UNLIMITED DISTRIBUTION

National Defence
Research and Development Branch

Défense nationale
Bureau de recherche et développement

UNLIMITED

TECHNICAL MEMORANDUM 94/201
January 1994

EVALUATION OF THE PROBABILISTIC DATA ASSOCIATION FILTER IN A REALISTIC SONAR ENVIRONMENT

William A. Roger

Defence Research Establishment Atlantic

Centre de Recherches pour la Défense Atlantique

Canada

DEFENCE RESEARCH ESTABLISHMENT ATLANTIC

9 GROVE STREET
P.O. BOX 1012
DARTMOUTH, N.S.
B2Y 3Z7

TELEPHONE
(902) 426-3100

CENTRE DE RECHERCHES POUR LA DÉFENSE ATLANTIQUE

9 GROVE STREET
C.P. 1012
DARTMOUTH, N.É.
B2Y 3Z7

UNLIMITED DISTRIBUTION

National Defence
Research and Development Branch

Défense nationale
Bureau de recherche et développement

EVALUATION OF THE PROBABILISTIC DATA ASSOCIATION FILTER IN A REALISTIC SONAR ENVIRONMENT

William A. Roger

January 1994

Approved by C.W. Bright
Director / Sonar Division

Distribution Approved by C.W. Bright

C W Bright
Director / Sonar Division

TECHNICAL MEMORANDUM 94/201

Defence Research Establishment Atlantic

Centre de Recherches pour la Défense Atlantique

Canada

Abstract

The problem of tracking an underwater narrowband signal that has experienced the many vagaries of open ocean transmission is of considerable interest. Usually the simpler signal followers developed for radar applications fail when attempting to handle the unique underwater environment encountered by a passive line array receiver. This paper investigates a particular recursive algorithm called the Probabilistic Data Association filter to determine if it can provide robust passive tracking of low signal-to-noise signals in a background of clutter caused by other sources and noise. It will be shown that two modifications to the basic algorithm substantially improve its tracking performance and its ability to estimate the target signal's trajectory. In addition, an effective technique for detecting when the filter is no longer following the signal will be presented. This loss-of-lock detector works for very low-level target signals.

Résumé

Le problème de poursuite d'un signal sous-marin à bande étroite soumis aux multiples aléas de la transmission océanique libre présente un intérêt considérable. Habituellement, les systèmes les plus simples de poursuite des signaux développés en vue d'applications radar se révèlent inefficaces dans le milieu sous-marin très particulier rencontré par les récepteurs passifs à réseau linéaire. Ce document examine un algorithme récursif particulier, appelé filtre probabiliste d'association de données, afin de déterminer s'il permet la poursuite passive robuste de signaux à faible rapport signal/bruit dans un clutter de fond provenant d'autres sources et de bruit. Il est démontré que deux modifications de l'algorithme de base améliorent considérablement le rendement de la poursuite ainsi que l'aptitude à évaluer la trajectoire du signal de la cible. De plus, le document présente une technique efficace pour détecter lorsque le filtre ne suit plus le signal. Ce détecteur de décrochage s'utilise avec les signaux de cible de très faible niveau.

Table of Contents

List of Figures

List of Tables

1. Introduction

Accurate estimation of the bearing to an acoustic source is a critical step in the narrowband passive localization problem. The estimates of bearing obtained from a towed array serve as inputs to contact motion analysis algorithms and the ultimate accuracy of the localization depends critically on the accuracy of the basic parameter extraction process. In a real ocean environment where observed signal bearings exhibit significant dynamic behavior, a standard method to improve the basic parameter extraction process is to apply a model-based, recursive filtering procedure to the time history of bearing observations. Two common examples are the Kalman filter[1] with nearest neighbour association[2] and its extension, the Probabilistic Data Association (PDA) filter[2] which better accounts for the spurious bearing measurements that result from background acoustic noise. The purpose of this paper is to analyze how well the PDA filter can follow an acoustic narrowband signal that is buried in ocean noise, and to illustrate how alterations can improve its performance. The Kalman filter will be used for comparison.

Such recursive filtering procedures have the potential for reducing the variance of the estimates from towed-array bearing and frequency measurements to provide more accurate post-processing analysis of target position. In addition, the PDA filter should provide more robust tracking of low Signal-to-Noise Ratio (SNR) signals in clutter due to preprocessing of the measurement space. Another advantage of this filter over the Kalman filter is its ability to detect when the signal is no longer being tracked. It will be shown that certain alterations to the basic PDA filter can improve its performance when the signal is barely detectable above the background noise. A Monte Carlo simulation was employed to provide an accurate comparison between filter estimates of the source bearing and frequency, and the actual values. The simulator provided measurement errors in bearing, frequency, and signal power that closely approximated those observed from an actual towed array. A practical method for detecting when the PDA filter has "lost lock" on the target signal will also be demonstrated.

The paper is organized in the following manner: In Section 2, a brief overview of the simulation of towed array measurements is presented, followed in Sections 3 and 4 by descriptions of the Kalman and Probabilistic Data Association filters. Section 5 describes modifications made to the PDA filter to enhance its tracking performance. The ability of these filters to follow a dynamic signal and the accuracy of their estimates is evaluated in Section 6. Finally, in Section 7, two methods are presented for detecting when a tracker is no longer following the signal.

2. Bearing Tracking

With towed arrays, a common approach for estimating the bearing to an acoustic source is to form multiple independent beams which span the angular space to be monitored. The beam which contains the maximum power from the signal provides an estimate of its angle from the array. However, the best accuracy one can expect with this approach is plus or minus one half of the beam-to-beam spacing, which often makes it difficult to achieve sufficient bearing accuracy for further processing. To partially circumvent this problem, knowledge of the beampattern shape of the beamformer can be used to interpolate between beams in order to provide more accurate bearing estimates[3]. Unfortunately, a towed line array does not measure the bearing to a source directly. Rather, a conical arrival angle relative to the array axis is available at the output of a phase-

to-plane beamformer[4,5] and source bearing is calculated by adding array heading which is obtained from a compass installed in the array. Note that the conical arrival angle may be a biased estimate of source bearing if there is a vertical component in the arrival angle. For the remainder of this paper it will be assumed that the vertical arrival angle is zero.

