



Australian Government
Department of Defence
Defence Science and
Technology Organisation

Distributed versus Centralised Tracking in Networked Anti-Submarine Warfare

J. M. Thredgold and M. P. Fewell

Maritime Operations Division
Defence Science and Technology Organisation

DSTO-TR-2373

ABSTRACT

This report describes a study of active sonar tracking, which explores the effect of networking sonars on tracking performance. We compare the tracking performance when sonars share detections (centralised tracking) with the performance when sonars share tracks (distributed tracking). Provided that the sonar layout and detection probabilities are such that multiple sonars have a reasonable probability (~30%) of obtaining detections from a target, we show that centralised tracking decreases the time to confirm a track on a target and improves the continuity of the target track. These improvements in target tracking occur at the expense of an increase in false track rate.

RELEASE LIMITATION

Approved for public release

Published by

*Maritime Operations Division
DSTO Defence Science and Technology Organisation
PO Box 1500
Edinburgh South Australia 5111 Australia*

Telephone: (08) 8259 5555

Fax: (08) 8259 6567

© Commonwealth of Australia 2010

AR-014-687

January 2010

APPROVED FOR PUBLIC RELEASE

Distributed versus Centralised Tracking in Networked Anti-Submarine Warfare

Executive Summary

A previous report (DSTO-TR-2086) suggests that appropriate networking of sonars may be able to make significant inroads on the problem of submarine detection. Of the many assumptions in this simple analytical approach, the avoidance of the issue of false alarms is perhaps the most serious. The present report seeks to address this by comparing distributed and centralised tracking in a scenario that explicitly includes false detections at a rate that leads to false tracks.

We compare tracking performance when sonars share detections (centralised tracking) with the performance when sonars share tracks (distributed tracking). The model builds on the earlier analytical work, which found that detection probabilities as low as 30% could be useful for track initiation, if there was a network of sonars such that these detections could be shared with other sonars with similar probabilities of detection.

The simulation model used for the analysis in this report was developed in order to see if this networking advantage carries through to other stages of the tracking process. The simulation model is documented in a separate report (DSTO-TR-2372). Using this model, we show that centralised tracking decreases the time to confirm a track on a target and can also improve the continuity of the target tracks.

These improvements in target tracking occur at the expense of an increase in false track rate. The majority of false tracks formed by the centralised tracker have similar characteristics to tracks formed by distributed tracking, and so are not more difficult to identify as false tracks. Classifying tracks by length alone, centralised tracking results in a small increase in the number of false tracks wrongly classified as potential targets (from 7 to 11%) but this is offset by the decrease (from 33 to 26%) in the number of target tracks which fail to be identified as such. These results indicate the potential for false detections and false tracks to cause difficulties when sonar systems are networked. It points to a need to improve detection classification techniques if networking is to achieve its full potential in anti-submarine warfare.

Authors



J.M. Thredgold
Maritime Operations Division

Jane Thredgold commenced working in the Maritime Operations Division at DSTO in early 2007, after completing a PhD in Mathematics at the University of South Australia. She is currently performing operations analysis work relevant to anti-submarine warfare.



M.P. Fewell
Maritime Operations Division

Matthew Fewell joined DSTO in 2001, coming from an academic physics background. He has worked and published in experimental nuclear structure physics, gaseous electronics, atom-photon interactions including coherent effects, laser physics, plasma processing of materials, the theory of network-centric warfare and its modelling at the operational level (including cognitive issues), human-in-the-loop experimentation, and weapon-target allocation in ship air defence. He is at present working on issues in anti-submarine warfare from surface-ship and maritime-patrol perspectives.

Contents

VARIABLES

1. INTRODUCTION.....	1
2. STUDY METHOD	2
2.1 Summary of the Simulation Model.....	2
2.2 Data Fusion	3
2.3 Scenario.....	4
2.4 Tracking Assessment Method	5
2.4.1 Track Classification Schemes	5
2.4.2 Metrics.....	6
2.4.3 Track Coalescence	7
3. COMPARISON OF CENTRALISED AND DISTRIBUTED TRACKING	8
3.1 Simulation Results	8
3.1.1 Track Establishment Delay	8
3.1.2 Number of False Tracks.....	9
3.1.3 Length of False Tracks	10
3.1.4 Length of Target Tracks.....	11
3.1.5 Number of Tracks Occurring at Each Time Step during a Simulation Run	12
3.1.6 System Confusion Matrix.....	13
3.2 False Track Behaviour.....	14
4. SUMMARY AND CONCLUSIONS.....	18
5. REFERENCES	20

Variables

L	track classification length threshold
P_d	probability of detection per ensonification
P_{fa}	probability of a false detection
P_{fti}	probability of false track initiation per 5 consecutive ensonifications

1. Introduction

The prevailing paradigm governing concepts of command and control – network-centric warfare [1,2] – makes it natural to look to networking for solutions to the antisubmarine-warfare (ASW) problem [3]. In a previous report [4], we described some analytical studies comparing the ASW effectiveness of forming tracks on detections from a group of sonars (centralised tracking) with the case where each individual sonar forms tracks using only its own detections (distributed tracking). This report describes a further study of the problem, motivated by limitations arising from some of the assumptions made in the earlier analytical work.

The previous work focussed on track initiation. Results from this analytical study indicated that networking sonars increases the area in which we can reasonably expect to start a track on a target, with the greatest networking benefits seen when the sonar probability of detection (P_d) versus range curves have long tails of low, but not too low, probability ($P_d \sim 0.3$). However the approach taken in the previous work did not consider measurement errors or the effect of false detections, nor the resulting difficulties in associating a detection with other detections to form tracks. In the analytic work, we assumed that, if the target is detected three times in five pings, then we always start a track. In reality this may not always occur, as measurement errors mean that the detections may not be close enough together or a false detection may interfere with the track formation process by confusing the picture. In short, the analytical study doesn't give sufficient consideration to the issue of *data association*.

Also the significant issue of the increased false track rate arising from centralised tracking was ignored. The previous work included an analytical study of false-detection rate in the centralised case. However this study was inconclusive and did not address at all the issue of the false detections leading to false tracks.

Because of these shortcomings, a simulation model was designed and implemented to address these issues. The model is described in detail in a companion paper [5]. It allows us to investigate the impact of measurement errors and false detections on ASW effectiveness. The false (i.e. non-target) detections in the model are assumed to be 'noise' detections which are not associated with any nearby real object. The presence of recurrent clutter type detections which result from objects, such as bottom features or fish schools, is a separate issue that is not addressed in this report. The simulation model includes data association, track maintenance and track termination, not just the track initiation step.

