A Randomized Sampling based Approach to Multi-Object Tracking

W. Faber and S. Chakravorty Department of Aerospace Engineering Texas A&M University College Station, TX Islam I. Hussein Applied Defense Solutions Columbia, MD

Abstract—In this paper, we present a randomized version of the finite set statistics (FISST) Bayesian recursions for multiobject tracking problems with application to the space situational awareness (SSA) problem. We introduce a hypothesis level derivation of the FISST equations that shows that the multiobject tracking problem may be considered as a finite state space Bayesian filtering problem, albeit with a growing state space. It further allows us to propose a randomized scheme, termed randomized FISST (R-FISST), where we choose the highly likely children hypotheses using Markov Chain Monte Carlo (MCMC) methods which allows us to keep the problem computationally tractable. We test the R-FISST technique on a fifty object birth and death SSA tracking and detection problem.

I. INTRODUCTION

In this paper, we present a randomized approach to approximate the full Bayesian recursions involved in solving the Finite Set Statistics (FISST) based approach to the problem of multi-object tracking and detection, in particular, to the problem of SSA. We show that the FISST recursions can essentially be considered as a discrete state space Bayesian filtering problem on "Hypothesis Space" with the only input from the continuous problem coming in terms of the likelihood values of the different hypotheses. The number of objects is implicit in this technique and can be a random variable. The "Hypothesis Space" perspective allows us to develop a randomized version of the FISST recursions where we sample the highly likely children hypotheses using a Markov Chain Monte Carlo (MCMC) technique thereby allowing us to keep the total number of possible hypotheses under control, and thus, allows for a computationally tractable implementation of the FISST equations, which otherwise grows at an exponential rate, and thus, can quickly lead to the problem becoming intractable. The method is applied to a fifty object SSA tracking and detection problem that has an unfixed number of objects throughout the simulation.

In the last 20 years, the theory of FISST-based multiobject detection and tracking has been developed based on the mathematical theory of finite set statistics [1], [2]. The greatest challenge in implementing FISST in real-time, which is critical to any viable SSA solution, is computational burden. The first-moment approximation of FISST is known as the Probability Hypothesis Density (PHD) approach [2], [3]. The PHD has been proposed as a computationally tractable approach to applying FISST. The PHD filter essentially finds the density of the probability of an object being at a given location, and thus, can provide information about the number of objects (integral of the PHD over the region of interest) and likely location of the objects (the peaks of the PHD). The PHD can further employ a Gaussian Mixture (GM) or a particle filter approximation to reduce the computational burden (by removing the need to discretize the state space). This comes at the expense of approximating the general FISST pdf with its first-moments [3]-[6]. The PHD filter does not attempt to solve the full FISST recursions, in particular, by considering the PHD, the filter gets rid of the data association problem inherent in these problems. In other previous work, a GM approximation was applied, not to the first-moment of the FISST pdfs, but to the original full propagation and update equations derived from FISST [7], [8]. This eliminates any information loss associated with using the first-moment PHD approximation, while at the same time increasing the computational tractability of the multi-object FISST pdfs. This approach is similar in spirit to the concept of the "para-Gaussian" pdf that was described in [9].

In this paper, in contrast, we introduce a hypothesis level derivation of the FISST equations that makes it clear as to how the full FISST recursions can be implemented without any approximation other than the approximations inherent in the underlying tracking filter, such as an extended Kalman filter. In order to ensure the computational tractability of the resulting equations, we introduce an MCMC based hypothesis selection scheme resulting in the Randomized FISST (R-FISST) approach that is able to scale the FISST recursions to large scale problems.

There are also non-FISST based approaches to multihypothesis tracking (MHT) such as the Hypothesis Oriented MHT (HOMHT) [10]–[13], and the track oriented MHT (TOMHT) techniques [14]. The MHT techniques can be divided into single-scan and multi-scan methods depending on whether the method uses data from previous times to distinguish the tracks [11], [13], [15]. The single-scan (recursive) methods such as joint probabilistic data association (JPDA) [13], [15] typically make the assumption that the tracks are independent which is not necessarily true. The multi-scan methods such as TOMHT [13], [14] are not recursive. The primary challenge in these methods is the management of the various different hypotheses related to the tracks which TOMHT does using an efficient tree structure, and the MCMCDA, and other related tracking techniques [16]–[18], do through the use of MCMC methods in sampling the data associations. We also use MCMC to sample children hypotheses given the parent hypothesis, however, our approach is a truly recursive technique which does not assume track independence as the above mentioned single scan methods. We essentially do an efficient management of the growing number of hypotheses at every generation through the judicious use of MCMC.

The rest of the paper is organized as follows. In Section II, we introduce the hypothesis level derivation of the FISST equations. In Section III, we introduce the MCMC based randomized hypothesis selection technique that results in the R-FISST algorithm. In Section IV, we show an application of the R-FISST technique on a fifty object SSA example and discuss the results.

II. A Hypothesis based Derivation of the FISST Equations

In this section, we shall frame the multi-object tracking equations at the discrete hypothesis level (which we believe are equivalent to the FISST equations) which then shows clearly as to how the full FISST recursions may be implemented. The derivation below assumes that the number of measurements is always less than the number of objects, which is typically the case in the SSA problem. We never explicitly account for the number of objects, since given a hypothesis, the number of objects and their probability density functions (pdf) are fixed, which allows us to derive the results without having to consider the random finite set (RFS) theory underlying FISST. Albeit the equations derived are not as general as the FISST equations, in particular, the birth and death models employed here are quite simple, we believe that the level of generality is sufficient for the SSA problem that is our application.