This paper utilizes computer simulations to emulate power beammaps obtained in a real world environment. At each time snapshot, a beammap was formed in the vicinity of the predicted signal position. Each frequency/beam cell had the ambient noise power modelled as an independent χ^2 sequence convolved with a noise model beampattern across beams and an FFT binpattern across frequency. The signal emulated the Swerling I model[6,7] by having the signal power weighted with a scaled χ^2 variate of unity mean and variance proportional to the time-bandwidth product associated with the simulated integration period. The bearing and frequency of the signal were corrupted by addition of variates from a zero-mean, Gaussian distribution having standard deviations similar to the measurement error found in an actual system. Then the noisy signal was mapped onto the beamformer/frequency measurement space by convolution with the appropriate beampattern and frequency binpattern. Next, the Local Maximum Data (LMD), consisting of all cells that contained simulated acoustic power greater than the surrounding eight neighbour cells, was obtained from the beammap. Parabolic interpolation[3] in both bearing and frequency was performed on each of the LMD to increase the accuracy of the measurements. Then, the normalized power, ρ, was calculated by taking the power in each cell designated as LMD and dividing it by the average background noise power. For the particular LMD generated by the signal, its normalized power ρ_s is the signal+noise-to-noise ratio expressed as $(P_s + P_n)/P_n$ in linear units, with P_s indicating the power from the signal and P_n the average noise power in an FFT bin. With this definition for the noise power a value of ρ_s of 0 dB represents a very weak signal. While the measured LMD are expressed in this fashion, the signal power for each simulation was known exactly before being perturbed by the emulated beamformer as discussed above. In the rest of the paper the term SNR will be used to denote the true signal-to-noise ratio of the signal before the beamformer, expressed as P_s/P_n (Note that ρ_s = SNR + 1).

Consequently, at each time snapshot a data set of bearing, frequency, and normalized power values served as input to the tracking filters. For some of the filters to be investigated, the input data was first preprocessed with a threshold to remove all LMD below a certain normalized power.

A simulated beammap from a single integration period is illustrated in Figure 1, where the tonal is clearly visible above the background clutter. The other peaks are caused by the ambient background noise. In Figure 2, the interpolated local maximum data obtained from this beammap are presented in a greyscale format where the levels of grey represent the normalized powers. It is the time sequence of these LMD that is processed by the recursive filters under study.

To investigate the tracking performance of the signal followers, a simulated target executed a track of straight line legs separated by large angle turns. Figure 3 illustrates the scenario, showing the constant course of the towing ship at the top and the zig-zag target track almost due South in bearing, with the signal arriving at the array with an angle in the vicinity of 90° relative to its axis. At 8-second intervals the beamformer provided measurements of bearing and frequency from the source, involving 4 non-overlapping averages, with a total of 80 samples available. In Figure 4 the true bearing and frequency of the target signal are shown for the 80 time snapshots. The scenario was repeated for a set of signal-to-noise ratios that spanned the range from very low to moderate, that is from -2 to +3 dB. At each SNR a Monte Carlo simulation involving a sample set of 200 independent examples of the scenario provided the data for a statistical analysis of tracking performance and accuracy of the bearing and frequency estimates.

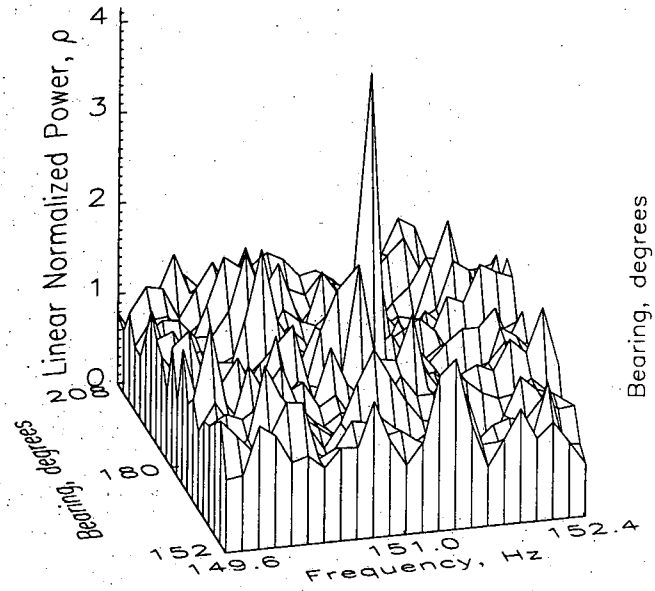

Fig. 1. Simulated beammap in bearing and frequency, showing a narrowband tonal with SNR = 5 dB

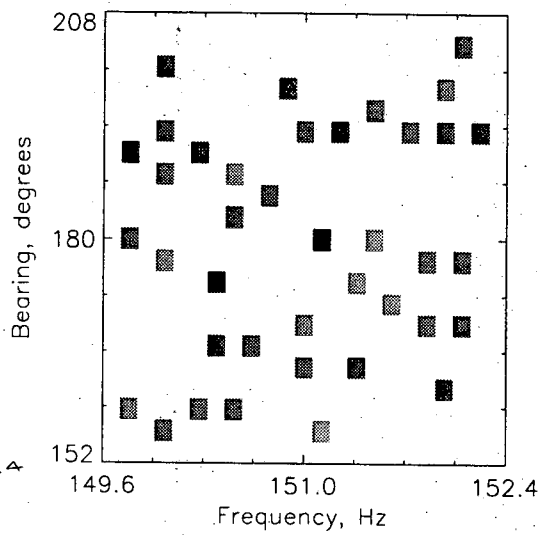

Fig. 2. LMD obtained from the beammap of Figure 1, showing bearing, frequency, and normalized power.

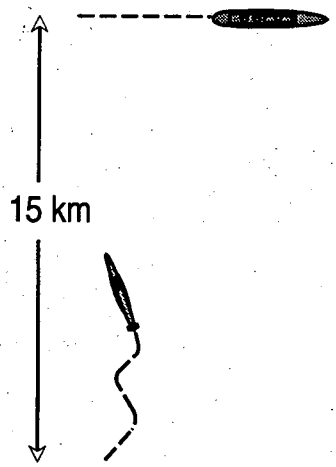

Fig. 3. Target and towing ship tracks used in Monte Carlo simulation.

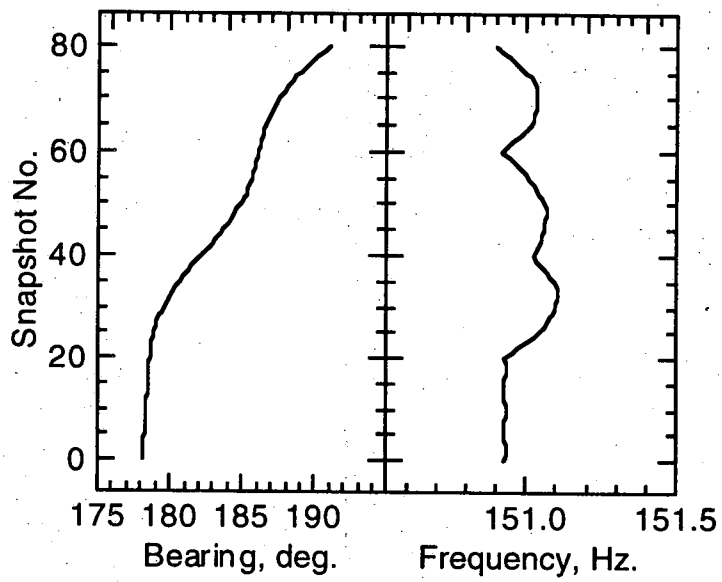

Fig. 4. True bearing and frequency of signal received at the line array over a time of 80 samples.

