PHD Filter Multi-target Tracking in 3D Sonar

Daniel Edward Clark, Judith Bell,
Yves de Saint-Pern and Yvan Petillot.

Abstract— **The Probability Hypothesis Density (PHD) Filter was developed as a method for tracking a time varying number of targets without data association. The first order statistical moment of the multiple target posterior distribution called the Probability Hypothesis Density which is represented by discrete samples or particles gives the expected locations of the targets. This property is used instead of the full multi-target posterior distribution as it is requires significantly less computation and particle filter implementations have demonstrated the potential of the algorithm to be used for real-time tracking applications. In this article, an application of the Particle PHD Filter is demonstrated to track a variable number of objects in three-dimensional sonar images estimating both the number of targets and their locations. The number of targets is estimated at each iteration by computing the mass of the particle weights. The locations of the targets are determined by extracting peaks of the PHD which is a distinct task from the computation of the particles. Previous approaches have used the Expectation Maximisation (EM) algorithm to fit a Gaussian mixture model whose time complexity is quadratic in the number of targets which is not ideal for a real-time tracking application and so alternative clustering techniques are considered here. A comparison is made between the methods for the accuracy of estimation, robustness and the time taken.** [1 2]

I. Introduction

One of the goals of the subsea research community is to develop Autonomous Underwater Vehicles (AUVs), self-navigating robots which operate underwater. Such vehicles can be equipped with a range of sensors including forward-look sonar, sidescan sonar and video to enable them to navigate autonomously and undertake a range of missions, for example mine countermeasures, pipeline inspection or seabed habitat mapping. In addition, a new range of high resolution 3D acoustic imaging sensors, such as Echoscope are emerging. To enable AUVs to operate successfully, methods for detecting and tracking objects on the seabed are required to aid path planning and navigation, as well as using these techniques as an integral part of the mission. The obvious initial application is to enable the vehicle to sense its environment and prevent collision with any object. To enable it to do this effectively it must track all objects within its field of view. This will include both stationary objects on the seabed and moving objects in the water column (including marine animals, divers and other underwater vehicles operating in the vicinity). However, since the vehicle is assumed to be moving at all times, even the stationary objects will move in reference to the AUV frame of reference.

[1]The authors are in the Ocean Systems Lab, Electrical and Computer Engineering, Heriot-Watt University, Edinburgh.

The ability to track the objects during the motion may also lead to improved classification of the scene, since it will essentially provide the 4th dimension and may provide additional information as to how the return from the target varies with time and with reference to the location of the vehicle.

This paper will present a tracking technique based on the Probability Hypothesis Density, illustrating the results using data obtained from the Echoscope forward looking 3D imaging sonar. The technique has previously been demonstrated on two dimensional forward-scan sonar [1]. The paper will also discuss the clustering methods used as an integral part of the technique to extract the location of the targets, comparing a range of techniques in terms of accuracy, robustness and computational complexity.

II. The Probability Hypothesis Density (PHD)

The theory for the multiple target tracking approach used in this paper was derived by Mahler [2] from Finite Set Statistics, a reformulation of point process theory, which provided a mathematical framework for multitarget multisensor data fusion. A recursive Bayesian approach for approximating the first order statistical moment of the joint multitarget probability distribution or Probability Hypothesis Density (PHD) was proposed as an efficient means of tracking a variable number of targets, this was defined as the PHD Filter. Data association techniques are avoided as the identities of the targets are not kept, this has a significant computational advantage over traditional methods of multiple target tracking which couple single target stochastic filters such as Kalman filters [3], extended Kalman filters or particle filters [4] [5] [6] with a data association strategy [7] for determining which measurements are most likely for each individual track.

Particle filter methods for the PHD-filter have been devised by Vo [8] and Zajic [9]. Practical applications of the filter include tracking vehicles in different terrains [10], tracking targets in passive radar located on an ellipse [11] and tracking a variable number of targets in forward scan sonar [1].

