Algorithms for the multi-sensor assignment problem in the δ-generalized labeled multi-Bernoulli filter

Jun Ye Yu, Augustin-Alexandru Saucan, Mark Coates, Michael Rabbat
Electrical and Computer Engineering, McGill University, Montreal, Quebec
Email: jun.y.yu, augustin.saucan@mail.mcgill.ca and mark.coates, michael.rabbat@mcgill.ca

***Abstract*—Previous adaptations of the δ-generalized labeled multi-Bernoulli (δ-GLMB) filter to the multi-sensor case involve the sequential application of the update step for each sensor or Gibbs sampling for multi-sensor data association. The practical usage of the sequential δ-GLMB filter is limited due to the number of hypotheses growing with each additional sensor. Similarly, the Gibbs method requires a large number of samples for each hypothesis. In this paper, in the aim of finding the optimal or near-optimal multi-sensor assignments, we propose two novel methods, the combination and the cross entropy methods. Numerical simulations are conducted to evaluate the proposed multi-assignment methods together with the standard sequential processing method and a stochastic optimization algorithm based on Gibbs sampling. The combination method is able to significantly reduce running time with respect to the sequential method while yielding competitive performance across a wide range of test scenarios.**

I. Introduction

The objective of multi-target tracking is to infer the target tracks in addition to estimating the number of targets and their kinematic states; but non-uniform detection probability, measurement origin uncertainty, false detection and target birth/death are all difficult obstacles to solving the problem.

Random finite set (RFS) filters [1] have emerged as a popular paradigm for solving the multi-target tracking problem in the Bayesian framework. Since the exact multi-target Bayes filter is computationally intractable, the *probability hypothesis density* (PHD) filter [2], *cardinalized PHD* (CPHD) [3] filter and *multi-Bernoulli* (MB) filter [4] have been proposed as tractable approximations, although they do not provide target tracks over time. Vo et al. [5] introduced the notion of a labeled RFS in which unique labels are appended to the elements of the RFS to identify their targets (and their estimates) across time and hence infer target tracks. They subsequently developed the *δ-generalized multi-Bernoulli* (δ-GLMB) density [5] and a single-sensor tracker based on δ-GLMB RFSs [6].

While all these filters have been initially developed for single-sensor tracking, multi-sensor extensions have also been proposed [7]–[14]. In the iterator-corrector PHD filter [7], [8], each sensor's measurements are processed sequentially and the output from one sensor is used as input for the next sensor. The recent work by Papi [10] presents the multi-sensor extension of the δ-GLMB filter, and its implementation also involves iterating through each sensor. Liu et al. [11] use an extended association table to generate the most likely associations between targets and measurements from all sensors, but no simulations are provided to validate the algorithm's performance. Both the multi-sensor CPHD filter [12] and the multi-sensor multi-Bernoulli filter [13] process all sensor measurements simultaneously by using a greedy algorithm to select the most likely associations. Vo et al. [14] have recently applied Gibbs sampling to find a number of likely multi-sensor assignments in the δ-GLMB filter.

The exact implementation of the multi-sensor δ-GLMB filter requires enumerating all multi-sensor assignments to compute the posterior multi-target density. In practice, we look for a number of likely multi-sensor assignments to construct a truncated density. Although the problem of finding the T best single-sensor assignments can be solved efficiently using Murty's algorithm [15], the multi-sensor counterpart is NP-hard [16]. In this paper, we present two approximation algorithms, the combination and the cross entropy methods, that yield a number of likely multi-sensor assignments without exhaustive enumeration. We compare their performance to the Gibbs method [14] and the sequential processing method [10]. The combination method first solves the assignment problem locally at each sensor and then combines the locally optimal solutions to form high-scoring multi-sensor assignments. The cross entropy method constructs a distribution on the space of all multi-sensor assignments with higher probability for more likely assignments. The algorithms' performances are compared via simulations and it is shown that the combination method greatly reduces computational time with respect to the sequential method while yielding near-optimal assignments.

II. Background

The study of finite set statistics [1] has led to the development of multi-target filters in the Bayesian framework where the target states are modeled as random finite sets. An RFS is a finite set with random cardinality and elements, and thus it conveniently captures the two unknown quantities of interest in multi-target tracking problems: the number of targets and their states. A labeled RFS appends a unique label to each element in the RFS. Elements with the same label correspond to the same target and allow the formulation of target tracks.

