

UNCLASSIFIED



Australian Government
Department of Defence
Defence Science and
Technology Organisation

Evaluating De-centralised and Distributional Options for the Distributed Electronic Warfare Situation Awareness and Response Test Bed

Kuba Kabacinski, Catherine Howard and Damian Hall

Electronic Warfare and Radar Division
Defence Science and Technology Organisation

DSTO-TR-2918

ABSTRACT

The Distributed Electronic Warfare Situation Awareness and Response Test Bed was designed to rapidly develop and evaluate distributed Electronic Warfare concepts through a networked set of heterogeneous, relatively unsophisticated sensors and effectors (deployed on ground based or aerial platforms) to detect, identify, locate, track or suppress stationary or slow moving surface based RF emitting targets. In the current test bed, little of the control or data processing occurs at the ground based or airborne nodes. This document evaluates a number of options for de-centralising the control functions and distributing the data processing from the ground station to the ground based or airborne nodes.

RELEASE LIMITATION

Approved for public release

UNCLASSIFIED

UNCLASSIFIED

Published by

*Electronic Warfare and Radar Division
DSTO Defence Science and Technology Organisation
PO Box 1500
Edinburgh South Australia 5111 Australia*

*Telephone: 1300 333 362
Fax: (08) 7389 6567*

*© Commonwealth of Australia 2013
AR-015-804
December 2013*

APPROVED FOR PUBLIC RELEASE

UNCLASSIFIED

UNCLASSIFIED

Evaluating De-centralised and Distributional Options for the Distributed Electronic Warfare Situation Awareness and Response Test Bed

Executive Summary

The Distributed Electronic Warfare Situation Awareness and Response (DEWSAR) Test Bed (TB) is a test bed developed by the Defence Science and Technology Organisation's (DSTO) Electronic Warfare and Radar Division which is designed to enable the rapid development and evaluation of distributed Electronic Warfare concepts by utilising a networked set of heterogeneous, relatively unsophisticated sensors and effectors (deployed on ground based or aerial platforms) to detect, identify, locate, track or suppress stationary or slow moving surface based RF emitting targets. In the current test bed, little of the control or data processing occurs at the (ground based or airborne) nodes. For example, data from the sensors is transmitted from the node to the ground station where it is collected, processed, stored and disseminated as appropriate.

The DEWSAR TB's control functions and data processing steps can be broken down in to the sensor Tuning Function, the aerial platform Steering Function, and the target Classification, information Aggregation, target Geolocation, target Tracking and information Reporting processes. A number of options were evaluated for migrating these control functions and data processing steps from the ground station to the ground based or airborne nodes. The evaluation was performed within the context of the purpose of the test bed, the size, weight and power constraints of the Aerosonde aerial platforms and the communication bandwidth limitations.

It was found to not be possible to make a judgement about distributing the Classification Process without information about the taxonomy and density of the radar target environment, which in a real world deployment may not be available. Whether distributing the Reporting Process would be advantageous to the test bed depends on the goal of the deployment. De-centralising the test bed's Tuning and Steering functions or distributing the Aggregation, Geolocation and Tracking processes would substantially increase the complexity of the test bed. The increased complexity of the test bed may result in increased development time, cost and increased maintenance overheads. Higher complexity could also increase the test bed's technical risk by increasing the likelihood and complexity of future technical problems. Any de-centralisation or distribution option that impinges a researcher's ability to rapidly develop and evaluate algorithms or concepts contravenes the purpose of the test bed.

Until recently the DEWSAR TB has been primarily focussed on the collection and generation of EW information of radar targets. The target set of the DEWSAR TB now includes communications emitters. The new focus on communications targets has

UNCLASSIFIED

UNCLASSIFIED

generated a requirement for new sensors and nodes in the test bed. Our research considered architectures required to network the sensors, and manage and process the information and discusses the scalability of the test bed for larger numbers of nodes.

Several strategies for improving the scalability of wireless networks have been developed including the hierarchical, location centric and information centric approaches. These three approaches aim to minimise network communications and hence improve the performance and scalability of ad hoc wireless networks. Each of these approaches has drawbacks for the DEWSAR TB:

- Implementing the hierarchical approach in the DEWSAR TB would be technically difficult and would substantially increase the complexity of the steering algorithms.
- The location centric approach may not provide the best sensor geometry for target geolocation.
- The data provided by the DEWSAR TB's passive ES sensors makes implementing an information centric approach difficult.

A better strategy to facilitate the scalability of wireless networks for the DEWSAR TB may be the 'networks of networks' strategy where the structure of a network may be virtual and adaptive. When addressing the question of how many sensor nodes should be deployed in a network, it is important to bear in mind that the efficiency of information exchange decreases as the number of sensors increases. Thus the number of sensors in a wireless network should be minimised in order to optimise the efficiency of information exchange. However, in practice, the number of sensor nodes in a network will be constrained by a number of factors, such as:

- the objective of the mission
- the desired detection accuracy
- the time available to achieve the desired detection accuracy
- the capabilities of the sensors including their mobility
- the capabilities of the targets including their mobility
- environmental factors
- the capacity of the communications network
- the need to exchange information between sensors within the sensor group and between other networks of sensors.

UNCLASSIFIED

Authors

Kuba Kabacinski

Consunet Pty Ltd (under contract to DSTO)

Kuba Kabacinski obtained a Bachelor in Information Technology from Flinders University in 2001 and specialises in network enabled software and Service Oriented Architecture design at Consunet Pty Ltd. His primary experience lies in research and development projects conducted by the Defence Science and Technology Organisation. This work has involved several multinational projects where he has worked with leading research teams from Australia and around the globe. The projects span a variety of areas such as logistics planning, autonomous UAV systems and embedded systems. Each project had a strong requirement for networked, distributed data processing and storage as well as multi-user access.

Catherine Howard

Electronic Warfare and Radar Division

Catherine Howard is a research scientist in Electronic Warfare and Radar Division's Distributed Electronic Warfare group. She has a BSc degree (1992) from University of Queensland, an Honours degree (1993) from the University of Tasmania, a Graduate Diploma in Computer and Information Science (2002) from the University of South Australia and a PhD in Information Technology (2010) from the University of South Australia. She has 16 years experience working at DSTO in a variety of fields including passive missile warning systems, radio frequency propagation in maritime environments, modelling and simulation, the exploitation and dissemination of Electronic Intelligence and the automation of higher level information fusion processes.

UNCLASSIFIED

Damian Hall

Electronic Warfare and Radar Division

Damian Hall has worked in Intelligence Surveillance and Reconnaissance Division and Electronic Warfare and Radar Division since joining DSTO in 1998. In ISRD he modelled integrated surveillance systems and analysed their performance before moving to Electronic Warfare and Radar Division in 2001 to work in the Land and Force Level Electronic Warfare areas. Since 2006 he has been the head of the Distributed Electronic Warfare Group leading the research and development of modelling, analysis and experimentation of distributed Electronic Warfare systems.

UNCLASSIFIED

Contents

LIST OF ACRONYMS

1. INTRODUCTION.....	1
2. DE-CENTRALISATION AND DISTRIBUTION OF DEWSAR FUNCTIONS AND PROCESSES	2
2.1 The DEWSAR Test Bed	2
2.2 De-centralisation versus Distribution	3
2.3 Wireless Network Topologies	3
2.4 The DEWSAR Functions and Processes that Could Potentially be De-centralised or Distributed	5
2.5 Considerations to be Kept in Mind When Evaluating the De-centralisation and Distribution Options	8
2.5.1 The Purpose of the DEWSAR Test Bed	8
2.5.2 The Size, Weight and Power Constraints of the Aerial Platforms....	8
2.5.3 Bandwidth	10
2.6 Advantages of De-centralisation and Distribution	10
3. DE-CENTRALISATION AND DISTRIBUTION OPTIONS	12
3.1 The Tuning Function.....	12
3.2 The Classification Process.....	13
3.3 The Aggregation Process	14
3.4 The Geolocation Process.....	15
3.5 The Tracking Process	17
3.6 The Steering Function.....	18
3.7 The Reporting Process	19
4. THE SCALABILITY OF WIRELESS NETWORKS	20
4.1 The Hierarchical Approach.....	22
4.2 The Location Centric Approach	23
4.3 The Information Centric Approach.....	24
4.4 Networks of Networks.....	24
5. CONCLUSIONS.....	26
6. ACKNOWLEDGEMENTS	28
7. REFERENCES	28

Figures

Figure 1. The current network topology of the DEWSAR TB. The topology is best described as a star topology..... 4

Figure 2. A mesh network topology..... 4

Figure 3. A general overview of the DEWSAR TB's low level control functions and data processing steps. In this diagram, the processing flows are illustrated as a Unified Modelling Language Activity Diagram. 6

Figure 4. The Aerosonde Miniature UAV. From (Devries 2007). 9

Figure 5. The results obtained by Li, Blake and D. D. Couto et al (Li, Blake and D. D. Couto, et al. 2001) using real hard ware. The plot shows the throughput achieved along a chain of communication nodes as a function of the length of the chain. In Li, Blake and D. D. Couto et al's experiment, each node was placed at the maximum distance from the previous node that allowed for low-loss communications. 22

Tables

Table I. The reduction in data volume achieved by the Aggregation Process for three recent experiments. In this table, the PDW word size was 54 bytes and the SDW word size was 164 bytes. 15

