



Australian Government
Department of Defence
Defence Science and
Technology Organisation

Benefits of Sharing Detections for Networked Track Initiation in Anti-Submarine Warfare

M.P. Fewell, J.M. Thredgold and D.J. Kershaw

Maritime Operations Division
Defence Science and Technology Organisation

DSTO-TR-2086

ABSTRACT

This report presents a quantitative study of a class of networking benefits in anti-submarine warfare. We show that sharing detections can produce an advantage over sharing only track-level information. We also indicate the conditions under which the advantage should be present, and estimate the magnitude of the advantage. This is achieved by focusing on the step of centralised track initiation, using metrics based mainly on sonar coverage area. We analyse multiple monostatic sonar, to give a concrete example aligned with current practice. The conclusions may be summarised in the statement that a 30% detection probability can be tactically useful, provided that there are other sonars with a similar P_d for the target concerned that are sharing information on detections. This result may provide a practical way around the great and continuing difficulty in obtaining acceptably high P_d values at tactically useful distances from a single sonar.

RELEASE LIMITATION

Approved for public release

Published by

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

*Telephone: (08) 8259 5555
Fax: (08) 8259 6567*

*© Commonwealth of Australia 2008
AR-014-075
January 2008*

APPROVED FOR PUBLIC RELEASE

Benefits of Sharing Detections for Networked Track Initiation in Anti-Submarine Warfare

Executive Summary

This report presents a quantitative investigation of a mechanism for networking advantage in anti-submarine warfare (ASW). The motivation for the work stems from current tasking of the Australian Defence Organisation's Rapid Prototyping, Development and Evaluation Program, which aims to demonstrate submarine detection with 95% confidence at a range exceeding the size of the torpedo danger zone using existing sonar equipment with minimal modifications. The concept is to achieve this through networking sonar data from several assets.

The mechanism consists of sharing detections with a centralised tracking processor for the complete system, as opposed to each individual sonar system performing tracking on its own detections only and then sharing the track information. We focus on the step of track initiation which, for the purposes of this initial analysis, we take to require 3 detections in 5 consecutive ensonifications. To provide a concrete example aligned with current practice, we model multiple monostatic active sonar. As a metric for the magnitude of the networking benefit, we mainly use the area enclosed by the 95% contour of track-initiation probability ('coverage area'). Other metrics are explored, with no change to the general conclusions.

The results depend on the behaviour of detection probability P_d for each sonar system. If P_d rolls off slowly with range, then the analysis shows considerable advantage in sharing detections and performing tracking centrally. Alternatively, if the P_d curve has a sharp cut-off, like a cookie cutter, there is still considerable advantage in centralised tracking if the P_d inside the cookie cutter is less than 100%. We also show how pre-track-level networking can cover sonar blind zones.

The overall conclusion can be summarised as:

A 30% probability of detection can be very useful in ASW; it can be made so by sharing detection-level data with neighbouring sonars that have similar P_d values for the target concerned.

This may provide a practical way around the great and continuing difficulty in obtaining acceptably high P_d values at tactically useful distances from a single sonar. Also, it should be possible to obtain the improved performance with only relatively minor alterations to existing sonar systems.

A disadvantage to networking lies in the sensitivity of the metrics examined in this report to the design of the ASW screen for the task group. Network-centric warfare (NCW) can enhance operational flexibility and capability, but the price is additional design complexity. As a simple example, the wrong choice of inter-sonar spacing can negate any networking benefit. As well as demonstrating the potential for this benefit, our analysis indicates the severity of the performance decrement that can arise with a poor choice of concept of operations (CONOPS) for the deployment of a sonar network.

Much of the existing analysis concerning NCW has been at the operational or strategic warfare levels with a focus on sharing and fusing tactical pictures from a variety of platforms, and providing access to reachback information. The work in this report gives a quantitative indication of how NCW should be able to produce benefits at the tactical level.

Authors



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.



J.M. Thredgold

Maritime Operations Division

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



D.J. Kershaw

Maritime Operations Division

David Kershaw started in Defence as a Cadet Engineer with Navy Material in 1987 and transferred to DSTO in 1989 following completion of a B.Sc(Hons) in Physics and a B.E in Electrical and Computer Systems Engineering. From 1989 though to 2002 he worked in the area of torpedoes, torpedo defence and undersea warfare and covered everything from circuit analysis through to operations research. He also completed a PhD in Tracking Systems in 1994. He was Head of Group from 1999 through to 2002 with a particular feature of that time being lead scientist for the inaugural Headmark military experimentation wargame in 2002. During 2003 and 2004 he was the Navy Scientific Adviser before returning to DSTO Edinburgh as the Air Warfare Destroyer S&T Adviser in 2005. He continued in this role during 2006 while assuming responsibility for Surface Ship Combat Systems in 2006. Since March 2007 he has been Head Torpedo Systems within the Undersea Sensors and Weapon Systems Branch in Maritime Operations Division.

Contents

1. INTRODUCTION.....	1
2. CONCEPT OF OPERATIONS AND METHOD OF ANALYSIS.....	2
2.1 Probability of Detection	2
2.1.1 Individual Sonars.....	2
2.1.2 Networked Detection Probability.....	4
2.2 Probability of Track Initiation	4
2.3 Geometry and Metrics.....	6
3. RESULTS.....	7
3.1 Exponential Detection Probability	7
3.2 Fermi-Function Detection Probability	10
3.2.1 With $P_d(0) = 1$	10
3.2.2 With $P_d(0) = 0.8$	11
3.3 Exponential Detection Probability with a Blind Zone	12
3.4 Other Metrics	12
3.4.1 Networked Coverage Area per Sonar.....	13
3.4.2 Two Lineal Metrics	14
3.4.3 Dimensionless Comparisons.....	15
4. ASSUMPTIONS AND LIMITATIONS	16
4.1 False-Alarm Rate	16
4.1.1 When False-Alarm Rate is Peripheral	17
4.1.2 Effect of Adjusting Detection Threshold	18
4.1.3 Detection-Fusion Rules	20
4.2 Independence of Ensonifications	21
4.3 Considering Five Ensonifications Only	21
4.3.1 Extension to More Than Five Ensonifications	21
4.3.2 Results for the 3-in-5 Rule	24
4.3.3 Varying the Track-Initiation Rule.....	25
4.3.4 Practical Limitations on Q	25
4.4 Sonar Geometric Arrangement.....	27
4.5 Isotropic P_d Characteristic	28
4.6 Schematic P_d Characteristic which is the Same for all Sonars	29
4.7 Other Assumptions.....	29
<i>Static model; zero speed of advance</i>	<i>29</i>
<i>Multiple monostatic active sonar</i>	<i>30</i>
<i>Sonars do not interfere with each other</i>	<i>30</i>
<i>Sufficient communications capacity exists to share detections</i>	<i>30</i>
<i>Flat 2-dimensional geometry</i>	<i>30</i>
5. SUMMARY AND CONCLUSIONS	30

APPENDIX A : OVERALL AREA OF OVERLAPPING CIRCLES	33
APPENDIX B : MATHEMATICAL AND CONCEPTUAL DETAILS ON PROBABILITY OF FALSE ALARM	34
B.1 Method of Computing Figures 22 and 23	34
B.1.1 Concept	34
B.1.2 Mathematical Detail	36
B.2 Practical Values of P_{fa}.....	37
<i>Derivation of Equation (29)</i>	<i>38</i>
APPENDIX C : MATHEMATICAL AND COMPUTATIONAL DETAILS FOR §4.3	39
C.1 Derivation of Equations (10) – (13)	39
C.2 Simulation for the Calculation of P_{ti} per Q Ensonifications	40
C.2.1 MATLAB Code	40
C.2.2 Convergence	42
REFERENCES.....	44

Figures

Figure 1:	Variation of detection probability P_d per ensonification with range for the four curves used in this work.....	3
Figure 2:	Track-initiation probability P_{ti} as a function of detection probability P_d for a selection of track-initiation rules.....	5
Figure 3:	Track initiation probability P_{ti} per group of q ensonifications for the four P_d curves in Figure 1.....	5
Figure 4:	Networked track-initiation probability P_{ti} for 1 to 10 sonar systems passing information on detections to a centralised tracker	6
Figure 5:	Contours of 80% and 95% track-initiation probability per group of 5 ensonifications when four sonars each perform track initiation separately using organic detections only	7
Figure 6:	As in Figure 5, but also showing contours of networked track-initiation probability, where detections from all the ensonifications are shared and tracking is performed centrally	8
Figure 7:	Coverage-area ratios for the exponential P_d of Figure 3(a) and sonars arranged close-packed on an equilateral-triangular grid with an inter-sonar spacing of 10 km.....	8
Figure 8:	As in Figure 7, but for variation of inter-sonar spacing in the 5-sonar case.....	9
Figure 9:	Contours of track-initiation probability in the 5-sonar case showing the evolution of the 95% networked contour from a single connected region to five pieces as the inter-sonar separation is increased	9
Figure 10:	As in Figure 6, but for a Fermi-function detection probability with $P_d(0)=1.0$	10
Figure 11:	Like Figure 7, but for five sonars only and the Fermi-function P_d of Figure 3(b).....	10
Figure 12:	As in Figure 10, but with $P_d(0) = 0.80$	11
Figure 13:	Networked coverage area per sonar for the Fermi-function P_d of Figure 3(c) and sonars arranged close-packed on an equilateral-triangular grid 10 km on a side.....	12
Figure 14:	Contours of networked track-initiation probability for an exponential P_d with a blind zone	13
Figure 15:	Comparison of networked coverage area per sonar for the three P_d curves of Figure 3(a-c)	14
Figure 16:	'Width of networked coverage area' and 'largest minimum screen distance' illustrated in the case of 4 sonars	14

Figure 17: Variation of (a) width of networked coverage area and (b) largest minimum screen distance with number of close-packed sonars	15
Figure 18: Figure 1 redrawn with the range axis for each curve scaled by the respective half distances	15
Figure 19: Comparison of (a) coverage area per sonar A_c scaled by the square of the half distance $r_{1/2}$ and (b) largest minimum screen distance D_{lms} scaled by half distance for three P_d characteristics	16
Figure 20: Schematic ROC curves	18
Figure 21: Illustration of the concept behind the construction of ROC curves	19
Figure 22: Contours of networked track-initiation probability for 5 sonars comparing networked P_{ti} at constant detection threshold with that at constant false-alarm rate	19
Figure 23: Comparison of coverage-area ratios at constant detection threshold and constant false-alarm rate	20
Figure 24: Concept of analysis for extending the calculation of track-initiation probability with a 3-in-5 rule to more than 5 ensonifications	22
Figure 25: Tree diagram for the 1-in-1 track-initiation rule	22
Figure 26: Tree diagram for the 2-in-3 track-initiation rule	23
Figure 27: Variation of the track-initiation probability P_{ti} for the 3-in-5 rule with detection probability P_d and total number Q of ensonifications	24
Figure 28: Same as in Figure 27, but displayed as contours of P_{ti} in the P_d - Q plane	25
Figure 29: Track-initiation probability for various initiation rules as a function of total number Q of ensonifications considered	26
Figure 30: Contours at which $P_{ti} = 95\%$ in the P_d - Q plane for several track-initiation rules	26
Figure 31: Examples of varying the geometrical arrangement with four and six sonars	27
Figure 32: Effect of geometrical arrangement on coverage area ratio	27
Figure 33: Contours of P_{ti} for an array of five sonars each comprising linear hydrophone arrays	28
Figure 34: (a) Two typical detection-probability-range curves from acoustic-propagation modelling; (b) resulting contours of track-initiation probability for a network of three sonars	29
Figure 35: Three equal, equi-spaced circles that (a) satisfy and (b) do not satisfy Equation (16)	33
Figure 36: Graphical depiction of the calculation of the effect on P_d of an adjustment in threshold	35
Figure 37: Variation of networked coverage-area ratio for 5 sonars and an exponential P_d with assumed value of P_{fa1}	37

Figure 38: Scatter plots showing, for each number of iterations, the results of 100 runs of the simulation code to calculate P_{ti}	42
Figure 39: Standard deviations of each bundle of results in the scatter diagram of Figure 38	43
Figure 40: Differences between analytic (i.e. exact) results and the P_{ti} values returned by the simulation with 10^6 iterations.....	43

Tables

Table 1: Distances at which $P_d = 0.5$ for the curves shown in Figure 1	3
Table 2: MATLAB code for calculation of track-initiation probability	40
Table 3: Statistical behaviour of the data in Figure 40	43

Acronyms and Initialisms

ASW	anti-submarine warfare
BARRA	a type of sonobuoy
CONOPS	concept of operations
cw	continuous wave
HVU	high-value unit
MPA	maritime patrol aircraft
NCW	network-centric warfare
ROC	receiver operating characteristic
RPDE	Rapid Prototyping, Development and Evaluation
TACCO	tactical commander
TADIL	tactical data information link (e.g. LINK 11)

Variables

a	parameter of either the exponential or Fermi-function models for P_d
A_c	networked coverage area per sonar
b	parameter of the Fermi-function model for P_d
C	number of sonar-display bins in a kinematic cluster
d	detection index
D_{lms}	largest minimum screen distance
G	complement of the usual cumulative standard normal probability distribution
m	number of sonars
N	number of bins in a sonar display
p	minimum number of detections required before a track is initiated
P_d	probability of detection per ensonification
P_{dn}	overall detection probability of a network of sonars
P_{dk}	detection probability for sonar k
P_{fa}	probability of false alarm
P_{ti}	probability of track initiation
P_0	detection probability at $r = 0$; a parameter of either the exponential or Fermi-function models for P_d
q	number of consecutive ensonifications within which at least p detections must be made before a track is initiated
Q	total number of ensonifications considered for track initiation
r	range or distance from a sonar
r_b	dimension of the blind zone in a model of P_d
$r_{1/2}$	half distance: range at which $P_d = 0.5$
s	inter-sonar spacing
T	detection threshold
θ	azimuthal angle from the axis of a linear array such as a towed array
σ	standard deviation of the noise

1. Introduction

The Australian Defence Organisation's RPDE (Rapid Prototyping, Development and Evaluation) Program currently has a focus on improving performance in anti-submarine warfare (ASW) through networking sonar data from several assets. The goal is to achieve 95% confidence in detecting a submarine at a range exceeding the size of the torpedo danger zone using existing equipment with minimal modifications.

