

Application of Multi-Hypothesis Sequential Monte Carlo for Breakup Analysis

W.R. Faber, Waqar Zaidi, Islam I. Hussein, Christopher W.T. Roscoe, Matthew P. Wilkins
*Aerospace Engineer, Applied Defense Solutions, 10440 Little Patuxent Pkwy, Columbia, MD
21044*

Paul W. Schumacher, Jr.
Air Force Research Laboratory, 550 Lipoa Parkway, Kihei, Hawaii 96753.

ABSTRACT

As more objects are launched into space, the potential for breakup events and space object collisions is ever increasing. These events create large clouds of debris that are extremely hazardous to space operations. Providing timely, accurate, and statistically meaningful Space Situational Awareness (SSA) data is crucial in order to protect assets and operations in space. The space object tracking problem, in general, is nonlinear in both state dynamics and observations, making it ill-suited to linear filtering techniques such as the Kalman filter. Additionally, given the multi-object, multi-scenario nature of the problem, space situational awareness requires multi-hypothesis tracking and management that is combinatorially challenging in nature. In practice, it is often seen that assumptions of underlying linearity and/or Gaussianity are used to provide tractable solutions to the multiple space object tracking problem. However, these assumptions are, at times, detrimental to tracking data and provide statistically inconsistent solutions. This paper details a tractable solution to the multiple space object tracking problem applicable to space object breakup events. Within this solution, simplifying assumptions of the underlying probability density function are relaxed and heuristic methods for hypothesis management are avoided. This is done by implementing Sequential Monte Carlo (SMC) methods for both nonlinear filtering as well as hypothesis management. This goal of this paper is to detail the solution and use it as a platform to discuss computational limitations that hinder proper analysis of large breakup events.

1. INTRODUCTION

This paper is an expansion of previous work, Faber et al. [1] to discuss impacts on computational burden caused by the size of a breakup event. The size of a breakup event such as an explosion or collision can be characterized by the minimum characteristic length. As new, more powerful sensors come on line the reward of detecting fragments with smaller characteristic length comes at the added expense of tracking the fragments. In Faber et al., a simulated GEO breakup event with eleven visible fragments was successfully tracked using the methodology presented therein. For the sake of making this paper self-contained, the results and discussion of that paper are included. The NASA Orbital Debris Program Office reported that within the first seven years of the 21st century there had been over 10 confirmed debris events such as a breakups, collisions, or intercepts. They also estimate that 94% of the GEO population are debris [2]. Increased debris may lead to cascading debris events otherwise known as the Kessler Syndrome. The aim of this research is to aid in preventing such a scenario by increasing debris and breakup analysis capabilities. There are many challenges to providing meaningful SSA data for highly ambiguous events such as RSO breakups. One of these challenges is debris fragment state initialization. This is a challenge caused by the available measurement types for SSA applications, limited, noisy sensor resources, and detection probability of the debris. Typically, only partial state measurements are available for debris fragments, that is, if they are observable at all. Considerable research has been done to help the SSA community initialize uncorrelated RSO with partial state observations. One method, for example, is the Constrained Admissible Region (CAR) [3] approach. This approach takes advantage of previous knowledge of RSO dynamics and conservation to provide a representation of the debris fragments Probability Density Function (PDF). In recent work, a probabilistic representation of the CAR called the Probabilistic Admissible Region (PAR) [4-6] has been developed for initial object PDF representation. Both CAR and PAR representations have been implemented for multiple measurement types. However, even after accurately initializing the PDF the challenge remains to maintain an accurate representation throughout prediction and update. The choice of how to represent an object's PDF can drastically impact tracking accuracy and computational burden. Typically, object PDFs are assumed to be Gaussian so that solving the Bayesian prediction and update equations is computationally tractable. However, Gaussian PDFs are known to become non-Gaussian through orbit propagation and may not accurately represent the true object PDF. Also, in situations like breakup events where there are many ambiguous measurement returns exhaustively generating the data association hypotheses can be computationally intractable. There is a significant amount of literature expanding the techniques to other PDF representations and

hypothesis generation techniques such as [7,8]. Many multiple hypothesis filters have shown promise in providing accurate tracking data for the breakup problem. For example, in the work by Kelecy et al [9], a multiple hypothesis filter called CAR-MHF is used to maintain the hypotheses and develop estimations of true debris fragment state. In that work, a CAR was used to initialize the underlying object PDF and a technique called MH-JPDA was used for Observation To Track Association (OTTA).

The following paragraph will highlight the techniques used within and the overall contributions to the SSA community. We utilize the PAR for fragment PDF initialization. In particular, we use a four-dimensional PAR that is capable of providing an initial PDF representation given raw angles only data. The benefit to using the PAR approach is that it allows for immediate computations once raw measurements are received. This avoids error caused by post measurement processes that relay data in a non-probabilistic fashion. The multi-object filtering aspect will be handled using a Random Finite Set (RFS) based technique called R-FISST [10]. This approach contains no assumptions of the underlying RSO PDF, allows for the choice of which continuous filter (UKF, PF, PGM, etc..) to use, and keeps the Data Association Problem (DAP) computationally tractable using an MCMC based sampling technique called Smart Sampling Markov Chain Monte Carlo (SSMCMC). SSMCMC is known to keep the problem tractable while sampling the set of the most probable hypotheses even in scenarios where the total number of possible hypotheses is on the order of 10^{36} [11]. Within the R-FISST technique the RSO PDF will be represented using the general (Sequential Monte Carlo) SMC methods presented in Hussein et al. [12] and the underlying PDF will be updated using the Particle Gaussian Mixture (PGM) technique [13].

The rest of the paper is structured as follows. Section 2 will contain a description of the overall solution framework within the multi-hypothesis filter R-FISST. This section will also discuss the advantages and disadvantages of the methods used within the solution framework. The application section 3 will describe the models used within the simulation. The results are shown in section 4 including a discussion on the computational burden associated with an increase in sensor capability. Lastly, section 5 will discuss the conclusions and future work.

