

A RANDOMIZED SAMPLING BASED APPROACH TO MULTI-OBJECT TRACKING WITH COMPARISON TO HOMHT

W. Faber^{*}, S. Chakravorty[†], and Islam I. Hussein[‡]

In this paper, we present a comparison between our recently published randomized version of the finite set statistics (FISST) Bayesian recursions for multi-object tracking with the commonly known Hypothesis Oriented Multiple Hypothesis Tracking (HOMHT) method. We start by revisiting our hypothesis level derivation of the FISST equations in order to appropriately introduce our randomized method, termed randomized FISST (RFISST). In this randomized method, we forgo the burden of having to exhaustively generate all possible data association hypotheses by implementing a Markov Chain Monte Carlo (MCMC) approach. This allows us to keep the problem computationally tractable. We illustrate the comparison by applying both methods to a space situational awareness (SSA) problem and show that as the number of objects and/or measurement returns increases, as does the computational burden. We then show that the RFISST methodology allows for accurate tracking information far beyond the limitations of HOMHT.

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 SSA tracking and detection problems and compared to the well-known method HOMHT.

In the last 20 years, the theory of FISST-based multi-object 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

^{*}Department of Aerospace Engineering, Texas A&M University, College Station, TX.

[†]Department of Aerospace Engineering, Texas A&M University, College Station, TX.

[‡]Applied Defense Solutions, Columbia, MD

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]. In previous work, a GM approximation was applied to the original full propagation and update equations derived from FISST [7, 8].

In this paper, in contrast, we introduce a hypothesis level derivation of the FISST equations that shows how the full FISST recursions can be implemented. In order to ensure the computational tractability of the resulting equations, we introduce an MCMC based hypothesis selection scheme resulting in the Randomized FISST (RFISST) approach that is able to scale the FISST recursions to large scale problems.

There are also non-FISST based approaches to multi-hypothesis tracking (MHT) such as the Hypothesis Oriented MHT (HOMHT) [9–12], and the track oriented MHT (TOMHT) techniques [13]. 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 [10, 12, 14]. The single-scan (recursive) methods such as joint probabilistic data association (JPDA) [12, 14] typically make the assumption that the tracks are independent which is not necessarily true. The multi-scan methods such as TOMHT [12, 13] 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 [15–17], 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 RFISST algorithm. In Section IV, we show an application of the RFISST technique on SSA examples and compare the results to the HOMHT technique. A related paper, [18], was presented at the International Conference of Information Fusion. This paper extends that paper with a new MCMC data association scheme and a full comparison of the methodology with the HOMHT technique. For the sake of the paper being self-contained, the next two sections detailing the hypothesis level derivation of the FISST equations are reproduced from reference [18].

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 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.

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 \mathbb{R}^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 $\{X\} = \{X_1, X_2, \dots, X_M\}$ where X_j denotes the state of the j^{th} object under hypothesis H_i and resides in \mathbb{R}^N . Let $p(\{X\}, i/\mathcal{F}^{t-1})$ denote the joint distribution of the state-hypothesis pair after time $t - 1$. Using the rule of conditional probability:

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

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.,

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

where $p_k(\cdot)$ 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

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

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:

$$\begin{aligned} p^-(\{X\}, (i, j)/\mathcal{F}^{t-1}) &= \\ &= \int p(\{X\}, (i, j)/\{X'\}, i) p(\{X'\}, i/\mathcal{F}^{t-1}) d\{X'\} = \\ &= \underbrace{\int p(\{X\}/(i, j), \{X'\}) p(\{X'\}/i, \mathcal{F}^{t-1}) d\{X'\}}_{p^-(\{X\}/(i, j), \mathcal{F}^{t-1})} \\ &= \underbrace{p(j/i)}_{p_{ij}} \underbrace{p(i/\mathcal{F}^{t-1})}_{w_i}, \end{aligned} \quad (4)$$

where $p^-(\cdot, (i, j)/\mathcal{F}^{t-1})$ 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

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

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:

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

where

$$\begin{aligned} \eta &= \\ \sum_{i', j'} \int p(\{Z_t\}/\{X'\}, (i', j'))p^-(\{X'\}, (i', j')/\mathcal{F}^{t-1})d\{X'\}, \end{aligned} \quad (7)$$