For such a program to succeed, it is necessary that the mechanism of networking advantage be identified—how exactly will data sharing among platforms enhance ASW performance? One of several strategies under consideration is the exchange of pre-track-level data such as detections, in contrast to the current practice of sharing tracks via tactical data links (TADILs). To explore this possibility, this report presents some simple quantitative operations analysis that

- shows that sharing data on detections can produce an advantage over the sharing of information on tracks,
- indicates the conditions under which the advantage may be present, and
- estimates the approximate magnitude of the advantage.

It should be stated at the outset that the concept of operations (CONOPS) analysed in this report is not new—it has been of interest to tracking theorists for well over a decade (e.g. [1–5]) and is essentially current practice for a maritime patrol aircraft monitoring a sonobuoy field. However, the theoretical studies are abstract or use unrealistic assumptions.^(a) Among operators, it does not seem to be widely recognised that the CONOPS is applicable to all types of sonars.

For the purpose of this report, the analysis considers multiple monostatic active sonars; these may be on ships, sonobuoys, dipped from a helicopter, lying on the seabed, or deployed in any other manner. The analysis of specific sonar systems is beyond the scope of this report.

In the next section, the networked ASW CONOPS to be analysed is outlined, along with the method of analysis. Section 3 provides the key results of this work and §4 identifies and looks at the implications of the assumptions and limitations of the analysis. Section 5 provides a summary of the paper and its conclusions.

^(a)e.g. that all sensors have the same detection probability independent of time and distance to the target [5].

2. Concept of Operations and Method of Analysis

The hypothesis being explored is:

Improved ASW performance is achieved in a networked force if sonar systems exchange detection information, rather than exchanging track information.

The analysis examines this hypothesis by looking at track initiation: ‘global’ or ‘networked’ or ‘centralised’ track initiation is compared with track initiation by individual platforms using organic data only. This requires a model for probability of detection P_d and a rule prescribing when a track is initiated. We used four simple models of P_d , chosen to illustrate a range of behaviour, as described in §2.1.1. For track initiation (§2.2), we consider rules of the form: at least p detections in q consecutive ensonifications. Any values of p and q are allowable in the modelling framework; for this initial analysis, we use 3 in 5 ($p = 3, q = 5$) for computing the results in §3. (For a real system, variation of this rule would be one of several design parameters to be investigated.)

To keep the analysis very simple, we make numerous assumptions, such as no mutual interference between sonars and the existence of appropriate communications channels. The analysis can be viewed as indicating the level of benefit to be gained if these practical issues can be managed, thereby providing a motivation for tackling the technical challenges.

2.1 Probability of Detection

2.1.1 Individual Sonars

One of the difficulties in undertaking sonar analysis is the complexity of the sonar detection curves for real sonar systems. These curves are normally very complicated, with detection probability P_d per ensonification typically varying non-monotonically with range. The details of the range dependence are heavily influenced by environmental conditions, which fluctuate in time, and also on the settings of the individual sonar processors. In order to isolate the key features of the range dependence that impact on networking benefit, it is desirable to select representative detection probability curves that capture the important characteristics of the real P_d curves while passing over less important details. Figure 1 shows the curves selected. The first representation chosen is a simple exponential form, as an example of a P_d with a long low-probability tail:

$$P_d(r) = P_0 \exp(-r/a). \quad (1)$$

For definiteness, we take $P_0 = 1.0$ and $a = 10.0$ km, though the value of a does no more than set the overall scale of the problem.

Although some sonar systems have a P_d characteristic with features similar to the exponential, others have a much sharper cut-off at a certain range. To explore the effect of this, we use the Fermi function:

$$P_d(r) = P_0 \frac{1 + \exp(-b/a)}{1 + \exp\frac{r-b}{a}}, \quad (2)$$

with $P_0 = 1.0$ or 0.8 , $a = 1.5$ km and $b = 9.0$ km.

All active sonars have a blind zone in their immediate neighbourhood. For mono-




Figure 1: Variation of detection probability P_d per ensonification with range for the four curves used in this work. (The exponential curve is used both with and without the blind zone; see text for details.)

static sonars, this arises because the sonar cannot receive while it is transmitting.^(b) The blind zone may become especially serious as the sonar operator increases the pulse length to raise the transmitted energy in an attempt to detect a distant target. We return to the exponential form for P_d to model blind zones:

$$P_d(r) = \begin{cases} 0 & r \leq r_b \\ P_0 \exp(-r/a) & r > r_b, \end{cases} \quad (3)$$

where r_b is the radius of the blind zone. We use $P_0 = 1.0$, $a = 10.0$ km, as before, and choose an extreme case: a blind zone 2 km in radius, equivalent to transmitting a ~ 2.7 s long pulse.

The parameter values quoted above were chosen to give P_d curves that ‘look’ similar in some sense. In comparing the performance of sonar systems, attention is often paid to the range at which $P_d = 0.5$; we term this the ‘half distance’ $r_{1/2}$. For the above functions, it is given by

$$r_{1/2} = \begin{cases} a \ln(2P_0) & \text{(exponential)} \\ b + a \ln[2P_0(1 + e^{-b/a}) - 1] & \text{(Fermi function).} \end{cases} \quad (4)$$

Half distance is a simplistic metric for comparing P_d curves but, in view of its operational use as a quick rule-of-thumb for indicating sonar performance, we quote in Table 1 values of $r_{1/2}$ for the P_d curves in Figure 1.

Table 1: Distances at which $P_d = 0.5$ for the curves shown in Figure 1.

	half distance (km)
exponential	6.93
Fermi, $P_0 = 1.0$	9.01
Fermi, $P_0 = 0.8$	8.24

^(b)For multistatic sonars, echoes from targets are swamped by the direct blast from the transmitter, which implies a blind zone around each receiver.

We neglect any angular variation of P_d . This is sufficient for modelling approximately spherical systems, like hull-mounted sonars, but not for linear systems like towed arrays. An indication of the effect of angular variations in beam pattern is given in §4.5.

2.1.2 Networked Detection Probability

Equations (1)–(3) give detection probabilities per ensonification by a single sonar. If m independent sonars emit one ensonification each, then the probability that at least one of them makes a detection – the ‘networked detection probability’ P_{dn} – is

$$P_{dn} = 1 - \prod_{k=1}^m (1 - P_{dk}), \quad (5)$$

where k enumerates the sonars and P_{dk} is the detection probability for sonar k , which, for the purposes of this study, we take to have one of the forms in §2.1.1.

Equation (5) assumes statistical independence of the ensonifications; this assumption is discussed in §4.2.

2.2 Probability of Track Initiation

Once a contact has been detected, it is necessary to decide when to commence tracking it. Track-initiation criteria often have the form of requiring a minimum number p of detections in a given number q of consecutive opportunities. This style of rule permits the derivation of track-initiation probability P_{ti} from detection probabilities and the assumption of statistical independence for successive ensonifications:

$$P_{ti} = \sum_{j=p}^q \binom{q}{j} P_d^j (1 - P_d)^{q-j}, \quad (6)$$

where j enumerates the number of detections in the q consecutive ensonifications. In this equation, P_{ti} means a probability per q ensonifications and P_d can refer either to a single-sonar value (Eqs 1–3) or the networked value of Equation (5). This provides the two sets of P_{ti} values with which to test the hypothesis stated on page 2.

The relationship between P_d and P_{ti} is shown in Figure 2 for a selection of values of p and q . The track-initiation rules shown are: 2 or more detections in 3 ensonifications, and 3 or more detections in 4 or 5 ensonifications.^(c) For all rules, P_{ti} exceeds P_d for large enough P_d and falls below it at low P_d . It is a property of both the 2-in-3 and the 3-in-5 rules that the crossing point occurs at $P_d = 0.5$. That is, P_{ti} exceeds P_d when $P_d > 0.5$ and P_{ti} is less than P_d when $P_d < 0.5$. In the case of the 3-in-4 rule, the crossing occurs at $P_d = (1 + \sqrt{13})/6 \approx 0.768$.

Figure 3 shows the range variation of single-sonar P_{ti} for each of the P_d curves in

^(c) For these cases, Eq. (6) reduces to

$$P_{ti} = \begin{cases} P_d^2 (3 - 2P_d) & p = 2, q = 3 \\ P_d^3 (4 - 3P_d) & p = 3, q = 4 \\ P_d^3 (10 - 15P_d + 6P_d^2) & p = 3, q = 5. \end{cases}$$




Figure 2: Track-initiation probability P_{ti} per q ensonifications as a function of detection probability P_d for a selection of track-initiation rules. Although the rules are expressed as ‘ p in q ’, the actual interpretation is ‘at least p detections in q consecutive ensonifications’. For reference, the dotted line shows $P_{ti} = P_d$ and the broken lines show $P_{ti} = 0.5$ and $P_d = 0.5$.




Figure 3: Track initiation probability P_{ti} per group of q consecutive ensonifications for the four cases in Figure 1: (a) exponential, (b) Fermi function with $P_0 = 1.0$, (c) Fermi function with $P_0 = 0.8$, (d) exponential with a blind zone. Each panel shows: P_d (chain curve), $p = 3$, $q = 5$ (red), $p = 2$, $q = 3$ (blue), $p = 3$, $q = 4$ (green).

Figure 1. Compared with P_d , P_{ti} values are larger at short range and smaller at long range.

To demonstrate the effect of networking on P_{ti} , we make—for the moment—the quite artificial assumption that all sonars in the network have the same value of P_d . Use of Equation (5) in Equation (6) then gives the results shown in Figure 4 for the 3-in-




Figure 4: Networked track-initiation probability P_{ti} for 1 to 10 sonar systems passing information on detections to a centralised tracker that uses the 3-in-5 track-initiation rule, assuming that all sonars have the same detection probability. The red curve for one sonar is the same as the red curve in Figure 2.

5 track-initiation rule. The key feature of this figure is the manner in which networked P_{ti} can attain large values at quite low detection probabilities. It is plain that a P_d value as low as 0.3 can be very useful in a networking context. Figure 4 also shows the significant gain available from networking as few as two sonars together. This situation is the opposite of that often claimed in NCW, where the gain from adding a node to the network increases non-linearly with the total number of nodes. Here, the greatest gain per added node is obtained with the mere establishment of the network.

For definiteness, we adopt the 3-in-5 rule for the calculations in the remainder of this report. Not only is it commonly used, but also the differences between the various rules are small enough that adopting a different rule changes only the detail of the results, without affecting the overall conclusions.

2.3 Geometry and Metrics

To turn networked P_{ti} into a useful metric, it is necessary to drop the assumption that all sonars in the network have equal detection probabilities. Since P_d depends on range, this requires specifying the physical disposition of the sonar systems relative to each other. For this report, the individual sonar systems are assumed to be positioned close-packed on an equilateral-triangular lattice. Figure 5 shows the layout for four sonars. We have not attempted to explore other geometric arrangements of the sonars beyond two indicative cases presented in §4.4.

Figure 5 also shows contours of constant P_{ti} , using the 3-in-5 rule, an exponential P_d and assuming that each sonar conducts tracking independently of the others using organic detections only (the non-NCW case). The results of the next section focus on the 95% contour (red in Fig. 5) because the brief to the RPDE Program emphasises this level of confidence. The 80% contour is also shown for comparison.

The following section presents contours of constant P_{ti} as the number and spacing of sonars is varied for the two cases of track initiation being performed either individually by each sonar—as shown in Figure 5—or centrally using shared detection data. As a measure of the relative benefit of the two concepts of operations, we mainly use ‘coverage area ratio’:




Figure 5: Contours of 80% (blue) and 95% (red) track-initiation probability P_{ti} per group of 5 ensonifications for 4 sonars spaced 10 km apart, each performing track initiation separately using organic detections only. The exponential detection probability (Fig. 3a) is assumed.

$$\text{coverage area ratio} = \frac{\text{area enclosed by 95\% networked } P_{ti} \text{ contour}}{\text{overall area enclosed by the 95\% single-sonar } P_{ti} \text{ contours}}. \quad (7)$$

However, some indicative examples of other metrics that may be of interest are presented in §3.4.