2. MULTIPLE OBJECT TRACKING USING R-FISST

2.1 Randomized Finite Set Statistics (R-FISST)

The SMC approach introduced in Hussein et al. [12]. and expands to multi-object scenarios. The multiple object aspect is handled using R-FISST. The R-FISST methodology and the overall solution framework is illustrated in figure [1]. Figure [1] and equations [1-4b] will be referenced throughout the manuscript to aid in understanding. R-FISST utilizes a hypothesis level derivation of the full FISST Bayesian recursions. To make this manuscript self-contained the derivation is summarized as follows. It starts by conditioning the FISST PDF to create a set of “parent” hypotheses. Assume that the number of objects is M , and each object state resides in \mathbb{R}^N . Consider some time instant t_0 , and the data available for the multi-object tracking problem until the current time \mathcal{F}^{t_0} . Let H_i denote the i^{th} hypothesis at time t_0 , and let $\{X\}$ denote the underlying continuous state. For instance, given the M -object hypothesis, the underlying state space would be $\{X\} = \{X_1, X_2, \dots, X_M\}$ where X_n denotes the state of the n^{th} object under hypothesis H_i and resides in \mathbb{R}^N . Let $p(\{X\}, i | \mathcal{F}^{t_0})$ denote the joint distribution of the state-hypothesis pair after time t_0 . Using the rule of conditional probability, the parent hypothesis is described as:

$$p(\{x\}, i | \mathcal{F}^{t_0}) = \underbrace{p(\{x\} | i, \mathcal{F}^{t_0})}_{\text{JPDF underlying, } H_i} \underbrace{p(i | \mathcal{F}^{t_0})}_{\text{prob. of } H_i}. \quad (1)$$

To propagate forward to a time when the next measurement is received, t_k , the R-FISST technique first considers a pseudo prediction step. Within this step the “parent” hypothesis is split into “children” hypotheses. Using the rules of total and conditional probability, it follows that the predicted multi-object PDF in terms of a child hypothesis, j , is:

$$p^-(\{X\}, (i, j) | \mathcal{F}^{t_0}) = \underbrace{\int p(\{X\} | (i, j), \{X'\}) p(\{X'\} | i, \mathcal{F}^{t_0}) d\{X'\}}_{p^-(\{X\} | (i, j), \mathcal{F}^{t_0})} \underbrace{p(j | i)}_{p_{ij}} \underbrace{p(i | \mathcal{F}^{t_0})}_{w_i}. \quad (2)$$

Note that l_{ij} is likelihood of the data $\{Z_{t_k}\}$ given the multi-object PDF underlying hypothesis H_{ij} , and the particular data association that is encoded in the hypothesis, i.e., $z_i \rightarrow x_{ji}$. This derivation provides the ability to separate the posterior density function Eq. [3] into two terms linked by the association likelihood Eqs. [4a, 4b].

$$\frac{p(\{Z_{t_k}\}|\{X\}, (i, j)) p^-(\{X\}|(i, j), \mathcal{F}^{t_0})}{l_{ij}} \quad (4a)$$

$$\frac{l_{ij} \overbrace{p_{ij} w_{ij}}^{w_{ij}}}{\sum_{i', j'} \underbrace{l_{i', j'} p_{i', j'} w_{i', j'}}_{w_{i', j'}}} \quad (4b)$$

Eq. [4b] is known given the association likelihoods, and allows one to compute the weight of a particular hypothesis without first enumerating the filter type update in Eq. [4a]. This is of particular importance for the implementation of random sampling techniques that would ideally sample based on the posterior probability of a hypothesis. Figure [1], illustrates the overall solution recursive frame work presented in this manuscript. The framework is organized in a top down fashion with five major component blocks. Blocks 1,2,3 on the right-hand side of the figure (light blue) describe the initialization, prediction, and update steps. These steps are visited in a recursive fashion throughout the simulation. Blocks 4 and 5 on the left (dark blue) are aspects inherent to the scenario and simulation environment and will be described in the applications section below.

