Sea Ice Tracking with a Spatially Indexed Labeled Multi-Bernoulli Filter

Jonatan Olofsson*, Clas Veibäck† and Gustaf Hendeby†

* Center for Autonomous Marine Operations and Systems (NTNU-AMOS)
Department of Engineering Cybernetics,
Norwegian University of Science and Technology, Trondheim, Norway.
jonatan.olofsson@ntnu.no

† Department of Automatic Control,
Linköping University, Linköping, Sweden.
Corresponding author: clas.veiback@liu.se

Abstract—In polar region operations, drift ice positioning and tracking is useful for both scientific and safety reasons. At its core is a Multi-Target Tracking (MTT) problem in which currents and winds make motion modeling difficult. One recent algorithm in the MTT field, employed in this paper, is the Labeled Multi-Bernoulli (LMB) filter. In particular, a proposed reformulation of the LMB equations exposes a structure which is exploited to propose a compact algorithm for the generation of the filter's posterior distribution. Further, spatial indexing is applied to the clustering process of the filter, allowing efficient separation of the filter into smaller, independent parts with lesser total complexity than that of an unclustered filter.

Many types of sensors can be employed to generate detections of sea ice, and in this paper a recorded dataset from a Terrestrial Radar Interferometer (TRI) is used to demonstrate the application of the Spatially Indexed Labeled Multi-Bernoulli filter to estimate the currents of an observed area in Kongsfjorden, Svalbard.

Figure 1: Example data from the Norut TRI dataset. The image shows backscattered intensity, measured by a terrestrial radar.

I. Introduction

Work in the polar regions of our planet is unavoidably linked with hazards such as drift ice. Increased presence fueled by economic interests in the Arctic, has for several decades [21] called for research in the field of Ice Management. A comprehensive overview of Ice Management in practical use is provided in [5]. The field deals with the detection, tracking and forecast of ice, but also the physical actions taken to avoid collisions [5]. While managing ice is of great importance to polar ventures, predicting ice movement has proven difficult [6], concluding that observations are essential for tracking.

For tracking sea ice movements, multiple sensors have been studied — such as satellite-carried Synthetic Aperture Radar (SAR) [18], Unmanned Aerial Systems (UAS's)[10, 9, 13] and, as studied in this paper, Terrestrial Radar Interferometer (TRI) [30].

At the core of tracking of individual sea ice objects from these data sources lies the problem of Multi-Target Tracking (MTT), inclu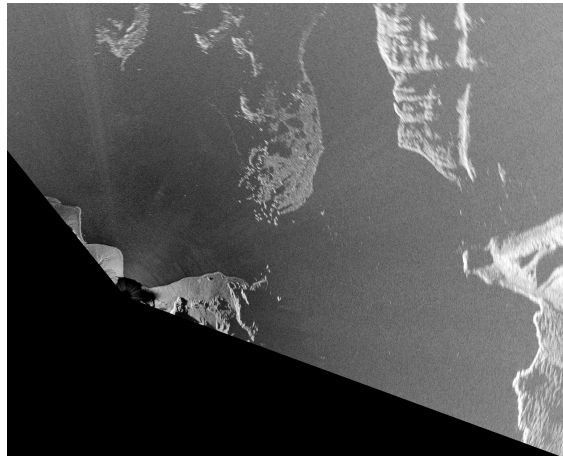ding the problem of associating incoming measurements to existing targets — or newly created ones. For this purpose we study the use of the Labeled Multi-Bernoulli (LMB) filter [23, 31] — presented in Section II — and also make a contribution with the first main result of the paper: a reformulation, in Section II-F, of the LMB equations to expose a structure which lends itself to an efficient algorithm for the filter implementation.

As background to the LMB presentation of Section II, the terminology is introduced in Section II-A. Following is a presentation of clustering and spatial indexing in Section II-B detailing their use for subdividing the filter correction into smaller parts, allowing the tracking algorithm to scale to hundreds of tracked targets. In Section III, we present the contribution of an Open Source implementation of the Spatially Indexed Labeled Multi-Bernoulli filter, which was implemented to demonstrate the aforementioned proposed algorithm.

The implementation is then applied, in the paper's

second main contribution, to a TRI dataset — exemplified in Figure 1 — provided by the research institute Norut. The application and dataset is described in Section IV, with results and conclusions presented in Section V and Section VI respectively.

II. The Labeled Multi-Bernoulli Filter

The LMB filter was proposed in [23] as a simplification of the δ-GLMB-filter [29, 28]. In this section, we review its general formulation, as well as its underlying algorithms and concepts.

A. Terminology and Notation

The terminology often employed in common Multi-Target Tracking (MTT) literature, and consequently here, is based on the following definitions.

At any given time instance, a sensor delivers a *scan* — an unordered but enumerated set of reports (measurements) detected at this instance. This set is exhaustive, i.e., the scan set Z_k contains all reports (indexed by i), $z_{k,i}$ for time k. In the following, time indices are dropped unless needed for clarity.

The purpose of MTT is to estimate a list of tracked *targets*. Each true target follows a "true track", the estimate of which is the result of filtering of data and assignments to the target. These, hypothetical, assignments can indicate association of the target to a specific report, but also a missed detection or even alternative motion models [12]. It is assumed that each true target will yield at most one report in each scan. Each track estimate starts with the event of a newly created target as hypothesized by a target birth model.

To accommodate the uncertainty of which report belongs to which target, the LMB filter considers multiple *hypotheses*. The definition for an hypothesis employed in this paper, is that of a set of hypothetical assignments which assigns each previously hypothesized target as either 1) existing and associated with a specific report in the incoming scan, 2) existing but missed in the scan ($z_\emptyset$ association), or 3) non-existent (indicated in equations by the letter F).