If the single-sonar contours do not overlap, as in Figure 5, and all sonars in the network have the same P_d vs range characteristic, then the denominator of Equation (7) can be simplified:

$$\text{coverage area ratio} = \frac{\text{area enclosed by 95\% networked } P_{ti} \text{ contour}}{m \times (\text{area within one 95\% single-sonar } P_{ti} \text{ contour})}, \quad (8)$$

where m is the number of sonars. Coverage areas are calculated using two MATLAB functions: ‘contour’ to determine the shape of the contours and ‘polyarea’ to calculate the area enclosed by a contour. The input to ‘contour’ was a lattice of P_{ti} values spaced by 0.07 km. Convergence checks for cases of interest here showed that this lattice spacing results in area values accurate to better than $\pm 0.5 \text{ km}^2$.

3. Results

3.1 Exponential Detection Probability

Figure 5 shows the platform-centric case, in which each sonar individually performs tracking using organic detections only. It corresponds to data fusion at the track level. That is, sonars do not report until a track is formed, and tracks are then fused. If, in contrast, we assume data fusion at the detection level, with centralised track initiation, then the picture looks like Figure 6. The outer, labelled, contours show P_{ti} values for centralised tracking (i.e. using Eq. 5 in Eq. 6).

The single-sonar contours in each panel of Figure 6 are identical, thereby providing a comparative scale across the panels. The key issue concerns the relative sizes of the single-sonar and networked contours: the arrows highlight how the 95% networked contour almost coincides with the 80% single-sonar contour in the 3-sonar case, but lies well outside it in the 5-sonar case. This is suggestive of the non-linear networking gain so often promised by proponents of network-centric warfare (NCW) [6–8], which seems to contradict the implication of Figure 4. Figure 7 explores this further by showing how the coverage-area ratio (Eq. 8) varies with the number of sonars for the geometries displayed in Figure 6. The value for one sonar is 1.0 (obviously), but rapidly rises with increasing number of sonars: with 3 sonars, the networked contour covers 6 times the area of the three single-sonar contours. It might seem that diminishing returns are setting in at about 6 sonars, which agrees qualitatively with the message of Figure 4,




Figure 6: As in Figure 5, but also showing contours of networked track-initiation probability (full labelled lines), where detections from all the ensonifications are shared and tracking is performed centrally. The panels show the geometrical arrangement chosen for 3–7 sonars, all equally spaced on a 10 km triangular grid. Note that the 95% networked contour progressively moves out from the single-sonar contours as the number of sonars increases, as the arrows in the top three panels indicate.




Figure 7: Coverage-area ratios for the exponential P_d of Figure 3(a) and sonars arranged as shown in Figure 6. The quantity plotted is the area enclosed by the 95% networked P_{ti} contour divided by the sum of the areas within the 95% single-sonar P_{ti} contours (Eq. 8).

though it must be remembered that other arrangements of the sonars may give bigger coverage areas. Two examples are presented in §4.4, but as yet this has not been systematically explored.

The results in Figure 7 apply to an inter-sonar spacing of 10 km. Figure 8 shows the behaviour as the separation between sonars is altered. In the case of 5 sonars (blue curve), the maximum in coverage-area ratio is reached at a spacing of about 14 km. The sharp decrease after 16 km separation is caused by the 95% contour breaking into pieces, as Figure 9 illustrates. As the inter-sonar separation s is increased, the 95%




Figure 8: As in Figure 7, but for variation of inter-sonar spacing in the 3-, 5- and 7-sonar cases. See text for interpretation of the arrows.




Figure 9: Contours of track-initiation probability in the 5-sonar case (networked – full lines; single-sonar – broken lines) showing the evolution of the 95% networked contour from a single connected region to five pieces as the inter-sonar separation s is increased: (a) $s = 16$ km, (b) $s = 16.5$ km, (c) $s = 17$ km, (d) $s = 18$ km.

networked P_{ti} contour develops scallops between the sonars; Figure 9(a) shows the extent of this at $s = 16$ km. As s is further increased, the contour around the two outer sonars necks off, producing a 95% contour with three pieces, as shown in Figure 9(c). By $s = 18$ km, the 95% contour has broken into 5 pieces (Fig. 9d). Further increases in s cause the networked contour to shrink to coincide with the single-sonar 95% contours as all networking advantage from sharing detections disappears. The 80% networked contour goes through an analogous evolution at larger s values.

The cases of 3 and 7 sonars behave similarly (green and red curves respectively in Fig. 8). The important feature shown by these curves is the increase in the maximum coverage-area ratio obtainable as the number of sonars is increased. This adds an additional dimension to the networking advantage. Figure 7 is constructed with constant inter-sonar spacing of 10 km, as indicated schematically by the vertical arrow in Figure 8. A similar figure using spacing that is optimal for the geometry chosen (oblique

arrow in Fig. 8) would show an even greater networking benefit than is apparent in Figure 7.

3.2 Fermi-Function Detection Probability

The results in Figures 5–9 are all obtained with a detection probability that has an exponential shape. Although some sonar systems have a P_d characteristic with similar features to the exponential, others have a much sharper cut-off at a certain range. To explore the extent to which the networking advantage shown in Figures 7 and 8 arises from the long low-probability tail of the exponential shape, this subsection repeats the calculation with a Fermi-function shape for P_d .

3.2.1 With $P_d(0) = 1$

Using Figure 3(b) for P_d and the 3-in-5 track initiation rule results in Figure 10, which is equivalent to Figure 6. The extent of the networking advantage is now much reduced. The overlapping single-sonar contours complicate the computation of the coverage-area ratio, for double counting must be avoided. The method is detailed in Appendix A and the results are shown in Figure 11. The minimum, near 12 km separation, is related to the sonar spacing beyond which there is no simultaneous overlap of three single-sonar 95% P_{ti} contours. It is remarkable that, near this minimum, there is a networking penalty rather than an advantage: the coverage-area ratio *decreases* (albeit it only




Figure 10: As in the top row of Figure 6, but for a Fermi-function detection probability with $P_d(0) = 1.0$.




Figure 11: Like Figure 8, but for the Fermi-function P_d of Figure 3(b). In calculating the single-sonar coverage areas, care was taken to avoid double counting where the 95% contours of neighbouring sonars overlap (see Appendix A for details).

slightly) in moving from 5 to 7 sonars. Networking penalties, each with a unique cause, have been observed in other studies of NCW (e.g. [9]).

Even at their maximum, the values in Figure 11 are strikingly small compared with the exponential P_d case.^(d) A little reflection shows the cause of this effect. For a P_d with a definite-range law (i.e. a ‘cookie-cutter’ shape with $P_d = 1.0$ inside the detection range), there can be no networking gain from sharing detections. This is because a sonar with this shape of P_d has

- 100% detection probability within its detection range, and so needs no assistance with detecting things in this region, and
- zero detection probability outside its detection range, and so cannot provide assistance with detections there.

It is the low-probability tail of the exponential shape that creates the opportunity for networking advantage, through the combination of low-probability regions from several sensors. This argument relies on the value of P_d inside the detection range being 100%. When a lesser value applies, then the opportunity for networking advantage from sharing detections again arises, as §3.2.2 below illustrates.

A second difference between Figure 11 and the exponential- P_d case (Fig. 8) is the lack of any variation of optimum sonar spacing with number of sonars. It is clear that the performance decrement in failing to use the optimum CONOPS for deployment of the sonar network can be severe if the shape of the P_d curve is unfavourable.

3.2.2 With $P_d(0) = 0.8$

To provide a stark illustration of the value of sharing detection-level data when P_d is less than 100%, we choose a modification of the definite-range law. We effectively retain the definite range law, but reduce the value of P_d inside the detection range. By using Equation (2) with $P_d(0) = 0.8$, we obtain a maximum value of track-initiation probability P_{ti} of 0.942. That is, no sonar by itself can achieve a 95% probability of track initiation anywhere. Yet, as Figure 12 shows, when detections are shared, appreciable areas are covered at the 95% networked- P_{ti} level and above.

Clearly, coverage-area ratio cannot be used in this case, because the single-sonar coverage areas are zero. Instead, we replace it with *networked coverage area per sonar*. Although this metric is of no value for comparing networked with platform-centric CONOPS, it does give a sense of the manner in which networking gain grows with the number of sonars. Figure 13 shows the results.

The value in diagrams like Figure 13 lies in their utility as an aid to deciding how many sonars to deploy. In the present case, the gain in coverage area is essentially constant at $\sim 90 \text{ km}^2$ for each additional sonar beyond about 4 sonars. Hence, if one has




Figure 12: As in Figure 10, but with $P_0 = 0.80$. Note the absence of the 95% single-sonar contours (red broken lines in other contour plots).

^(d)Recall that a coverage-area ratio of 1.0 means precisely zero networking gain.




Figure 13: Networked coverage area **per sonar** for the Fermi-function P_d of Figure 3(c) and sonars arranged close-packed on an equilateral-triangular grid 10 km on a side. The quantity plotted is the area enclosed by the 95% networked P_{ti} contour divided by the number of sonars. (The value for one sonar is zero because networked P_{ti} never reaches 95% for this system.)

7 sonars available, it is almost as good, from the point of view of total coverage area, to split them into two groups of 4 and 3 as to keep them close-packed in a single group. The choice between the arrangements then comes down to operational requirements of the situation at hand.

3.3 Exponential Detection Probability with a Blind Zone

As a final example, we illustrate the value of sharing detections for covering blind zones. These arise because a sonar cannot receive while it is transmitting. They may be especially serious as the sonar operator increases the pulse length to raise the transmitted energy in an attempt to detect a distant target. In the current example, we choose an extreme case: a blind zone 2 km in radius (Figs 1, 3d), equivalent to transmitting a ~ 2.7 -s long pulse.

Figure 14 shows a selection of results for arrays of 3, 4 and 5 sonars. In the bottom row of this figure, the inter-sonar separation s is 10 km, the value used in Figures 6, 10 and 12. With this separation, the 3-sonar array shows scarcely any filling in of the blind zones, even at the 80% level, although rather less than half of the 80% blind-zone contour remains in the outer sonars for the 5-sonar array. Reducing the inter-sonar separation improves matters, with just two small segments of the blind zones remaining at the 95% level for 5 sonars and an inter-sonar spacing of 7 km. If two more sonars are added, giving a 7-sonar array (not shown in Fig. 14), then all blind zones are covered at the 95% level for an inter-sonar spacing of 9.7 km.

3.4 Other Metrics

Most of the results above use coverage-area ratio (Eq. 7) as the performance metric. This seems appropriate for choke-point monitoring, but may be less important in other situations. Indeed, it could not be used in §3.2.2. In some cases, there may be more interest in the overall width of the 95% P_{ti} contour, or the shortest distance from this contour to a high-value unit. This section looks at these and other metrics from the points of view of displaying their features and asking whether any of them give a




Figure 14: Contours of networked track-initiation probability for an exponential P_d with a blind zone (red: $P_{ti} \geq 0.95$; blue: $P_{ti} \geq 80\%$). Top row – inter-sonar spacing of 7 km; middle row – 9 km; bottom row – 10 km. Note that single-sonar contours are **not** shown – the contours near each sonar show the effects of the blind zones.

picture of the occurrence of networking advantage different from that implied by coverage-area ratio.

3.4.1 Networked Coverage Area per Sonar

Networked coverage area per sonar is defined as the area enclosed by the 95% networked P_{ti} contour divided by the number of sonars in the pattern. It was used in §3.2.2 because coverage-area ratio could not be defined in that situation. It can readily be calculated for the cases of §§3.1 and 3.2.1 also, with results shown in Figure 15. The variation of networked coverage area per sonar with the number of sonars is the same as that of coverage-area ratio, where the second can be defined, if the single-sonar 95% contours do not overlap, for then the two are different only by a factor of the coverage area of a single sonar, which is constant as the number of sonars changes. Hence, the trend of the results in Figure 15 for the exponential P_d is the same as in Figure 7.

The Fermi P_d with $P_0 = 1.0$ shows a coverage area per sonar that decreases with increasing sonar numbers. At face value, it indicates that sonars with this P_d characteristic should be operated stand-alone rather than networked – the value for a single sonar is the largest. However, this effect is caused by the overlapping of the single-sonar 95% contours; the real message is that a 10 km inter-sonar spacing is too small given the shape of the P_d characteristic.




Figure 15: Comparison of networked coverage area per sonar for the three P_d curves of Figure 3(a–c) and an intersonar spacing of 10 km. The values for Fermi, $P_0 = 0.8$ are identical to those in Figure 13; they are included here solely for ease of comparison. (The value for one sonar, Fermi, $P_0 = 0.8$ is zero because networked P_{ti} never reaches 95% for this system.)

3.4.2 Two Lineal Metrics

When the task is ASW in support of a transiting task group, coverage area may be less important than the dimension of the screened area perpendicular to the mean line of advance, or the distance from the edge of the screen to a high-value unit (HVV). This suggests metrics based on length rather than area; the two considered here are illustrated in Figure 16. The ‘width of networked coverage area’ is the maximum linear extent of the 95% P_{ti} contour. The value of ‘largest minimum screen distance’ is determined by choosing a location for the HVV that maximises the shortest distance to the 95% P_{ti} contour. For the present work, both lineal metrics were estimated by measurement on large copies of the relevant contour plots. Results are shown in Figure 17.

A close-packed arrangement is obviously not optimal for maximising the width of the networked coverage area. For example, the gain in width in going from 5 to 7 sonars is small or zero in all cases (Fig. 17a). Clearly, a more linear arrangement would be better when coverage-area width is the pre-eminent metric.

