

REPORT DOCUMENTATION PAGE					<i>Form Approved</i> OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.						
1. REPORT DATE (DD-MM-YYYY) 03-09-2014		2. REPORT TYPE Final			3. DATES COVERED (From - To) 31 Jan 13 – 29 Apr 14	
4. TITLE AND SUBTITLE A Random Finite Set Approach to Space Junk Tracking and Identification				5a. CONTRACT NUMBER FA2386-13-1-4010		
				5b. GRANT NUMBER Grant AOARD-134010		
				5c. PROGRAM ELEMENT NUMBER 61102F		
6. AUTHOR(S) Professor Ba Ngu Vo				5d. PROJECT NUMBER		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Curtin University of Technology GPO BOX U1987 Perth 6845 Australia					8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AOARD UNIT 45002 APO AP 96338-5002					10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/AFOSR/IOA(AOARD)	
					11. SPONSOR/MONITOR'S REPORT NUMBER(S) AOARD-134010	
12. DISTRIBUTION/AVAILABILITY STATEMENT Distribution A: Approved for public release. Distribution is Unlimited.						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT The report summarizes the feasibility study performed for AFRL/AOARD regarding the use of Random Finite Set (RFS) filtering approach for tracking orbital space debris. Specifically, we investigate the capability of the RFS approach to accommodate: nonlinear motion and measurement models; unknown and time varying target number; and limited sensor field-of-view. This report shows the capability of the random finite set approach to provide large scale multi-target tracking. In particular it is shown that an approximate filter known as the labeled multi-Bernoulli filter can simultaneously track one thousand five hundred targets in clutter on a standard laptop computer.						
15. SUBJECT TERMS Data Fusion, Space Debris, Statistics						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 15	19a. NAME OF RESPONSIBLE PERSON Ingrid J. Wysong, Ph.D.	
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code) +81-3-5410-4409	

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 03 SEP 2014		2. REPORT TYPE Final		3. DATES COVERED 31-01-2013 to 29-04-2014	
4. TITLE AND SUBTITLE A Random Finite Set Approach to Space Junk Tracking and Identification			5a. CONTRACT NUMBER FA2386-13-1-4010		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER 61102F		
6. AUTHOR(S) Ba Ngu Vo			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Curtin University of Technology,GPO BOX U1987,Perth 6845,Australia,AU,6845			8. PERFORMING ORGANIZATION REPORT NUMBER N/A		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AOARD, UNIT 45002, APO, AP, 96338-5002			10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/AFOSR/IOA(AOARD)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) AOARD-134010		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The report summarizes the feasibility study performed for AFRL/AOARD regarding the use of Random Finite Set (RFS) filtering approach for tracking orbital space debris. Specifically, we investigate the capability of the RFS approach to accommodate: nonlinear motion and measurement models; unknown and time varying target number; and limited sensor field-of-view. This report shows the capability of the random finite set approach to provide large scale multi-target tracking. In particular it is shown that an approximate filter known as the labeled multi-Bernoulli filter can simultaneously track one thousand five hundred targets in clutter on a standard laptop computer.					
15. SUBJECT TERMS Data Fusion, Space Debris, Statistics					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 15	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

A Random Finite Set Approach to Space Junk Tracking and Identification

Ba-Ngu Vo¹, Ba-Tuong Vo¹,

¹Department of Electrical and Computer Engineering
Curtin University, Perth, WA, Australia.

ABSTRACT

The report summarizes the feasibility study performed for AFRL/AOARD regarding the use of Random Finite Set (RFS) filtering approach for tracking orbital space debris. Specifically, we investigate the capability of the RFS approach to accommodate: nonlinear motion and measurement models; unknown and time varying target number; and limited sensor field-of-view. This report shows the capability of the random finite set approach to provide large scale multi-target tracking. In particular it is shown that an approximate filter known as the labeled multi-Bernoulli filter can simultaneously track one thousand five hundred targets in clutter on a standard laptop computer.

1. INTRODUCTION

Space junk is the term commonly used to refer to the many millions of pieces of waste products and defunct objects currently in orbit around the earth, which are largely a result of human activities in relation to space exploration.^{1,9} Although the size and mass of individual fragments varies from millimetres to metres and grams to tonnes, even the most seemingly insignificant pieces can pose a real risk of collision and destruction for space related infrastructure and exploration. However it is currently understood that the most dangerous debris are those with diameters between 1 and 10 centimetres, of which it is estimated that there are several hundred thousand, since they are small enough to be difficult to track but large enough to cause significant damage. Fragments larger than 10cm can generally be tracked and avoided for time being, and those smaller than 1cm can generally be shielded from with appropriate apparatus.

A prerequisite for space debris management is to understand and maintain awareness of orbital space objects and the space environment, which, in the space domain, is called Space Situational Awareness (SSA). Utilising available data from multiple sources to keep track of multiple space objects is a core component of SSA.⁴ There are three major approaches to tracking multiple objects: Multiple Hypotheses Tracking (MHT);³ Joint Probabilistic Data Association (JPDA);² and Random Finite Set (RFS).¹² The comprehensive study of astrodynamics standards in,⁴ identified data association, information fusion and nonlinear estimation as three key research areas for tracking multiple space objects. MHT and its variations involve the propagation of data association hypotheses in time. The JPDA approach weights individual observations by their association probabilities. MHT and JPDA are classical approaches that dominated the field of multi-target tracking and are well documented in.^{2,3} The RFS approach provides a general systematic treatment of multi-target system by modeling the multi-target state as an RFS and accommodates

- nonlinear motion and measurement models
- unknown time varying target number
- multi-sensor data with support for heterogeneous sensor types
- limited sensor field-of-view

Scalable multi-target tracking solutions are needed in SSA where the number of targets is large.⁹ However, multi-target tracking is intrinsically an NP-hard problem and the complexity of multi-target tracking solutions

Further author information: Ba-Ngu Vo, E-mail: ba-ngu.vo@curtin.edu.au
Department of Electrical and Computer Engineering
Curtin University, Kent Street, Bentley, Perth Western Australia, 6102.

usually do not scale gracefully with problem size. The RFS approach provides the flexibility to derive multi-target tracking algorithms that bypass the data association problem e.g the Probability Hypothesis Density (PHD),¹⁰ Cardinalized PHD,¹¹ and Multi-Bernoulli filters.^{15,16} Although RFS approaches have been extensively applied in aerospace or the traditional defence related problems, there has been very little activity in the space situational awareness and specifically in the space junk monitoring domain. Current works include an application to tracking in the 0-1 case, or Bernoulli filtering, where at most one target is present with significant nonlinearities.^{7,7} Additionally a conceptual solution has been derived assuming a fixed a known number of targets.⁷ Still the feasibility of RFS/FISST approaches for space junk monitoring remains an open question. The broad objectives of this study are to

- Investigate feasibility of RFS based approach for large scale real time space junk monitoring
- Investigate methodological and computational strategies for practical implementation

The report is organized as follows. Background material on the RFS approach to multi-target tracking needed for this work is given in Section 2. In Section 3 we present the labeled multi-Bernoulli filter—an approximation to the optimal Bayes multi-target tracking filter—that has the potential to address large scale multi-target problems. In Section 4 we present a numerical example involving 1500 targets. To the best of our knowledge this is the only algorithm that can track such a large number of targets in clutter using only a standard laptop computer.