For the rest of the paper, unlabeled single-object states are denoted by lower case letters (e.g., x) and multi-object states (realizations of an RFS) by upper case letters (e.g., X). Their labeled counterparts are bold letters (e.g., $\mathbf{x} = (x, l)$, $\mathbf{X} = \{\mathbf{x}_1, ...\mathbf{x}_n\}$). The blackboard bold letters (e.g., $\mathbb{X}$, $\mathbb{L}$) denote the corresponding state space. The projection function $\mathcal{L}(\mathbf{X}) = \{l, (x, l) \in \mathbf{X}\}$ returns the labels of a labeled RFS. The distinct

label indicator function $\Delta(\mathbf{X})$ returns 1 if all labels in the RFS are distinct and 0 otherwise. For a real-valued function $h(x)$ and a set $\mathbf{X}$, the multi-object exponential is defined as $[h]^{\mathbf{X}} = \prod_{x\in\mathbf{X}} h(x)$. Finally, we adopt a generalized Kronecker delta function $\delta_Y(X)$ where X and Y can be scalars, vectors, or unordered sets.

The δ-GLMB RFS is a labeled RFS with state space $\mathbb{X}$ at time k and has the following distribution [5]:

$$\pi_{k|k}(\mathbf{X}) = \Delta(\mathbf{X}) \sum_{(I,\epsilon)\in\mathcal{F}(\mathbb{L}_{0:k})\times\Xi} w_{k|k}^{(I,\epsilon)} \delta_I(\mathcal{L}(\mathbf{X}))[p_{k|k}^{(\epsilon)}]^{\mathbf{X}}, \quad (1)$$

where $\mathbb{L}_{0:k}$ is the label space of targets that exist from time 0 to k, $\mathcal{F}(\mathbb{L}_{0:k})$ is the space of all finite subsets of $\mathbb{L}_{0:k}$, Ξ is a discrete space of target-measurement association histories, and $w_{k|k}^{(I,\epsilon)}$ are normalized weights. Each I represents a set of track labels, and ϵ represents a history of measurement-target associations. Each pair (I,ϵ) thus represents the hypothesis that the labels I have the measurement association history ϵ and hypothesis weight $w_{k|k}^{(I,\epsilon)}$. The term $p_{k|k}^{(\epsilon)}(\cdot,l)$ is the single-target density of label $l \in I$ given ϵ.

The δ-GLMB filter is a two-step Bayesian filter which models the multi-object target states as δ-GLMB RFSes.

1) Predict: If the posterior multi-object distribution at time k is a δ-GLMB of the form (1), then the predicted multi-object distribution is also a δ-GLMB of the following form [5]:

$$\pi_{k+1|k}(\mathbf{X}) = \Delta(\mathbf{X}) \sum_{(I,\epsilon)\in\mathcal{F}(\mathbb{L}_{0:k+1})\times\Xi} w_{k+1|k}^{(I,\epsilon)} \delta_I(\mathcal{L}(\mathbf{X}))[p_{k+1|k}^{(\epsilon)}]^{\mathbf{X}} \quad (2)$$

The predicted multi-object density (2) is a weighted mixture of multi-target exponentials. Each component comprises birth and surviving labels. A birth label is a new target that appears at time $k+1$. A surviving label is a target that exists at time k and continues to exist at time $k+1$. The predicted weights $w_{k+1|k}$ and predicted single-target densities $p_{k+1|k}$ are given in [5]. Exact computation of the predicted δ-GLMB distribution requires exhaustive enumeration of all possible combinations of surviving and birth labels, which is computationally intractable in general. In practice, the K-shortest-paths algorithm [18] is used to the find K label sets with the highest mixture weight $w_{k+1|k}^{(I,\epsilon)}$ [6].