Section 2 of this report describes the analytical method adopted for the study and Section 3 gives results, including some analysis of how the number of sensors in the network and the probability of a false detection affect the false track rate.

2. Study Method

The aim of the study is to examine the extent to which the conclusions of the earlier analytical study [4] stand up to the inclusion of false detections and measurement errors. This was approached by constructing a simulation model explicitly including these features. The model is summarised in Section 2.1; details appear elsewhere [5]. Compared with an analytical study, a simulation-based study has the disadvantage that it must work with a physical scenario, leading to the potential criticism that the results are scenario-specific. This criticism is difficult to counter, short of running many different scenarios. Our approach is to choose a scenario that represents a generic ASW task, as described in Section 2.3. We choose the simplest (from the modelling point of view) of the four generic ASW scenarios [6], open-ocean transit.

The conclusions of an operations-research study can be strongly influenced by the metrics chosen. The earlier analytical work [4] examined a range of metrics, none of which are well adapted to analysis of scenarios with false alarms. Since the present work is exploratory, we decided that its main thrust would involve a comparison of tracking metrics. The metrics chosen, detailed in Section 2.4.2, were based on metrics used in the surveillance-radar domain [7].

2.1 Summary of the Simulation Model

Given a scenario, such as that described in Section 2.3 below, the elements that define the model can be grouped under the following topics:

- type and operation of the sonars modelled
- model of sonar performance
- manner in which target and false detections are generated
- method of initiating a track
- how to decide whether a detection should be associated with an existing track
- the tracking algorithm
- track termination

These are summarised in the following paragraphs. Full details are given in a companion report [5]. Since our interest is not in improving tracker performance but rather in comparing ASW effectiveness when tracking is performed centrally as opposed to fusing the tracks from individual sonars, we use standard algorithms for tracking and data association.

We model a field of active sonars with multiple monostatic operation. The performance of the modelled sonars is defined by probability of detection versus range curves, with no bearing or depth dependence. For simplicity and ease of comparison with the analytical study, and to keep the modelling unclassified, in the work reported here we use exponential functions as the probability of detection versus range curves.

Whether or not a sonar ping produces a target detection is determined by a draw from a uniform random distribution compared with the detection probability at the relevant range. If the target is detected, then the recorded position is generated using additional random numbers to introduce measurement error. Measurement errors have a zero-mean bivariate normal distribution in range and bearing. The covariance matrix of the errors is set so that the

standard deviations correspond to the Cramér–Rao lower bound [8], approximately 1 metre in range and 2.5° in bearing.

The false detections generated differ for each run of the simulation. The number of false detections occurring at each ping is drawn from a Poisson distribution with the expected value depending on the probability P_{fa} of a false detection per range–bearing cell per ping set by the user. In practice an operator who is concerned by a high false alarm rate would raise the detection threshold (hence decreasing the probability P_{fa} of a false detection) to reduce the number of tracks they need to consider. One option in our simulation model would be to make the false track rates with the two different tracking options equal by adjusting the probability of obtaining a false detection. However we have not done this, as we preferred to keep the detection lists passed to the trackers identical for the two tracking options. We used $P_{fa} = 10^{-5}$ with the scenario described in Section 2.3 below to give the results presented in Section 3.

The track-initiation rule specifies three detections in five consecutive pings of the active-sonar field. False detections and measurement errors complicate the application of this rule: we require a way of deciding whether a detection ought to be associated with an existing track, or whether it is a candidate for forming a new track. For this, we use Mahalanobis distance [9] from the detection to the current track position. We require this distance to be small enough that there is a 95% probability of the detection being a target detection. Where more than one detection fulfils this criterion, we pick the one with the smallest distance. This is done with each existing track, which means that a detection might be associated with more than one track.

After all tracks have been processed, any detection that is not associated with an existing track is considered as a potential source of a new track. Given an isolated detection, the algorithm will search for a second detection within the next three time steps using an expanding, almost square, gate centred on the initial detection. The size of the gate is primarily determined by the length of time elapsed since the initial detection and the maximum relative velocity with which it is estimated potential targets could be travelling.

The tracking algorithm used is relatively simple as our interest is not in improving tracker performance but in comparing the tracks produced when tracking is performed centrally with those obtained by fusing the tracks from individual sonars. We use a standard Kalman filter [10]. The 3-in-5 rule is used not only for track initiation, but also for track maintenance and termination.

2.2 Data Fusion

Within each simulation run the same detections (both target and false) are used by the two tracking options. The only difference is in how these detections are processed to form tracks, as described below.

For centralised tracking, the detection lists from each sensor are combined and passed to the tracker as a single set. The combination rule is the equivalent of a simple logical ‘or’, except for the case where two detections occur at the same time very close to one another. Detections

from two sonars are fused into a single detection if the Mahalanobis distance between them is consistent with zero at the 95% confidence level.

For distributed tracking, the tracking algorithm is applied separately to the detections recorded by each sensor. The tracks are then fused into a single list. If two sensors are tracking the target at the same time, and the track position estimates are close enough together, the tracks will be fused into a single track. For example if sonar 1 has a track on the target from ping 10 to ping 14 and sonar 2 has a track on the target from ping 12 to ping 17, for analysis purposes this is regarded as a single track on the target from ping 10 to ping 17, as Figure 1 illustrates.




Figure 1: Illustration of track fusion for distributed tracking

2.3 Scenario

The scenario used in this study is shown in Figure 2. This scenario comprises a network of three sonars. There is a single submarine target present. The positions of the sonars and the submarine at the beginning of the simulation are shown in the figure, together with their relative motion, which is constant throughout the simulation. If the sonars were hull mounted sonars attached to ships, then this scenario could be considered as representative of a task




Figure 2: The positions of the three sensors and the submarine at the beginning of the simulation. The relative velocity is constant throughout the simulation.

group in the process of running over an adversary submarine.

Each sonar pings once every 60 seconds, which is also the time step of the simulation. The simulation runs for 40 minutes (41 pings).

For the example shown in Figure 2, we give the rear ship a more powerful sonar system than the other two. This is done so that we have a reasonable probability of obtaining detections from more than the closest ship to the submarine. We use exponential P_d curves with a half distance⁽¹⁾ of 20 km for the rear ship, compared with 15 km for the other two vessels.