List of Acronyms

Acronym	Definition
ADF	Australian Defence Force
CPU	Central Processing Unit
DEW	Distributed Electronic Warfare
DEWSAR	Distributed Electronic Warfare Situational Awareness and Response
DPDM	Discrete Probability Density Method
ES	Electronic Support
EO	Electro-Optic
FPGAs	Field Programmable Gate Arrays
IR	Infra-red
LADAR	Laser Detection and Ranging
OSX	Mac OS X; the apple operating system
PDWs	Pulse Descriptor Words
RADAR	Radio Detection and Ranging
RF	Radio Frequency
SDWs	Signal Description Words
SEI	Specific Emitter Identification
SWAP	Size, Weight and Power
TB	Test Bed
UAV	Uninhabited Aerial Vehicle

UNCLASSIFIED

DSTO-TR-2918

This page is intentionally blank

UNCLASSIFIED

1. Introduction

Technological advances in miniaturised Electronic Warfare (EW) systems and network connectivity are enabling a paradigm change for military commanders. In this new paradigm (Brown, Bailey, Drake et al. 2005) the focus has changed from survivability systems for self protection of a single asset against terminal threats to a whole of system approach that utilises elements including counter surveillance jamming, EW data exploitation and fusion to enhance real-time situation awareness, and human factors. These elements combined with stand-in autonomous systems for proactive shaping of the electromagnetic battle-space are enabling advances to the tactical commander's ability to achieve mission success. To investigate how this paradigm shift can be leveraged by the Australian Defence Force (ADF), in conjunction with Australian industry and academia, The Defence Science and Technology Organisation's (DSTO) Electronic Warfare and Radar Division has developed the Distributed Electronic Warfare Situation Awareness and Response (DEWSAR) Test Bed (TB).

The DEWSAR TB was designed to enable the rapid development and evaluation of distributed Electronic Warfare (DEW) concepts by utilising a networked set of heterogeneous, relatively unsophisticated EW sensors and effectors (deployed on ground based or aerial platforms) to detect, identify, locate, track or suppress stationary or slow moving surface based RF emitting targets. In the current DEWSAR TB, little data processing occurs at the (ground based or airborne) nodes; sensor data is transmitted from the node to the ground station where it is collected, processed, stored and disseminated as appropriate.

Previous demonstrations and experiments with the DEWSAR Testbed include the coordinated detection, geolocation and surveillance of adversary radar systems, coordinated detection, geolocation and jamming of adversary radar systems (Mason, Brown and Kabacinski 2009), and the coordinated detection and geolocation of adversary radar systems by indigenous and third party sensors to produce fused, targeting coordinates that were of sufficient quality for simulated interdiction (Gibard, Kabacinski, Drake et al. (in preparation)). The latter trial examined and demonstrated the DEWSAR TB's ability to enhance the suppression of adversary air-defences (SEAD) capability of a simulated FA-18 asset to engage radar targets in real-time. Furthermore, two mobile, ground-based ADF EW sensors, the Passive Radar Identification System (PRISM) were integrated into the DEWSAR TB to supplement the EW data provided by the airborne sensors.

Until recently the DEWSAR TB has been primarily focussed on the collection and generation of EW information of radar targets. The target set of the DEWSAR TB now includes communications emitters. The new focus on communications targets has generated a requirement for new sensors and nodes in the test bed. Consequently architectures required to network the sensors, and manage and process the information must be considered. A number of options were evaluated for migrating some of the DEWSAR TB's control functions and data processing from the ground station to the

ground based or airborne nodes and considers these options along with the scalability of the test bed for larger numbers of nodes.

The rest of this document is structured as follows; Section 2 discusses the test bed in further detail and differences between de-centralisation and distribution, wireless network topologies, and the DEWSAR functions and processes that could potentially be de-centralised or distributed. Section 3 evaluates the de-centralisation and distribution options. Section 4 considers architectures required to network the sensors including scalability of the DEWSAR TB. Section 5 presents the conclusions.

2. De-centralisation and Distribution of DEWSAR Functions and Processes

2.1 The DEWSAR Test Bed

The DEWSAR TB was designed to enable the rapid development and evaluation of DEW concepts. The test bed utilises a networked set of heterogeneous, relatively low cost, unsophisticated EW sensors and effectors which can be deployed on ground based or aerial platforms. The test bed uses its sensor and effectors to detect, identify, locate, track or suppress stationary or slow moving surface based RF emitting targets such as air defence and maritime navigation radars.

Sensors which have been developed for the DEWSAR TB include wide band and narrow band electronic support (ES) sensors, a stabilised pan-tilt-zoom electro-optic (EO) sensor, a high resolution, still-image EO camera, multi-aspect, fixed orientation and gimballed infra red (IR) (Division) sensors, acoustic sensors and a high resolution scanning laser radar (LADAR). Effectors which have been developed for the test bed include autonomous electronic attack (Division) payloads.

Consider an example of how the DEWSAR TB can be used. An operator, who is located at the ground station, issues a command to the airborne sensor nodes to search a particular geographic area. Based on the tasking instructions issued by the operator, the payloads onboard each of the aerial platforms and any available a priori information about the potential targets in the search area, the DEWSAR TB autonomously selects a sub-set of the nodes to perform the task/s. For example, the test bed could use aerial platforms with broadband receivers to monitor the electromagnetic spectrum. Detections from the broadband receivers could then be used to cross-cue nodes with narrow band receivers or sensors with shorter ranges and/or more restricted fields of view such as EO or IR sensors in order to identify, geolocate or track the target.

Past demonstrations and experiments with the DEWSAR Testbed include the coordinated detection, geolocation and surveillance of adversary radar systems, coordinated detection, geolocation and jamming of adversary radar systems (Mason, Brown and Kabacinski 2009), and the coordinated detection and geolocation of adversary radar systems by

indigenous and third party sensors to produce fused, targeting coordinates that were of sufficient quality for simulated interdiction (Gibard, Kabacinski, Drake et al. (in preparation)). The latter trial examined and demonstrated the DEWSAR TB's ability to enhance the suppression of adversary air-defences (SEAD) capability of a simulated FA-18 asset to engage radar targets in real-time. Furthermore, two mobile, ground-based ADF EW sensors, the Passive Radar Identification System (PRISM) were integrated into the DEWSAR TB to supplement the EW data provided by the airborne sensors.

In summary, the DEWSAR TB consists of ground and aerial sensor/effector nodes and platforms, communications and network infrastructure, low and high level data exchange protocols and low and high level data processing. In the current DEWSAR TB, little data processing occurs at the (ground based or airborne) nodes; only basic data conditioning¹ is performed at the nodes. The sensor data is transmitted from the sensor to the ground station where it is collected, processed, stored and disseminated as appropriate. Tasking instructions are issued by the operator at the ground station. From the ground station, the operator can control the individual aerial platforms and/or teams of the aerial platforms by specifying waypoints, targets, search zones, exclusion zones or other spatially meaningful instructions. The operator can also control the individual payloads deployed on the ground or aerial platforms.

2.2 De-centralisation versus Distribution

This subsection discusses the differences between de-centralisation and distribution. In this document, **de-centralisation** is defined as controlling the test bed's sensors or aerial platforms from a location other than the ground station while **distribution** is defined as the sharing of data, information or data processing algorithms between multiple nodes.

2.3 Wireless Network Topologies

The current network topology of the DEWSAR TB, which is shown in Figure 1, is best described as a star topology.

¹ Data conditioning includes removing aberrant data reports and performing any required time stamp corrections (these corrections are required if there is a drift in the payload clock).

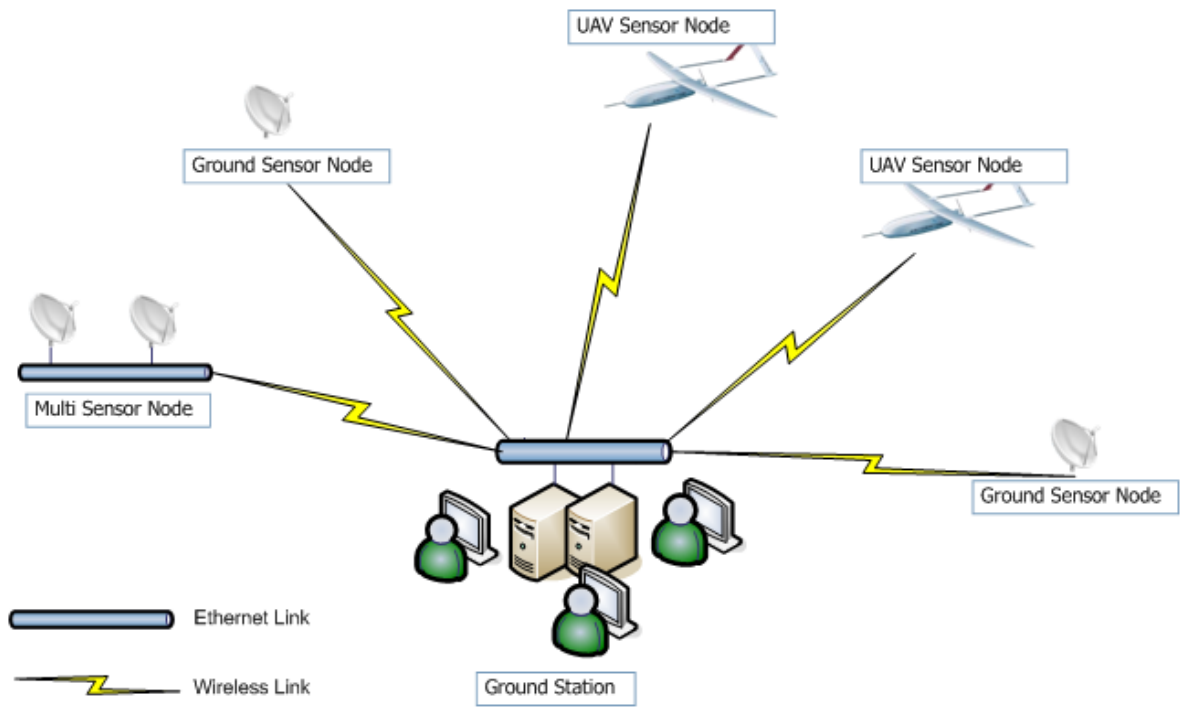


Figure 1. The current network topology of the DEWSAR TB. The topology is best described as a star topology.