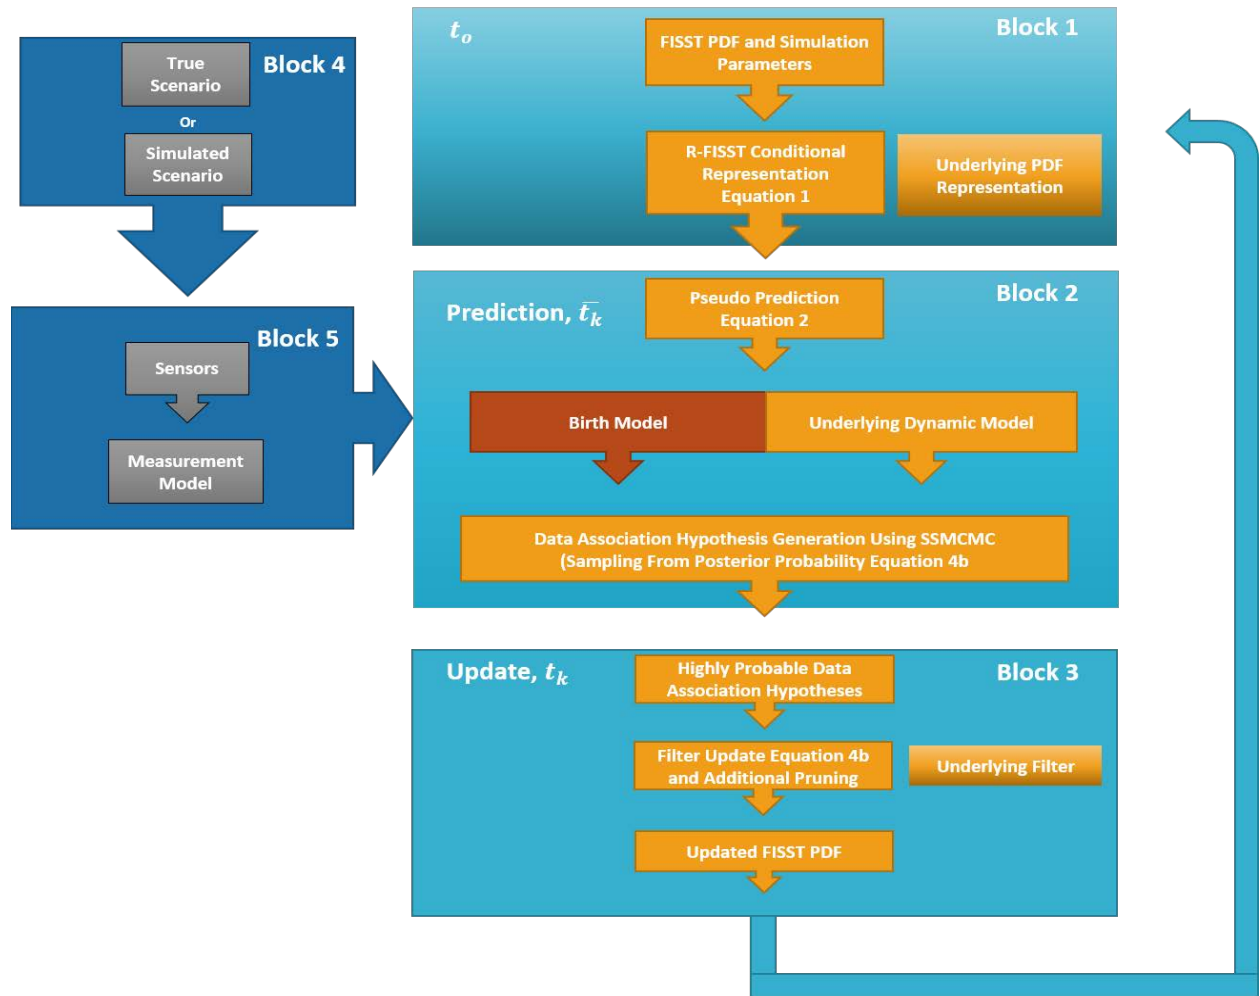


Figure 1. Block diagram of the overall solution recursive framework

2.2 Sequential Monte Carlo Object PDF Representation

Within the parent hypothesis, Eq. [1], $p(\{X\}, i | \mathcal{F}^{t_0})$ is the multi-object PDF constructed of M independent object PDF. Solving the equations presented in the previous section is easiest when the object PDF is approximated as a multivariate Gaussian distribution. This, coupled with normally distributed measurement noise, provides a normal joint distribution whose update can be expressed easily using many extensions of the common Kalman Filter. However, it is known that this assumption falls apart because the object PDF becomes non-Gaussian through propagation. This can lead to type one and two errors and tracking inaccuracy. In the R-FISST framework, there are no assumptions about the form of the underlying PDF (figure [1], block 1), $p_k(\cdot)$; however, it may be computationally intractable to solve the prediction and update equations for some PDFs. Thus, we approximate the PDF using a SMC particle representation. This has the benefit of properly representing the PDF through propagation at the added expense of computational burden. The expense is due to the requirement of propagating all particles instead of just propagating the mean and covariance of a PDF as in a Gaussian case.

2.3 Data Association Using SSMCMC

When tracking scenarios involving closely spaced objects, the number of ambiguous measurement returns can cause the problem to become computationally intractable. This is caused by the Data Association Problem (DAP) that has combinatorial growth dependent on the number of ambiguous measurement returns and is further complicated by allowing birth and death hypotheses. Exhaustively generating all the possible hypotheses is both expensive and unnecessary since an overwhelming majority of the hypotheses will have very low probability. R-FISST uses (figure [1], block 2) a Markov Chain Monte Carlo (MCMC) based technique called Smart Sampling MCMC (SSMCMC) one can determine the highly probable hypotheses without exhaustively generating all possible hypotheses. SSMCMC was developed to account for situations where the total number of possible hypotheses is too large to search through using standard random sampling techniques. For instance, when using standard MCMC implementations, a uniform proposal distribution within the MCMC is too vague and the chances of sampling a probable hypothesis is vanishingly small. SSMCMC uses a biased proposal distribution that allows it to capture the highly probable hypotheses even in situations where the number of possible hypotheses is $\sigma(10^{36})$. For more details on SSMCMC please refer to Faber et al. [11].

2.4 Birth Model and Overview of the Probabilistic Admissible Region

When an uncorrelated measurement is received it can be hypothesized to originate from one of three possible sources: 1) a clutter source, 2) a previously known object that was not accurately represented using the base dynamic model, or 3) a previously unknown object. In the case where we hypothesize that the observation came from a previously unknown object, we are faced with the issue of initializing its PDF. In previous implementations, the birth model predicted where objects would be born in the FOV and assigned a probability of birth to represent the likelihood of occurrence. In this paper, we only consider birth from uncorrelated measurement returns. To initialize the PDF given an angles-only measurement, we use the Four-Dimensional PAR (4DPAR) approach. The approach is described in detail and developed in a multiple hypothesis framework in Hussein et al. [5,12].

2.5 PGM Filter

The last block in figure [1] is the update step. The update of Eq. [4a] can be performed using any appropriate continuous filter. The choice of filter is dependent on the representation of the RSO PDF. In this paper, a Particle Gaussian Mixture (PGM) filter is employed [13] due to its ability to update a particle representation without strictly forcing Gaussian assumptions. This filter was created to take advantage of the particle ensemble representation for propagation but alleviates the issues caused by particle updates such as particle depletion in the Particle Filter (PF). The PGM filter works in three steps. The first is to approximate the PDF represented by the particles using a Gaussian Mixture (GM). This is accomplished by clustering the particles into an optimal number of components using an optimal K-means algorithm. The second step is to update the GM representation using in our case an Unscented Kalman Filter update. The last step is performed by re-sampling the resulting posterior GMM. Figure [2] illustrates the use of a PGM filter update on a newly initialized PAR4D PDF [12].




Figure 2. PAR representation updated throughout time using PGM

3. APPLICATION

This section describes the application used to test the methods presented in this paper. First, the test scenario is outlined and justified. Then the underlying breakup, dynamic, and measurement models are described in detail. The platform specifications used to run this simulation and tests are detailed in the last sub section. The goal of this application is to show the capability of tracking a breakup scenario using the solution framework presented in this paper

3.1 GEO Breakup Scenario

To test the aforementioned techniques, the authors simulate the explosion of a well tracked GEO object. The breakup will be simulated using the underlying breakup and dynamic models below. The breakup is simulated to occur on the 18th of September year 2020. Due to the random nature of the breakup simulation, many tests are conducted and considered for analysis. Between tests, any change of simulation parameters will be documented and described in the results section. All the tests assume that the RSO is initially in the field of view (FOV) of a single sensor and will be tracked in stare mode continuously for a minimum of four hours prior to breakup. Data will be received in ten-minute intervals according to the measurement model described below. This simulation setup means that it can be reasonably assumed that the RSO statistics before the breakup are well known. Additionally, we have constrained the problem so that the breakup occurs in the sensor FOV, which obviates the need to approximate the breakup time and location or to search for all fragments from the event.




Figure 3 GEO breakup over a one-day period plotted in thirty-minute intervals in J2000ECI

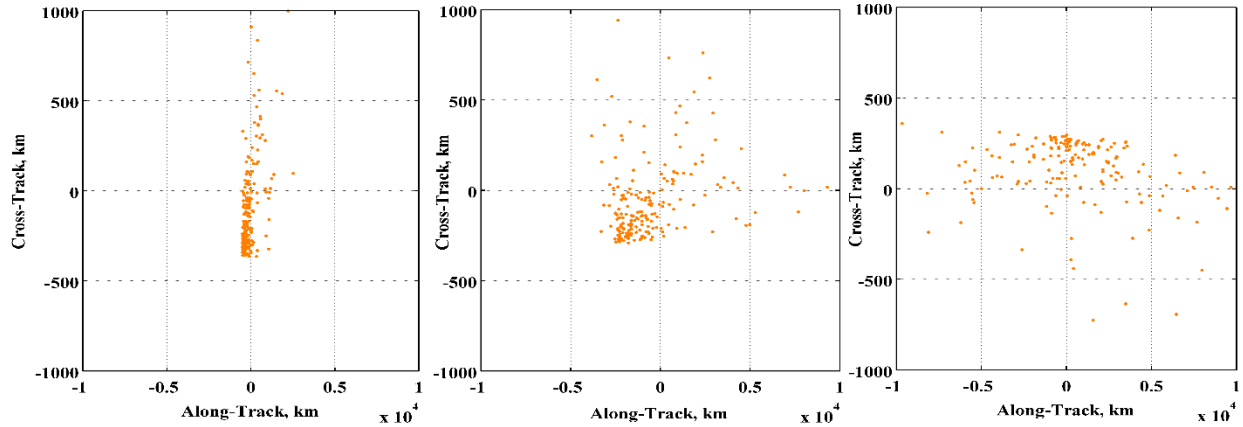


Figure 4. In Track and Cross Track positions for GEO Breakup fragments six hours (left), twelve hours (center), and one day (right) after the breakup.