2.4 Tracking Assessment Method

In order to analyse the results of the simulation, tracks need to be classified as either false tracks or target tracks. As we are using simulated data, for each detection we know whether it is a target or false detection. This knowledge can be used to classify the tracks as target or false. Where a track is constructed entirely from one type of detection, it is obvious how to classify the track. However some tracks contain a combination of false and target detections, and tracks from the centralised tracker may also contain detections which were created by fusing detections of different types.

2.4.1 Track Classification Schemes

We classify the tracks by giving each track a classification score, calculated at the end of a run of the simulation as the number of target detections minus the number of false detections in that track,⁽²⁾ normalised by dividing by the total number of detections. The value of the track score will be between -1 and 1, inclusive. If the track score is positive, then the track is classified as a target track, otherwise the track is classified as false. Note that a track containing an equal number of target and non-target detections (with a classification score of zero) is classified as a false track.

To check that this track classification scheme does not unduly affect the metrics discussed in later sections, we produced a histogram of the track classification scores from one thousand runs of the scenario described in Section 2.3 with $P_{fa} = 10^{-5}$. As the top plot in Figure 3 shows, the vast majority of tracks are either composed entirely of false measurements (79.3%) or of target measurements (13.4%). Only just over 7% of all tracks are mixed, and of these many are composed of one type of measurement with only a single instance of the other type. In addition to the extreme peaks at -1 and 1, the histogram shows two smaller peaks around $\pm 1/3$, these peaks reflect the number of tracks which contain just the three measurements required for a track to be confirmed. The peak at $+1/3$ corresponds to tracks which contain one false and two target measurements and the peak at $-1/3$ corresponds to tracks with a single target and two false measurements.

(1) Half distance' is the range at which $P_d = 0.5$.

(2) With centralised tracking, a detection passed to the tracker may be formed from the fusion of true and false detections, as described in Section 2.2. These make a contribution to the track classification score that reflects their composition. That is, a detection formed from the fusion of one true and one false detection contributes zero to the track classification score; a detection formed by fusing one true and two false detections contributes $-1/3$, and so on.




Figure 3: Histograms of track classification scores. The lower plot excludes those tracks with track classification scores of ± 1 in order to better show the detail of the distribution between these two dominant extremes. Note the different scales on the y axes.

2.4.2 Metrics

The metrics used to compare the performance of centralised and distributed tracking were derived from metrics chosen for studying tracker performance in the surveillance-radar domain [7]. They are:

- *Track establishment delay*: Defined as the number of pings between the first target detection and confirmation of a track (obtaining a third detection within a five ping window), this metric can be calculated for both the centralised tracker and for each individual sensor.
- *Number of false tracks*: A simple count of the total number of false tracks occurring in the simulation run for each individual sensor and for the centralised tracker, this metric allows comparison of the false track rates.
- *Length of false tracks*: In order to see if the false tracks formed by the centralised tracker exhibited different behaviour to those formed by the individual sensors, the length of each false track was recorded. We define track length to be the total number of pings

between the first and last detection of the track, counting all intermediate pings regardless of whether they produced a detection.

- *Length of target tracks:* The lengths of all target tracks were also recorded. Target track length provides an indication of the track continuity.
- *Number of tracks occurring at each time step during a simulation run:* For an operator monitoring a sonar display, the total number of false tracks occurring over some period of time may be less important than how many tracks appear on the display at a single time instant. The total number of tracks (target and false) in existence at each time during the simulation was recorded. Average, maximum and minimum numbers of tracks will be presented, together with the 5% and 95% percentiles of the track-number distribution.
- *System confusion matrix:* This is an analysis tool used in the classification literature which effectively measures the ‘confusion’ in the tracking picture, defined as how often target tracks are missed compared with how often false tracks are mistaken for real. That is, in the current context, the confusion matrix can be thought of as

% false tracks identified as such	% false tracks where action taken
% missed target tracks	% declared target tracks

Where some action is taken to further investigate a false track (top right value), this is considered a system false alarm. A confusion matrix must be defined in terms of a chosen classification rule. The results in the next section suggest that a useful rule may be based on track length, with false tracks in general being shorter than target tracks.

Hence, we examine confusion matrices of the form:

$\frac{U\#false\ tracks < L}{\#false\ tracks}$	$\frac{\#false\ tracks \geq L}{\#false\ tracks}$
$\frac{\#true\ tracks < L}{\#true\ tracks}$	$\frac{\#true\ tracks \geq L}{\#true\ tracks}$

for a specified value of the threshold track length L . In the ideal situation where all tracks are correctly identified, the confusion matrix would be the identity matrix.

2.4.3 Track Coalescence

During the analysis of the simulation results it was observed that in some runs of the scenario, multiple target tracks existed at the same time. This occurred when a track initiated using a false detection close to the target position was updated at subsequent times using target detections, as illustrated in Figure 4. In a small number of runs, there were as many as four simultaneous target tracks. This phenomenon is known as track coalescence [11] and is not uncommon in automated tracking systems. It arises because a detection may be associated with more than one track. In the situation where a sonar display is being monitored by an operator, the additional tracks would normally be recognised as such, and ignored or fused manually.

As we needed to automate the analysis of the tracks to allow Monte Carlo simulation, we needed to eliminate duplicate target tracks formed as described in the preceding paragraph as they could potentially bias a number of the metrics, particularly those which rely on track length and track count. Specialised tracking algorithms have been developed to deal with the track coalescence problem but, rather than complicate the tracking algorithm, we chose a simpler solution. Instead of attempting to prevent the formation of these tracks, during the analysis we simply ignore duplicate tracks which are consistently updated with the same measurements.




Figure 4: *Illustration of track coalescence—formation of duplicate target tracks initiated by false detections.*

3. Comparison of Centralised and Distributed Tracking

The main differences between centralised and distributed tracking are expected to relate to the timeliness, continuity and accuracy of the tracks and the numbers of false tracks which form. The earlier analytical work [4] indicated that centralised tracking has a higher probability of tracks starting sooner (when the target is at a greater distance from the sonars) than distributed tracking. It was also thought that centralised tracking may provide more continuous tracking of the target, however it is expected to have a higher false track rate.

3.1 Simulation Results

The simulation model was run 1000 times for the scenario discussed in Section 2.3. The following Subsections detail the comparison between centralised and distributed tracking as quantified by each metric.

