

Tracking Multiple Objects with Particle Filtering

C. Hue*

J.-P. Le Cadre

P. Pérez

IRISA/Université de Rennes1 IRISA/CNRS

Microsoft Research

Campus universitaire de Beaulieu

7 J.J. Thomson Avenue

35042 Rennes Cedex, France

Cambridge, CB3 0FB, UK

e-mail: {chue,lecadre}@irisa.fr

pperez@microsoft.com

Tel : 33-2.99.84.25.21 Fax : 33-2.99.84.71.71 Tel : 44.1223.479.844 Fax : 44.1223.479.999

<http://www.irisa.fr/vista><http://research.microsoft.com/users/pperez/>

* Corresponding author

Abstract

We address the problem of multitarget tracking encountered in many situations in signal or image processing. We consider stochastic dynamic systems detected by observation processes. The difficulty lies on the fact that the estimation of the states requires the assignment of the observations to the multiple targets. We propose an extension of the classical particle filter where the stochastic vector of assignment is estimated by a Gibbs sampler. This algorithm is used to estimate the trajectories of multiple targets from their noisy bearings, thus showing its ability to solve the data association problem. Moreover this algorithm is easily extended to multireceiver observations where the receivers can produce measurements of various nature with different frequencies.

Keywords

Multitarget tracking, particle filter, Gibbs sampler, multireceiver, bearings-only tracking.

I. INTRODUCTION

Multitarget tracking (MTT) deals with the state estimation of an unknown number of moving targets. Available measurements may both arise from the targets if they are detected, and from clutter. Clutter is generally considered as a model describing false alarms. Its (spatio-temporal) statistical properties are quite different from those of the target, which makes the extraction of target tracks from clutter possible. To perform multitarget tracking the observer can rely on a huge amount of data, possibly collected on multiple receivers. Elementary measurements are receiver outputs, e.g., bearings, ranges, time-delays, Dopplers, etc. The main difficulty, however, comes from the assignment of a given measurement to a target model. These assignments are generally unknown, as are the true target models. This is a neat departure from classical estimation problems. Thus, two distinct problems have to be solved *jointly*: the data association and the estimation.

The simplest approach is probably the Nearest Neighbor approach. Using only the observation the closest to the predicted state, this algorithm is not robust enough in many situations. As long as the association

is considered in a deterministic way, the possible associations must be exhaustively enumerated. This leads to a NP-hard problem because the number of possible associations increases exponentially with time, as in the Multiple Hypotheses Tracker (MHT) [1]. To cope with this problem, pruning and gating eliminate the less likely hypotheses but can unfortunately eliminate good ones as well. In the Joint Probabilistic Data Association Filter (JPDAF) [2], the association variables are considered as stochastic variables and one needs only to evaluate the validated association probabilities at each time step. However, the dependence assumption on the associations implies the exhaustive enumeration of all possible associations at the current time step. When the association variables are instead supposed statistically independent like in the Probabilistic MHT (PMHT [3], [4]), the complexity is reduced. For instance in [3] or [4], the algorithm is presented as an incomplete data problem solved by an EM algorithm. There is no measurement gating as in the JPDAF and all the associations are considered. The results are then satisfactory when the measurement equation is linear and when the trajectories are deterministic. In [5] the algorithm is extended to the tracking of maneuvering targets with an hidden “model-switch” process controlled by a Markov probability structure. A comparison of the PMHT with the JPDAF is described in a practical two-target scenario in [6], focusing on the mean-square estimation errors and the percentage of lost tracks. Unfortunately, the above algorithms do not cope with non linear models and non Gaussian noises.

Under such assumptions (stochastic state equation and non linear state or measurement equation, non Gaussian noises), sequential Monte Carlo methods, also called particle filtering methods, are particularly appealing. They mainly consist of propagating, in a possibly non-linear way, a weighted set of particles which approximates the probability density of the state conditioned on the observations according to Monte Carlo integration principles. The weights of the particles are updated using Bayes’s formula. Particle filtering can be applied under very general hypotheses, is able to cope with heavy clutter, and is very easy to implement. Such filters have been used in very different areas for Bayesian filtering, under different names: the bootstrap filter for target tracking in [7] and the Condensation algorithm in computer vision [8] are two examples among others. In earliest studies, the algorithm was only composed of two periods: the particles were predicted according to the state equation during the prediction step; then their weights were calculated with the likelihood of the new observation combined with the former weights. A resampling step has rapidly been added to dismiss the particles with lower weights and avoid the degeneracy of the particle set into a unique particle of high weight [7]. Many ways have been developed to accomplish this resampling whose final goal is to enforce particles in areas of high likelihood. The frequency of this resampling has also been studied. Also the use of kernel filters [9] has been introduced to regularize the sum of Dirac densities associated to the particles when the dynamic noise of the state equation was too low [10]. Despite this long history of studies, in which the ability of particle filter to track multiple posterior modes is claimed, the extension of the particle filter to multiple target tracking has progressively received attention only in the last five years. Such extensions have first been claimed theoretically feasible in [11] and [12] but the examples chosen only dealt with one single target. In computer vision a probabilistic exclusion principle has been developed in

[13] to track multiple objects but the algorithm is very dependent of the observation model and is only applied for two objects. In the same context, a Bayesian multiple-blob tracker called BraMBLe [14] has just been proposed. It deals with a varying number of objects which are depth-ordered thanks to a 3-D state space. Lately, in mobile robotic [15], a set of particle filters for each target connected by a statistical data association has been proposed. We propose here a general algorithm for multitarget tracking in the passive sonar context.

This work is organized as follows. In Section II, we describe the basic particle filter for a single target with two versions for the resampling step, which can be systematic or adaptive. The superiority of adaptive resampling towards systematic resampling is clearly demonstrated in a passive sonar application in Section III. Section IV, the central part of this work, deals with an extension of the basic filter to multiple objects. The new algorithm combines the two major steps of prediction and weighting of the classical particle filter with a Gibbs sampler computing an estimation of the vector of the assignment probabilities. We introduce two different versions of this Gibbs sampler which take into account the past information in a different way. An extension to multi-receiver data in the context of multiple targets ends this Section and highlights the versatility of our approach. Finally, Section V is devoted to an application to bearings-only MTT which enables us to compare the results obtained with these two versions. In each case, the data association problem is overcome and the first version is clearly superior.

As far as the notational conventions are concerned, we always use the index i to refer to one among the M tracked objects. The index j designates one of the m_t observations obtained at instant t . The index n is reserved for the N particles denoted by s . The index τ is used for indexing the iterations in the Gibbs Sampler and r is used for the different receivers. Finally, the probability densities are denoted by p if they are continuous and by \mathbb{P} if they are discrete.

II. THE BASIC PARTICLE FILTER

For the sake of completeness, the basic particle filter is now briefly reviewed. The general principle of sampling it relies on will be used throughout the paper. We consider a dynamic system represented by the stochastic process $(X_t) \in \mathbb{R}^{n_x}$ whose temporal evolution is given by the state equation:

$$X_t = F_t(X_{t-1}, V_t). \quad (1)$$

We want to estimate the state vector (X_t) at discrete times with the help of system's observations which are realizations of the stochastic process $(Y_t) \in \mathbb{R}^{n_y}$ governed by the measurement equation:

$$Y_t = H_t(X_t, W_t). \quad (2)$$

The two processes $(V_t) \in \mathbb{R}^{n_v}$ and $(W_t) \in \mathbb{R}^{n_w}$ are only supposed to be independent white noises. Note that the functions F_t and H_t are not assumed linear.

We will denote by $Y_{0:t}$ the sequence of the random variables (Y_0, \dots, Y_t) and by $y_{0:t}$ one realization of this sequence. Note that throughout the paper, the first subscript of any vector will always refer to the time.

Our problem consists in computing at each time t the conditional density $L_t = p(X_t|Y_0 = y_0, \dots, Y_t = y_t)$ of the state X_t given all the observations accumulated up to t , and also in estimating any functional $g(X_t)$ of the state via the expectation $\mathbb{E}(g(X_t)|Y_{0:t})$. The Recursive Bayesian filter, also named Optimal Filter, resolves exactly this problem in two steps at each time t .

Suppose we know L_{t-1} . The *prediction step* is done according to the following equation:

$$p(X_t = x_t|Y_{0:t-1} = y_{0:t-1}) = \int_{\mathbb{R}^{n_x}} p(X_t = x_t|X_{t-1} = x)L_{t-1}(x)dx. \quad (3)$$

Using (1), we can calculate $p(X_t = x_t|X_{t-1} = x)$:

$$p(X_t = x_t|X_{t-1} = x) = \int_{\mathbb{R}^{n_v}} p(X_t = x_t|X_{t-1} = x, V_t = v)p(V_t = v|X_{t-1} = x)dv. \quad (4)$$

$$(5)$$

The observation y_t enables us to correct this prediction using Bayes's rule:

$$L_t(x_t) = \frac{p(Y_t = y_t|X_t = x_t)p(X_t = x_t|Y_{0:t-1} = y_{0:t-1})}{\int_{\mathbb{R}^{n_x}} p(Y_t = y_t|X_t = x)p(X_t = x|Y_{0:t-1} = y_{0:t-1})dx}. \quad (6)$$

Under the specific assumptions of Gaussian noises V_t and W_t and linear functions F_t and H_t , these equations lead to the Kalman filter's equations. Unfortunately this modeling is not appropriate in many problems in signal and image processing, which makes the calculation of the integrals in equations (3) and (6) infeasible (no closed-form).

The original particle filter, which is called the bootstrap filter [7], proposes to approximate the densities $(L_t)_t$ by a finite weighted sum of N Dirac densities centered on elements of \mathbb{R}^{n_x} , which are called particles.

The application of the bootstrap filter requires that one knows how to do the following:

- sample from initial prior marginal $p(X_0)$;
- sample from $p(V_t)$ for all t ;
- compute $p(Y_t = y_t|X_t = x_t)$ for all t through a known function l_t such that $l_t(y; x) \propto p(Y_t = y|X_t = x)$ where missing normalization must not depend on x .

The first particle set S_0 is created by drawing N independent realizations from $p(X_0)$ and assigning uniform weight $1/N$ to each of them. Then, suppose we have at our disposal at time $t - 1$ the weighted particle set $S_{t-1} = (s_{t-1}^n, q_{t-1}^n)_{n=1, \dots, N}$ where $\sum_{n=1}^N q_{t-1}^n = 1$. The *a posteriori* marginal L_{t-1} is then estimated by the probability density $L_{S_{t-1}} = \sum_{n=1}^N q_{t-1}^n \delta_{s_{t-1}^n}$. The *prediction step* consists of propagating each particle of S_{t-1} according to the evolution equation (1).

The weight of each particle is updated during the *correction step*. Up to a constant, equation (6) comes down to adjust the weight of predictions by multiplying it by the likelihood $p(y_t|x_t)$. In the most general setting of sequential Monte Carlo methods [16], the displacement of particles is obtained by sampling from an appropriate density f which might depend on the data as well. The general algorithm is summarized in Fig. 1. The density L_{S_t} is often multi-modal as several hypotheses about the position of the object can be made at one time. It is for instance the case when one object is tracked in the presence of significant

- Initialization: $\begin{cases} s_0^n \sim p(X_0) \\ q_0^n = 1/N \end{cases} \quad n = 1, \dots, N.$
- For $t = 1, \dots, T$:
 - Proposal: sample s_t^n from $f(X_t|X_{t-1} = s_{t-1}^n, Y_t = y_t)$ for $n = 1, \dots, N.$
 - Weighting: $\begin{cases} \text{Compute un-normalized weights: } \tilde{q}_t^n = q_{t-1}^n \frac{p(s_t^n|s_{t-1}^n) l_t(y_t; s_t^n)}{f(s_t^n|s_{t-1}^n, y_t)} \text{ for } n = 1, \dots, N. \\ \text{Normalize weights: } q_t^n = \frac{\tilde{q}_t^n}{\sum_{n=1}^N \tilde{q}_t^n} \text{ for } n = 1, \dots, N. \end{cases}$
 - Return $\hat{\mathbb{E}}g(X_t) = \sum_{n=1}^N q_t^n g(s_t^n).$

Fig. 1. *Basic particle filter without resampling.*

clutter. Several hypotheses about the object position can then be kept if the set of particles splits into several subsets. This is where the great strength of this filter lies. In [17] for instance, the measurement vector Y_t consists of a set of detected features along line measurements. The assumed underlying generative model affects each feature either to the target boundary, or to its interior or to the background. The likelihood function is built from this generative model and then takes into account the clutter model. An extension of the algorithm in Fig. 1, called the hybrid bootstrap filter, has also been proposed to deal with significant clutter and spurious objects for target tracking [12] and guidance [18]. The weighted sum of Dirac laws is then approximated by a Gaussian mixture obtained by a clustering method [19].