A. Framing FISST at the Hypothesis Level

We consider first the case when the number of objects is fixed, which we shall then generalize to the case when the number of objects is variable, i.e, there is birth and death in the object population. Assume that the number of objects is M, and each object state resides in \Re^N . Consider some time instant t - 1, and the data available for the multi-object tracking problem till the current time \mathcal{F}^{t-1} . Let H_i denote the i^{th} hypothesis at time t - 1, and let $\{X\}$ denote the underlying continuous state. For instance, given the N- object hypothesis, the underlying state space would be $\{X\} = \{X_1, X_2, \cdots X_M\}$ where X_j denotes the state of the j^{th} object under hypothesis H_i and resides in \Re^N . Let $p(\{X\}, i/\mathcal{F}^{t-1})$ denote the joint distribution of the statehypothesis pair after time t - 1. Using the rule of conditional probability:

$$p(\{X\}, i/\mathcal{F}^{t-1}) = \underbrace{p(\{X\}/i, \mathcal{F}^{t-1})}_{\text{MT-pdf underlying}_{H_i}} \underbrace{p(i/\mathcal{F}^{t-1})}_{W_i = \text{prob. of}_{H_i}}, (1)$$

where MT-pdf is the multi-object pdf underlying a hypothesis. Given the hypothesis, the MT-pdf is a product of independent individual pdfs underlying the objects, i.e.,

$$p(\{X\}/i, \mathcal{F}^{t-1}) = \prod_{k=1}^{M} p_k(x_k),$$
 (2)

where $p_k(.)$ is the pdf of the k^{th} object. Next, we consider the prediction step between measurements. Each hypothesis H_i splits into A_M children hypotheses, and let us denote the j^{th} child hypothesis as H_{ij} . The children hypotheses correspond to the different data associations possible given a measurement of size m, i.e., m returns, and

$$A_M = \sum_{n=0}^{\min(m,M)} \binom{M}{n} \binom{m}{n} n!.$$
 (3)

We want to note here that this is a pseudo-prediction step since we assume that we know the size of the return m. However, it allows us to fit the MT-tracking method nicely into a typical filtering framework. Using the rules of total and conditional probability, it follows that the predicted multi-object pdf in terms of the children hypotheses is:

$$p^{-}(\{X\}, (i, j)/\mathcal{F}^{t-1}) = \int p(\{X\}, (i, j)/\{X'\}, i)p(\{X'\}, i/\mathcal{F}^{t-1})d\{X'\} = \underbrace{\int p(\{X\}/(i, j), \{X'\})p(\{X'\}/i, \mathcal{F}^{t-1})d\{X'\}}_{p^{-}(\{X\}/(i, j), \mathcal{F}^{t-1})} \underbrace{p(j/i)}_{p_{ij}} \underbrace{p(i/\mathcal{F}^{t-1})}_{w_{i}},$$
(4)

where $p^{-}(.,(i,j)/\mathcal{F}^{t-1})$ is the joint distribution of the state and hypothesis pairs before the measurement at time t. We have used the fact that $p((i,j)/\{X'\},i) = p(j/i) = p_{ij}$, and p_{ij} is the transition probability of going from the parent i to the child j and w_i is the probability of the parent hypothesis H_i . Let $p_k(x_k/x'_k)$ denote the transition density function of the k^{th} object. Expanding the predicted MT-pdf, we obtain:

$$p^{-}(\{X\}/(i,j),\mathcal{F}^{t-1}) \equiv \int p(\{X\}/(i,j),\{X'\})p(\{X'\}/(i),\mathcal{F}^{t-1})d\{X'\}, \quad (5)$$

where

$$p(\{X\}/(i,j),\{X'\}) \equiv \prod_{k=1}^{M} p_k(x_k/x'_k)$$
$$\int p(\{X\}/(i,j),\{X'\})p(\{X'\}/(i),\mathcal{F}^{t-1})d\{X'\}$$
$$\equiv \int \prod_k p_k(x_k/x'_k) \prod_{k'} p_{k'}(x'_k)dx'_1 \cdots dx'_M$$
$$= \prod_k \int p_k(x_k/x'_k)p_k(x'_k)dx'_k = \prod_k p_k^-(x_k), \quad (6)$$

where $p_k^-(x_k)$ is the prediction of the k^{th} object pdf underlying the hypothesis H_{ij} .

Remark 1: Eq. 4 has a particularly nice hybrid structure: note that the first factor is the multi-object continuous pdf underlying the child hypothesis H_{ij} , while the second factor $p_{ij}w_i$ is the predicted weight of the hypothesis H_{ij} . For the no birth and death case, all p_{ij} are equal to $\frac{1}{A_M}$, where recall that A_M is the total number of data associations possible (Eq. 3). Note that the MT-pdf underlying H_{ij} is simply the product of the predicted individual object pdf, and in the case of no birth and death, it is the same for all children hypothesis H_{ij} .