3.1.1 Track Establishment Delay

The track establishment delay is not affected by the track coalescence problem as only the first target track formed is used in the calculation of this metric. As the histogram in Figure 5 shows, tracks are formed sooner with centralised tracking. Values for mean and median track establishment delay for the two trackers are listed in Table 1.

As Figure 5 shows, the distribution of the track establishment delay has a much longer tail for distributed tracking when compared to the distribution of the delay for centralised tracking. As well as reducing the average time to start a track, centralised tracking greatly reduces the number of cases in which a track is started a long time after the first target detection. For example, for the thousand runs from which the mean and median values above were

obtained, there were only 23 cases (or 2.3% of the runs) where the centralised tracker did not confirm a target track within 15 pings of the first target detection. Contrasted with the 188 cases (or 18.8% of runs) where the distributed tracker took 16 or more pings to confirm a track, this highlights the improved performance of the centralised tracker. Recalling that the ping interval for the simulation is 60 seconds, 15 pings equates to 15 minutes.

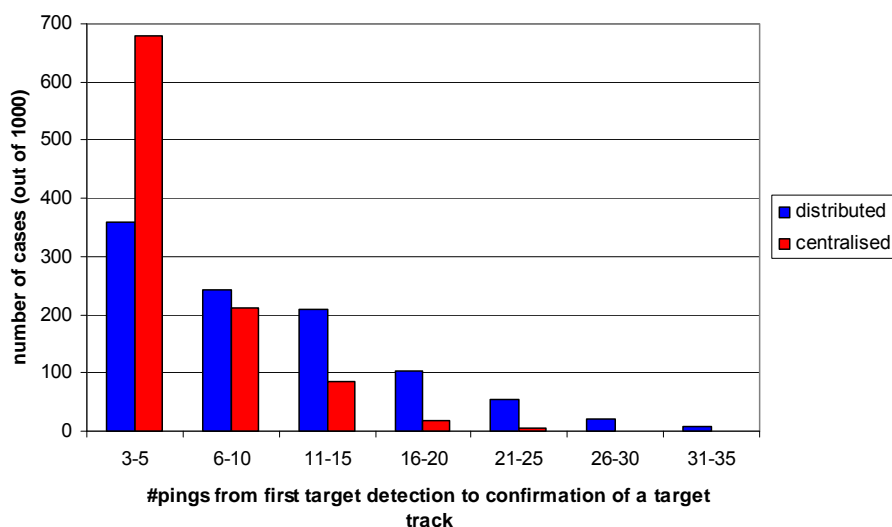


Figure 5: Track establishment delay, the number of pings between the first target detection (by any sensor) and the formation of a track, for both centralised and distributed tracking. All histograms show results obtained from 1000 runs of the scenario.

Table 1: Mean and median track-establishment delays for the two tracking methods

	Track establishment delay (pings)	
	mean	median
centralised tracking	5.2	4
distributed tracking	9.8	8

3.1.2 Number of False Tracks

As expected, the number of false tracks formed by the centralised tracker is much higher than the number of false tracks produced by combining the tracks from the individual sensors.

Figure 6 is a histogram of the total number of false tracks occurring in each of the thousand runs. It appears that the number of false tracks per run is approximately normally distributed. The mean values are listed in Table 2, together with the maximum number of false tracks obtained in the 1000 runs. The highest number of false tracks obtained on a single sensor is 10.

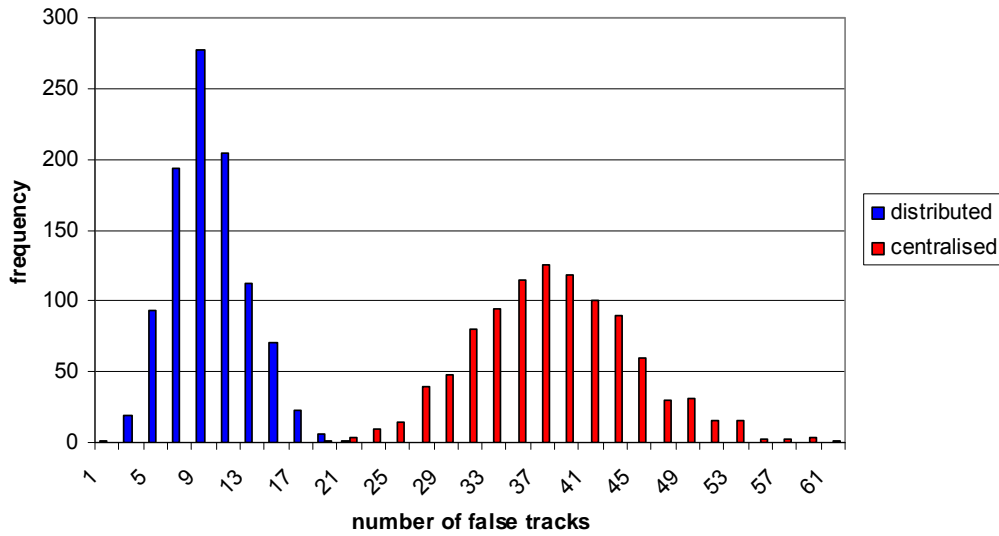


Figure 6: Histogram of the number of false tracks in a scenario run

Table 2: Mean number of false tracks per run of the scenario for the two tracking methods and the maximum observed in 1000 runs

	Number of false tracks	
	mean	maximum
centralised tracking	37.2	61
distributed tracking	9.1	21

3.1.3 Length of False Tracks

While there are many more false tracks formed in the centralised tracking case, there do not appear to be significant differences in the way the length of these tracks is distributed, as Figure 7 shows. The mean length of a false track is 4.83 pings in the distributed case and 4.84