To provide a reference, the standard Kalman filter was employed using the so-called nearest neighbour approach[2] to determine the appropriate input. At each time snapshot, the sample of bearing/frequency/normalized power closest to the position predicted by the Kalman filter was assumed to be from the signal. It will be seen that this filter is capable of tracking the target reasonably well. However, large errors in the estimation of the signal trajectory were observed due to its susceptibility to outliers, which occur when measurements from the signal fall below a threshold level and are discarded. Then the nearest neighbour algorithm incorrectly accepts noise spikes as being from the signal, even if they are distant from the filter-predicted position.

3. The Kalman Filter

To use the Kalman filter in a signal following application one must assume that the measurements $\mathbf{z}_k$ of bearing, frequency and normalized power at time k can be modelled by the equations,

$$\mathbf{x}_k = \mathbf{F}\mathbf{x}_{k-1} + \mathbf{w}_{k-1} \tag{1}$$

$$\mathbf{z}_k = \mathbf{H}\mathbf{x}_k + \mathbf{v}_k \tag{2}$$

where $\mathbf{x}_k$ is a state vector of parameters to be estimated. Equation 1 is a difference equation which describes the time-evolution of the arriving signal via the matrix $\mathbf{F}$ and $\mathbf{w}_k$ is a random vector included to account for the uncertainty in the dynamic model. This vector is assumed to be zero-mean and Gaussian and is described by the covariance matrix $\mathbf{Q}$. Whereas Equation 1 describes the evolution of the signal, the signal observations are assumed to be generated from the state according to Equation 2 where $\mathbf{H}$ represents a linear function converting the signal state to towed array observations, and $\mathbf{v}_k$ is a zero-mean Gaussian random vector process with covariance $\mathbf{R}$ which accounts for the measurement noise. Under the above assumptions, the Kalman filter is a recursive procedure that is an optimal state estimator in the minimum least-square-error sense.

Since the detection ranges of towed arrays can be quite long, the actual bearing to a detected source generally changes slowly with time. The time-evolution of bearing can thus be described by a low-order model. The Doppler effect occasionally causes a more dynamic behavior on the received frequency, but it is not possible to exactly model these effects from the towed array measurements. Consequently the time-evolution of the signal frequency was also described by a low-order model with adequate process noise $\mathbf{Q}$ to allow tracking through these dynamic periods, as illustrated in Figure 4. Thus for the purposes of this paper the models will be limited to first order in both bearing and frequency. To further distinguish the narrowband signal from background noise spikes, the normalized power is measured and included as a state. Consequently, the state is a five-dimensional vector consisting of the bearing θ, bearing rate $\dot{\theta}$, frequency f, frequency rate $\dot{f}$, and normalized power, ρ_s, which allows Equation 1 to be written as

$$\mathbf{x}_k = \begin{bmatrix} \theta \\ \dot{\theta} \\ f \\ \dot{f} \\ \rho_s \end{bmatrix}_k = \begin{bmatrix} 1 & T & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & T & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \theta \\ \dot{\theta} \\ f \\ \dot{f} \\ \rho_s \end{bmatrix}_{k-1} + \mathbf{w}_{k-1} \tag{3}$$

where T is the time interval between beammap samples and the transition matrix $\mathbf{F}$ defines a linear model. The state-to-measurement conversion provided by Equation 2 can be written for the present case as

$$\mathbf{z}_k = \begin{bmatrix} \theta \\ f \\ \rho_s \end{bmatrix}_k = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \theta \\ \dot{\theta} \\ f \\ \dot{f} \\ \rho_s \end{bmatrix} + \mathbf{v}_k \tag{4}$$

The Kalman filter is a recursive procedure defined by the following set of equations:

$$\hat{\mathbf{x}}_k^+ = \hat{\mathbf{x}}_k^- + \mathbf{K}_k\left[\mathbf{z}_k - \mathbf{H}\hat{\mathbf{x}}_k^-\right] \tag{5}$$

$$\mathbf{P}_k^+ = (\mathbf{I} - \mathbf{K}_k\mathbf{H})\mathbf{P}_k^- \tag{6}$$

$$\mathbf{K}_k = \mathbf{P}_k^-\mathbf{H}^{\mathrm{T}}\mathbf{S}_k^{-1} \tag{7}$$

$$\mathbf{S}_k = \mathbf{H}\mathbf{P}_k^-\mathbf{H}^{\mathrm{T}} + \mathbf{R} \tag{8}$$

$$\hat{\mathbf{x}}_{k+1}^- = \mathbf{F}\hat{\mathbf{x}}_k^+ \tag{9}$$

$$\mathbf{P}_{k+1}^- = \mathbf{F}\mathbf{P}_k^+\mathbf{F}^{\mathrm{T}} + \mathbf{Q}. \tag{10}$$

In the above equations "hatted" quantities such as $\hat{\mathbf{x}}_k^-$ represent estimated values and quantities with a "-" superscripted symbol represent quantities estimated from previously estimated values based on knowledge about the signal dynamics and model uncertainty. On the other hand, symbols with a superscripted "+" represent updated values which are corrected based on the most recent observations. The symbol $\mathbf{P}_k$ denotes the state covariance matrix defined by

$$\mathbf{P}_k^- := E\left[\left(\mathbf{x}_k - \hat{\mathbf{x}}_k^-\right)\left(\mathbf{x}_k - \hat{\mathbf{x}}_k^-\right)^{\mathrm{T}}\right] \tag{11}$$

and $\mathbf{S}_k$ represents the measurement prediction covariance matrix

$$\mathbf{S}_k := E\left[\left(\mathbf{z}_k - \mathbf{H}\hat{\mathbf{x}}_k^-\right)\left(\mathbf{z}_k - \mathbf{H}\hat{\mathbf{x}}_k^-\right)^{\mathrm{T}}\right]. \tag{12}$$

Here $\mathbf{K}_k$ designates the so-called Kalman gain which is a time-varying gain matrix.

In order to run the Kalman filter on a sequence of bearing and frequency observations, initial values must be provided for $\hat{\mathbf{x}}_0$ and $\mathbf{P}_0^-$ and the quantities $\mathbf{Q}$ and $\mathbf{R}$ must be supplied. Variation in these parameters can have very pronounced effects on the performance of a Kalman filter so care must be taken in setting them. The particular values of $\mathbf{Q}$ used in this paper will be discussed in more detail when the results are presented. Since initialization of recursive filters is not being investigated here, the components of $\hat{\mathbf{x}}_0$ have been equated to the true bearing, frequency, and normalized power of the signal, while the bearing and frequency rates have been set to zero. For the next three measurements the true values of θ, f, and ρ_s were provided to allow convergence of the filter before commencement of the full noise-corrupted observations. The initialization of the covariance was handled by setting the diagonal of $\mathbf{P}_0^-$ to large values that converged to the filter-predicted estimates within a couple of time samples.