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:

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

We may then factor the above equation as follows:

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

where

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

Note that l_{ij} is likelihood of the data $\{Z_t\}$ given the multi-object 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: $z_1 \rightarrow X_{j_1}, \dots, z_m \rightarrow X_{j_m}$. Then, we define the MT-likelihood function:*

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

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

$$\begin{aligned} p(\{Z_t\} / \{X\}, (i, j)) p^-(\{X\} / (i, j), \mathcal{F}^{t-1}) &= \\ [\prod_{k=1}^m p_D p(z_k / X_{j_k} = x_{j_k}) p_{j_k}^-(x_{j_k})] [\prod_{l \neq j_k} (1 - p_D) p_l^-(x_l)], & \end{aligned} \quad (12)$$

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:

$$\begin{aligned}
l_{ij} &= \int p(\{Z_t\}/\{X'\}, (i, j)) p^-(\{X'\}/(i, j), \mathcal{F}^{t-1}) d\{X'\} \\
&\equiv \int \left[\prod_{k=1}^m p_D p(z_k/X_{j_k} = x'_{j_k}) p_{j_k}^-(x'_{j_k}) \right] \times \\
&\quad \left[\prod_{l \neq j_k} (1 - p_D) p_l^-(x'_l) \right] dx'_1 \dots dx'_M \\
&= \left[\prod_{k=1}^m (p_D \int p(z_k/X_{j_k} = x'_{j_k}) p_{j_k}^-(x'_{j_k}) dx'_{j_k}) \right] \times \\
&\quad \left[\prod_{l \neq j_k} (1 - p_D) \int p_l^-(x'_l) dx'_l \right] \\
&= (1 - p_D)^{M-m} \prod_{k=1}^m p_D p(z_k/X_{j_k}), \tag{13}
\end{aligned}$$

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,

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

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:

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

where

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

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 not 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 hypotheses H_{ij} , the FISST updates can be written purely at the hypothesis level as follows:

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

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.

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:

$$\begin{aligned}
 & p(\{X\}, (i, j) / \mathcal{F}^t) \\
 &= \underbrace{\frac{p(\{Z_t\} / \{X\}, (i, j)) p^-(\{X\} / (i, j), \mathcal{F}^{t-1})}{\int p(\{Z_t\} / \{X'\}, (i, j)) p^-(\{X'\} / (i, j), \mathcal{F}^{t-1}) d\{X'\}}}_{\text{updated pdf underlying } H_{ij}} \times \\
 & \quad \frac{\overbrace{l_{ij} \hat{p}_{ij} w_i}^{w_{ij}}}{\sum_{i', j'} \underbrace{l_{i', j'} \hat{p}_{i', j'} w_{i'}}_{w_{i', j'}}}. \tag{18}
 \end{aligned}$$

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 $1 - M_t^b\alpha - M_t^d\beta, \beta$ or α respectively. It behooves 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:

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

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:

$$\begin{aligned} p^-(\{X\}/(i, j), \mathcal{F}^t) &= \\ &= \int (\prod_{k=1}^{M-1} p(x_k/x'_k)p(x'_k/i, \mathcal{F}^t))\delta(\phi/x'_M)dx'_1..dx'_M \\ &= \prod_{k=1}^{M-1} p^-(x_k/i, \mathcal{F}^t), \end{aligned} \quad (20)$$

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

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

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 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:

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

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

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*




Figure 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.

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.

A RANDOMIZED FISST (RFISST) 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.

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., $l_{ij} \gg 0$.

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 $k \rightarrow \infty$, the sequence $\{j_k\}$ generated by the MCMC procedure above would sample

Algorithm 1 MCMC Hypothesis Sampling

Generate child hypothesis j_0 , set $k = 0$.

Generate $j_{k+1} = \pi(j_k)$ where $\pi(\cdot)$ 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 $\frac{\bar{p}_{ij_{k+1}}}{\bar{p}_{ij_k}}; k = k + 1$.

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:

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

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 $\pi(\cdot)$ 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 RFISST algorithm.*

Smart Sampling Markov Chain Monte Carlo

In this section, we reveal the process used to perform the MCMC sampling discussed in the previous section. This process is performed at every scan to generate the highly likely children hypotheses. Consider the following SSA scenario depicted in figure 2. In this scenario the belief is that there are ten objects in the field of view. The sensor then detects five measurement returns.