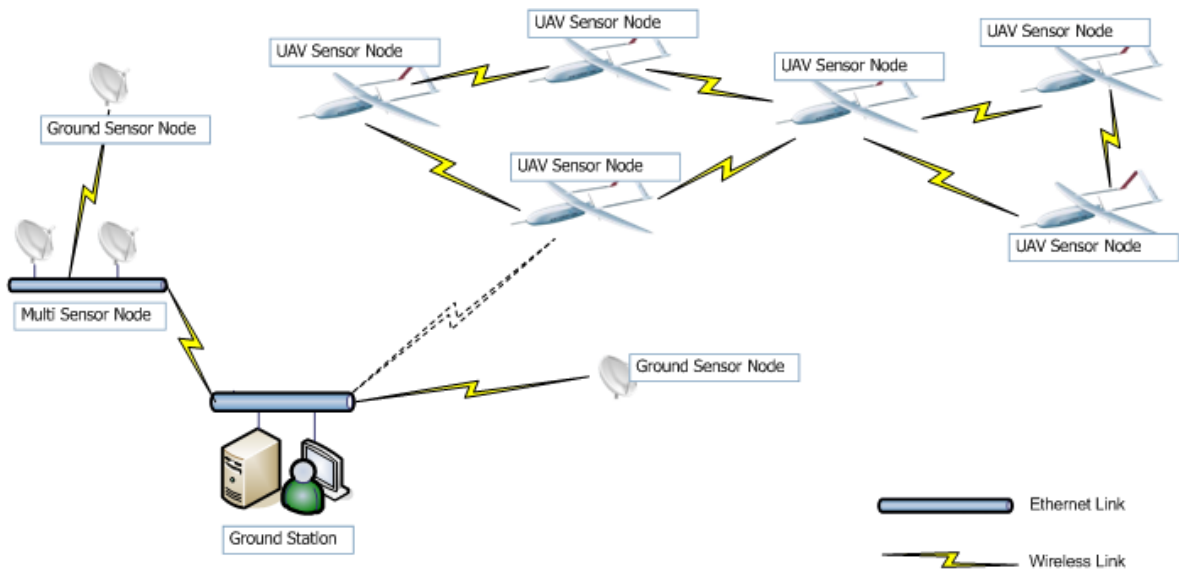


Figure 2. A mesh network topology.

De-centralising the DEWSAR TB's control functions² or distributing the DEWSAR TB's data, information or data processing algorithms entails migrating these functions, data, information or algorithms from the ground station to the ground based or airborne nodes, either onboard the platform or as part of the sensor/effector itself. This would mean that the Ground Station as shown in Figure 1 is no longer a single point of failure for communication, control or data processing. A topology such as the mesh topology shown in Figure 2 could be used to achieve this. In a mesh topology, the network nodes are capable of performing data forwarding and routing.

The processing currently performed at the ground station itself could also be distributed among a number of computers at the ground station; the processing capability of the ground station contains configurable components that can be orchestrated into parallel and distributed configurations, provided there is network connectivity. However, as distributing the processing in this fashion would not change the fact that little of the control or data processing occurs at the (ground based or airborne) nodes, this concept is not explored in this document.

2.4 The DEWSAR Functions and Processes that Could Potentially be De-centralised or Distributed

This subsection outlines the DEWSAR functions and processes that could potentially be de-centralised or distributed. Figure 3 presents a general overview of the test bed's low level control functions and data processing steps. The overview represents the identification, geolocation, tracking and reporting of a target by the test bed's ES sensors.

² For the system to be useful in a military sense, high level control must be centralised such that there is a single commander; similar to the military chain of command. However, low level control (for example: fine grained tasking which requires a quick turnaround) could be de-centralised.

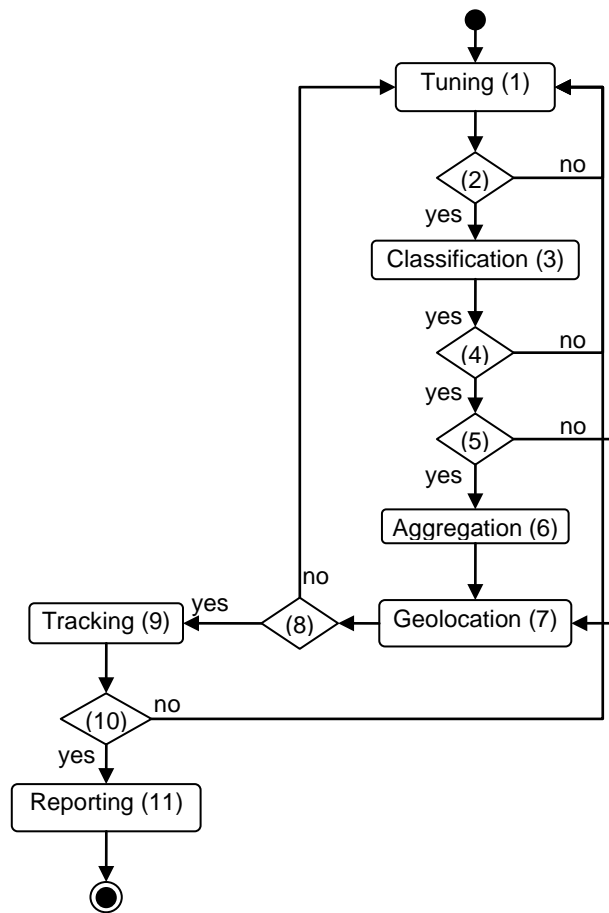


Figure 3. A general overview of the DEWSAR TB's low level control functions and data processing steps. In this diagram, the processing flows are illustrated as a Unified Modelling Language Activity Diagram.

- Step 1. The ES sensors accept tuning instructions from the test bed. Tuning is a low level function that directly affects the data captured by the sensor. The Tuning Function does not apply to all sensors; wide band ES sensors, for example, do not require tuning. Narrow band ES sensors, however, do require tuning to ensure that they can detect a signal of interest. When tuning is required, a trigger or schedule based strategy must be used to tune the sensors. These strategies are discussed further in Section 3.1.
- Step 2. Determines whether the narrow band ES sensor has been successfully tuned.
- Step 3. The Classification Process involves matching the received signal against a library of known radars in order to identify the target.
- Step 4. Determines whether the signal belongs to a known radar type. Unknown pulses are not currently processed by the test bed.
- Step 5. Determines whether the processing further downstream can handle the output from the sensor or whether the sensor data needs to be aggregated. The ES sensors represent the parameters of the received signal using Pulse Descriptor

Words (PDWs). PDWs include information such as the frequency, pulse width and amplitude of the detected radar and the latitude, longitude, altitude, azimuth angle and elevation angle of the sensor.

- Step 6. During the Aggregation Process, the PDWs generated by a particular sensor are aggregated to form a Signal Descriptor Word (SDW). SDWs include information such as the latitude, longitude, altitude, scan peak, scan period, scan type, frequency, pulse width, pulse repetition interval and amplitude of the detected radars and the latitude, longitude, altitude, azimuth angle and azimuth error and the elevation angle and elevation error of the sensor.
- Step 7. During the Geolocation Process, the test bed uses all relevant data to Geolocate the targets. The test bed has several geolocation techniques; the selection of a suitable geolocation technique depends on several factors such as the capability of the sensors, the geographic distribution of the sensors and targets and the quality of the sensor data.
- Step 8. Determines whether the quality of the target geolocation is satisfactory. Because the test bed can execute multiple geolocation algorithms in parallel, the geolocation results can be compared and the best results can be selected for reporting.
- Step 9. The Tracking Process involves the temporal and geo-spatial tracking of targets. The test bed has several tracking algorithms and, as will be discussed in Section 3.6, the Tracking Process and the Steering Functions are coupled.
- Step 10. Determines whether the quality of the target track is satisfactory.
- Step 11. The Reporting Process packages information as a command and control level product and distributes the information through the test bed's server to provide situation awareness and facilitate decision making.

Steps 1 to 10 require sequential execution for each target. However each step can be performed independently to others and in parallel for each target. Each of the steps can have operator participation.

It must be borne in mind that Figure 3 is a general overview of the DEWSAR TB low level data processing steps; for some of the DEWSAR TB's geolocation algorithms, some of the processes are more tightly coupled than indicated in this figure. For the Discrete Probability Density Method (DPDM) geolocation algorithm, for example, the geolocation and tracking processes are merged. Also because Figure 3 presents a general overview of the test bed's low level control functions and data processing steps it does not include the Steering Function. The Steering Function is used to control the trajectory of the aerial platforms during deployment. It is a high level control function which occurs asynchronously and uses a priori information (such as the intent of the mission) and newly discovered information to position the platforms and sensors.

In summary, the Tuning and Steering functions are functions which are used to control the test bed's sensors or aerial platforms while the Classification, Aggregation, Geolocation, Tracking and Reporting processes are data processing steps which convert the raw sensor

data into information products, such as target tracks or geolocations. This paper will evaluate options for de-centralising the Tuning and Steering functions and distributing the Classification, Aggregation, Geolocation, Tracking and Reporting processes.