3.2 Breakup and Dynamic Model

To accurately simulate the breakup event, a breakup model was developed using a version of the NASA standard breakup model [14]. This model characterizes the size distributions, area-to-mass and impact velocity assignments and distributions, and delta-velocity distributions where relevant for both colliding or exploding rocket body and spacecraft. Thus, for a spacecraft explosion in the GEO orbit regime, the number of debris fragments, distribution of area-to-mass, and the distribution of delta-velocities were developed directly from the NASA standard breakup model. The true RSO orbital parameters before the event are displayed in table [1]. The breakup was simulated using the scaling factor $s = 1$. An example of a GEO breakup event simulated using this model is provided in figures [3,4]. Due to the computational complexity of breakup scenarios and the limited ability to optically track GEO fragments, only fragments with characteristic length $70\text{cm} < L_c < 100\text{cm}$ were considered. This is further motivated by the fact that objects of $L_c < 70\text{cm}$ will have a much lower probability of detection. Also, objects of $L_c > 100\text{cm}$ may have dynamic effects that are not accounted for in this proof of concept simulation. All spacecraft and fragment dynamics were modeled using Orekit analytical propagation.

Element	Value
Semi-Major axis, a	42111123.663 (m)
Eccentricity, e	0.039907 (rad)
Inclination, i	0.0017450 (rad)
Argument of Perigee, ω	2.5587×10^{-5} (rad)
Right Ascension of Ascending Node, Ω	3.5079 (rad)
True Anomaly, ν	6.0153 (rad)

Table 1. RSO Orbital Elements

3.3 Measurement Model

A single angles-only optical sensor located on the island of Maui, HI was simulated to capture the event. For each fragment, the sensor provides raw measurements of the topocentric J2000 right ascension and declination. The sensor suffers from observation noise that is assumed zero-mean Gaussian with a standard deviation of 2 arc second in both right ascension and declination. The sensor has a non-zero rate of false alarm. False alarms are modeled as a

Bernoulli distribution with a probability of $p_f = 0.2$. The look direction of the sensor is set to track the expected mean of the initial RSO. Fragments have a probability of detection equal to zero when they are outside the FOV and at times when there is too much ambient light on the optical sensor. When fragments are within the FOV they have probability of detection equal to $p_d = 0.9$. This detection model is shown in Eq. [5].

$$p_d(x) = \begin{cases} 0.9 & \text{If } x \in \mathcal{T} \\ 0 & \text{otherwise} \end{cases}, \quad (5)$$

where x is the fragment's state and \mathcal{T} represents the region where line of site (LOS) calculations for both object sun and object ground sensor are satisfied and LOS for sun ground sensor is not.

3.4 Platform

The test cases performed in the following section were completed using a Dell Precision 7510 running Windows 10 with an Intel(R) Core(TM) i7-6920HQ CPU at 2.90GHz. The platform contained 32GB of Random Access Memory (RAM) and operated on a 64-bit system. The test code was written in the SCALA functional programming language using the ECLIPSE Integrated Development Environment.

4. RESULTS

4.1 Performance

The methods presented in this paper were tested in a variety of ways using the simulated scenario provided above. The tests were designed to show proof of concept and to determine performance in terms of filter accuracy, and computational efficiency. To show proof of concept and provide a lower bound on the computational burden, a first test case was implemented without using the PAR approach for fragment PDF initialization. In place of PAR, a simple placeholder birth PDF was created using an assumed full state measurement. This allowed for the ability to show that the core aspects of the methodology were performing as expected. Specifically, it shows that for the given scenario the R-FISST technique can keep the DAP computationally tractable while maintaining an accurate representation of the true fragment states. Also, it allows us to analyze the SMC particle representation of debris fragment PDF and the PGM filter without seeing the PAR's effect on accuracy. The test is initialized with a FISST PDF that contained a set of three hypotheses with zero objects, one object, and two objects respectively. The object PDFs were represented by particle ensembles containing 1000 particles. The particles were initially sampled from a Gaussian distribution with a mean value near the initial RSO true ECI J2000 position and velocity with position standard deviation of $\sigma_p = 1000 \text{ m}$ and velocity standard deviation $\sigma_v = 10 - 2 \text{ m/s}$. The total simulation time was a twelve-hour period around the object breakup (7 hours before and 5 hours after). Recall from the simulation setup that measurements of the initial RSO and post breakup debris fragments were provided and analyzed in ten-minute intervals. The breakup occurs in between simulation time steps fifty-three and fifty-four and results in a total of eleven debris fragments.

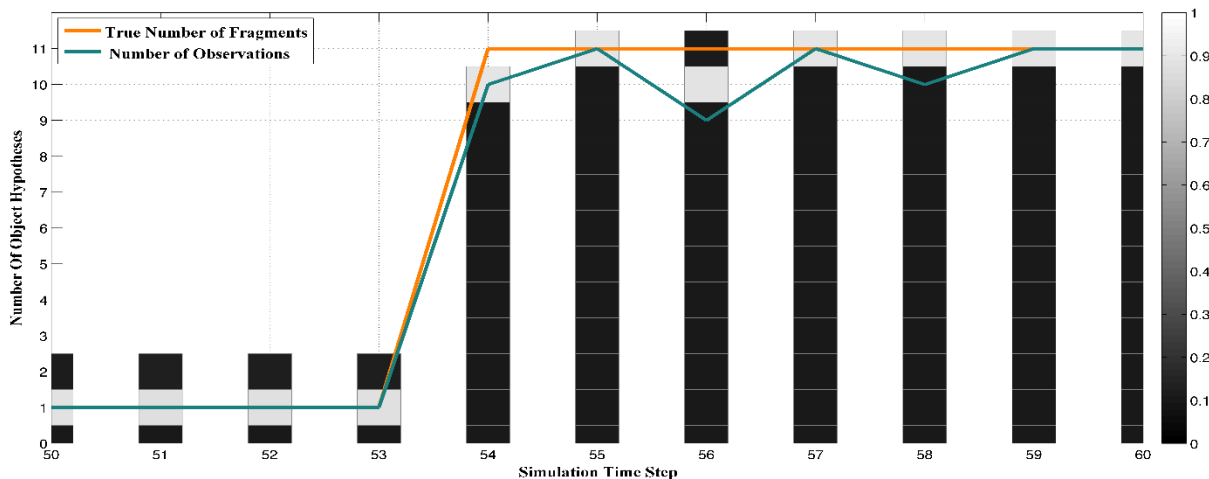


Figure 5. Probabilities from the top hypotheses sorted by the number of objects within.