Each hypothesized target is assigned a label at its creation. The space of hypothetical associations (in the sense of 1) and 2) above) — denoted Θ_I — is parametrized by the set I of labeled targets assumed to be existent. All assignments in Θ_I assign all targets in I to either a specific report or as a missed detection. Thus, given a set $\mathcal{L}$ of labeled targets, $I \cap \mathcal{L}$ are considered as existing, and $\mathcal{L} - I$ as non-existing.

Example: Given the set of tracked labeled targets $\mathcal{L} = \{\ell_1, \ell_2, \ell_3\}$, hypothesis n assigns the target with label ℓ_1 to report z_1, presumes that ℓ_2 has been missed, and that ℓ_3 does not exist:

$$\theta_n = \{(\ell_1, z_1), \quad (\ell_2, z_\emptyset)\},$$
$$I = \{\ell_1, \ell_2\},$$
$$\mathcal{L} - I = \{\ell_3\}.$$

The space of hypothetical associations is, for each time step, combinatorially large. It can, however, be enumerated in order of decreasing probability and truncated to the k-best hypotheses [17]. For each hypothesis, it is assumed that each track can be modeled as a standard single-target estimation problem, conditioned on the validity of the hypothetical assignments, and that its evolution is described by the single-target transition density $f_k(x_{k,\ell}|x_{k-1,\ell}, \ell)$ for target ℓ at time step k.

Hypothetically, any report can be associated with any target, but in practice many of those associations will be very unlikely to be true. That is, targets and reports which are sufficiently unlikely to be associated are approximately independent. A *gating* [22, 2, 26] function can be defined to limit the number of assignments that are considered. Targets which, as decided through the gating function, no hypothesis will associate to the same reports can be treated in separate *clusters* — independent groups of targets which share ambiguous reports [22].

Notation: In the paper, notation common to the LMB literature is used [23], including the inner product,

$$\langle f, g \rangle \triangleq \int f(x)\, g(x)\, dx, \tag{1}$$

as well as the multi-object exponential notation

$$h^X \triangleq \prod_{x \in X} h(x) \tag{2}$$

($h^\emptyset = 1$ by convention).

Further, the Kronecker delta function is defined by

$$\delta_Y(X) \triangleq \begin{cases} 1, & \text{if } X = Y, \\ 0, & \text{otherwise,} \end{cases} \tag{3}$$

and is used to select summands relevant to exactly the set Y. Similarly, the inclusion function,

$$1_Y(X) \triangleq \begin{cases} 1, & \text{if } X \subseteq Y, \\ 0, & \text{otherwise,} \end{cases} \tag{4}$$

is used to select summands relevant to a set Y of which X is a subset. If X is singular, $X = \{x\}$, the notation $1_Y(x)$ is used. Further, the association indicator $A^\theta_{z \leftrightarrow \ell}$ is defined as

$$A^\theta_{z \leftrightarrow \ell} \triangleq \begin{cases} 1, & \text{if } \theta \text{ assigns label } \ell \text{ to report } z, \\ 0, & \text{otherwise.} \end{cases} \tag{5}$$

To exhaustively iterate all hypothetical report-target associations, the operator $\mathcal{F}(\mathbb{X})$ is used to denote the collection of all subsets of set $\mathbb{X}$.

B. Target-Report Matching and Clustering

The likelihood of association between a report and a target is based on the probability density of a true report being in the (estimated) target probability space [3]. In the case of Gaussian reports and targets, this is synonymous with the probability density function (PDF) of the innovation, itself a Gaussian. The area bounded

by any Gaussian iso-probability limit corresponds to an ellipse, the size and orientation of which can be calculated from its covariance matrix and a limit parameter, K [24].

Hence, gating through a limit of the association probability is referred to as ellipsoidal gating [3]. A report $\bar{z}$ with associated covariance R is positively gated if it fulfils the condition

$$\bar{z} \in \left\{ z \quad : \quad (z - \hat{z})^T (P + R)^{-1} (z - \hat{z}) \leq K \right\}, \quad (6)$$

where $\hat{z}$ and P represent the mean and covariance, respectively, of the target state transformed into measurement space. In [3], Collins and Uhlmann show that a necessary condition for (6) is that

$$\exists z : \quad \begin{cases} (z - \hat{z})^T P^{-1} (z - \hat{z}) \leq K \ \wedge \\ (z - \bar{z})^T R^{-1} (z - \bar{z}) \leq K, \end{cases} \quad (7)$$

that is, that the ellipses enscribed respectively by the covariances of the transformed state estimate and the report overlap. This motivates the use of intersections as a method for gating, and since the ellipsoids of the tracks and reports are independent, the gate equation in (6) does not need to be recomputed for each association pair.

To further simplify the gating, we use the axis-aligned bounding box of each of the ellipses in (7). Since the association probability at the square bounds is lower than or equal at the ellipsoid edge, the intersection of their bounding boxes is a conservative estimate of (7): if the bounding boxes do not overlap, neither do the ellipses.

The concept may be extended to other unimodal distributions, as well as mixtures by using the minimum bounding box of the desired probability limit.

From the matching of reports to targets, a connection graph can be formed by connecting pairs of targets which both are a potential match for a common report. The resulting graph will contain one or more groups of connected components which represent the targets that must be kept in the same cluster. Algorithms for finding graph connected components are studied in e.g. [27, 20].