2) Update: Let $Z_j = \{z_j^1, z_j^2, ...\}$ denote the measurements of sensor j and let $Z = \{Z_1, ...Z_S\}$ denote the measurements from all S sensors. We assume that $P(Z_i, Z_j|\mathbf{X}) = P(Z_i|\mathbf{X})P(Z_j|\mathbf{X}) \quad \forall i \neq j$. Each sensor j is characterized by a likelihood function $g_j(z|x)$, a probability of detection $P_{D,j}(x)$ and a clutter intensity $\kappa_j(z)$. Denote by θ_j a mapping between the track labels and measurement indices of sensor j. In a valid association map, each label is associated to at most one measurement or miss detected, and each measurement is associated to at most one label. If the prior distribution is a δ-GLMB of the form (2), then the multi-sensor posterior distribution is also a δ-GLMB given by [14]

$$\pi_{k+1|k+1}(\mathbf{X}|Z) = \Delta(X) \sum_{(I,\epsilon)} \sum_{\theta_1} \cdots \sum_{\theta_S} w_{k+1|k+1}^{(I,\epsilon,\theta_1,...,\theta_S)}(Z) \times \delta_I(\mathcal{L}(\mathbf{X}))[p_{k+1|k+1}^{(\epsilon,\theta_1,...,\theta_S)}(\cdot|Z)]^{\mathbf{X}} \quad (3)$$

$$w_{k+1|k+1}^{(I,\epsilon,\theta_1,...,\theta_S)}(Z) \propto w_{k+1|k}^{(I,\epsilon)} [\eta_Z^{(\epsilon,\theta_1,...,\theta_S)}]^I \quad (4)$$

$$\eta_Z^{(\epsilon,\theta_1,...,\theta_S)}(l) = \int p_{k+1|k}^{(\epsilon)}(x,l) \prod_{j=1}^{S} \psi_j(x,l;\theta_j)dx \quad (5)$$

$$\psi_j(x,l;\theta_j) = \begin{cases} 1 - P_{D,j}(x) & \theta_j(l) = 0 \\ \frac{P_{D,j}(x) g_j(z_j^i|x,l)}{\kappa_j(z_j^i)} & \theta_j(l) = i > 0 \end{cases} \quad (6)$$

$$p_{k+1|k+1}^{(\epsilon,\theta_1,...,\theta_S)}(x,l|Z) = \frac{p_{k+1|k}^{(\epsilon)}(x,l) \prod_{j=1}^{S} \psi_j(x,l;\theta_j)}{\eta_Z^{(\epsilon,\theta_1,...,\theta_S)}(l)} \quad (7)$$

The exact multi-sensor δ-GLMB update (3) requires generating all multi-sensor association maps for each predicted hypothesis which is infeasible in all but the simplest scenarios. Practical solutions involve exploring only a limited number of associations for each predicted hypothesis, leading to a truncated δ-GLMB posterior.

III. Implementation of δ-GLMB filter

Given a valid multi-sensor association map $(\theta_1, ..., \theta_S)$ for the predicted hypothesis (I,ϵ) with weight $w_{k+1|k}^{(I,\epsilon)}$, the posterior hypothesis weight is proportional to $w_{k+1|k}^{(I,\epsilon)} \prod_{l\in I} \eta_Z^{(\epsilon,\theta_1,\cdots,\theta_S)}(l)$. Therefore, maximizing the posterior hypothesis weight (score) for a given predicted hypothesis is equivalent to solving:

$$\begin{aligned} \underset{\theta_1 \cdots \theta_S}{\text{minimize}} \quad & \sum_{l\in I} -\log\left(\eta_Z^{(\epsilon,\theta_1,...,\theta_S)}(l)\right), \quad (8) \\ \text{subject to} \quad & \theta_j \text{ valid association map } \forall\, j = 1,\ldots,S. \end{aligned}$$

Equation (8) is known as the multi-sensor assignment problem and is NP-hard for $S > 1$ [16]. We can interpret the term $-\log\left(\eta_Z^{(\epsilon,\theta_1,...,\theta_S)}(l)\right)$ as the cost of assigning measurements $\theta_1(l), ...\theta_S(l)$ to label l and we can thus define the cost of a multi-sensor map for label set I as the sum of assignment costs for individual labels $l \in I$. The objective is to minimize the total cost of $|I|$ assignments without any conflict. In the following, we present two novel approximation algorithms that produce T high-scoring (low-cost) multi-sensor associations for the label set I without exhaustive enumeration.

