15th ICCRTS

“The Evolution of C2”

June 22-24, 2010
Fairmont Miramar Hotel & Bungalows
Santa Monica, CA, U.S.A.

The Development of Advanced Recognition Concepts for the HALIFAX Class Command and Control System

TOPIC: Experimentation and Analysis

Bruce McArthur
Defence R&D Canada – Atlantic
P.O. Box 1012, 9 Grove Street, Dartmouth, Nova Scotia, Canada B2Y 3Z7

Sharon McFadden
Defence R&D Canada – Toronto
P.O. Box 2000, 1133 Sheppard Avenue West, Toronto, Ontario, Canada M3M 3B9

Bruce Chalmers
Defence R&D Canada – Atlantic
P.O. Box 1012, 9 Grove Street, Dartmouth, Nova Scotia, Canada B2Y 3Z7

POINT OF CONTACT: Bruce McArthur
Defence R&D Canada - Atlantic
P.O. Box 1012, 9 Grove Street, Dartmouth, Nova Scotia, Canada B2Y 3Z7
Phone: (902) 426-3100 (x359); Fax: (902) 426-9654
E-mail: bruce.mcarthur@drdc-rddc.gc.ca

The Development of Advanced Recognition Concepts for the HALIFAX Class Command and Control System

Bruce McArthur, Bruce Chalmers

Defence R&D Canada - Atlantic
Dartmouth, Nova Scotia, Canada
(902) 426-3100

bruce.mcarthur@drdc-rddc.gc.ca,
bruce.chalmers@drdc-rddc.gc.ca

Sharon McFadden

Defence R&D Canada - Toronto
Toronto, Ontario, Canada
(416) 635-2189

sharon.mcfadden@drdc-rddc.gc.ca

Abstract

In modern shipboard C2 systems, the traditional tactical display is being replaced by the Maritime Tactical Picture (MTP) that integrates the output of sensor sources with geographic information, contact attribute data, and the wide area picture. The recognition process, whereby input data are processed to determine a contact's identity within a classification hierarchy, is one of the key operational processes involved in the compilation of the MTP. The current paper describes challenges encountered and processes used for the design, development and evaluation of advanced recognition concepts for the HALIFAX Class frigate, developed under Command Decision Aid Technology (COMDAT), a Defence R&D Canada Technology Demonstration Project. Aspects of the work discussed in the paper include: the development of models of the information flows and decision processes used by the HALIFAX Class operations room team in performing contact recognition tasks; an assessment of where data fusion technology might provide the most effective support to operators; automated recognition capabilities based on the truncated Dempster-Shafer fusion of attribute data from ownship and remote data sources; user interface concepts and a concept of operations for operator involvement with automated recognition processes; and, measures and methods for assessing operator and system performance in carrying out recognition tasks.

1. Introduction

Shipboard command teams perform a continuous sequence of data and information processing tasks which together constitute a tactical decision loop. These tasks include: the integration of data from the ship's sensors and other sources to compile a tactical picture of the situation; monitoring the tactical picture to assess threats; the identification and selection of courses of action in response to anticipated or actual threats to the mission, in accordance with rules of engagement; and action implementation once a decision to act has been made and is being carried out. In modern shipboard C2 systems, the traditional tactical display is being replaced by the Maritime Tactical Picture (MTP) that integrates the output of sensor sources with geographic information, contact attribute data, and the wide area picture. A timely and accurate integrated MTP is essential to provide the command team with as complete an understanding of the situation and threats as possible.

As part of the Command Decision Aid Technology (COMDAT) Technology Demonstration Project, Defence R&D Canada (DRDC) carried out the demonstration of data fusion technologies and associated concepts for the development of an integrated Above Water Warfare (AWW)

Maritime Tactical Picture for the Canadian Navy's HALIFAX Class frigate. The project focused on two key operational processes involved in the compilation of the MTP, tracking and recognition, and drew upon DRDC research programs in multi-source data fusion (MSDF) and human factors. The work conducted under COMDAT included research studies, the implementation of a sea-going testbed, as well as ship and shore-based experimentation.

Using NATO terminology, the *recognition* process [12] involves the interpretation of data from sensor and information sources to determine the characteristics of a contact, which are compared against reference data. Contact characteristics or attributes may include physical attributes, such as size, dynamic attributes, such as course, speed, or altitude, relational attributes such as movement history relative to geographic features or other platforms, emission attributes such as detected radar, radio, or acoustic emitters, as well as recognition attributes, such as platform type or flag. A contact's recognition is a declaration about its identity within a classification hierarchy and normally includes an associated level of confidence. Recognition can range from general (for example, surface platform) to specific (ship name) and is usually determined both by operational requirements and by data availability.