2.5 Considerations to be Kept in Mind When Evaluating the De-centralisation and Distribution Options

This subsection discusses a number of considerations which must be kept in mind when evaluating the de-centralisation and distribution options for the DEWSAR TB.

2.5.1 The Purpose of the DEWSAR Test Bed

The most important thing which must be kept in mind when evaluating de-centralisation and distribution options is the purpose of the test bed. The test bed was designed to enable the rapid development and evaluation of DEW concepts. Hence any de-centralisation or distribution option which impinges a researcher's ability to rapidly develop and evaluate algorithms or concepts contravenes the purpose of the test bed.

2.5.2 The Size, Weight and Power Constraints of the Aerial Platforms

The aerial platform currently utilised by the DEWSAR TB is the Aerosonde Miniature Uninhabited Aerial Vehicle (UAV), shown in Figure 4. The Aerosonde is a small (the approximate wingspan is 2.9 meters), stealthy, low cost, long endurance UAV. Because the Aerosonde is small, its physical characteristics will limit the:

- Size, weight and configuration of the payloads. The maximum payload size is 5.5 litres (which must be distributed over two payload compartments) and the maximum payload weight is 2.5 kilograms. For radio frequency (RF) payloads, for example, the size and weight of the Aerosonde platform represent severe physical constraints; RF payload components can easily exceed these limits.
- The amount of power available onboard the UAV for:
 - payloads
 - communication³
 - onboard data processing.

Hereafter, these constraints will be collectively referred to as the Aerosonde's Size, Weight and Power (SWAP) constraints.

³ Communicating over large distances requires a lot of energy due to path loss. Path loss, or path attenuation, is the reduction in power density of an electromagnetic signal as it propagates through the atmosphere. Path loss may be due to many effects, such as refraction, diffraction, reflection and absorption and can be influenced by terrain, environment, the distance between the transmitter and receiver and the height and location of the antennas. When taking into account effects of interference, obstacles and hardware imperfections, path loss can become a significant problem for distributed sensor network communications where the power available for communication is limited.

Due to the Aerosonde's SWAP constraints, the sensors and effectors which can be deployed on these platforms are generally significantly less capable than the sensors and effectors which can be deployed on larger, more strategic UAVs such as Global Hawk or Predator. However, this may be offset by the:

- Affordability of the platform, sensors and effectors and hence the capacity to withstand losses due to conflict or due to malfunction. The Aerosonde is 'affordably expendable'. The concept of affordable expendability relies upon the notion that the useful life of the platform is a function of its constituent payloads and technologies rather than a function of the physical lifespan of the airframe (Devries 2007).
- Ability to network the UAVs, sensors and effectors to derive process gain. The observation system which is created through the cooperation and adaptive networking of the UAVs, sensors and effectors provides sufficient process gain to achieve results similar to those of significantly more expensive platform-centric systems, but with the added advantage of achieving robustness to an adversary's counter-measures whilst simultaneously maintaining operational capability at a reduced cost. Even if it were possible for an adversary to target the individual UAVs, the distributed, autonomous, adaptive, and robust nature of the system makes it difficult to counter effectively (Finn, Brown and Lindsay 2002).

Due to their affordable expendability and robustness, adaptive, distributed sensor and targeting networks such as the DEWSAR TB, offer the prospect of a capability that allows exploration of the high value, high risk missions that are beyond the justifiability or capability of more strategic systems such as Global Hawk or Predator (Elsaesser 2007).



Figure 4. The Aerosonde Miniature UAV. From (Devries 2007).

2.5.3 Bandwidth

Bandwidth for communication between the Ground Station and the ground and aerial nodes is limited⁴. A sensor's capability to distribute data in a wireless network is limited by the available RF bandwidth (Hz) and the sensor's communication spectral efficiency (bits/sec/Hz). It is communication constraints which prevent sensors in a network from passing arbitrarily large volumes of data to one another. High bandwidth communication conduits are needed in sensor networks to facilitate scalability and low latency communications. However, high bandwidth between any two nodes in a network is not an indicator of the overall network's capacity to communicate data between all of the nodes.

Due to the fact that there is a finite amount of energy available on the aerial platforms, irrespective of the method of implementation, there is an upper limit to the amount of data processing which can be done on-board⁵. Because of the Aerosonde's SWAP constraints, any processing capability available onboard the aerial platforms would not be comparable to the processing capability currently available at the ground station. Hence any functions or algorithms which are de-centralised or distributed to the aerial nodes must be efficient in utilizing the limited power and bandwidth available.

2.6 Advantages of De-centralisation and Distribution

There are a number of advantages of de-centralising the test bed's control functions or distributing the test bed's data, information or data processing algorithms. For example:

- Adopting a mesh topology (shown in Figure 2) rather than a star topology (shown in Figure 1):
 - Would enable the aerial platforms to be operated at larger distances from the ground station. This would provide a higher level of protection to the ground station operators without sacrificing capability⁶.
 - Could potentially make the DEWSAR TB more scalable. The maximum number of nodes used in the test bed to date is ten. Even this number of sensors occasionally stretches the capabilities of the communications network. Using a mesh topology could mean that a larger number of nodes

⁴ An alternative to transmitting sensor data to the ground station or other nodes is to store the data onboard the sensor node. However, an important limitation of this type of approach is that the data can not be processed until the UAV returns to its base, which vastly reduces the types of scenarios for which such a system would be useful (for example, the detection and interrogation of mobile targets would not be possible with onboard data storage). As such, this alternative is not considered in this paper.

⁵ There is also a theoretical lower limit on how much energy is needed to perform a computation. The minimum power for computation limit is described by information theory as the Landauer-Shannon limit which identifies the minimum energy required to process one bit of data. Current technology is some way off achieving the Landauer-Shannon limit; it is anticipated that technology will reach this limit in approximately twenty years (Izydorzcyk and Cionaka).

⁶ Other strategies, such as the use of communication relays, could also be used to minimize the deployment risk to ground station operators.

could be employed, increasing geographic and RF coverage and providing potential for redundancy. The increased geographic and RF coverage could improve target detection rates and the redundancy could improve the overall robustness of the system. However, perhaps the most important limitation of wireless networks is their scalability. Section 4 discusses the scalability of wireless networks in further detail.

- Distributing the test bed's data, information or data processing algorithms:
 - May increase the overall accuracy of the generated information products. In the current DEWSAR TB, due to the limitations of the data communications channels, some sensor data is discarded⁷.
 - Should reduce the network bandwidth required for communications, as more data processing would be carried onboard the nodes. For example, current ES sensors used in test bed have a sampling rate of 10 nanoseconds, resulting in peak internal data rate of 10^8 samples/second. The sensors then perform basic processing and report at a maximum external rate of 115,200 baud resulting in a sample rate of roughly 1200 samples/second. This is further reduced to 19,200 baud for transport over a wireless modem communication channel which equates to a maximum of 184 samples/second. If data processing up to and including the Aggregation Process were performed onboard of the aerial nodes, the data transfer rate requirement could be reduced by up to a factor of 10. If more processing was carried onboard the nodes:
 - Lower volume information products could then be exchanged with other nodes to achieve a cooperative functionality. If the neighbouring nodes were relatively close to each other, the local network bandwidth would be larger than a long distance link to a ground station. The lower signal strength required for local communication between neighbouring nodes would also mean that the vulnerability of the nodes to adversary surveillance would be reduced.
 - The reduction in the bandwidth required for communications would enable sensors with better capabilities to be more readily incorporated into the test bed (the current utilisation of relatively unsophisticated EW sensors and effectors in the test bed does not preclude use of higher quality sensors and effectors in the future). Because the test bed currently transmits the sensor data to the ground station where it is processed, the maximum data output rates of the test bed sensors are effectively limited by the bandwidth of the wireless network. While this limitation has not been a major cause of concern for past experiments (because of the relatively simple RF environments utilised during these experiments), in real

⁷ Loss of sensor data has not been a significant issue in the DEWSAR TB experiments to date because these experiments have always involved controlled RF environments containing a limited number of transmitters.

world, complex RF environments, this limitation may present a challenge.

- The test bed could operate in more complex RF environments.

3. De-centralisation and Distribution Options

Keeping the considerations outlined in Section 2.5 in mind, this section evaluates a number of options for de-centralising the Tuning and Steering functions and distributing the Classification, Aggregation, Geolocation, Tracking and Reporting processes outlined in Section 2.4.

3.1 The Tuning Function

This subsection evaluates the de-centralisation options for the Tuning Function. As discussed in Section 2.4, tuning does not apply to all sensor types. However, when tuning is required, a suitable tuning strategy must be used, for example trigger based or schedule based tuning.

The current DEWSAR TB uses a trigger based tuning strategy. A trigger based strategy is where an observed or internal event is used to perform a tuning operation. Such events could include, for example, the received signal falling below a specified minimum threshold. The DEWAR TB's tuning commands are issued by an agent located at the ground station. The agent determines the tuning requirements using information such as the intent of the mission, the priority of the target, the proximity of the sensor to other sensors, the UAV teaming behaviour, target detections from wide band receivers etc. Hence the agent uses both high level and low level information to determine the tuning requirements. De-centralising the test bed's trigger based tuning implies that, instead of being centrally located at the ground station, the tuning agent would be installed on any of the ground based and aerial platforms carrying sensors which require tuning. It also implies that the low level and high level information required by the agent to determine the tuning requirements would need to be distributed to each sensor which requires tuning. This would create a data distribution overhead, substantially increasing the complexity of the resultant system. Hence de-centralising the test bed's current agent controlled trigger based tuning strategy may not be practical.