On the other hand, a close-packed geometry performs much better with largest minimum screen distance as the metric, as Figure 17(b) shows: for all P_d shapes, 7 sonars perform significantly better than e.g. 5 sonars. Indeed, this metric rates the 7-sonar case comparatively more highly than any other of the metrics considered herein, owing




Figure 16: ‘Width of networked coverage area’ and ‘largest minimum screen distance’ in the case of 4 sonars. The largest minimum screen distance is determined by choosing a location for the high-value unit that maximises the shortest distance from it to the 95% contour of track-initiation probability.




Figure 17: Variation of (a) width of networked coverage area and (b) largest minimum screen distance with number of close-packed sonars (inter-sonar spacing of 10 km) for the three P_d characteristics of Figure 3(a-c).

to the high degree of symmetry of the configuration—the place to put the HVU is on top of the central sonar, that is, in the middle of Figure 6(e).

Largest minimum screen distance also shows interesting comparisons among the P_d shapes: with a small number of sonars, the Fermi shape is better than the exponential, but this is reversed for 5–7 sonars. This reflects the importance of the long tail of the P_d shape in providing networking advantage, at least so far as this metric is concerned.

3.4.3 Dimensionless Comparisons

The distance scales shown in, for example, Figures 8–12 are arbitrary in the sense that they are all related to the distance scale chosen in Figure 1. This arbitrariness can be removed by scaling by a characteristic distance, so allowing intrinsic features to be emphasised. There are several possible choices for the characteristic distance; to indicate the effect, we choose to use the half distances $r_{1/2}$ listed in Table 1. Half distance is a quantity frequently used as an indicator of sonar performance. Figure 18 shows Figure 1 with the range axis rescaled by the respective half distances. The quantity on the abscissa of Figure 18 is ‘dimensionless’ in the sense used in physical science; that is, it is independent of units of measurement.




Figure 18: Figure 1 (p. 3) redrawn with the range axis for each curve scaled by the respective half distances (Table 1).

Coverage-area ratio is already a dimensionless quantity but, as §3.2.2 suggests, it can be subject to distortion: if the P_{ti} level chosen for the contours is so high that the area within the single-sonar contours is small, then the coverage-area ratio becomes artificially large. Figure 19 shows the effect of scaling on two of the other metrics. In Figure 19(a), the networked coverage area per sonar A_c (Fig. 13) is scaled by the square of $r_{1/2}$. Compared with Figure 15, Figure 19(a) suggests that the exponential P_d characteristic becomes increasingly advantageous as the number of sonars in the network grows. This shows the effect of the long tail and the low value of $r_{1/2}$ relative to the Fermi function. Figure 19(a) also gives the interesting insight that a reduced value of P_0 becomes increasingly unimportant as the size of the network grows: the network of seven sonars with a Fermi-function characteristic shows little difference in the value of $A_c/r_{1/2}^2$ with either value of P_0 . Figure 19(b), showing a scaled largest minimum screen distance, tells essentially the same story.

4. Assumptions and Limitations

The preceding analysis depends crucially on the assumption that ‘detection probability’ is a useful concept for describing the physical reality of sonar systems. If this is not so, then the analysis is meaningless. If, however, the utility of detection probability is accepted, as it seems widely to be among the operator community, then the above results largely follow as a logical consequence. There are, of course, other assumptions in the analysis, many of which this section lists and comments on. Most of the more significant of these can be addressed by elaborations to the analysis, as sketched out in the ensuing subsections.

To keep the discussion as simple as practicable, the assumptions are relaxed one at a time, despite the obvious potential for interactions between several of them.

4.1 False-Alarm Rate

False-alarm rate is so much a preoccupation of sonar-system design that the lack of its mention in the body of this report may seem strange. In fact, depending on how the




Figure 19: Comparison of (a) networked coverage area per sonar A_c scaled by the square of the half distance $r_{1/2}$ and (b) largest minimum screen distance D_{lms} scaled by half distance for three P_d characteristics.

sonar systems are operated, this may be neither an oversight nor a limitation of the current analysis. Rather, false-alarm rate would be a peripheral consideration in the circumstances described in the next sub-section. When these circumstances do not hold, §4.1.2 describes an extension of the analysis that covers the situation by variation of detection threshold. Another approach involves the rule under which detections from more than one sensor are fused; this is touched upon in §4.1.3.

The concern over false-alarm rate may partly come down to terminology. We take ‘detection’ and ‘false alarm’ to have their engineering meanings; that is, a detection occurs when a signal rises above a predetermined detection threshold and a false alarm is a detection caused by a noise fluctuation. Detections of physical non-threat objects are not considered here to be false alarms. Another issue concerns the psychology of infrequent detections. If, in the real world of sonar practice, persistent detections engender confidence in the reality of a contact, then the reverse is also true: infrequent detections lead to doubt and, from there, to a widening of the notion of ‘false alarm’. Would a sonar operator be willing to report a contact to higher command if it is giving a return in only one in ten ensonifications? Yet if the sonar operator has four neighbours and all five of them have one return per ten ensonifications from the same location, then it would seem clear that there is advantage to be gained in sharing the information.

4.1.1 When False-Alarm Rate is Peripheral

According to the standard engineering model of detection (e.g. §§12.1–4 in [10]), a threshold is chosen that trades off false-alarm rate against detection probability. Suppose that this trade-off is conducted in the normal manner, and that the operator has selected a detection threshold to give a false-alarm rate with which he or she is comfortable. The detection probability is thereby set and detections occur. This is the point of departure of our analysis, with Figure 1 (p. 3) representing possible types of resulting P_d curves. That is, we assume that every step in the ASW process up to and including detection remain as currently performed. Our focus is to explore the questions: What happens to a detection once it is made? Do sonar operators post it on the network immediately, or do they keep it to themselves and wait for 3 detections in 5 consecutive ensonifications before posting it?

The above argument effectively assumes that thresholds are not altered from their platform-centric settings when detections are passed over a network for centralised tracking. This is plausible with automated detection and tracking, since the number m of networked sonars is unlikely in the foreseeable future to be more than a few tens at the very most. However, with a manual system, one may ask whether the operator performing the centralised tracking role is willing to tolerate a false-alarm rate that is m times the rate from a single sonar. If the tracking operator is unable to cope with this and responds by asking individual sonar operators to raise their detection thresholds, then the individual P_d values in the networked case will be lower than in the single-sonar-tracking case, so reducing the coverage area ratio. That is, the networking advantage would be reduced. The next subsection shows how this situation can be analysed.

The previous paragraph assumes that the effect of fusing detections from m sonars is to increase the false-alarm rate by a factor of m . This is the case if the fusion rule is a logical ‘or’; that is, the fused picture contains every detection that appears in any sonar. Other rules are possible. For example, the logical ‘and’ would place a detection in the

fused picture only if every sonar shows a contact in the same physical location. This would heavily suppress false alarms, to the extent that false-alarm rate may become unimportant, but it would also reject many real detections in situations of low P_d . Fusion rules are treated in §4.1.3; the next subsection works with the simple ‘or’ rule.

4.1.2 Effect of Adjusting Detection Threshold

To determine the effect of detection-threshold adjustment on coverage-area ratio, one requires:

- the relationship between false-alarm probability and detection probability in terms of detection threshold, and
- a rule for deciding what false-alarm rate the tracking operator will tolerate.

The relationship between detection probability P_d , false-alarm probability P_{fa} and detection threshold T is conventionally displayed as a ‘receiver operating characteristic’ (ROC) curve. Figure 20 shows the main features of ROC curves and Figure 21 illustrates the concept underlying their construction (e.g. §12.2 of [10]). The top panel of Figure 21 shows schematically the distribution of sensor outputs due to random or quasi-random noise only. The noise may arise from sources so distant that they can not be localised, or from distributed closer sources, such as waves, rain or reverberation, or may be hard to ascribe to any obvious physical source [11]. The lower panels of Figure 21 show the change when localisable objects, either distant or near, produce a sonar return in addition to the noise. The detection threshold is the operator-adjustable output level above which a detection is declared to have occurred. If there is no localisable object present, then an output above the threshold produces a false alarm, otherwise it is a true detection. The shaded areas in Figure 21 give the values for P_{fa} (top panel) and P_d (lower panels); it is clear how these vary as the threshold level is changed. It is conventional to depict this information as lines of constant target range r , giving ROC curves like those shown in Figure 20. On such a line, low threshold levels correspond to P_{fa} and P_d near 1.0, and high threshold levels to P_{fa} and P_d near zero.




Figure 20: Schematic ROC curves – variation of detection probability P_d with false-alarm probability P_{fa} for at a given target distance r as the detection threshold T is altered. Curves for two target distances r are shown.




Figure 21: Illustrating the concept behind the construction of ROC curves. The top panel shows schematically the distribution of sensor outputs when there is no physical object present and the lower panels show the same but with a physical object present. The operator-adjustable threshold level is the same in all panels.

For present purposes, we need constant threshold with varying target range, not the reverse. This is indicated by the vertical arrow in Figure 20. That is, we do not work on a single ROC curve, but rather across ROC curves. The method is described in Appendix B.1; results are shown in Figures 22 and 23.

Figure 22 shows the effect on networked P_{ti} contours and Figure 23 compares coverage-area ratios for constant detection threshold with those for constant false-alarm rate. Threshold adjustment has a large effect when the P_d is exponential. For the Fermi-function P_d , the effect is to reduce the coverage-area ratio below 1.0, which means that the networked system performs worse than the sonars forming tracks independently.




Figure 22: Contours of networked track-initiation probability for 5 sonars with (a) exponential P_d (Fig. 3a) and (b) Fermi-function P_d (Fig. 3b). Full curves show networked P_{ti} at constant detection threshold (the same as in Figs 6 and 10); chain curves show the effect of keeping false-alarm rate constant. Blue curves – 80% networked P_{ti} ; red curves – 95% networked P_{ti} . As before, sonars are arranged close packed on an equilateral-triangular grid 10 km on a side. Single-sonar contours are not shown, but are the same as in Figures 6 and 10 respectively.




Figure 23: Comparison of coverage-area ratios at constant detection threshold and constant false-alarm rate for sonars arranged as in Figure 22 with (a) exponential P_d and (b) Fermi-function P_d with $P_0 = 1.0$. The legend applies to both panels.

This is an example of a mechanism by which a networking decrement can arise. Figure 23 indicates the advantage to be gained with automated track initiation or any other method that can tolerate the higher false-alarm rate that comes with networking. It seems clear that, if the actual threshold adjustment is as large as assumed here, then the optimum spacing is rather less than 10 km if the P_d is exponential and rather more than 10 km if it is Fermi.

The impact of false-alarm rate on networked coverage-area requires further exploration, but the results in Figures 22 and 23 indicate that the modelling method presented here should provide a useful framework for this.

4.1.3 Detection-Fusion Rules

As well as adjusting detection thresholds, the false alarm rate can also be controlled by varying the rule under which detections qualify to be passed to the central tracker. Section 4.1.2 assumed a logical 'or' rule, which can be expressed as:

- ♦ Record a detection at location x, y in the fused picture if any sonar has a detection at that location.^(e)

This means that, if the m sonars all have the same false-alarm rate, then the field false-alarm rate is m times the single-sonar rate.

The analogy with logical operations suggests another rule, based on the logical 'and' operation:

- ♦ Record a detection at location x, y in the fused picture only if all sonars have a detection at that location.

This rule would suppress false alarms because they are not associated with localisable objects and so appear at randomly distributed locations. The degree of suppression depends on the false-alarm rate, being highest at low rate. However, this is not a practical rule in the sonar realm because it also suppresses real detections, unless P_d is very high for all sensors.

^(e)This and the other fusion rules are expressed as if the location of a detection is determined with arbitrarily high precision. They can be modified to take account of the effects of measurement uncertainty, but this adds a level of detail that is unnecessary for the present discussion.

The form of the above rules suggests a generalisation:

- ♦ Record a detection at location x, y in the fused picture if k of the m sonars have a detection at that location.

It has been shown theoretically that, for any desired value of field false-alarm rate, there is a value of k that maximises the field detection probability [2]. The form of the third rule above applies only if all sonars have the same detection probability and false-alarm rate; a generalisation to systems composed of sensors with different detection probabilities and false-alarm rates has also been derived [2,12,13]. The effect of using different detection-fusion rules on networked coverage-area ratios has not been explored, to our knowledge. It is likely that it would need to be studied on a case-by-case basis for specific operational systems.

4.2 Independence of Ensonifications

We have assumed that probabilities can be combined as though each ensonification is wholly independent, in the statistical sense, from any other. This may not be true if the sonar environment is changing slowly. Correlations between successive ensonifications are believed to be common in practical situations (e.g. [14], pp. 107 & 133 ff), although cases are known where there was no detectable correlation between ensonifications just seconds apart (e.g. [15]). Where correlations do occur, a pair of successive ensonifications each with, for example, $P_d = 50\%$ will have a combined P_d of less than the value of 75% that would hold were they independent. This affects the calculation of track-initiation probability.

We expect the possibility of ensonification-to-ensonification correlation to be much greater for consecutive ensonifications from one sonar than for two ensonifications from widely separated sonars. That is, the effect will impact on the single-sonar coverage areas much more than on the networked coverage area. Since the effect acts to reduce P_{ti} from the values that we calculate, the consequence is that the real coverage-area ratio would be larger than we have calculated.