Given the prediction step above, let us consider the update step given the measurements $\{Z_t\} = \{z_{1,t}, \dots, z_{m,t}\}$, where there are *m* measurement returns. We would like to update the weights of each of the multi-object hypotheses to obtain $p(\{X\}, (i, j)/\{Z_t\}, \mathcal{F}^{t-1})$ by incorporating the measurement $\{Z_t\}$. Using Bayes rule:

$$p(\{X\}, (i, j)/\{Z_t\}, \mathcal{F}^{t-1}) = \eta p(\{Z_t\}/\{X\}, (i, j))p^-(\{X\}, (i, j)/\mathcal{F}^{t-1}),$$

where

$$\eta = (7)$$

$$\sum_{i',j'} \int p(\{Z_t\}/\{X'\}, (i',j')) p^-(\{X'\}, (i',j')/\mathcal{F}^{t-1}) d\{X'\},$$

where the MT-likelihood function $p(\{Z_t\}/\{X\}, (i, j))$ and the Bayes normalizing factor $\int p(\{Z_t\}/\{X'\}, (i', j'))p^-(\{X'\}, (i', j')/\mathcal{F}^{t-1})d\{X'\}$ are defined in Eqs. 11 and 13 below. Using the prediction

equation 4, it follows that:

$$\underbrace{p(\{X\}, (i, j)/\{Z_t\}, \mathcal{F}^{t-1})}_{p(\{X\}, (i, j)/\mathcal{F}^t)} = \\ \eta p(\{Z_t\}/X, (i, j)) p^-(\{X\}/(i, j), \mathcal{F}^{t-1}) p_{ij} w_i.$$
(8)

We may then factor the above equation as follows:

$$p(\{X\}, (i, j) / \mathcal{F}^{t}) = \frac{p(\{Z_{t}\} / \{X\}, (i, j)) p^{-}(\{X\} / (i, j), \mathcal{F}^{t-1})}{l_{ij}} \frac{l_{ij} \widetilde{p_{ij} w_{i}}}{\sum_{i', j'} l_{i', j'} \underbrace{p_{i'j'} w_{i'}}_{w_{i'j'}}}$$
(9)

where

$$l_{ij} = \int p(\{Z_t\}/\{X'\}, (i,j))p^-(\{X'\}/(i,j), \mathcal{F}^{t-1})d\{X'\}.$$
(10)

Note that l_{ij} is likelihood of the data $\{Z_t\}$ given the multiobject pdf underlying hypothesis H_{ij} , and the particular data association that is encoded in the hypothesis.

Remark 2: It behooves us to understand the updated pdf underlying the child hypothesis H_{ij} , the first factor on the right hand side of Eq. 9. Let p_D denote the probability of detection of a object given that it is in the field of view (FOV) of the monitoring sensor(s). Let $p_F(z)$ denote the probability that the observation z arises from a clutter source. Let H_i denote an M-object hypothesis with object states $\{X\} = \{X_1, \dots, X_M\}$ governed by the pdfs $p_1(x_1), \dots, p_M(x_M)$. Let the child hypothesis H_{ij} correspond to the following data association hypothesis: $z_1 \to X_{j_1}, \dots, z_m \to X_{j_m}$. Then, we define the MT-likelihood function:

$$p(\{Z_t\}/\{X\}, (i, j)) \equiv p(\{z_1 \cdots z_m\}/\{X_1 = x_1, \cdots X_M = x_M\}, (i, j))$$
$$= [\prod_{k=1}^m p_D p(z_k/X_{j_k} = x_{j_k})](1 - p_D)^{M-m}, \quad (11)$$

where $p(z_k/X_{j_k} = x_{j_k})$ is simply the single object observation likelihood function for the sensor. Thus,

$$p(\{Z_t\}/\{X\}, (i, j))p^-(\{X\}/(i, j), \mathcal{F}^{t-1}) = [\prod_{k=1}^m p_D p(z_k/X_{j_k} = x_{j_k})p^-_{j_k}(x_{j_k})][\prod_{l \neq j_k} (1 - p_D)p^-_l(x_l)],$$
(12)

where $l \neq j_k$ denotes all objects X_l that are not associated with a measurement under hypothesis H_{ij} . Further, defining the MT-Bayes factor as:

$$l_{ij} = \int p(\{Z_t\}/\{X'\}, (i, j))p^-(\{X'\}/(i, j), \mathcal{F}^{t-1})d\{X'\}$$

$$\equiv \int [\prod_{k=1}^m p_D p(z_k/X_{j_k} = x'_{j_k})p^-_{j_k}(x'_{j_k})] \times [\prod_{l \neq j_k} (1 - p_D)p^-_l(x'_l)]dx'_1..dx'_M$$

$$= [\prod_{k=1}^m (p_D \int p(z_k/X_{j_j} = x'_{j_k})p^-_{j_k}(x'_{j_k})dx'_{j_k})] \times [\prod_{l \neq j_k} (1 - p_D) \int p^-_l(x'_l)dx'_l]$$

$$= (1 - p_D)^{M-m} \prod_{k=1}^m p_D p(z_k/X_{j_k}), \qquad (13)$$

where $p(z_k/X_{j_k}) \equiv \int p(z_k/X_{j_k} = x'_{j_k})p_{j_k}^-(x'_{j_k})dx'_{j_k}$. Hence,

