Track-Before-Detect for an Active Towed Array Sonar

Han X. Vu (1,2), Samuel J. Davey (1,2), Fiona K. Fletcher (1), Sanjeev Arulampalam (1,2), Richard Ellem (1) and Cheng-Chew Lim (2)

(1) Defence Science and Technology Organisation, AUSTRALIA
(2) School of Electrical and Electronic Engineering, University of Adelaide, AUSTRALIA

ABSTRACT

Conventional active sonar processing systems typically reduce the sensor data from an intensity map to a point-measurement form via a detection thresholding process. This approach is often sufficient for detecting and tracking high signal-to-noise-ratio (SNR) targets but becomes more challenging for low SNR targets. Track-Before-Detect (TkBD) is an alternative tracking technique that supplies raw sensor data to the tracker to determine the presence and kinematic state of a target. This paper considers the application of TkBD to an active sonar tracking problem using the Histogram-Probabilistic Multi-Hypothesis Tracking (H-PMHT) algorithm. To aid in the detection of targets in the sensor image, the measurement model for the standard H-PMHT algorithm is extended to incorporate a bearing-dependent point spread function. Using real data gathered from sonar trials, the performance of the resulting multi-target TkBD algorithm is shown to be more robust at low SNR levels when compared against a conventional point measurement tracker based on Integrated Probabilistic Data Association.

1 INTRODUCTION

The detection and tracking of low signal-to-noise-ratio (SNR) targets using an active towed array sonar is a non-trivial problem due to the complex nature of the underwater acoustic environment. Active sonar detection and tracking is particularly challenging in littoral environments where performance can be degraded by high levels of acoustic clutter that result in false alarms, fluctuating target returns due to complex time-varying acoustic conditions, and a relatively low sonar data update rate.

Traditionally, a sonar system's detection and tracking capabilities have been considered separate functions. Conventional active sonar processing systems use beamforming and matched filter correlation with a replica of the transmitted pulse to generate a sensor image of the reflected acoustic intensity as a function of range and bearing. Typically, this sensor image is normalised to remove mean background variations and a fixed threshold is applied to produce detections that are then provided to the tracker. The role of the tracker is to associate point-measurements from a common target across time and return estimates of the target's trajectory. This approach is often sufficient for detecting and tracking high SNR targets but becomes more challenging for low SNR targets, as the process of reducing the sensor image to thresholded detections discards valuable target information.

An example of active sonar tracking using conventional non-linear Kalman filtering techniques on simulated measurements is provided by [2]. A comparison of a number of conventional multi-target point-measurement trackers based on Global Nearest Neighbour and Probabilistic Data Association (PDA) using a simulated active sonar environment was addressed in [14]. Techniques based on the inclusion of target and clutter amplitude information within conventional point-measurement trackers to improve tracking performance for active sonar have also been considered [1, 10].

The desire for more robust trackers against low SNR targets in clutter-rich environments has resulted in the development of alternative techniques that perform concurrent detection and tracking, commonly referred to as Track-Before-Detect (TkBD). TkBD algorithms eliminate the thresholding process and directly use the sensor intensity data to determine the presence and kinematic state of a target. A review of early TkBD algorithms for applications in image sequences, radar and sonar is provided by [18]. These techniques have the potential to provide significant gains for both detecting and tracking low SNR targets in high clutter scenarios [8].

The first applications of TkBD to active sonar were based on dynamic programming techniques, which use a fixed grid to model the propagation of target states with time [4, 9, 18]. Furthermore, most of these techniques have been demonstrated with simulated data and the application of TkBD to real sea trials data has been limited [15].

This paper considers the application of an alternative TkBD technique called Histogram-Probabilistic Multiple Hypothesis Tracking (H-PMHT) [11, 16] to an active sonar tracking problem. In addition, the standard H-PMHT algorithm is extended to incorporate a bearing-dependent point spread function. The performance of the H-PMHT is compared with a conventional point-measurement tracker based on Integrated Probabilistic Data Association (IPDA) [6, 12]. The benefits of TkBD over conventional tracking is analysed in two representative acoustic environments using trials data from an active towed array sonar system.

The paper is arranged as follows. Section 2 outlines the tracking problem for both the conventional and TkBD case. Section 3 gives a brief review of the IPDA approach to conventional tracking as well as an introduction to the H-PMHT algorithm and its modification to incorporate a bearing-dependent point spread function. A comparative study of the two algorithms is presented in Section 4 using real data gathered from an active towed array sonar system. Section 5 summarises and discusses avenues for future work.