The particle sets enable one to estimate any functional of X_t in particular the two first moments with $g(x) = x$ and $g(x) = x^2$ respectively. The mean can be used to estimate the position of one object but it can be a bad estimator if the posterior is highly multimodal. In such cases the ideal would be to calculate the mean only over the particles that contribute to the principal mode but such an estimator has not been developed for the moment.

In practice, the number of particles is finite and the major drawback of this algorithm is the degeneracy of the particle set: only few particles keep high weights and the others have very small ones. The former carry the information, whereas the latter are mostly useless. The resampling is a good way to remedy this drawback because it eliminates the particles of smallest weight. The stochastic resampling consists of sampling N particles with replacement in the particle set with the probability q^n to draw s^n . The new particles have uniform weights equal to $1/N$. A first solution, adopted in [7] for example, consists of applying the resampling step at each time period. The corresponding algorithm of particle filter with systematic resampling is described in Fig. 2.

To measure the degeneracy of the algorithm, the effective sample size N_{eff} has been introduced in [20], [21]. We can estimate this quantity by $\hat{N}_{eff} = 1 / \sum_{n=1}^N (q_t^n)^2$ which measures the number of meaningful particles. As advocated in [16], the resampling can be done only if $\hat{N}_{eff} < N_{threshold}$. This enables the particle set to better learn the process and to keep its memory during the interval where no resampling occurs. The algorithm of the basic particle filter with adaptive resampling is described in Fig. 3. Details can be found in [16], [20] or [21].

- Initialization: $\begin{cases} s_0^n \sim p(X_0) \\ q_0^n = 1/N \end{cases} \quad n = 1, \dots, N.$
- For $t = 1, \dots, T$:
 - Proposal: sample \tilde{s}_t^n from $f(X_t|X_{t-1} = s_{t-1}^n, Y_t = y_t)$ for $n = 1, \dots, N.$
 - Weighting: $\begin{cases} \text{Compute un-normalized weights: } \tilde{q}_t^n = q_{t-1}^n \frac{p(\tilde{s}_t^n | s_{t-1}^n) l_t(y_t; \tilde{s}_t^n)}{f(\tilde{s}_t^n | s_{t-1}^n, y_t)} \text{ for } n = 1, \dots, N. \\ \text{Normalize weights: } q_t^n = \frac{\tilde{q}_t^n}{\sum_{n=1}^N \tilde{q}_t^n} \text{ for } n = 1, \dots, N. \end{cases}$
 - Return $\hat{\mathbb{E}}g(X_t) = \sum_{n=1}^N q_t^n g(\tilde{s}_t^n).$
 - Resampling: $\begin{cases} \text{Draw } s_t^n \text{ from } \sum_{k=1}^N q_t^k \delta_{\tilde{s}_t^k} \\ \text{Set } q_t^n = 1/N \end{cases} \quad n = 1, \dots, N$

Fig. 2. Basic particle filter with systematic resampling.

Some convergence results of the empirical distributions to the posterior distribution on the path space have

- Initialization: $\begin{cases} s_0^n \sim p(X_0) \\ q_0^n = 1/N \end{cases} \quad n = 1, \dots, N.$
- For $t = 1, \dots, T$:
 - Proposal: sample \tilde{s}_t^n from $f(X_t|X_{t-1} = s_{t-1}^n, Y_t = y_t)$ for $n = 1, \dots, N.$
 - Weighting: $\begin{cases} \text{Compute un-normalized weights: } \tilde{q}_t^n = q_{t-1}^n \frac{p(\tilde{s}_t^n | s_{t-1}^n) l_t(y_t; \tilde{s}_t^n)}{f(\tilde{s}_t^n | s_{t-1}^n, y_t)} \text{ for } n = 1, \dots, N. \\ \text{Normalize weights: } q_t^n = \frac{\tilde{q}_t^n}{\sum_{n=1}^N \tilde{q}_t^n} \text{ for } n = 1, \dots, N. \end{cases}$
 - Return $\hat{\mathbb{E}}g(X_t) = \sum_{n=1}^N q_t^n g(\tilde{s}_t^n).$
 - Calculate $\hat{N}_{eff} = \frac{1}{\sum_{n=1}^N (q_t^n)^2}.$
 - Resampling: if $\hat{N}_{eff} < N_{threshold}$: $\begin{cases} s_t^n \sim \sum_{k=1}^N q_t^k \delta_{\tilde{s}_t^k} \\ q_t^n = 1/N \end{cases} \quad n = 1, \dots, N, \text{ else } s_t^n = \tilde{s}_t^n \text{ for } n = 1, \dots, N.$

Fig. 3. Basic particle filter with adaptive resampling.

been proved when the number N of particles tends towards infinity [22], [23]. In the path space $(\mathbb{R}^{n_x})^{t+1}$, each particle s_t^n at time t can be considered to be a discrete path of length $t+1$. Compared with the particle filter presented in Fig. 1, particle filtering in the space of paths consists of incrementing the particle state space at each time step and representing each particle by the concatenation of the new position at time t and the set of previous positions between times 0 and $t-1$. In [22], the fluctuations on path space of the so-called interacting particle systems are studied. In the context of sequential Monte Carlo methods [23], which cover most of the particle filtering methods proposed in the last years, the convergence and the rate of convergence of order $1/N$ of the average mean square error is proved. Under more restrictive hypotheses, the almost sure convergence is proved as well [23].

III. APPLICATION TO BEARINGS-ONLY PROBLEMS

To illustrate the previous algorithms, we first deal with classical bearings-only problems. The object is then a “point-object” in the $x - y$ plane. In the context of a slowly maneuvering target, we have chosen a nearly-constant-velocity model (see [24] for a review of the principal dynamical models used in this domain).

A. The model

The state vector X_t represents the coordinates and the velocities in the $x - y$ plane: $X_t = (x_t, y_t, vx_t, vy_t)$. The discretized state equation associated with time period Δt is:

$$X_{t+\Delta t} = \begin{pmatrix} I_{2 \times 2} & \Delta t I_{2 \times 2} \\ 0 & I_{2 \times 2} \end{pmatrix} X_t + \begin{pmatrix} \frac{\Delta t^2}{2} I_{2 \times 2} \\ \Delta t I_{2 \times 2} \end{pmatrix} V_t, \quad (7)$$

where $I_{2 \times 2}$ is the identity matrix in dimension 2 and V_t is a Gaussian zero-mean vector of covariance matrix $\Sigma_V = \begin{pmatrix} \sigma_x^2 & 0 \\ 0 & \sigma_y^2 \end{pmatrix}$. Let \hat{X}_t be the estimate of X_t computed by the particle filters with $g(x) = x$, i.e., $\hat{X}_t = \sum_{n=1}^N q_t^n \hat{s}_t^n$. The posterior covariance matrix of \hat{X}_t is

$$P_t = \sum_{n=1}^N q_t^n (\hat{s}_t^n - \hat{X}_t)(\hat{s}_t^n - \hat{X}_t)^T \quad (8)$$

where T denotes the transposition.

The observations are available at discrete times according to:

$$Y_t = \arctan \left(\frac{y_t - y_t^{obs}}{x_t - x_t^{obs}} \right) + W_t, \quad (9)$$

where W_t is a zero-mean Gaussian noise of covariance σ_w^2 independent of V_t , and x_{obs} and y_{obs} are the Cartesian coordinates of the observer, which are known.

B. Results of particle filters with systematic and adaptive resampling respectively

The initial position of the target and of the observer are

$$X_0 = \begin{pmatrix} 200 \text{ m} \\ 1500 \text{ m} \\ 1.0 \text{ ms}^{-1} \\ -0.5 \text{ ms}^{-1} \end{pmatrix}; \quad X_0^{obs} = \begin{pmatrix} 200 \text{ m} \\ -3000 \text{ m} \\ 1.2 \text{ ms}^{-1} \\ +0.5 \text{ ms}^{-1} \end{pmatrix}.$$

The dynamic noise is a normal zero-mean Gaussian vector with $\sigma_x = \sigma_y = 0.001 \text{ ms}^{-2}$. The observer is following a leg by leg trajectory. Its velocity vector is constant on each leg and modified at the following instants:

$$\begin{pmatrix} vx_{200}^{obs} \\ vy_{200}^{obs} \end{pmatrix} = \begin{pmatrix} -0.6 \\ 0.3 \end{pmatrix}; \quad \begin{pmatrix} vx_{400}^{obs} \\ vy_{400}^{obs} \end{pmatrix} = \begin{pmatrix} 2.0 \\ 0.3 \end{pmatrix}; \quad \begin{pmatrix} vx_{600}^{obs} \\ vy_{600}^{obs} \end{pmatrix} = \begin{pmatrix} -0.6 \\ 0.3 \end{pmatrix}$$

$$\begin{pmatrix} vx_{800}^{obs} \\ vy_{800}^{obs} \end{pmatrix} = \begin{pmatrix} 2.0 \\ 0.3 \end{pmatrix}; \begin{pmatrix} vx_{900}^{obs} \\ vy_{900}^{obs} \end{pmatrix} = \begin{pmatrix} -0.6 \\ 0.3 \end{pmatrix}.$$

The bearings measurements are simulated with a Gaussian noise of standard deviation $\sigma_w = 0.05$ radians (about 3 degrees) every time period, i.e., every 6 seconds. The measurements set used and the trajectories of the observer and of the target are presented in Fig. 4.




Fig. 4. *Single target bearings-only experiment in Section III (a) Target and observer trajectories; (b) Bearings simulated at each time period.*




Fig. 5. *The values of the effective sample size \hat{N}_{eff} obtained with 1000 particles and adaptive resampling.*

First, we have studied the impact of adaptive or systematic resampling on the estimate posterior covariance defined by (8). We have used bootstrap filters, i.e, the importance function f is in fact the prior law $p(X_t|X_{t-1})$. The initialization of the filters has been done in both cases according to a Gaussian law whose mean vector is the true vector X_0 and covariance matrix X_{cov} is:

$$X_{cov} = \begin{pmatrix} 20.0^2 & 0 & 0 & 0 \\ 0 & 100^2 & 0 & 0 \\ 0 & 0 & 0.05^2 & 0 \\ 0 & 0 & 0 & 0.05^2 \end{pmatrix}.$$

For the filter with adaptive resampling, the threshold $N_{threshold}$ is fixed to 0.9. The average frequency of resampling obtained for such a threshold is 0.132. The values of \hat{N}_{eff} are presented in Fig. 5 for one particular run.