$$\frac{p(\{Z_t\}/\{X\},(i,j))p^{-}(\{X\}/(i,j),\mathcal{F}^{t-1})}{\int p(\{Z_t\}/\{X'\},(i,j))p^{-}(\{X'\}/(i,j),\mathcal{F}^{t-1})d\{X'\}} \\
= \frac{\prod_{k=1}^{m} p_{D}p(z_k/X_{j_k} = x_{j_k})p_{j_k}^{-}(x_{j_k})}{\prod_{k=1}^{m} p_{D}\int p(z_k/X_{j_k} = x'_{j_k})p_{j_k}^{-}(x'_{j_k})dx'_{j_k}} \times \frac{(1-p_D)^{M-m}\prod_{l\neq j_k} p_l^{-}(x_l)}{(1-p_D)^{M-m}} \\
= \prod_{k=1}^{m} p_{j_k}(x_{j_k}/z_k) \times \prod_{l\neq j_k} p_l^{-}(x_l), \quad (14)$$

where $p_{j_k}(x_{j_k}/z_k)$ denotes the updated object pdf of X_{j_k} using the observation z_k and the predicted prior pdf $p_{j_k}^-(x_{j_k})$, and $p_l^-(x_l)$ is the predicted prior pdf of X_l whenever $l \neq j_k$, i.e., the pdf of object X_l is not updated with any measurement. In the above, we have assumed that all the measurements are assigned to objects, however, some of the measurements can also be assigned to clutter, in which case, the object pdfs are updated exactly as above, i.e., all objects' predicted prior pdfs associated with data are updated while the unassociated objects' predicted priors are not updated, except now the likelihoods l_{ij} of the children hypothesis H_{ij} are given by:

$$l_{ij} = (1 - p_D)^{M - m'} \prod_{i=1}^m p(z_i/X_{j_i}),$$
(15)

where

$$p(z_i/X_{j_i}) = \begin{cases} p_D \int p(z_i/x) p_{j_i}(x) dx & \text{if } X_{j_i} \in \mathcal{T} \\ p_F(z_i) & \text{if } X_{j_i} \in \mathcal{C} \end{cases}$$
(16)

where \mathcal{T} is the set of all objects and \mathcal{C} is clutter, m' is the number of objects associated to measurements, and the above equation implies that the measurement z_i was associated to clutter if $X_{j_i} \in \mathcal{C}$. Note that in the above equation the FOV is assumed to cover the entire set of objects, if it does not do so, the factor $(1-p_D)^{M-m}$ is replaced by $(1-p_D)^{M_t-m}$ where M_t is the number of objects in the FOV of the sensor.

Remark 3: The recursive equation 9 above has a particularly nice factored hybrid form. The first factor is just a continuous multi-object pdf that is obtained by updating the predicted multi-object pdf obtained by associating the measurements in $\{Z_t\}$ to objects according to the data association underlying H_{ij} . The second factor corresponds to the update of the discrete hypothesis weights.

Remark 4: Given that there is an efficient way to predict/ update the multi-object pdfs underlying the different hypotheses, Eq. 9 actually shows that the FISST recursions may essentially be treated as a purely discrete problem living in the "Hypothesis level" space. The "hypothesis level" weights are updated based on the likelihoods l_{ij} which is determined by the continuous pdf underlying H_{ij} . Also, the continuous pdf prediction and updates are independent of the hypothesis level prediction and updates, i.e, the hypothesis probabilities do no affect the multi-object pdfs underlying the hypotheses. Thus, given that the likelihoods of different hypothesis l_{ij} arise from the underlying multi-object pdf and the encoded data association in the hypothesis H_{ij} , the FISST updates can be written purely at the hypothesis level as follows:

$$w_{ij} := \frac{l_{ij}w_{ij}}{\sum_{i',j'} l_{i'j'}w_{i'j'}},\tag{17}$$

where $w_{ij} = p_{ij}w_i$. Thus, we can see that the FISST update has a particularly simple Bayesian recursive form when viewed at the discrete hypothesis level, given that the multi-object pdfs underlying the hypotheses H_{ij} are tracked using some suitable method. We can summarize the above development of the Bayesian recursion for multi-object tracking as follows:

Proposition 1: Given an M-object hypothesis H_i , and its children hypotheses H_{ij} , that correspond to the data associations $\{z_i \rightarrow X_{j_i}\}$, the joint MT-density, hypothesis weight

update equation is:

$$p(\{X\}, (i, j) / \mathcal{F}^t) = p(\{X\} / (i, j), \mathcal{F}^t) \frac{w_{ij} l_{ij}}{\sum_{i', j'} w_{i'j'} l_{i'j'}},$$

where $w_{ij} = p_{ij}w_i$, l_{ij} is given by Eq. 15, and the MT-pdf underlying H_{ij} :

$$p(\{X\}/(i,j),\mathcal{F}^t) = \prod_{k=1}^m p_{j_k}(x_{j_k}/z_k) \prod_{l \neq j_k} p_l^-(x_l)$$

where $p_{j_k}(X_{j_k}/z_{j_k})$ denotes the predicted prior of object X_{j_k} , $p_{j_k}^-(x_k)$, updated by the observation z_{j_k} , and $p_l^-(x_l)$ is the predicted prior for all objects X_l that are not associated.

We may renumber our hypothesis H_{ij} into a parent of the next generation of hypothesis through a suitable map F((i, j)) that maps every pair (i, j) into a unique positive integer i', and start the recursive procedure again. However, the trouble is that the number of hypotheses grows combinatorially at every time step since at every step the number of hypotheses grow by the factor A_M (Eq. 3), and thus, the above recursions can quickly get intractable.