2 ACTIVE SONAR PROBLEM

In target tracking, the main objective is to identify the number of targets and estimate their trajectories over time using a sequence of noisy measurements. In practice, it is common to assume stochastic models for the system target dynamics and sensor data.

Define the state vector $\boldsymbol{x}_t$, which evolves with time $t \in \mathrm{N}$, where N is the set of all natural numbers. For conventional active sonar target tracking, it is sufficient to describe the target state using position and velocity in two-dimensions. However in the TkBD case, it is common to supplement the state vector with the target amplitude,

$$x_t^m = \left[x_t^m \;\; \dot{x}_t^m \;\; y_t^m \;\; \dot{y}_t^m \right]^T , \quad (1)$$

where $m = 1,\ldots, M$ denotes the target index in the case of a multiple target scenario and M denotes the total number of targets.

Generally, when analysing a dynamic system, two models are required; the target and measurement models. The target model describes the target state evolution with time and can be expressed in terms of a linear discrete-time stochastic model

$$\boldsymbol{x}_t^{\boldsymbol{m}} = \boldsymbol{F}_{t-1} \boldsymbol{x}_{t-1}^{\boldsymbol{m}} + \boldsymbol{v}_{t-1} , \quad (2)$$

where $\boldsymbol{F}_{t-1}$ is a known matrix describing a linear state transition from $\boldsymbol{x}_{t-1}^{\mathrm{m}}$ to $\boldsymbol{x}_t^{\mathrm{m}}$, and $\boldsymbol{v}_{t-1}$ is an independent identically distributed (i.i.d.) system noise sequence representing uncertainties in the target motion.

For the active sonar tracking problem, the target model needs to capture the dynamics of an underwater target. In this paper, it is assumed that a nearly-constant-velocity model is sufficient.

The second model required is the measurement model, which relates the noisy measurements to the state $\boldsymbol{x}_t^{\mathrm{m}}$. In conventional sonar tracking, it is assumed that point-measurements are received in range and bearing, and are given by

$$\zeta_t = \eta_t (x_t^{\boldsymbol{m}}) + \mu_t , \quad (3)$$

where $\boldsymbol{\mu}_t$ is an i.i.d. measurement noise sequence and $\eta(\boldsymbol{x}_t^{\mathrm{m}})$ denotes the measurement model that maps the state into the measurement space. The set of received measurements which satisfy the threshold crossing requirements may originate from either clutter or target components and it is not guaranteed that returns from each target will satisfy the threshold.

In the TkBD case, the measurement model relates images $\boldsymbol{z}_t$ in range and bearing to the target state $\boldsymbol{x}_t$. Let $\boldsymbol{z}_t^{\mathrm{i}}$ denote the ith pixel in the measurement image at time t, and let $\boldsymbol{z}_t = \{\boldsymbol{z}_t^{\mathrm{i}}\}_{\mathrm{i}=1,\ldots\mathrm{I}}$ represent a stacked vector of all the pixels in the image, where i is the pixel index and I is the total number of pixels in the measurement image. For ease of presentation, we have used a stacked vector to represent the image to allow single index referencing. A two dimensional representation could just as easily have been used. We assume a point-scatterer target, such that the target contribution to the measurement image can be described purely in terms of the point spread function (psf), $h\,(\boldsymbol{x}_t^{\mathrm{m}})$

$$z_t^m = \sum_{m=1}^{M} A_t^m h(x_t^m) + w_t , \quad (4)$$

where $\boldsymbol{w}_t$ is an i.i.d. measurement noise sequence. Note that the psf is a property of the sensor and is the same for all targets, but can vary with different sensors. Assuming independent pixel noise, the likelihood for the image $\boldsymbol{z}_t$ can be factorised as

$$p(z_t \mid x_t^m) = \prod_{i=1}^{I} p(z_t^i \mid x_t^m) . \quad (5)$$

For active sonar tracking, the measurement process only observes the position component of the target, thus the measurement model assumes that the likelihood is independent of the target velocity component. To prevent numerical problems, the log of the likelihood is also calculated.

Under linear Gaussian assumptions, the optimal finite dimensional solution to the discrete-time recursive Bayesian state estimation problem is the Kalman Filter (KF). However as the sensor image is a function of range and bearing, approximations or suboptimal solutions must be considered to accommodate the resulting non-linear measurement model. Furthermore, the TkBD measurement model defined by (4) is clearly a highly non-linear function of the target state. Thus only non-linear tracking methods can be considered when using TkBD techniques. In the next section, a brief review of the IPDA approach for conventional tracking and the H-PMHT approach for TkBD will be given.