The target labels in $\mathcal{L}_+ = \mathcal{L} \cup \mathcal{B}$ (the sets of old and newborn target labels) are therefore correspondingly partitioned into disjoint sets

$$\mathcal{L}_+ = \bigcup_{i=1}^{N} \mathcal{L}_+^{(i)}$$

with $\mathcal{L}_+^{(i)} \cap \mathcal{L}_+^{(j)} = \emptyset$ for $i \neq j$. Similarly, the set of current reports Z can be partitioned into the corresponding clusters

$$Z = Z^{(0)} \bigcup_{i=1}^{N} Z^{(i)}$$

with $Z^{(i)} \cap Z^{(j)} = \emptyset$ for $i \neq j$, and $Z^{(0)}$ is the set of measurements not associated with any previously known target.

C. Spatial Indexing

Unlike in, e.g., Multiple Hypothesis Tracking (MHT) filter, the targets in the LMB filter are not connected to each other through restrictions of track history compatibility — each target is independent following each update. This means that for an incoming scan it can easily be determined, on a target-level, which targets that fall in the sensor field-of-view and thus will be affected by the correction update, simply by looking at the PDF's of the current time-step. Similarly, the associations between reports is a novel matching process between the PDF's of each report and that of each target.

As previously noted, the area in which a report may match a target may be limited through gating, and any gate may be approximated by its axis-aligned bounding box (regardless of distribution). Such rectangular areas are well suited for fast intersection lookups, when stored in structures suitable for spatial indexing, such as the R-tree [8]. By storing the targets in a database indexed by their bounding-boxes, the affected targets of each scan can as such be efficiently loaded using bounding box intersections and grouped into clusters as per the previous section — leaving unaffected targets entirely unloaded from the database.

D. Labeled Multi-Bernoulli Filter Outline

The LMB filter is defined in the framework of Finite Set Statistics (FISST) [23], of which the Random Finite Set (RFS) is an integral part. An RFS is a set with a probabilistic cardinality distribution, i.e., each potential element is included in the set with a given probability. Specifically, a Bernoulli RFS is a random set which is empty with probability $1 - r$, and with probability r is a singleton. For an element x with probability $p(\cdot)$, the Bernoulli RFS PDF is given by

$$\pi(X) = \begin{cases} 1 - r, & \text{if } X = \emptyset, \\ r \cdot p(x), & \text{if } X = \{x\}. \end{cases} \quad (8)$$

A multi-Bernoulli RFS is the resulting set of the union of M independent Bernoulli-distributed random finite sets $X^{(i)}$, given by $X = \bigcup_{i=1}^{M} X^{(i)}$. Consequently, the multi-Bernoulli RFS is parametrized by the set $\left\{ (r^{(i)}, p^{(i)}) \right\}_{i=1}^{M}$, and its PDF is given by [15]

$$\pi(\{x_1, \ldots, x_n\}) = \prod_{j=1}^{M} \left(1 - r^{(j)} \right)$$

$$\times \sum_{1 \leq i_1 \neq \cdots \neq i_n \leq M} \prod_{j=1}^{n} \frac{r^{(i_j)} p^{(i_j)}(x_j)}{1 - r^{(i_j)}}. \quad (9)$$

The LMB RFS is obtained by the augmentation of each Bernoulli RFS with a unique label, $\ell \in \mathcal{L}$. The RFS can thus be described by the set

$$\left\{ \left(r^{(\ell)}, p^{(\ell)} \right) \right\}_{\ell \in \mathcal{L}}.$$

As this set fully describes a multi-target probability density, $\pi(X)$, the shorthand notation $\pi = \left\{\left(r^{(\ell)}, p^{(\ell)}\right)\right\}_{\ell \in \mathcal{L}}$ will be used in the following.

The LMB filter follows the classical predict/correct filter recursion, each step outlined below.

1) LMB Prediction: Given an LMB PDF $\pi(X)$, the prediction step of the LMB filter and the updated distribution, $\pi_+(X)$, is obtained by the application of the standard prediction update of a Bayesian filter, the Chapman-Kolmogorov equation,

$$\pi_+(X_+) = \int f(X_+)\,\pi(X)\,\delta X, \tag{10}$$

extended for an RFS as defined in [29]. This gives the following set of surviving and new-born targets [23],

$$\pi_+ = \left\{\left(r^{(\ell)}_{+,S}, p^{(\ell)}_{+,S}\right)\right\}_{\ell \in \mathcal{L}} \cup \left\{\left(r^{(\ell)}_B, p^{(\ell)}_B\right)\right\}_{\ell \in \mathcal{B}}, \tag{11}$$

where

$$r^{(\ell)}_{+,S} = \eta_S(\ell)\,r^{(\ell)}, \tag{12}$$

$$p^{(\ell)}_{+,S} = \frac{\langle p_S(\cdot,\ell)\,f(x|\cdot,\ell), p(\cdot|\ell)\rangle}{\eta_S(\ell)}, \tag{13}$$

$$\eta_S(\ell) = \langle p_S(\cdot,\ell), p(\cdot,\ell)\rangle, \tag{14}$$

$p_S(\cdot,\ell)$ is the distribution of target survival probability and $f(x|\cdot,\ell)$ is the transition density. The set $\left\{\left(r^{(\ell)}_B, p^{(\ell)}_B\right)\right\}_{\ell \in \mathcal{B}}$ is given by the birth model, further discussed in Section II-E.

2) LMB Correction: Drawn from the update of the δ-GLMB [29], the LMB correction is derived in [23].