Furthermore, consider the scenario where the test bed has deployed a number of narrow band ES sensors on aerial platforms distributed in a particular area of interest. If a target begins to emit, it would be desirable for all of the deployed sensors to collect data. However, for this to occur, the tuning of the sensors needs to occur in an orchestrated fashion. Because sophisticated radars may only emit for brief periods, the orchestrated tuning of multiple sensors can be difficult to achieve as the time required to distribute the tuning information may exceed the radar's emission time. An orchestrated trigger based tuning strategy is best achieved using a deterministic algorithm. Such an algorithm would

operate on either equal or varying inputs. However, utilising equal inputs requires the dissemination of large data sets, resulting in a communication overhead. Utilising varying inputs may reduce the communication overhead. However, it is difficult to guarantee a deterministic and consistent result when utilising varying inputs.

Now considering schedule based tuning. Schedule based strategies tune the sensor based on a list; this list can either be a predetermined list or a list which is dynamically adapted. Because the information required for schedule based tuning is simple and doesn't need to be re-communicated often, de-centralisation of a schedule based strategy has low technical requirements. However, schedule based tuning is not as adaptive as the test bed's current agent controlled trigger based tuning strategy. The differences between adaptive schedule based tuning and predetermined schedule tuning approaches would need to be investigated further before a recommendation can be made as to whether the improvement would be significant enough to justify the cost.

3.2 The Classification Process

This subsection evaluates the distribution options for the Classification Process. As discussed in Section 2.4, the Classification Process involves matching the received signal against a library of known radars in order to identify the target. The library is a catalogue of radars and their identifying signal characteristics. The test bed's current Classification Process uses a relatively simple strategy; the parameters of the received signal are compared against every entry in the library using an exhaustive search strategy.

A number of techniques could be used to distribute the Classification Process. For example each sensor could be:

- assigned only a subset of intercepts to classify
- assigned all the intercepts to classify using a predetermined subset of the library
- a combination of both.

However, for all of these techniques, the distribution of the Classification Process would come at a cost; network communication would be increased as the data is disseminated and the results collected. Furthermore, if each sensor were assigned only a subset of intercepts to classify, further computation would be required to determine how the intercepts should be divided among the sensors. Consider the situation where two sensors observe the same signal almost simultaneously. Which sensor should be assigned to classify this intercept? Also consider the situation where a sensor is classifying some intercepts which it has collected and, in addition to its own intercepts, the sensor is then assigned some intercepts from other distributed sensors. Due to the SWAP constraints, the aerial platforms may not have sufficient processing capability to cope with this situation. These examples highlight the fact that attempting to distribute the Classification Process may only serve to increase the network communication overheads without necessarily providing any benefits.

A number of other issues also indicate that attempting to distribute the Classification Process may not prove to be beneficial. For example:

- The Classification Process is computationally expensive; it is exponential in the number of intercepted signals. As the library increases in size (and hence becomes more useful for classifying radars), the computational complexity of the Classification Process also increases. In past experiments the DEWSAR TB has been deployed in sparse radars environments. In such environments, the current computational facilities available onboard the aerial platforms might be able to satisfactorily perform the Classification Process. However, this is unlikely to be true for dense radars environments.
- The accuracy of DEWSAR TB sensors. The accuracy of the sensors currently used within the DEWSAR TB is low and the current classification strategy works well for pulsed radars. As the accuracy of the sensor increases, different classification strategies, for example Specific Emitter Identification (SEI), become more appropriate. In SEI, precise signal measurement is used to determine the radar's identity. A SEI strategy could be more computationally tractable than the strategy currently implemented in the DEWSAR TB because SEI allows radars to be identified using a single intercept. However, in order to implement a SEI strategy in DEWSAR TB, new sensors and data representation and exchange techniques would be required. Also, given the SWAP constraints of the Aerosonde UAVs it may not be possible to deploy sensors capable of SEI on these platforms.

Quantifying the computational resources required by the aerial platforms in the DEWSAR TB to perform the Classification Process is not possible without information about the classification strategy to be used (e.g. the current exhaustive library matching strategy or SEI) and the taxonomy (for example: pulsed, continuous wave or low probability of intercept radars) and density of the radar target environment. In a real world deployment, information about the taxonomy and density of the radar target environment may not be available.

3.3 The Aggregation Process

This subsection evaluates the distribution options for the Aggregation Process. As discussed in Section 2.4, during the Aggregation Process the PDWs generated by a particular sensor are aggregated to form a SDW. The amount of processing required to aggregate PDWs into SDWs depends on:

- The sensor's capabilities (type⁸ and volume of measurements taken and the sensitivity of the sensor)
- Proximity of the sensor to the target/s of interest

⁸ The type of measurements taken by the sensor has a major impact on the amount of processing required to aggregate PDWs into SDWs. For example, consider the situation where the sensor has Direction Finding capability and the *Classification Process* has resulted in an unambiguous match. Producing a SDW which aggregates the PDWs from this sensor over a brief time period would not require complex data analysis and hence the *Aggregation Process* would not require large amounts of processing. However, consider the situation where the sensor does not have Direction Finding capability. Without line of bearing information from the sensor, the radars line of bearing must be calculated during the *Aggregation Process* by analyzing the sequence of PDWs, hence incurring an additional computational cost.

The main benefit in migrating the Aggregation Process from the ground station to the ground based or airborne nodes would be the reduction in the volume of data transmitted over the wireless network. The reduction in the data volume would lead to a reduction in the bandwidth required to communicate information around the network. The amount of reduction would depend on the number and types of sensors in the test bed and the taxonomy and density of the radar environment. A reduction in bandwidth requirements could result in savings in the communications power budget, the ability to add more sensor nodes to the network or the ability to increase the endurance of the aerial platforms in the test bed. Table 1 illustrates the reduction in data volume achieved by the Aggregation Process for three recent experiments.

While distributing the Aggregation Process may have significant benefits in terms of network utilisation, it may increase the complexity of the test bed.

Table 1. The reduction in data volume achieved by the Aggregation Process for three recent experiments. In this table, the PDW word size was 54 bytes and the SDW word size was 164 bytes.

Experiment	Sensor	PDW Word Count (word/min)	SDW Word Count (word/min)	Reduction in Data Volume	Reduction in Bandwidth Usage
Lakehurst , New Jersey, USA, September 2006	SELEX (UK) ground based ES sensor	85,852	2,665	97%	90%
Fort Bliss , New Mexico, USA, September 2007	Heterogeneous ES sensors	7,579	503	93%	80%
Woomera , South Australia, May 2009	Heterogeneous ES sensors	316	22	93%	98%

3.4 The Geolocation Process

This subsection evaluates the distribution options for the Geolocation Process. As discussed in Section 2.4, during the Geolocation Process the test bed uses all relevant data to geolocate the targets. In addition to the general advantages of de-centralisation and distribution discussed in Section 2.6, distributing the Geolocation Processes would have a number of other advantages. Distributing the Geolocation Processes would:

- Lead to a reduction in communication bandwidth; for a target radar, the Geolocation Process combines the SDWs produced in the Aggregation Process into a geolocation report.
- create a shared awareness of the locations of targets

- Lead to improved steering behaviour. As will be discussed in Section 3.6, the current test bed uses a steering agent to control the flight paths of the aerial platforms on a macro scale. If each aerial platform had ready access to target location estimates, micro scale alterations could be made to the platform's flight path or sensor configuration to enable better target detection. For example, the platform's relative aspect to the target could be changed to improve the orientation of the antenna or to minimise the radar cross section of the platform. Such fine grained control of the platforms and sensors is not readily achievable in the current test bed due to communication latency and bandwidth limitations.

Two approaches can be taken to distributing the Geolocation Processes:

- distribute all the available data to all of the nodes
- only distribute data about a subset of the targets to predetermined nodes.

In the case where all the available data is distributed to all of the nodes, each node is performing geolocation for all the targets; a (perhaps unnecessary) duplication of effort. Assuming that the geolocation algorithms used by the nodes are deterministic and time invariant and that each node uses the same algorithm and that each node receives the same data then each node should produce equivalent geolocation results. However, many of these assumptions are difficult to guarantee. For example, guaranteeing delivery of data transmitted around the wireless network will increase the network communication costs (due to retransmissions).

In the case where only data about a subset of targets is distributed to predetermined nodes, only one node performs geolocation for the subset of the targets, avoiding the duplication of effort. While this approach would require a more sophisticated data routing scheme, it may reduce the required communication bandwidth. Also, dividing the targets into subsets requires that they be unambiguously classified, which may not always be possible. Furthermore, if an aerial node was lost due to conflict or malfunction, the subset of targets assigned to this node would need to be reassigned.

Geolocation algorithms typically require an ordered data stream, but in a network with a mesh topology, data arriving at a node to be fused may travel along non uniform data paths⁹. Non-uniform data paths may cause data to arrive out of synch. Thus networks with a mesh topology may require each node to have a buffering facility to ensure that the data arriving from different sensors can be aligned. However, the buffering and sorting processes can become expensive when there are large differences between the shortest and longest paths in the network. It may also be difficult to synchronise the buffering across all the nodes in the test bed; attempting to do so may result in wait deadlock conditions which would be difficult to detect in a distributed system.