This paper demonstrates an application of the Particle PHD Filter to tracking a variable number of targets in a sequence of three-dimensional sonar images. One of the advantages of the PHD Filter is its ability to track objects in heavy clutter, which is often the case in sonar data where there are many spurious measurements due to noise and reverberation. The measurements are taken in the sonar reference plane so that a stationary object in the global or world reference plane will be moving with respect to the underwater vehicle. Whilst many of the objects to be tracked will be in the world reference plane, there could also be moving objects which it may be necessary to track such as fish. Thus the ability to track a variable number of targets in the presence of missed detections and spurious measurements is advantageous in this application.

A. PHD Filter Equations

The Probability Hypothesis Density (PHD) is the first moment of the multiple target posterior distribution [2]. The PHD represents the expectation, the integral of which in any region of the state space S is the expected number of objects in S. The PHD is estimated instead of the multiple target posterior distribution as it is much less computationally expensive to do so. The time required for calculating joint multi-target likelihoods grows exponentially with the number of targets and is thus not very practical for sequential target estimation as this may need to be undertaken in real time. The model used here only calculates single target likelihoods and so is a significant improvement on explicitly calculating joint multi-target likelihoods[12].

The PHD is defined as the density, $D_{t|t}(x_t|Z_{1:t})$, whose integral:

$$\int_S D_{t|t}(x_t|Z_{1:t})\mu(dx_t) = \int |X_t \cap S| f_{t|t}(X_t|Z_{1:t})\mu(dX_t) \qquad (1)$$

on any region S of the state space is the expected number of targets in S. The estimated object states can be detected as peaks of this distribution. The dominating measure μ is an extended Lebesgue measure, described in [8].

The derivation for the PHD equations is provided by Mahler [2], the prediction and update equations are given by:

$$D_{t|t-1}(x) = \gamma_t(x) + \int \phi_{t|t-1}(x,\xi) D_{t-1|t-1}(x_{t-1})\mu(dx_{t-1}), \quad (2)$$

$$D_{t|t}(x) = \left[\nu(x) + \sum_{z \in Z_t} \frac{\psi_{t,z}(x)}{\kappa_t(z) + \langle D_{t|t-1}, \psi_{t,z} \rangle} \right] D_{t|t-1}(x), \quad (3)$$

where $\phi_{t|t-1}(x,\xi) = P_S(\xi) f_{t|t-1}(x|\xi) + b_{t|t-1}(x|\xi), \nu(x) = 1 - P_D(x)$, $\kappa_t(z) = \lambda_t c_t(z)$ and $\psi_{t,z} = P_D(x) g(z|x)$.

In the prediction equation, b_t is the PHD for spontaneous birth of a new target at time t, P_S is the probability of target survival and $f_{t|t-1}(x_t|x_{t-1})$ is the single target motion distribution. In the data update equation, g is the single target likelihood function, P_D is the probability of detection, λ_t is the Poisson parameter specifying the expected number of false alarms and c_t is the probability distribution over the state space of clutter points.

B. Particle Filter Implementation

The implementation of the PHD Particle filter employed here is an adaptation of the method described by Vo et al [8] based on a sequential Monte Carlo algorithm for multitarget tracking. The state vector for the tracker is defined as the 3D position and velocity vector: $(x, \dot{x}, y, \dot{y}, z, \dot{z})$. The algorithm proceeds as follows, where steps 1 to 4 are repeated for each iteration:

Step 0: Initialisation
In the initialisation stage, particles are uniformly distributed across the field of view.

Step 1: Prediction
The particles are propagated using the dynamic model: $x_{k+1} = F_k x_k + \omega_k$ where F_k is the transition matrix for the motion and ω_k is Gaussian noise. In addition, particles are added to allow for incoming targets into the field of view.

Step 2: Data Update
When the measurements are received, weights are calculated for the particles based on their likelihoods determined by the distance of the particles to the set of observations. The sum of the weights gives the estimated number of targets.

Step 3: Resampling
Particles are resampled from the weighted particle set. The particle distribution is an unweighted representation of the PHD.

Step 4: Target Extraction
Target locations are found by clustering the data using the estimated number of targets as the number of clusters and taking the centroids of the clusters.