The measurement covariance $\mathbf{R}$ can be obtained from an analysis of the errors expected when measuring bearing, frequency, and normalized power. Consequently, a Monte Carlo simulation was invoked to provide the standard deviation of the parameter estimates as a function of the SNR. For bearing measurements, Figure 2 of [3] illustrates the SNR dependence for the standard deviation of the arrival angle estimates, σ_{aa}. The Monte Carlo

analysis showed that the standard deviation σ_{FFT} when estimating frequency through the FFT algorithm follows σ_{aa} but with a constant offset when expressed in logarithmic form. This occurs since the FFT can be viewed as a conventional beamformer [4,5], and consequently has the same dependence on SNR. When using the Hanning window this constant becomes 0.16:

$$\log_{10}(\sigma_{FFT}) = \log_{10}(\sigma_{aa}) + 0.16$$

Here σ_{aa} is relative to the half-power width of the beampattern mainlobe and σ_{FFT} is relative to an FFT bin, where the spectra have been convolved with a Hanning spectral window and beams by a Taylor window with the free parameter set to p=4. The standard deviation on the measurements of ρ under the Swerling model can be approximated by $\sigma_\rho = \rho P_n/\sqrt{BT}$ [8] where P_n is the background noise power and BT is the time-bandwidth product associated with the simulated integration period.

One important aspect of the Kalman filter is that it accepts a single stream of measurements. After each time snapshot, the filter accepts one measurement representing the signal, $\mathbf{z}_k$, updates the state estimate through Equation 5, then performs a one step prediction using Equation 9 that provides the data coordinates to be used in the selection of the next measurement. However, in a noisy environment containing low level signals in clutter, choosing the correct measurement may be difficult. In such cases the nearest neighbour procedure is often used, where that observation closest to the filter-predicted position is assumed to arise from the source of interest. In this study, a large window centred on the position predicted by the Kalman filter was used to limit the distance to the nearest measurement when thresholding of the data was invoked. This meant that, with very high threshold levels, occasionally no measurement was available and the filter provided a one-step prediction.

4. The Probabilistic Data Association Filter

The simplistic nearest neighbour procedure is prone to incorrect choice of the signal measurement when the signal level is low. An alternative algorithm that employs Bayesian techniques to determine the correct measurement is the Probabilistic Data Association filter[2]. It is assumed that there is a single target of interest, modelled by Equation 1, in a background clutter that can be modelled as independent identically distributed random variables with uniform spatial and spectral distributions. A hyper-ellipsoidal gate or window centred on the filter-predicted position of the next signal measurement is maintained by the PDA filter, and only observations within the gate are considered as valid. The size of the gate is determined by the level of confidence that the signal is included. The probability P_G that the signal is within the gate places an upper bound G on the square of the distance[1] d_{ik}^2 to the "edge" of the gate [9]:

$$d_{ik}^2 = \left(\mathbf{z}_{ik} - \mathbf{H}\hat{\mathbf{x}}_k^-\right)^{\mathrm{T}} \mathbf{S}_k^{-1} \left(\mathbf{z}_{ik} - \mathbf{H}\hat{\mathbf{x}}_k^-\right) \le G \qquad (13)$$

for observation i at time k. $\mathbf{z}_{ik} - \mathbf{H}\hat{\mathbf{x}}_k^-$ is the Kalman filter residual describing the difference between actual measurement i and the predicted measurement position. When P_G=0.99 it can be shown that G=11.84 for the three dimensions being considered here [9]. This value is used in the remainder of the paper. Although it may appear that the gate is of a fixed size, it is actually data-dependent due to $\mathbf{S}_k^{-1}$.

[1] Although in reality a distance squared, d^2 is commonly referred to as a distance and will be so called in the rest of this paper.

Only those observations of bearing, frequency, and normalized power that fall within the gate are considered as valid. Each is assigned a probability β_{ik} of being from the signal of interest, while β_{0k} represents the alternate hypothesis that none of the measurements originated from the target[2]. At each time snapshot k, given N_k observations within the gate and a time-dependent probability of detection P_{dk}:

$$\beta_{ik} = \frac{e_{ik}}{b_k + \sum_{i=1}^{N_k} e_{ik}}, \tag{14}$$

$$\beta_{0k} = \frac{b_k}{b_k + \sum_{i=1}^{N_k} e_{ik}} \tag{15}$$

where $$e_{ik} = e^{-d_{ik}^2/2} \tag{16}$$

and $$b_k = \frac{\left(\frac{2\pi}{G}\right)^{n_z} N_k\left(1 - P_G P_{dk}\right)}{P_{dk}\left(\frac{4\pi}{3}\right)}. \tag{17}$$

Here n_z is the number of dimensions in the measurement vector $\mathbf{z}_{ik}$, and e_{ik} is proportional to the likelihood function for observation i, given the distance d_{ik}^2 described by Equation 13.

Once the probabilities β_{0k} and β_{ik} have been calculated, the hypotheses are merged to provide a weighted sum of the residuals $\tilde{\mathbf{z}}_{ik}$ associated with the N_k observations:

$$\tilde{\mathbf{z}}_k = \sum_{i=1}^{N_k} \beta_{ik} \tilde{\mathbf{z}}_{ik} \tag{18}$$

where $$\tilde{\mathbf{z}}_{ik} = \mathbf{z}_{ik} - \mathbf{H}\hat{\mathbf{x}}_k^-. \tag{19}$$

When the signal has high SNR it is easily distinguished from the clutter, and β_{ik} for the signal approaches unity while the remaining measurements from clutter have little influence. However, for a weak signal several observations may have roughly equal β_{ik} which in turn results in a residual $\tilde{\mathbf{z}}_k$ that may not "point" to any one measurement. This increases the possibility that the filter may lose track of the signal as it moves in bearing and frequency.

5. Modifications to the PDAF Filter for Sonar Applications

Two ad hoc modifications were made to the standard PDA filter to provide enhanced tracking of low-level narrowband signals. The first involves altering the probabilities β_{ik} to improve detection of the signal measurement, and the second involves improving those parts of the state $\hat{\mathbf{x}}_k$ that involve rates of change.