Figure 2. A possible SSA event where there exists ten objects in the field of view and five measurement returns.

Typically when generating the hypotheses exhaustively one would create a matrix where each row represents a particular hypothesis. The columns of the matrix represent the measurement returns provided by the sensor. Each column entry represents the object that measurement is being associated to in the particular hypothesis. The hypothesis matrix for our example scenario would look like figure 3. However, if all objects and measurements within the field of view can be associated

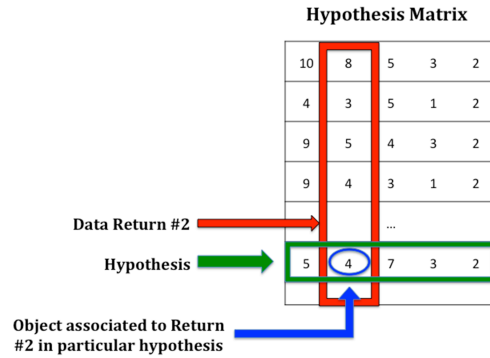


Figure 3. An example of a typical Hypotheses Matrix used when exhaustively generating hypotheses. This particular matrix represents a portion of the hypothesis matrix that would be generated for the scenario in figure 2.

then according to Eq. (3), with $m = 5$ and $M = 10$, the total number of possible hypotheses would be $A_M = 63,591$. Thus, the hypothesis matrix actually has 63,591 rows. This illustrates the importance of a randomized approach. One can see that even with relatively low numbers of objects and measurement returns exhaustively generating and maintaining the hypotheses will cause a large computational burden. In our randomized approach we sample the rows of the hypothesis matrix based on hypothesis probability. We do this by creating a matrix we call the data association

matrix, figure 4. The data association matrix lists the objects as the columns and the measurement

		Object Number + Clutter										
		1	2	3	4	5	6	7	8	9	10	C
Return Number	M1	.1	.2	0	.03	.05	.004	.002	.001	.01	.005	.02
	M2	.04	.03	.05	.01	.04	.014	.2	.004	.005	.006	.02
	M3	.09	.05	.04	.03	.01	.004	.005	.006	.1	.001	.02
	M4	.09	.04	.1	.01	.02	.04	.3	.004	.005	.006	.02
	M5	.001	.003	.002	.004	.005	.006	.007	.5	.01	.03	.02

Figure 4. The Data Association Matrix. Each row represents a particular measurement return. Each column is a particular association. The elements represent the likelihood of the corresponding association. Green boxes here show a visual representation of an example hypothesis.

return as the rows. The entries of the matrix contain the likelihood value of that particular measurement to object assignment. The last column of the matrix is dedicated to clutter and contains the likelihood that a particular measurement is associated to clutter. The dimensions of this matrix are $m \times (M + 1)$ which is much smaller than the dimensions of the hypothesis matrix. This makes it much more practical to explore using the MCMC technique.

Remark 10 *The numbering of the objects and measurement returns in the data association matrix is done strictly for organization and is redone at random each time step with no record of previous numbering or labeling kept throughout scans.*

Remark 11 *Computing the data association matrix does not add any computational burden because the object to measurement likelihood is necessary in every tracking method.*

We start the randomized technique by creating a row vector of length m containing a permutation of column numbers. The green boxes in figure 4 are a visual representation of such a row vector $[5 \ 4 \ 2 \ 1 \ 7]$. This row vector is used as our first hypothesis. We then take a step in the MCMC by generating a proposed hypothesis. This is done by randomly choosing a row (Measurement) of the data association matrix to change. We then randomly sample a column (object) to associate the measurement to. If there is a conflicting assignment (i.e. a measurement is already assigned to that object) then we automatically assign the conflicting measurement to clutter. We then compare the proposed hypothesis to the current hypothesis in an MCMC fashion using a criteria which stems from the Metropolis condition $U[0, 1] < \min(1, \frac{P_{(i,j)_{k+1}}}{P_{(i,j)_k}})$ where $P_{(i,j)_k}$ is the probability of the hypothesis at step k . In words, if the proposed hypothesis has a higher probability then we keep it, if not, we keep it with probability proportional to the ratio of the hypothesis probabilities. These steps are then repeated until assumed stationary distribution. We then continue walking for a user defined amount of steps and record all hypotheses sampled during these steps. The recorded hypotheses represent the highly likely hypotheses.