3 TRACKING ALGORITHMS

3.1 Integrated Probabilistic Data Association (IPDA) for Conventional Tracking

One of the main difficulties in point-measurement tracking is determining which measurements arise due to a particular target and which measurements are the result of false alarms or due to objects that are not of interest. This problem is referred to as data association. Under a particular data association hypothesis, the target state can be estimated using a KF or non-linear counterpart. Probabilistic Data Association (PDA) expresses the target state probability density function (pdf) as a sum over data association hypotheses and provides expressions to determine the probabilities of these hypotheses [6]. The resulting pdf is a Gaussian mixture, which is then approximated by a single Gaussian. The Integrated PDA (IPDA) extends the target state-space by defining a binary existence variable that indicates whether or not there is actually a target present and assumes that this variable evolves according to a Markov Chain. The IPDA provides equations for recursively updating the target states and the probability of target existence based on the PDA approach for data association [12]. The probability of existence can then be used to automate track management.

The advantages of IPDA are that it is computationally inexpensive and can be implemented as a modified KF: the target description consists only of a mean vector, a covariance matrix and a scalar existence probability. However, the algorithm uses a single Gaussian component to approximate the Gaussian mixture arising from the sum of data association hypotheses.

This approximation can be poor, especially when the mixture has more than one dominant component. IPDA also assumes the existence of at most one target.

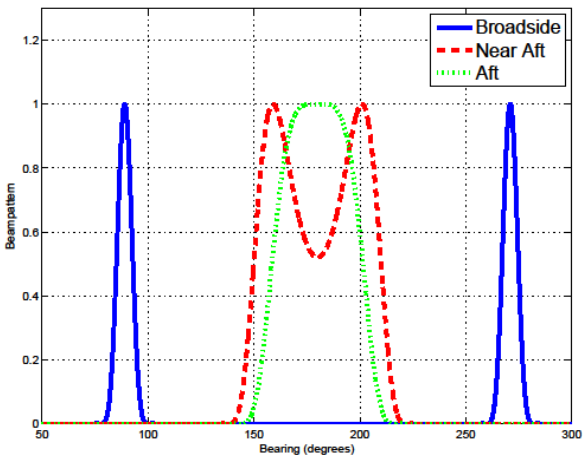

Figure 1. Beam pattern vs. bearing for transmissions at broadside, near aft and aft of the array.

In this paper, a single target IPDA tracker based on [12] is implemented. However, a multi-target track management logic is imposed based on the assumption that targets in the image do not overlap or interact with each other. Joint Integrated Probabilistic Data Association (JIPDA) can be used for scenarios featuring interacting targets but it is more computationally expensive than IPDA [13] and is not considered here. Tracks are initiated when the estimated probability of track existence rises above a certain threshold. Likewise, track termination occurs when the probability of track existence falls below another threshold.

3.2 H-PMHT for TkBD

The H-PMHT algorithm is an efficient multi-target approach to the TkBD problem. The technique is based on the generation of a synthetic histogram by quantising the energy in the sensor data followed by the application of Expectation-Maximisation (EM) mixture modelling to describe the underlying data sources. The algorithm name stems from the interpretation of the quantised data as a histogram and the use of the Probabilistic Multi-Hypothesis Tracking (PMHT) algorithm [17] to perform the association of the resultant measurement counts to target objects. In the final step of the derivation, the limit of the quantisation is taken and the original sensor data is recovered.

Unlike other TkBD approaches, the H-PMHT algorithm employs a parametric fitting approach and has been shown to give performance that is close to the optimal Bayesian filter at a fraction of the computational cost [8]. The H-PMHT can be naturally extended to track in a multi-target scenario but still has linear complexity with the number of targets. The H-PMHT can also be applied to a wide range of problems as long as an appropriate state estimator exists to perform the EM step of the algorithm.

The EM algorithm operates by iteratively optimising an auxiliary function. In the case of H-PMHT, this function consists of a target state evolution term, a mixing term dependent on the target SNRs, and a measurement term that couples the target states with the intensity map. This measurement term is the logarithm of the measurement likelihood weighted by data association probabilities, i.e. the logarithm of equation (5) with pixel-dependent scaling terms. When the sensor psf $h(\mathbf{x}_t^m)$ can be approximated as a Gaussian then the logarithm