The test bed has several geolocation algorithms. Each of the algorithms has unique processing requirements and the algorithms are selected using factors such as the capability of the sensors, the geographic distribution of the sensors and targets and the

⁹ In a star topology, every node is one data hop away and hence the data paths are uniform.

quality of the sensor data. Because the test bed can execute multiple geolocation algorithms in parallel, the geolocation results can be compared and the best results can be selected for reporting. When distributing the geolocation algorithms, some algorithms face more challenges than others. For example, the DPDM (Elsaesser 2006; Elsaesser 2007) is a geolocation technique which relies on matrix multiplication. Hence this algorithm is computationally expensive (Baldock, Drake, Howard et al. 2009) and the computational resources required by this algorithm are unlikely to be available on the test bed's aerial platforms in the foreseeable future. It may be possible to reduce the algorithm's overall processing requirements by implementing some numerical efficiency gains on the algorithm's matrix multiplications. However the DPDM algorithm is still under development and hence it is impractical to implement these gains at this time.

Bearing in mind the Aerosonde's SWAP constraints, there are several processing platforms which could be used to perform processing onboard the ground based or airborne nodes, for example, Field Programmable Gate Arrays (FPGAs) or embedded processors. FPGAs are integrated circuits which are designed to be configured after manufacture while embedded processors are computer systems designed to do one or more dedicated and/or specific functions, often with real time computing constraints. Both FPGAs and embedded processors improve execution efficiency and hence power consumption at the expense of development time and development and maintenance costs. Due to these increased costs, it is not appropriate to implement an algorithm which is still under development, such as the DPDM geolocation algorithm, in FPGAs or embedded processors.

In summary, distribution of the test bed's Geolocation Process would substantially increase the complexity of the test bed.

3.5 The Tracking Process

This subsection evaluates the distribution options for the Tracking Process. As discussed in Section 2.4, the Tracking Process involves the temporal and geo-spatial tracking of targets. The test bed has a suite of tracking algorithms and multiple algorithms can be executed in parallel. This allows single target tracking algorithms to operate in a multi target environment and also allows the performance of multiple algorithms to be compared.

Two techniques could be used to distribute the test bed's Tracking Process; the information centric approach or the location centric approach.

The information centric approach leverages the information generated from the Classification Process; the Classification Process matches the received signal against a library of known radars in order to identify the target. The temporal and geo-spatial identification of targets could be used to form a track. This process could be distributed to the individual test bed nodes. However, the information centric approach relies on the targets being unambiguously classified, which may not always be possible. Also the information centric approach would require significant network resources; each node

would be responsible for tracking a particular sub-set of targets and hence a sophisticated data routing scheme would be required to ensure the node received all the relevant data.

The location centric approach (Brooks, Ramanathan and Sayeed 2003) dynamically allocates nodes into spatial cells. The nodes within each cell are then responsible for tracking the targets within the cell, the 'local' targets. If a mobile target leaves a cell, nodes adjacent cells are alerted so that the target can be successfully handed over. In this case, the location centric approach aims to reduce the communication requirements by ensuring that nodes exchange data directly rather than routing it through proxy nodes. The location centric approach may be difficult to implement for the DEWSAR TB. The large detection ranges of wide band ES receivers means that a large number of targets could potentially be visible to the node. Without accurate geolocation, it may be difficult for the node to determine if a target is within its designated spatial cell.

In summary, both the information and location centric approaches have limitations for the DEWSAR TB. Also distribution of the Tracking Processes would substantially increase the complexity of the test bed and yet may not yield savings in network utilisation.

3.6 The Steering Function

This subsection evaluates the de-centralisation options for the Steering Function. As discussed in Section 2.1, the test bed utilises a networked set of heterogeneous, relatively low cost, unsophisticated EW sensors and effectors which can be deployed on ground based or aerial platforms. Because they are relatively unsophisticated, the individual sensors are not capable of individually geolocating a target. Data from several sensors must be fused to provide an estimate of the target's location. The distance between the sensors and the target and the relative geometry of the sensors and the target play an important role in geolocation.

The DEWSAR TB uses a steering agent to control the flight paths of the aerial platforms on a macro scale. For RF angle of arrival sensors, the steering agent must ensure that the target is observed from multiple, sufficiently different angles in order to reduce triangulation error and minimise the uncertainty in the target geolocation (Dogancay and Hmam 2008). Where sensors are directional, the relative aspect of the platform and target is also important. The steering agent must also manage the connectivity of the network; the aerial platforms must remain within the line of sight of the ground station. When considering the de-centralisation of the Steering Function, it should be borne in mind that it takes time to change the node distribution within the network.

Non RF sensors, such as optical or acoustic, will need different steering strategies as their proximity to the target is more important due to propagation distance of measured emission and increased occurrence of environmental occlusions. Effector systems will have additional steering constraints dictated by mission specific requirements as well as competing goals.

When considering the de-centralisation of the Steering Function, for situations where the targets are stationary, it may well be the case that the de-centralisation of the Steering Function will allow the aerial platforms to more quickly improve their relative distributions, enabling the sensors to improve the target geolocations. However, for situations where the targets are mobile and operating in a 'pop-up' manner, even with a de-centralised Steering Function, it is unlikely that the platforms would be able to react quickly enough to improve their relative distributions. However, as previously mentioned, for directional sensors, the relative aspect of the platform and target is also important. Changing the relative aspect of the platform and target requires relatively little time and could provide significant information gain for sensors that do not have an omnidirectional field of view. For these sensors, de-centralisation of the Steering Function would be beneficial. However, de-centralisation of the test bed's Steering Function would substantially increase the complexity of the test bed.

Recent research into distributed tracking for mobile sensor networks (Olfati-Saber 2007) has shown that tracking and steering are coupled. This research has also shown that the geographic optimisation of sensor distribution based on an information centric cooperative tracking approach yields better targeting results. Interestingly, this research also showed that the distributed sensors tended to assemble into a regular pattern or 'flock'. Given this result and the fact that de-centralisation of the test bed's Steering Function would substantially increase the complexity of the test bed, instead of de-centralising the Steering Function, it may make more sense to simply deploy the test bed's aerial platforms in this flock pattern, retaining steering control at the ground station.

3.7 The Reporting Process

This subsection evaluates the distribution options for the Reporting Process. As discussed in Section 2.4, the Reporting Process packages information as a command and control level product and distributes the information through the test bed's server to provide situation awareness and facilitate decision making.

There are a number of data distribution paradigms including the data pull, data push, subscription based and localised broadcast paradigms. The test bed currently has provisions for the data pull and subscription based distribution paradigms. Each of these paradigms faces different challenges when evaluating distribution options.

While the data pull paradigm requires little in way of special planning (other than being able to address each of the data producing nodes in the network), it is not particularly useful in distributed sensor networks because nodes requiring data would need to continuously interrogate other nodes. This has a dual effect of inducing latency and increasing network bandwidth utilisation (through redundant requests). The pull frequency dictates the induced latency; as the pull frequency increases, the number of redundant requests which yield no new information increases.

In the subscription based paradigm, nodes register their interest in certain data, which is transmitted to them as it becomes available. While this method of distributing data is more

efficient than the data pull paradigm, it requires continuous monitoring of the subscriptions and the network topology to ensure that older subscriptions are still relevant.

Whether distributing the Reporting Process would be advantageous to the DEWSAR TB or not would depend on the goal of the deployment. If, for example, the goal of the deployment was to optimise the coverage of the sensor network, then distributing the Reporting Process would place an extra burden on the network which would reduce the coverage of the sensors and the endurance of the platforms during the deployment. If, however, the goal of the deployment was to optimise the speed of detection, distributing the Reporting Process would be advantageous. If the Reporting Process was to be distributed, it is likely that a mixture of the data pull, subscription based and localised broadcast paradigms¹⁰ would be required.

4. The Scalability of Wireless Networks

Section 2.6 stated that de-centralising the DEWSAR TB's control functions or distributing the test bed's data, information or data processing algorithms would improve the scalability of the test bed. Using a mesh topology would mean that a larger number of nodes could be employed, increasing geographic and RF coverage and providing potential for redundancy. The increased geographic and RF coverage could improve target detection rates and the redundancy could improve the overall robustness of the system. However, perhaps the most important limitation of wireless networks is their scalability. This section discusses the scalability of wireless networks in further detail.

The scalability of a sensor network establishes an upper limit of the number of sensor nodes which can be deployed in the network; if sensor nodes cannot exchange information, they cannot work cooperatively to identify, geolocate and track targets.

Given the network's communication channel bandwidth, Equation 1 (Gupta and Kumar 2000) can be used to establish the upper limit of the number of sensor nodes for an ad hoc wireless network.

$$\lambda(n) = \Theta\left(\frac{W}{\sqrt{n \log n}}\right) \quad (1)$$

¹⁰ In the localised broadcast paradigm, nodes transmit information only to their neighbours. The (physical or network hop) distance between neighbouring nodes is used as a heuristic to determine data relevance. The localised broadcast paradigm is simpler to implement than the subscription based paradigm and has reduced network resource requirements. However, being heuristic based, the localised broadcast paradigm may lead to information not reaching nodes which most require it.

In this equation, $\lambda(n)$ is the maximum throughput obtainable by each node for a randomly chosen destination (under a non-interference protocol), n is the number of identical, randomly located nodes in the wireless network and each of these nodes is capable of transmitting at W bits per second using a fixed range. Equation 1 indicates that the throughput available to each node approaches zero as the number of nodes increases, which implies that ad hoc wireless sensor networks are fundamentally non-scalable. However, as argued by Li, Blake and D. D. Couto et al (Li, Blake and D. D. Couto, et al. 2001) this result may not reflect reality for a number of reasons. Firstly, Equation 1 assumes a random communication pattern, which is a reasonable assumption for small networks. However in larger networks, nodes may communicate mostly with other nodes which are physically close leading to locally focused communications. And if local communication predominates, Li, Blake and D. D. Couto et al argue that communication path lengths could remain practically constant as the network grows in size, which would lead to a constant throughput being available per node. Li, Blake and D. D. Couto et al performed a number of experiments with real hardware. Figure 5 shows some of their results. It shows the throughput achieved along a chain of communication nodes as a function of the length of the chain where each node was placed at the maximum distance from the previous node that allowed for low-loss communications. Figure 5 shows that there is a significant reduction in network resources as the number of nodes in the network increases. The scalability of wireless networks has been examined extensively. Several strategies for improving the scalability of wireless networks have been developed including the hierarchical approach (Zhao, Seskar and Raychaudhuri 2004), the location centric approach (Ramanathan 2002) and the information centric approach (Goodman, Seed and Kiefer 2004). These three approaches aim to minimise network communications and hence improve the performance and scalability of ad hoc wireless networks. The next three subsections discuss these approaches in further detail.