The results are plotted in Fig. 6: the observer trajectory, the target trajectory and the estimate obtained with adaptive and systematic resampling are plotted with the 2σ confidence ellipses on position and the line

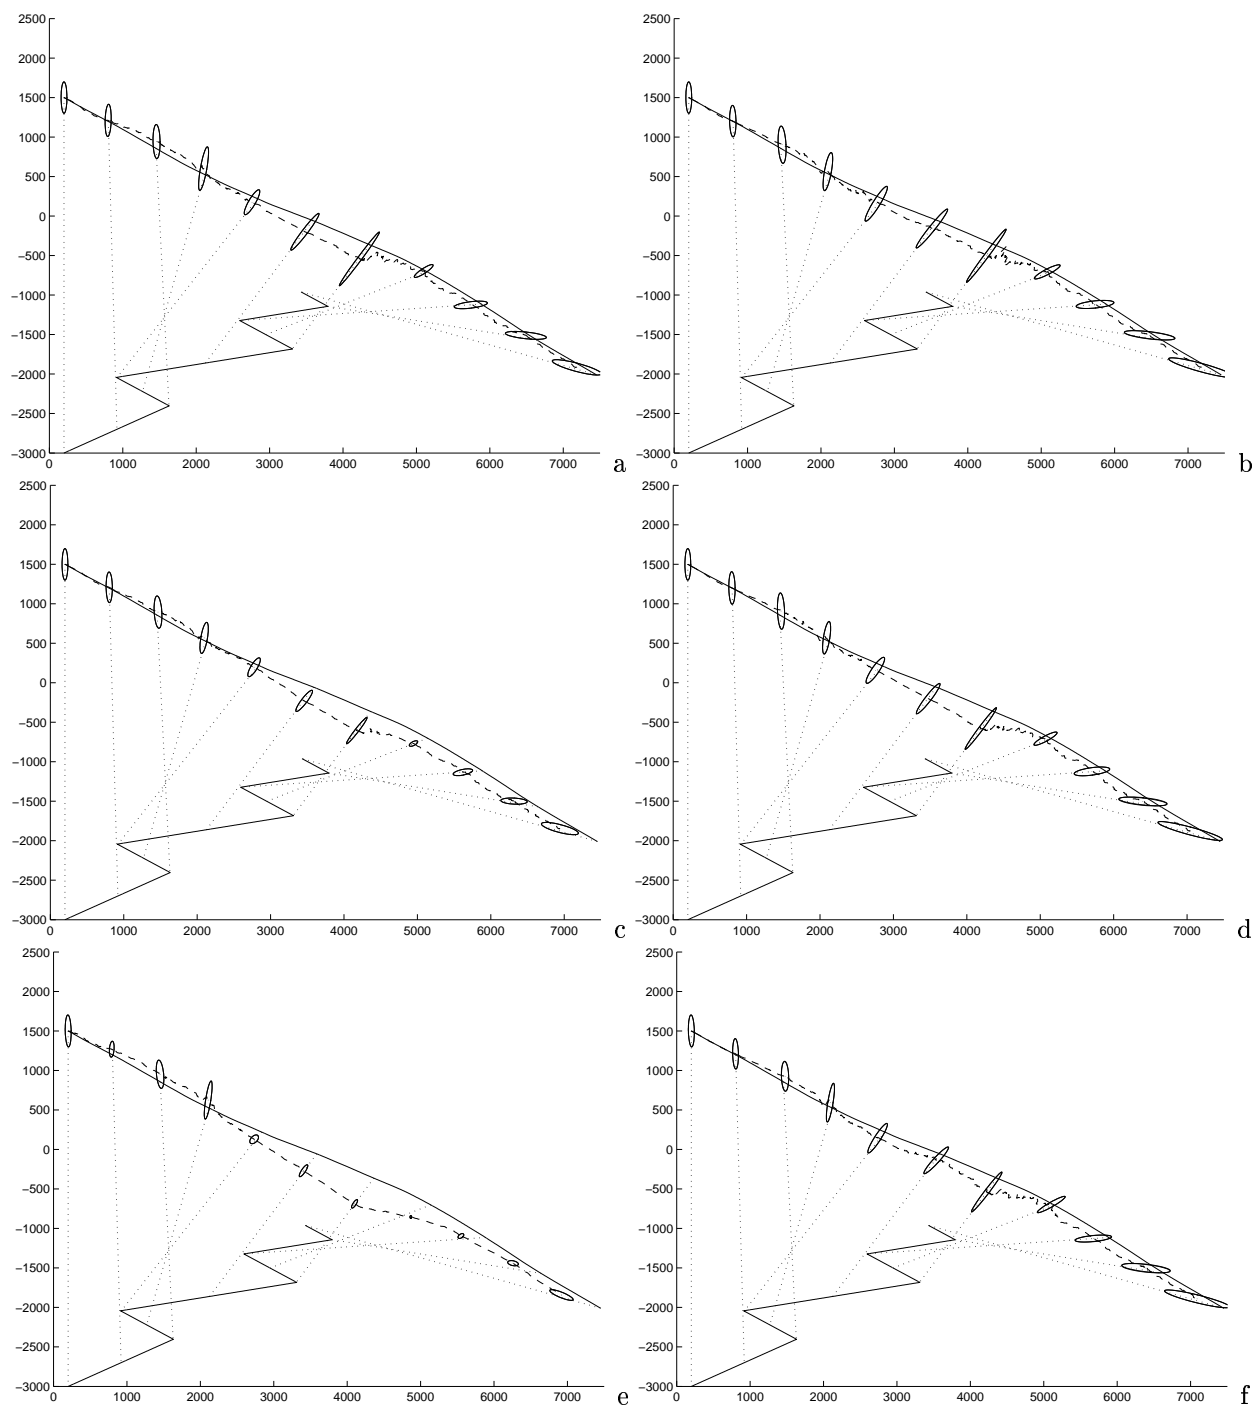


Fig. 6. Estimates for one particular run (dashed lines) and 2σ confidence ellipses obtained with the bootstrap filter. The solid lines stand for the target and observer trajectories. The dotted lines indicate the line of sight of the observer every hundred times. Top: 10000 particles, Middle: 5000 particles, Bottom: 1000 particles; Left column: systematic resampling, Right column: adaptive resampling.

of sight of the observer every hundred times for 1000, 5000 and 10000 particles. The ellipses represent the regions containing 95% of the particles assuming they are Gaussian distributed. At each time, it is interesting to check that the ellipses are oriented according to the line of sight. Moreover with only 1000 particles, the advantage of using adaptive resampling is emphasized: resampling impoverishes the particle set and when applied systematically, the particle set can become so reduced that the ellipse does not contain the true trajectory any more. It is the case from instant 400 in Fig. 6.e whereas with adaptive resampling in Fig. 6.f the confidence ellipses still contain the true trajectory.

The estimates obtained with basic particle filters using adaptive or systematic resampling are optimal for a infinite particle number. In practice, the particle number is of course finite, and the estimation is determined by the realizations of the random processes used in the proposal. To further assess the accuracy of the two filters, we have performed 100 different runs of the particle filters with systematic resampling on the one hand and adaptive resampling on the other hand, for a given realization of the measurement process.¹ From these different runs, we have computed the averaged estimate and the 2σ confidence ellipses on position containing 95% of the estimates. They are represented in Fig. 7 for $N = 1000, 5000, 10000$ particles and adaptive or systematic resampling. The averaged estimates are similar but the confidence ellipses are globally twice as small with adaptive than with systematic resampling whatever the particle number. This indicates that adaptive resampling significantly reduces the standard deviation on the particle filter estimates.

As for the computational cost, Table I contains the cost for one time iteration with systematic resampling and the time spent on the resampling procedure according to the particle number on a 863 Mhz Pentium III.

TABLE I

Computational cost for one iteration on a 863 Mhz Pentium III.

particle number	1000	5000	10000
systematic resampling	7.9ms	41.9ms	88.5ms
resampling procedure	4ms	28ms	64 ms

C. Overestimated state noise

It is to be noted that in the case of real data, the variance of the state noise is not always known precisely. Some simulations have then been performed with an overestimated state noise. Intuitively, the particles are thus predicted in a larger area but are also more prone to be eliminated in the resampling step. In practice this is a great strength of particle filtering. Fig. 8 shows the averaged estimates and the 2σ confidence ellipses obtained with adaptive resampling, and with 1000 and 5000 particles. The estimation is less smooth than

¹It is also interesting to compute the particle filter estimation for different realizations of the measurement process, as it is often done to evaluate the performance of deterministic algorithms such as the Kalman filter or the extended Kalman filter. It is done in Section III-E.




Fig. 7. Averaged estimates (dashed lines) and 2σ confidence ellipses obtained with the bootstrap filter. Top: 10000 particles, Middle: 5000 particles, Bottom: 1000 particles; Left column: systematic resampling, Right column: adaptive resampling.

with a correct state noise. This jittering increases the bias and the std and more particles are needed for results comparable with those obtained with a correct state noise: with 5000 particles the standard deviation between estimates is drastically lower than with 1000 particles.




Fig. 8. *Averaged estimates (dashed lines) and 2σ confidence ellipses obtained with the bootstrap filter with an overestimated state noise. (a) 1000 particles; (b) 5000 particles.*

D. Shifted initialization

We have also performed particle filters runs with a shifted initialization: the initialization is done according to a Gaussian law whose mean vector X_{mean} and covariance matrix X_{cov} are:

$$X_{mean} = X_0 + \begin{pmatrix} 50.0 \text{ m} \\ 500 \text{ m} \\ 0.0 \text{ ms}^{-1} \\ -0.0 \text{ ms}^{-1} \end{pmatrix}; X_{cov} = \begin{pmatrix} 40.0^2 & 0 & 0 & 0 \\ 0 & 400.0^2 & 0 & 0 \\ 0 & 0 & 0.05^2 & 0 \\ 0 & 0 & 0 & 0.05^2 \end{pmatrix}. \quad (10)$$

After around 200 time periods, the particles have recovered from their shifted initialization and provide satisfactory estimates as shown in Fig. 9. The ellipses size is considerably reduced with 5000 particles compared to only 1000 particles.

E. Comparison with a Modified Polar coordinate Extended Kalman Filter

We have finally compared particle filtering methods with a linearized filter, the Modified Polar coordinate Extended Kalman Filter (MP-EKF). Let us denote by X_t^C , $X_t^{C,Rel}$ and X_t^{MP} the state vector in the fixed Cartesian system, in the Cartesian system relative to the observer, and in the modified polar system respectively.

$$X_t^C = (x_t, y_t, \dot{x}_t, \dot{y}_t), X_t^{C,Rel} = (x_t - x_t^{obs}, y_t - y_t^{obs}, \dot{x}_t - \dot{x}_t^{obs}, \dot{y}_t - \dot{y}_t^{obs}), \text{ and}$$

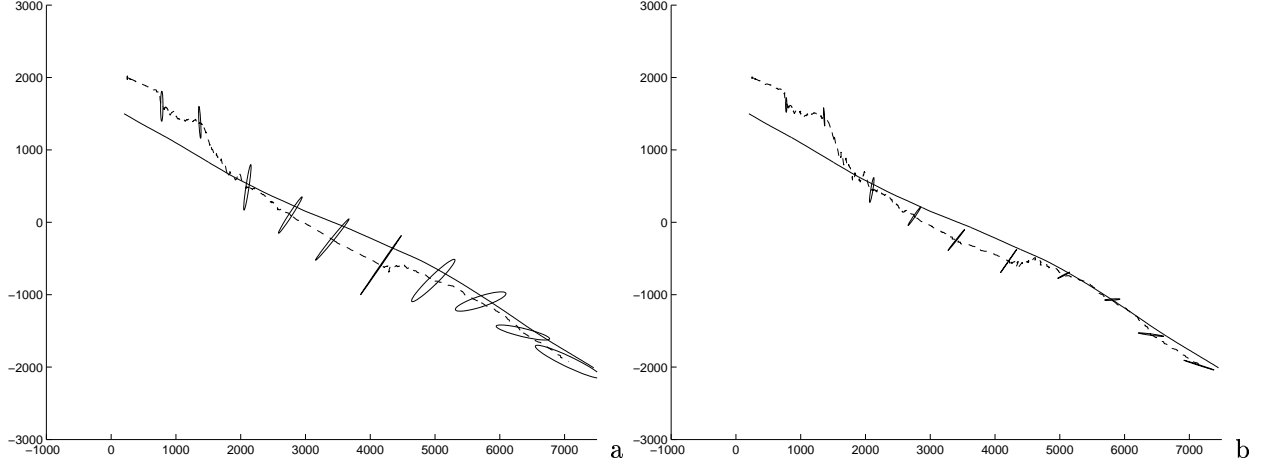


Fig. 9. Averaged estimates (dashed lines) and 2σ confidence ellipses obtained with the bootstrap filter with shifted initialization. (a) 1000 particles; (b) 5000 particles.