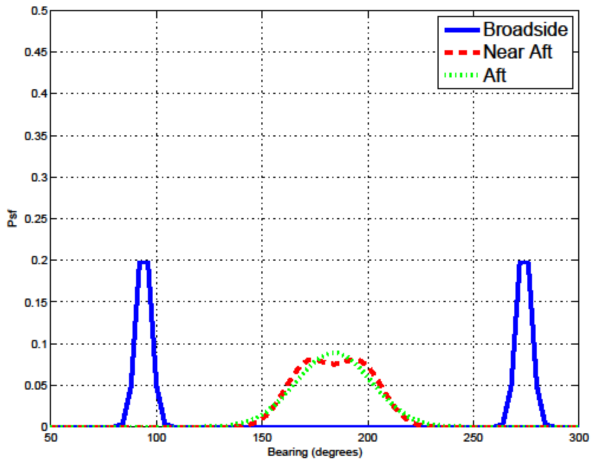

Figure 2. Point spread function $h_\theta(\theta)$ vs. bearing for transmissions at broadside, near aft and aft of the array.

of equation (5) is a sum of quadratics and this can be factorised into a single quadratic, equivalent to the log-likelihood of a point-measurement [7]. Thus H-PMHT can be implemented as an iterated data association step and point-measurement estimation step. It is important to emphasise that this is not an approximation. In this paper, the maximisation step of the H-PMHT is performed using an extended Kalman Filter (EKF). However, the linearisation point for the EKF is modified with each EM iteration so the result is similar to the Iterated EKF, which is known to provide a more accurate estimate than the standard EKF [3].

Next, we will describe how the standard measurement model for the H-PMHT algorithm can be modified to include a bearing-dependent psf to model the variation in the beam patterns for an active towed array sonar.

Due to left-right ambiguity issues characteristic of linear array systems, when a target is detected by a towed array, it will appear as two identical targets symmetrically placed on either side of the towed array in the sensor image. An ownship manoeuvre is required to identify a real target from the ambiguous one. Figure 1 shows representative beam patterns and the effect of left-right ambiguity issues for an active towed array sonar system at the following receive directions:

- Broadside: defined as 90 degrees from the heading of the array. The beam pattern consists of a narrow beam and the two peaks generated by left-right ambiguity are well-separated,
- Near Aft of the array: where the beams are wider and the left-right ambiguous peaks begin to overlap, and
- Aft of the array: defined as close to 180 degrees from the heading of the array, where the left-right beam patterns merge into a single wide peak.

It can be seen that as the receive direction moves from broadside to aft of the array, the spread in the beam pattern increases and the ambiguous target in the image merges with the real target to form a single target smeared across multiple bearing bins in the image sensor. The half-height beamwidth for the beam pattern is approximately 8 degrees at broadside, 50 degrees (across both peaks) in the near aft direction and 42 degrees at aft.

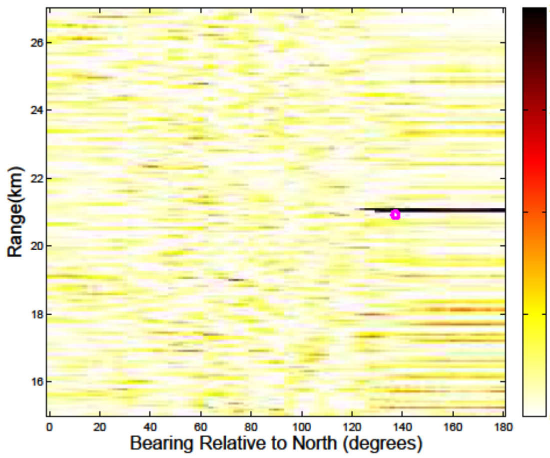

Figure 3. TkBD measurement image for a target SNR return value of 24 dB.

In this paper, the variation in the beam pattern with receive direction will be modelled by assuming a bearing-dependent psf. Recall that a sensor's psf can be used to describe the appearance of the target in the sensor image. For the sonar problem, the sensor outputs a measurement in range r and bearing θ space such that

$$\begin{aligned} h(x_t^m) &= h(r,\theta) \\ &= h_r(r)h_\theta(\theta) \end{aligned} \quad (6)$$

where $h_r(r)$ and $h_\theta(\theta)$ are defined as the psfs for range and bearing space respectively, and assumed to be independent of each other. Both psfs can be approximated using Gaussians. For an active towed array sonar system, the psf function $h_r(r)$ can be assumed to be consistent across all bearings, however the psf in the bearing space $h_\theta(\theta$ will be dependent on the receive direction due to the variation of beam patterns with bearing.

In this paper, a bearing-dependent psf using a Gaussian approximation will be assumed. For a given receive direction, the Gaussian psf can be calculated by setting the half-height width of the Gaussian pdf to be equal to the half-height beamwidth of the current receive direction. The corresponding variance for the Gaussian psf can then be easily calculated. Figure 2 shows the resulting psfs corresponding to broadside, near aft and aft receive directions. Note that for the near aft direction, a sum of two Gaussians, one for each of the left-right beams is used to model the beam pattern. In addition, the images used have been truncated such that when the left-right beams are well separated, only the correct beam is used.