In general, as noted in [23], the LMB distribution is not closed under the Bayesian filter update. However, the resulting δ-GLMB distribution — which can represent multiple disjoint hypotheses — may be approximated as in (16) with an LMB PDF through the collapse of its hypotheses, weighted by their probabilities. That is, the correction updates the set

$$\pi_+ = \left\{\left(r^{(\ell)}_+, p^{(\ell)}_+\right)\right\}_{\ell \in \mathcal{L}_+} \tag{15}$$

by the following approximation:

$$\pi(\cdot|Z) \approx \left\{\left(r^{(\ell)}, p^{(\ell)}\right)\right\}_{\ell \in \mathcal{L}_+}$$
$$= \bigcup_{i=1}^N \left\{\left(r^{(\ell,i)}, p^{(\ell,i)}\right)\right\}_{\ell \in \mathcal{L}^{(i)}_+}, \tag{16}$$

in which parameters are given by

$$r^{(\ell,i)} = \sum_{(I_+,\theta)\in\mathcal{F}\left(\mathcal{L}^{(i)}_+\right)\times\Theta_{I_+}} w^{(I_+,\theta)}\left(Z^{(i)}\right)1_{I_+}(\ell), \tag{17}$$

$$p^{(\ell,i)}(x) = \frac{1}{r^{(\ell,i)}}\sum_{(I_+,\theta)\in\mathcal{F}\left(\mathcal{L}^{(i)}_+\right)\times\Theta_{I_+}} w^{(I_+,\theta)}\left(Z^{(i)}\right)$$
$$\times 1_{I_+}(\ell)\,p^{(\theta)}\left(x,\ell|Z^{(i)}\right), \tag{18}$$

where Θ_{I_+} is the space of mappings of tracks $\theta : I_+ \to \{0,1,\ldots,|Z^{(i)}|\}$ such that $\theta(\iota) = \theta(\iota') > 0$ implies that

$\iota = \iota'$, i.e., the mapping is unique for all values except those mapped to zero [23]. Also, for $I_+ \subseteq \mathcal{L}^{(i)}_+$

$$w^{(I_+,\theta)}\left(Z^{(i)}\right) \propto w^{(I_+)}_{+,i}\left[\eta^{(\theta)}_{Z^{(i)}}\right]^{I_+}, \tag{19}$$

$$w^{(I_+)}_{+,i} = \prod_{\ell\in\mathcal{L}^{(i)}_+ - I_+}\left(1 - r^{(\ell)}_+\right)\prod_{\ell'\in I_+} r^{(\ell)}_+, \tag{20}$$

$$\eta^{(\theta)}_{Z^{(i)}}(\ell) = \left\langle p^{(\ell,i)}_+(x), \psi_{Z^{(i)}}(\cdot,\ell;\theta)\right\rangle, \tag{21}$$

$$\psi_{Z^{(i)}}(x,\ell;\theta) = \begin{cases} \frac{p_D(x,\ell)p_G\, g\left(z_{\theta(\ell)}|x,\ell\right)}{\kappa\left(z_{\theta(\ell)}\right)}, & \theta(\ell)\neq z_\emptyset, \\ q_{D,G}(x,\ell), & \theta(\ell) = z_\emptyset, \end{cases} \tag{22}$$

$$q_{D,G}(x,\ell) = 1 - p_D(x,\ell)\,p_G, \tag{23}$$

$$p^{(\theta)}\left(x,\ell|Z^{(i)}\right) = \frac{p^{(\ell,i)}_+(x)\,\psi_{Z^{(i)}}(x,\ell;\theta)}{\eta^{(\theta)}_{Z^{(i)}}(\ell)}, \tag{24}$$

where p_G is the gating probability, $g\left(z_{\theta(\ell)}|x,\ell\right)$ is the likelihood and $\kappa\left(z_{\theta(\ell)}\right)$ is the (Poisson) clutter intensity [23]. Note that while the δ-GLMB representation in [23] is an intermediate representation in the theoretical derivation of the filter, its construction in implementation is not necessary to reach the collapsed LMB representation of (16).

To calculate the weight from (19) for an hypothesis we start by making the distinction between associated and non-associated targets by splitting the hypothesis label set I_+:

$$I^a_+ = \{\ell: \quad \theta(\ell)\neq z_\emptyset\}_{\ell\in I_+}, \tag{25}$$

$$I^n_+ = \{\ell: \quad \theta(\ell) = z_\emptyset\}_{\ell\in I_+}, \tag{26}$$

(implying $I_+ = I^a_+ \cup I^n_+$ and $I^a_+ \cap I^n_+ = \emptyset$). We can then rewrite (19)–(23) as

$$w^{(I_+,\theta)}\left(Z^{(i)}\right) \propto w^{(I_+)}_{+,i}\left[\eta^{(\theta)}_{Z^{(i)}}\right]^{I_+}$$
$$= \prod_{\ell\in\mathcal{L}^{(i)}_+ - I_+}\left(1 - r^{(\ell)}_+\right) \tag{27}$$
$$\times \prod_{\ell'\in I^a_+} r^{(\ell')}_+ \eta^{(\theta)}_{Z^{(i)}}(\ell')\prod_{\ell''\in I^n_+} r^{(\ell'')}_+ \eta^{(\theta)}_{Z^{(i)}}(\ell''),$$

The product of (27) can be efficiently expressed using the Negative Log Likelihoods (NLL's), Λ_ℓ;

$$e^{-\Lambda_\ell} = \begin{cases} 1 - r^{(\ell)}_+, & \text{if } \ell\in\mathcal{L}^{(i)}_+ - I_+, \\ r^{(\ell)}_+\eta^{(\theta,a)}_{Z^{(i)}}(\ell), & \text{if } \ell\in I^a_+, \\ r^{(\ell)}_+\eta^{(\theta,n)}_{Z^{(i)}}(\ell), & \text{if } \ell\in I^n_+, \end{cases} \tag{28}$$

yielding

$$w^{(I_+,\theta)}\left(Z^{(i)}\right) \propto \exp\left(-\sum_{\ell\in\mathcal{L}^{(i)}_+}\Lambda_\ell\right). \tag{29}$$

Each hypothesis (I_+,θ) is generated for each cluster in order of probability using Murty's algorithm [17, 16], to create a truncated approximation of the full sums of (17), (18). The truncation is achieved through the termination

of Murty's algorithm based on either a maximum number of drawn hypotheses, or a minimum hypothesis probability.