A process closely related to recognition, *identification* [12], includes the assignment of a standard identity to a contact (from the set: hostile, suspect, unknown, neutral, assumed friend, friend). Identification was addressed to a lesser extent as part of the COMDAT project but is not a focus of this paper.

The multi-source data fusion community has developed a variety of techniques (set theoretic, probabilistic, evidential reasoning, rule-based) that have been applied to aspects of the recognition process in the military domain. These techniques can be broadly categorized as normative (recognition decisions are based on a set of formal criteria) or descriptive (attempts to emulate human decision making) [1]. Prior to the initiation of the COMDAT project, DRDC, together with Canadian industry, had conducted research demonstrating the potential utility of Dempster-Shafer evidential reasoning algorithms (a normative technique) for automated recognition in the maritime domain (e.g. [19, 15]). These algorithms work in conjunction with other MSDF system functions for data association and kinematic fusion.

In contrast, relatively little work had previously been conducted on how best to integrate technologies, such as MSDF, into the shipboard environment. Although the algorithms had been tested extensively, most of the evaluations involved simulated sensors in a relatively benign environment. Also, at the time that the COMDAT project was initiated, the recognition and identification processes were carried out almost entirely by the human operator. Thus, it was anticipated that the Navy could face several challenges in an attempt to transition these automated recognition algorithms, which involved little or no human interaction in their processing, from the laboratory to a more realistic shipboard Command and Control system environment, including:

- understanding how the recognition algorithms would enhance the decision processes used by operations room personnel;
- the integration of the technology into the work processes of operational personnel carrying out recognition tasks so as to gain operator acceptance;
- the requirement for the algorithms to handle input data that was often incomplete, unreliable, ambiguous, or conflicting; and
- the limited availability of facilities and military test subjects to evaluate the technology using realistic scenarios.

The following sections of this paper describe in more detail the challenges encountered and processes used for the design, development and evaluation of advanced recognition concepts developed under the COMDAT project. Aspects of the work discussed include: models of the information flows and decision processes used by the HALIFAX Class operations room team in performing contact recognition tasks; an assessment of where data fusion technology might provide the most effective support to operators; automated recognition capabilities based on the truncated Dempster-Shafer fusion of attribute data from ownship and remote data sources; user interface concepts and concept of operations for operator interaction with automated recognition processes; and, measures and methods for assessing operator and system performance in carrying out recognition tasks.

2. Models of the information flows and decision processes used by the HALIFAX Class operations room team in performing recognition tasks

For reasons such as the complexity of the shipboard C2 environment, technological limitations and the need to adopt a human systems integration perspective, MSDF technologies must, for the foreseeable future, be integrated to assist, not replace, operators in performing their picture compilation processes. This required conducting work in COMDAT to develop an in-depth understanding of those processes. In addition to consulting current doctrine on picture compilation processes, knowledge elicitation sessions were conducted over two days with two sub-teams of key operations room personnel from the command team and various sensor operators.

Subject Matter Expert (SME) interviews focused on Above Water Warfare (AWW). Knowledge elicitation interviews were conducted in three phases:

1. A structured interview with a number of pre-determined but open-ended questions surrounding the detect-to-recognize (D-R) process for contacts.
2. A structured interview in which the SMEs were presented with a list of contact types and asked to describe the process by which the team might detect the contact and subsequently arrive at a recognition or identification.
3. A case study approach, using the Critical Decision Method (CDM) [7], that was intended to focus the team on one critical incident in order to gain more insight into problem solving and decision making for non-routine cases.

These interviews helped to highlight situations where recognition and identification (R/ID) are more challenging for operators. Examples include multi-threat environments, in which the team must regularly switch their focus of attention as they build both the air and surface pictures simultaneously, and complex surface picture scenarios. SMEs specifically described the difficulties in regaining awareness of the ‘big picture’ following a period of focusing on a particular contact, such as recognizing and identifying an enemy that is hiding amongst a fishing fleet. Discussion in the second of the three phases in the knowledge elicitation interviews also provided useful insights into the potential utility of data fusion technology in operator D-R processes (see Section 4). The SME data was used to conduct two types of analysis, a descriptive analysis and a formative analysis. We now discuss these two types of analysis.

In the descriptive analysis, data collected from the knowledge acquisition process was documented in a series of tables, and/or flow diagrams, for D-R event sequences for specific critical types of contact. The general intent of this approach was to tap into many different dimensions or aspects of these event sequences, including:

- The timeline
- The flow of data
- Communications
- Decision points
- Actions/tasks of different operations room team members

- Information used
- Sources of information and their associated uncertainty
- Operator strategies
- Operator assessment of potential utility and role of data fusion functionality.