Tracks are initiated based on a peak detection thresholding process at a given SNR level and the well known M/N logic [5] is imposed to upgrade tracks from tentative to confirmed status. The transition matrix for the target dynamics is given by a constant velocity model which varies depending on the time between consecutive transmissions.

3 COMPARATIVE STUDY USING SONAR DATA

In this section, the performance of the H-PMHT is compared with that of IPDA using archived sonar data. Both the H-PMHT and IPDA performance is quantified for a set of SNR

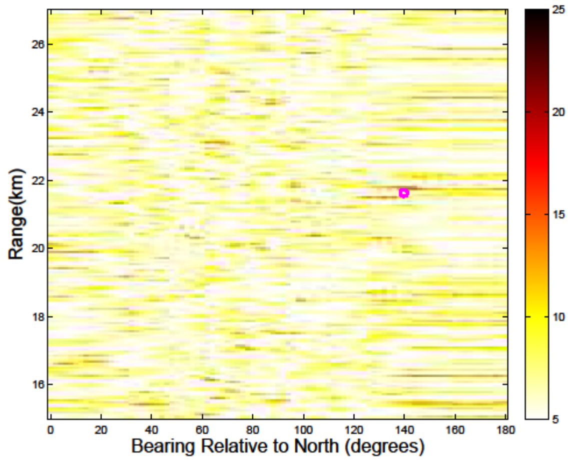

Figure 4. TkBD measurement image for a target SNR return value of 13 dB.

thresholds. For the IPDA, the SNR threshold will determine the number of detections that are passed to the tracker. In the case of the H-PMHT, the SNR threshold will be used in a peak detector for track initiation.

The data used in this paper originates from a series of sonar trials conducted by the Defence Science and Technology Organisation (DSTO) from May to August 2003 using a containerised active towed array demonstration sonar system called CASSTASS. The sonar trials featured a line array towed behind a moving surface ship at two different locations in the Western Australian eXercise Area (WAXA). The two datasets feature characteristics that are unique to the sonar detection and tracking problem in a shallow environment with water depths from 150-250m and for an intermediate ocean environment with depths from 800-1400m. Transmissions were set to detect possible targets with a maximum range of 60 km with the majority of transmissions being in the aft direction. The datasets feature a fluctuating target and persistent clutter detections that are the result of reflections from bathymetric features along the continental shelf. Both datasets consist of approximately 20 transmissions with the time duration between transmissions being approximately 90 seconds.

The target is an echo-repeater (ER) made to simulate the returns from a simple point-like target in the ocean environment. During the trial, the ER platform was observed to have an average travelling speed of 0-4 knots with varying target strengths of 9, 19 and 29 dB. The SNR for the two datasets is approximately 29 dB, which is relatively high, and it is expected that both the IPDA and TkBD will be able to form tracks on the ER. However, variations in performance are also expected as the SNR for the ER did fluctuate with time, with SNR levels as low as 13 dB being observed.

Both the IPDA and H-PMHT tracked using global Cartesian coordinates with an origin fixed at the last recorded ownship position in each trial.

The IPDA algorithm used point-measurement detections which were generated from an automated detection scheme featuring clustering. To gain an idea of the sensitivity of conventional tracking to thresholding, the IPDA algorithm was implemented using three detection thresholding levels at 11, 13 and 15 dB. The expected number of clutter measurements was modelled using a Poisson distribution with parameter $\lambda = N_k / V_k$ where V_k denotes the area of surveillance and N_k was estimated using the average number of point detections which varied with the SNR thresholding level. The IPDA algorithm initiated tentative tracks using two-point differencing. When the probability of existence for a track rose above 0.5, tentative tacks were upgraded to confirmed status. Tracks were terminated when the probability existence fell below a threshold of 0.3.

For the TkBD case, the magnitude of the sensor returns in terms of power was collected in bearing and range cells consisting of 181 beam bins at 2 degree intervals from 0 to 360 degrees and approximately a few hundred range bins, with range intervals of approximately 60m. Figure 3 and Figure 4 show examples of TkBD measurement images across range and bearing space for target SNR values of 24 dB and 13 dB respectively. The true target position is indicated by the red circle. Clearly, when the target SNR is high, the target return is easily distinguishable from the background clutter. However when the target SNR drops, it is more difficult to detect the target from the background clutter. It is also important to realise that as the resolution of the bearing bins is finer than the beamwidth of the sensor, it is expected that the target returns will be spread over multiple bearing bins.