III. Cluster Analysis

In Step 4 of the algorithm, the locations of the targets are determined by using a clustering algorithm to partition the data and the target positions are taken to be the centroids of the partitions. The different classes of clustering algorithms are considered here for this task.

A. Overview of different clustering algorithms

The aim of cluster analysis is to find a classification of data sets into homogeneous groups or clusters based on some discriminative criteria. Three main classes of technique are used in the literature[13]: agglomerative hierarchical clustering techniques, optimisation methods and mixture models. In hierarchical clustering, the classification contains a hierarchy of clusters from a single cluster containing all the data down to n clusters containing one item of data using a criterion for discriminating the data. If n is the number of particles, then the time complexity of the algorithm is $O(n^2 \log n)$.

The second commonly used approach to clustering data is to partition the data based on some optimisation criterion, for example minimising the sum of square errors within each group. For large data sets, comparing every pair of points is impractical so usually a finite number of iterations are computed until some convergence criterion is satisfied. An example of this kind of technique is the k-means algorithm. If n is the number of particles, k is the number of targets and t is the number of iterations, then the time complexity is $O(tkn)$.

The third approach is to fit a finite mixture to the data based on a probability model. In the case here, the number of components in the mixture would be the expected number of targets. An example of this type of technique is using the Expectation Maximisation (EM) algorithm to fit a Gaussian Mixture Model to the data. Previous approaches to target extraction with the PHD filter have focused mainly on this method[10][11][1]. The time complexity for this algorithm is $O(tk^2n)$.

If $tk << n\log n$ then k-means is the best from a purely computational point of view, clearly if we have a large number of particles then the hierarchical approach is not practical.

In the literature it is reported that the disadvantages of k-means are that it often converges to a local optimum with poor quality, it is not robust and it is statistically biased and inconsistent. The EM algorithm is theoretically unbiased, however the runtime is quadratic in the number of targets which means that the time taken for each iteration is highly dependent on the number of targets present which has proved to be the bottleneck in current implementation of the PHD filter. The implementation of k-means considered here has been shown to give better results than the standard approach using a more complex algorithm [14] and this is shown to outperform the EM algorithm for the criteria chosen for comparison. Examples of k-means and mixture model techniques have been implemented and tested on the clouds of particles within the execution of the particle PHD filter.

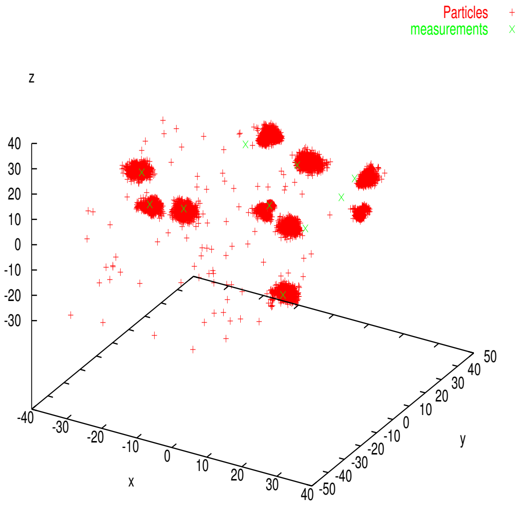

Fig. 1. Particle cloud for 10 simulated targets

B. Estimating the number of clusters

If the number of clusters is not known a-priori then it would be useful to be able to determine this directly from the data. In Bouman's unsupervised algorithm for modelling Gaussian mixtures [15], a measure of goodness of fit is found called the Rissanen criterion or Minimum Description Length (MDL) estimator. This works by attempting to find the model order which minimizes the number of bits required to code the data samples and parameter vector. The final number of clusters chosen is the value which minimises the MDL over possible values of k.

In the k-means or Lloyds algorithm, given a set of n data points the problem is to minimise the mean squared distance from each data point to its nearest centre called the average distortion. A method of determining the correct number of clusters is called v-fold cross validation which computes the average distortion for each value of k. To analyse the data the average distortion is plotted against the number of k, which exhibits a scree-plot pattern, decreases rapidly as the number of clusters increases and levels off around the true value. The correct value of k can be estimated by looking at the gradient of this graph although some judgement needs to be exercised for the specific application.