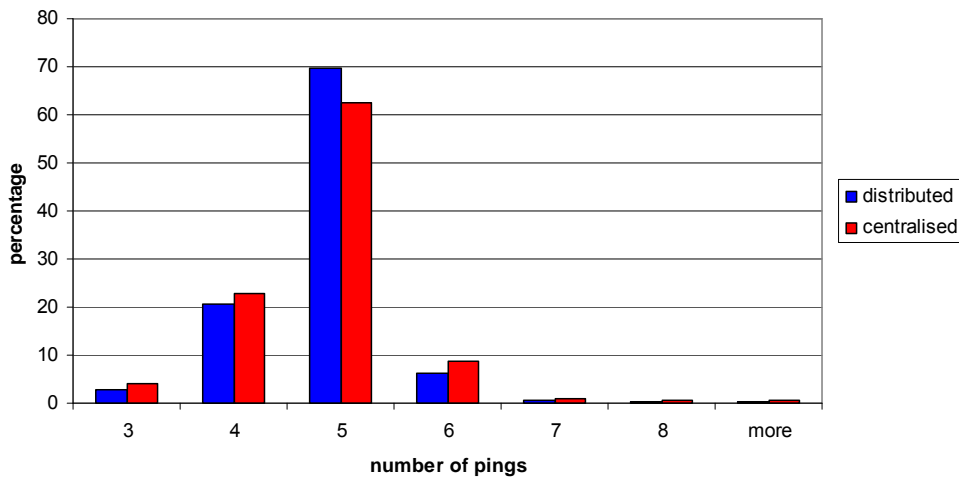


Figure 7: Histogram showing the distribution of the false track lengths for both centralised and distributed tracking

pings in the centralised case. While there is little difference between the means of the distributions and the medians are identical (5 pings), the variances of the two distributions do differ: 0.44 and 0.66 for the distributed and centralised tracking cases respectively.

While 99% of false tracks for both tracking options last for 7 or fewer pings, there are a small number of longer tracks. The longest false track formed in the distributed case – using only detections from a single sensor – is 11 pings long. For the centralised tracker there exists one track (out of the 37238 formed) which is 14 pings long. Closer examination of this track shows that it is a mixture of target and false detections. The longest track consisting entirely of false detections is 12 pings long.

3.1.4 Length of Target Tracks

Unlike the false track lengths, there is a definite, although small, difference in the distributions of the target track lengths for the different tracking options. The two histograms (shown in Figure 8) look similar. Both are skewed with long tails to the right and the most common target track length (the mode) in both distributions is 5 pings. However, the mean and median target track lengths are noticeably different, as Table 3 shows, with centralised tracking giving the longer track lengths on average.




Figure 8: Histogram showing the distribution of the target track lengths for both centralised and distributed tracking

Table 3: Mean and median target-track lengths for the two tracking methods

	Target track length (pings)	
	mean	median
centralised tracking	11.32	9
distributed tracking	9.5	7

Table 4: Comparison of the long track length tail of the two track length histograms (Fig. 8)

	% of target tracks longer than:	
	15 pings	30 pings
centralised tracking	25	3
distributed tracking	16	2

In fact, the mean or median behaviour is not the most important aspect of these distributions, for the ideal situation comprises a single very long target track. That is, the behaviour of the distributions out at long track lengths is the critical feature. In an attempt to quantify this, Table 4 lists the percentages of all target tracks recorded in the 1000 scenario runs that exceed 15 pings and 30 pings in length. Centralised tracking results in a higher number of tracks in both categories.

3.1.5 Number of Tracks Occurring at Each Time Step during a Simulation Run

Figure 9 shows the average, maximum and minimum numbers of all tracks, both target and false, existing at each ping during the scenario for both tracking options, and also the 5th and 95th percentiles of these distributions. Start and end effects are visible in both plots for about 5 pings at the beginning and end.

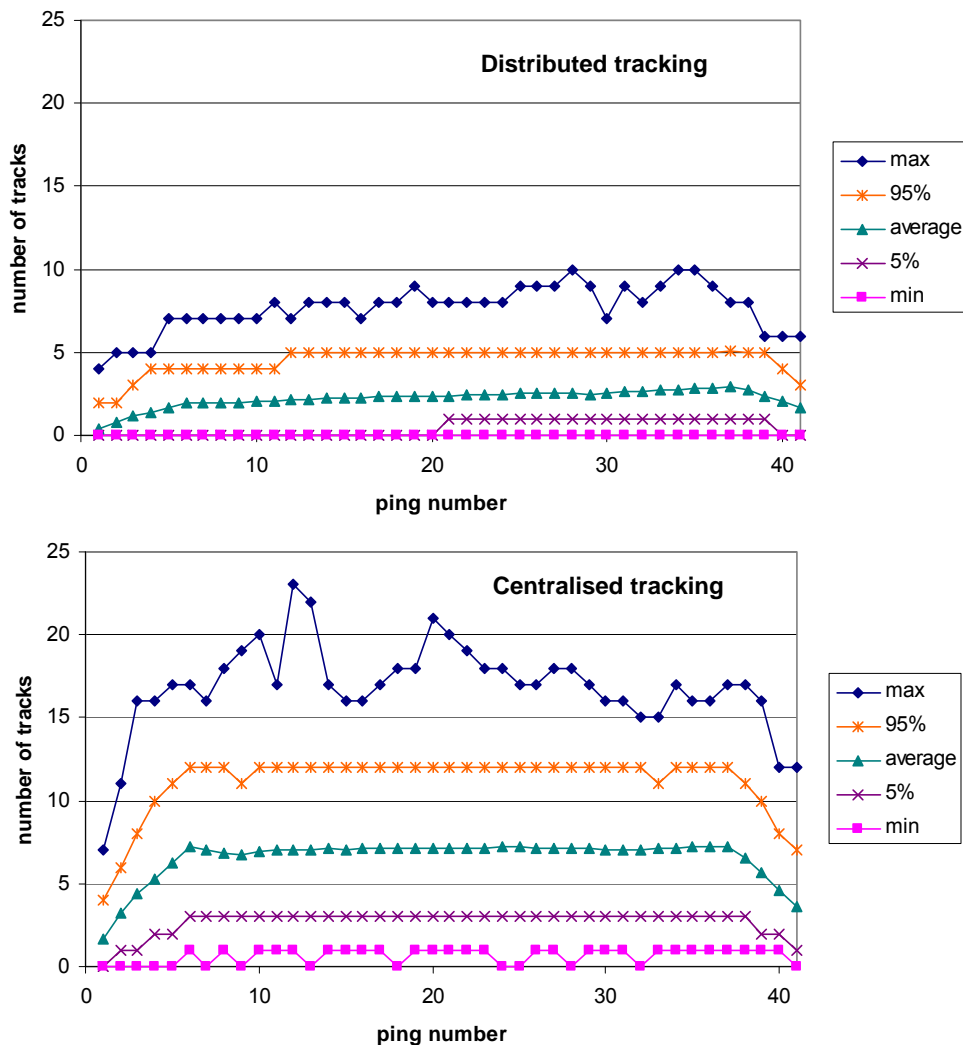


Figure 9: Number of tracks in existence at each ping for both distributed and centralised tracking
 In the plot for the distributed tracking case, there is an upward trend in both the average and maximum number of tracks. This is due to the increased likelihood of target tracks towards

the end of the scenario as the target gets closer to the sensors. This trend is not seen in the lower plot for centralised tracking, simply because the higher number of false tracks diminishes the effect that a single target track will have on the number of tracks.