For track initiation, the H-PMHT algorithm applied a peak detection threshold to the sensor image according to a certain SNR threshold and then used two-point differencing to initiate tentative tracks. In the case when the current position estimates of several tentative tracks were within 250 m in position, the highest SNR track was retained and all other tracks were discarded. A measure for the track quality was derived using target SNR calculated from the mixing proportions estimates. Tracks were confirmed and terminated using a 3/5 logic rule requiring this track quality measure to be above a threshold for 3 out of 5 returns.

Multi-target implementations were used to compare TkBD to conventional tracking. In both algorithms, it was assumed that background noise was uniform and measurement noise followed a Gaussian distribution. Due to the high volume of data, a time-recursive H-PMHT filter was used in the analysis rather than processing the sequence as a batch. For the H-PMHT, a maximum of ten EM iterations were performed at each time scan.

4 RESULTS

The tracking outputs of the H-PMHT and the IPDA algorithm for the two sets of trials data is now presented. The tracks from the two algorithms were compared with truth data provided by GPS logs.

Figures 5, 6 and 7 show the tracking results in both datasets at SNR thresholding levels of 11, 13 and 15 dB respectively. In these figures, the IPDA algorithm used detections thresholded at the respective SNR level while the H-PMHT algorithm used the same SNR level to initiate tracks using a peak detection thresholding scheme. It is important to realise that although the H-PMHT initiated tracks on thresholded peaks, state estimates were updated using the entire raw image data. For the intermediate dataset, manual detections which have been visually picked out with the naked eye by a sonar operator are shown in green. Observe that there is a bias between the ground truth given by the GPS logs shown in black and the manual detections shown in green in the intermediate dataset. The differences are due to effects of current on the towed array, which have not been taken into account. Tables 1 and 2 show the average number of measurements at each time frame received by the IPDA at each thresholding level for the shallow and intermediate dataset respectively. The tables also collate the number of false and divergent tracks formed by the IPDA in each dataset. A track is considered divergent if it successfully initiated a track on the target, but was seduced away due to spurious measurements.

Table 1. Number of false and divergent IPDA tracks at different SNR thresholding levels for the shallow dataset.

Threshold Level (dB)	*Average number of detections per frame*	*Number of divergent tracks*	*Number of false tracks*
11	115.4	0	8
13	4.4	0	1
15	0.8	0	0

Table 2. Number of false and divergent IPDA tracks at different SNR thresholding levels for the intermediate dataset.

Threshold Level (dB)	*Average number of detections per frame*	*Number of divergent tracks*	*Number of false tracks*
11	118.5	1	12
13	5.1	1	0
15	0.7	0	0

The H-PMHT algorithm was able to track the ER target with zero false and divergent tracks formed in all scenarios.

It can be seen that at SNR thresholding level of 11 dB, on average the IPDA processed over 100 detections at each time frame. The IPDA algorithm was unable to form sensible tracks because of the large number of detections, which served to increase the size of the covariance and hence the validation gate at each iteration. As a result, the IPDA was unable to form a track on the ER target in the shallow dataset and although it did form one target track in the intermediate dataset, the track rapidly diverged after several time scans. Figure 5(a) only displays tracks from the H-PMHT algorithm as the IPDA algorithm produced 8 spurious false tracks which were not related to the ER target. For similar reasons, Figure 5(b) only shows the track from the H-PMHT algorithm and the single divergent track from IPDA, and none of the false IPDA tracks.

At a SNR thresholding level of 13 dB, the IPDA algorithm formed a track on the ER target in the shallow dataset but was only able to form a divergent track in the intermediate dataset. When the SNR threshold level was raised to 15 dB, the IPDA was able to form a track on the ER target in both datasets with zero false and divergent tracks. The likelihood of initiating and maintaining a track on a low SNR target at higher thresholds would be negligible as no detections would be obtained. Even if the SNR threshold was lowered, it would still be difficult to form a track on a low SNR target due to the sheer number of detections as observed at the low SNR thresholds here.