B. Incorporating Birth and Death in Hypothesis level FISST

The development thus far in this section has assumed (implicitly) that there are a fixed and known number of objects. However, this is not necessarily true since new objects can arrive while old objects can die. Thus, we have to incorporate the possibility of the birth and death of objects. In the following, we show that this can be done in quite a straightforward fashion using Eqs. 4, 9 and 17.

Let α denote the birth probability of a new object being spawned and β denote the probability that an object dies in between two measurements. We will assume that $\alpha^2, \beta^2 \approx 0$. This assumption implies that exactly one birth or one death is possible in between measurement updates. Consider the time instant t, and consider an M-object hypothesis at time t, H_i . Depending on the time t, let us assume that there can be M_t^b birth hypotheses and M_t^d death hypothesis corresponding to one of M_t^b objects being spawned or one of M_t^d objects dying. In particular, for the SSA problem, we can divide the FOV of the sensor into M_t^b parts and the births correspond to a new object being spawned in one of these FOV parts. The death hypotheses correspond to one of the M_t^d objects expected to be in the FOV dying. Hence, a child hypothesis H_{ij} of the parent H_i can be an M+1 object hypothesis with probability α in exactly M_t^b different ways. The child H_{ij} could have M-1 objects with probability β each in M_t^d different ways corresponding to the M_t^d different objects dying. Thus, the child H_{ij} could have M objects with probability $(1 - M_t^b \alpha M_t^d\beta$) in exactly one way (the no birth/ death case). Please see Fig. 1 for an illustration of the process.

Further, the child hypothesis H_{ij} can then split into further children H_{ijk} where the total number of children is A_M , A_{M+1} or A_{M-1} depending on the number of objects underlying the hypothesis H_{ij} , and corresponding to the various different data associations possible given the measurement $\{Z_t\}$. Note that the above process degenerates into the no birth and death case when $\alpha = \beta = 0$. Thus, we can see that the primary consequence of the birth and death process is the increase in the total number of children hypotheses. However, the equations for the multi-object filtering (with a little effort, due to the fact that the child hypotheses may have different number of objects than the parent hypothesis thereby complicating the integration underlying the prediction step) can be shown to remain unchanged. Recall Eq. 9, which is reproduced below for clarity:

$$p(\{X\}, (i, j)/\mathcal{F}^{t}) = \underbrace{p(\{Z_{t}\}/\{X\}, (i, j))p^{-}(\{X\}/(i, j), \mathcal{F}^{t-1})}_{\substack{\int p(\{Z_{t}\}/\{X'\}, (i, j))p^{-}(\{X'\}/(i, j), \mathcal{F}^{t-1})d\{X'\}}_{\text{updated pdf underlying}_{ij}} \times \underbrace{\frac{l_{ij} p_{ij} w_{i}}{\sum_{i', j'} l_{i', j'} p_{i'j'} w_{i'}}}_{\substack{w_{i', i'}}}.$$
(18)

The only difference from the no birth and death case is, given H_i is an M- object hypotheses, the children hypotheses H_{ij} can have M, M-1 or M+1 objects underlying them, and the corresponding p_{ij} value is $1 - M_t^b \alpha - M_t^d \beta$, β or α respectively. It behaves us to look closer at the prediction equations in the birth and death case as that is the source of difference from the no birth and death case.

First, consider the case of a death hypothesis. Consider an M-object hypothesis, H_i , with underlying MT-pdf $\prod_k p_k(x_k)$. Suppose without loss of generality that the M^{th} object dies. Then, the transition density for the multi-object system is defined as:

$$p(\{X\}/\{X'\},(i,j)) = \left[\prod_{k=1}^{M-1} p_k(x_k/x'_k)\right] \delta(\phi/x_M), \quad (19)$$

where $\delta(\phi/x_M)$ denotes the fact that the M^{th} object becomes the null object ϕ with probability one. Thus, the predicted MT-transition density underlying H_{ij} is:

$$p^{-}(\{X\}/(i,j),\mathcal{F}^{t}) = \int (\prod_{k=1}^{M-1} p(x_{k}/x_{k}')p(x_{k}'/i,\mathcal{F}^{t}))\delta(\phi/x_{M}')dx_{1}'..dx_{M}'$$
$$= \prod_{k=1}^{M-1} p^{-}(x_{k}/i,\mathcal{F}^{t}),$$
(20)

i.e., the predicted MT-pdf is simply the predicted pdfs of all the objects that do not die.

Next, consider the case of a birth hypothesis H_{ij} where the birthed pdf has a distribution $p_b^l(x_{M+1})$. The transition pdf is now

$$p(\{X\}/\{X'\},(i,j)) = \left[\prod_{k=1}^{M} p_k(x_k/x'_k)\right] p_{M+1}(x_{M+1}/\phi),$$
(21)

where $p_{M+1}(x_{M+1}/\phi) = p_b^l(x_{M+1})$ denotes that the null object ϕ spawns an $M + 1^{th}$ object with underlying pdf $p_b^l(x_M)$. It can be shown similar to above that the predicted distribution in this case is:

$$p^{-}(\{X\}/(i,j),\mathcal{F}^{t}) = [\prod_{k=1}^{M} p_{k}^{-}(x_{k}/i,\mathcal{F}^{t})]p_{b}^{l}(x_{M+1}), \quad (22)$$

i.e., the predicted distribution of all the objects with the addition of the birth pdf $p_b^l(x_{M+1})$.