A. Combination of single-sensor assignments

Consider Eq. (5). If $\eta_Z^{(\epsilon,\theta_1...\theta_S)}(l)$, the multi-sensor assignment weight for label l, can be approximated as a product of single-sensor assignment weights,

$$\eta_Z^{(\epsilon,\theta_1...\theta_S)}(l) \approx \prod_{j=1}^{S} \eta_Z^{(\epsilon,\theta_j)}(l) \quad (9)$$

then the objective function (8) becomes

$$\sum_{j=1}^{S} \left[\sum_{l\in I} -\log\left(\eta_Z^{(\epsilon,\theta_j)}(l)\right) \right]. \quad (10)$$

The approximation in (9) is exact in two different cases. On one hand, if all of the $\psi(x,l)$ are constant with respect to x, then they factor out. On the other hand, in the limit as $p_{k+1|k}^{(\epsilon)}(x,l) \to \delta_\mu(x)$, then

$\int \prod_{j=1}^{S} g_j(z_j^{\theta_j(l)}|x,l)p_{k+1|k}^{(\epsilon)}(x,l)dx \approx \prod_{j=1}^{S} g_j(z_j^{\theta_j(l)}|\mu,l) \approx \prod_{j=1}^{S}\left[\int g_j(z_j^{\theta_j(l)}|x,l)p_{k+1|k}^{(\epsilon)}(x,l)dx\right]$ with μ being the mean of $p_{k+1|k}^{(\epsilon)}(\cdot)$. Note that the same approximation is applied in [14] to simplify the sampling process.

Minimizing (10) is equivalent to solving S independent single-sensor assignment problems using only the sensor's local measurements. Using this assumption, we propose a two-step algorithm to produce T multi-sensor assignment maps which minimize $\sum_{l\in I}\sum_{j=1}^{S} -\log\left(\eta_Z^{(\epsilon,\theta_j)}(l)\right)$.

First, each sensor generates T best (lowest cost) single-sensor association maps using Murty's algorithm [15]. These maps are then combined into multi-sensor maps. Let $\theta_{1...s}^{(1)}, ...\theta_{1...s}^{(T)}$ denote the T best multi-sensor maps for sensors 1 to s. At sensor $s+1$, we obtain T^2 multi-sensor maps by concatenating $\theta_{1...s}^{(1)}, ...\theta_{1...s}^{(T)}$ with $\theta_{s+1}^{(1)}, ...\theta_{s+1}^{(T)}$ (the maps of sensor $s+1$). We compute the cost of the new maps using Eq. (10), retain the T maps with lowest costs and propagate them to sensor $s+2$ and so on.

B. Cross-entropy

The cross-entropy method constructs a distribution on the space of all valid multi-sensor assignments and draws samples from this distribution. The samples with lowest costs are in turn used to update the distribution parameters. With each iteration, the distribution approaches convergence in the sense that assignments with low costs have higher probability. Unlike the combination method, the generated association maps are not necessarily distinct since the same map may be resampled.

We adapt the algorithm proposed in [16]. Consider label set I with n labels and sensor j with measurements $\{z_j^1, ...z_j^{m_j}\}$. We define a row-stochastic probability matrix

$$\pi_j = \begin{bmatrix} p_j(l^1|0) & p_j(l^1|1) & ... & p_j(l^1|m_j) \\ & ... & & \\ p_j(l^n|0) & p_j(l^n|1) & ... & p_j(l^n|m_j) \end{bmatrix} \quad (11)$$

where $p_j(l|i)$ is the probability of assigning the i^{th} measurement of sensor j to label l. To sample from π_j, we go through each label and perform categorical sampling. If z_j^i is assigned to label l, then we set $p(l'|i) = 0, \forall l' \neq l$ and renormalize to ensure the map remains valid. We can generate a multi-sensor assignment by sequentially sampling from π_1, π_2 and so on.

We draw T samples, compute their costs using (8) and select the $\lceil \gamma T \rceil$ samples with lowest costs to update the parameters of π. Let $\theta_j^{(1)}, ...\theta_j^{(\lceil \gamma T \rceil)}$ denote sensor j's assignments from the $\lceil \gamma T \rceil$ $(0 < \gamma \leq 1)$ best samples. We update π_j as follows [16]:

$$p_j(l|z) = \frac{|t \in \{1,2,...,\lceil \gamma T \rceil\} : \theta_j^{(t)}(l) = z|}{\lceil \gamma T \rceil} \quad (12)$$

The numerator in Eq. (12) counts the number of samples in which sensor j assigns measurement z to label l and the denominator is a normalization constant.