Figure [5] provides an illustration of how the probabilities of the hypotheses shift throughout the simulation. The Y-axis of figure [5] is a label for the top hypotheses containing the corresponding number of objects. The X-axis is simulation time step truncated to emphasize the time of interest. Hypothesis weights are shown using a color map with lighter colors corresponding to higher probability. Also, in figure [5] the true number of objects and the number of observations at each time step are plotted using the orange and blue lines respectively. Before the breakup, as expected the hypothesis containing one object is most probable. After the event, the most probable hypothesis varies between ten fragments and eleven fragments until it settles on the true correct number of fragments forty minutes after the breakup. The number of observations varies due to the non-zero false alarm rate and the non-zero probability of misdetection. The conclusion drawn from figure [5] is that the R-FISST framework is capable of maintaining a set of hypotheses including the correct hypothesis even though the number of possible data association hypotheses at each step after breakup is $\sigma(10^8)$.

Figures [6,7] are used to discuss accuracy and uncertainty. Figure [6] shows the particle cloud of a single fragment just after breakup (left) and the same fragment at the end of the simulation (right). The particle cloud is shown as a heat map with the higher number of particles represented by darker shades.

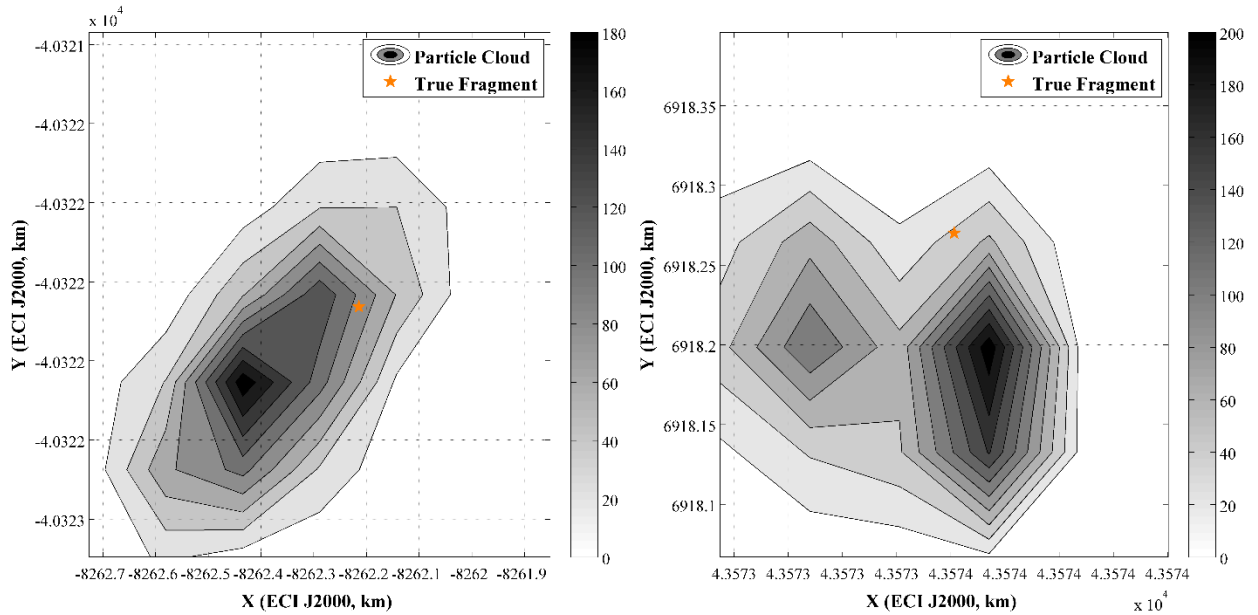


Figure 6. Particle cloud represented as a heat map for a single fragment belief from the top hypothesis just after breakup (left) and at the end of the simulation (right) plotted against the true fragment position

As an illustration to further show the reduction of uncertainty figure [7] shows the total entropy for all objects or fragments within the top hypothesis throughout time. Entropy of a particular particle cloud, S_l is represented,

$$S_l = \frac{1}{2} \lg(|P_l| (2\pi e)^N), \quad (6)$$

Thus, the total entropy for a particular hypothesis, S_{tot} , is as follows

$$S_{tot} = \sum_{l=1}^M S_l = \frac{1}{2} \sum_{l=1}^M \lg(|P_l| (2\pi e)^N), \quad (7)$$

where P_l is the LU factorization of the ensemble covariance of the particle cloud and N is the dimension of the covariance. It can be seen in figure [7] that as measurements are received the entropy reduces. Once the breakup occurs the initial entropy spikes however as more measurements are received and correct associations are made the total entropy begins to trend downward again.

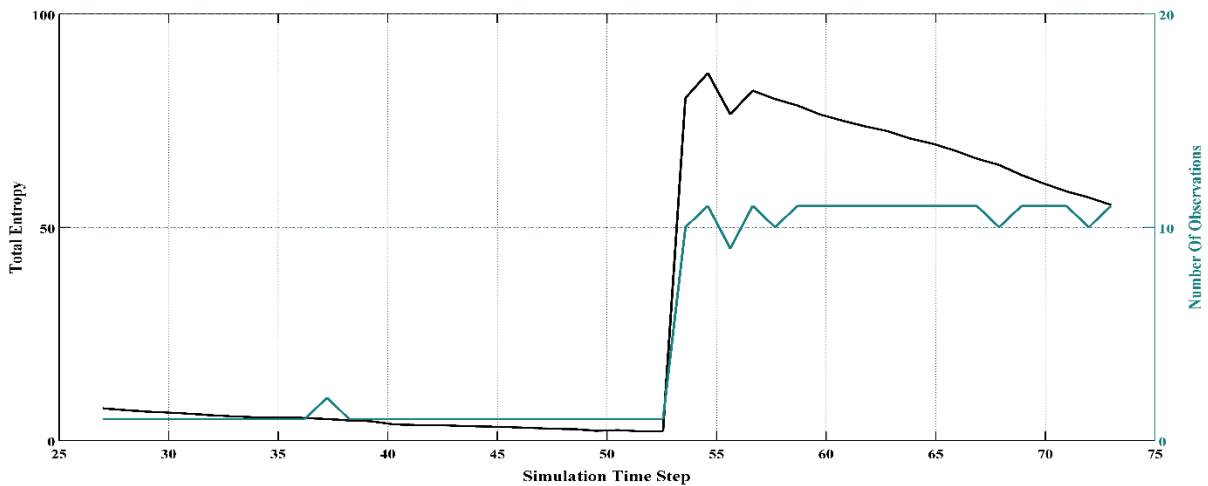


Figure 7 The left Y axis represents the total particle cloud entropy for all fragment PDFs within the top hypothesis plotted as a function of time. The Y axis on the right-hand side represents the number of observations as a function of time