The different graphical techniques for representing the data compiled included: information flow and processing analysis diagrams [13] to illustrate the sequential steps involved in decision making (e.g., what options are considered, who makes the decision, and what are the consequences); timeline diagrams to depict the sequence of tasks performed by team members along a common time base; and flow sequence diagrams of critical D-R tasks for key operator positions in a HALIFAX Class operations room.

Although the various descriptive analyses provided the basis for a user-centred investigation of ways to incorporate data fusion technologies into current operations room D-R processes, a second approach, based on a formative analytical approach, was also implemented. The goal of the formative analysis was to model R/ID processes from a perspective that would provide designers with a broader functional perspective of how R/ID *could* be done (as opposed to how it *is* currently done), thereby providing a design lens that would not be unnecessarily clouded by the current process views. One analytical framework that adopts that particular perspective is Cognitive Work Analysis (CWA) [13]. CWA focuses on analyzing constraints that shape work, thereby delimiting both current work trajectories, as well as work trajectories that are difficult to specify *a priori* (e.g., as a result of integrating new technologies in the work environment). Although, a full CWA would have involved conducting five separate phases of analysis, we concentrated on one specific phase, a Control Task Analysis, with the specific goal of developing a formative model of the cognitive activities involved in performing R/ID processes.

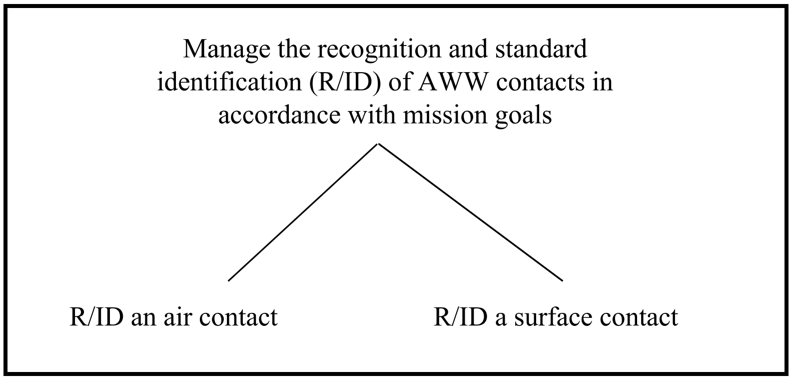

Figure 1. Functional hierarchy of R/ID functions.

Figure 1 shows the functional hierarchy used to separate out the information processing activity related to recognizing and identifying individual contacts from the information processing activity related to managing all of the streams of R/ID activity in accordance with mission goals. This latter type of information processing activity would be associated with any of the following: prioritizing, timesharing, or scheduling the R/ID processing of individual contacts; changing processing focus among individual contacts or warfare areas; dealing with multiple contacts in one or more AWW warfare areas that compete for R/ID processing; or adapting the R/ID activity of individual contacts based on factors such as context, a contact's tactical importance, workload, and requests or commands from superiors. The representation of the R/ID of air and surface contacts as separate activities reflects the greater complexity of surface picture compilation relative to air picture compilation as was highlighted in the knowledge elicitation interviews. Analysis of the recognition and standard identification processes for air and surface contacts was based on developing information processing activity models for each of the node functions of the functional hierarchy shown in Figure 1.

The specific information processing activity models developed were three decision ladder (DL) models, one for the function “Manage R/ID of AWW contacts in accordance with mission goals”, and one for each of the two sub-functions: “R/ID an air contact” and “R/ID a surface contact”, shown in Figure 1. The generic DL model template, shown in Figure 2, was used to parse the information processing associated with a function according to an opportunistic, contextually tailored collection of data/information processing activities (represented by rectangles) and states of knowledge (represented by circles or ellipses). The template provides for eight types of information-processing activity and a corresponding number of types of states of knowledge. It also supported the representation of some aspects of expertise or adaptive information processing, associated with a work function, that are the result of characteristics of the decision-making context in which the function is performed (e.g., time pressure, uncertainty, information availability, information quality). This is done by indicating various shortcuts from one part of the ladder to another, including *shunts* that connect an information-processing activity to a state of knowledge, and *leaps* that connect two states of knowledge.