In [16], the authors suggest initializing π to be uniform, running multiple iterations and retaining only the T samples from the last iteration. In our simulations, we find that a poor initialization leads to a majority of samples having very high cost in the first few iterations and in turn to a slow convergence towards the optimal distribution. We instead initialize π using probabilistic gating. Each sensor computes the likelihood that label l generates a measurement z. The likelihoods are then normalized and used as initial values of $p(l|z)$. Furthermore, we retain all distinct samples (from all sampling iterations) to maximize the number of output maps.

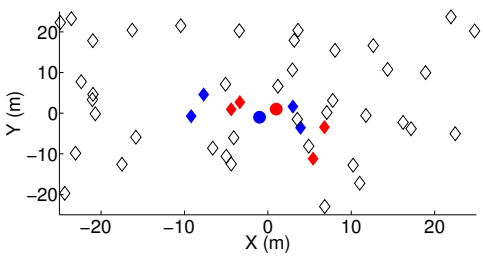

Fig. 1: Target positions (blue and red dots), target-originated measurements (blue and red diamonds) and clutter measurements (empty black diamonds)

IV. Performance evaluation

In this section we evaluate and compare the performance of the two proposed algorithms with the sequential method [10] and an adaptation of the Gibbs sampler described in [14]. The sequential multi-sensor filter involves the iterative application of single-sensor δ-GLMB update where the number of hypotheses is capped after each update step. In [17], a Gibbs sampler is proposed in the joint δ-GLMB framework in which multi-sensor association maps are extended to account for target birth and death. This method is later extended to the multi-sensor setting [14]. We do not consider target birth and death to be a part of the multi-sensor assignment problem and adapt the Gibbs sampler for the separate predict-update δ-GLMB framework. In addition, we reduce the computational complexity of Gibbs sampling by approximating the multi-sensor assignment costs as in the combination method (Eq. (9)). As in the cross entropy method, the sampled maps are not guaranteed to be distinct. We note that our adapted sampler loses the benefits of operating in the joint δ-GLMB framework which is shown to yield robust performance [14].

We consider a network of $S = 4$ sensors. Each sensor detects target(s) with probability $P_D = 0.9$. We consider a linear measurement model (i.e., sensors measure target positions directly) so sensor positions do not affect the algorithms' performance. All measurements are corrupted by white Gaussian noise with covariance matrix $\sigma_{xy}^2 I_2$ where I_2 is the 2×2 identify matrix and $\sigma_{xy} = 3$m. At each time step, each sensor also receives $N_c = 6$ clutter measurements which are uniformly distributed over a $50\text{m} \times 50\text{m}$ tracking area.

For the combination and sequential methods, each sensor generates a maximum of $T = 500$ single-sensor maps. For the cross entropy method, at each sampling iteration, a maximum of T multi-sensor maps are generated and the best $\lceil 0.3T \rceil$ maps are used to update the sampling distribution parameters. Five sampling iterations are run, and all distinct samples from all iterations are output as the multi-sensor assignment maps. Finally, we run two instances of Gibbs sampling. The first

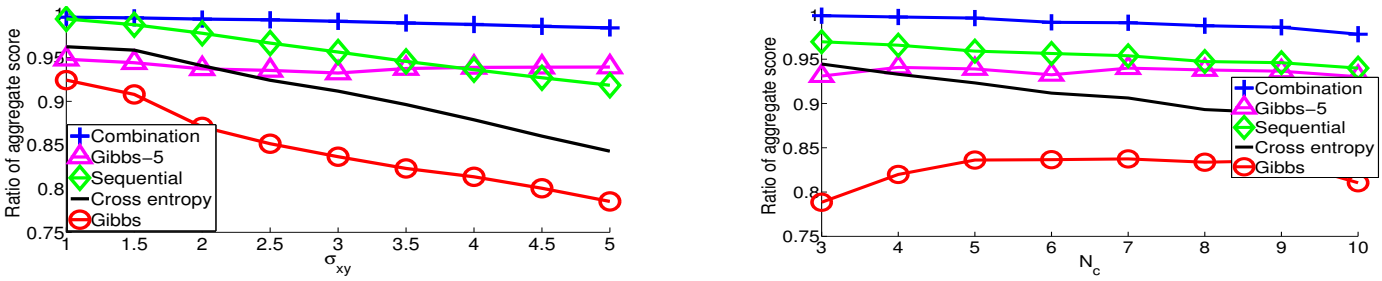

(a) (b)

Fig. 2: a) Average ratio of aggregate map score with respect to σ_{xy} for $N_c = 6$. b) Average ratio of aggregate map score with respect to N_c for $\sigma_{xy} = 3$.