$$X_t^{MP} = \left(\beta, \frac{1}{R}, \dot{\beta}, \frac{\dot{R}}{R} \right)^T = \left(\arctan \left(\frac{x_t^{C,Rel}}{y_t^{C,Rel}} \right), \frac{1}{\sqrt{x_t^{C,Rel}{}^2 + y_t^{C,Rel}{}^2}}, \frac{y_t^{C,Rel} \dot{x}_t^{C,Rel} - x_t^{C,Rel} \dot{y}_t^{C,Rel}}{x_t^{C,Rel}{}^2 + y_t^{C,Rel}{}^2}, \frac{x_t^{C,Rel} \dot{x}_t^{C,Rel} + y_t^{C,Rel} \dot{y}_t^{C,Rel}}{x_t^{C,Rel}{}^2 + y_t^{C,Rel}{}^2} \right)^T.$$

Let f_C^{MP} and f_{MP}^C be the functions such that:

$$X_t^{MP} = f_C^{MP}(X_t^{C,Rel}) \text{ and } X_t^{C,Rel} = f_{MP}^C(X_t^{MP}). \quad (11)$$

It must be noted that these two functions are non linear. The dynamic model in the Cartesian system is defined by:

$$X_{t+\Delta t}^C = F X_t^C + V_t \quad (12)$$

where $F = \begin{pmatrix} I_{2 \times 2} & \Delta t I_{2 \times 2} \\ 0 & I_{2 \times 2} \end{pmatrix}$ and V_t is a Gaussian noise of constant covariance matrix Σ_V . In the MP system it reads:

$$X_{t+\Delta t}^{MP} = f_C^{MP} (F f_{MP}^C(X_t^{MP}) - U_{t+\Delta t}^{obs} + V_t) \quad (13)$$

where $U_{t+\Delta t}^{obs} = (x_{t+\Delta t}^{obs} - x_t^{obs} - \Delta t \dot{x}_t^{obs} \quad y_{t+\Delta t}^{obs} - y_t^{obs} - \Delta t \dot{y}_t^{obs} \quad \dot{x}_{t+\Delta t}^{obs} - \dot{x}_t^{obs} \quad \dot{y}_{t+\Delta t}^{obs} - \dot{y}_t^{obs})^T$. If the observer follows a leg by leg trajectory, the two first components of U_t^{obs} are always null. The two last components are null except if t is a manoeuvre time.

In the Cartesian system, the measurement equation is defined by:

$$Y_t = H X_t^C + W_t \quad (14)$$

where W_t is a Gaussian noise of covariance σ_w^2 . In the MP system, it reads:

$$Y_t = Q X_t^{MP} + W_t \quad (15)$$

where $Q = (1 \ 0 \ 0 \ 0)$.

Let us now describe the propagation and updating steps of the MP-EKF used. The predicted state is given by:

$$\tilde{X}_{t+\Delta_t}^{MP} = f_C^{MP} \left(F f_{MP}^C(\hat{X}_t^{MP}) - U_{t+\Delta_t}^{obs} \right). \quad (16)$$

The predicted covariance matrix in the Cartesian system is:

$$\tilde{P}_{t+\Delta_t}^C = F \hat{P}_t^C F^T + \Sigma_V \quad (17)$$

and in the MP system:

$$\tilde{P}_{t+\Delta_t}^{MP} = J_C^{MP} \tilde{P}_t^C (J_C^{MP})^T \quad (18)$$

where J_C^{MP} is the Jacobian matrix of f_C^{MP} taken around the predicted state.

As the measurement equation is linear in the MP system, the updating step is the same as in the classical Kalman filter:

$$K = \tilde{P}_{t+\Delta_t}^{MP} Q^T (Q \tilde{P}_{t+\Delta_t}^{MP} Q^T + \sigma_w^2)^{-1}. \quad (19)$$

$$\hat{X}_{t+\Delta_t}^{MP} = \tilde{X}_{t+\Delta_t}^{MP} + K (Y_{t+\Delta_t} - Q \tilde{X}_{t+\Delta_t}^{MP}). \quad (20)$$

$$\hat{P}_{t+\Delta_t}^{MP} = (I - KQ) \tilde{P}_{t+\Delta_t}^{MP} (I - KQ)^T + K \sigma_w^2 K^T = \tilde{P}_{t+\Delta_t}^{MP} - KQ \tilde{P}_{t+\Delta_t}^{MP}. \quad (21)$$

$$\hat{P}_{t+\Delta_t}^C = J_{MP}^C \hat{P}_{t+\Delta_t}^{MP} (J_{MP}^C)^T \quad (22)$$

where J_{MP}^C is the Jacobian matrix of f_{MP}^C taken around the estimated state.

More details on the MP-EKF can be found in [25], [26] and [27]. As the MP-EKF is a deterministic algorithm, we have computed the estimates obtained for several realizations of the measurement process. For a given realization, the MP-EKF has to be compared with the mean of the particle filter estimates over several runs. We have compared the performance in the case of a shifted initialization like in (10) in the previous Section over 10 realizations of the measurement process. For each of them, 20 runs of the particle filter with adaptive resampling and 5000 particles are computed. The final bias and standard deviation are similar but the particle filter recovers faster from the shifted initialization than the MP-EKF (see Fig. 10).

IV. MULTITARGET PARTICLE FILTERS

The purpose of the two precedent Sections was to provide a general representation of particle filtering. We can now turn to its extension to multi-target tracking.




Fig. 10. Mean estimates (dotted lines) obtained over 10 realizations of the measurement process and 2σ confidence ellipses with a shifted initialization; (a) Estimates obtained with MP-EKF; (b) Estimates obtained with particle filters with adaptive resampling and 5000 particles.

A. The MTT problem and its classical treatment

Let M be the number of targets to track. This number is assumed to be known and fixed for the moment (the case of a varying unknown number will be addressed in another work). The index i designates one among the M targets and is always used as first superscript. Multitarget tracking consists in estimating the state vector made by concatenating the state vectors of all targets. It is generally assumed that the targets are moving according to independent Markovian dynamics. At time t , $X_t = (X_t^1, \dots, X_t^M)$ follows the state equation (1) decomposed in M partial equations:

$$X_t^i = F_t^i(X_{t-1}^i, V_t^i) \quad \forall i = 1, \dots, M. \quad (23)$$

The noises (V_t^i) and $(V_t^{i'})$ are supposed only to be white both temporally and spatially, and independent for $i \neq i'$.

The observation vector collected at time t is denoted by $y_t = (y_t^1, \dots, y_t^{m_t})$. The index j is used as first superscript to refer to one of the m_t measurements. The vector y_t is composed of detection measurements and clutter measurements. The false alarms are assumed to be uniformly distributed in the observation area. Their number is assumed to arise from a Poisson density of parameter λV where V is the volume of the observation area and λ the number of false alarms per unit volume. As we do not know the origin of each measurement, one has to introduce the vector K_t to describe the associations between the measurements and the targets. Each component K_t^j is a random variable that takes its values among $\{0, \dots, M\}$. Thus, $K_t^j = i$ indicates that y_t^j is associated with the i^{th} target. In this case, y_t^j is a realization of the stochastic process:

$$Y_t^j = H_t^i(X_t^i, W_t^j) \text{ if } K_t^j = i. \quad (24)$$

Again, the noises (W_t^j) and $(W_t^{j'})$ are supposed only to be white noises, independent for $j \neq j'$. We assume that the functions H_t^i are such that they can be associated to functional forms l_t^i such that $l_t^i(y; x) \propto p(Y_t^j = y | K_t^j = i, X_t^i = x)$.

We dedicate the model 0 to false alarms. Thus, if $K_t^j = 0$, the j^{th} measurement is associated to the clutter,

but we do not associate any kinematic model to false alarms.

As the indexing of the measurements is arbitrary, all the measurements have the same *a priori* probability to be associated with a given model i . At time t , these association probabilities define the vector $\pi_t = (\pi_t^0, \pi_t^1, \dots, \pi_t^M) \in [0, 1]^{M+1}$. Thus, for $i = 1, \dots, M$, $\pi_t^i \triangleq \mathbb{P}(K_t^j = i)$ for all $j = 1, \dots, m_t$ is the discrete probability that any measurement is associated with the i^{th} target.

To solve the data association some assumptions are commonly made [28]:

A1: One measurement can originate from one target or from the clutter.

A2: One target can produce zero or one measurement at one time.

The assumption (A1) expresses that the association is exclusive and exhaustive. Consequently, $\sum_{i=0}^M \pi_t^i = 1$.

The assumption (A2) implies that m_t may differ from M and above all that the association variables K_t^j for $j = 1, \dots, m_t$ are dependent.

Under these assumptions, the MHT algorithm [1] builds recursively the association hypotheses. One advantage of this algorithm is that the appearance of a new target is hypothesized at each time step. However, the complexity of the algorithm increases exponentially with time. Some pruning solutions must be found to eliminate some of the associations.

The JPDAF begins with a gating of the measurements. Only the measurements which are inside an ellipse around the predicted state are kept. The gating assumes that the measurements are distributed according to a Gaussian law centred on the predicted state. Then, the probabilities of each association $K_t^j = i$ are estimated. As the variables K_t^j are assumed dependent by (A2), this computation requires the exhaustive enumeration of all the possible associations K_t^l for $l \neq j$.

The novelty in the PMHT algorithm [3], [4], [5] consists of replacing the assumption (A2) by (A3):

A3: One target can produce zero or several measurements at one time.

This assumption is often criticized because it does not match the physical reality. However, from a mathematical point of view it ensures the stochastic independence of the variables K_t^j and it drastically reduces the complexity of the π_t vector estimation.

The assumptions (A1) and (A3) will be kept in the joint filters presented later. Let us present now the existing works solving MTT with particle filtering methods.

B. Related work: MTT with particle filters

In the context of multitarget tracking, particles filters are appealing: as the association needs only to be considered at a given time iteration, the complexity of data association is reduced. First, two extensions of the bootstrap filter have been considered. In [11], a bootstrap-type algorithm is proposed in which the sample state space is a “(multitarget) state space”. However, nothing is said about the association problem that needs to be solved to evaluate the sample weights. It is in fact the ability of the particle filtering to deal with multimodality due to (high) clutter that is pointed out compared with deterministic algorithms like the nearest neighbor filter or the PDA filter. No examples with multiple targets are presented: the simulations only deal with a single target in clutter with a linear observation model. In [12], a hybrid bootstrap filter

is presented where the particles evolve in a single-object state space. Each particle gives an hypothesis on the state of one object. Thus, the *a posteriori* law of the targets given the measurements is represented by a Gaussian mixture. Each mode of this law then corresponds to one of the objects. However, as pointed out in [12], the likelihood evaluation is possible only under the availability of the “prior probabilities of all possible associations between” the measurements and the targets. Even if the likelihood could be evaluated, the way to represent the *a posteriori* law by a mixture can lead to the loss of one of the targets during occlusions. The particles tracking an occluded target get very small weights and are therefore discarded during the resampling step. This fact has been pointed out in [15].

In image analysis, the Condensation algorithm has been extended to the case of multiple objects as well: in [13], the case of two objects is considered. The hidden state is the concatenation of the two single-object states and of a binary variable indicating which object is closer to the camera. This latter variable solves the association during occlusion because the measurements are affected to the foreground object. Moreover, a probabilistic exclusion principle is integrated to the likelihood measurement to penalize the hypotheses with the two objects overlapping. In [14], the state is composed of an integer equal to the number of objects and of a concatenation of the individual states. A 3-D representation of the objects gives access to their depth ordering, thus solving the association issue during occlusions. Three other works combine particle filtering with the JPDA. In mobile robotic [15], a particle filter is used for each object tracked. The likelihood of the measurements is written like in a JPDAF. Thus, the assignment probabilities are evaluated according to the probabilities of each possible association. Given these assignment probabilities, the particle weights can be evaluated. The particle filters are then dependent through the evaluation of the assignment probabilities. The algorithms presented in [29] and [30] are both applied to target tracking. The state space of each particle is the concatenation of the state space of each target and the likelihood of the measurements given a particle is derived according to the JPDAF. Independently of the latter works [14], [15], [29] and [30], we have developed the joint filters where the data association is approached in the same probabilistic spirit as the basic PMHT [3], [4].