In summary, if there is ensonification-to-ensonification correlation that affects P_{ti} values, then the actual networking advantage will be even greater than our results indicate.

4.3 Considering Five Ensonifications Only

The modelling methodology used in this paper examines track-initiation probabilities per group of 5 ensonifications. However, if a track is not initiated by the fifth ensonification, perhaps the sixth will be sufficient, and so on. Practically, this is akin to asking the question: how many ensonifications are required in order to gain a given degree of confidence that the submarine is not in this region? The implied extension to the present methodology can be handled as a combinatorial problem, as the next subsection shows.

4.3.1 Extension to More Than Five Ensonifications

For the purposes of this work, we concentrate on the 3-in-5 track-initiation rule. This can be maintained as more ensonifications are sent out by dropping early ensonifications. That is, we employ a sliding-window approach, as illustrated in Figure 24. The




Figure 24: Concept of analysis for extending the calculation of track-initiation probability with a 3-in-5 rule to more than 5 ensonifications. The vertical lines represent regularly spaced ensonifications and the long brackets represent successive applications of a sliding window. The track initiation rule is held to be satisfied with the first window that contains 3 detections.

track is initiated with the first window that contains 3 or more detections. If ensonifications are not dropped, then the window grows in length and the rule becomes 3 in more than 5, but these cases are beyond the scope of this work.

The sliding-window concept can obviously be generalised to other track-initiation rules: if one requires at least p detections in q consecutive ensonifications from a total of Q ensonifications, then the width of the sliding window is q and clearly $p \leq q \leq Q$. (Results in §3 have $q = Q$.) It is also clear from the concept, and the following results confirm, that $P_{ti} \rightarrow 1$ as $Q \rightarrow \infty$, so one must recognise practical upper limits to the value of Q . Some sources of such upper limits are explored in §4.3.4 below.

The simplest track-initiation rule is 1-in-1, which means that a track is initiated at every detection. For this rule, it is straightforward to derive an analytic expression for the variation of P_{ti} with Q . The logical progression implied by the rule can be displayed in the ‘tree diagram’ shown in Figure 25. At each ensonification, the probability of not initiating a track is $1 - P_d$, that is, the same as the probability of not making a detection. The probability of not making a detection after Q ensonifications is therefore $(1 - P_d)^Q$. Hence, the track-initiation probability $P_{ti,1|1}$ per group of Q ensonifications is

$$P_{ti,1|1} = 1 - (1 - P_d)^Q. \quad (9)$$

This shows the result mentioned above: provided that $P_d \neq 0$, $P_{ti} \rightarrow 1$ as $Q \rightarrow \infty$.

Unfortunately, the line of reasoning leading to Figure 25 and Equation (9) is not easy to generalise to other track-initiation rules. For example, Figure 26 shows the tree




Figure 25: Tree diagram for the 1-in-1 track-initiation rule carried to 3 ensonifications.




Figure 26: Top part of the infinitely deep tree diagram for the 2-in-3 track-initiation rule, carried to 7 ensounifications. Arrows marked '2 | 3, stop' show places where the 2-in-3 rule is fulfilled. ' p ' is used in place of P_d for the sake of brevity. Arrows leading from the 'ping 7' boxes are not shown. The ovals enclose instances of the fundamental repeating unit, which forms the basis of the derivation of the recurrence relation (Appendix C.1).

diagram for a 2-in-3 rule, which is perhaps the simplest interesting case. Analysis of this diagram, presented in Appendix C.1, leads to a recurrence relation for the probability $\bar{P}_{ti,2|3}(Q)$ of *not* initiating a track after Q ensonifications:^(f)

$$\bar{P}_{ti,2|3}(Q) = (1 - P_d) \bar{P}_{ti,2|3}(Q - 1) + P_d (1 - P_d)^2 \bar{P}_{ti,2|3}(Q - 3), \quad (10)$$

with the initial values

$$\bar{P}_{ti,2|3}(0) = 1, \quad \bar{P}_{ti,2|3}(1) = 1, \quad \bar{P}_{ti,2|3}(2) = 1 - P_d^2. \quad (11)$$

We also obtained a recurrence relation for the simpler 2-in-2 rule, of interest when manual track initiation is performed with a significant false-alarm rate (§4.1.2). That recurrence relation is (Appendix C.1)

$$\bar{P}_{ti,2|2}(Q) = (1 - P_d) \bar{P}_{ti,2|2}(Q - 1) + P_d (1 - P_d) \bar{P}_{ti,2|2}(Q - 2), \quad (12)$$

with the initial values

$$\bar{P}_{ti,2|2}(0) = 1, \quad \bar{P}_{ti,2|2}(1) = 1. \quad (13)$$

To perform general calculations, we developed a MATLAB simulation that carries out many trials and computes the probability as the mean success rate. The code is reproduced in Appendix C.2.1, together with tests of its accuracy and convergence. The following subsections present some results.

4.3.2 Results for the 3-in-5 Rule

The behaviour of the 3-in-5-rule P_{ti} as Q is increased beyond 5 is illustrated in Figures 27 and 28. Figure 27(a) shows P_{ti} as a function of detection probability. The black curve ($Q = 5$) is the same behaviour as shown in Figure 3, with $P_{ti} = 0.5$ when $P_d = 0.5$. As further ensonifications are considered, P_{ti} rises for any given P_d value, but Figure 27(a)




Figure 27: Variation of the track-initiation probability P_{ti} for the 3-in-5 rule with (a) detection probability P_d and (b) total number Q of ensonifications. In panel (b), the lines serve merely to guide the eye.

^(f) Castella [16] has published recurrence relations for $P_{ti,p|q}(Q)$, i.e. the complement of the probability in Eqs (10) and (12), for all valid combinations of p and q with $q \leq 4$. Equations (10) and (12) agree with, and are much simpler than, his expressions. A previous DSTO report on sliding-window methods [17] appears to contain some erroneous expressions (e.g. its Eqs 2 and 3).




Figure 28: Same as in Figure 27, but displayed as contours of P_{ti} in the P_d - Q plane.

suggests diminishing returns: the curves for $Q = 9$ and 11 are closer together than the curves for $Q = 5$ and 7 .

Figure 27(b) is a plot of similar data to Figure 27(a), but as a function of Q . Diminishing returns are quite evident at high P_d , but not apparent at $P_d = 0.2$. Figure 28 shows the same data again, but this time as a contour plot in the P_d - Q plane. The right-most contour in this plot shows the values of Q that are required to give $P_{ti} = 0.95$ as a function of P_d .

4.3.3 Varying the Track-Initiation Rule

The analysis of this subsection (§4.3) provides a context for comparing track-initiation rules, in the manner shown in Figure 29. The three panels of this figure show values of P_{ti} as a function of Q for a variety of track-initiation rules. Each rule is expressed as p -in- q , which means p or more detections in q consecutive ensonifications. As well as the commonly used rules 2-in-3, 3-in-4 and 3-in-5, Figure 29 shows 1-in-1, since this is the only curve for which we have an analytic expression in closed form (Eq. 9) and 2-in-2, because of the importance of this rule in manual track initiation when there is an appreciable false-alarm rate (§4.1.2).

It is seen that the highest value of P_{ti} always occurs for the 1-in-1 rule, but this means starting a track on every detection, and so will produce many false tracks when the false-alarm rate is high. The 2-in-3 and 2-in-2 rules have the next highest values of P_{ti} , though Figure 29(c) shows that the 3-in-5 rule gives essentially the same P_{ti} values as the 2-in-2 rule when detection probabilities are high.

Figure 30 compares track-initiation rules using a contour plot. The lines are loci of $P_{ti} = 95\%$ in the P_{ti} - Q plane for the rules shown. That is, for each Q value, the curves show P_d values that must be obtained for a 95% probability of initiating a track in the Q ensonifications. This plot tells the same story as Figure 29: the 1-in-1 rule gives a 95% P_{ti} at the lowest P_d value for all values of Q .

4.3.4 Practical Limitations on Q

The diminishing returns explored in the last section provide one source of a practical upper limit on the number Q of ensonifications considered. At some point, the gain in waiting for yet one more ensonification becomes too small given the length of time involved. This is much more of an issue in sonar than in radar because of the very low

repetition rate enforced by the slow speed of sound in water. Unless de-interleaving techniques are used, which are uncommon in sonar, there must be many seconds between successive ensonifications when searching for a distant target.

A second source of practical limitation to the value of Q derives from the scenario, which we effectively assume to be static. This assumption is commented on further in §4.7. Here, we note only that any significant relative motion of own sonars and potential targets means that early ensonifications in a sequence steadily lose value as the sequence is lengthened.




Figure 29: Track-initiation probability for various initiation rules as a function of total number Q of ensonifications considered for $P_d =$ (a) 0.2, (b) 0.5 and (c) 0.8. Each curve starts at the window length of the rule (i.e. at $Q = q$). In panel (c), the results for the 3-in-5 rule – starting at $Q = 5$ – lie essentially on top of those for the 2-in-2 rule. For reference, broken lines mark $P_{ti} = 0.95$. Note the suppressed zero in panel (c).




Figure 30: Contours at which $P_{ti} = 95\%$ in the P_d - Q plane for several track-initiation rules. As in Figure 29, each curve starts at $Q = q$.

4.4 Sonar Geometric Arrangement

Sonars are taken to be positioned close-packed on an equilateral-triangular lattice. This assumption was made for the sake of definiteness, to provide a concrete model to calculate. Any geometric arrangement can be accommodated in the model. As an indicative example, Figure 31 shows other geometries in the 4- and 6-sonar cases (exponential P_d without a blind zone) and Figure 32 compares coverage ratios for the geometries of Figure 31(b,d) with the close-packed arrangements. The wedge-shaped geometry of Figure 31(b) has essentially the same coverage ratio as four sonars close-packed, so that this change in geometry brings no improvement in the value of the metric. On the other hand, the hexagonal arrangement of Figure 31(d) has a rather larger coverage area than six sonars close-packed. In fact, it has a larger coverage-area ratio even than seven sonars close-packed, but that is due to the m factor in the denominator of Equation (8) (p. 7); its total networked coverage area (as distinct from coverage-area ratio) is less than that of seven sonars close-packed.




Figure 31: As in Figure 6 (p. 8) but showing examples of different geometric arrangements: (a) four sonars close-packed, (b) four sonars wedge-shaped, (c) six sonars close-packed, (d) six sonars symmetric.




Figure 32: Effect of geometric arrangement on coverage area ratio.

The exploration of other arrangements would be a useful extension of the present work. For example, in a certain sense linear geometries provide natural opposites to close-packed arrangements. These may be relevant to operational contexts involving task-group transit.

4.5 Isotropic P_d Characteristic

This is another assumption not required by the structure of the model. For example, Figure 33 shows a calculation in which the sonars have a characteristic typical of linear arrays of hydrophones (towed array or remotely deployed array), with zero P_d at the end-fire directions. This was achieved by writing

$$P_d(r, \theta) = P_d(r) \sin \theta, \quad (14)$$

where θ is an azimuthal angle from the axis of the hydrophone line ($0 \leq \theta \leq 180^\circ$) and $P_d(r)$ has one of the forms given in §2.1.1. The angular dependence in Equation (14) is intended to be schematic only, though it is reminiscent of a dipole antenna.

To obtain Figure 33, one must specify sonar orientations as well as sonar positions; these are indicated by the short green lines at the sonar locations; no attempt has been made to optimise these, since the purpose of Figure 33 is merely to indicate how azimuth dependence can be treated.




Figure 33: Contours of P_{ti} for an array of five sonars each comprising linear hydrophone arrays oriented as shown by the short green lines: panels on the left – single-sonar P_{ti} ; panels on the right – corresponding networked P_{ti} . In all plots, red shows regions where $P_{ti} \geq 95\%$ and blue where $P_{ti} \geq 80\%$. The exponential form is used for the range dependence of P_d and the angular dependence is given by a simple sine function (Eq. 14).

4.6 Schematic P_d Characteristic which is the Same for all Sonars

By ‘schematic’, we mean that the P_d shapes chosen are not derived from the sonar equation or other more detailed environmental modelling. This assumption is also not a requirement of the model— P_d characteristics derived from the sonar equation (e.g. Fig. 12.10 in [10]), or in any other manner, could be used if desired. Similarly, the assumption that all sonar systems in the network behave the same was made for the sake of simplicity and definiteness and is not required by the structure of the model. If desired, each sonar in the array can have a different characteristic; all that is needed is to specify them.

Figure 34 shows an example of these points: Figure 34(a) shows two range dependences that are similar to curves calculated by acoustic-propagation modelling of real sonars in specified environments and Figure 34(b) shows a coverage plot for a network of three sonars, two with one P_d characteristic and one with the other.^(g)

4.7 Other Assumptions

Static model; zero speed of advance

This assumption is necessary to give a simple analytical model and is fit for purpose in a first quick look at the problem. It is appropriate if the sonar field comprises sonobuoys, or if the task is choke-point monitoring, but not in the case of monitoring ahead of a transiting fleet. Our present view is that analysis of the transiting-fleet scenario cannot be done simply, but rather requires a simulation. Nevertheless, we believe that such a simulation would show a networking advantage analogous to that displayed by the current analysis for the static scenario.




Figure 34: (a) Two typical detection-probability-range curves from acoustic-propagation modelling; (b) resulting contours of track-initiation probability (red: $\geq 95\%$, blue: $\geq 80\%$) for a network of three sonars, two with the green P_d curve and one with the black. The distance scales are suppressed for reasons of classification.