Due to the dynamics involved when tracking moving contacts, it is necessary to use short integration times which in turn cause large variances on the measurement of normalized power. As a result, incorporation of ρ_s into the filter state does not greatly help to distinguish signal from clutter at low levels. To assist in the identification of the correct observation an ad hoc modification in the calculation of β_{ik} has been incorporated. This alteration improves the recognition of the signal measurement from an SNR perspective by enhancing those β_{ik} that correspond to LMD with high normalized power values:

$$\beta_{ik} = \frac{\Phi_{ik} e_{ik}}{b_k + \sum_{i=1}^{N_k} \Phi_{ik} e_{ik}} \tag{20}$$

$$\beta_{0k} = \frac{b_k}{b_k + \sum_{i=1}^{N_k} \Phi_{ik} e_{ik}}, \tag{21}$$

where $$\Phi_{ik} = \begin{cases} 0 & \text{if } \rho_i < \rho_T, \text{ or} \\ P_{nn,ik} & \text{if } \rho_i \geq \rho_T; \text{ otherwise} \\ 1 & \text{if all } \rho_i < \rho_T \end{cases}$$

First, a threshold ρ_T is established with only a small probability, say 10^{-2}–10^{-4}, that an observation from noise alone will exceed this normalized power level. If there are one or more observations within the gate and the normalized power ρ of at least one exceeds the threshold, then all measurements falling below ρ_T are deleted as arising from clutter. The likelihood functions of the remainder are modified by a probability $P_{nn,ik}$ that will be explained below. If no measurement exceeds ρ_T all observations are incorporated and the normal PDA filter is used. Consider the likely situation when only the observation from the signal exceeds the threshold. In this case, independent of the position within the window, the signal measurement is the only one used in the residual $\tilde{z}_k$, thus providing enhanced tracking.

With the PDA filter one assumes that there is one signal in background clutter within the gate from which measurements are extracted. Thus one can define the probability P_{nn} that a particular measurement did *not* arise from noise alone (and in fact contains the desired signal), while the alternate and exclusive hypothesis is that this measurement originated from clutter with probability P_n. Figure 5 illustrates this concept where an LMD measurement having a particular normalized power ρ_0 is compared to a population of noise-only LMD with the probability density shown. The probability P_{nn} that it originated from the source (not due to noise) is the area under the probability density as shown, while the probability of false alarm $P_{fa} = P_n$. The value for the threshold ρ_T can be obtained in a similar manner by choosing that normalized power that yields the desired P_{fa}.

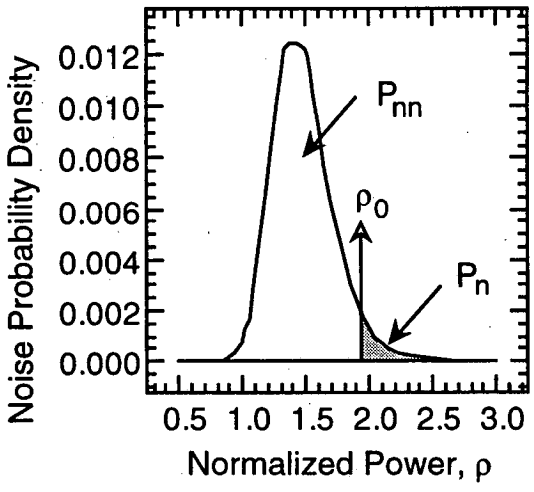

Fig. 5. Probability density distribution for noise-only LMD with BT=4, along with P_n and P_{nn} for the observation with normalized power ρ_0.

Table I. Signal Followers and their modifications.

Filter Type	Thresholded Input	β_{ik} Modified	$\hat{\theta}, \hat{f}$ Smoothed
PDAF–MS	No	Yes	Yes
PDAF–M	No	Yes	No
PDAF–T	Yes	No	No
PDAF	No	No	No
Kalman	Yes		

The second modification to the PDA filter follows from the observation that the state estimates of θ and f are noisy, since they are inferred only indirectly from the measurements of bearing and frequency and the correct model for a maneuvering target is not known. At all levels of SNR the tracking capability of the filter was improved by smoothing these two states following the update step of Equation 5, and using the smoothed estimates in the one-step prediction described by Equation 9. A median filter was chosen for the smoothing operation rather than a running average because it is less susceptible to outliers. For the following simulations the best filter length was found to be three past samples, which balanced the conflicting requirements of smoothness versus quick response to a maneuver.

6. Tracking Performance

The attribute of paramount importance for any signal following filter is its ability to accurately track, despite changes in signal characteristics caused by such effects as varying propagation conditions and target maneuvers. The function of tracking can be viewed from two perspectives. The first involves the ability of a filter to follow a weak signal without loosing lock. The second involves the ability to accurately estimate that signal's trajectory. A particular filter may be robust from one perspective but do poorly from the other. To evaluate the ability of each filter from these two aspects, the analysis was separated into two parts. The first assessed the capability of each filter to remain locked onto the signal up to the end of the scenario, while the second compared each filter's ability to estimate the true signal trajectory over the entire scenario. The following two subsections describe the results of the analysis.

6.1 Maintaining Lock on the Signal

Four forms of the PDA filter along with the Kalman filter were compared to ascertain which best provided robust signal following, and whether this tracking performance extended to a low enough SNR to be usable in an operational sonar system. Table I lists the five filter forms to be compared, with the following designations: PDAF-MS represents the PDAF with both the modifications to β_{ik} and β_{0k} described by Equations 20 and 21, and the smoothing of θ and f; PDAF-M represents the PDAF with just the modifications to β_{ik} and β_{0k}; PDAF-T represents the PDAF that receives LMD values which have first been pre-processed by a threshold filter; PDAF represents the PDAF with no modifications or input threshold; and Kalman represents the standard Kalman filter with an input threshold and nearest neighbour association.

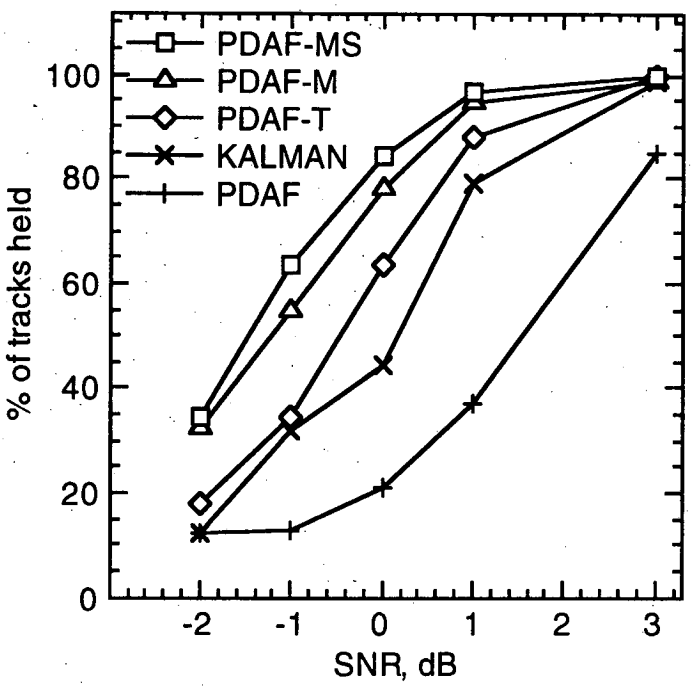

Fig. 6. Percentage of signal followers that still held lock on the signal at the end of the 80 sample scenario.