First, to estimate the density $L_t = p(X_t = (X_t^1, \dots, X_t^M) | Y_{0:t} = y_{0:t})$ with particle filtering methods we must choose the state space for the particles. We could use particles of X_t^i 's dimension and distribute them among the objects (providing that the objects have the same dimension which might not be the case for complex objects in image processing). Each of them would then be known to represent a given object. The weight of particle n is then:

$$q_t^n \propto p(Y_t = (y_t^1, \dots, y_t^{m_t}) | X_t^i = s_t^n). \quad (25)$$

As the observations cannot be supposed independent conditioned on only one of the M objects, the likelihood cannot be decomposed in the product of likelihoods of each observation. One solution could be to use Bayes's rule and to compute $p(X_t^i = s_t^n | Y_t = (y_t^1, \dots, y_t^{m_t}))$, but this requires that every object produces at least one measurement (to be able to use the Total Probability Theorem), which seems too restrictive. Moreover, as

mentioned before, a unique particle filter with a single-target state space seems inappropriate as the particles tracking an occluded object would be quickly discarded. We have considered to use one particle filter per object but without finding a consistent way to make them dependent. The stochastic association vector K_t introduced in Section IV-A could also be considered as an additional particle component. However, as the ordering of the measurements is arbitrary, it would not be possible to devise a dynamic prior on it. Moreover, the state space would increase further making particle filter less effective. Finally, we have chosen to use particles whose dimension is the sum of those of the individual state spaces corresponding to each target, as in [13] and [14]. Each of these concatenated vectors then gives jointly a representation of all targets. We will call the filters associated with this representation the joint filters.

C. The MTPF algorithms

Let us first study the joint density $p(X_t, Y_t, \Pi_t, K_t)$ (implicitly conditioned on X_{t-1}) where Π_t is the stochastic variable associated with the probability π_t .

We will drop the subscript t for a few lines.

$$p(X, Y, \Pi, K) = p(Y|K, \Pi, X)p(K|\Pi, X)p(\Pi, X) \quad (26)$$

$$= p(Y|K, X) p(K|\Pi) p(\Pi)p(X) \quad (27)$$

where the last equality is based on the following independence assumptions:

- Y and Π are independent given K and X ;
- the association vector K is independent from the state vector X given Π .

The factors in (27) write:

- $p(Y|K, X) = \prod_j p(Y^j|K^j, X) = \prod_j \begin{cases} l^i (y^j|x^i) & \text{with } i = K^j \in \{1 \dots M\}; \\ 1/V & \text{if } K^j = 0. \end{cases}$
- To simulate according to $p(K|\Pi)$ it is sufficient to generate $N(K)$ where $N^i(K) \triangleq \#\{j : K^j = i\}$. The vector $(N^0(K), \dots, N^M(K))$ follows a multinomial law of size m_t and of parameters (Π^0, \dots, Π^M) ;
- $p(\Pi) = p(\Pi^0)p(\Pi^1, \dots, \Pi^M|\Pi^0)$ where $p(\Pi^1, \dots, \Pi^M|\Pi^0)$ is uniform on the $[0, 1 - \pi^0]^M$'s hyperplane $\sum_{i=1}^M \pi_i = 1 - \pi^0$. π^0 is a constant that can be computed:

$$\pi^0 = \mathbb{P}(K_t^j = 0) \quad (28)$$

$$= \sum_{l=0}^{m_t} \mathbb{P}(K_t^j = 0 | N_t^0 = l) \mathbb{P}(N_t^0 = l) \quad (29)$$

$$= \sum_{l=0}^{m_t} \frac{l}{m_t} \exp(-\lambda V) \frac{(\lambda V)^l}{l!} \quad (30)$$

Assuming that there are l measurements due to clutter among the m_t measurements collected at time t , the *a priori* probability that any measurement comes from the clutter is equal to $\frac{l}{m_t}$ hence the equality $\mathbb{P}(K_t^j = 0 | N_t^0 = l) = \frac{l}{m_t}$ used to derive (30) from (29).

- $p(X)$, i.e., $p(X_t|X_{t-1})$ conditional on X_{t-1} .

This independence structure of this distribution is illustrated in Fig. 11.

The initial particle set $S_0 = (s_0^n, 1/N)_{n=1, \dots, N}$ is such that each component $s_0^{n,i}$ for $i = 1, \dots, M$ is sampled




Fig. 11. Independence graph of the stochastic variables X, Y, K, Π .

from $p(X_0^i)$ independently from the others. Assume we have obtained $S_{t-1} = (s_{t-1}^n, q_{t-1}^n)_{n=1, \dots, N}$ with $\sum_{n=1}^N q_{t-1}^n = 1$. Each particle is a vector of dimension $\sum_{i=1}^M n_x^i$ where we denote by $s_{t-1}^{n,i}$ the i^{th} component of s_{t-1}^n and where n_x^i designates the dimension of object i .

The prediction can be done according to the following equation:

$$\text{For } n = 1, \dots, N \quad \tilde{s}_t^n = \begin{pmatrix} F_t^1(s_{t-1}^{n,1}, v_t^{n,1}) \\ \vdots \\ \vdots \\ F_t^M(s_{t-1}^{n,M}, v_t^{n,M}) \end{pmatrix}. \quad (31)$$

Examine now the computation of the likelihood of the observations conditioned on the n^{th} particle. We can write for all $n = 1, \dots, N$:

$$p(Y_t = (y_t^1, \dots, y_t^{m_t}) | X_t = \tilde{s}_t^n) = \prod_{j=1}^{m_t} p(y_t^j | \tilde{s}_t^n), \quad (32)$$

$$\propto \prod_{j=1}^{m_t} \left[\frac{\pi_t^0}{V} + \sum_{i=1}^M l_t^i(y_t^j; \tilde{s}_t^{n,i}) \pi_t^i \right]. \quad (33)$$

It must be noted that equation (32) is true only under the assumption of conditional independence of the measurements, which we will make. Moreover, the normalization factors between l_t^i and $p(Y_t^j = y | K_t^j = i, X_t^i = x)$ must be the same for all i to write the last equality (33).

We still need to estimate at each time step the association probabilities $(\pi_t^i)_{i=1, \dots, M}$, which can be seen as the stochastic coefficients of the M -component mixture. The simplest way would consist in estimating the π vector from the problem parameters like the number of targets M , the obtained number of measurements m_t and the clutter density λ . Assuming, the π_t^i are equal for all $i \neq 0$, we can empirically estimate $\pi_t^i = \frac{m_t - \lambda V}{M}$. However, these estimates do not take into account the current measurements. In particular, if a target is not

The initialization of the Gibbs sampler consists of assigning uniform association probabilities, i.e., $\pi_{t,0}^i = \frac{1-\pi_t^0}{M}$ for all $i = 1, \dots, M$, and taking $X_{t,0} = \sum_{n=1}^N q_{t-1}^n \tilde{s}_t^n$, i.e., the centroid of the predicted particle set. The K_t variables do not need initializing because at the first time step of the Gibbs sampler they will be sampled conditioned on $\pi_{t,0}^i$, $i = 1, \dots, M$ and $X_{t,0}$. Then, suppose that at instant t we have already simulated $(\theta_{t,1}, \dots, \theta_{t,\tau})$. The $\tau + 1^{th}$ iteration is handled as follows.

- As the $(K_t^j)_{j=1, \dots, m_t}$ are supposed to be independent, their individual conditional density reads:

$$p(K_t^j | Y_{0:t}, X_t, (K_t^l)_{l \neq j}, \Pi_t) = p(K_t^j | Y_t^j, X_t, \Pi_t). \quad (35)$$

(K_t^j) are discrete variables and we can write:²

$$\mathbb{P}(K_t^j = i | Y_t^j = y_t^j, X_t, \Pi_t) = \frac{p(Y_t^j = y_t^j | K_t^j = i, X_t, \Pi_t) \mathbb{P}(K_t^j = i | X_t, \Pi_t)}{p(Y_t^j = y_t^j | X_t, \Pi_t)} \quad (36)$$

$$\propto \begin{cases} \pi_t^i l_t^i(y_t^j; x_t^i) & \text{if } i = 1, \dots, M; \\ \pi_t^0 / V & \text{if } i = 0. \end{cases} \quad (37)$$

The realizations $k_{t,\tau+1}^j$ of the vector $K_{t,\tau+1}$ are then sampled according to the weights

$p_{t,\tau+1}^{j,0} = \pi_t^0 / V, p_{t,\tau+1}^{j,i} = \pi_{t,\tau}^i l_t^i(y_t^j; x_{t,\tau}^i)$ for $i = 1, \dots, M$.

- Mixture proportion vector $\Pi_{t,\tau+1}^{1:M}$ is drawn from the conditional density:

$$p(\Pi_{t,\tau+1}^{1:M} | K_{t,\tau+1}, X_{t,\tau}, Y_{0:t}) = p(\Pi_t^1, \dots, \Pi_t^M | K_{t,\tau+1}^1, \dots, K_{t,\tau+1}^M, X_{t,\tau}, Y_{0:t}) \quad (38)$$

$$= p(\Pi_t^1, \dots, \Pi_t^M | K_{t,\tau+1}^1, \dots, K_{t,\tau+1}^M) \quad (39)$$

$$\propto \text{Dirichlet}(\Pi_t^{1:M} | \{1 + N^i(K_{t,\tau+1})\}_{i=1, \dots, M}) \quad (40)$$

where we denote by $N^i(K)$ the number of k^j equal to i and where $\text{Dirichlet}(\delta_1, \dots, \delta_M)$ denotes the Dirichlet distribution on the simplex $\{(\pi_t^1, \dots, \pi_t^{M-1}, 1 - \pi_t^1 - \dots - \pi_t^{M-1}) : \pi_t^1 + \dots + \pi_t^{M-1} \leq 1\}$ with density proportional to $\pi_t^{1(\delta_1-1)} \times \dots \times \pi_t^{M-1(\delta_{M-1}-1)} \times (1 - \pi_t^1 - \dots - \pi_t^{M-1})^{\delta_M-1}$. The vector $\Pi_{t,\tau+1}^{1:M}$ is first drawn according to (40) and then normalized to restore the sum $\sum_{i=1}^M \pi_{t,\tau+1}^i$ to $1 - \pi_t^0$.

- $X_{t,\tau+1}$ has to be sampled according to the density

$$p(X_t | Y_{0:t}, K_{t,\tau+1}, \Pi_{t,\tau+1}) = \prod_{i=1}^M p(X_t^i | Y_{0:t}, K_{t,\tau+1}, \Pi_{t,\tau+1}). \quad (41)$$

The values of $K_{t,\tau+1}$ can imply that one object is associated with zero or several measurements. Let us define $\mathcal{M}_{t,\tau+1}^i = \{j \in [1 \dots m_t] : K_{t,\tau+1}^j = i\}$. Hence we decompose the preceding product in two products:

$$\prod_{i: \mathcal{M}_{t,\tau+1}^i \neq \emptyset} p(X_t^i | Y_{0:t-1}, y_t^{\mathcal{M}_{t,\tau+1}^i}, \Pi_{t,\tau+1}) \prod_{i: \mathcal{M}_{t,\tau+1}^i = \emptyset} p(X_t^i | Y_{0:t-1}, \Pi_{t,\tau+1}) \quad (42)$$

where $y_t^{\mathcal{M}_{t,\tau+1}^i} = \{y_t^j, j \in \mathcal{M}_{t,\tau+1}^i\}$. The first product contains the targets that are associated to at least one measurement under the association $K_{t,\tau+1}$. In this case, the measurements are denoted by $y_t^{\mathcal{M}_{t,\tau+1}^i}$. The

²Using Bayes's rule, $p(a|b, c) = \frac{p(b|a, c)p(a|c)}{p(b|c)}$.

second product contains the targets that are associated with no measurements under $K_{t,\tau+1}$.