2. LABELED RANDOM FINITE SET AND MULTI-TARGET TRACKING

The random finite set (RFS) or finite set statistics (FISST) paradigm¹² for multi-object estimation is an ideal candidate platform for developing top-down algorithmic solutions for space junk monitoring. The RFS paradigm was initially developed as a completely new paradigm for multi-object estimation, encompassing a principled mathematical foundation as well as appropriate calculus like tools for manipulating probability densities of finite-set-valued random variables. The framework is general enough to accommodate non-linear motion and measurement models, including image measurement models, unknown and time varying number of target, arbitrary sensor field-of-view, and multi-sensor data support for heterogeneous sensor types (and even non-standard data), see.¹² Thus the RFS approach is ideally suited to networked multi-sensor multi-object estimation problems, such as space situational awareness.

RFS-based multi-target filters such as the Probability Hypothesis Density (PHD),¹⁰ Cardinalized PHD (CPHD)¹¹ and multi-Bernoulli filters^{15,16} can avoid data association. Indeed, under low clutter, it was demonstrated that a version of CPHD filter can simultaneously handle over a thousand targets on a standard laptop computer. However these filters, in principle, are not multi-target trackers because they rest on the premise that targets are indistinguishable.

In,¹⁷ the notion of *labeled RFSs* is introduced to address target trajectories and their uniqueness. Moreover, an analytic solution to the Bayes multi-target tracking filter known as the δ -generalized labeled multi-Bernoulli (δ -GLMB) filter was developed. Apart from producing tracks, the δ -GLMB filter does not assume high detection probability and low clutter and outperforms the PHD/CPHD and multi-Bernoulli filters. In this report we show that the labeled RFS technique is capable of tracking simultaneously thousands of targets in heavy clutter with different birth and death times, on a standard laptop computer. The key innovation is the family of labeled RFS conjugate priors in¹⁷ and a principled and efficient and highly parallelizable approximation, called the labeled multi-Bernoulli filter.¹⁹

2.1 Random Finite Sets

In a multiple target system, the number of targets varies with time due to the appearance and disappearance of targets, and the number of measurements received at each time step does not necessarily match the number of targets due to missed detections and clutter. The objective of multiple target filtering is to jointly estimate the number of targets and their states from the accumulated observations. Suppose that at time k , there are $N(k)$ target states $x_{k,1}, \dots, x_{k,N(k)}$, each taking values in a state space $\mathbb{X} \subseteq \mathbb{R}^{n_x}$, and $M(k)$ observations $z_{k,1}, \dots, z_{k,M(k)}$ each taking values in an observation space $\mathbb{Z} \subseteq \mathbb{R}^{n_z}$. Define the finite sets

$$\begin{aligned} X_k &= \{x_{k,1}, \dots, x_{k,N(k)}\} \subset \mathbb{X}, \\ Z_k &= \{z_{k,1}, \dots, z_{k,M(k)}\} \subset \mathbb{Z}, \end{aligned}$$

to be the *multiple target state* and *multiple target observation* respectively.

In the Bayesian estimation paradigm, the state and measurement are treated as realizations of random variables. Since the (multi-object) state X_k is a finite set, the concept of a random finite set (RFS) is required to cast the multi-object estimation problem in the Bayesian framework. An RFS is simply a random variable that take values as (unordered) finite sets, i.e. a *finite-set-valued random variable*. The space of finite subsets of $\mathbb{X}$ does not inherit the usual Euclidean notion of integration and density. Hence, standard tools for random vectors are not appropriate for RFSs. Mahler's Finite Set Statistics (FISST) provides powerful yet practical mathematical tools for dealing with RFSs.^{10,12} The notion of probability densities, generating functionals, and calculus like tools for are provided by the theory of *finite set statistics*.^{10,12} For connections between aspects of the FISST notion of probability densities and standard measure theoretic probability theory we refer the reader to.¹⁴

2.2 Labeled Random Finite Sets

To incorporate target identity, each state $x \in \mathbb{X}$ is augmented with a unique label $\ell \in \mathbb{L} = \{\ell_i : i \in \mathbb{N}\}$, where $\mathbb{N}$ denotes the set of positive integers and the ℓ_i 's are distinct. More detail on labeled RFS can be found in.^{17,18}

DEFINITION 2.1. A labeled RFS with state space $\mathbb{X}$ and (discrete) label space $\mathbb{L}$ is an RFS on $\mathbb{X} \times \mathbb{L}$ such that each realization has distinct labels.

Let $\mathcal{L} : \mathbb{X} \times \mathbb{L} \rightarrow \mathbb{L}$ be the projection $\mathcal{L}((x, \ell)) = \ell$, then a finite subset set $\mathbf{X}$ of $\mathbb{X} \times \mathbb{L}$ has distinct labels if and only if $\mathbf{X}$ and its labels $\mathcal{L}(\mathbf{X}) = \{\mathcal{L}(\mathbf{x}) : \mathbf{x} \in \mathbf{X}\}$ have the same cardinality, i.e. $\delta_{|\mathbf{X}|}(|\mathcal{L}(\mathbf{X})|) = 1$. The function $\Delta(\mathbf{X}) \triangleq \delta_{|\mathbf{X}|}(|\mathcal{L}(\mathbf{X})|)$ is called the *distinct label indicator*.

The unlabeled version of a labeled RFS is obtained by simply discarding the labels. Consequently, the cardinality distribution (the distribution of the number of objects) of a labeled RFS the same as its unlabeled version.¹⁷

For the rest of the paper, single-object states are represented by lowercase letters, e.g. $x, \mathbf{x}$ while multi-object states are represented by uppercase letters, e.g. $X, \mathbf{X}$, symbols for labeled states and their distributions are bolded to distinguish them from unlabeled ones, e.g. $\mathbf{x}, \mathbf{X}, \boldsymbol{\pi}$, etc., spaces are represented by blackboard bold e.g. $\mathbb{X}, \mathbb{Z}, \mathbb{L}, \mathbb{N}$, etc., and the class of finite subsets of a space $\mathbb{X}$ is denoted by $\mathcal{F}(\mathbb{X})$. The inner product of two functions f and g is denoted by $\langle f, g \rangle \triangleq \int f(x)g(x)dx$. The integral of a function $\mathbf{f} : \mathbb{X} \times \mathbb{L} \rightarrow \mathbb{R}$ is given by

$$\int \mathbf{f}(\mathbf{x})d\mathbf{x} = \sum_{\ell \in \mathbb{L}} \int_{\mathbb{X}} \mathbf{f}((x, \ell))dx.$$