Figure 5. The results obtained by Li, Blake and D. D. Couto et al (Li, Blake and D. D. Couto, et al. 2001) using real hard ware. The plot shows the throughput achieved along a chain of communication nodes as a function of the length of the chain. In Li, Blake and D. D. Couto et al's experiment, each node was placed at the maximum distance from the previous node that allowed for low-loss communications.

Given that there is a minimum required throughput that a sensor network requires in order to achieve its goals, the rapid degradation of throughput effectively limits the size of the network. Depending on the communication channel technology used, a usable network size may be less than ten nodes. For DEWSAR TB, the demand on network resources will also increase with the number of detections that is expected to result from increased network coverage.

4.1 The Hierarchical Approach

The hierarchical approach utilises three different classes of wireless devices in the network; sensor nodes, forwarding nodes and access points. The sensor nodes are the lowest tier in the hierarchy; they are low power nodes with limited functionality which route packets via the higher tier nodes. The forwarding nodes are the middle tier in the hierarchy; they are high power nodes which offer a multi-hop routing capability to nearby sensor or forwarding nodes. The access points are the highest tier in the hierarchy; they provide multi-hop routing for packets from nearby sensor or forwarding nodes in addition to routing data to and from other networks.

While simulations indicate that the hierarchical approach offers substantial improvements to the performance and scalability of ad hoc wireless networks (Zhao, Seskar and Raychaudhuri 2004), there would be a number of problems implementing this approach in the DEWSAR TB. Two of the key aspects of the hierarchical approach are that the network nodes are not homogeneous and the physical topology of the network is constant. In the DEWSAR TB, however, it is preferable that the network nodes are homogeneous in terms of network capability and the test bed's sensor nodes are often mobile and hence the physical topology of the network is not constant. If the hierarchical approach were implemented in the DEWSAR TB, the steering algorithms used to position the sensors deployed on the aerial platforms would need to take network capability and topology into account. This would substantially increase the complexity of the steering algorithms. In addition, during a deployment, the requirement of the steering algorithms to keep the sensor network in a functional topology (e.g. the requirement to keep the sensor nodes within transmitting distance of the forwarding nodes) may compete with the mission objectives (e.g. geolocation of the targets).

A modified version of the hierarchical approach could be implemented in the DEWSAR TB where sensor nodes could be dynamically assigned different networking tasks as required. This would require the development of an algorithm to dynamically perform the assignment. This algorithm would be quite complex as it would need to reason about the continuous changes in the physical network topology and the corresponding information routing requirements, the physical capabilities of the sensor platforms, the prevailing environmental conditions in addition to ensuring that the mission objectives are efficiently met. Whilst the approach is quite complex, numerous wireless protocols have been implemented with this approach. However for the DEWSAR TB¹¹ implementing such an approach would be technically difficult and bring few benefits.

4.2 The Location Centric Approach

In a conventional node-centric approach, collaboration and information exchange is between an arbitrarily specified set of devices. Even if these devices move, the collaboration typically continues between the same set of devices. In contrast, the location centric approach is based on the premise that sensor networks typically require collaboration among devices in a certain geographic area. Thus a device begins/ceases to participate in an ongoing collaboration if it enters/leaves the corresponding defining region. The location centric approach makes use of regions, instead of individual devices, as addressable entities.

As mentioned in Section 3.6, the geometry of the sensor nodes in the DEWSAR TB can play an important role in the geolocation of targets. The test bed's steering agent is used to ensure that the target is observed from multiple, sufficiently different angles in order to reduce triangulation error and minimise the uncertainty in the target geolocation. Consider the situation where there are several of the test bed's sensor nodes within

¹¹ Note that, in a sense, the DEWSAR TB currently uses a form of hierarchical topology; the DEWSAR TB currently employs a star topology, where the central node (i.e. the ground station) is more capable than the nodes.

relatively close proximity to the target and (using the location centric approach) these nodes are all assigned to the same sensor region. While these sensor nodes are relatively close to the target, they may have a poor geometry relative to the target. The resulting target geolocation may be inferior to the geolocation obtained had the sensors collaborated with sensors outside their region which, although they may be further from the target, may have a better relative geometry. Thus while implementing a location centric approach in the DEWSAR TB may bring data exchange benefits, sensor nodes within a particular geographic region may not provide the best geometry for target geolocation. Thus in order for the location centric approach to be beneficial for the DEWSAR TB, sensors must be able to collaborate with sensors outside their region and the data exchange between such regions must be optimised.

4.3 The Information Centric Approach

The information centric approach formulates the distribution of data amongst sensors in a network, with a limited communication capacity, as a constrained optimization problem; the problem is one of minimizing the cost of sensor communications while simultaneously achieving a desired accuracy (information threshold). Each sensor chooses, not only what data to share, but to which other sensors to shares the data with by comparing the cost of the communication against the benefit of sharing the data.

Goodman et al (Goodman, Seed and Kiefer 2004) used an information centric approach to autonomously form sensor subnets for ground based target tracking. Using ground based and airborne radars (which could individually resolve a target location), Goodman et al's algorithm was able to optimise the information gain obtained by sharing information between nodes to attain a desired track accuracy whilst minimising network utilisation. Goodman et al achieved an order of magnitude reduction in network bandwidth usage whilst sacrificing little in overall target detection accuracy. However, the ground based and airborne radars used in their network are active sensors. The DEWSAR TB, however, employs passive ES sensors. The data provided by such passive sensors makes an information centric approach difficult; the information required to make collaboration decisions (for example: whether to share information in order to geolocate a target) may not be available until some communication has taken place.

4.4 Networks of Networks

It can be seen that the general strategy employed by the hierarchical, location centric and information centric approaches is to create subgroups of sensors based on a criterion. Thus in essence, these approaches create networks of networks; rather than using a large number of sensors simultaneously these approaches select smaller subsets of sensors based on a criterion. This small sensor subset then operates autonomously without exceeding the limitations of the local wireless network. The sensor groups then exchange higher level information products between themselves at a lower data rate in order to achieve broader goals. Such a 'networks of networks' strategy facilitates scalability. These networks of networks could be based on any combination of the hierarchical, location

centric or information centric approaches. Furthermore, the structure of a network of networks may be virtual and adaptive.

When addressing the question of how many sensor nodes should be deployed in a network, it is important to bear in mind that the efficiency of information exchange decreases as the number of sensors increases. Thus the number of sensors in a wireless network should be minimised in order to optimise the efficiency of information exchange. However, this must be balanced with the ability of the network to accomplish its mission. For example, the mobility of the aerial platforms will play significant role in the geolocation of targets in networks where a small number of sensors are used to cover large geographic areas. Using a small number of sensor nodes in the network means that the network may take longer to geolocate a target than a network with a larger number of sensors; the platforms may have to travel significant distances in order to decrease the distance between, or improve the relative geometry of, the sensors and the target.

Sadaphal and Jain (Sadaphal and Jain 2005) investigated the optimal sensor density required to maximise the accuracy of an estimate of target position. They showed that an

optimum node density of $\rho = \frac{8}{\pi r_0^2}$, for a network of homogeneous sensors each with a

detection range of r_0 , results in the maximum position estimation accuracy. They also showed that increasing the density beyond ten neighbours provided marginal gains in accuracy. This result allows us to draw the conclusion that, although the DEWSAR TB differs from the work of Sadaphal and Jain in that the sensors are non-homogeneous rather than homogeneous, a DEWSAR TB network should be relatively small, with possibly no more than eight nodes for a given geographic area. Using more than eight nodes will not significantly improve the accuracy of target geolocations and will suffer from significantly decreased network efficiency as shown by Gupta and Kumar (Gupta and Kumar 2000) and Li, Blake and D. D. Couto et al (Li, Blake and D. D. Couto, et al. 2001). It must also be borne in mind that, because energy is a finite resource and because communication is energy intensive, using a larger number of nodes may also decrease the endurance (aka deployment time) of the network. At a minimum, the DEWSAR TB requires a minimum of two to three sensors to geolocate a target.

In practise, the number of sensor nodes in a network will be constrained by a number of factors, such as:

- the objective of the mission
- the desired detection accuracy
- the time available to achieve the desired detection accuracy
- The capabilities of the sensors including their mobility. A sensor network may benefit from non homogenous mobility behaviour of individual nodes. For example, some sensors may be used to provide coverage “hot spots” by exhibiting higher mobility than the average node. This means the sensor network could proactively increase sensor density in areas where targets are most likely to appear or reactively pursuing targets which do not have sufficient sensors nearby.