- Let i be an integer in the first product. We propose two approaches to sample $X_{t,\tau+1}$:

★ Without making any additional assumption we can write

$$p(X_t^i | Y_{0:t-1}, \mathbf{y}_t^{\mathcal{M}_{t,\tau+1}^i}, \Pi_{t,\tau+1}) = \frac{p(\mathbf{y}_t^{\mathcal{M}_{t,\tau+1}^i} | X_t^i) p(X_t^i | Y_{0:t-1})}{p(\mathbf{y}_t^{\mathcal{M}_{t,\tau+1}^i} | Y_{0:t-1})}. \quad (43)$$

We are not able to sample directly from the density $\frac{p(\mathbf{y}_t^{\mathcal{M}_{t,\tau+1}^i} | X_t^i) p(X_t^i | Y_{0:t-1})}{p(\mathbf{y}_t^{\mathcal{M}_{t,\tau+1}^i} | Y_{0:t-1})}$, for the same reasons as those exposed in Section II to justify the use of the particle filter (intractability of the integrals). A first solution consists in building the particle set $\Sigma_{\tau+1} = (\sigma_{\tau+1}^n, \chi_{\tau+1}^n)_{n=1,\dots,N}$ whose weights $\chi_{\tau+1}^n$ measure the likelihood of the observations affected by $K_{t,\tau+1}$ to object X_t^i . More precisely, we let:

$$\begin{cases} \sigma_{\tau+1}^n = \tilde{s}_t^{n,i}; \\ \chi_{\tau+1}^n = \frac{p(\mathbf{y}_t^{\mathcal{M}_{t,\tau+1}^i} | X_t^i = \sigma_{\tau+1}^n) q_{t-1}^n}{\sum_{n=1}^N p(\mathbf{y}_t^{\mathcal{M}_{t,\tau+1}^i} | X_t^i = \sigma_{\tau+1}^n) q_{t-1}^n}. \end{cases} \quad (44)$$

As the predicted empirical law $L_{\tilde{S}_t} = \sum_{n=1}^N q_{t-1}^n \tilde{s}_t^{n,i}$ is “close” from the predicted law $p(X_t^i | Y_{0:t-1})$, we expect the empirical distribution $\Lambda_{\tau+1} = \sum_{n=1}^N \chi_{\tau+1}^n \delta_{\sigma_{\tau+1}^n}$ to be close to $p(X_t^i | \mathbf{y}_t^{\mathcal{M}_{t,\tau+1}^i}, Y_{0:t-1})$. However the weak convergence of $\Lambda_{\tau+1}$ to $p(X_t^i | \mathbf{y}_t^{\mathcal{M}_{t,\tau+1}^i}, Y_{0:t-1})$ when N tends towards infinity remains to be proved. Not being able to sample from this last density, $X_{t,\tau+1}^i$ is drawn as a realization from $\Lambda_{\tau+1}$.

★ The second solution assumes the measurement equation enables us to sample from the density $P(X_t = x | Y_t = y)$ and to forget the observations from the past. The likelihood $p(X_t^i | Y_{0:t-1}, \mathbf{y}_t^{\mathcal{M}_{t,\tau+1}^i}, \Pi_{t,\tau+1})$ is then reduced to $p(X_t^i | \mathbf{y}_t^{\mathcal{M}_{t,\tau+1}^i})$. We do not assume that the observations $\mathbf{y}_t^{\mathcal{M}_{t,\tau+1}^i}$ are independent but we use their centroid \bar{y} and replace $p(X_t^i = x | \mathbf{y}_t^{\mathcal{M}_{t,\tau+1}^i}, K_t^{j^1} = i, \dots, K_t^{j^i} = i)$ by $p(X_t^i = x | \bar{y})$.

As far as the complexity of these two solutions is concerned, it is to be noticed that the first one depends linearly on the total number of particles whereas the second is independent of it. On the other hand, the second solution requires the ability to sample from $p(X_t | Y_t)$.

- Now let i be an integer in the second product. As we do not have any measurement to correct the predicted particles we draw a realization from the density $\sum_{n=1}^N q_{t-1}^n \delta_{\tilde{s}_t^n}$ for $X_{t,\tau+1}^i$.

After a finite number of iterations, we estimate the vector π_t by the average of its last realizations:

$$\hat{\pi}_t^i = \frac{1}{\tau_{beg} - \tau_{end}} \sum_{\tau=\tau_{beg}}^{\tau_{end}} \pi_{t,\tau}^i. \quad (45)$$

Finally the weights can be computed according to (33) using the estimate $\hat{\pi}_t^i$ of π_t^i . By construction, $\hat{\pi}_t^i$ follows the law $p(\Pi_t^i | Y_{0:t})$. Thus, the use of a Gibbs Sampler enables to take into account the current measurements. Consequently the estimates measure in a way the *a posteriori* detection probability of each target. It improves the estimation of the targets because the measurements contribute to the estimation proportionally to these probabilities $\hat{\pi}_t^i$. Moreover, the *a priori* probability of detecting a target, which is usually denoted by P_d is not needed in the MTPF. This probability is needed when the associations are

considered in classical algorithms like the PMHT or the JPDAF.

The resampling step is performed in an adaptive way when the estimated effective sample size \hat{N}_{eff} is under the threshold $N_{threshold}$.

Due to the estimation of the π_t vector needed for the computation of the particles likelihood, the convergence of the MTPF could be affected. It could be interesting to evaluate the error on the estimate of X_t implied by the error made on the estimate of π_t . This is not addressed in this work. Figs. 12 and 13 summarize the whole algorithm.

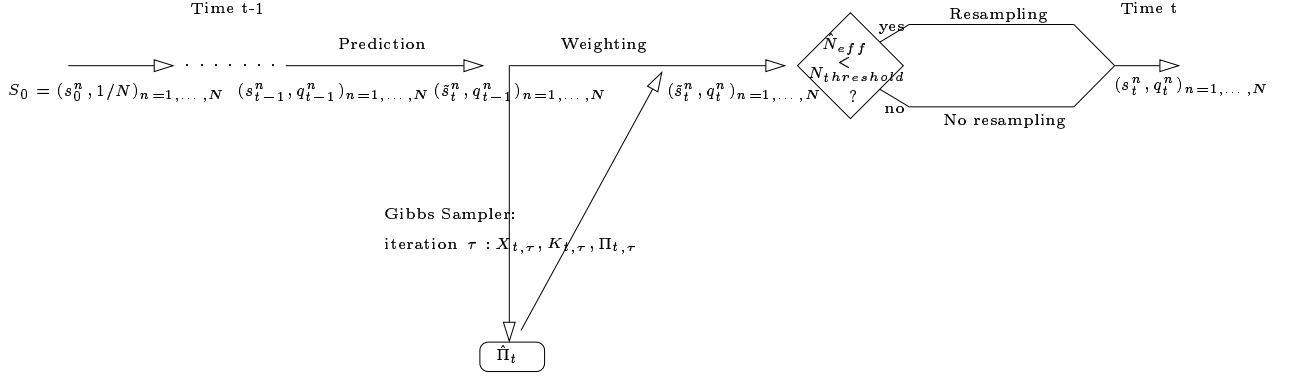


Fig. 12. Particle Filter for multiple objects with adaptive resampling.

Before presenting the results of some simulations, let us detail the approach used in [29] and [30] which is the closest to our approach.

D. The SIR-JPDA algorithm

The initialization of the particle set is done in the same way as for the joint filters and the prediction step as well, according to (31). The likelihood is then computed in a different way:

$$p(Y_t = (y_t^1, \dots, y_t^{m_t}) | X_t = \tilde{s}_t^n) = \sum_{\text{all associations}} p(Y_t = (y_t^1, \dots, y_t^{m_t}) | X_t = \tilde{s}_t^n, K_t^u) p(K_t^u) \quad (46)$$

$$= \sum_{\text{all associations}} \prod_{j=1}^{m_t} p(y_t^j | \tilde{s}_t^n, K_t^u) p(K_t^u) \quad (47)$$

where K_t^u denotes the u^{th} association hypothesis at time t . The associations between the measurements and the targets are established under the assumptions (A1)-(A2) exposed in Section IV-A. The *a priori* probability of this association is:

$$p(K_t^u) = \frac{\Phi^u!}{m_t!} p_F(\Phi^u) \prod_{i=1}^M P_d^{D^u(i)} \prod_{i=1}^M (1 - P_d)^{1-D^u(i)} \quad (48)$$

where Φ^u is the number of false alarms in K_t^u , $p_F(\Phi^u)$ the probability to have Φ^u false alarms, P_d the probability for a target to be detected and $D^u(i) = \{1, 0\}$ depending on whether target i is detected or not.

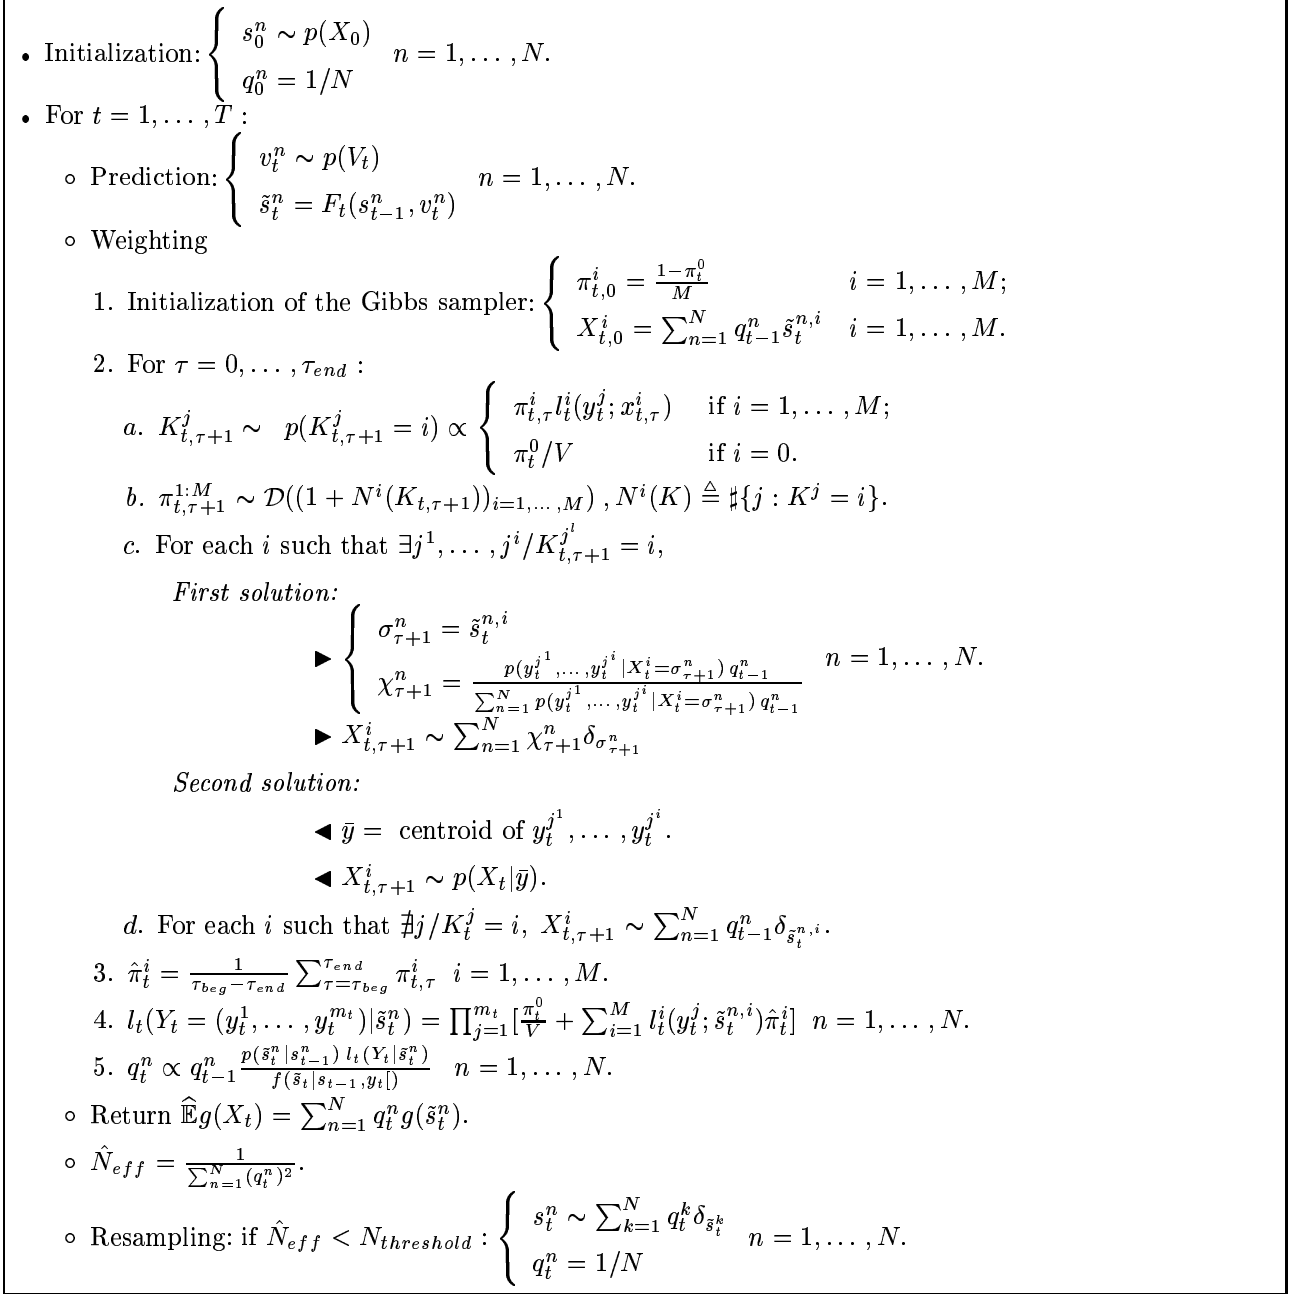


Fig. 13. Particle filter algorithm for multiple objects with adaptive resampling.