After establishing a base test case and confirming the methodology presented in this paper performs as expected, another test case was created to incorporate the PAR fragment PDF initialization. The total simulation time was an eight-hour period around the object breakup (four before and four after). In this example, the PAR was initiated at only one-time step that correlated to when the first measurement interval after breakup. At that time, eleven uncorrelated angles only measurements were received and processed to form eleven PAR PDF representations. Figure [8], provides an illustration of a PAR for a single debris fragment. The bi-modality in ECI J2000 velocity space is a natural consequence of creating the PAR from raw angles only data due to the fact that no information on the direction of velocity is provided. With this in mind, the accuracy of the approach will be affected drastically when considering the mean error until enough data is received to collapse the density to a unimodal PDF. Figures [9-11] are included to show a comparison to the first test case. Figure [9] illustrates the shift of probability between hypotheses throughout the simulation and shows x, y that the hypothesis containing the true number of fragments is most probable.

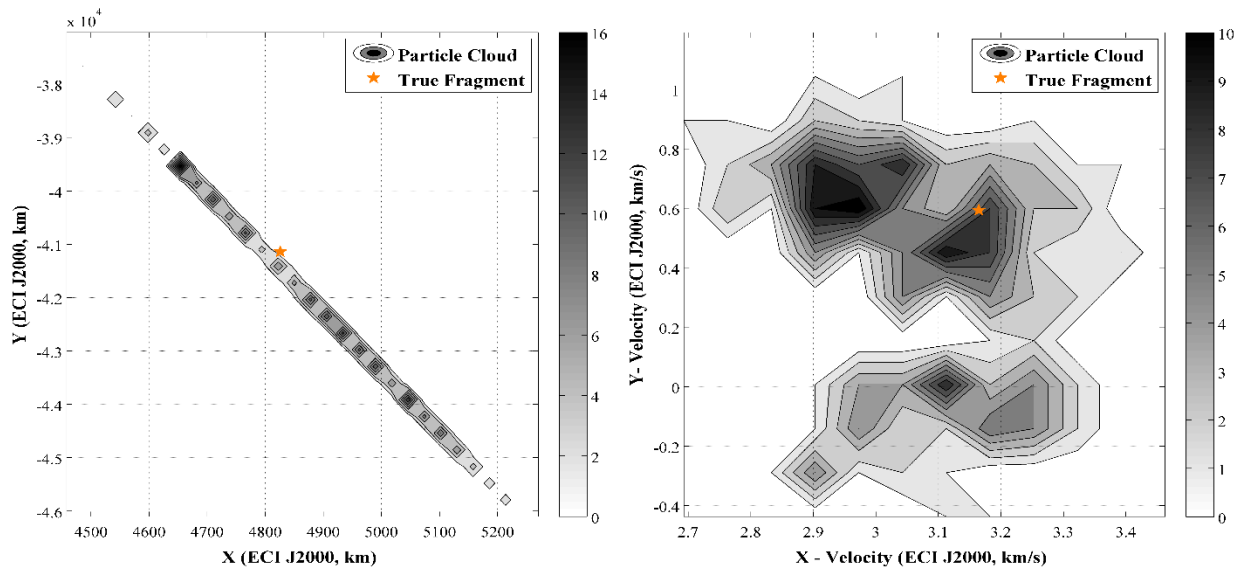


Figure 8. PAR PDF representation from angles only measurement plotted in ECI J2000 position (left) and velocity (right)

Figure [10] shows a particle ensemble's position for a single debris fragment after breakup (left) and at the end of the simulation (right) and the reduction in uncertainty is shown using the concept of total entropy in figure [11]. It is

important to note that in figure [11] the total entropy immediately after the breakup is higher than in figure [7] due to the higher initial uncertainty caused by generating the initial PDF using the PAR approach. However, the total entropy over time is reduced. Overall, these figures are consistent with those seen in the first test case and show that the PAR method together with the methods presented are capable of tracking breakup events.

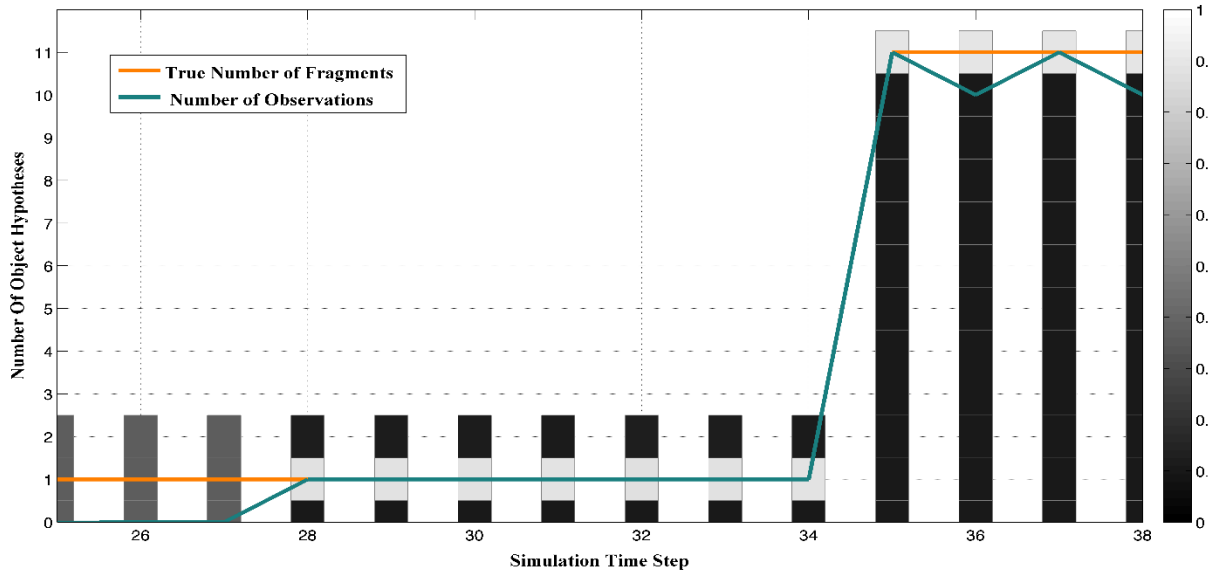


Figure 9. Test case where the PAR is used for fragment PDF initialization. Probabilities from the top hypotheses sorted by the number of objects within.