- the capabilities of the targets including their mobility
- Environmental factors such as terrain, weather conditions and zone restrictions (such as no fly zones for the aerial platforms). Different environments have varying emitter characteristics and therefore will require modification to sensor numbers. Land and sea environments typically don't have fast moving vehicles whereas air environment does (thus air environments may require a larger number of sensors). Sea and air will typically have unobstructed visibility, however land based geographic features can create obstructions for RF signals.
- the capacity of the communications network
- The need to exchange information between sensors within the sensor group and between other networks of sensors. Depending on network topology, information may need to be relayed over multiple hops.

5. Conclusions

The DEWSAR TB is a test bed designed to enable the rapid development and evaluation of DEW concepts by utilising a networked set of heterogeneous, relatively unsophisticated EW sensors and effectors (deployed on ground based or aerial platforms) to detect, identify, locate, track or suppress stationary or slow moving surface based RF emitting targets. In the current DEWSAR TB, little of the control or data processing occurs at the (ground based or airborne) nodes. For example, data from the sensors is transmitted from the node to the ground station where it is collected, processed, stored and disseminated as appropriate. This document evaluated a number of options for de-centralising the test bed's Tuning and Steering control functions and distributing the Classification, Aggregation, Geolocation, Tracking and Reporting processes. The evaluation was performed within the context of the purpose of the test bed, the size, weight and power constraints of the Aerosonde aerial platforms and the communication bandwidth limitations.

It was found that it is not possible to make a judgement about distributing the Classification Process without information about the taxonomy and density of the radar target environment, which in a real world deployment may not be available. Whether distributing the Reporting Process would be advantageous to the test bed depends on the goal of the deployment. De-centralising the test bed's Tuning and Steering functions or distributing the Aggregation, Geolocation and Tracking processes would substantially increase the complexity of the test bed. Although de-centralisation and distribution is likely to improve the scalability of the test bed (and through improving scalability, increase the coverage, robustness and accuracy of the system), the increased complexity of the test bed¹² may result in increased development time and cost and increased

¹² Much of the increased complexity is in the implementation of the wireless networking systems. Commercial of the shelf systems were considered but not implemented due to the costs and risks associated with this sort of technology. While it is relatively easy to acquire the radio technology,

maintenance overheads. Higher complexity could also increase the test bed's technical risk by increasing the likelihood and complexity of future technical problems. Any de-centralisation or distribution option which impinges a researcher's ability to rapidly develop and evaluate algorithms or concepts contravenes the purpose of the test bed. Thus, at this stage, the benefits afforded by de-centralising the test bed's control functions or distributing the test bed's data processing do not offset the increased development and maintenance costs and increased technical risk.

Until recently the DEWSAR TB has been primarily focussed on the collection and generation of EW information of radar targets. The target set of the DEWSAR TB now includes communications emitters. The biggest difference between communication and radar emitters is that radar emitters typically have a much higher output power (that is, the standoff distance for the DEWSAR TB's sensor nodes can be larger). The new focus on communications targets has generated a requirement for new sensors and nodes in the test bed. This document also considered architectures required to network the sensors, and manage and process the information and discussed the scalability of the test bed for larger numbers of nodes.

Several strategies for improving the scalability of wireless networks have been developed including the hierarchical, location centric and information centric approaches. Each of these approaches has drawbacks for the DEWSAR TB. Rather than attempt to de-centralise/distribute for a single large network, a better strategy to facilitate the scalability of the test bed may be the 'networks of networks' strategy where the structure of a network may be virtual and adaptive. But the test bed must have enough sophistication to deal with its environment. The ultimate decision between de-centralisation/distribution of a single large network and the network of networks approach depends on balancing the demands of the environment, network bandwidth and the SWAP constraints.

When addressing the question of how many sensor nodes should be deployed in a network, it is important to bear in mind that the efficiency of information exchange decreases as the number of sensors increases. Thus the number of sensors in a wireless network should be minimised in order to optimise the efficiency of information exchange. In practice, the number of sensor nodes in a network will be constrained by a number of factors such as the objective of the mission, the capabilities of the sensors and the complexity of the environment. However, the DEWSAR TB should have around eight sensor nodes in the network; increasing the number of sensor nodes above eight significantly constrains the capacity of the network and does not provide increased accuracy for a coverage area.

the network behavior and protocols need to be developed so that they are fit for purpose. Also existing WiFi mesh topology networks have typically been designed for low mobility applications, not high mobility applications such as the DEWSAR TB's UAV network.

6. Acknowledgements

The authors wish to thank Dr Ian Grivell and Dr Iain Jameson for their constructive comments on the preparation of this document.

7. References

- Baldock, A., S. P. Drake, C. Howard, et al. (2009). Evaluation of the Performance of the Discrete Probability Density Method using Data from Trident Warrior 2009 **TR-2524**.
- Brooks, R. R., P. Ramanathan and A. M. Sayeed (2003). "Distributed Target Classification and Tracking in Sensor Networks." Proceedings of the IEEE **91**(8): 1163 - 1171.
- Brown, K., A. Bailey, S. Drake, et al. (2005). A Demonstration of Distributed Sensing and Situational Awareness Systems. Land Warfare Conference, Gold Coast, Queensland, Australia.
- Devries, C. (2007). Mk 4.5 Aerosonde AV Technical Manual.
- Division, N. A. W. C. W. (1999). Electronic Warfare And Radar Systems Engineering Handbook, Revision 2. Point Mugu, California, USA.
- Dogancay, K. and H. Hmam (2008). "Optimal angular sensor separation of AOA localization." Signal Processing **88**(5): 1248-1260
- Elsaesser, D. (2006). The Discrete Probability Density Method for Emitter Geolocation. Canadian Conference on Electrical and Computer Engineering., Ottawa, Canada.
- Elsaesser, D. (2007). Sensor Data Fusion Using a Probability Density Grid. Proceedings of the 10th International Conference on Information Fusion, Quebec, Canada.
- Finn, A., K. Brown and T. Lindsay (2002). Miniature UAV's and Future Electronic Warfare. Land Warfare Conference, Brisbane, Queensland, Australia.
- Gibard, J.-P., K. Kabacinski, S. Drake, et al. ((in preparation)). Distributed Electronic Warfare Situation Awareness and Response Testbed 2009 Trial Summary and Initial Results.
- Goodman, J., R. Seed and P. Kiefer (2004). Dynamic Subnets for Sensor Networks. Proceedings of the IEEE Sixtieth Conference on Vehicular Technology Los Angeles, California, USA.
- Gupta, P. and P. R. Kumar (2000). "The Capacity of Wireless Networks." IEEE Transactions on Information Theory **46**(2): 388-404.
- Izydorczyk, J. and L. Cionaka (2008). A Practical Low Limit on Energy Spent on Processing of One Bit of Data. Proceedings of the Eighth International IEEE Conference on Computer and Information Technology Workshops. Sydney, New South Wales, Australia.
- Li, J., C. Blake, et al. (2001). Capacity of Ad Hoc Wireless Networks. Proceedings of the 7th Annual International Conference on Mobile Computing and Networking. New York, NY, USA.
- Mason, K., K. Brown and K. Kabacinski (2009). Coalition Information Exchange at the 2004 and 2005 C4ISR On The Move Experiments. DSTO-TN-0873.

- Olfati-Saber, R. (2007). Distributed Tracking for Mobile Sensor Networks with Information-Driven Mobility. Proceedings of the American Control Conference 2007, New York, USA.
- Ramanathan, P. (2002). Location-centric Approach for Collaborative Target Detection, Classification, and Tracking. Proceedings of the IEEE CAS Workshop Wireless Communications and Networking, Pasadena, California.
- Sadaphal, V. and B. Jain (2005). Localization Accuracy and Threshold Network Density for Tracking Sensor Networks. Proceedings of the IEEE International Conference on Personal Wireless Communications.
- Zhao, S., I. Seskar and D. Raychaudhuri (2004). Performance and Scalability of Self-Organizing Hierarchical Ad Hoc Wireless Networks. Proceedings of the IEEE Wireless Communications and Networking Conference.

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA				1. PRIVACY MARKING/CAVEAT	
2. TITLE Evaluating De-centralised and Distributional Options for the Distributed Electronic Warfare Situation Awareness and Response Test Bed			3. SECURITY CLASSIFICATION (NEXT TO DOCUMENT CLASSIFICATION) Document (U) Title (U) Abstract (U)		
4. AUTHOR(S) Kuba Kabacinski, Catherine Howard and Damian Hall			5. CORPORATE AUTHOR Defence Science and Technology Organisation PO Box 1500 Edinburgh South Australia 5111 Australia		
6a. DSTO NUMBER DSTO-TR-2918		6b. AR NUMBER AR-015-804		6c. TYPE OF REPORT Technical Report	
				7. DOCUMENT DATE December 2013	
8. FILE NUMBER -		9. TASK NUMBER -	10. TASK SPONSOR -		11. NO. OF PAGES 29
				12. NO. OF REFERENCES 19	
13. DOWNGRADING/DELIMITING INSTRUCTIONS To be reviewed three years after date of publication			14. RELEASE AUTHORITY Chief, Electronic Warfare and Radar 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-vic.dsto.defence.gov.au/workareas/library/resources/dsto_thesaurus.shtml Wireless sensor networks; Sensor systems; Electronic warfare systems; Test beds					
19. ABSTRACT The Distributed Electronic Warfare Situation Awareness and Response Test Bed was designed to rapidly develop and evaluate distributed Electronic Warfare concepts through a networked set of heterogeneous, relatively unsophisticated sensors and effectors (deployed on ground based or aerial platforms) to detect, identify, locate, track or suppress stationary or slow moving surface based RF emitting targets. In the current test bed, little of the control or data processing occurs at the ground based or airborne nodes. This document evaluates a number of options for de-centralising the control functions and distributing the data processing from the ground station to the ground based or airborne nodes.					