2.3 Labeled multi-Bernoulli RFS

A labeled multi-Bernoulli RFS $\mathbf{X}$ with state space $\mathbb{X}$, label space $\mathbb{L}$ and parameter set $\{(r^{(\zeta)}, p^{(\zeta)}) : \zeta \in \Psi\}$, is a multi-Bernoulli RFS on $\mathbb{X}$ augmented with labels corresponding to the successful (non-empty) Bernoulli components, i.e. if the Bernoulli component $(r^{(\zeta)}, p^{(\zeta)})$ yields a non-empty set, then the label of the corresponding state is given by $\alpha(\zeta)$, where $\alpha : \Psi \rightarrow \mathbb{L}$ is a 1-1 mapping. A labeled multi-Bernoulli RFS has (multi-target) density given by¹⁷

$$\boldsymbol{\pi}(\{(x_1, \ell_1), \dots, (x_n, \ell_n)\}) = \delta_n(|\{\ell_1, \dots, \ell_n\}|) \prod_{\zeta \in \Psi} \left(1 - r^{(\zeta)}\right) \prod_{j=1}^n \frac{1_{\alpha(\Psi)}(\ell_j) r^{(\alpha^{-1}(\ell_j))} p^{(\alpha^{-1}(\ell_j))}(x_j)}{1 - r^{(\alpha^{-1}(\ell_j))}}.$$

where

$$1_Y(X) \triangleq \begin{cases} 1, & \text{if } X \subseteq Y \\ 0, & \text{otherwise.} \end{cases}$$

is the inclusion function. For convenience we use the abbreviation $\boldsymbol{\pi} = \{(r^{(\zeta)}, p^{(\zeta)})\}_{\zeta \in \Psi}$ for the density of a labeled multi-Bernoulli RFS. Although the formulation allows for a general mapping α for the labels, in this work we assume that α is an identity mapping in order to simplify notations. Thus superscripts for component

indices correspond directly to the label in question. The density of a labeled multi-Bernoulli RFS with parameter set $\boldsymbol{\pi} = \{r^{(\ell)}, p^{(\ell)}\}_{\ell \in \mathbb{L}}$ is given more compactly by

$$\boldsymbol{\pi}(\mathbf{X}) = \Delta(\mathbf{X}) w(\mathcal{L}(\mathbf{X})) p^{\mathbf{X}} \quad (1)$$

where $p^{\mathbf{X}} = \prod_{\mathbf{x} \in \mathbf{X}} p(\mathbf{x})$,

$$w(L) = \prod_{i \in \mathbb{L}} \left(1 - r^{(i)}\right) \prod_{\ell \in L} \frac{1_{\mathbb{L}}(\ell) r^{(\ell)}}{1 - r^{(\ell)}}, \quad (2)$$

$$p(x, \ell) = p^{(\ell)}(x). \quad (3)$$

2.4 Generalized Labeled Multi-Bernoulli RFS

A generalized labeled multi-Bernoulli RFS is a labeled RFS on $\mathbb{X} \times \mathbb{L}$ distributed according to

$$\boldsymbol{\pi}(\mathbf{X}) = \Delta(\mathbf{X}) \sum_{c \in \mathbb{C}} w^{(c)}(\mathcal{L}(\mathbf{X})) \left[p^{(c)}\right]^{\mathbf{X}} \quad (4)$$

where $\mathbb{C}$ is a discrete index set, $w^{(c)}(L)$ and $p^{(c)}$ satisfy

$$\sum_{L \subseteq \mathbb{L}} \sum_{c \in \mathbb{C}} w^{(c)}(L) = 1, \quad (5)$$

$$\int p^{(c)}(x, \ell) dx = 1. \quad (6)$$

The cardinality distribution of a generalized labeled multi-Bernoulli RFS is given by

$$\rho(n) = \sum_{L \in \mathcal{F}_n(\mathbb{L})} \sum_{c \in \mathbb{C}} w^{(c)}(L) \quad (7)$$

The labeled multi-Bernoulli RFS distributed by (1) is a special case of the generalized labeled multi-Bernoulli RFS with

$$\begin{aligned} p^{(c)}(x, \ell) &= p^{(\ell)}(x) \\ w^{(c)}(L) &= \prod_{i \in \mathbb{L} - L} \left(1 - r^{(i)}\right) \prod_{\ell \in L} 1_{\mathbb{L}}(\ell) r^{(\ell)} \end{aligned}$$

comprising a single component in which case the superscript (c) is omitted.

2.5 GLMB Multi-target Tracking Filter

In labeled RFS multi-target tracking, targets are identified by an ordered pair of integers $\ell = (k, i)$, where k is the time of birth and $i \in \mathbb{N}$, with $\mathbb{N}$ denoting the set of natural numbers, is a unique index to distinguish objects born at the same time. The label space for objects born at time k , denoted as $\mathbb{L}_k$, is then $\{k\} \times \mathbb{N}$. An object born at time k , has state $\mathbf{x} \in \mathbb{X} \times \mathbb{L}_k$. The label space for targets at time k (including those born prior to k), denoted as $\mathbb{L}_{0:k}$, is constructed recursively by $\mathbb{L}_{0:k} = \mathbb{L}_{0:k-1} \cup \mathbb{L}_k$ (note that $\mathbb{L}_{0:k-1}$ and $\mathbb{L}_k$ are disjoint). A multi-object state $\mathbf{X}$ at time k , is a finite subset of $\mathbb{X} \times \mathbb{L}_{0:k}$.

Let $\boldsymbol{\pi}_k(\cdot|Z_k)$ denotes the *multi-target posterior density* at time k , and $\boldsymbol{\pi}_{k+1|k}$ denotes the *multi-target prediction density* to time k . Then, the *multi-target Bayes recursion* propagates $\boldsymbol{\pi}_k$ in time¹² according to the following update and prediction

$$\boldsymbol{\pi}_k(\mathbf{X}_k|z_k) = \frac{g_k(z_k|\mathbf{X}_k) \boldsymbol{\pi}_{k|k-1}(\mathbf{X}_k)}{\int g_k(z_k|\mathbf{X}) \boldsymbol{\pi}_{k|k-1}(\mathbf{X}) \delta \mathbf{X}}, \quad (8)$$

$$\boldsymbol{\pi}_{k+1|k}(\mathbf{X}_{k+1}) = \int \mathbf{f}_{k+1|k}(\mathbf{X}_{k+1}|\mathbf{X}_k) \boldsymbol{\pi}_k(\mathbf{X}_k|z_k) \delta \mathbf{X}_k, \quad (9)$$

where $g_k(\cdot|\cdot)$ is the *multi-target likelihood function* at time k , $\mathbf{f}_{k+1|k}(\cdot|\cdot)$ is the *multi-target transition density*, to time $k+1$, and the integral is a set integral defined for any function $\mathbf{f} : \mathcal{F}(\mathbb{X} \times \mathbb{L}) \rightarrow \mathbb{R}$ by

$$\int \mathbf{f}(\mathbf{X}) \delta \mathbf{X} = \sum_{i=0}^{\infty} \frac{1}{i!} \int \mathbf{f}(\{\mathbf{x}_1, \dots, \mathbf{x}_i\}) d(\mathbf{x}_1, \dots, \mathbf{x}_i).$$

The multi-target posterior density captures all information on target number, and individual target states.¹² The multi-target likelihood function encapsulates the underlying models for detections and false alarms while the multi-target transition density encapsulates the underlying models of target motions, births and deaths.