These plots tell us that, with the chosen false detection rate and distributed tracking, a sonar operator would have to deal with two to three tracks on average after each ping, but could have as many as ten.

3.1.6 System Confusion Matrix

The confusion matrix indicates the proportion of target and false tracks that would be classified correctly by considering track length alone. In reality this would not be the sole classification technique used, but track length may be used by an operator to prioritise the tracks investigated by, for example, listening to the sonar echoes.

For the classification rule examined here, calculation of the confusion matrix requires the selection of a value L , the track length threshold. Based on the results in Figure 7 and Figure 8, we used a threshold length of 5 pings, leading to the results in Figure 10. With distributed tracking, 33% of target tracks and 7% of false tracks would then be wrongly classified. With centralised tracking, we are better able to identify target tracks, with 26% too short to be flagged as such, but have a higher number of longer false tracks with 11% wrongly classified by length alone. The increase in the misclassified false tracks arises because of the small number of longer false tracks which are formed by the centralised tracker.

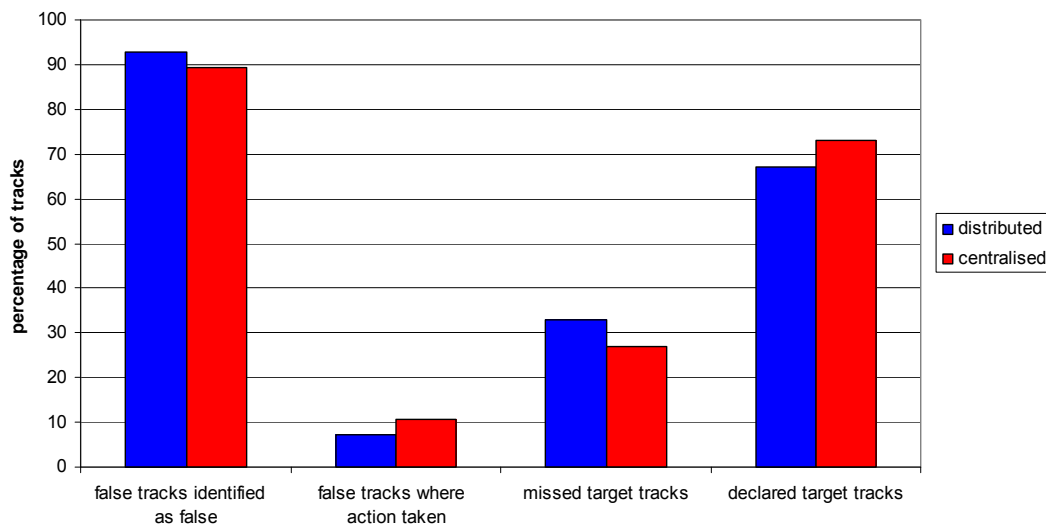


Figure 10: Graphical representation of confusion matrix showing the success of classifying tracks by track length alone, for both distributed and centralised tracking

To illustrate why we chose the threshold length of 5 pings, Figure 11 shows the percentage of tracks which are wrongly classified using this method with all values of L between 3 and 10 inclusive. The optimal value of L minimises both the number of target tracks and the number of false tracks that are wrongly classified. This value can be found from Figure 11 for a particular mode of tracking (centralised or distributed) as the point at which the relevant lines

cross. The crossing point is not exactly at an integer value, but is closest to 5 for both centralised and distributed tracking.




Figure 11: The percentage of false and target tracks which are wrongly classified (middle two pairs of bars in Figure 10), as a function of the track classification length threshold L , for both distributed and centralised tracking. Note that L can only take integer values – the lines connecting the markers at the data points are included to more clearly show the trends.

3.2 False Track Behaviour

The false detection rate can be adjusted by changing the probability P_{fa} of false alarm per range-bearing cell per ping. In addition to this probability, the false track rate also depends on the number and spacing of sensors in the network and on the type of tracking. As an example of the variation of the false track rate with the probability of obtaining a false detection, Figure 12 shows the probability of starting a false track using centralised tracking and distributed tracking with different probabilities of a false detection for a network of three sensors laid out as in Figure 2, but without the target submarine present.

To generate false detections, the area scanned by the sonar is broken up into cells or segments with the range increment and bearing span specified by the resolution that can be achieved by the sonar. Each of the cells is modelled as having a fixed false alarm probability P_{fa} , but the cells do not have equal area. Figure 13 illustrates the way that two cells which have the same probability of false alarm associated with them correspond to different areas on the scan.

Consequently there will be fewer false alarms per unit area a long way from the sensor than close to the sensor. This may mean that the number of false tracks formed depends to some extent on the spacing of the sensors in the network. If sensors are closer together there may be more false detections falling within each others' data-association gates, and hence more false tracks.




Figure 12: Probability P_{fti} of false track initiation for a network of three sensors, given the probability P_{fa} of a false detection. The probability of a false detection is per range-bearing cell per ping; the probability of false track initiation is per ping. The values in this plot are calculated as averages over 200 runs of 100 pings in length, using the scenario layout described in Section 2.3.

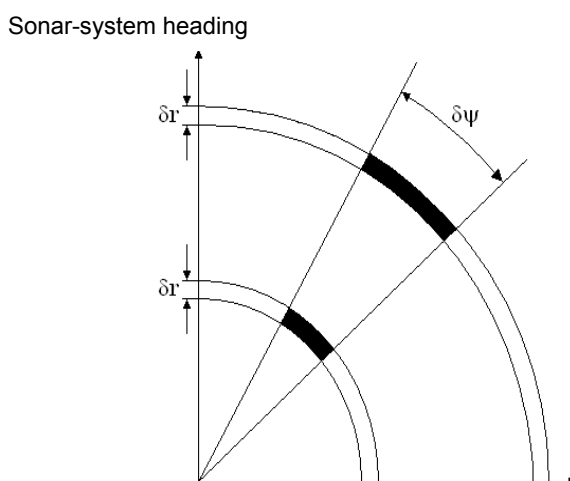


Figure 13: Variation of detection cell size with range

The number of false tracks formed will also vary depending on the number of sensors in the network. We explored this by placing sensors 10 km apart on a triangular grid. As Figure 14 (for $P_{fa} = 10^{-5}$ with 2 through to 7 sensors) shows, there is an increase in the number of false tracks formed when additional sensors are added to the network. This increase is linear for distributed tracking and approximately quadratic for centralised tracking. The values shown are average values calculated from 1000 runs.