Further, each of these hypothesis split into children H_{ijk} based on the possible data associations: if H_{ij} is a birth hypothesis the number of children is A_{M+1} , if its a death hypothesis the number of children is A_{M-1} and if it is no birth or death, the number of children is A_M . In particular, using the development outlined above (where we have replaced the child notation H_{ijk} by H_{ij} for simplicity), we can see that the transition probability p_{ij} of a child hypothesis H_{ij} is:

$$p_{ij} = \begin{cases} \frac{\alpha}{A_{M+1}}, & \text{if } j \in B_{M+1} \\ \frac{1-M_t^b \alpha - M_t^d \beta}{A_M}, & \text{if } j \in B_M \\ \frac{\beta}{A_{M-1}}, & \text{if } j \in B_{M-1} \end{cases}$$
(23)

where B_M refers to the set of all M object hypothesis, and recall that $A_M = \sum_{k}^{\min(m,M)} {m \choose k} {M \choose k} k!$.



Fig. 1. A schematic of the splitting of the hypothesis due to birth/ death of objects and data associations. Underlying each blob is a continuous MT-pdf.

The above development can be summarized as the following result:

Proposition 2: Given an M-object hypothesis H_i and its children H_{ij} , the update equation for joint MT-pdf-hypothesis density function is given by Eq. 18, where the only differences from the no birth or death case is that p_{ij} in the equations is different according as the hypothesis H_{ij} birth, death or a no birth or death hypothesis and is given by Eq. 23, and the predicted priors required in Eq. 18 is calculated from Eq. 20

if H_{ij} is a death hypothesis, Eq. 22 if it is a birth hypothesis and Eq 6 if it is a no birth or death hypothesis.

III. A RANDOMIZED FISST (R-FISST) TECHNIQUE

In the previous section, we have introduced the hypothesis level FISST equations and shown that they are particularly easy to comprehend and implement. However, the number of children hypothesis increase exponentially at every iteration and thus, can get computationally intractable very quickly. However, it can also be seen that most children hypotheses are very unlikely and thus, there is a need for intelligently sampling the children hypotheses such that only the highly likely hypotheses remain. In the following, we propose an MCMC based sampling scheme that allows us to choose the highly likely hypotheses.

A. MCMC based Intelligent Sampling of Children Hypothesis

Recall Eq. 17. It is practically plausible that most children j of hypothesis H_i are highly unlikely, i.e., $l_{ij} \approx 0$ and thus, $w_{ij} \approx 0$. Hence, there is a need to sample the children H_{ij} of hypothesis H_i such that only the highly likely hypotheses are sampled, i.e., $l_{ij} >> 0$.

Remark 5: Searching through the space of all possibly hypotheses quickly becomes intractable as the number of objects and measurements increase, and as time increases.

Remark 6: We cannot sample the hypothesis naively either, for instance, according to a uniform distribution since the highly likely hypothesis are very rare under the uniform distribution, and thus, our probability of sampling a likely hypothesis is vanishingly small under a uniform sampling distribution.

Thus, we have to resort to an intelligent sampling technique, in particular, an MCMC based approach.

Given a hypothesis H_i , we want to sample its children according to the probabilities $\bar{p}_{ij} = w_{ij}l_{ij}$. This can be done by generating an MCMC simulation where the sampling Markov chain, after enough time has passed (the burn in period), will sample the children hypotheses according to the probabilities \bar{p}_{ij} . A pseudo-code for setting up such an MCMC simulation is shown in Algorithm 1. In the limit, as $k \to \infty$, the sequence

Algorithm 1 MCMC Hypothesis Sampling

Generate child hypothesis j_0 , set k = 0.

Generate $j_{k+1} = \pi(j_k)$ where $\pi(.)$ is a symmetric proposal distribution

If $\bar{p}_{ij_{k+1}} > \bar{p}_{ij_k}$ then $j_k := j_{k+1}; k := k+1;$

else $j_k := j_{k+1}$ with probability proportional to $\frac{\bar{p}_{ij_{k+1}}}{\bar{p}_{ij_k}}$; k = k+1.

 $\{j_k\}$ generated by the MCMC procedure above would sample the children hypotheses according to the probabilities \bar{p}_{ij} . Suppose that we generate *C* highest likely distinct children hypothesis H_{ij} using the MCMC procedure, then the FISST recursion Eq. 17 reduces to:

$$w_{ij} := \frac{l_{ij}w_{ij}}{\sum_{i',j'} l_{i'j'}w_{i'j'}},$$
(24)

where i' and j' now vary from 1 to C for every hypothesis H_i , instead of the combinatorial number A_M .

Given these M * C hypotheses, i.e. C children of M parents, we can keep a fixed number H_{∞} at every generation by either sampling the H_{∞} highest weighted hypotheses among the children, or randomly sampling H_{∞} hypotheses from all the children hypotheses according to the probabilities w_{ij} .

Remark 7: The search for the highly likely hypotheses among a very (combinatorially) large number of options is a combinatorial search problem for which MCMC methods are particularly well suited. Thus, it is only natural that we use MCMC to search through the children hypotheses.

Remark 8: The choice of the proposal distribution $\pi(.)$ is key to the practical success of the randomized sampling scheme. Thus, an intelligent proposal choice is required for reducing the search space of the MCMC algorithm. We show such an intelligent choice for the proposal in the next section.