Figure 5. Visualization of a single MCMC step using the Data Association Matrix. Green boxes represent the current hypothesis while blue boxes represent the changes made for the proposed hypothesis. This particular example contains a conflicting assignment with measurement return two and shows how the association is then changed to clutter.

COMPARISONS

In this section, we show a comparison between our randomized method RFISST and a well-known tracking method called HOMHT. This comparison helps us illustrate two main points. The first being that the accuracy of the estimation provided by the RFISST method is either equal to or better than that of HOMHT but never worse. We will achieve this by showing side by side estimations from both methods. The second point is seen when the number of hypotheses rapidly escalates due to a large number of objects and/or a large number of measurement returns. Such occurrences happen often in SSA and for many reasons, for example, when a debris field crosses the sensor's field of view. In these situations HOMHT fails because it is computationally impossible to generate such a large number of hypotheses. Due to the randomized scheme, the RFISST methodology continues to perform in such scenarios.

The following figure, figure 6, 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.






Figure 6. An example of how estimation data is presented throughout the paper. Green circles represent the correct estimates while poor estimates are represented by red stars. All axes are in tens of thousands of kilometers

Comparison between RFISST and HOMHT: SSA Tracking


In order to compare both methods we simulated a fifteen-space object tracking and detection problem. Each object was given a random planar orbit ranging between LEO and MEO with unique orbital properties and zero-mean Gaussian process noise appropriate for SSA models. The objects were simulated for long enough to where the object with the largest period would be able to complete at least one orbit. Thus each object was allowed to pass completely through the field of view at least one time. In this particular simulation we initialized all orbits to begin within the field of view. In order to achieve an apples to apples comparison we used this simulation to test both methods. The goal of each method would be to accurately track each object given only an imperfect initial hypothesis containing the objects' mean and covariance as well as measurement returns from a single noisy sensor. State vectors for this problem consisted of the objects' position along the x and y axes as well as the magnitude of their velocity in the x and y directions. The 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. An Extended Kalman Filter (EKF) was used in conjunction with each method to compute the underlying state and covariance updates. That being said both methods will produce the same estimation given the correct hypotheses were generated throughout the simulation.




(a) Estimation at 10 percent completion



(b) Estimation at 50 percent completion



(c) Estimation at 75 percent completion



(d) Estimation at 100 percent completion

Figure 7. Side by side snapshots of the true positions (blue) and the estimated states of the top hypothesis from both RFISST and HOMHT (green). Snapshots are taken at ten, fifty, seventy-five, and one hundred percent simulation time. All axes in tens of thousands of kilometers.

Figures 7(a)-7(d) are snapshots of the simulation at ten, fifty, seventy-five, and one hundred percent of the simulation time. Each snapshot shows the true positions of the objects and the position estimations provided by both the HOMHT method and the RFISST method. These figures are provided to illustrate that the methods accurately track the objects. Furthermore, it shows that each method maintained the correct hypothesis throughout the simulation. If either method was unable to generate the correct hypothesis then the position estimations would be incorrect. These incorrect position estimates would be seen as red stars in the snapshots.




Figure 8. The hypotheses generation time for both methods as shown on a log y-axis scale.

Comparison between RFISST and HOMHT: Computation Time

When tracking large numbers of objects the majority of the computational burden lies in the hypothesis generation. HOMHT uses an exhaustive approach to generate the hypotheses. This exhaustive approach has its pros and cons. For example, generating all the hypotheses guarantees that the hypothesis containing the correct data associations is sampled. Also this exhaustive approach is very easy to implement. On the other hand, as the number of hypotheses grows so does the burden placed on generating them. This can be seen as an increase of computation time. Even with proper gating and pruning methods the number of hypotheses can be so large that generating them would exceed the computers memory heap space making it computationally intractable. It is at this point that we say the HOMHT method breaks. Using our randomized approach we never generate all hypotheses, which allows us to handle scenarios with very high number of possible hypotheses. However this method is more difficult to implement and must be tuned to guarantee that the hypothesis containing the correct data association is sampled. Figure 8 shows a comparison of the computation times for hypothesis generation of HOMHT and RFISST. The y-axis is the log of the computation time in seconds while the x-axis is the log of the total number of possible hypotheses being generated. HOMHT is represented by the blue line and resembles an exponential curve. RFISST is represented by the green line and resembles a linear growth. It can be seen that at first for