Out of the 200 Monte Carlo simulations the percentage of times each filter maintained lock on the target signal is presented in Figure 6, where for the purposes of this analysis "maintained lock" is defined to mean a filter was still tracking over the last 10 time samples in a simulation. In practice the vast majority of signal followers that lost lock at some point over the 80 time samples did not re-acquire the target; however, any filter that did re-acquire the signal before the last 10 samples is considered to have maintained lock. Note, however, that any filter prone to losing and re-acquiring the signal will provide very poor estimates of the true signal trajectory. In this paper loss-of-lock indicates that a filter is no longer tracking the signal but is following noise measurements. Since initialization of the filters was not part of the study, each filter was initialized with the true signal values and allowed to settle before the background noise was added.

Figure 6 illustrates the SNR-dependence on tracking capability for each of the filters under study. The modification to β_{ik} provides a marked improvement in tracking over the standard PDA and Kalman filters, with over 80% following a signal with SNR = 0 dB when the estimated states for bearing and frequency rates were also smoothed. It is immediately apparent from the relatively poor tracking by the standard PDA filter that some form of thresholding of the input data is crucial. This is the only filter that had no thresholding of any kind, for although both PDAF-MS and PDAF-M used the entire set of LMD on input, a threshold function was performed by the modification to β_{ik} described by Equation 20. The essential difference between the two types of thresholding occurs when all the LMD have normalized powers below the threshold value. Then the input-threshold method provides *no* measurements while the β_{ik} modification uses *all* the LMD within the filter gate. This allows the PDA filter to assess which observation is from the signal through Bayesian statistics when it is not apparent from the signal power alone.

Table II. Best threshold ρ_T and P_{fa} as a function of SNR for BT = 4.

SNR,	PDAF-MS		PDAF-M		PDAF-T		Kalman	
dB	ρ_T	P_{fa}	ρ_T	P_{fa}	ρ_T	P_{fa}	ρ_T	P_{fa}
3	2.41	10^{-3}	2.41	10^{-3}	2.66	10^{-4}	2.66	10^{-4}
1	2.41	10^{-3}	2.41	10^{-3}	2.09	10^{-2}	2.66	10^{-4}
0	2.41	10^{-3}	2.09	10^{-2}	2.09	10^{-2}	2.41	10^{-3}
-1	2.09	10^{-2}	2.09	10^{-2}	1.72	10^{-1}	2.41	10^{-3}
-2	2.09	10^{-2}	2.09	10^{-2}	1.72	10^{-1}	2.41	10^{-3}

This study showed that tracking performance can be enhanced by allowing the threshold level to be SNR-dependent. For this analysis, a threshold was defined by a given probability of false alarm, P_{fa}, which describes the probability that a measurement above the threshold ρ_T actually came from background noise alone. The optimum values of ρ_T for the scenario under study are presented in Table II, along with the corresponding P_{fa}. As expected, at higher SNR a threshold P_{fa} in the range of $10^{-3} - 10^{-4}$ provided mostly signal measurements to the filter resulting in enhanced tracking. But when following signals of lower power it proved beneficial to also lower ρ_T in order to capture sufficient measurements arising from the signal as time progressed. The exception is the Kalman filter whose susceptibility to outliers was very evident. This signal follower always required a high ρ_T in order to track the signal, which meant for lower level signals it only occasionally glimpsed an observation from the source then continued on without any input until the signal again exceeded the threshold setting. The result is a filter that meets our criterion for tracking but which provides a very poor estimate of the true signal trajectory. On the other hand, the PDAF with input thresholding, PDAF-T, required significant changes in threshold P_{fa} as SNR was varied in order to have sufficient observations that included those from the source of interest. The critical dependence of threshold P_{fa} on SNR makes this filter less desirable in practice since the large variance on the measured SNR (ρ_s) makes an accurate estimate difficult or impossible. The two modified filters were much less susceptible to threshold level since they always had measurements that included those originating from the signal, with the thresholds listed in Table II simply providing better tracking performance than with other P_{fa}.

The tracking ability of the filters shown in Figure 6 was obtained at the cost of allowing the plant covariance $\mathbf{Q}$ to be SNR-dependent. Under the constraint that the bearing and frequency components in $\mathbf{Q}$ had to be monotonic in SNR, the values in Table III represent the optimal values for $\mathbf{Q}_{ii}$ at each SNR for this scenario. In an operational system, ρ_T and $\mathbf{Q}$ could be updated through a table look-up scheme since an estimate of SNR is provided in the state vector $\hat{\mathbf{x}}_k^+$.

6.2 Estimating the Trajectory of the Signal

The accuracy of the estimation of the signal trajectory can be quantified through the Root Mean Square (RMS) error between the filter estimate and true position in bearing and frequency over time. Since the important signal parameters are bearing and frequency rather than ρ_s its estimate was not included in the RMS error calculation. At each SNR, those signal followers in the 200 simulations that were *not* tracking the signal over the last 10 time snapshots were discarded. Using the remaining simulations the mean and standard

Table III. Best variance **Q** for the filters under study at each SNR. Variance on ρ_s, $\mathbf{Q}_{55} = 10.0$ for all filters and $\mathbf{Q}_{ij} = 0$ for $i \neq j$.

SNR,	PDAF-MS		PDAF-M		PDAF-T		PDAF		Kalman	
dB	θ	f	θ	f	θ	f	θ	f	θ	f
3	0.01	0.001	1.0	0.1	0.1	0.001	0.1	0.001	0.1	0.01
1	0.01	0.001	1.0	0.1	0.1	0.001	0.1	0.01	0.1	0.01
0	0.01	0.001	1.0	0.1	0.1	0.001	0.1	0.01	0.1	0.01
-1	0.01	0.001	1.0	0.1	0.01	0.001	0.1	0.01	1.0	0.1
-2	0.03	0.003	0.1	0.01	0.01	0.01	0.01	0.01	1.0	0.1
	$\dot{\theta}$	$\dot{f}$	$\dot{\theta}$	$\dot{f}$	$\dot{\theta}$	$\dot{f}$	$\dot{\theta}$	$\dot{f}$	$\dot{\theta}$	$\dot{f}$
all	0.01	0.001	0.001	0.001	0.01	0.001	0.01	0.001	0.01	0.001

deviation of the difference between estimated and true bearing were combined into an RMS error at each of the 80 time snapshots. In a like manner, the RMS error for frequency was determined and the two were combined into a global RMS error Σ_k. Based on the beam spacing in degrees and the FFT bin resolution in Hz, the number of beams represented by an RMS error of 1° was equivalent to the FFT bins represented by an RMS error of 0.1 Hz. To give bearing and frequency errors approximately equal weight the frequency RMS errors were first multiplied by 10 before being combined into Σ_k. Finally, the global RMS error was averaged over the 80-sample track to provide a single number Σ that described the accuracy of each filter.