Remark 9: The discrete hypothesis level update Eq. 17 is key to formulating the MCMC based sampling scheme, and, hence, the computational efficiency of the R-FISST algorithm.

B. Smart Sampling Markov Chain Monte Carlo

In the previous section, MCMC was proposed to intelligently sample the children hypotheses based on their likelihood. However, as the number of possible hypotheses grows so does the time it takes for the MCMC to converge to the stationary distribution. Also having to enumerate large numbers of possible hypotheses is computationally strenuous. Therefore, we propose a new method that allows for faster time to convergence and alleviates the burden of having to enumerate large numbers of hypotheses. This method is called Smart Sampling Markov Chain Monte Carlo, SSMCMC.

SSMCMC utilizes a two level Markov Chain Monte Carlo algorithm. The first level establishes a Markov Chain whose states are the possible data associations for a single object. The chain is then run using a uniform proposal distribution. The criteria for the MC to move to a new state stems from the Metropolis-Hastings condition: $U[0,1] < min(1, \frac{p(z_{i_{k+1}}/x)}{p(z_{i_k}/x)}),$ where $p(z_{i_k}/x)$ is the likelihood of the data association at step k. The stationary distribution of this Markov Chain represents the highly likely object to measurement data associations. This process is repeated for all objects within the gating parameters, i.e., objects that are expected to be in the FOV. The children hypotheses are then formed using only these highly likely data associations. These children hypotheses have very high weights and are a lot fewer in number than the total amount of possible children hypotheses. These children hypotheses are then used as the states of the chain for a second MCMC. They are again explored using a uniform proposal distribution. The criteria on which a new state is assumed for this MCMC is the Metropolis-Hastings condition: $U[0,1] < min(1, \frac{l_{(i,j)_{k+1}}}{l_{(i,j)_k}})$, where $l_{(i,j)_k}$ is the likelihood of the child hypothesis at step k. The stationary distribution of this MC is then guaranteed to sample the most highly likely data associations.

IV. APPLICATIONS

This section illustrates the application of the results from the previous sections. We illustrate the R-FISST based approach to the multi-object tracking and detection problem inherent in SSA applications. In particular, we will discuss the results from a fifty-space object birth and death scenario. Our goal is to show that the aforementioned methodology allows for accurate estimation while determining the correct number of objects in an environment where the number of objects is not fixed. This will allow for the methodology to be used in both catalog update and catalog maintenance.

A. R-FISST Application to a Fifty Object Birth and Death Scenario

In order to test the methods discussed in this paper a fiftyspace object tracking and detection problem was simulated using a planar orbit model. These fifty objects were in orbits ranging from LEO to MEO and had varying orbital properties as well as zero-mean Gaussian process noise appropriate for SSA models. The objects were simulated for at least one orbital period. That being said each object was allowed to pass completely through the field of view at least one time. The objective was to accurately track all objects given only an imperfect initial hypothesis containing some of their means and covariances. Also, to simulate a birth and death environment, the correct number of objects is initially unknown to the algorithm. In this particular example, the initial hypothesis only contains information on forty five of the fifty objects. The five left over will be seen as objects that are randomly introduced to the environment or simply "births". This is often described as the launch of a new satellite into orbit and is not to be confused with "spawns" in which an object currently in orbit divides into two or more pieces. Spawns can be accounted for by our methodology but are not explicitly programed in this example. The R-FISST methodology must recognize the five births and provide accurate estimations of all of the objects' states and covariances. State vectors for this particular problem contain the objects' position along the xand y axes as well as the magnitude of their velocity in the x and y directions. A single noisy sensor was positioned at a fixed look direction of 15 degrees above the positive x-axis with a field of view of 30 degrees. The sensor was used to measure objects' position in the x - y plane with zero-mean Gaussian measurement noise appropriate for the application.

Figure 2 shows snapshots of the hypotheses' weights throughout the simulation. The snapshots were taken at ten, fifty, seventy-five, and one hundred percent of the total simulation time. From this figure, one can see that, in the beginning, the initial hypothesis caries all the weight. How-ever, throughout the simulation the number of maintained hypotheses (shown on the x-axis of the graphs in Figure 2) varies as does the weights of those hypotheses. The number of hypotheses maintained has a direct correlation to the number of recent ambiguities in the field of view. Ambiguities occur when one or more measurement returns can be associated to multiple objects in the field of view.







(c) Hyptheses' weights at 75 percent (d) Hyptheses' weights at 100 percompletion cent completion

Fig. 2. Snapshots of the hypotheses' weights throughout the simulation

Before discussing the state estimations, figure 3 shows an example of how the estimation data is to be presented. The figure shows the actual positions of the objects labeled "Current Position" and the estimated positions from two separate hypotheses. If the estimated position for a particular object is within an error bound then the object's position will be represented by a green circle otherwise the object's position will be represented by a red star.



Fig. 3. An example of how estimation data is presented throughout the paper. Correct estimates are represented by green circles while poor estimates are represented by red stars. All axes are in tens of thousands of kilometers

In figure 4, the snapshots show the actual positions (blue) against the estimated position from the top hypotheses (green). The black lines bound the field of view. These snapshots were taken at the same time intervals as in figure 2 and thus the estimates throughout figure 4 are taken from the hypotheses with the highest weights in figure 2. Notice in figure 4 there are no instances of red stars. This is particularly important because it shows that the hypotheses accurately estimated object positions throughout the simulation. Hence, the R-FISST approach accurately tracked and detected the fifty objects.

