and the number of targets are estimated by propagating the first-moment approximation of the multitarget probability density function (pdf) and the target cardinality distribution, respectively. A stereoscopic imaging system onboard the inspector satellite produces noisy inspector-to-target angular measurements and sporadic false alarms, both of which are processed jointly to refine the state and cardinality estimates. A measurement-driven model for introducing previously untracked targets into the tracking solution is formulated which makes use of the stereoscopic geometry and relative Lambert targeting.

To evaluate the performance of the presented frameworks, a satellite debris cloud tracking problem is simulated using synthetic measurements, in which stereoscopic angle measurements are corrupted with 750 arcseconds standard deviation noise. A total of nine targets are tracked, and the Gaussian mixture implementations of the PHD and CPHD filters are compared based on their state estimate accuracy, cardinality estimate accuracy, and recursion execution time. In the presented example, the CPHD filter consistently produces solutions that are far more accurate than the PHD filter in their state and cardinality estimates. Furthermore, when the Gaussian mixture representations of the intensity are subjected to the same pruning/merging rules, the resulting CPHD intensity solution often contains significantly less components, and thus benefits from shorter total execution times in its subsequent recursions. In the presented example, the highest execution time savings of the CPHD filter over the PHD filter

are observed when new targets are introduced into the tracking solutions times.

REFERENCES

- [1] S. Segal, P. Gurfil, and K. Shahid, "In-orbit tracking of resident space objects: A comparison of monocular and stereoscopic vision," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 50, no. 1, pp. 676–688, January 2014.
- [2] D. Fourie, B. E. Tweddle, S. Ulrich, and A. Saenz-Otero, "Flight Results of Vision-Based Navigation for Autonomous Spacecraft Inspection of Unknown Objects," *Journal of Spacecraft and Rockets*, pp. 2016–2026, 2014.
- [3] S. D. G. Gaias and J. Ardaens, "Angles-only navigation to a noncooperative satellite using relative orbital elements," *Journal of Guidance, Control, and Dynamics*, vol. 37, no. 2, pp. 439–451, 2014.
- [4] K. A. LeGrand and K. J. DeMars, "Bearing-Only Initial Relative Orbit Determination," to appear, *Journal of Guidance, Control, and Dynamics*.
- [5] United Nations, "Technical report on space debris."
- [6] A. Poore and N. Rijavec, "A Lagrangian Relaxation Algorithm for Multidimensional Assignment Problems Arising from Multitarget Tracking," *SIAM Journal on Optimization*, vol. 3, pp. 544–563, 1993.
- [7] K. J. DeMars, I. I. Hussein, C. Frueh, M. K. Jah, and R. S. Erwin, "Multiple Object Space Surveillance Tracking using Finite Set Statistics," to appear, *Journal of Guidance, Control, and Dynamics*.
- [8] B. A. Jones, S. Gehly, and P. Axelrad, "Measurement-based Birth Model for a Space Object Cardinalized Probability Hypothesis Density Filter," in *Proceedings of the 2014 AIAA/AAS Astrodynamics Specialist Conference*, no. AIAA 2014-4311, 2014.
- [9] R. Mahler, "PHD Filters of Higher Order in Target Number," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 43, no. 4, pp. 1523–1543, October 2007.
- [10] B. Vo, B. Vo, and A. Cantoni, "Analytic implementations of the cardinalized probability hypothesis density filter," *IEEE Transactions on Signal Processing*, vol. 55, no. 7, pp. 3553–3567, July 2007.
- [11] G. Inalhan, M. Tillerson, and J. P. How, "Relative Dynamics and control of spacecraft formations in eccentric orbits," *Journal of Guidance, Control, and Dynamics*, vol. 25, no. 1, pp. 48–59, 2002.
- [12] D. Gim and K. T. Alfriend, "State transition matrix of relative motion for the perturbed noncircular reference orbit," *Journal of Guidance, Control, and Dynamics*, vol. 26, no. 6, pp. 956–971, 2003.
- [13] R. P. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Artech House Boston, 2007.
- [14] B. Vo and W. Ma, "The gaussian mixture probability hypothesis density filter," *IEEE Transactions on Signal Processing*, vol. 54, no. 11, pp. 4091–4104, 2006.
- [15] B. Vo, B. Vo, and A. Cantoni, "The Cardinalized Probability Hypothesis Density Filter for Linear Gaussian Multi-Target Models," in *Information Sciences and Systems, 2006 40th Annual Conference on*, March 2006, pp. 681–686.
- [16] B. Vo, S. Singh, and A. Doucet, "Sequential monte carlo implementation of the phd filter for multi-target tracking," in *Proceedings of the Sixth International Conference of Information Fusion*, vol. 2, July 2003, pp. 792–799.
- [17] E. Rosten and T. Drummond, "Machine Learning for High-speed Corner Detection," in *Computer Vision—ECCV 2006*. Springer, 2006, pp. 430–443.
- [18] H. Bay, A. Ess, T. Tuytelaars, and L. Van Gool, "Speeded-up Robust Features (surf)," *Computer Vision and Image Understanding*, vol. 110, no. 3, pp. 346–359, 2008.
- [19] K. M. Jorgensen, *Using Reflectance Spectroscopy to Determine Material Type of Orbital Debris*, 2000.
- [20] S. Julier and J. Uhlmann, "Unscented Filtering and Nonlinear Estimation," *Proceedings of the IEEE*, vol. 92, no. 3, pp. 401–422, 2004.
- [21] Y. Wang, J. Wu, A. Kassim, and W. Huang, "Data-Driven Probability Hypothesis Density Filter for Visual Tracking," *Circuits and Systems for Video Technology, IEEE Transactions on*, vol. 18, no. 8, pp. 1085–1095, Aug 2008.
- [22] O. Erdinc, P. Willett, and S. Coraluppi, "The gaussian mixture cardinalized phd tracker on mstw and seabar'07 datasets," in *FUSION*, 2008, pp. 1–8.
- [23] D. Schuhmacher, B. Vo, and B. Vo, "A Consistent Metric for Performance Evaluation of Multi-Object Filters," *IEEE Transactions on Signal Processing*, vol. 56, no. 8, pp. 3447–3457, 2008.