For convenience, in what follows we omit explicit references to the time index k , and denote $\mathbb{L} \triangleq \mathbb{L}_{0:k}$, $\mathbb{B} \triangleq \mathbb{B}_{k+1}$, $\mathbb{L}_+ \triangleq \mathbb{L} \cup \mathbb{B}$, $\boldsymbol{\pi} \triangleq \boldsymbol{\pi}_k$, $\boldsymbol{\pi}_+ \triangleq \boldsymbol{\pi}_{k+1|k}$, $g \triangleq g_k$, $\mathbf{f} \triangleq \mathbf{f}_{k+1|k}$. In this work we assume the space of birth label $\mathbb{B}$ is finite.

2.6 GLMB Prediction

Given the current multi-object state $\mathbf{X}$, each state $(x, \ell) \in \mathbf{X}$ either continues to exist at the next time step with probability $p_S(x, \ell)$ and evolves to a new state (x_+, ℓ_+) with probability density $f(x_+|x, \ell) \delta_{\ell}(\ell_+)$, or dies with probability $1 - p_S(x, \ell)$. The set of new targets born at the next time step is distributed according to

$$\mathbf{f}_B(\mathbf{Y}) = \Delta(\mathbf{Y}) w_B(\mathcal{L}(\mathbf{Y})) [p_B]^{\mathbf{Y}} \quad (10)$$

The birth density $\mathbf{f}_B$ is defined on $\mathbb{X} \times \mathbb{B}$ and $\mathbf{f}_B(\mathbf{Y}) = 0$ if $\mathbf{Y}$ contains any element $\mathbf{y}$ with $\mathcal{L}(\mathbf{y}) \notin \mathbb{B}$. The birth model (10) covers both labeled Poisson and labeled multi-Bernoulli. For a labeled multi-Bernoulli birth model:

$$w_B(J) = \prod_{i \in \mathbb{B} - J} \left(1 - r_B^{(i)}\right) \prod_{\ell \in J} 1_{\mathbb{B}}(\ell) r_B^{(\ell)}, \quad (11)$$

$$p_B(x, \ell) = p_B^{(\ell)}(x). \quad (12)$$

The multi-target state at the next time $\mathbf{X}_+$ is the superposition of surviving targets and new born targets. Assuming that targets evolve independently of each other and that births are independent of surviving targets, it was shown in¹⁷ that the multi-target transition kernel is given by

$$\mathbf{f}(\mathbf{X}_+|\mathbf{X}) = \mathbf{f}_S(\mathbf{X}_+ \cap (\mathbb{X} \times \mathbb{L})|\mathbf{X}) \mathbf{f}_B(\mathbf{X}_+ - (\mathbb{X} \times \mathbb{L})) \quad (13)$$

where

$$\begin{aligned} \mathbf{f}_S(\mathbf{W}|\mathbf{X}) &= \Delta(\mathbf{W}) \Delta(\mathbf{X}) 1_{\mathcal{L}(\mathbf{X})}(\mathcal{L}(\mathbf{W})) [\Phi(\mathbf{W}; \cdot)]^{\mathbf{X}} \\ \Phi(\mathbf{W}; x, \ell) &= \begin{cases} p_S(x, \ell) f(x_+|x, \ell), & \text{if } (x_+, \ell) \in \mathbf{W} \\ 1 - p_S(x, \ell), & \text{if } \ell \notin \mathcal{L}(\mathbf{W}) \end{cases} \end{aligned}$$

PROPOSITION 2.2. *If the current multi-object prior is a generalized labeled multi-Bernoulli of the form (4), then the predicted multi-object density is also a generalized labeled multi-Bernoulli given by*

$$\boldsymbol{\pi}_+(\mathbf{X}_+) = \Delta(\mathbf{X}_+) \sum_{c \in \mathbb{C}} w_+^{(c)}(\mathcal{L}(\mathbf{X}_+)) \left[p_+^{(c)}\right]^{\mathbf{X}_+} \quad (14)$$

where

$$w_+^{(c)}(L) = w_B(L - \mathbb{L}) w_S^{(c)}(L \cap \mathbb{L}), \quad (15)$$

$$p_+^{(c)}(x, \ell) = 1_{\mathbb{L}}(\ell) p_S^{(c)}(x, \ell) + (1 - 1_{\mathbb{L}}(\ell)) p_B(x, \ell), \quad (16)$$

$$p_S^{(c)}(x, \ell) = \frac{\langle p_S(\cdot, \ell) f(x|\cdot, \ell), p^{(c)}(\cdot, \ell) \rangle}{\eta_S^{(c)}(\ell)}, \quad (17)$$

$$\eta_S^{(c)}(\ell) = \langle p_S(\cdot, \ell), p^{(c)}(\cdot, \ell) \rangle \quad (18)$$

$$w_S^{(c)}(J) = [\eta_S^{(c)}]^J \sum_{I \subseteq \mathbb{L}} 1_I(J) [q_S^{(c)}]^{I-J} w^{(c)}(I), \quad (19)$$

$$q_S^{(c)}(\ell) = 1 - \langle p_S(\cdot, \ell), p^{(c)}(\cdot, \ell) \rangle. \quad (20)$$

2.7 GLMB Update

For a given multi-target state $\mathbf{X}$, at time k , each state $(x, \ell) \in \mathbf{X}$ is either detected with probability $p_D(x, \ell)$ and generates a point z with likelihood $g(z|x, \ell)$, or missed with probability $1 - p_D(x, \ell)$. The multi-object observation $Z = \{z_1, \dots, z_{|Z|}\}$ is the superposition of the detected points and Poisson clutter with intensity function κ .

DEFINITION 2.3. *An association map (for the current time) is a function $\theta : \mathbb{L} \rightarrow \{0, 1, \dots, |Z|\}$ such that $\theta(i) = \theta(i') > 0$ implies $i = i'$. The set Θ of all such association maps is called the association map space. The subset of Θ with domain I is denoted by $\Theta(I)$.*

An association map describes which tracks generated which measurements, i.e. track ℓ generates measurement $z_{\theta(\ell)} \in Z$, with undetected tracks assigned to 0. The condition $\theta(i) = \theta(i') > 0$ implies $i = i'$, means that a track can generate at most one measurement at any point in time.