	Mapping time in seconds (number of maps)				
	Comb	CE	Gibbs	Gibbs-5	Seq
Baseline	0.067(500)	0.314(99)	*0.058*(76)	0.271(176)	0.321(329)
T=750	0.113(750)	0.461(122)	*0.088*(95)	0.423(215)	0.390(432)
T=1000	0.133(1000)	0.597(139)	*0.112*(112)	0.579(245)	0.415(518)
N_c=2	*0.039*(500)	0.265(44)	0.052(26)	0.261(65)	0.092(135)
N_c=4	*0.053*(500)	0.288(73)	0.055(52)	0.266(121)	0.195(239)
N_c=8	0.084(500)	0.348(137)	*0.061*(108)	0.276(259)	0.471(404)
N_c=10	0.104(500)	0.365(163)	*0.063*(131)	0.280(311)	0.636(447)

TABLE I: Average runtime in seconds for map generation. The number in bracket is the number of distinct generated maps. For baseline, $\sigma_{xy} = 3$, $T = 500$, $N_c = 6$, $P_D = 0.9$. Shortest runtimes are marked in red and highest number of maps are marked in blue.

runs for 500 iterations (Gibbs) and the second runs for 2500 iterations (Gibbs-5). Both instances output at most T distinct highest-scoring assignments among the samples generated.

The objective is to evaluate and compare the algorithms' ability to generate optimal assignment maps. In our test, we generate a single hypothesis comprised of true target labels. We model the target state density using Gaussian distribution with mean equal to the true target state and covariance matrix $\text{diag}([10, 10, 10, 10])$. We then apply the association algorithms to generate the assignment maps. We run a number of tests in which we vary a single simulation parameter. Each test consists of 200 random Monte Carlo trials. The sensor measurements change at each trial; but the target positions remain fixed.

Fig. 1 shows the positions of targets and all sensor measurements for one trial (at default parameters values). Note that the two targets are in close proximity and several clutter measurements are difficult to distinguish from true target-originated measurements, making this a challenging scenario.

Table I shows the average running time and the number of generated maps for all algorithms. As expected, the running time increases when the number T of maps or the number N_c of clutter measurements increases. The combination and Gibbs method have comparable running time for $N_c \leq 4$ and are significantly faster than the other two methods. The cross entropy has the longest running time. This can be attributed to two factors: running multiple sampling iterations and computing the exact assignment cost for all distinct samples. The high running time of Gibbs-5 is due to the higher number of samples. In the case of the sequential δ-GLMB, the number of posterior hypotheses grows with each application of the single-sensor update procedure which also accounts for the high computational load of the filter.

The combination method consistently generates T distinct maps. The Gibbs and cross entropy methods generate significantly fewer maps as expected. The Gibbs-5 method generates more maps at the cost of higher computational overhead.

Our second metric is the ratio of aggregate map score. We construct a combined pool of distinct maps from all algorithms. Note that the algorithms can generate different maps so the pool may contain as many as $5T$ maps. We then select the T highest-scoring maps from the pool. These T maps are the best maps that the algorithms have produced and their aggregate score, C_{optimal}, is a suitable baseline for measuring different algorithms' performance. For each algorithm, we compute the aggregate of its generated maps and report its ratio to C_{optimal} with higher ratio denoting better performance. A ratio of 0 means that the algorithm did not generate any of the top T maps and a ratio of 1 means that algorithm has produced all T highest-scoring maps.

Fig. 2(a) shows the average ratio of aggregate map score with respect to σ_{xy}. As σ_{xy} increases, the ratio decreases for all algorithms with a bigger drop for the Gibbs, cross entropy and sequential methods. The combination method has the highest ratio for $\sigma_{xy} \geq 1.5$. The Gibbs-5 method outperforms the sequential method for $\sigma_{xy} \geq 4$.

Fig. 2(b) shows the average ratio of aggregate map score with respect to N_c. The combination method has the best performance for all values of N_c followed by the Gibbs-5 and sequential methods. For the Gibbs method, higher N_c leads to more maps and better performance; although the performance does degrade for sufficiently high N_c.