5 SUMMARY

This paper presents a comparison of the tracking performance of a TkBD algorithm using H-PMHT with the outputs of a conventional IPDA tracker with multi-target logic using point detections. The algorithms were tested on archived data from

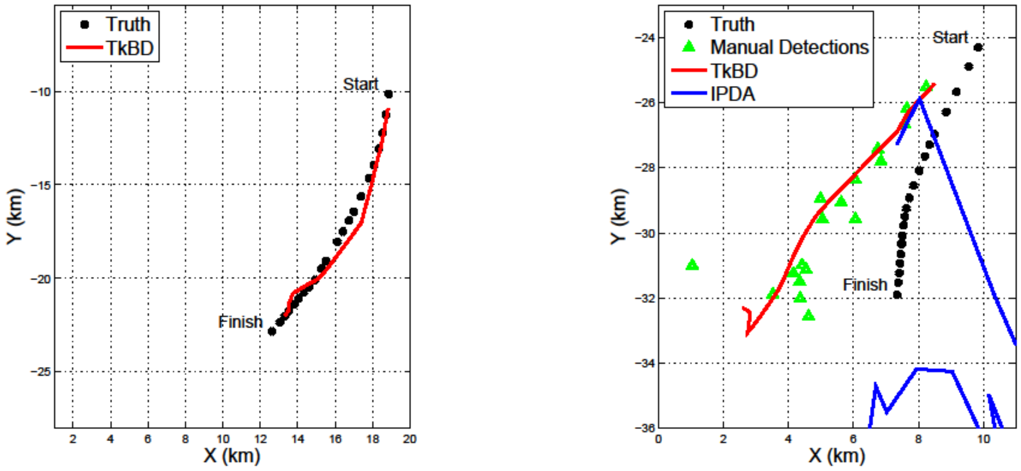

(a) Shallow (b) Intermediate

Figure 5. Tracking Results for H-PMHT vs. IPDA using a SNR thresholding level at 11 dB.

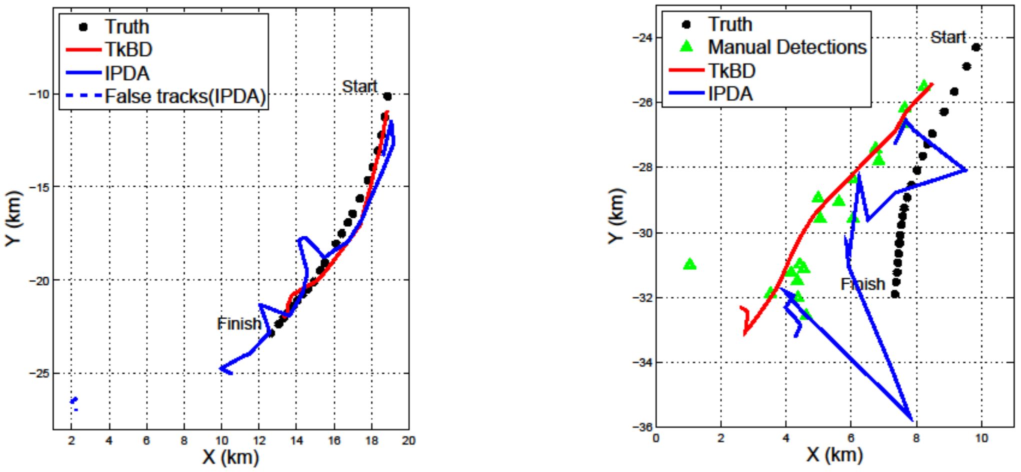

(a) Shallow (b) Intermediate

Figure 6. Tracking Results for H-PMHT vs. IPDA using a SNR thresholding level at 13 dB.

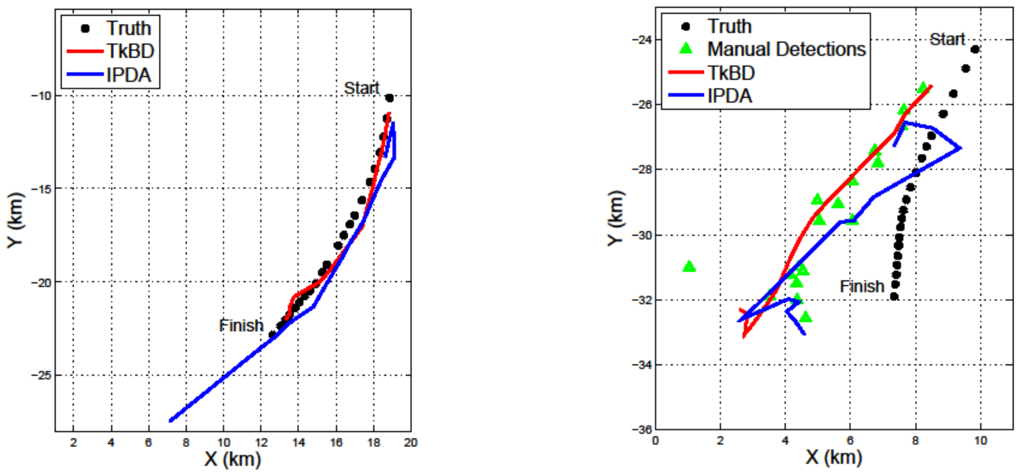

(a) Shallow (b) Intermediate

Figure 7. Tracking Results for H-PMHT vs. IPDA using a SNR thresholding level at 15 dB.