^(g)A technical problem arises with the method for calculating coverage area described at the end of §2.3: the MATLAB function `polyarea` has difficulties when the contour is very irregular, as in Figure 34(b). This is circumvented by making use of the discretisation of the calculation. Since contours are determined by filling the area with a lattice of points, with P_{ti} computed at every point, the area inside a contour can be determined by counting the number of lattice points at which P_{ti} exceeds the desired contour level.

Multiple monostatic active sonar

This assumption was made with a view to alignment with current operating practice as much as possible. Multiple monostatic is the way in which active sonar is currently performed in the RAN and RAAF. The extension to passive sonar is simple: P_d is interpreted as detection probability per integration period rather than P_d per ensonification. The rest of the analysis then follows unchanged. Extension to multistatic active is more involved, for then P_d becomes a function of two distances, but this is in principle possible within the structure of the model.

Sonars do not interfere with each other

Sonobuoy arrays currently achieve this by appropriate scheduling of ensonifications, so the assumption is not an in-principle limitation.

Sufficient communications capacity exists to share detections

Similarly, the example of current practice with sonobuoys shows that communications capacity cannot be an in-principle limitation. A maritime patrol aircraft (MPA) monitoring a BARRA sonobuoy field receives essentially raw sonar data, that is, data with a much lower level of processing than detection-level information. Hence, the communications requirements for a sonobuoy field are more demanding than would be required by the CONOPS proposed here. Nevertheless, communication capacity is an important consideration in net-centric operations; passage of tactical-level sonar data will not be the only demand on communications bearers. Future studies should be carried out to examine the additional impact of the data traffic engendered by the CONOPS envisaged herein.

Flat 2-dimensional geometry

We ignore effects of the curvature of the Earth and of target depth. For the overall system dimensions envisaged here, we expect curvature-of-the-Earth effects to be very minor. As to submarine depth, the probability-of-detection approach that we adopt includes its effect in principle. That is, one could adopt a different P_d characteristic for each submarine depth of interest, in which case coverage-area plots would vary with target depth. A more explicit description would require detailed sound-propagation modelling. This is possible, and brings the potential to include environmental data, but at the expense of a very substantial increase in model complexity.

5. Summary and Conclusions

This report presents a quantitative investigation of a mechanism for networking advantage in anti-submarine warfare (ASW). The mechanism consists of sharing pre-track-level data at the level of single detections, with a centralised tracking processor for the overall system, as opposed to each sonar system performing tracking on its own detections only and then sharing the track information. We focus on the step of track initiation, in particular the probability P_{ti} that a track is initiated after 5 ensonifications from each sonar. As a metric for the magnitude of the networking benefit, we mainly use the area enclosed by the 95% contour of P_{ti} ('coverage area'), but we also consider

other metrics with a view to exploring how their use may change the overall conclusions.

The results depend on the behaviour of detection probability P_d for each individual sonar system. If P_d rolls off slowly with range, then there is considerable advantage in sharing detections and performing tracking centrally (Figs 6–8); for 5 sonars each with exponentially shaped P_d curves, the coverage area is over 8 times larger with centralised tracking than with individual tracking. This gain in coverage area is obtained without any changes to individual sonar performance: all steps up to the recognition of a contact are assumed to be the same with both concepts of operation; the difference lies in how detections, once made, are handled.^(h)

In the above example, it is the low-probability tail of the exponential shape that creates the opportunity for networking advantage, through the combination of low-probability regions from several sensors. If the P_d curve has a sharp cut-off, with values close to 100% at short range, then the gain in coverage area is smaller (Figs 10 & 11), reducing to zero for a definite-range law. This result, however, depends on the P_d inside the detection range being 100%. If the value is lower, then the possibility re-emerges of gain in coverage area from sharing information on detections. For example, if $P_d = 80\%$ inside the detection range, then the single-sonar P_{ti} does not reach 95% anywhere (assuming that track initiation requires 3 detections in 5 consecutive ensonifications), but geometrical arrangements of several sonars can be found for which the networked P_{ti} exceeds 95% over a sizeable area (e.g. Fig. 12).

An example is also presented of using pre-track-level networking to cover sonar blind zones (Fig. 14).

Essentially similar results are obtained with other metrics of networking benefit, such as the width of the networked coverage area or the 'depth of defence' – the largest minimum distance from an HVU location to the boundary of the networked coverage area. These lineal metrics tend to assess cookie-cutter-like P_d characteristics as closer in performance to exponential detection probabilities than does networked coverage area, but the general conclusion – that a substantial networking benefit is available – is still obtained.

An important consequence of the existence of networking benefit can be succinctly stated as:

A 30% probability of detection can be very useful in ASW; it can be made so by sharing detection-level data with neighbouring sonars that have similar P_d values for the target concerned.

This finding would seem to provide a practical way around the great and continuing difficulty in obtaining acceptably high P_d values at tactically useful distances from a single sonar. Further, it should be possible to obtain the improved performance with only relatively minor alterations to existing sonar systems.

A disadvantage to networking lies in the sensitivity of the metrics examined in this report to the design of the ASW screen for the task group. NCW can enhance operational flexibility and capability, but the price is additional design complexity.

^(h)Although this assumption may seem benign, there could be an issue with false-alarm rate if the centralised tracking is performed manually and the operator is unable to tolerate the increased rate of false alarms arising from the fusing of detections from many sensors. If the operator responds by raising the detection thresholds of the sensors, then the networking gain is attenuated by an amount that depends on the magnitude of the change to the detection threshold and the initial false-alarm rate (§4.1, Appendix B.1).

Figures 9 and 11 illustrate how the wrong choice of inter-sonar spacing can negate networking benefit. Figures 31, 33 and 34 show some of the complexities that may be involved in designing ASW screens using real systems. And this is before the consideration of multistatic effects. These examples make clear the severity of the performance decrements that can arise with a poor choice of CONOPS for the deployment of the sonar network.

The concept of operations analysed herein—the sharing of detection-level information—is not new; there are previous theoretical analyses of it [1–5] and it is similar to the manner in which a maritime patrol aircraft operates when monitoring a field of sonobuoys. Because of power constraints, sonobuoys perform minimal on-board processing and stream essentially raw sonar data to the aircraft. On the aircraft, sensor operators process the data from each sonobuoy, deciding on detections, which are passed to the tactical commander (TACCO) who makes decisions on initiating tracking. The TACCO acts as a centralised tracking operator, with the additional feature that he or she is aggregating detections not only from multiple sonars, but also from other sensors on the aircraft. It would be an interesting further study to explore the extent to which the workload of both the TACCO and the aircraft as a whole may be impacted by a higher level of automation of sensor networking, along the lines suggested in this report.

Further work on the benefits of sharing detection-level data could involve the exploration of those assumptions used herein (§4) that can be relaxed with relatively simple additions to the model. Of these, perhaps the most promising, from the point of view of gaining further networking advantage, may be the assumption of multiple monostatic operation. It seems very likely that further networking benefit could be gained by multistatic processing, where each sonar processes returns from any ensonification, regardless of its source.

Much of the existing analysis concerning network-centric warfare (NCW) has been concerned with the operational and strategic warfare levels, with a focus on sharing and fusing tactical pictures from a variety of platforms, and providing access to reach-back information. On this view, NCW is about building and disseminating a common relevant operating picture. Quantitative analysis of the higher levels of warfare is difficult, which may explain why there are so few studies exploring the benefits of NCW quantitatively [18].⁽ⁱ⁾ The present work provides a quantitative indication that NCW should be able to produce benefits at the tactical level.

(i) Ref. 18 is now over 3 years old (dated November 2004). Since then, the only other quantitative studies published, to our knowledge, are contained in four RAND reports [9,19–21] and a handful of more abstract studies [22–26].

Appendix A: Overall Area of Overlapping Circles

Evaluation of coverage-area ratios (Eq. 7, p. 7) requires determining the overall area of single-sonar contours. These overlap for small enough inter-sonar spacing, in which case one must avoid counting the overlap areas twice. This appendix gives details of how this was done for the construction of Figure 11 (p. 10). The assumptions about the single-sonar P_d curves described in §§4.4–4.6 mean that

- ♦ all single-sonar coverage areas are circles,
- ♦ the circles all have the same radius r , and
- ♦ the centres of pairs of neighbouring circles are all separated by the same distance s .

When $s > 2r$, the circles do not overlap and hence the overall area A_t equals NA_o , where N is the number of circles and $A_o = \pi r^2$ is the area of one circle. For $s < 2r$, circles intersect and the quantity

$$y_i = \frac{1}{2} \sqrt{4r^2 - s^2} \quad (15)$$

is the distance of an intersection point of two circles from the line joining their centres, as Figure 35 illustrates. If

$$\frac{s\sqrt{3}}{2} - y_i > r \Rightarrow s > r\sqrt{3} \quad (16)$$

(the case shown in Fig. 35a), then no point is inside more than two circles. If not, then there are regions where three or more circles overlap (Fig. 35b). Expressions for the area of common overlap of three circles are available [27], so we are able to calculate this case. We have not attempted to treat situations where four or more circles have an area of common overlap; this is why the abscissa of Figure 11 has a suppressed zero.

For the case shown in Figure 35(a), the overlap areas are lens shapes, each of which has an area

$$A_l = 2r^2 \arccos \frac{s}{2r} - \frac{s}{2} \sqrt{4r^2 - s^2}. \quad (17)$$

Inspection of a series of diagrams like Figure 35(a) for differing numbers of circles shows that the overall area A_t is

$$A_t = \begin{cases} NA_o - (2N-3)A_l & 2 \leq N < 7 \\ 7A_o - 12A_l & N = 7, \end{cases} \quad (18)$$




Figure 35: Three equal, equi-spaced circles that (a) satisfy and (b) do not satisfy Equation (16).

since, for the arrangement of circles adopted herein (Fig. 6, p. 8), each additional circle from the third to the sixth adds two lenticular areas, but the addition of the seventh circle adds three.

For circles spaced more closely, so that Equation (16) does not hold, as in Figure 35(b), there are regions inside three circles (provided that $N \geq 3$). The area of each resulting equilateral circular triangle can be expressed in terms of the chord length c between vertices of the circular triangle [27], which is given by

$$c^2 = 3r^2 - s^2/2 - s\sqrt{3r^2 - 3s^2/4}. \quad (19)$$

The area A_Δ of the circular triangle is [27]

$$A_\Delta = \frac{\sqrt{3}}{4}c^2 + 3\left(r^2 \arcsin \frac{c}{2r} - \frac{c}{4}\sqrt{4r^2 - c^2}\right). \quad (20)$$

Now each additional circle up to the seventh adds two lenses and a circular triangle, which is the region of overlap of lenses. Hence, for every overlapping pair of lenticular areas subtracted, a circular-triangle area must be added back, giving the overall area as

$$A_t = \begin{cases} NA_o - (2N - 3)A_l + (N - 2)A_\Delta & 3 \leq N < 7 \\ 7A_o - 12A_l + 6A_\Delta & N = 7. \end{cases} \quad (21)$$

As with Equation (18), the last line of Equation (21) arises because of the symmetry of the geometrical arrangement (Fig. 6e) makes it a special case: adding the seventh circle adds three lenses and two circular triangles.

To reiterate: these expressions apply only if the circles are not so closely spaced that there are points simultaneously inside four of them, which occurs when $s < 2r/\sqrt{3}$.

Appendix B: Mathematical and Conceptual Details on Probability of False Alarm

B.1 Method of Computing Figures 22 and 23

The aim is to develop a method for quantifying the effect of changing the detection threshold (Figs 22 & 23, pp. 19 & 20) on coverage-area ratio. The next subsection describes the concept of the method by way of a schematic sample calculation. Section B.1.2 rewrites this mathematically.

B.1.1 Concept

Figure 36 shows graphically a sample calculation.⁽ⁱ⁾ The method is as follows:

(i) The construction of Fig. 36 involves the following assumptions:

- The distribution of sensor outputs arising from noise is Gaussian.
- Signals from sonar returns are deterministic; that is, neither the width nor the shape of the sensor-output distribution is changed by the addition of the sonar return.

Neither of these assumptions is essential; any shape of sensor-output distribution can be used. In particular, it is known that fluctuations in phase as well as amplitude mean that the




Figure 36: Graphical depiction of the calculation of the effect on P_d of an adjustment in threshold. Both axes use normal-probability coordinates, so that the ROC curves (broken lines) are straight. The two vertical lines represent two threshold values that correspond to false-alarm rates a factor of 5 apart. The scale on the line at $P_{fa} = 0.5\%$ show range values for the Fermi-function P_d behaviour of Figure 3(b). The arrows show how the range values translate with the change in threshold from that with $P_{fa} = 0.5\%$ to that with $P_{fa} = 0.1\%$.

- One first decides on the P_{fa} value that an operator would choose in the platform-centric case. In Figure 36, we picked 0.5% .^(k)
- Choosing a P_d curve then allows a range scale to be placed on the diagram at the selected value of P_{fa} . The range scale shown in Figure 36 corresponds to the Fermi-function P_d characteristic (Figs 1 & 3b). The exponential P_d characteristic gives a different range scale.
- Next, one needs a rule for deciding what false-alarm rate the operator will tolerate in the networked situation. We suppose that the operator seeks to keep the total false-alarm rate constant as sonar systems are added to the network, and that this is done

noise distribution of Fig. 21(a) (p. 19) is actually Rayleigh and the distributions with target returns (Figs 21b,c) are Ricean. There is no in-principal difficulty in using these more realistic distributions.