E. Extension to multireceiver data

A natural extension of the MTPF is to consider that observations can be collected by multiple receivers. Let R be their number. We will see that we can easily adapt the particle filter to this situation. We always consider that the M targets (their number is fixed again) obey the state equation (23). Some useful notations must be added to modify the measurement equations. The observation vector at time t will be denoted by $y_t = (y_{t,r^1}^1, \dots, y_{t,r^{m_t}}^{m_t})$ where r^j refers to the receiver which received the j^{th} measurement. This

measurement is then a realization of the stochastic process:

$$Y_{t,r_j}^j = H_{t,r_j}^i(X_t^i, W_t^j) \text{ if } K_t^j = i. \quad (49)$$

We assume the independence of the observations collected by the different receivers. We denote by $l_{t,r_j}^i(y; x)$ the known functions which are proportional to $p(Y_{t,r_j}^j = y | K_t^j = i, X_t^i = x)$. The likelihood of the observations conditioned on the n^{th} particle is readily obtained:

$$p(Y_t = (y_{t,r^1}^1, \dots, y_{t,r^{m_t}}^{m_t}) | X_t = s_t^n) = \prod_{j=1}^{m_t} p(y_{t,r^j}^j | s_t^n) \quad (50)$$

$$\propto \prod_{j=1}^{m_t} \left[\frac{\pi_t^0}{V} + \sum_{i=1}^M l_{t,r^j}^i(y_{t,r^j}^j; s_t^{n,i}) \pi_t^i \right]. \quad (51)$$

There is no strong limitation on the use of the particle filter for multireceiver and multitarget tracking: the MRMTPF is obtained from the MTPF by replacing the likelihood functions $l_t^i(y; x)$ by the functions $l_{t,r^j}^i(y; x)$. Moreover it can deal with measurements of varied periodicities.

V. APPLICATION TO BEARINGS-ONLY PROBLEMS

The following multitarget scenario has been considered for illustrating our algorithm. Three targets follow a near-constant-velocity model defined by (7) with $\sigma_x = \sigma_y = 0.0005 \text{ms}^{-2}$. The initial positions and velocities of the targets and of the observer are the following:

$$X_0^1 = \begin{pmatrix} 200 \text{ m} \\ 1500 \text{ m} \\ 1 \text{ ms}^{-1} \\ -0.5 \text{ ms}^{-1} \end{pmatrix}; X_0^2 = \begin{pmatrix} 0 \text{ m} \\ 0 \text{ m} \\ 1 \text{ ms}^{-1} \\ 0 \text{ ms}^{-1} \end{pmatrix};$$

$$X_0^3 = \begin{pmatrix} -200 \text{ m} \\ -1500 \text{ m} \\ 1 \text{ ms}^{-1} \\ 0.5 \text{ ms}^{-1} \end{pmatrix}; X_0^{obs} = \begin{pmatrix} 3500 \text{ m} \\ -2500 \text{ m} \\ -0.5 \text{ ms}^{-1} \\ 0 \text{ ms}^{-1} \end{pmatrix}.$$

The observer is following a leg by leg trajectory. Its velocity vector is constant on each leg and modified as follows:

$$\begin{aligned}
\begin{pmatrix} vx_{30}^{obs} \\ vy_{30}^{obs} \end{pmatrix} &= \begin{pmatrix} 1.2 \\ 0.3 \end{pmatrix}; & \begin{pmatrix} vx_{100}^{obs} \\ vy_{100}^{obs} \end{pmatrix} &= \begin{pmatrix} -3.0 \\ 0.3 \end{pmatrix}; \\
\begin{pmatrix} vx_{200}^{obs} \\ vy_{200}^{obs} \end{pmatrix} &= \begin{pmatrix} 1.2 \\ 0.3 \end{pmatrix}; & \begin{pmatrix} vx_{500}^{obs} \\ vy_{500}^{obs} \end{pmatrix} &= \begin{pmatrix} -3.0 \\ 0.3 \end{pmatrix}; \\
\begin{pmatrix} vx_{600}^{obs} \\ vy_{600}^{obs} \end{pmatrix} &= \begin{pmatrix} 1.2 \\ 0.3 \end{pmatrix} & \begin{pmatrix} vx_{800}^{obs} \\ vy_{800}^{obs} \end{pmatrix} &= \begin{pmatrix} -4.0 \\ 0.3 \end{pmatrix}; \\
\begin{pmatrix} vx_{900}^{obs} \\ vy_{900}^{obs} \end{pmatrix} &= \begin{pmatrix} 1.2 \\ 0.3 \end{pmatrix}.
\end{aligned} \tag{52}$$

The trajectories of the three targets and of the observer are plotted in Fig. 14.a.

The targets produce one measurement at each time period according to (9) with $\sigma_w = 0.02$ radians except during the time interval [600 700] where the first object does not produce any measurement. The simulated bearings are plotted in Fig. 14.b. As soon as the difference between two bearings issued from two different

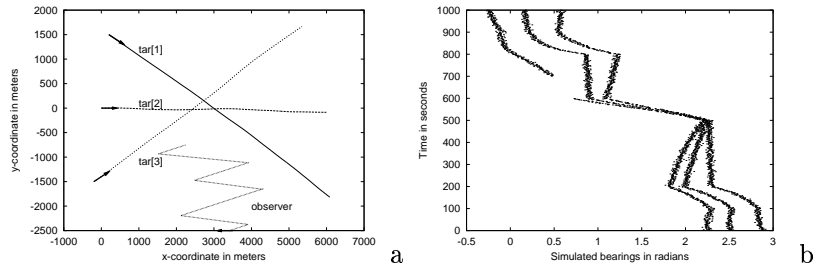


Fig. 14. *Three targets bearings-only experiment in Section V (a) Trajectories of the three targets and of the observer; (b) Simulated bearings at each time with a detection hole for one target between times 600 and 700.*

targets is lower than the standard deviation of the observation noise, the two measurements cannot be distinguished, which makes the data association of this scenario very difficult. This difficulty is increased by the detection hole for the first object. We compare the estimated trajectories when the assignment probabilities are respectively estimated by the two versions of the Gibbs sampler in p. 22: In both cases, the particle filters have been initialized according to a Gaussian law whose mean vector X_{mean} and covariance matrix X_{cov} are:

$$X_{mean}^1 = X_0^1 + \begin{pmatrix} 200 \text{ m} \\ 200 \text{ m} \\ 0 \text{ ms}^{-1} \\ 0 \text{ ms}^{-1} \end{pmatrix}; \quad X_{mean}^2 = X_0^2 + \begin{pmatrix} -200 \text{ m} \\ -200 \text{ m} \\ 0 \text{ ms}^{-1} \\ 0 \text{ ms}^{-1} \end{pmatrix}; \quad X_{mean}^3 = X_0^3 + \begin{pmatrix} -100 \text{ m} \\ 200 \text{ m} \\ 0 \text{ ms}^{-1} \\ 0 \text{ ms}^{-1} \end{pmatrix}; \tag{53}$$

$$X_{cov}^i = \begin{pmatrix} 200^2 & 0 & 0 & 0 \\ 0 & 200^2 & 0 & 0 \\ 0 & 0 & 0.05^2 & 0 \\ 0 & 0 & 0 & 0.05^2 \end{pmatrix} \text{ for } i = 1, 2, 3. \quad (54)$$

The burn-in period of the Gibbs sampler has been fixed to $N_{beg} = 100$ and the total amount of iterations to $N_{end} = 500$.

The averaged estimates with the associated 2σ confidence ellipses obtained with 1000 particles and adaptive resampling are presented in Figs. 15.a and 15.b using respectively the first and the second version of the Gibbs sampler to estimate the vector π_t . The two plots show that the data association is overcome in the two

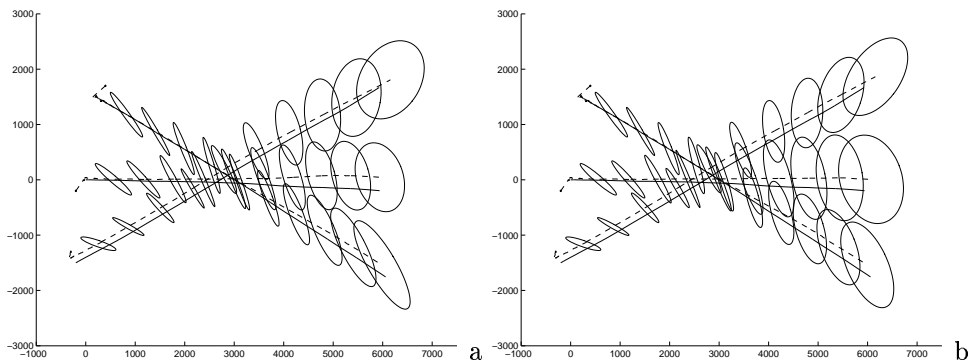


Fig. 15. Averaged estimates (dotted lines) and 2σ confidence ellipses obtained with 1000 particles. (a) estimation of π_t with the first version of the Gibbs sampler; (b) estimation of π_t with the second version of the Gibbs sampler.

versions of the algorithm. There is no trajectory reversal and the estimates are quite satisfactory for the two versions. The confidence ellipses increase when the simulated bearings are very close, i.e., especially between times 450 and 600. After this time period, their size stabilizes. The obtained estimates are very similar for both versions of the algorithm. However, the estimation of the π_t vector is much more satisfactory for the first version. Figs. 16 and 17 show the results of the estimation of the three components of π_t , respectively with the first and the second versions of the algorithm. Fig. 18 represents the average of each component




Fig. 16. The estimated components of the vector π_t obtained with the first version of the Gibbs sampler and 1000 particles: (a) $\hat{\pi}_t^1$; (b) $\hat{\pi}_t^2$; (c) $\hat{\pi}_t^3$.




Fig. 17. The estimated components of the vector π_t obtained with the second version of the Gibbs sampler and 1000 particles: (a) $\hat{\pi}_t^1$; (b) $\hat{\pi}_t^2$; (c) $\hat{\pi}_t^3$.




Fig. 18. The average of the estimated components of the vector π_t over the consecutive ten time intervals of length 100 and over 20 trials: (a) first version; (b) second version of the Gibbs sampler. The averaged estimates are marked with '+' for π^1 , 'o' for π^2 , and '*' for π^3 .