In the context of LMB, the cost matrix C for Murty's algorithm is created for each target cluster, from the NLLs of the track assignments outlined in (28), resulting in a matrix exemplified by

$$C = \begin{pmatrix} {}^{z_1}\Lambda_{\ell_1} & {}^{z_2}\Lambda_{\ell_1} & {}^{n}\Lambda_{\ell_1} & \infty & {}^{F}\Lambda_{\ell_1} & \infty \\ {}^{z_1}\Lambda_{\ell_2} & {}^{z_2}\Lambda_{\ell_2} & \infty & {}^{n}\Lambda_{\ell_2} & \infty & {}^{F}\Lambda_{\ell_2} \end{pmatrix}, \quad (30)$$

for a two-track hypothesis and a two-report scan (with costs for each target being assocated, non-associated and false targets respectively).

E. Adaptive Birth Model

To include new targets in the tracker, the LMB filter relies on a birth distribution. Different birth-models has been discussed in e.g. [31], but here, following [23], the selected birth model for time $k+1$ is based on the reports of time k:

$$\pi_{B,k+1} = \left\{ \left(r_B^{(\ell)}, p_B^{(\ell)} \right) \right\}_{\ell \in \mathcal{B}_k} \quad (31)$$

for new labels in $\mathcal{B}_k$ generated for each report in Z_k.

The existence probabilities of new targets in this model are proportional to the probability of the report not being associated with any previously known target. For a report, the association probability is given by

$$r_{U,k}(z) = \sum_{(I_+,\theta) \in \mathcal{F}\left(\mathcal{L}_+^{(i)}\right) \times \Theta_{I_+}} w^{(I_+,\theta)}\left(Z^{(i)}\right) 1_\theta(z). \quad (32)$$

Given an expected number of new targets in each scan, $\lambda_{B,k+1}$, the existence probability of new targets is thus given by

$$r_{B,k+1}(z) = \min\left(r_B^{\max}, \frac{\left(1 - r_{U,k}(z)\right) \cdot \lambda_{B,k+1}}{\sum_{\xi \in Z_k} 1 - r_{U,k}(\xi)} \right). \quad (33)$$

Note that, for $\lambda_{B,k+1} > 1$, the existence probability may need to be limited by the min()-clause to a maximum value of $r_B^{\max} \leq 1$.

F. Reformulation

The classic LMB filter formulation of [23] carries a heritage from the δ-GLMB implementation, leading to the necessity of the artificial reconstruction of the δ-GLMB distribution using the k-shortest-path algorithm in the correction update stage. Here we propose a formulation that does not require the δ-GLMB reconstruction due to its immediate collapse into the LMB-distribution approximation.

In any hypothetical association, a target may be assigned as either associated to a specific report, missed, or assumed non-existent. In the following, being missed is defined as being associated with the null report, $z_\emptyset$. However, unlike for standard reports, any number of targets may be assigned to $z_\emptyset$. This leads to the following definition:

$$Z^\dagger = Z \cup \{z_\emptyset\}. \quad (34)$$

With this definition, we note that $p^{(\theta)}(x, \ell|Z)$ in (18), with the cluster index omitted for notational convenience, belongs to a limited set:

$$p^{(\theta)}(x, \ell|Z) \in \left\{ p^{(\ell)}(x|z) \right\}_{z \in Z^\dagger}. \quad (35)$$

Thus, the sums in (17)–(18) may be partitioned as follows, abbreviating with $w^\theta = w^{(I_+,\theta)}\left(Z^{(i)}\right)$ and denoting the inner sums as ${}^z w_\ell$. Again, $A_{z \leftrightarrow \ell}^\theta$ is the indicator function for hypothesis θ assigning report z (or $z_\emptyset$) to label ℓ (also implying the inclusion function $1_{I_+}(\ell)$).

$$r^{(\ell)} = \sum_{z \in Z^\dagger} \left[\sum_{(I_+,\theta) \in \mathcal{F}(\mathbb{L}_+) \times \Theta_{I_+}} w^\theta A_{z \leftrightarrow \ell}^\theta \right]$$
$$= \sum_{z \in Z^\dagger} {}^z w_\ell, \quad (36)$$

$$p^{(\ell)}(x) = \frac{1}{r^{(\ell)}} \sum_{z \in Z^\dagger} \left[\sum_{(I_+,\theta) \in \mathcal{F}(\mathbb{L}_+) \times \Theta_{I_+}} w^\theta A_{z \leftrightarrow \ell}^\theta \right] p^{(\ell)}(x|z)$$
$$= \frac{1}{r^{(\ell)}} \sum_{z \in Z^\dagger} {}^z w_\ell p^{(\ell)}(x|z). \quad (37)$$

We see that ${}^z w_\ell$ corresponds to *the sum of weights of all hypotheses that assign report z to label ℓ*. The outer sum includes all reports (and the null report), altogether covering the same summands as the original sums.