C. K-means vs EM

The k-means and EM clustering algorithms [14] [15] have been tested on particle cloud output for four iterations in the PHD algorithm for different numbers of components k in the clustering algorithms.

Ten targets have been simulated and particle outputs of the algorithm have been analysed for four of the iterations, see figure 1 for an example of the particle cloud output. To compare the accuracy of the algorithms, the particle data for the four iterations have been input into the clustering algorithms for a range of k to assess how many clusters have been correctly identified. In each of the examples, there are 10 different clusters with additional particles uniformly spread for the introduction of new targets which can be viewed as outliers for the clustering algorithms. The maximum numbers identifiable when $k < 10$ is k and 10 when $k \geq 10$, the results of this empirical evaluation are in the table. In every case, the k-means algorithm has identified at least as many of the targets as the EM algorithm and in most cases it has identified more. The run time for the algorithms is shown in figure 2. As expected, k-means is significantly faster than the EM algorithm except when $k = 1$ (in which case there is no need for a clustering algorithm at all as one can just take the mean location).

The two statistics used for estimating k are the Minimum Description Length (MDL) for the EM algorithm and the Average Distortion for k-means which calculates the mean squared distance for the data points in each partition to its centre. The minimum value for the MDL has been taken over a range of values of k for the estimated value in the EM algorithm and the derivative of the average distortion graph has been taken and the result has been tuned for the data. The estimated value for the number of clusters has been given in the table. The k-means has outperformed the EM here, getting the correct number 3 times out of 4, although the k-means has been tuned specifically for this case and the EM could possibly be tuned for a better approximation.

IV. 3D Sonar Data

The 3D sonar data was supplied and obtained by QinetiQ from an Echoscope forward-looking sonar with 64×64 beams to provide instantaneous 3D data created from a single ping rather than the 3D imagery obtained by stacking a sequence of images from a conventional 2D sonar. This results in highly

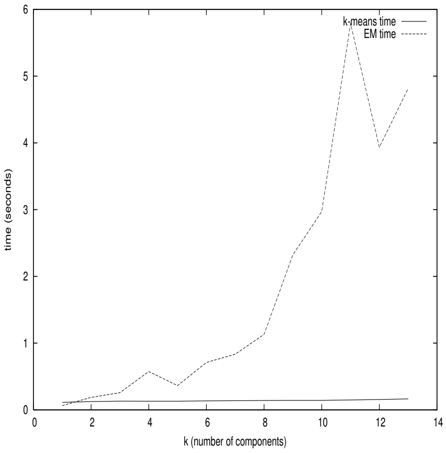

Fig. 2. Run time plot: k-means vs EM

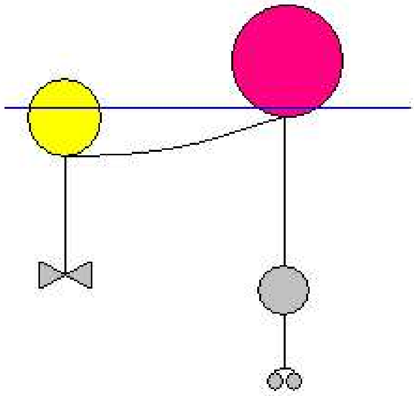

Fig. 3. Midwater target

Number of correctly determined centres								
	Iteration 2		Iteration 5		Iteration 10		Iteration 19	
k	EM	k-m	EM	k-m	EM	k-m	EM	k-m
1	0	0	0	0	0	0	0	0
2	0	0	0	0	1	0	0	0
3	0	0	0	0	1	0	0	0
4	0	1	0	0	1	1	0	0
5	0	1	0	1	1	2	1	1
6	1	2	1	3	1	2	1	3
7	1	4	2	4	3	3	2	4
8	3	6	4	6	4	5	3	6
9	6	8	6	8	3	7	6	8
10	3	10	6	8	3	7	5	10
11	7	9	10	10	1	9	6	10
12	6	10	10	10	6	10	8	10
13	7	10	9	10	8	8	8	10
?	15	10	14	10	12	10	12	10