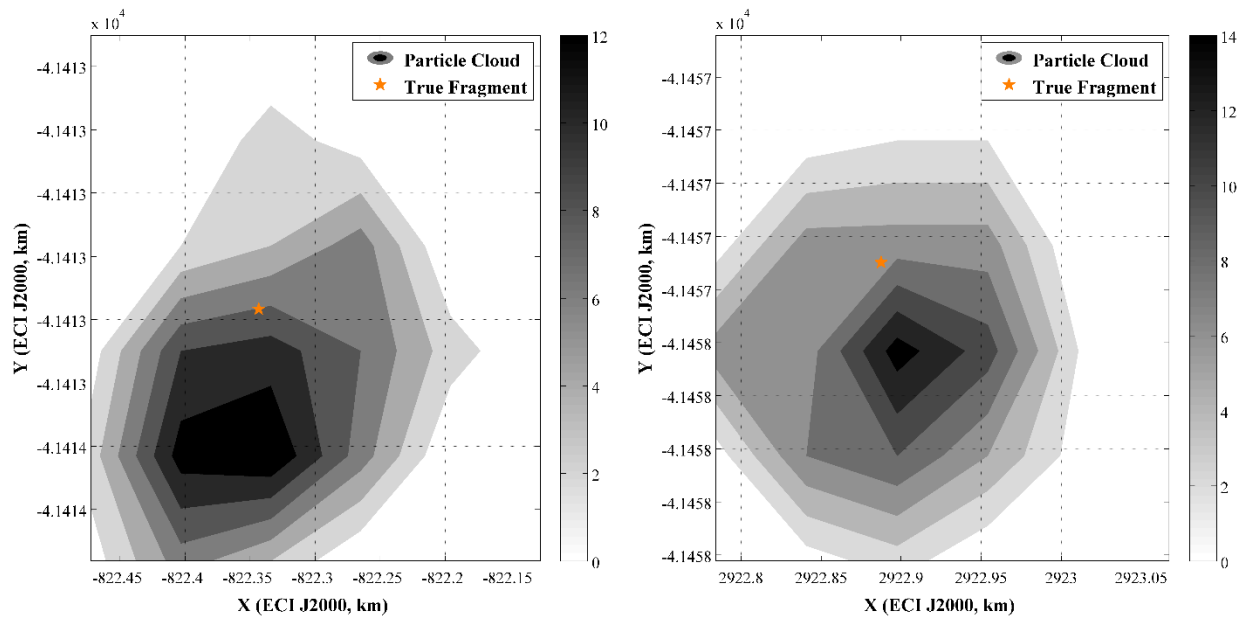


Figure 10. Test case where the PAR is used for fragment PDF initialization. Particle cloud represented as a heat map for a single fragment belief from the top hypothesis after breakup (left) and at the end of the simulation (right) plotted against the true fragment position.

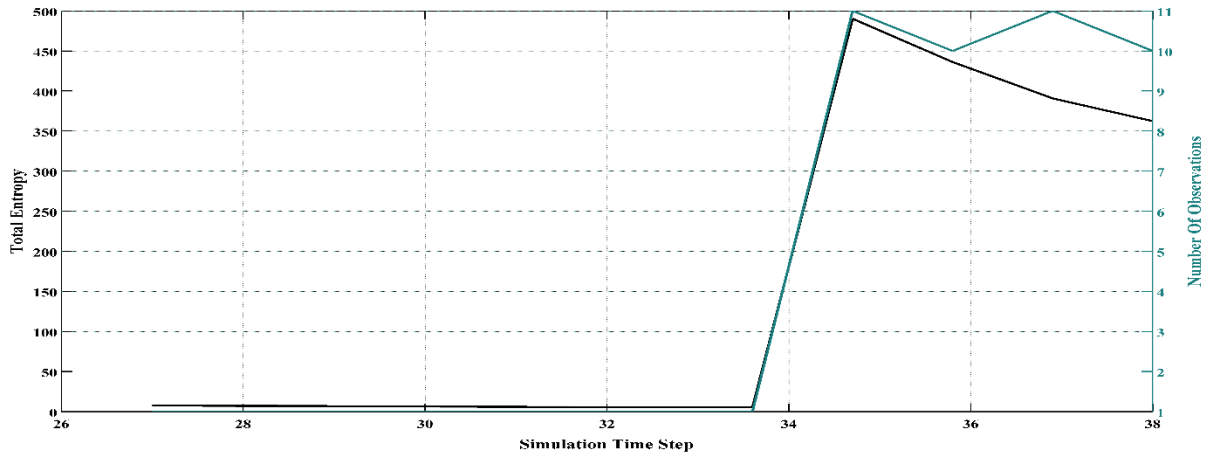


Figure 11. Test case where the PAR is used for fragment PDF initialization. The left Y axis represents the total particle cloud entropy for all fragment PDF within the top hypothesis plotted as a function of time. The Y axis on the right-hand side represents the number of observations as a function of time.

4.2 Computational Burden

As more powerful sensors, such as The Space Fence, begin to operate, the capability of tracking smaller debris is both a blessing and a curse. According to NASA’s Orbital Debris Program Office the Fengyen 1C satellite fragmentation created hundreds of thousands of debris fragments of which only a few thousand have been cataloged. The estimated number of debris greater than 1cm characteristic length is over 150,000 fragments. Consider the case where n newly spawned fragments are detected. The number of possible hypotheses A_n without data association overlap is,

$$A_n = \sum_{k=0}^n \binom{n}{k}. \quad (8)$$

Figure 12 shows the number of possible hypotheses against the minimum characteristic length. This figure is meant to show that as sensors get more powerful the capability of detecting fragments with lower characteristic length improves which in turn causes the number of possible hypotheses in a breakup scenario to become computationally intractable. Exhaustively generating all possible hypotheses is not possible without sampling methods such as SSMCMC or waiting for data to disambiguate. Waiting for data to disambiguate could be detrimental to other RSO and lead to lost debris fragments. The technique discussed in this paper provides a tractable solution that allows for immediate processing of raw measurements.