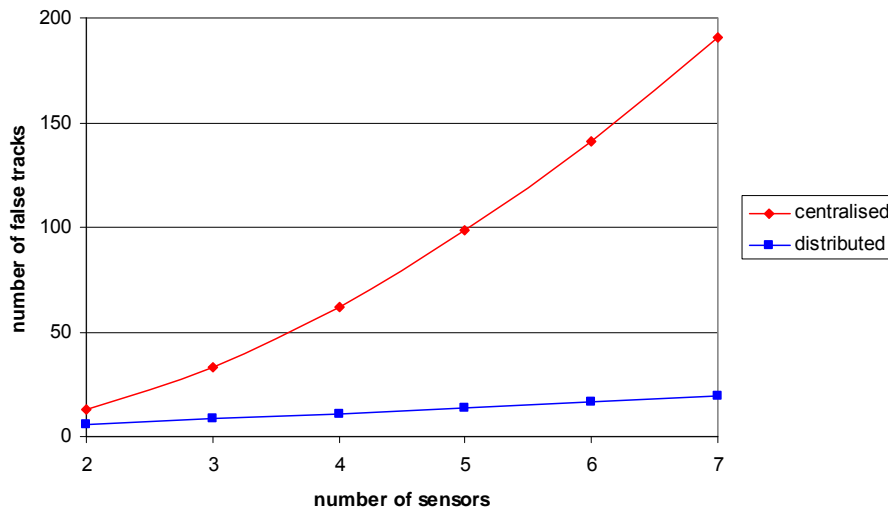


Figure 14: Average number of false tracks formed in one run of the scenario for networks of 2 to 7 sensors, with a probability P_{fa} of false detection per range bearing cell per ping of 10^{-5} . The sensors are equispaced on a 10 km triangular grid.

The length of the false tracks formed by different numbers of sensors was also recorded, and comparing the results for distributed tracking with those for centralised tracking is an interesting exercise. As Figure 15 shows, the false track lengths for distributed tracking do not display any particular trends as the number of sensors increases. However the false tracks formed by the centralised tracker show clear trends, with the percentage of longer false tracks increasing as the number of sensors increases. This of course occurs in conjunction with a decrease in the number of shorter false tracks, shown most obviously by the downward trend in the percentage of false tracks which are 5 pings long.

The increase in the number of longer false tracks as the number of sensors increases indicates that identifying false tracks by using track length as a classification tool will be more difficult in larger networks if centralised tracking is used. This is highlighted by Figure 16 which shows the percentage of false tracks longer than 5 pings, or the number of false tracks which would be wrongly classified as target tracks on the basis of track length alone with $L = 5$, for networks of 2 to 7 sensors, with $P_{fa} = 10^{-5}$. As the number of sensors in the network increases it is probable that the optimal value of L would increase. If track length were to be used as a classification tool, then this would increase the number of missed target tracks.




Figure 15: Distribution of false track lengths for distributed and centralised tracking, showing the variation with the number of sensors in the network.




Figure 16: Percentage of false tracks longer than 5 pings, as a function of the number of sensors in the network

4. Summary and Conclusions

This report documents results of a study of networked tracking in anti-submarine warfare. Previous analytical work [4] had indicated that passing detections to a centralised tracker significantly increases the area over which a target will be tracked compared to the situation in which each sonar in the network performs tracking separately and then pools tracks. The earlier study suggested that significantly improved performance can be obtained with a small network; in fact it was concluded networking just two sonars gives a substantial improvement. However, this previous work focussed on track initiation, without considering the rest of the tracking process. It also did not analyse the effect of false detections beyond a consideration of the effect of networking on false-detection rate, nor the difficulties in associating new detections with existing tracks. The purpose of the present work is to analyse all of these issues.

Because the purpose is to explore the impact of networking, rather than tactics or tracker development, it was not necessary to use a complicated scenario or a sophisticated tracking algorithm. The scenario chosen is a realisation of one of the four standard ASW scenarios [6], involving a small task group of three sonar systems that encounters an adversary submarine. The sonars generate true detections of the submarine and false detections that are randomly distributed over their fields of view. The tracker is a standard Kalman filter with data association by nearest neighbour in Mahalanobis distance. All detections, true and false, are treated equally by the tracker. The '3 detections in 5 consecutive pings' rule is used both for track initiation and track termination. Potentially crossing tracks are handled by allowing a detection to be associated with more than one existing track.

As in the analytical work, we compare the tracking performance when tracks are formed centrally using pooled detections from a group of sonars (centralised tracking) with the case where each individual sonar forms tracks using only its own detections (distributed tracking). When pooling detections, we tested for and eliminated duplicate detections using Mahalanobis distance. In the distributed-tracking case, we fused tracks from the sonars to give a single track list. Thus we obtain two track lists for comparison, one from centralised tracking and one from distributed tracking. Both track lists are formed from the same set of detections.

The two track lists were repeatedly generated in a Monte-Carlo manner. Parameters such as sonar spacing and range dependence of detection probability were chosen to give a high likelihood that the target track would not be continuous, but come and go. That is, the parameters chosen meant that the 3-in-5 rule would not be fulfilled on an ongoing basis. The resulting instances of track-list pairs were compared using a range of metrics, with the following results:

- Centralised tracking produces a true target track earlier than distributed tracking, with a median of 4 pings between the first target detection and the initialisation of the track, compared with 8 pings for distributed tracking. Centralised tracking also greatly reduces the percentage of cases in which a large number of pings are required to start a track, for example reducing the number of cases in which there are more than 15 pings between the first target detection and the confirmation of a track, from 19% to about 2%.