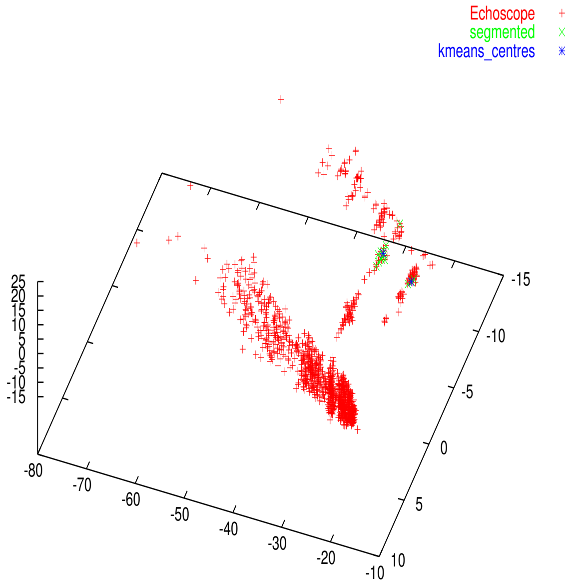

Fig. 4. Segmented Echoscope image

accurate positioning. The data is in the form of x,y,z and intensity. The sequence of images used here demonstrates the imaging of a complex midwater target involving weights of different shapes suspended below surface buoys (see figure 3). The targets which are used for the tracking are the metal objects below the buoys as they return strong intensities which can easily be found by thresholding the data on intensity. The measurements for the tracker are the centroids of the data above the threshold: an example of the 3D data with thresholded measurements is given in figure 4.

V. Target Tracking Results

The measurements obtained by thresholding the data on the intensity are input into the tracker. The sequence is of a stationary target in the water but the target positions in the sequence will change due to a slight movement of the target in the water. Initially, the particles are distributed uniformly across the region where targets can be detected. The distribution of particles represents the multimodal PHD measure from which the estimated locations are found with additional particles to allow for incoming targets. The number of particles is adapted to be proportional to the expected number of targets so that it can accurately track a variable number of targets. After the first set of measurements, the particles are given weights according to their distance from the measurements. During resampling, particles with large weights are represented more and those with low weights are killed off so that the distribution more accurately represents the positions of the targets. Figure 5 shows the initial distribution of the particles across the bounded 3D region where measurements are expected, after the first iteration the particles accumulate around the target measurements (see figure 6). As only a small proportion of the particles will represent the target location in the first iteration, these particles will be resampled heavily and