Assuming that, conditional on $\mathbf{X}$, detections are independent, and that clutter is independent of the detections, the multi-object likelihood is given by

$$g(Z|\mathbf{X}) = e^{-\langle \kappa, 1 \rangle} \kappa^Z \sum_{\theta \in \Theta(\mathcal{L}(\mathbf{X}))} [\psi_Z(\cdot; \theta)]^{\mathbf{X}} \quad (21)$$

where

$$\psi_Z(x, \ell; \theta) = \begin{cases} \frac{p_D(x, \ell) g(z_{\theta(\ell)}|x, \ell)}{\kappa(z_{\theta(\ell)})}, & \text{if } \theta(\ell) > 0 \\ 1 - p_D(x, \ell), & \text{if } \theta(\ell) = 0 \end{cases} \quad (22)$$

PROPOSITION 2.4. *If the prior distribution is a generalized labeled multi-Bernoulli of the form (4), then, under the multi-object likelihood function (21), the posterior distribution is also a generalized multi-Bernoulli given by*

$$\pi(\mathbf{X}|Z) = \Delta(\mathbf{X}) \sum_{c \in \mathbb{C}} \sum_{\theta \in \Theta(\mathcal{L}(\mathbf{X}))} w_Z^{(c, \theta)}(\mathcal{L}(\mathbf{X})) [p^{(c, \theta)}(\cdot|Z)]^{\mathbf{X}} \quad (23)$$

where

$$w_Z^{(c, \theta)}(L) \propto w^{(c)}(L) [\eta_Z^{(c, \theta)}]^L, \quad (24)$$

$$p^{(c, \theta)}(x, \ell|Z) = \frac{p^{(c)}(x, \ell) \psi_Z(x, \ell; \theta)}{\eta_Z^{(c, \theta)}(\ell)}, \quad (25)$$

$$\eta_Z^{(c, \theta)}(\ell) = \langle p^{(c)}(\cdot, \ell), \psi_Z(\cdot, \ell; \theta) \rangle, \quad (26)$$

For a valid label set L , the updated weight $w_Z^{(c, \theta)}(L)$ is proportional to the prior weight $w^{(c)}(L)$ scaled by the product $[\eta_Z^{(c, \theta)}]^L$ of single-object normalizing constants.

The above results are the prediction and update step of the δ -generalized labeled multi-Bernoulli tracking filter,¹⁷ also known as the Vo-Vo filter (see Mahlerbook2014). Tractable techniques for truncating the posterior and prediction densities were proposed based on the k -shortest paths and ranked assignment algorithms, see¹⁸ for further details.

3. LABELED MULTI-BERNOULLI FILTER FOR LARGE SCALE TRACKING

The labeled multi-Bernoulli filter is an approximation of the δ -generalized labeled multi-Bernoulli filter using labeled multi-Bernoulli RFS. The number of components in the the Vo-Vo filter grows exponentially, whereas the growth is linear for a labeled multi-Bernoulli representation. In this section we outline the labeled multi-Bernoulli filter and its implementation, further detail can be found in.¹⁹

3.1 LMB Prediction

It was shown in¹⁹ that the labeled multi-Bernoulli is closed under the prediction step.

PROPOSITION 3.1. *Suppose that the multi-target posterior density is a labeled multi-Bernoulli with parameter set $\pi = \{r^{(\ell)}, p^{(\ell)}\}_{\ell \in \mathbb{L}}$, and that the multi-target birth model is a labeled multi-Bernoulli with parameter set $\pi_B = \{r_B^{(\ell)}, p_B^{(\ell)}\}_{\ell \in \mathbb{B}}$ (where $\mathbb{L} \cap \mathbb{B} = \emptyset$), then the multi-target predicted density is also a labeled multi-Bernoulli with parameter set*

$$\pi_+ = \left\{ \left(r_{+,S}^{(\ell)}, p_{+,S}^{(\ell)} \right) \right\}_{\ell \in \mathbb{L}} \cup \left\{ \left(r_B^{(\ell)}, p_B^{(\ell)} \right) \right\}_{\ell \in \mathbb{B}}, \quad (27)$$

where

$$r_{+,S}^{(\ell)} = \eta_S(\ell) r^{(\ell)}, \quad (28)$$

$$p_{+,S}^{(\ell)} = \langle p_S(\cdot, \ell) f(x|\cdot, \ell), p(\cdot, \ell) \rangle / \eta_S(\ell), \quad (29)$$

$$\eta_S(\ell) = \langle p_S(\cdot, \ell), p(\cdot, \ell) \rangle \quad (30)$$

REMARK 3.2. *The multi-target prediction for a labeled multi-Bernoulli actually coincides with the prediction for the unlabeled multi-Bernoulli and interpreting the component indices as track labels.*

3.2 LMB Update

While the labeled multi-Bernoulli family is closed under the prediction operation, it is not closed under the update operation. Inspired by the multi-Bernoulli filter,¹⁵ we seek the labeled multi-Bernoulli approximation that matches the first moment of the multi-target posterior density. One such labeled multi-Bernoulli approximation is given in.¹⁹

PROPOSITION 3.3. *Suppose that the multi-target predicted density is a labeled multi-Bernoulli with parameter set $\pi_+ = \{r_+^{(\ell)}, p_+^{(\ell)}\}_{\ell \in \mathbb{L}_+}$. The labeled multi-Bernoulli that matches exactly the first moment of the multi-target posterior density is $\pi(\cdot|Z) = \{(r^{(\ell)}, p^{(\ell)}(\cdot))\}_{\ell \in \mathbb{L}_+}$ where*

$$r^{(\ell)} = \sum_{(I_+, \theta) \in \mathcal{F}(\mathbb{L}_+) \times \Theta(I_+)} w^{(I_+, \theta)}(Z) 1_{I_+}(\ell), \quad (31)$$

$$p^{(\ell)}(x) = \frac{1}{r^{(\ell)}} \sum_{(I_+, \theta) \in \mathcal{F}(\mathbb{L}_+) \times \Theta(I_+)} w^{(I_+, \theta)}(Z) 1_{I_+}(\ell) p^{(\theta)}(x, \ell|Z), \quad (32)$$

$$w^{(I_+, \theta)}(Z) \propto \left[\eta_Z^{(\theta)} \right]^{I_+} \prod_{\ell \in \mathbb{L}_+ - I_+} \left(1 - r_+^{(\ell)} \right) \prod_{i \in I_+} 1_{\mathbb{L}_+}(i) r_+^{(i)}, \quad (33)$$

$$p^{(\theta)}(x, \ell|Z) = \frac{p_+(x, \ell) \psi_Z(x, \ell; \theta)}{\eta_Z^{(\theta)}(\ell)}, \quad (34)$$

$$\eta_Z^{(\theta)}(\ell) = \langle p_+(\cdot, \ell), \psi_Z(\cdot, \ell; \theta) \rangle. \quad (35)$$

The key advantage of the labeled multi-Bernoulli update is that it only involves one approximation of the multi-target posterior density. In contrast the multi-Bernoulli filter requires two approximations on the multi-target posterior probability generating functional.¹⁵

The mean cardinality of the labeled multi-Bernoulli approximation is identical to that of the full posterior. However, the cardinality distributions themselves differ, since the cardinality distribution of the labeled multi-Bernoulli approximation follows the one for a multi Bernoulli RFS while the cardinality distribution of the full posterior follows that of a Generalized labeled multi-Bernoulli (GLMB) RFS. Although there are many possible choices of approximate labeled multi-Bernoulli posteriors, the above choice preserves the estimated spatial densities of each track and matches exactly the first moment.¹⁹