The average RMS error for each filter is presented in Figure 7, where it is seen that the Kalman filter provides by far the poorest estimate of signal history independent of signal level. The PDA filter with input thresholding, PDAF-T, gives good estimates for

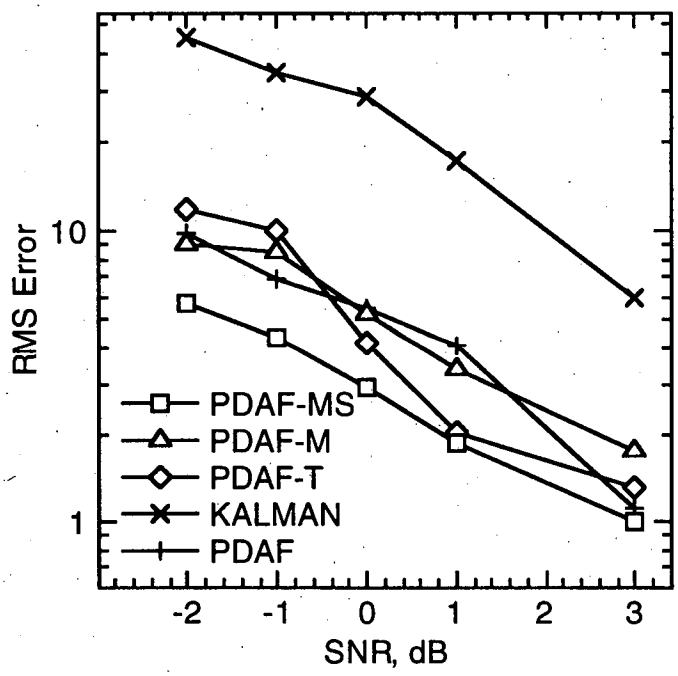

Fig. 7. The average RMS error for those signal followers that maintained lock, as a function of SNR.

moderate SNR but degrades more quickly than the others as the signal level decreases. The effect of smoothing the bearing and frequency rates is illustrated by the difference between PDAF-M and PDAF-MS, which shows that for this particular scenario smoothing approximately halves the error at all signal levels. Finally, the PDA filter with no thresholding of the input LMD exhibited similar tracking errors to the one with thresholding, but as discussed earlier it had a substantially lower tracking capability.

To give some feel to the actual RMS errors, consider the best filter (PDAF-MS) following a signal with an SNR of -2 dB. From Figure 7 the average difference between the estimated and true track represented by the RMS error is approximately 6. For an equally weighted bearing and frequency error this is roughly equivalent to an average deviation of two beams and two FFT bins from the true signal position. For this very weak signal, the filter is often unable to discern the signal from the background noise measurements. The filter tends to wander until the signal appears somewhat above the background, whereupon the modification of Equation 20 causes an enhancement of the β_{ik} associated with the signal measurement, and the filter jumps over to the signal position. When the filter is following a signal with SNR at or above 0 dB, the average error over the 80-sample track falls to less than one beam and one FFT bin.

7. The Detection of Signal Loss-of-Lock

A critical aspect of any signal following system is its ability to alert an operator when there has been loss-of-lock on the target signal. In most experimental towed array systems the signal-following function is a background operation which does not require operator intervention once started. Whenever a filter fails to follow the signal it will continue on but process just background random noise. If there is no mechanism to flag this condition then all post-processing algorithms depending on the signal follower data will in turn generate erroneous information. With the Kalman filter, monitoring of the innovations $\tilde{\mathbf{z}}_k$ is often used. However, this can be inconclusive since a deviation from a white, zero-mean Gaussian sequence can indicate either loss of the signal or detection of a maneuver. In addition, for low level signals simulation has shown that this method is very unreliable for either purpose. On the other hand, the PDA filter accepts both signal and clutter and assigns a probability to each observation through β_{ik}. The modification described in Equation 20 substantially improves the discrimination between signal and noise, and so enhances the detection of a loss-of-lock condition. Another parameter that can indicate whether the filter is tracking the signal or just noise is P_{nn}, which provides the probability that a particular measurement has not originated from clutter. By monitoring the history of P_{nn} assuming the filter's state estimate of ρ_s, $\hat{\mathbf{x}}_\rho$, one can deduce when the filter has lost track of the target.

A reliable loss-of-lock detector can be formed by comparing either the maximum β_i at time k, $max_k(\beta_i)$, or $P_{nn,k}(\hat{\mathbf{x}}_\rho)$ against a corresponding threshold. When either parameter falls below its threshold a loss-of-lock condition is announced. To improve performance both $max_k(\beta_i)$ and $P_{nn,k}(\hat{\mathbf{x}}_\rho)$ are smoothed with a median filter over the preceding time snapshots. For moderate SNRs either detector can be used, but for signals that are barely discernible, logical *or*-ing the output from these two detectors produces a more sensitive detection mechanism. Improved response is obtained by adjusting both the median filter length and the threshold levels as a function of SNR. Table IV lists the values for filter length and threshold used in the study, which where chosen to meet the conflicting requirements of low probability of false alarm and rapid and robust detection of loss-of-lock.

Table IV. Length of median filter and probability threshold for *max*(β) and P_{nn} detectors as a function of the SNR.

SNR, dB	Filter Length	β_k Threshold	P_{nn} Threshold
≥3	7	0.65	0.85
1	11	0.60	0.80
0	23	0.53	0.77
≤ -1	33	0.53	0.77

A different scenario was constructed to analyze how well the best filter, PDAF-MS, could detect loss-of-lock over a range of SNR. Both the target and towing ship maintained a steady course and speed to minimize the dynamics on θ and f, and the simulated true SNR received by the line array remained constant until sample time k reached 30. Then the SNR decreased at a linear rate until it reached the unobservable level of -6 dB at sample time 40, remaining at this value until the end of the 80 sample scenario. The Monte Carlo simulation included 200 runs but those signal followers that lost contact before k=30 were eliminated. Thus the remainder maintained lock over the time the signal was observable. Figure 8 presents the percentage of filters whose β_k-or-P_{nn} detectors declared loss-of-lock at each time snapshot k. Only at the lowest SNR did one filter erroneously declare loss-of-lock before the signal faded. With heavier smoothing at the lower signal levels the detectors were more sluggish in declaring loss of the signal, but all did detect the absence of the target by the end of the scenario. Even for the lowest SNR, 95% of the filters had declared loss-of-lock within 20 samples of the signal falling below the average noise level. For an integration time of 8 seconds this represents detection of signal loss within 2.6 minutes, while for a signal initially at 3 dB the time lag for 95% of the filters was about 5 samples or 40 seconds.