Lastly, figure 5 shows a visual representation of how weight shifts from the forty five object assumption to the fifty object assumption. The x-axis represents the simulation time in percent completed. The y-axis represents the expected number of objects. The expected number of objects is found by summing the weights of all hypotheses containing the same number of objects. The magnitudes of these summations are then compared to determine the expected number of objects.



(a) Estimation from the top hypotheses at 10 percent completion (b) Estimation from the top hypotheses at 50 percent completion



(c) Estimation from the top hypotheses at 75 percent completion (d) Estimation from the top hypotheses at 100 percent completion

Fig. 4. Snapshots of the actual states (blue) and the estimated states from the top hypotheses (green) throughout the simulation. Axes in tens of thousands of kilometers



Fig. 5. The expected number of objects throughout the simulation. The graph shows intermediate values during the transition periods. Over time it can be seen that weight shifts from the forty five object hypotheses to the fifty Object hypotheses

It is important to note that weight seems to be handed off in a single file fashion until the fifty object assumption accumulates all the weight toward the end of the simulation.

V. CONCLUSION

In this paper, we have presented an alternate hypothesis based derivation of the FISST recursions for recursive multiobject tracking. We have also introduced a randomized version of the FISST recursion, called R-FISST, which provides a computationally efficient randomized solution to the full FISST recursions via an MCMC sampling of likely children hypotheses. We have also shown the application of the R-FISST technique on a fifty object SSA problem as well as problems with birth and death. It is our belief that the randomized technique is extremely parallellizable, and thus, our next step will be to look at large scale GPU based implementation of the method that can scale to realistic SSA scenarios. We also intend to look at the integration of sensor tasking with the R-FISST technique such that the ambiguities inherent in the problem can be minimized. Also, comparisons and research on implementing the randomized technique with a HOMHT framework as well as a more detailed discussion of the randomized technique shall be presented in the proceedings of the 2015 AAS/AIAA Astrodynamics Specialist Conference to be held at Vail, CO, August 9-13 of 2015 [19].

VI. ACKNOWLEDGMENTS

This work is funded by AFOSR grant number: FA9550-13-1-0074 under the Dynamic Data Driven Application Systems (DDDAS) program.

REFERENCES

- I. R. Goodman, R. P. S. Mahler, and H. T. Nguyen, *Mathematics of Data Fusion*. Kluwer Academic Publishers, 1997.
- [2] R. P. S. Mahler, Statistical Multisource-Multitarget Information Fusion. Artec House, 2007.
- [3] B. Vo and W. Ma, "The gaussian mixture probability hypothesis density filter," *IEEE TRANSACTIONS ON SIGNAL PROCESSING*, vol. 54, no. 11, November 2006.
- [4] B. Vo et al., "Sequential monte carlo methods for multi target filtering with random finite sets," *IEEE Tr. Aerosp. Electronic Systems*, vol. 4, pp. 1224–1245, 2005.
- [5] B. T. Vo et al., "An analytic implementation of the cardinalized probability hypothesis density filter," *IEEE Tr. Signal Processing*, vol. 55, pp. 3553–3567, 2007.
- [6] B. Ristic *et al.*, "Improved smc implementation of the phd filter," in *PRoc. Conf. Info. fusion (FUSION)*, 2010.
- [7] I. Hussein, K. DeMars, C. Frueh, R. S. Erwin, and M. Jah, "An aegisfisst algorithm for joint detection and tracking in space situational awareness," AAS/AIAA Astrodynamics Specialist Conference, August 2012.
- [8] ——, "An aegis-fisst integrated detection and tracking approach to space situational awareness," *International Conference of Information Fusion*, July 2012.
- [9] R. P. S. Mahler, "Multitarget markov motion models," SPIE Conference on Signal Processing, Sensor Fusion, and Target Recognition, vol. 3720, pp. 47–56, 1999.
- [10] D. B. Reid, "An algorithm for tracking multiple targets," *IEEE Trans*actions on Automatic Control, vol. 24, pp. 843–854, 1979.
- [11] Y. B. Shalom and X. Li, Multitarget-Multisensor Tracking: Principles and Techniques, New York, 1995.
- [12] Y. Blair and E. W. D. Blair, Multitarget Multisensor tracking: Applications and advances, vol. III. Norwood, MA: Artech House, 2000.
- [13] S. Oh et al., "Markov chain monte carlo data association for multi-target tracking," IEEE Tr. Automatic Control, vol. 54, pp. 481–497, 2009.
- [14] S. S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*. Norwood, MA: Artech House, 1999.
- [15] Y. B. Shalom and T. Fortmann, *Tracking and Data Association*. San Diego, CA: Academic Press, 1988.
- [16] N. Bergman and A. Doucet, "Markov chain monte carlo data association for target trading," in *Brog. ICASSB*, 2000
- for target tracking," in *Proc. ICASSP*, 2000.
 [17] H. Pasula *et al.*, "Tracking many objects with many sensors," in *Proc. IJCAI*, 1999.
- [18] S. Cong *et al.*, "Markov chain monte carlo approach fro association probability evaluation," *Proc. IEE*, vol. 151, pp. 185–193, 2004.
- [19] W. Faber, S. Chakravorty, and I. Hussein, "A randomized sampling based approach to multi-object tracking with comparison to homht," AAS/AIAA Astrodynamics Specialist Conference, August 2015 Vail, Colorado.