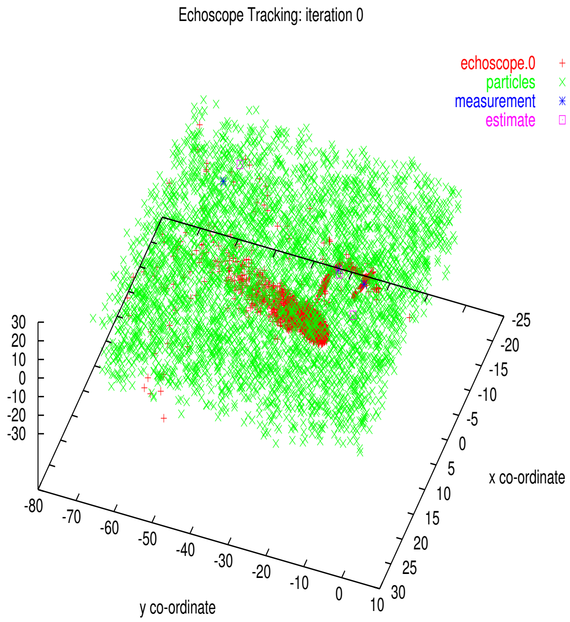

Fig. 5. Initial Particle Distribution

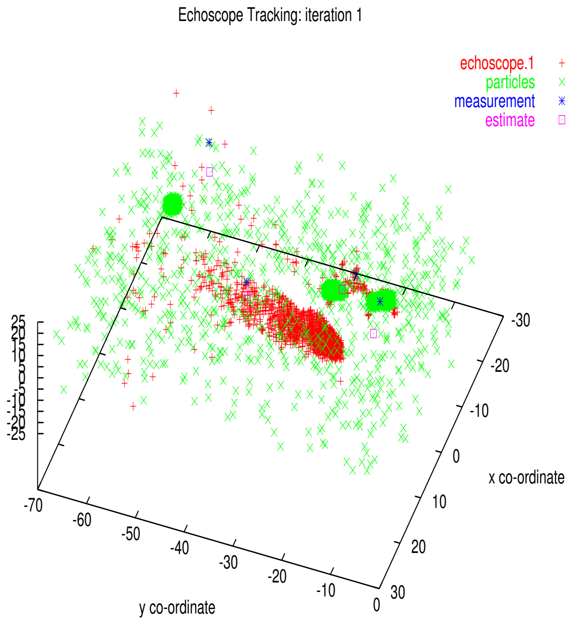

Fig. 6. Particle Distribution after first iteration

so the distribution of the particles is not particularly good. After a couple of iterations however, noise is added to the particles in the system model which will then better represent the different targets (see figure 7). Due to the noisy nature of sonar, spurious measurements are almost inevitable. The clutter parameter κ is tuned so that measurements from regions which were not predicted to have a target are not given as much prominence in the PHD distribution. This means that they will not immediately be detected as new targets and their positions extracted in the clustering. Figure 8 illustrates an example where a false alarm has entered and the particle cloud is smaller than that of the two targets.

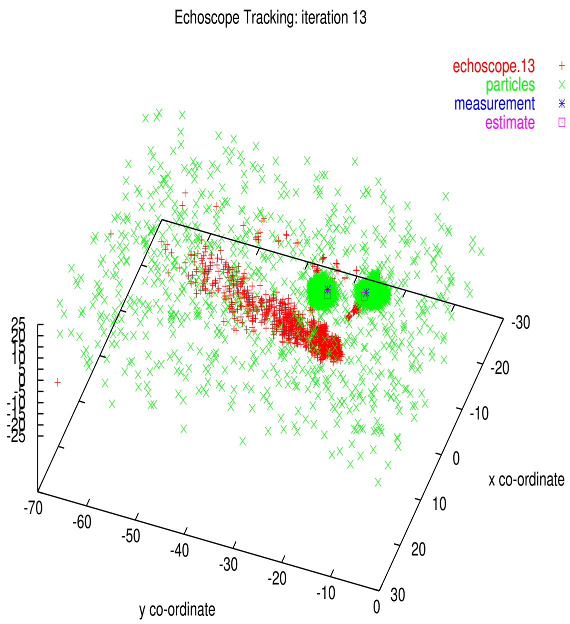

Fig. 7. Particle Distribution after a few iterations

VI. Discussion

This paper has shown the potential of using a multimodal particle filter for tracking objects in 3D Sonar data. The sonar data has been thresholded on intensity to locate the targets and the measurements for the tracker have been found by finding the centroids of the thresholded data which have been used as input to the tracking algorithm. Previous implementations of the PHD technique used Gaussian mixture models to extract the target locations. A faster, more robust technique for extracting the target locations based on k-means has been used here which enables the algorithm to track more targets faster. Future work will include investigating data association between iterations in the PHD filter for track continuity.

VII. Acknowledgements

The authors wish to thank QinetiQ for sponsoring the work and providing the data and in particular Douglas Carmichael for his valuable contribution to this work.