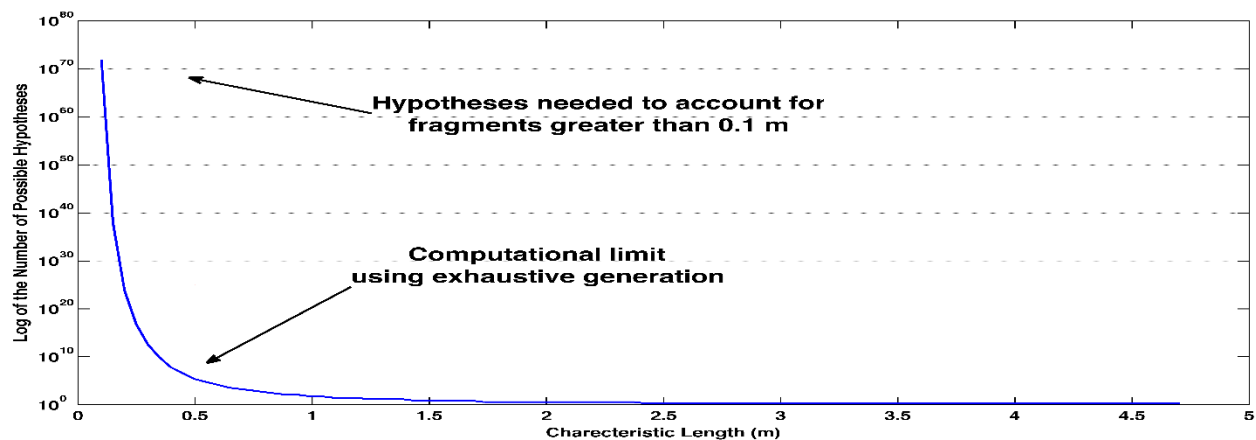


Figure 12. Characteristic length against the number of possible hypotheses needed without data association ambiguities. As visible characteristic length decreases the number of possible hypotheses explodes.

5. CONCLUSIONS

In this paper, we have presented a solution framework for SSA multi-object tracking problems applicable to RSO breakup events. The contributions are that the underlying object PDFs within the framework are free of Gaussian assumptions and are represented by an SMC particle ensemble. The framework also incorporates a four-dimensional PAR that is capable of initializing a birth PDF using raw angles-only observations. This allows for the immediate computation of probabilistic tracking and prevents fragment loss. Also, using raw observations allows us to avoid the computationally burdensome process of observation compression and errors created therein. The framework was tested using a GEO breakup event simulated using a version of NASA's standard breakup model. Results show that the framework was capable of keeping the inherent DAP computationally tractable while converging to the true number of fragments. In future work, topics such as filter performance metrics and breakup size effects on computational complexity will be investigated along with techniques to improve overall performance.

6. ACKNOWLEDGEMENTS

This work would not be possible without the contributions from our colleagues in the SSA community. We would like to thank Professor Suman Chakravorty at Texas A&M for his contribution to the development of R-FISST and the PGM filter. We would also like to thank Tom Kececy for his insight application of CAR-MHF for initial orbit determination of breakup event.

7. REFERENCES

- 1 W. Faber, W. Zaidi, I. Hussein, C. W. T. Roscoe, M. P. Wilkins, and J. Paul W. Schumacher, "Application Of Multi-Hypothesis Sequential Monte Carlo For Breakup Analysis," 2017. (Proceedings of the AAS/AIAA Astrodynamics Specialist Conference, Columbia River Gorge, Stevenson, WA Paper AAS 17-579
- 2 N. L. Johnson, E. Stansbery, D. O. Whitlock, K. J. Abercromby, and D. Shoots, "History of On Orbit Satellite Fragmentations," *Advances in the Astronautical Sciences*, Vol. 14th Edition, 2008. NASA's Orbital Debris Program Office.
- 3 C. W. T. Roscoe, P. W. Schumacher, Jr., and M. P. Wilkins, "Parallel Track Initiation for Optical Space Surveillance using Range and Range-Rate Bounds," *Advances in the Astronautical Sciences*, Vol. 150, 2014, pp. 989–1008. (Proceedings of the AAS/AIAA Astrodynamics Specialist Conference, Hilton Head, SC, August 11–15 2013, Paper AAS 13-767).
- 4 C. W. T. Roscoe, I. I. Hussein, M. P. Wilkins, and P. W. Schumacher, Jr., "The Probabilistic Admissible Region with Additional Constraints," *Advances in the Astronautical Sciences*, Vol. 156, 2015, pp. 117–130. (Proceedings of the AAS/AIAA Astrodynamics Specialist Conference, Vail, CO, August 9–13 2015, Paper AAS 15-577).
- 5 I. I. Hussein, C. W. T. Roscoe, P. W. Schumacher, Jr., and M. P. Wilkins, "Probabilistic Admissible Region for Short-Arc Angles-Only Observation," *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*, Wailea, HI, September 9–12 2014.
- 6 I.I. Hussein, C.W.T. Roscoe, M.P. Wilkins, and P.W. Schumacher, Jr., "Probabilistic Admissibility in Angles-Only Initial Orbit Determination," *Proceedings of the 24th International Symposium on Space Flight Dynamics*, Laurel, MD, May 5–9 2014.
- 7 H. G. Hoang, B. T. Vo, and B. N. Vo, "A Fast Implementation of the Generalized Labeled Multi Bernoulli Filter with Joint Prediction and Update," *Proceedings of the 18th International Conference on Information Fusion*, Piscataway, NJ, IEEE, 2015, pp. 999–1006.
- 8 B.-N. Vo, S. Singh, and A. Doucet, "Random Finite Sets and Sequential Monte Carlo Methods in Multi-Target Tracking," *Proceedings of the International Conference on Radar*, Piscataway, NJ, IEEE, September 2003, pp. 486–491, 10.1109/RADAR.2003.1278790.
- 9 T. Kececy, M. Shoemaker, and M. Jah, "Application of the Constrained Admissible Region Multiple Hypothesis Filter to Initial Orbit Determination of a Break-up," *Proceedings of the 6th European Conference on Space Debris Darmstadt, Germany, 08/2013*.
- 10 W. R. Faber, S. Chakravorty, and I. Hussein, "Multiple Space Object Tracking with R-FISST," *Proceedings of the AAS/AIAA Space Flight Mechanics Meeting Paper AAS 17-477*, San Antonio TX, February 5-9th 2017.

- 11 W. R. Faber, S. Chakravorty, and I. Hussein, "R-FISST and the Data Association Problem with Applications to Space Situational Awareness," Proceedings of the AIAA/AAS Astrodynamics Specialist Conference, Long Beach, CA, September 13-16, 2016.
- 12 I.Hussein, W. Zaidi, W. R. Faber, C.W.T. Roscoe, M.P. Wilkins, J. Paul W. Schumacher, and M. Bolden, "Application of Sequential Monte Carlo Methods for Space Object Tracking," 2017. (Proceedings of the AAS/AIAA Space Flight Mechanics Meeting, San Antonio, TX, Paper AAS 17-415).
- 13 D. Raihan and S. Chakravorty, "Particle Gaussian Mixture (PGM) Filters," Proceedings of the 19th International Conference on Information Fusion, Piscataway, NJ, IEEE, 2016, pp. 1369–1376.
- 14 N.L. Johnson, P.H. Krisko, J. C.Lieu, and P.D.Am-Meador, "NASAs New Breakup Model of Evolve 4.0," Adv. Space Res., Vol. Vol. 28, No. 9, 2001. pp. 1377-1384.