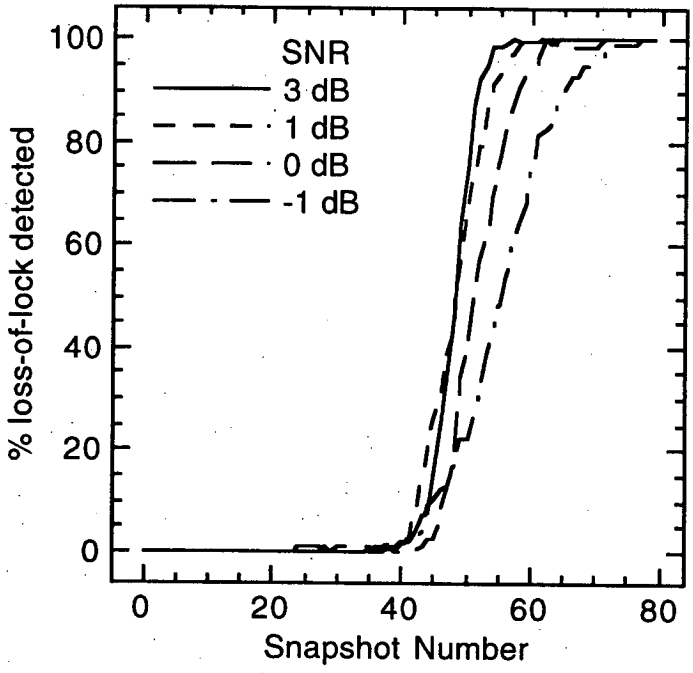

Fig. 8. The percentage of detectors that declared loss-of-lock at each sample time k, as a function of SNR.

8. Conclusions

In this paper we have examined the performance of several variations of the Probabilistic Data Association filter. It has been demonstrated that two modifications can enhance both tracking ability and estimation of the target signal trajectory. By modifying the parameter β_{ik} that assesses the probability that observation i is from the signal, one can reduce the clutter within the gate and hence provide reasonable tracking even for barely detectable signals. By also smoothing those filter states that describe rates of change in bearing and frequency the accuracy of the estimate of the signal's true position over time is further increased.

A prime requirement for any signal following filter is the ability to recognize when the signal of interest has been lost. An effective loss-of-lock detector has been demonstrated for the PDA filter using the modified Bayesian probability β_{ik} and the probability P_{nn}. This detector was shown to be sensitive even for very low level signals.

From the analysis described in this paper it has been demonstrated that the Probabilistic Data Association filter, with certain modifications, can provide robust tracking of a very low level signal emanating from a maneuvering target.

References

[1] Gelb, A., Ed., *Applied Optimal Estimation*, MIT Press, Cambridge MA 1974.

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

[3] Roger, W.A. and Walker, R.S., "Accurate Estimation of Source Bearing from Line Arrays," *Proc. Thirteenth Biennial Symposium on Communications*, June 2-4, 1986, pp. B.4.1–B.4.4, Queens University, Kingston, Ontario

[4] Van Veen, B.D. and Buckley, K.M., "Beamforming: A Versatile Approach to Spatial Filtering," *IEEE ASSP Magazine*, pp. 4–24, April 1988.

[5] Burdic, W.S., *Underwater Acoustic System Analysis*, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1984.

[6] Walker, R.S., "The Detection Performance of FFT Processors for Narrowband Signals," *Defence Research Establishment Atlantic Technical Memorandum 82/A*, 1982.

[7] DiFranco, J.V. and Rubin, W.L., *Radar Detection*, Prentice-Hall, Englewood Cliffs, NJ, 1968.

[8] Bendat, J.S. and Piersol, A.G., *Random Data: Analysis and Measurement Procedures,* Wiley-Interscience, New York, NY, 1971.

[9] Blackman, S.S., *Multiple-Target Tracking with Radar Applications*, Artech House Inc., Norwood, MA, 1986.

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1. ORIGINATOR (The name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Establishment sponsoring a contractor's report, or tasking agency, are entered in section 8.)

Defence Research Establishment Atlantic
P.O. Box 1012, Dartmouth, N.S. B2Y 3Z7

2. SECURITY CLASSIFICATION (Overall security of the document including special warning terms if applicable.)

UNCLASSIFIED

3. TITLE (The complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S,C,R or U) in parentheses after the title.)

Evaluation of the Probabilistic Data Association Filter in a Realistic Sonar Environment

4. AUTHORS (Last name, first name, middle initial. If military, show rank, e.g. Doe, Maj. John E.)

Roger, Willian A.

5. DATE OF PUBLICATION (Month and year of publication of document.)

January 1994

6a. NO. OF PAGES (Total containing information. Include Annexes, Appendices, etc.)

21

6b. NO. OF REFS. (Total cited in document.)

9

6. DESCRIPTIVE NOTES (The category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

Technical Memorandum

8. SPONSORING ACTIVITY (The name of the department project office or laboratory sponsoring the reseach and development. Include the address.)

Defence Research Establishment Atlantic
P.O. Box 1012, Dartmouth, N.S. B2Y 3Z7

9a. PROJECT OR GRANT NUMBER (If appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant.)

DRDM 12

9b. CONTRACT NUMBER (If appropriate, the applicable number under which the document was written.)

10a. ORIGINATOR'S DOCUMENT NUMBER (The official document number by which the document is identified by the originating activity. This number mrst be unique to this document.)

DREA Technical Memorandum 94/201

10b. OTHER DOCUMENT NUMBERS (Any other numbers which may be assigned this document either by the originator or by the sponsor.)

11. DOCUMENT AVAILABILITY (Any limitations on further dissemination of the document, other than those imposed by security classification)

(X) Unlimited distribution
( ) Distribution limited to defence departments and defence contractors; further distribution only as approved
( ) Distribution limited to defence departments and Canadian defence contractors; further distribution only as approved
( ) Distribution limited to government departments and agencies; futher distribution only as approved
( ) Distribution limited to defence departments; further distribution only as approved
( ) Other (please specify):

12. DOCUMENT ANNOUNCEMENT (Any limitation to the bibliographic annoucement of this document. This will normally correspond to the Document Availability (11). However, where futher distribution (beyond the audience specified in 11) is possible, a wider announcement audience may be selected.)

13. ABSTRACT (A brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both offical languages unless the text is bilingual)

The problem of tracking an underwater narrowband signal that has experienced the many vagaries of open ocean transmission is of considerable interest. Usually the simpler signal followers developed for radar applications fail when attempting to handle the unique underwater environment encountered by a passive line array receiver. This paper investigates a particular recursive algorithm called the Probabilistic Data Association filter to determine if it can provide robust passive tracking of low signal-to-noise signals in a background of clutter caused by other sources and noise. It will be shown that two modifications to the basic algorithm substantially improve its tracking performance and ability to estimate the target signal's trajectory. In addition, an effective technique for detecting when the filter is no longer following the signal will be presented. This loss-of-lock detector works for very low-level target signals.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (Technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus. e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus-identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

Passive Sonar
Probabilistic Data Association Filter
Recursive Filtering
Signal Following
Tracking Problem
Sound Localization
Kalman Filter
State Estimation
Tracking Filters
Underwater Acoustics
Signal to Noise Ratios
Towed Array

D
R
E
A

C
R
D
A