- For a network of three sensors, with $P_{fa} = 10^{-5}$, centralised tracking produces, on average, almost 4 times the number of false tracks as distributed tracking. However, with this network size and false detection probability, false tracks are short for both tracking concepts, with 99% of all false tracks lasting for no more than 7 pings. Thus in this example track length is of relatively equal utility in classifying tracks from both tracking options. This is not necessarily true for larger networks.
- The number of false tracks formed depends on the type of tracking, the number of sensors in the network and the probability P_{fa} of false detection per range-bearing cell per ping. As the size of the network increases, the false track lengths for centralised tracking increase. With larger networks, this makes track length a less efficient classification tool for centralised tracking.
- The total number of false tracks over the whole scenario is less important as a measure of operator workload than the number of tracks at any given time during the scenario. Both tracking types give reasonably steady distributions of track number over the duration of the scenario, although the number distribution for centralised tracking shows a tendency to greater fluctuations than for distributed tracking, but only in the maximum number of tracks at any given time. The mean, 5th and 95th percentiles are all as little varying with time for centralised tracking as for distributed tracking.
- Centralised tracking gives longer target tracks, on average, than distributed tracking. For this metric, the mean behaviour is not as important as the extremes of the distribution, for the ideal situation comprises a single very long target track. Centralised tracking shows a greater fraction of runs with target tracks exceeding 15 pings in length.
- The system confusion matrix provides a way to quantify the utility of a classification rule. The rule we have considered was the use of track length, with the track length threshold set to 5 pings. Classification of tracks using this rule gives a higher success rate for false tracks for distributed tracking, and for target tracks for centralised tracking, reflecting the previous observations made about track lengths.

Taken as a whole, these metrics indicate that the benefits of centralised tracking over distributed tracking survive the more stringent test of the simulation model, particularly for small networks. However the use of centralised tracking with larger networks or higher false detection rates than we have considered in this example will result in a significant increase in the false track rate. In order to fully realise the benefits of centralised tracking with a network of sonars, we will need to develop techniques and tools to help sonar operators to deal with the increased false alarm rate. These may include filtering on other sonar-signal properties (e.g. Doppler shift, signal level), fusion with other information, such as a tactical picture, or the active classification methods currently under development in MOD [12, 13].

5. References

- 1 Bell, J., Clothier, J. and O'Neill, J (2003) 'Options for Australian Defence in Network Centric Warfare', Client Report DSTO-CR-0339 of the Defence Science and Technology Organisation.
- 2 Fewell, M.P. and Hazen, M.G. (2003) 'Network-Centric Warfare: Its Nature and Modelling', Research Report DSTO-RR-0262 of the Defence Science and Technology Organisation.
- 3 Davies, A. (2007) 'The Enemy Below: Anti-Submarine Warfare in the ADF', Special Report no. 2 of the Australian Strategic Policy Institute, available at <www.aspi.org.au >.
- 4 Fewell, M.P., Thredgold, J.M. and Kershaw, D.J. (2008) 'Benefits of Sharing Detections for Networked Track Initiation in Anti-Submarine Warfare', Technical Report DSTO-TR-2086 of the Defence Science and Technology Organisation.
- 5 Thredgold, J.M., Lourey, S.J., Vu, H.X. and Fewell, M.P. (2010) 'A Simulation Model of Networked Tracking for Anti-Submarine Warfare', Technical Report DSTO-TR-2372 of the Defence Science and Technology Organisation.
- 6 Forward, T. and Gamble, G. (2004) 'Anti-Submarine Warfare Scenarios for Operations Research', General Document DSTO-GD-0387 of the Defence Science and Technology Organisation.
- 7 Colegrove, S.B., Davey, S.J. and Cheung, B. (2005) 'A Tracker Assessment Tool for Comparing Tracker Performance', Technical Report DSTO-TR-1694 of the Defence Science and Technology Organisation.
- 8 Burdic, W.S., (1991) *Underwater Acoustic System Analysis*, 2nd Edition, Prentice Hall, Englewood Cliffs NJ.
- 9 Duda, R.O., Hart, P.E. and Stork, D.G. (2001) *Pattern Classification* (2nd edn), Wiley, New York.
- 10 Bar-Shalom, Y. and Li, X.-R. (1995) *Multitarget-Multisensor Tracking: Principles and Techniques*, YBS, Storrs CT.
- 11 Blom, H.A.P. and Bloem, E.A. (2000) 'Probabilistic Data Association Avoiding Track Coalescence', *IEEE Trans. Automatic Control*, vol. 45, pp. 247-259.
- 12 von Trojan, A. and Kouzoubov, A. (2007) 'Active Sonar Classification Research Tool Capabilities and Implementation Outcomes', Technical Report DSTO-TR-2042 of the Defence Science and Technology Organisation
- 13 von Trojan, A. and Kouzoubov, A. (2008) 'Active Sonar Classification using a Binary Kernel Machine Classifier', Technical Report DSTO-TR-2110 of the Defence Science and Technology Organisation.

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA				1. PRIVACY MARKING/CAVEAT (OF DOCUMENT)	
2. TITLE Distributed versus Centralised Tracking in Networked Anti-Submarine Warfare			3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION) Document (U) Title (U) Abstract (U)		
4. AUTHOR(S) J. M. Thredgold and M. P. Fewell			5. CORPORATE AUTHOR DSTO Defence Science and Technology Organisation PO Box 1500 Edinburgh South Australia 5111 Australia		
6a. DSTO NUMBER DSTO-TR-2373		6b. AR NUMBER AR-014-687		6c. TYPE OF REPORT Technical Report	7. DOCUMENT DATE January 2010
8. FILE NUMBER 2009/1140439	9. TASK NUMBER ERP 07/310	10. TASK SPONSOR CMOD	11. NO. OF PAGES 20		12. NO. OF REFERENCES 13
13. URL on the World Wide Web http://www.dsto.defence.gov.au/corporate/reports/DSTO-TR-2373.pdf				14. RELEASE AUTHORITY Chief, Maritime Operations Division	
15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT <i>Approved for public release</i> <small>OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE, PO BOX 1500, EDINBURGH, SA 5111</small>					
16. DELIBERATE ANNOUNCEMENT No Limitations					
17. CITATION IN OTHER DOCUMENTS Yes					
18. DSTO RESEARCH LIBRARY THESAURUS http://web-vic.dsto.defence.gov.au/workareas/library/resources/dsto_thesaurus.shtml Antisubmarine warfare, Networked active sonar, Tracking, False alarms					
19. ABSTRACT This report describes a study of active sonar tracking, which explores the effect of networking sonars on tracking performance. We compare the tracking performance when sonars share detections (centralised tracking) with the performance when sonars share tracks (distributed tracking). Provided that the sonar layout and detection probabilities are such that multiple sonars have a reasonable probability (~30%) of obtaining detections from a target, we show that centralised tracking decreases the time to confirm a track on a target and improves the continuity of the target track. These improvements in target tracking occur at the expense of an increase in false track rate.					