^(k)The concept of 'false-alarm probability' is more complicated than this description pretends. Details of a more realistic view are discussed in Appendix B.2.

by raising the detection thresholds of all sonar systems equally. Since false-alarm rate is proportional to P_{fa} , this means dividing the single-sonar P_{fa} value by the number m of sonars in the network. Figure 36 shows the situation for a 5-sonar network, for which the single-sonar P_{fa} needs to be reduced to 0.1%.

- All points for a given target range lie on one ROC curve. Hence, the P_d value for a given range after threshold adjustment is found by following the appropriate ROC curve to the new P_{fa} value. For example, Figure 36 shows $P_d = 34\%$ at 10 km before threshold adjustment. The broken line near 10 km in Figure 36 indicates that, after adjustment, $P_d \sim 18\%$ at this target range. By carrying out this procedure for all target ranges, one effectively rescales the abscissa of the P_d characteristic (e.g. Figs 1 & 3).
- Coverage-area ratios are then calculated using the rescaled ranges.

B.1.2 Mathematical Detail

The translation of Figure 21(a) (p. 19) into mathematics is (e.g. Fig. 12.8 in [10])

$$P_{fa} = G(T/\sigma), \quad (22)$$

where $G(x)$ is the complement of the usual cumulative standard normal probability distribution:

$$G(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-t^2/2} dt, \quad (23)$$

T is the value of the detection threshold and σ is the standard deviation of the noise. Similarly, the lower panels in Figure 21 result in

$$P_d = G\left(\frac{T}{\sigma} - \sqrt{d}\right), \quad (24)$$

where d is the ‘detection index’ (§12.2 in [10]). A ROC curve is characterised by constant d ; the broken lines in Figure 36 correspond to, from the bottom, $d = 0, 1, 5, 10, 20, 30$. Equations (22) and (24) can be formally solved for T/σ and $\sqrt{d}$ by invoking the inverse function of G ; that is, if $P = G(x)$, then $x = \text{inv}G(P)$. The results written in terms of two P_d, P_{fa} pairs can be rearranged to eliminate T/σ and $\sqrt{d}$, giving

$$P_{d2} = G\left(\text{inv}G(P_{fa2}) - \text{inv}G(P_{fa1}) + \text{inv}G(P_{d1})\right), \quad (25)$$

where P_{d1}, P_{fa1} are values before threshold adjustment and P_{d2}, P_{fa2} are values after.

Evaluation of Equation (25) was implemented in MATLAB. MATLAB has functions for the cumulative standard normal distribution and its inverse, but these are contained in the ‘statistics toolbox’, which can be subject to restrictive licensing controls. To avoid unnecessarily occupying a statistics-toolbox licence, we used ‘erf’ and ‘erfinv’ instead. The transformation is

$$G(x) = \frac{1}{2} - \frac{1}{2} \text{erf}\left(\frac{x}{\sqrt{2}}\right), \quad \text{inv}G(P) = \sqrt{2} \text{erfinv}(1 - 2P). \quad (26)$$

Figures 22 and 23 are constructed assuming $P_{fa1} = 0.5\%$ and $P_{fa2} = P_{fa1}/5$. If the scale on the abscissa of Figure 36 were logarithmic, then the results in Figures 22 and 23 would depend on the two false-alarm probabilities only through the ratio P_{fa1}/P_{fa2} , but this is not the case, rather the scale is normal-probability. Hence the results depend on P_{fa1} by itself, in addition to the dependence through the ratio of

false-alarm probabilities.^(l) Figure 37 shows the magnitude of the dependence in the case of 5 sonars ($P_{fa1}/P_{fa2} = 5$) with a P_d characteristic of exponential form (Fig. 3a). The dependence on the assumed value of P_{fa1} is significant, with networking advantage reducing to zero (i.e. coverage-area ratio = 1) at $P_{fa1} \approx 2\%$. The next subsection considers the question of realistic values of P_{fa1} .

B.2 Practical Values of P_{fa}

False-alarm probability was not explicitly defined in §4.1.2, with the intention that the reader would interpret the term by analogy with detection probability, that is, as a probability per ensonification. Although this may be sufficient for qualitative discussion, the real situation is more complicated because, unlike a genuine detection, there is no physical location correlated with the false alarm. The P_{fa} value quoted for active sonar systems is actually a false-alarm probability per ensonification *per range-bearing-Doppler bin*. That is, the display of the sonar processing unit is divided into 'bins'. The bin size in the range direction may be 2–150 m, depending on the transmitted waveform. In bearing, the bin size is given by the beamwidth, typically 5–10°. If a continuous-wave (cw) waveform is used, then there will also be a subdivision into increments of Doppler shift. Altogether, modern sonar processing units have of order 1 million bins over the whole display. Since typical P_{fa} values may be 10^{-5} per ensonification per bin, it is clear that there will be many false alarms every ensonification. Hence, if interpreted literally, the real P_{fa} per ensonification is closer to 100% than the 0.5% used in Figures 22 and 23.^(m)

However, this is where track-initiation rules enter the argument. Although many bins may show detections every ensonification, false alarms are not correlated with real-world objects, so the chance that any one bin will experience a false alarm in




Figure 37: Variation of networked coverage-area ratio for 5 sonars and an exponential P_d with assumed value of P_{fa1} for $P_{fa2} = P_{fa1}/5$.

^(l)That is, the normal-probability scale is required to turn the ROC curves into straight lines; and straight-line ROC curves is the second requirement for the results to depend on P_{fa1} only through the ratio of the two false-alarm probabilities (the first being a logarithmic axis scale).

^(m)Explicitly, the probability of *no* false alarms in *any* bin per ensonification is $(1 - P_{fa})^N$, which equals 0.0045% for $P_{fa} = 10^{-5}$ and $N = 10^6$. Since the occurrence of false alarms is equivalent to a binomial distribution for N trials, the mean number of bins showing a false alarm in each ensonification is $P_{fa}N = 10$.

successive ensonifications, or 3 false alarms in 5 consecutive ensonifications, is very small. This suggests that, rather than false-alarm probability per ensonification, the quantity of actual interest is the rate at which tracks are initiated solely from false alarms. We focus on the 3-in-5 track-initiation rule in this report. In the context of false alarms, this might be appropriate for automatic track initiation, but an operator is unlikely to be able to keep track of detections over 5 ensonifications, particularly if some tens of them occur every ensonification, so the 2-in-2 rule may be the more reasonable description of the effect of false alarms on manual track initiation.

We denote by $P_{fa,2|2}$ the probability that at least one bin has a pair of false alarms in successive ensonifications. The usual rules for combining probabilities give

$$P_{fa,2|2} = 1 - \left(1 - P_{fa}^2\right)^N, \quad (27)$$

where N is the total number of bins across the display. The equivalent quantity for 3 or more false alarms in one bin in 5 ensonifications is (cf. footnote c on p. 4)

$$P_{fa,3|5} = 1 - \left[1 - P_{fa}^3(10 - 15P_{fa} + 6P_{fa}^2)\right]^N. \quad (28)$$

For $P_{fa} = 10^{-5}$ per ensonification per bin and $N = 10^6$ bins these quantities are $P_{fa,2|2} = 0.01\%$ and $P_{fa,3|5} = 10^{-6}\%$. From this point of view, the P_{fa} of 0.5% used in Figures 22, 23 and 36 seems too large.

There is a further consideration that mitigates the last tentative conclusion. Successive detections of a real object will not usually occur in the same bin, owing either to relative motion between target and sonar, or to measurement uncertainty. That is, for real detections, the track-initiation rule must be applied to a cluster of bins, rather than to a single bin. The size of the cluster depends on the time interval between ensonifications and the maximum speed, relative to the sonar system, expected for targets of interest, or may be correlated with Doppler shift if this is measured. Where the maximum relative speed is small, measurement uncertainty becomes a consideration. Hence cluster size varies widely with the tactical situation, but a ‘typical’ value may be of order 100.

The generalisation of Equation (27) to the case where the second detection can occur anywhere within a cluster of C bins is (see below)

$$P_{fa,2|2} = 1 - \left[1 - P_{fa} + P_{fa}(1 - P_{fa})^C\right]^N \quad (29)$$

For $C = 100$ bins, $N = 10^6$ bins and $P_{fa} = 10^{-5}$ per ensonification per bin, Equation (29) evaluates to $P_{fa,2|2} = 0.99\%$. The corresponding value of $P_{fa,3|5}$ would be smaller. The identification of these quantities with the ‘false-alarm probability’ used in §4.1.2 is the justification for choosing 0.5% for the construction of Figures 22 and 23.

Derivation of Equation (29)

Let $\bar{P}_{ft,i}$ be the probability of not forming a false track in a particular bin i after two successive ensonifications. Since the track-initiation rule in question is 2-in-2,

$$\bar{P}_{ft,i} = \bar{P}(\text{false alarm in bin } i \text{ for ping 1}) + P(\text{false alarm in bin } i \text{ for ping 1}) \times \bar{P}(\text{false alarm in } C \text{ bins around bin } i \text{ for ping 2}) \quad (30)$$

where $\bar{P}$ denotes the complement of the described probability. We assume that the false-alarm probability P_{fa} is the same in all bins for each ensonification, and so

$$\bar{P}_{ft,i} = (1 - P_{fa}) + P_{fa} (1 - P_{fa})^C \quad (31)$$

Since all bins are equivalent as regards the occurrence of false alarms, the probability of no false track in any of N bins is given by Equation (31) raised to the N th power. Equation (29) follows directly from this.

Appendix C: Mathematical and Computational Details for §4.3

C.1 Derivation of Equations (10) – (13)

Equation (10) (p. 24), which is a recurrence relation for the probability $\bar{P}_{ti,2|3}(Q)$ of not initiating a track after Q ensonifications using the 2-in-3 rule, can be derived as follows. In the tree diagram (Fig. 26, p. 23), branches terminate when a track is initiated. For any ensonification ('ping') in the tree numbered 4 or higher, if a detection is made with that ping but a track is not initiated, then the previous 2 pings cannot have resulted in detections. So, if a ping does not result in a track being initiated, then either that ping did not result in a detection or the previous two pings did not. Instances of this behaviour are enclosed in ovals in Figure 26. Thus, if we represent the result of a sequence of pings using 0s and 1s, with '1' for a detection and '0' for a miss, then any sequence in the tree diagram where a track is not initiated ends either with a zero or with the string '001'. An unsuccessful sequence of pings of length Q (i.e. a sequence in which a track is not initiated) can therefore be constructed from a shorter unsuccessful sequence in one of two ways:

- (i) by adding 0 to the end of an unsuccessful sequence of length $Q - 1$, or
- (ii) by adding 001 to the end of an unsuccessful sequence of length $Q - 3$.

These two points provide the key for writing down the recurrence relation. Writing P_d for the detection probability, the contribution to $\bar{P}_{ti,2|3}(Q)$ from the mechanism described in point (i) is

$$(1 - P_d) \bar{P}_{ti,2|3}(Q - 1), \quad (32)$$

since $\bar{P}_{ti,2|3}(Q - 1)$ is the probability of obtaining an unsuccessful sequence of length $Q - 1$ and $1 - P_d$ is the probability of adding a zero to it. Similarly, point (ii) contributes

$$(1 - P_d)^2 P_d \bar{P}_{ti,2|3}(Q - 3). \quad (33)$$

These two mechanisms are both exhaustive and mutually exclusive—either the Q th ping produces a detection or it does not—so $\bar{P}_{ti,2|3}(Q)$ is obtained as their sum.

Equation (11) (p. 24) gives the initial values for commencing the recurrence relation in the case of the 2-in-3 rule. The first two are obvious—if we require 2 detections in 3 pings to initiate a track, then a track cannot be initiated with fewer than 2 pings. Hence $\bar{P}_{ti,2|3}(0) = \bar{P}_{ti,2|3}(1) = 1$. The third initial value can be seen from the tree diagram. If the first two pings both result in detections, then we do not need a third. The probability of this occurring is P_d^2 , so $\bar{P}_{ti,2|3}(2) = 1 - P_d^2$.

This derivation assumes, as we have throughout this report, that P_d is the same for all pings, an assumption that is consistent with the static scenario. However, it is not an essential assumption. If P_d were to vary from ping to ping, then one would write $P_d(Q)$

in the above derivation, taking care to select the correct argument. For example, the contribution from point (ii) above would become

$$[1 - P_d(Q-3)][1 - P_d(Q-2)]P_d(Q-1)\bar{P}_{ti,2|3}(Q-3). \quad (34)$$

Equation (12) (p. 24), which is the recurrence relation for the 2-in-2 track-initiation rule, is derived similarly. That is, the probability that a track will *not* be initiated with the Q th ping is equal to the sum of

- the probability of not obtaining a detection with the Q th ping, given that we had not previously initiated a track, which is $(1 - P_d)\bar{P}_{ti,2|2}(Q-1)$, and
- the probability of obtaining a detection on the previous ping but not on the ping immediately prior (remembering that 2 consecutive detections are required), given that a track had not been initiated with the prior sequence of $Q-2$ pings, which is $P_d(1 - P_d)\bar{P}_{ti,2|2}(Q-2)$.