In both figures, the low aggregate score of the Gibbs method comes from the combined effect of using the approximate scoring function and more importantly having few distinct multi-sensor assignments. On the other hand, the Gibbs-5 method which generates more distinct maps is able to yield robust performance at the cost of higher computational overhead.

V. Conclusion

In this paper we present two novel algorithms to solve the multi-sensor assignment problem in the δ-GLMB filter. The slow running time of the cross entropy method renders it unsuitable for applications with strict timing requirements. The combination method offers fast running time and competitive performance across a wide range of test scenarios, making it an attractive solution over the standard sequential method.

REFERENCES

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

[2] ——, "PHD filters of higher order in target number," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 43, no. 4, pp. 1523–1543, 2007.

[3] B.-T. Vo, B.-N. Vo, and A. Cantoni, "Analytic implementation of the cardinalized probability hypothesis density filter," *IEEE Trans. Signal Processing*, vol. 55, no. 7, pp. 3553–3567, Jul. 2007.

[4] ——, "The cardinality balanced multi-target multi-Bernoulli filter and its implementations," *IEEE Trans. Signal Processing*, vol. 57, no. 2, pp. 409–423, 2009.

[5] B. T. Vo and B. N. Vo, "Labeled random finite sets and multi-object conjugate priors," *IEEE Trans. Signal Processing*, vol. 61, no. 13, pp. 3460–3475, 2013.

[6] B.-N. Vo, B.-T. Vo, and D. Phung, "Labeled random finite sets and the Bayes multi-target tracking filter," *IEEE Trans. Signal Processing*, vol. 62, no. 24, pp. 6554–6567, Dec. 2014.

[7] D. Clark and S. Nagappa, "On the ordering of the sensors in the iterated-corrector probability hypothesis density PHD filter," in *Proc. SPIE Signal Proc., Sensor Fusion, Target Recognition*, vol. 8050, Orlandao, Florida, Apr. 2011, pp. 80 500M–80 500M–6.

[8] R. Mahler, "The multisensor PHD filter: I, general solution via multitarget calculus," in *Proc. SPIE*, vol. 7336, 2009, pp. 73 360E–73 360E–12.

[9] B. Wang, W. Yi, S. Li, L. Kong, and X. Yang, "Distributed fusion with multi-Bernoulli filter based on generaized covariance intersection," in *IEEE Radar Conf.*, May 2015, pp. 958–962.

[10] F. Papi, "Multi-sensor δ-GLMB filter for multi-target tracking using Doppler only measurements," in *Europen Intell. Security Inform. Conf.*, Sept. 2015, pp. 83–89.

[11] W. F. Liu, B. S. Wei, and S. J. Zhu, "A multi-sensor generalized labeled multi-Bernoulli filter via extended association map," in *Proc. Intl. Conf. Control, Automation, Inform. Sci.*, Oct. 2015, pp. 225–230.

[12] S. Nannuru, S. Blouin, M. Coates, and M. Rabbat, "Multisensor CPHD filter," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 52, no. 4, pp. 1834–1854, Aug. 2016.

[13] A.-A. Saucan, M. Coates, and M. Rabbat, "A multi-sensor multi-Bernoulli filter," 2016, `arXiv:1609.05108 [stat.ME]`.

[14] B. N. Vo and B. Vo, "Multi-sensor multi-object tracking with the generalized labeled multi-Bernoulli filter," *IEEE Trans. Signal Proc.*

[15] M. Miller, H. Stone, and I. Cox, "Optimizing Murty's ranked assignment method," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 33, no. 3, pp. 851–862, Jul. 1997.

[16] D. M. Nguyen, H. A. L. Thi, and T. P. Dinh, "Solving the multidimensional assignment problem by a cross-entropy method," *J. Combinatorial Optimization*, vol. 27, no. 4, pp. 808–823, May 2014.

[17] B. N. Vo, B. Vo, and H. G. Hoang, "An efficient implementation of the generalized labeled multi-Bernoulli filter," *IEEE Trans. Signal Processing*, vol. 65, pp. 1975–1987, Dec. 2016.

[18] D. Eppstein, "Finding the K shortest paths," *SIAM J. Comput.*, vol. 28, no. 2, pp. 652–673, Feb. 1999.