low numbers of possible hypotheses HOMHT performs faster. However, as the number of possible hypotheses grows into the tens of thousands the RFISST method becomes more efficient. Furthermore, as the number of possible hypotheses grows into the hundred millions HOMHT struggles to generate the hypotheses and eventually breaks. RFISST can generate the correct hypotheses even as the number of possible hypotheses grows to the order of 10^{33} . In each of these simulations the RFISST MCMC methodology was able to sample the correct hypotheses using only 100,000 steps in the MCMC.

CONCLUSION

In this paper, we have presented a comparison between our randomized sampling based approach to multi object tracking, RFISST, and the well-known method HOMHT. We introduced the RFISST method by presenting a hypothesis based derivation of the FISST recursions. Next we showed that by using MCMC we can sample the highly likely hypotheses without having to generate all possible hypotheses. We then showed that given the same underlying filter that HOMHT and RFISST produce the same accuracy of tracking. Lastly, we showed that in situations of high numbers of possible hypotheses that HOMHT breaks down while the RFISST method continues to perform. Our current research interests includes making our randomized approach adaptive and applying it to real data. This includes comparing with other methods and developing hybrid methods as well as expanding to large scale GPU based implementation. We are also looking into the integration of sensor tasking to minimize ambiguities and to maximize gain.

ACKNOWLEDGMENTS

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

REFERENCES

- [1] 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, 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, 2005, pp. 1224–1245.
- [5] B. T. Vo *et al.*, "An analytic implementation of the cardinalized probability hypothesis density filter," *IEEE Tr. Signal Processing*, Vol. 55, 2007, pp. 3553–3567.
- [6] B. Ristic *et al.*, "Improved SMC implementation of the PHD filter," *Proc. Conf. Info. fusion (FUSION)*, 2010.
- [7] I. Hussein, K. DeMars, C. Frueh, R. S. Erwin, and M. Jah, "An AEGIS-FISST Algorithm for Joint Detection and Tracking in Space Situational Awareness," *AAS/AIAA Astrodynamics Specialist Conference*, August 2012.
- [8] I. Hussein, K. DeMars, C. Frueh, R. S. Erwin, and M. Jah, "An AEGIS-FISST Integrated Detection and Tracking Approach to Space Situational Awareness," *International Conference of Information Fusion*, July 2012.
- [9] D. B. Reid, "An Algorithm for Tracking Multiple Targets," *IEEE Transactions on Automatic Control*, Vol. 24, 1979, pp. 843–854.
- [10] Y. B. Shalom and X. Li, *Multitarget-Multisensor Tracking: Principles and Techniques*. 1995.
- [11] Y. Blair and E. W. D. Blair, *Multitarget Multisensor tracking: Applications and advances, vol. III*. Norwood, MA: Artech House, 2000.
- [12] S. Oh *et al.*, "Markov Chain Monte Carlo Data Association for Multi-Target Tracking," *IEEE Tr. Automatic Control*, Vol. 54, 2009, pp. 481–497.

- [13] S. S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*. Norwood, MA: Artech House, 1999.
- [14] Y. B. Shalom and T. Fortmann, *Tracking and Data Association*. San Diego, CA: Academic Press, 1988.
- [15] N. Bergman and A. Doucet, “Markov Chain Monte Carlo data association for target tracking,” *Proc. ICASSP*, 2000.
- [16] H. Pasula *et al.*, “Tracking many objects with many sensors,” *Proc. IJCAI*, 1999.
- [17] S. Cong *et al.*, “Markov chain Monte Carlo approach for association probability evaluation,” *Proc. IEEE*, Vol. 151, 2004, pp. 185–193.
- [18] W. Faber, S. Chakravorty, and I. Hussein, “A Randomized Sampling based Approach to Multi-Object Tracking,” *International Conference of Information Fusion*, July 2015.