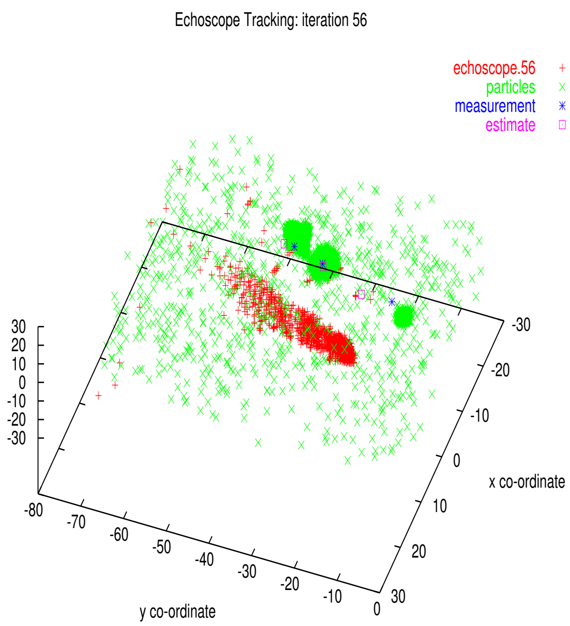

Fig. 8. Data with false alarm

REFERENCES

[1] D.E. Clark and J. Bell. Bayesian Multiple Target Tracking in Forward Scan Sonar Images Using the PHD Filter. *provisionally accepted to IEE Radar, Sonar and Navigation*, 2005.

[2] R. Mahler. Multitarget Bayes filtering via first-order multitarget moments. *IEEE Transactions on Aerospace and Electronic Systems*, 39, No.4:1152–1178, 2003.

[3] I. Tena Ruiz, Y. Petillot, D. Lane, and J. Bell. Tracking objects in underwater multibeam sonar images. *Motion Analysis and Tracking (Ref. No. 1999/103), IEE Colloquium on , 10 May 1999*, pages 11/1 – 11/7, 1999.

[4] D. Schulz, W Burgard, D. Fox, and A. B. Cremers. People tracking with a mobile robot using sample-based Joint Probabilistic Data Association Filters. *International Journal of Robotics Research*, pages 99–116, 2003.

[5] R. Karlsson and F. Gustafsson. Monte Carlo data association for multiple target tracking. *Target Tracking: Algorithms and Applications (Ref. No. 2001/174), IEE*, 1:13/1– 13/5, 2001.

[6] C. Hue, J.P. Le Cadre, and P. Perez. Tracking multiple objects with particle filtering. *IEEE Trans. on Aerospace and Electronic Systems*, 38(3):791–812, 2002.

[7] Y. Bar-Shalom and T.E. Fortmann. *Tracking and Data Association*. Academic Press, 1988.

[8] B.N. Vo, S. Singh, and A. Doucet. Sequential Monte Carlo Implementation of the PHD filter for Multi-target Tracking. *Proc. FUSION 2003*, pages 792–799, 2003.

[9] T. Zajic and R. Mahler. A particle-systems implementation of the PHD multitarget tracking filter. *SPIE Vol. 5096 Signal Processing, Sensor Fusion and Target Recognition*, pages 291–299, 2003.

[10] H. Sidenbladh. Multi-target particle filtering for the Probability Hypothesis Density. *International Conference on Information Fusion*, pages 800–806, 2003.

[11] M. Tobias and A.D. Lanterman. A Probability Hypothesis Density-based multitarget tracker using multiple bistatic range and velocity measurements. *System Theory, 2004. Proceedings of the Thirty-Sixth Southeastern Symposium on , March 14-16, 2004*, pages 205–209, 2004.

[12] D.L. Hall and J. Llinas, editors. *Handbook of Multisensor Data Fusion*, chapter 7. CRC Press, 2001.

[13] B. S. Everitt and G. Dunn. *Applied Multivariate Data Analysis*. Arnold, 2 edition, 2001.

[14] T. Kanungo, D. M. Mount, N. Netanyahu, C. Piatko, R. Silverman, and A. Y. Wu. A Local Search Approximation Algorithm for k-means Clustering. *Proc. of the 18th Annual ACM Symp. on Computational Geometry*, pages 10–18, 2002.

[15] C. A. Bouman. Cluster: An unsupervised algorithm for modeling Gaussian mixtures. Available from http://www.ece.purdue.edu/~bouman, April 1997.