an active towed array sonar system collected off the coast of Western Australia. This paper has demonstrated that for multi- target tracking in shallow and intermediate scenarios, TkBD can provide significant performance advantages over the IPDA implementation. The test data contained a target which gave relatively high SNR returns on most scans so the IPDA was able to form a reliable track when a high SNR threshold was applied during the point-measurement extraction stage. In contrast the H-PMHT processed the intensity map data directly without an explicit SNR threshold (except artificially imposed for track initiation) and was able to detect the target without forming any false tracks.

In future work, we will consider testing the algorithms on datasets featuring low SNR targets. An extension of the TkBD algorithm to include a multi-path target model to aid in the detection of targets in deep water scenarios will also be considered.

REFERENCES

Aughenbaugh, JM & La Cour, BR 2009, 'Use of Prior Information in Active Sonar Tracking', *Proceedings of the 12th International Conference on Information Fusion.*

Babu, VS, Shankar, MR, Majumdar, S & Rao, SK 2001, 'IMM Unscented Kalman Filter Based Tracking of Maneuvering Targets using Active Sonar Measurements', *2011 International Conference on Communications and Signal Processing (ICCSP)*, pp.126-130.

Bar-Shalom, Y, Li, XR & Kirubarajan, TE 2001, *Estimation with Applications to Tracking and Navigation Theory*, Algorithm and Software, Wiley, New York.

Blanding, WR, Willett, PK , Bar-Shalom, Y & Coraluppi, S 2009, 'Multisensor Track Management for Targets with Fluctuating SNR', *IEEE Transactions on Aerospace and Electronic System*, vol. 45, no.4, pp.1275-1292.

Castella, FR 1976, 'Sliding Window Detection Probabilities', *IEEE Transactions on Aerospace and Electronic Systems*, vol. 12, pp. 815-819.

Colegrove, SB, Davis, AW & Ayliffe, JK 1986, 'Track Initiation and Nearest Neighbours Incorporated into Probabilistic Data Association', *Journal of Electrical and Electronics Engineering, Australia*, vol. 6, no, 3, pp. 191-198.

Davey, SJ 2011, 'Histogram PMHT with Particles', *Proceedings of the 14th International Conference on Information Fusion*, Chicago, USA.

Davey, SJ, Rutten, MG & Cheung, B 2008, 'A comparison of detection performance for several track-before-detect Algorithms', EURASIP *Journal on Advances in Signal Processing*.

El-Jaber, M, Osman, A, Mellema,G & Nourledin, A 2009, 'Target Tracking in Multi-Static Active Sonar Systems using Dynamic Programming and Hough Transform', *Proceedings of the 14th International Conference on Information Fusion*, Chicago, USA.

Hempel, CG & Doran, SL 2004, 'A PMHT algorithm for active sonar', *Proceedings of SPIE 5430, Acquisition, Tracking, and Pointing XVIII*, 132.

Luginbuhl, TE & Willett, P 2004, 'Estimating the Parameters of General Frequency Modulated Signals', *IEEE Transactions on Signal Processing*, vol. 52, no. 1, pp. 117-131.

Musicki, D, Evans, R & Stankovic, S 1994, 'Integrated Proabilistic Data Association', *IEEE Transactions on Automatic Control*, vol. 39, no. 6, pp.1237-1241.

Musicki, D & Evans, R 2004, 'Joint Integrated Probabilistic Data Association: JIPDA', *IEEE Transactions on Aerospace and Electronic Systems*, vol. 40, no. 3, pp.1093-1099.

Musicki, D, Wang, X, Ellem, R & Fletcher, F 2007, 'Efficient Active Sonar Multitarget Tracking', *OCEANS 2006*, pp.1-8.

Orlando, D, Ehlers, F & Ricci, G 2010, 'Track-before-detect Algorithms for Bistatic Sonars', *International Workshop on Cognitive Information Processing*, pp. 180-185.

Streit, RL 2000, 'Tracking on intensity-modulated data streams', Technical report 11221, NUWC, RI, USA.

Streit, RL & Luginbuhl, TE 1995, 'Probabilistic multi-hypothesis tracking', Technical report 10428, NUWC, RI, USA.

Tonissen, SM & Evans, RJ 1996, 'Performance of dynamic programming techniques for Track-Before-Detect', *IEEE Transactions on Aerospace and Electronic Systems*, vol 32, no. 4, pp. 1440-1451.