Equation (13) gives the initial values for the 2-in-2 case, which are obvious from the nature of the rule.

C.2 Simulation for the Calculation of P_{ti} per Q Ensonifications

C.2.1 MATLAB Code

The Matlab code for calculating P_{ti} by simulation is listed in Table 2 (next page). Perhaps the simplest way to understand the algorithm is via a sample calculation. The following notes, which are intended to be read in parallel with the code, present a calculation for a 3-in-5 rule applied to 8 pings ($p = 3, q = 5, Q = 8$).

- Assume that a large number is specified for iterations, so that the returned value is (an approximation to) P_{ti} .

- For the parameters above, $V = \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}$ and, after the first for loop, $Test = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{pmatrix}$.

- Consider the first iteration of the main (second) for loop. Assume that the values of P_d and probs are such that, this time round, $detected = (0 \ 1 \ 0 \ 0 \ 1 \ 0 \ 1 \ 1)$. Then

$$Initiation_Intermediate = (0 \ 1 \ 0 \ 0 \ 1 \ 0 \ 1 \ 1) \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{pmatrix} = (2 \ 2 \ 2 \ 3),$$

Table 2: MATLAB code for calculation of track-initiation probability

```

function Pti = track_initiation(Q, p, q, Pd, Iterations)

% Probabilities of track initiation. If a single iteration, function returns
% a one or a zero. Otherwise the Pti probability averaged over the number
% of iterations is returned.

% Problem: What is the probability a track is initiated within Q pings
%          using a p | q track initiation criterion?

% Input Variables:
% Q — Number of Pings
% p — Minimum number of detections required
% q — Window length for track initiation
% Pd — Detection Probability
% Iterations — Number of iterations (default = 1).

% David Kershaw

if nargin < 5
    Iterations = 1;
end

% Create the Test Matrix for the p | q detection
V = ones(q,1);
Test = zeros(Q,Q-q+1);
for i = 1:(Q-q+1)
    Test(i:(i+q-1),i) = V;
end

% Now set up the iteration loop
Initiated = zeros(1,Iterations);
for i = 1:Iterations

    % Create the probability set
    probs = rand(1,Q);

    % Check if they have been detected (ie: random number <= Pd)
    detected = probs <= Pd;

    % Now conduct the p | q test. Multiplying detected by Test will provide
    % a row vector with each entry being the number of detections in a
    % given q-length window within the Q pings. The final test is then to
    % check to see if any of the entries are > p
    Initiation_Intermediate = detected * Test;
    Initiated_Element = sum(Initiation_Intermediate >= p);
    Initiated(1,i) = Initiated_Element > 0;
end

% Final step is to take the mean of the Initiated vector.
Pti = mean(Initiated);

```

and hence $\text{Initiated_Element} = 1$ (since only one element of $\text{Initiation_Intermediate}$ is greater than the required value given by $p = 3$), and $\text{Initiated}(1,1) = 1$.

- The main for loop is then iterated to obtain values of $\text{Initiated}(1,i)$, which comprise a sequence of ones or zeros.
- Finally, P_{ti} is computed as the mean value of $\text{Initiated}(1,i)$.

C.2.2 Convergence

The accuracy of the calculation of P_{ti} clearly depends on the number of iterations used in the simulation. This was explored in two ways:

- by choosing several numbers of iterations in the range 1000– 10^6 and running the code 100 times at each selected value, and
- by comparing the results of the code with one of the cases for which we have algebraic results.

The results of the first test are summarised in Figures 38 and 39. Figure 38 shows a scatter diagram for the 100 runs at each number-of-iterations value, using the 3-in-5 track-initiation rule. For a given number of iterations, the greatest spread in returned P_{ti} values occurs when $P_d = 0.5$. As the number of iterations is increased, the variability in the results is reduced. This is quantified in Figure 39, which shows the standard deviations of the data in Figure 38. All of the data presented in this report that were computed with this code (i.e. Figs 27–30) were generated using 1 million iterations for each result.

The second test involved comparing results from the simulation code with the formulae available for the 2-in-3 rule (Eqs 10 & 11, p. 24). The code was run 100 times each for a range of P_d values, using $Q = 3$ and 10^6 iterations in each run, and the differences between the values returned by the code and the results of the recurrence relation were tabulated. Figure 40 shows these differences, categorised by P_{ti} value. We also computed the standard deviation for each sample of 100 data, with the results shown by




Figure 38: Scatter plots showing, for each number of iterations, the results of 100 runs of the simulation code to calculate P_{ti} using P_d values of (a) 0.9, (b) 0.5, (c) 0.1 and a track-initiation rule of 3-in 5 out of 5 ensonifications.




Figure 39: Standard deviations of each bundle of results in the scatter diagram of Figure 38.

the lines in Figure 40. Of the 1100 data shown in Figure 40, the numbers that lie more than one, two and three standard deviations from the respective mean P_{ti} values are listed in Table 3 and compared with expected numbers for a sample of 1100 data drawn from a normal distribution. We conclude that the data are close to normally distributed. Examination of normal-probability plots for each data set (i.e. each of the 100 data at each P_{ti} value considered separately) support this. We then carried out t tests on each data set, with the result that, for each data set, the mean difference




Figure 40: Modulus of differences between analytic (i.e. exact) results and the P_{ti} values returned by the simulation with 10^6 iterations. The simulation was run 100 times at each chosen P_d value; results are shown sorted by P_{ti} . The full lines show (green) one, (blue) two and (red) three standard deviations of the samples of simulation P_{ti} values.

Table 3: Statistical behaviour of the data in Figure 40. Of the total of 1100 data points, entries show the number found to lie more than one, two or three standard deviations from the mean (i.e. above the green, blue and red lines respectively in Figure 40) compared with the number expected for an exactly normally distributed sample.

	observed	expected
σ	360	349
2σ	43	50
3σ	4	3

between the simulation result and the exact value of P_{ti} is statistically indistinguishable from zero.

References

- [1] J.A. Roecker & C.D. McGillem (1988) 'Comparison of two-sensor tracking methods based on state vector fusion and measurement fusion', *IEEE Trans. Aero. Elec. Sys.* **24**, 447–9.
- [2] A. Ashraf Mambouh (1997) 'Multiple-sensor distributed detection systems with data fusion', *Proc. 13th Internat. Conf. Digital Signal Processing*, 1031–4.
- [3] H. Chen. T. Kirubarajan & Y. Bar-Shalom (2000) 'Comparison of centralised and distributed tracking algorithms using air to air scenarios', *Proc. SPIE* **4048**, 440–451.
- [4] H. Chen. T. Kirubarajan & Y. Bar-Shalom (2003) 'Performance limits of track-to-track fusion versus centralised estimation: theory and application', *IEEE Trans. Aero. Elec. Sys.* **39**, 386–400.
- [5] S. Coraluppi (2005) 'Analysis of tracker performance models for centralized and distributed tracking', *Proc. 7th Internat. Conf. on Information Fusion*, pp. 1404–11.
- [6] D.S. Alberts, J.J. Gartska & F.P. Stein (1999) *Network Centric Warfare: Developing and Leveraging Information Superiority*, Vienna VA, CCRP Publications.
- [7] J.R. Cares, R.J. Christian & R.C. Manke (2002) 'Fundamentals of Distributed, Networked Military Forces and the Engineering of Distributed Systems', Technical report 11366 of the Naval Undersea Warfare Center Division Newport.
- [8] J. Thackray (2003) 'The Holy Grail' in D. Potts (ed.) *The Big Issue: Command and Control in the Information Age*, Vienna VA, CCRP Publications.
- [9] I.R. Porche III & B. Wilson (2006) *The Impact of Network Performance on Warfighter Effectiveness*, technical report of RAND Arroyo Center, available at <www.rand.org/pubs/>.
- [10] R.J. Urick (1983) *Principles of Underwater Sound*, 3rd edn, N.Y., McGraw Hill.
- [11] P.H. Dahl, J.H. Miller, D.H. Cato & R.K. Andrew (2007) 'Underwater ambient noise', *Acoustics Today* **3**, 23–33.
- [12] Z. Chair & P.K. Varshney (1986) 'Optimal data fusion in multiple sensor detection systems', *IEEE Trans. Aero. Elec. Sys.* **AES-22**, 98–101.
- [13] S.C.A. Thomopoulos, R. Viswanathan & D.C. Bougoulas (1987) 'Optimal decision fusion in multiple sensor systems', *IEEE Trans. Aero. Elec. Sys.* **AES-23**, 644–53.
- [14] D.H. Wagner, W.C. Mylander & T.J. Sanders (eds) (1999) *Naval Operations Analysis*, 3rd edn, Annapolis MD, Naval Institute Press.
- [15] R.C. Gauss, J.M. Fialkowski & R. Menis (2003) 'Time spread and signal coherence measurements during T-MAST 02', *Proc. 32nd TTCP MAR TP-9 Meeting*, DSTO, Sept 2003.
- [16] F.R. Castella (1976) 'Sliding window detection probabilities', *IEEE Trans. Aero. Elec. Sys.* **AES-12**, 815–9.
- [17] P. Williams (1998) *Evaluating the state probabilities of M out of N sliding window detectors*, Technical note DSTO-TN-0722 of the Defence Science and Technology Organisation.
- [18] S. Dickinson, M.P. Fewell, P. Sutherland & C. Davis (2004) *Bibliography of Quantitative Analyses of Network-Centric Warfare*, Document DOC-MAR-9-2004 of The Technical Co-operation Program.
- [19] D. Gonzales, M. Johnson, J. McEver, D. Leedom, G. Kingston & M. Tseng (2005) *Network-Centric Operations Case Study: the Stryker Brigade Combat Team*, monograph of RAND National Defense Research Institute, available at <www.rand.org/pubs/>.

- [20] D. Gonzales, J. Hollywood, G. Kingston & D. Signori (2005) *Network-Centric Operations Case Study: Air-to-Air Combat with and without LINK 16*, monograph of RAND National Defense Research Institute, available at <www.rand.org/pubs/>.
- [21] C.G. Pernin & L.R. Moore (2005) *The Weapons Mix Problem: a Math Model to Quantify the Effects of Internetting of Fires to the Future Force*, technical report of RAND Arroyo Center, available at <www.rand.org/pubs/>.
- [22] M.F. Ling, T. Moon & E. Kruzins (2005) 'Proposed network centric warfare metrics: from connectivity to the OODA cycle', *Mil. Op. Res.* **10** (1), 5–13.
- [23] A. Kaufman (2005) 'Critical factors affecting the military utility of networking' *J. Battlefield Tech.* **8** (3), 27–32.
- [24] J.T. Buontempo, H.G. Potrykus & A.I. Kaufman (2006) 'A symmetric adaptive model of combat' *J. Math. Sociol.* **30** 113–36.
- [25] A. Kaufman (2006) 'Strategic consequences of networking', *J. Battlefield Tech.* **9** (1), 15–18.
- [26] J. Moffat (2007) 'Modelling human decision-making in simulation models of conflict', *Internat. C2 J.* **1** (1), 31–60.
- [27] M.P. Fewell (2006) *Area of Common Overlap of Three Circles*, Technical note DSTO-TN-0722 of the Defence Science and Technology Organisation.

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA							
				1. PRIVACY MARKING/CAVEAT (OF DOCUMENT) None			
2. TITLE Benefits of Sharing Detections for Networked Track Initiation in 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) M.P. Fewell, J.M. Thredgold and D.J. Kershaw				5. CORPORATE AUTHOR DSTO Defence Science and Technology Organisation PO Box 1500 Edinburgh South Australia 5111 Australia			
6a. DSTO NUMBER DSTO-TR-2086		6b. AR NUMBER AR 014-075		6c. TYPE OF REPORT Technical Report		7. DOCUMENT DATE January 2008	
8. FILE NUMBER 2007/1087654/1		9. TASK NUMBER NAV 05-097	10. TASK SPONSOR DGMD	11. NO. OF PAGES 45		12. NO. OF REFERENCES 27	
13. URL on the World Wide Web http://www.dsto.defence.gov.au/corporate/reports/DSTO-TR-2086.pdf				14. RELEASE AUTHORITY Chief, Maritime Operations Division			
15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT <i>Approved for public release</i> OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE, PO BOX 1500, EDINBURGH, SA 5111							
16. DELIBERATE ANNOUNCEMENT No limitations							
17. CITATION IN OTHER DOCUMENTS Yes							
18. DSTO RESEARCH LIBRARY THESAURUS http://web-sa.dsto.defence.gov.au/workareas/library/resources/dsto_thesaurus.htm Antisubmarine warfare, Sonar systems, Tracking, Sensor fusion, Data fusion, Network-centric warfare							
19. ABSTRACT This report presents a quantitative study of a class of networking benefits in anti-submarine warfare. We show that sharing detections can produce an advantage over sharing only track-level information. We also indicate the conditions under which the advantage should be present, and estimate the magnitude of the advantage. This is achieved by focusing on the step of centralised track initiation, using metrics based mainly on sonar coverage area. We analyse multiple monostatic sonar, to give a concrete example aligned with current practice. The conclusions may be summarised in the statement that a 30% detection probability can be tactically useful, provided that there are other sonars with a similar P_d for the target concerned that are sharing information on detections. This result may provide a practical way around the great and continuing difficulty in obtaining acceptably high P_d values at tactically useful distances from a single sonar.							