3.3 Efficient Implementation of the Labeled Multi-Bernoulli Filter

The labeled multi-Bernoulli form can naturally be decomposed into products of smaller labeled multi-Bernoullis that allows for parallel update of ‘groups’ of closely spaced targets and their associated measurements, which drastically reduces computation time. Target grouping is based on a standard gating procedure which also partitions the measurement set, resulting in groups of targets and measurements which can be considered statistically independent.⁵ Each resultant group is usually much smaller than the entire multi-target state and measurement, and can then be updated in parallel, which is usually much simpler and faster than updating the entire multi-target state as a single group. The update for each group is performed by expanding the labeled multi-Bernoulli prediction into δ -GLMB form, and performing a standard δ -GLMB update resulting in a δ -GLMB posterior.¹⁷ The posterior for each group is then collapsed to a matching labeled multi-Bernoulli approximation after which groups are recombined and the recursion continues. This proposed implementation applies to both Gaussian Mixture and Sequential Monte Carlo based representations for the underlying single target densities.¹⁹

3.3.1 Grouping and Gating

Let $\{\mathbb{L}_+^{(1)}, \dots, \mathbb{L}_+^{(N)}\}$ be a partition of the label set $\mathbb{L}_+ = \mathbb{L} \cup \mathbb{B}$, and $\{Z^{(0)}, Z^{(1)}, \dots, Z^{(N)}\}$ be a partition of the measurement set Z . A grouping $\mathcal{G}^{(n)}$ is defined as the set of pairs $\{(\mathbb{L}_+^{(1)}, Z^{(1)}), \dots, (\mathbb{L}_+^{(N)}, Z^{(N)})\}$ and each pair $\mathcal{G}^{(n)} = (\mathbb{L}_+^{(n)}, Z^{(n)})$ is referred to as a group. Note that $Z^{(0)}$ is the set of measurements which are not assigned to any target labels.

We are interested in groupings of target labels and likely corresponding measurements. To generate such groupings, we start with an initial grouping of each labeled Bernoulli component ℓ and any associated measurements which fall within its prediction gates:

$$\tilde{\mathcal{G}}^{(\ell)} = (\{\ell\}, \{z : d_{\text{MHD}}(\hat{z}^{(\ell)}, z) < \sqrt{\gamma}\}). \quad (36)$$

where $d_{\text{MHD}}(\hat{z}^{(\ell)}, z)$ is the Mahalanobis distance (MHD) between the predicted measurement for track ℓ and the received measurement $z \in Z$, and γ is the gating distance threshold calculated using the inverse Chi-squared cumulative distribution corresponding to the desired σ -gate size for gating of measurements from tracks. Then, tentative groups with common measurements

$$Z^{(i)} \cap Z^{(j)} \neq \emptyset \quad (37)$$

are merged as follows

$$\tilde{\mathcal{G}}^{(i,j)} = (\mathbb{L}_+^{(i)} \cup \mathbb{L}_+^{(j)}, Z^{(i)} \cup Z^{(j)}). \quad (38)$$

The merging is repeated for all tentative groups until there are no common measurements. Finally, a total of N groups $\mathcal{G}^{(1)}, \dots, \mathcal{G}^{(N)}$ of tracks and associated measurements is obtained. Consequently, the predicted multi-target density can be rewritten as

$$\pi_+ = \bigcup_{i=1}^N \pi_+^{(i)} \quad (39)$$

where

$$\pi_+^{(i)} = \left\{ (r_+^{(\ell)}, p_+^{(\ell)}) \right\}_{\ell \in \mathbb{L}_+^{(i)}} \quad (40)$$

3.3.2 Parallel δ -GLMB Group Updates

Since the predicted multi-target density for each group is a labeled multi-Bernoulli of the form (39), we need to express it in δ -GLMB form, in order to apply the data update. The δ -GLMB form for the i -th group $\mathcal{G}^{(i)} = (\mathbb{L}_+^{(i)}, Z^{(i)})$ is given by

$$\pi_+^{(i)}(\mathbf{X}_+) = \Delta(\mathbf{X}_+) \sum_{I_+ \in \mathcal{F}(\mathbb{L}_+^{(i)})} w_{+,i}^{(I_+)} \delta_{I_+}(\mathcal{L}(\mathbf{X}_+)) [p_+]^{\mathbf{X}_+} \quad (41)$$

where

$$w_{+,i}^{(I_+)} = \prod_{\ell \in \mathbb{L}_+^{(i)} - I_+} (1 - r_+^{(\ell)}) \prod_{\ell' \in I_+} 1_{\mathbb{L}_+^{(i)}}(\ell') r_+^{(\ell')},$$

Due to the smaller label space within one group, the number of hypotheses across all groups is significantly smaller than for the number of hypotheses in case of a single big group.

A brute-force way to enumerate the sum is to generate all possible combinations for a set of labels $\mathbb{L}_+^{(i)}$ and cardinalities $n = 0, 1, \dots, |\mathbb{L}_+^{(i)}|$. The number of combinations for each cardinality is given by the binomial coefficient $C(|\mathbb{L}_+^{(i)}|, n) = |\mathbb{L}_+^{(i)}|! / (n!(|\mathbb{L}_+^{(i)}| - n)!)$ and the number of combinations for a set of track labels is $2^{|\mathbb{L}_+^{(i)}|}$. Consequently explicit evaluation of all combinations is only possible for small $|\mathbb{L}_+^{(i)}|$. For larger $|\mathbb{L}_+^{(i)}|$, the sum can be approximated to its k most significant terms by use of the k -shortest paths algorithm without enumerating all possible terms.¹⁷ Consequently, I_+ only consists of the most significant hypotheses.

The δ -GLMB update for each group i is given by

$$\pi^{(i)}(\mathbf{X}_+ | Z^{(i)}) = \Delta(\mathbf{X}_+) \sum_{(I_+, \theta) \in \mathcal{F}(\mathbb{L}_+^{(i)}) \times \Theta_{I_+}} w^{(I_+, \theta)}(Z^{(i)}) \delta_{I_+}(\mathcal{L}(\mathbf{X}_+)) \left[p^{(\theta)}(\cdot | Z^{(i)}) \right]^{\mathbf{X}_+}$$

where

$$w^{(I_+, \theta)}(Z^{(i)}) \propto \left[\eta_{Z^{(i)}}^{(\theta)} \right]^{I_+} \prod_{\ell \in \mathbb{L}_+^{(i)} - I_+} (1 - r_+^{(\ell)}) \prod_{\ell' \in I_+} 1_{\mathbb{L}_+^{(i)}}(\ell') r_+^{(\ell')} \quad (42)$$

$$p^{(\theta)}(x, \ell | Z^{(i)}) = \frac{p_+(x, \ell) \psi_{Z^{(i)}}(x, \ell; \theta)}{\eta_{Z^{(i)}}^{(\theta)}(\ell)}, \quad (43)$$

$$\eta_{Z^{(i)}}^{(\theta)}(\ell) = \langle p_+(x, \ell), \psi_{Z^{(i)}}(\cdot, \ell; \theta) \rangle, \quad (44)$$

$$\psi_{Z^{(i)}}(x, \ell; \theta) = \begin{cases} \frac{p_D(x, \ell) p_G g(z_{\theta(\ell)}^{(i)} | x, \ell)}{\kappa(z_{\theta(\ell)}^{(i)})}, & \text{if } \theta(\ell) > 0 \\ 1 - p_D(x, \ell) p_G, & \text{if } \theta(\ell) = 0 \end{cases}. \quad (45)$$

Observe, that (45) has to incorporate the gating probability p_G ,³ since small gates increase the probability of missed detection.