Further, the birth model of (32) may be rewritten as follows (for time index k):

$$r_{U,k}(z) = \sum_{\ell \in L_+^{(i)}} \left[\sum_{(I_+,\theta) \in \mathcal{F}\left(L_+^{(i)}\right) \times \Theta_{I_+}} w^\theta A_{z \leftrightarrow \ell}^\theta \right] \quad (38)$$
$$= \sum_{\ell \in L_+^{(i)}} {}^z w_\ell. \quad (39)$$

To exploit this reformulation, consider a cluster of N targets and M reports, and an $N \times (M + 2)$ matrix W. Further, consider a hypothesis assignment mapping $R_\theta(i)$ to be used for mapping each row index of W (corresponding to a target) to a column index (corresponding to an assignment). I.e., for all known targets (rows), $R_\theta(i)$

1) maps associated targets to its report's integer position in an ordered enumeration of the reports,
2) maps missed targets to the integer index $M + 1$ and
3) maps false targets to the integer index $M + 2$.

Using this mapping, Algorithm 1 will work on the assignments θ and the hypothesis score w^θ (of (29)), drawn from each iteration of Murty's algorithm, to readily form the relevant sums of the LMB algorithm. Note that the addition to the matrix in the algorithm adds the value w^θ once to each row in the column specified by $R_\theta(i)$.

The result is exemplified in (40) for a problem of three targets and two reports.

Recalling the definitions of $r^{(\ell)}$ and $r_{U,k}(z)$ from (36) and (38), respectively, we see that these correspond to

Algorithm 1 Weight matrix

$W \leftarrow N \times (M + 2)$ zero matrix.
$s \leftarrow 0$
for $(w^\theta, \theta) \in \mathrm{murty}(C)$ **do**
  $W[i, R_\theta(i)] \leftarrow W[i, R_\theta(i)] + w^\theta, \ \forall i \in \{1, \ldots, N\}$
  $s \leftarrow s + w^\theta$
end for
$W \leftarrow \frac{W}{s}$

---

the column (excluding the "False" column) and row sum respectively, as illustrated in (40).

$$W = \begin{pmatrix} {}^{z_1}w_{\ell_1} & {}^{z_2}w_{\ell_1} & {}^{\emptyset}w_{\ell_1} & {}^F w_{\ell_1} \\ {}^{z_1}w_{\ell_2} & {}^{z_2}w_{\ell_2} & {}^{\emptyset}w_{\ell_2} & {}^F w_{\ell_2} \\ {}^{z_1}w_{\ell_3} & {}^{z_2}w_{\ell_3} & {}^{\emptyset}w_{\ell_3} & {}^F w_{\ell_3} \end{pmatrix}. \quad (40)$$

Thus, with this proposed reformulation and algorithm, it is possible to easily extract the existence probabilities. Moreover, through (37), this formulation clarifies that the target PDFs may be attained simply through the weighted sum of $M + 1$ PDFs, instead of one per hypothesis.

III. Implementation

This section introduces the implementation of the Spatially Indexed LMB filter implemented for this article. The implementation, written in the Python language, is available under a Free and Open-Source Software (FOSS) license at https://github.com/jonatanolofsson/lmb.

A few remarks regarding the implementation;

Particle Implementation For now, the implementation is based on particle filter distributions, where the general equations of (21)–(24) are specialized as in [28].

Target Storage and Indexing In the implementation, targets are serialized post-update and stored in an SQLite database. The database is R-tree indexed based on the axis-aligned bounding boxes of each target PDF, allowing for fast extraction of relevant targets in the gating process.

IV. Sea Ice Tracking

In April 2016, the partners of the Norwegian Centre for Integrated Remote Sensing and Forecasting for Arctic Operations — CIRFA — conducted a field campaign in Kongsfjorden on Svalbard. One of the datasets collected is from a GAMMA Portable Radar Interferometer (GPRI). As an application, we study the extraction of ice detections from this data and its tracking with the LMB filter.

A. Pre-processing

Each scan is delivered as the intensity of the response along range and azimuth transformed into Cartesian coordinates. Detections are extracted from the raw signal in a pre-processing step. The aim is to process the signal such that the resulting detections are as likely as possible to be generated by true sea ice. Considering the format of the signal the pre-processing step mainly constitutes of the heuristic application of standard image processing methods.

Two types of ice are considered; 1) large regions of stationary sea ice with high signal-to-noise ratio that can be segmented independently for each scan and 2) drift ice with low signal-to-noise ratio that require pre-processing over several scans for detection. A land mask is applied to the image to ensure that detections are only obtained in water regions. The segmentation methods are presented below.

1) Detection of Stationary Sea Ice: Large areas of stationary ice are visible in the water, in particular in proximity to land. However, due to speckle noise and varying intensity over the image a simple threshold results in poor performance. To improve the signal-to-noise ratio, a sequence of standard image segmentation methods [7] are applied.

The stationary regions are first stabilised by averaging the ten latest frames. A Gaussian filter and a top-hat transform are then applied to reduce speckle noise and intensity variations in the image, respectively [7]. The segmentation of the stationary ice regions is performed using Otsu's optimal thresholding [19]. To further reduce noise in the segmentation, closing and hole filling are also applied [7]. The connected components adjacent to land are finally extracted as detections of stationary sea ice regions.

2) Detection of Drift Ice: Drift ice is generally small and often difficult to distinguish in the speckle noise. A method from the image processing field that can be employed is to adaptively estimate a background model in water regions. The background for each pixel is modelled as a mixture of Gaussian distributions [11, 25]. A simplified expectation-maximization method [4] is then used to continuously estimate the means and covariances in the model over time. For an incoming scan, all pixels that are significantly different from their background models are segmented as foreground. This implies many false detections, which is mitigated by extracting connected components of a minimum size of $150\,\mathrm{m}^2$. The reports are then obtained as the centroid of each connected component.

B. Model

The drift ice is modelled as having a nearly constant velocity subject to zero-mean white-noise acceleration. The states in the model are position and velocity in two dimensions. The model is derived by discretizing the continuous-time nearly-constant-velocity model [14]. The sensor is modelled to directly measure the position of the drift ice with zero-mean Gaussian noise.