π_t^i over successive intervals of 100 time steps and over the 20 trials. At time 100 u , the average over the time interval $[100 u \ 100(u + 1)]$ and over the 20 trials is plotted for each component. When there is an ambiguity about the origin of the measurements (i.e., when the difference between the bearings is lower than the standard deviation noise), the first version makes π vary in average around $1/3$ for $M = 3$ objects and produces uniform estimation ($1/3$ for $M = 3$ objects) when the ambiguity disappears (before time 400 or after time 700 for instance). Second, the momentary measurement gap for the first object is correctly handled: the first component π_t^1 is instantaneously estimated as 0.26 in average from instant 600 to 700 and the second and third components as 0.37.

The obtained estimates with the second version of Gibbs sampler are much less satisfactory. First, the measurement gap for the first object is not detected by this version as the π_t^1 estimate does not decrease during the time interval $[600 \ 700]$. Second, as shown in Fig. 18.b, the averaged estimates diverge from instant 100 from the expected value $1/3$, whereas all the targets are detected at that time.

The expected superiority of the first algorithm seems to hold in practice, with a better estimation of π_t . However, one can be surprised that the bad estimation of π_t obtained with the second algorithm has no impact on the estimation of the targets. On this subject, it must be noticed that during the hundred first instants, the vector π_t is correctly handled by both versions. Moreover the particles recover very quickly

from their shifted initialization. Thus, from time 100, as the dynamic noise on the targets is low enough, the predicted particles are all “good”. Moreover, during the time interval [400 700], the targets produce very close bearings. Consequently, the likelihood $l_t^i(y_t^j; x_t^i)$ of a measurement given a target is almost independent from the target index: $l_t^i(y_t^j; x_t^i) \simeq l$ for all i and then $\sum_{i=0}^M l_t^i(y_t^j; x_t^i) \hat{\pi}_t^i \simeq l \sum_{i=0}^M \hat{\pi}_t^i = l$ whatever the $\hat{\pi}_t$ estimates. The estimation quality of π_t has then little impact on the likelihood of the particles during this time interval.

As for the computational cost, it takes almost 450ms per iteration with $N = 1000$ particles, $N_{beg} = 100$ and $N_{end} = 500$ on a 863 Mhz Pentium III.

VI. CONCLUSION

The multitarget tracking has been investigated in the framework of particle filtering and Gibbs sampling. Target state vectors and association probabilities are estimated jointly without enumeration, pruning or gating, by means of particle sets representing the joint *a posteriori* law of the target states. Two versions, derived from a common formalism, have been considered. This framework is sufficiently versatile to handle a wide variety of situations like multitarget tracking for multireceivers, including non-linear models.

ACKNOWLEDGMENTS

The authors are grateful to the reviewers and the editor for their helpful suggestions.

APPENDIX

Gibbs sampler

By construction, $(\theta_\tau)_{\tau \in \mathbb{N}}$ is a Markov chain. Let us prove it admits the stationary distribution $\lambda = p(\theta | Y_{0:t})$. For that, let us prove that if θ_τ is distributed according to λ then $\theta_{\tau+1}$ is also distributed according to λ . Assume $\theta_\tau = (\theta_\tau^1, \dots, \theta_\tau^P) \sim \lambda$. Then, for all (z^1, \dots, z^P) :

$$\begin{aligned} p(\theta_{\tau+1}^1 = z^1, \theta_\tau^2 = z^2, \dots, \theta_\tau^P = z^P) &= p(\theta_{\tau+1}^1 = z^1 | \theta_\tau^2 = z^2, \dots, \theta_\tau^P = z^P) p(\theta_\tau^2 = z^2, \dots, \theta_\tau^P = z^P) \\ &= p(\theta^1 = z^1 | Y_{0:t}, \theta^2 = z^2, \dots, \theta^P = z^P) p(\theta^2 = z^2, \dots, \theta^P = z^P | Y_{0:t}) \\ &= p(\theta^1 = z^1, \dots, \theta^P = z^P | Y_{0:t}). \end{aligned} \quad (55)$$

The second equality is obtained from the first in (55) as, by construction, $\theta_{\tau+1}^1 \sim p(\theta^1 | \theta_\tau^{2:P}, Y_{0:t})$. $(\theta_{\tau+1}^1, \theta_\tau^2, \dots, \theta_\tau^P)$ is then distributed according to the law $p(\theta^1, \theta^2, \dots, \theta^P | Y_{0:t})$. In the same way we can show that $(\theta_{\tau+1}^1, \theta_{\tau+1}^2, \dots, \theta_\tau^P), \dots, (\theta_{\tau+1}^1, \dots, \theta_{\tau+1}^P)$ are distributed according to $p(\theta^1, \dots, \theta^P | Y_{0:t})$. Moreover, provided that the conditional distributions $p(\theta^1 | x, \theta^2, \dots, \theta^P), \dots, p(\theta^P | x, \theta^1, \dots, \theta^{P-1})$ are strictly positive, θ_τ is irreducible. These two conditions imply the convergence for π -almost all θ_0 of $(\theta_\tau)_{\tau \in \mathbb{N}}$ to λ .

REFERENCES

- [1] D. Reid. An algorithm for tracking multiple targets. *IEEE Transactions on Automation and Control*, 24(6):84–90, 1979.

- [2] T. E. Fortmann, Y. Bar-Shalom, and M. Scheffe. Sonar tracking of multiple targets using joint probabilistic data association. *IEEE Journal of Oceanic Engineering*, 8:173–184, July 1983.
- [3] H. Gauvrit, J.-P. Le Cadre, and C. Jauffret. A formulation of multitarget tracking as an incomplete data problem. *IEEE Trans. on Aerospace and Electronic Systems*, 33(4):1242–1257, October 1997.
- [4] R. L. Streit and T. E. Luginbuhl. Maximum likelihood method for probabilistic multi-hypothesis tracking. In *Proceedings of SPIE International Symposium, Signal and Data Processing of Small Targets 1994*, volume 2235, Orlando, FL, 5-7 April 1994.
- [5] P. Willett, Y. Ruan, and R. Streit. The PMHT for maneuvering targets. In *Conference on Signal and Data Processing of Small Targets*, volume 3373, SPIE Annual international symposium on aerosense, Orlando, FL, April 1998.
- [6] C. Rago, P. Willett, and R. Streit. A comparison of the JPDA and PMHT tracking algorithms. In *Proc. Int. Conf. Acoust., Speech, Signal Processing*, pages 3571–3573, May 1995.
- [7] N. Gordon, D. Salmond, and A. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proc.F, Radar and signal processing*, 140(2):107–113, April 1993.
- [8] M. Isard and A. Blake. CONDENSATION — Conditional density propagation for visual tracking. *Int. J. Computer Vision*, 29(1):5–28, 1998.
- [9] M. Hürzeler and H.R. Künsch. Monte Carlo approximations for general state space models. *Journal of Computational and Graphical Statistics*, 7:175–193, 1998.
- [10] C. Musso and N. Oudjane. Particle methods for multimodal filtering. In *The 2nd international Conference on information fusion IEEE*, Silicon Valley, CA, 6-8 July 1999.
- [11] D. Avitzour. Stochastic simulation Bayesian approach to multitarget approach. *IEE Proc.-Radar. Sonar Navig*, 142(2):41–44, April 1995.
- [12] N. Gordon. A hybrid bootstrap filter for target tracking in clutter. *IEEE Trans. on Aerospace and Electronic Systems*, 33(1):353–358, 1997.
- [13] J. MacCormick and A. Blake. A probabilistic exclusion principle for tracking multiple objects. In *Proceedings of the 7th International Conference on Computer Vision, Kerkyra, Greece*, pages 572–578, 20-27 September 1999.
- [14] J. Mac Cormick and M. Isard. Bramble: a Bayesian multiple-blob tracker. In *Proceedings of the 8th International Conference on Computer Vision, Vancouver, Canada*, volume 2, pages 34–41, 9-12 July 2001.
- [15] D. Schulz, W. Burgard, D. Fox, and A.B. Cremers. Tracking multiple moving targets with a mobile robot using particle filters and statistical data association. In *Proceedings of IEEE International Conference on Robotics and Automation, Seoul, Korea*, pages 1665–1670, 21-26 May 2001.
- [16] A. Doucet. On sequential simulation-based methods for Bayesian filtering. Technical report, CUED/F-INFENG/TR 310, Signal Processing Group, Departement of Engineering, University of Cambridge, 1998.
- [17] MacCormick. *Probabilistic modelling and stochastic algorithms for visual localisation and tracking*. PhD thesis, University of Oxford, January 2000.
- [18] *Sequential Monte Carlo methods in practice*, chapter 25, pages 525–532. Springer, 2001.
- [19] D. J. Salmond. Mixture reduction algorithms for target tracking in clutter. In *SPIE Signal and Data Processing of small targets*, volume 1305, pages 434–445, 1990.
- [20] A. Kong, J.S. Liu, and W.H Wong. Sequential imputation method and Bayesian missing data problems. *J. Amer. Statist. Assoc.*, 89:278–288, 1994.
- [21] J.S Liu. Metropolisized independent sampling with comparison to rejection sampling and importance sampling. *Statistics and Computing*, 6:113–119, 1996.
- [22] P. Del Moral and A. Guionnet. Central limit theorem for nonlinear filtering and interacting particle systems. *The Annals of Applied Probability*, 9(2):275–297, 1999.
- [23] A. Doucet. Convergence of sequential Monte Carlo methods. Technical report, CUED/F-INFENG/TR 381, Signal Processing Group, Departement of Engineering, University of Cambridge, 2000.
- [24] X. Li and V. Jilkov. A survey of maneuvering target tracking: dynamic models. In *Conference on Signal and Data Processing of Small Targets*, Orlando, April 2000.
- [25] J.C. Hassab. *Underwater signal and data processing*, pages 221–243. CRC Press, 1989.

- [26] N. Peach. Bearings-only tracking using a set of range parameterised extended Kalman filters. *IEE Proc.-Control theory appl.*, 142(1), January 1995.
- [27] W. Grossman. Bearings-only tracking: a hybrid coordinate system approach. *Journal of Guidance, Controls, and Dynamics*, 17(3):451–457, May-June 1994.
- [28] Y. Bar-Shalom and T.E. Fortmann. *Tracking and data association*. Academic Press, 1988.
- [29] R. Karlsson and F. Gustafsson. Monte Carlo data association for multiple target tracking. In *IEE Seminar - Target Tracking: algorithms and applications*, pages 13/1–13/5, 2001.
- [30] M. Orton and A. Marrs. A Bayesian approach to multi-target tracking and data fusion with out-of-sequence measurements. In *IEE Seminar - Target Tracking: algorithms and applications*, pages 15/1–15/5, 2001.
- [31] G. Celeux and J. Diebolt. A stochastic approximation type EM algorithm for the mixture problem. *Stochastic and Stochastics Reports*, 41:119–134, 1992.
- [32] A. E. Gelfand and A. F. M. Smith. Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association*, 85:398–409, 1990.
- [33] G. Casella and E. I. George. Explaining the Gibbs sampler. *The American Statistician*, 46(3):167–174, 1992.
- [34] A. Smith and G. Roberts. Bayesian computation via the Gibbs Sampler and related Markov chain Monte Carlo methods. *J. Royal Statist. Soc., Series B* 55(1):3–24, 1993.
- [35] J. Diebolt and C. P. Robert. Estimation of finite mixture distributions through Bayesian sampling. *Journal of the Royal Statistical Society series B*, 56:363–375, 1994.
- [36] M. Stephens. *Bayesian methods for mixtures of normal distributions*. PhD thesis, Magdalen College, Oxford, 1997.
- [37] J. F. G. de Freitas. *Bayesian methods for neural networks*. PhD thesis, Trinity College, University of Cambridge, 1999.
- [38] D. Lurie and H.O. Hartley. Machine-generation of order statistics for Monte Carlo computations. *The American Statistician*, 26:26–27, February 1972.
- [39] I. Geronditis and R.L. Smith. Monte Carlo generation of order statistics from general distributions. *J. Royal Statist. Soc.*, 31(3):238–243, 1982.
- [40] J. Carpenter, P. Clifford, and P. Fearnhead. Improved particle filter for nonlinear problems. *IEE Proc.-Radar. Sonar Navig.*, 146(1), February 1999.