By representing the complete predicted distribution using several groups, the number of track labels $|\mathbb{L}_+^{(i)}|$ within each group is significantly lower than the total number of possible labels $|\mathbb{L}_+|$. Since the update is combinatorial, the number of components or hypotheses within each group grows exponentially in the number of track labels $|\mathbb{L}_+^{(i)}|$. Thus, for large $|\mathbb{L}_+^{(i)}|$ a truncation of the δ -GLMB distribution is required. The truncation can be realized using ranked assignment algorithm which evaluates only the M most significant hypotheses without evaluating all possible solutions.¹⁸ Since the complexity of ranked assignment algorithm is cubic, the computational costs for the evaluation of multiple groups is generally cheaper, compared to the evaluation for a single large group. Moreover, the evaluation for each group can be performed in parallel.

3.3.3 Approximation of the updated δ -GLMB as labeled multi-Bernoulli

After performing the group updates across all groups $\mathcal{G}^{(i)}$ for $i = 1, \dots, N$ the δ -GLMB form for each group is then collapsed back to labeled multi-Bernoulli form

$$\pi^{(i)}(\cdot|Z^{(i)}) \approx \tilde{\pi}^{(i)}(\cdot|Z^{(i)}) = \left\{ \left(r^{(\ell,i)}, p^{(\ell,i)} \right) \right\}_{\ell \in \mathbb{L}_+^{(i)}}$$

where $r^{(\ell,i)}, p^{(\ell,i)}$ are calculated according to (31)-(32) and taking the union of the approximate labeled multi-Bernoulli groups given by $\tilde{\pi}^{(i)}$ yields the labeled multi-Bernoulli approximation to the multi-target posterior

$$\pi(\cdot|Z) \approx \tilde{\pi}(\cdot|Z) = \bigcup_{i=1}^N \tilde{\pi}^{(i)}(\cdot|Z^{(i)}). \quad (46)$$

REMARK 3.4. *Theoretically in the worst case where all targets are close together, it may not be possible to partition into smaller groups. In this case, the complexity of the update step is the same as that of the δ -GLMB filter. However, our experience with empirical data suggests that in most scenarios we can partition the targets into many groups with small numbers of targets.*

3.4 Track Extraction

Since tracks are represented after update by a labeled multi-Bernoulli, an obvious track extraction scheme is to pick all tracks with an existence probability higher than an application specific threshold:

$$\hat{\mathbf{X}} = \left\{ (\hat{x}, \ell) : r^{(\ell)} > \vartheta \right\} \quad (47)$$

where $\hat{x} = \arg_x \max p^{(\ell)}(x)$. On the one hand, choosing a high value for ϑ significantly reduces the number of clutter tracks at the cost of a delayed inclusion of a new-born track. On the other hand, low values for ϑ report new-born tracks immediately at the cost of a higher number of clutter tracks. Choosing a high value for ϑ , additional care has to be taken for missed detections. In case of $p_D \approx 1$, a missed detection considerably reduces the existence probability. Consequently, one missed detection might suppress the output of a previously confirmed track with $r^{(\ell)} \approx 1$.

To mitigate this issue, a hysteresis is used, which activates the output if the maximum existence probability $r_{max}^{(\ell)}$ of a track ℓ has once exceeded an upper threshold ϑ_u and if the current existence probability $r^{(\ell)}$ is higher than a lower threshold ϑ_l :

$$\hat{\mathbf{X}} = \{ (\hat{x}, \ell) : r_{max}^{(\ell)} > \vartheta_u \wedge r^{(\ell)} > \vartheta_l \} \quad (48)$$

4. A LARGE SCALE TRACKING DEMONSTRATION

This demonstration of the labeled multi-Bernoulli (LMB) filter involves 3D states and observations in x, y, z coordinates, and features 1500 targets and various births and deaths throughout. Figure 1 shows the true tracks in 3D-space to indicate the density and number of targets in this scenario. We note the the implementation for this example is a standard serial implementation and has not been parallelised. Parallelisation would significant improve teh speed of the as well as scalability.

The following dynamical and measurement models are used. The target state variable is a vector of 3D position and velocity $x_k = [p_{x,k}, \dot{p}_{x,k}, p_{y,k}, \dot{p}_{y,k}, p_{z,k}, \dot{p}_{z,k}]^T$. The single-target transition model is linear Gaussian with transition density $f_{k|k-1}(x|\zeta) = \mathcal{N}(x; F_{k-1}\zeta, Q_{k-1})$ where

$$F_{k-1} = \begin{bmatrix} 1 & \Delta & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & \Delta & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & \Delta \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \quad Q_{k-1} = \sigma_\nu^2 \begin{bmatrix} \frac{\Delta^4}{4} & \frac{\Delta^3}{2} & 0 & 0 & 0 & 0 \\ \frac{\Delta^3}{2} & \Delta^2 & 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{\Delta^4}{4} & \frac{\Delta^3}{2} & 0 & 0 \\ 0 & 0 & \frac{\Delta^3}{2} & \Delta^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{\Delta^4}{4} & \frac{\Delta^3}{2} \\ 0 & 0 & 0 & 0 & \frac{\Delta^3}{2} & \Delta^2 \end{bmatrix},$$

$\Delta = 1s$ is the sampling period, and $\sigma_\nu = 5km/s^2$ is the standard deviation of the process noise. The probability of survival $p_{S,k} = 0.99$. The cardinality distributions involved have a maximum support of $N_{\max} = 2000$ terms. The birth process is modeled as a Poisson RFS with an average of 15 births per scan, coming from 3 different positions in the 3D surveillance region.

The probability of detection $p_{D,k} = 0.98$. The single-target measurement model is also linear Gaussian with likelihood $g_k(z|x) = \mathcal{N}(z; H_k x, R_k)$ where

$$H_k = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}, \quad R_k = \sigma_\varepsilon^2 \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

$\sigma_\varepsilon = 10km$ is the standard deviation of the measurement noise. Clutter is modeled as a Poisson RFS with uniform distribution at an average of 100 points per scan.




Figure 1. 3D tracks for approximately 1500 targets.

Note that we use the linear Gaussian model for this example as a proof-of-concept, to show that the proposed LMB filter can handle a large number of targets. Non-linear models that include non-uniform clutter, non-uniform sensor field-of-view can be easily accommodated using a particle implementation. However, the cost of single-target filtering increases, which in turn increase the computational requirement for the multi-target tracker.