The sampling time of the motion model is $180\,\mathrm{s}$. The motion model covariance parameter is chosen in both dimensions as $1.7 \times 10^{-5}\,\mathrm{m}^2/\mathrm{s}^3$ and the measurement

(a) Tracks after 2 h. (b) Tracks after 4 h. (c) Tracks after 7 h.
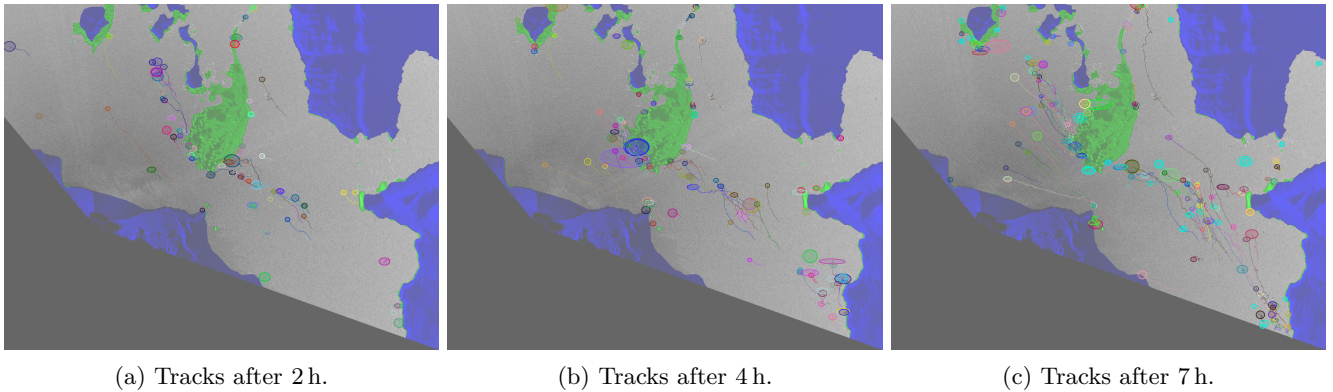
Figure 2: Drift ice tracks over time, showing the land mask in blue and stationary detections in green. Tracks and targets retain an individual randomly assigned color over time.

noise standard deviation is chosen as 12.2 m in each dimension.

C. Estimation of Currents

Apart from tracking individual sea ice objects for collision avoidance, the movements of drift ice can also be used to estimate the water currents in the region over time. A Gaussian kernel smoother [1] is applied to a grid over the water region. The position and time of the drift ice, obtained from the tracks, are used to estimate the local velocity, which depends on the current. The length scale is chosen at 750 m for the distance component and 1 h for the time component. The grid spacing used is 600 m in the water. Note that the method can easily be adapted to an on-line method by considering only causal measurements.

Figure 3: Estimation of currents in the water region.

V. Results

Following the procedure outlined in Section IV, the LMB implementation was used to track the ice movements in Kongsfjorden, Svalbard, over a period of seven hours, with scans delivered every three minutes. In Figure 2, we see the tracks of tracked ice objects build up over time.

The stationary sea ice, shown in green, changes only slightly over the course of the experiment, suggesting it would remain largely undetected if treated as drift ice. The detections of drift ice, shown in red, suffer many false alarms, but the LMB filter manages to confirm the targets, shown as ellipses, and maintain their tracks, shown as lines, over large stretches of water.

Looking at the estimated tracks from the LMB filter, there is a trend for the drift ice to move in the north-west direction. This is further corroborated in Figure 3 where the currents are estimated from the velocity estimates of the targets.

As can be seen in Figure 2, each scan consists of hundreds of reports, approximately 150 to 400, and about 50 to 190 targets are being maintained over time by the LMB filter. However, the low sampling rate of the TRI allows the filter to run in real-time, despite such conditions.
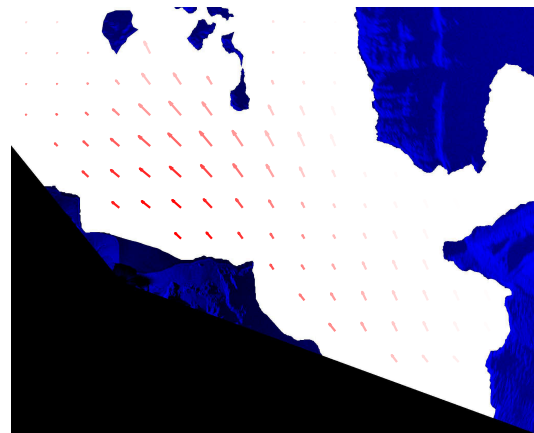
VI. Conclusions

In this paper we presented a simplifying reformulation of the Labeled Multi-Bernoulli filter, which highlights that the LMB filter largely can be reduced to a summation of hypothesis weights according to the proposed algorithm. This is exploited in an algorithm which simultaneously calculates both the update weights and the weights for the adaptive birth model. Further, we presented its implementation in which also the proposed application of spatial indexing aids in the on-line formation of target clusters. This clustering is critical for the scalability of the LMB algorithm, and for the continuation of the studied application of tracking a large number of ice objects. Future work also include further study and comparison of e.g. Gaussian fields for the current estimation, from which also uncertainty measures can be extracted.

The current information we extract from this tracking scenario could potentially be used for prediction and modeling. However, this might cause information looping which would need to be studied further prior to a practical application. The information carried in the LMB filter is essentially the same as other MTT-filters — some of which might be preferred for application-

specific advantages in specific cases. In related experiments, Global Nearest Neighbour (GNN) tracking has been used to verify, with good results, the LMB solution studied in this paper. Note that GNN corresponds to limiting the LMB algorithm to use only the single best hypothesis, and as such is expected to perform faster but less robust to incorrect associations.

For the specific task of current estimation (without positioning) one may also compare the results with techniques such as optical flow. This is, however, outside the scope of this paper.

ACKNOWLEDGMENTS

REFERENCES

[1] C. M. Bishop, *Pattern Recognition and Machine Learning.* New York, NY, USA: Springer, 2006.

[2] S. S. Blackman, "Multiple Hypothesis Tracking For Multiple Target Tracking," *IEEE Aerospace and Electronic Systems Magazine*, vol. 19, no. 1 January, 2004.

[3] J. Collins and J. Uhlmann, "Efficient gating in data association with multivariate Gaussian distributed states," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 28, no. 3, pp. 909 – 916, 1992.

[4] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the Royal Statistical Society, Series B*, vol. 39, no. 1, pp. 1–38, 1977.

[5] K. Eik, "Review of Experiences within Ice and Iceberg Management," *Journal of Navigation*, vol. 61, no. 04, p. 557, 2008.

[6] ——, "Iceberg drift modelling and validation of applied metocean hindcast data," *Cold Regions Science and Technology*, vol. 57, no. 2-3, pp. 67–90, 2009.

[7] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, 3rd ed. Upper Saddle River, NJ, USA: Pearson Prentice Hall, 2008.

[8] A. Guttman, "R-trees: A Dynamic Index Structure for Spatial Searching," in *Proceedings of the 1984 ACM SIGMOD International Conference on Management of Data*, Boston, Massachusetts, USA, 1984, pp. 47–57.

[9] J. Haugen, "Autonomous Aerial Ice Observation," Ph.D. dissertation, Norwegian University of Science and Technology, 2014.

[10] T. A. Johansen and T. Perez, "Unmanned Aerial Surveillance System for Hazard Collision Avoidance in Autonomous Shipping," 2015.

[11] P. Kaewtrakulpong and R. Bowden, "An improved adaptive background mixture model for real-time tracking with shadow detection," in *Proceedings of 2nd European Workshop on Advanced Video Based Surveillance Systems*, Sep. 2001.

[12] T. Kurien, "Issues in the Design of Practical Multitarget Tracking Algorithms," in *Multitarget-Multisensor Tracking: Advanced Applications*, Y. Bar-Shalom, Ed. Artech House, 1990, pp. 43–83.

[13] F. S. Leira, "Object Detection and Tracking With UAVs," Ph.D. dissertation, Norwegian University of Science and Technology, 2017.

[14] X. R. Li and V. P. Jilkov, "Survey of maneuvering target tracking. part i. dynamic models," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1333–1364, Oct 2003.

[15] R. P. S. Mahler, *Statistical Multisource-Multitarget Information Fusion.* Norwood, MA, USA: Artech House, Inc., 2007.

[16] M. L. Miller, H. S. Stone, and I. J. Cox, "Optimizing murty's ranked assignment method," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 33, no. 3, pp. 851–862, 1997.

[17] K. G. Murty, "An Algorithm for Ranking all the Assignments in Order of Increasing Cost," *Operations Research*, vol. 16, no. 3, pp. 682–687, 1968.

[18] J. Olofsson, E. Brekke, T. I. Fossen, and T. A. Johansen, "Spatially Indexed Clustering for Scalable Tracking of Remotely Sensed Drift Ice," in *IEEE Aerospace Conference Proceedings*, Big Sky, MT, USA, 2017.

[19] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 9, no. 1, pp. 62–66, Jan. 1979.

[20] D. J. Pearce, "An Improved Algorithm for Finding the Strongly Connected Components of a Directed Graph," Victoria University, Wellington, NZ, Tech. Rep., 2005.

[21] C. Randell, F. Ralph, D. Power, and P. Stuckey, "OTC 20264 Technological Advances to Assess , Manage and Reduce Ice Risk in Northern Developments," in *Offshore Tecdhnology Conference*, 2009.

[22] D. B. Reid, "An algorithm for tracking multiple targets," in *1978 IEEE Conference on Decision and Control including the 17th Symposium on Adaptive Processes*, vol. 17, no. 6, San Diego, California, USA, 1978, pp. 843–854.

[23] 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.

[24] M. I. Ribeiro, "Gaussian Probability Density Functions : Properties and Error Characterization," Institute for Systems and Robotics, Lisboa, Portugal, Tech. Rep. February, 2004.

[25] C. Stauffer and W. E. L. Grimson, "Adaptive background mixture models for real-time tracking," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2, Ft. Collins, CO, USA, Jun. 1999, pp. 2246–2252.

[26] J. J. Stein and S. S. Blackman, "Generalized Correlation of Multi-Target Track Data," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-11, no. 6, November, pp. 1207–1217, 1975.

[27] R. Tarjan, "Depth-first Search and Linear Graph Algorithms," *SIAM Journal on Computing*, vol. 1, no. 2, pp. 146–160, 1972.

[28] 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*, vol. 62, no. 24, pp. 6554–6567, 2014.

[29] 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.

[30] D. Voytenko, T. H. Dixon, M. E. Luther, C. Lembke, I. M. Howat, and S. de la Peña, "Observations of inertial currents in a lagoon in southeastern Iceland using terrestrial radar interferometry and automated iceberg tracking," *Computers and Geosciences*, vol. 82, pp. 23–30, 2015.

[31] J. L. Williams, "Marginal multi-bernoulli filters: RFS derivation of MHT, JIPDA, and association-based member," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 51, no. 3, pp. 1664–1687, 2015.