For performance evaluation, the Optimal Sub-Pattern Assignment (OSPA) distance between the estimated and true multiple target state is used as the estimation error.¹³ The OSPA metric $\bar{d}_p^{(c)}$ is defined as follows. Let $d^{(c)}(x, y) := \min(c, \|x - y\|)$ for $x, y \in \mathcal{X}$, and Π_k denote the set of permutations on $\{1, 2, \dots, k\}$ for any positive integer k . Then, for $p \geq 1$, $c > 0$, and $X = \{x_1, \dots, x_m\}$ and $Y = \{y_1, \dots, y_n\}$,

$$\bar{d}_p^{(c)}(X, Y) := \left(\frac{1}{n} \left(\min_{\pi \in \Pi_n} \sum_{i=1}^m d^{(c)}(x_i, y_{\pi(i)})^p + c^p(n-m) \right) \right)^{\frac{1}{p}}$$

if $m \leq n$, and $\bar{d}_p^{(c)}(X, Y) := \bar{d}_p^{(c)}(Y, X)$ if $m > n$; and $\bar{d}_p^{(c)}(X, Y) = \bar{d}_p^{(c)}(Y, X) = 0$ if $m = n = 0$. The OSPA distance is interpreted as a p -th order per-target error, comprised of a p -th order per-target localization error

and a p -th order per-target cardinality error. The order parameter p determines the sensitivity of the metric to outliers, and the cut-off parameter c determines the relative weighting of the penalties assigned to cardinality and localization errors.

Plots of true and estimated multiple target cardinality as well as the OSPA miss-distance (for $p = 1$ and $c = 30km$) from the labeled multi-Bernoulli filter are shown in figures 2 and 3. From the plot of the cardinality in figure 2, it appears that the labeled multi-Bernoulli filter determines the number of targets within a reasonable error. The plots of the OSPA miss distance,¹³ from figure 3, confirm the correct operation of the filter; the errors are consistent with the measurement noise in the model and the number of target as seen from the cardinality plot. Execution time per cluster for each iteration of the filter on a single standard laptop computer is shown in figure 4. Consequently, with modest computational resources, the labeled multi-Bernoulli filter can potentially deliver real time performance in large scale tracking.




Figure 2. LMB estimated cardinality versus time




Figure 3. OSPA miss distance given by total, localization component and cardinality component

5. CONCLUSION

The RFS approach is ideally suited to networked multi-sensor multi-object estimation problems, such as space situational awareness. It provides a general systematic treatment of multi-target system and accommodates: nonlinear motion and measurement models; unknown time varying target number; multi-sensor data with support for heterogeneous sensor types; limited sensor field-of-view. More importantly, it provides scalable algorithms that can be used to track a large number of targets. A 3D scenario using Gaussian mixture implementations has shown that the labeled multi-Bernoulli filter can simultaneously track up to 1500 targets, with different birth and death times, on a single standard laptop computer. Performance analysis of the filter also suggests that it accurately estimates target numbers and states in dense situations. We note the the implementation for this example is a standard serial implementation and has not been parallelised. Without parallelization, execution




Figure 4. Average execution time per cluster for each iteration

times for each iteration was of the order of tens of seconds. Parallelisation would significantly improve the speed of the algorithm as well as scalability. Consequently, with modest computational resources, the labeled multi-Bernoulli filter can potentially scale to real time large scale tracking.

REFERENCES

- [1] M. Ansdell "Active Space Debris Removal: Needs, Implications, and Recommendations for Today's Geopolitical Environment," Technical Report, School of International Affairs, George Washington University, November 2010.
- [2] Y. Bar-Shalom and T. E. Fortmann, *Tracking and Data Association*. Academic Press, San Diego, 1988.
- [3] S. Blackman, *Multiple Target Tracking with Radar Applications*. Artech House, Norwood, 1986.
- [4] Committee for the Assessment of the US Air Forces Astrodynamics Standards, Aeronautics and Space Engineering Board, Division on Engineering and Physical Sciences, & National Research Council (US). *Continuing Kepler's Quest: Assessing Air Force Space Command's Astrodynamics Standards*. National Academies Press, 2012.
- [5] J. Dezert and Y. Bar-Shalom, "Joint probabilistic data association for autonomous navigation," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 29, no. 4, pp. 1275–1286, 10 1993.
- [6] K. DeMars, I. Hussein, M. Jah, and E. Scott "The Cauchy-Schwarz divergence for assessing situational information gain," *Int. Conf. Information Fusion*, pp. 1126–1133, Singapore, July 2012.
- [7] I. Hussein, K. DeMars, C. Fruh, R. Erwin, and M. Jah, "An AEGIS-FISST integrated detection and tracking approach to Space Situational Awareness," *Proc. Int. Conf. Information Fusion*, pp. 2065–2072, Singapore, July 2012.
- [8] I. Hussein, C. Frueh, and K. DeMars, "An AEGIS-FISST Algorithm for Multiple Object Tracking in Space Situational Awareness," *AIAA/AAS Astrodynamics Specialist Conference*, Minneapolis, August 2012.
- [9] J. A. Kennewell, B.-N. Vo, "An overview of space situational awareness," *Proc. Int. Conf. Information Fusion (FUSION)*, 1029–1036, 2013.
- [10] R. Mahler, "Multi-target Bayes filtering via first-order multi-target moments," *IEEE Trans. Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1152–1178, 2003.
- [11] R. Mahler, "PHD filters of higher order in target number," *IEEE Trans. Aerospace & Electronic Systems*, Vol. 43, No. 3, July 2007.
- [12] R. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Artech House, 2007.
- [13] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, "A consistent metric for performance evaluation of multi-object filters," *IEEE Trans. Signal Processing*, Vol. 56, no. 8, pp. 3447–3457, Aug. 2008.
- [14] B.-N. Vo, S. Singh, and A. Doucet, "Sequential Monte Carlo methods for multi-target filtering with random finite sets," *IEEE Trans. Aerospace and Electronic Systems*, vol. 41, no. 4, pp. 1224–1245, 2005.
- [15] B.-T. Vo, B.-N. Vo, and A. Cantoni, "The cardinality balanced multi-target multi-Bernoulli filter and its implementations," *IEEE Trans. Signal Processing*, Vol. 57, No. 2, pp. 409–423, Feb. 2009.

- [16] B.-N. Vo, B.-T. Vo, N.-T. Pham and D. Suter, "Joint detection and estimation of multiple objects from image observations," *IEEE Trans. Signal Procesing*, Vol. 58, No. 10, pp. 5129–5241, 2010.
- [17] B.-T. Vo, and B.-N. Vo, "Labeled random finite sets and multi-object conjugate priors," *IEEE Transactions on Signal Processing*, vol. 61, no. 13, pp. 3460–3475, 2013.
- [18] B.-N. Vo, B.-T. Vo, and D. Phung, "Labeled random finite sets and the Bayes multi-target tracking filter," *IEEE Transactions on Signal Processing*, (under review). [Online]. Available: <http://arxiv.org/abs/1312.2372>
- [19] S. Reuter, B.-T. Vo, B.-N. Vo, and K. Dietmayer, "The labeled multi-Bernoulli filter," *IEEE Transactions on Signal Processing*, Vol. 62, No. 12, pp. 3246–3260, 2014.