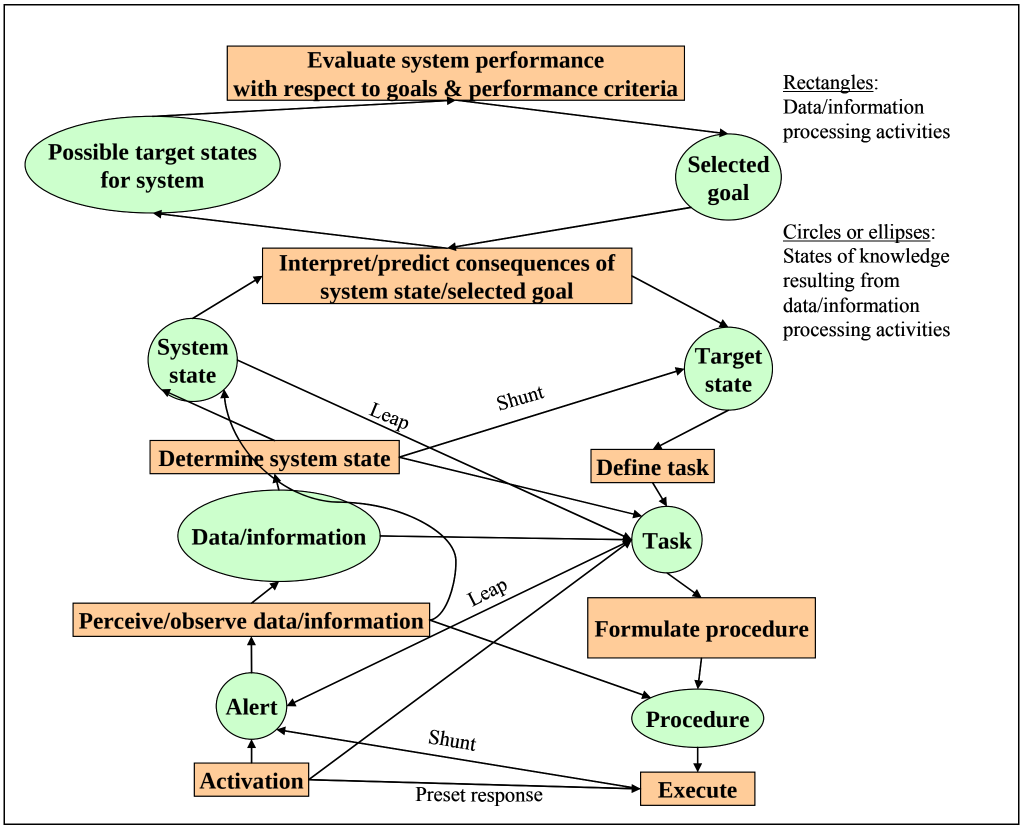

Figure 2. Generic Decision Ladder template (Adapted from [21].)

The formative modeling approach provided a very coherent and detailed cognitive description of the recognition and identification processes, including the management of those processes, and specific information used at each step of the process, together with the sources of that information. The analyses showed in a generic sense which processes are required and their associated decisions, as well as which could be skipped over (shortcuts) in conducting R/ID. The formative nature of the analyses permits developing an understanding of what *generic support* could be provided in the future to enhance the HALIFAX Class Command and Control system.

However, for detailed system design work, the Control Task Analysis will need to be augmented by additional types of work analyses (e.g., a formative organizational analysis to identify new ways of allocating information processing roles and knowledge states in the DLs across the operations room team) as design options are considered for augmenting capability by incorporating data fusion technology.

3. Automated recognition capabilities based on the truncated Dempster-Shafer fusion of attribute data

The implementation of an automated recognition capability within the COMDAT testbed's MSDF subsystem is based on a computationally-efficient variant of Dempster-Shafer (D-S) evidential reasoning, termed *truncated Dempster-Shafer* [19, 15]. D-S evidential reasoning is a generalization of Bayesian subjective probability that supports the representation and combination of evidence from multiple sources. D-S has been widely used as a method for handling uncertain, incomplete and/or conflicting information based on characteristics that include: an explicit representation for ignorance; no requirement for prior probabilities; and, the assignment of evidence directly to sets. Using D-S, the recognition domain is defined in terms of:

- a *frame of discernment* Θ, comprised of n air and surface platforms, at a recognition level of ship class or aircraft type (for example, HALIFAX Class ship or F-18 aircraft), and stored in the MSDF subsystem's *platform database* (PDB);
- a set of *propositions*, $\{P_i\}$, each of which corresponds to a subset of the platforms in Θ, and which includes the null proposition $\varnothing$ (representing the set of no platforms) and the *ignorance* proposition Θ (representing the set of all platforms); and,
- a *basic probability assignment* (BPA) that assigns a *mass* $m(P_i)$ to each proposition, the mass representing the degree of belief accorded to a proposition.

For each track processed by COMDAT's MSDF subsystem, D-S is used to maintain a set of propositions and associated masses, which together represent the track's recognition estimate. In theory, D-S may generate each of the 2^n possible propositions in the frame of discernment (which for even a modest database of 100 platforms would exceed 10^{30} propositions). However, in order to achieve a feasible implementation of a real-time automated recognition capability, only a user-defined maximum number of propositions is maintained (typically eight, in the case of COMDAT), using the truncated D-S algorithm. This reduced set of propositions and associated masses is termed a *proposition set*. More detailed descriptions of the truncated D-S algorithm and D-S evidential reasoning are presented in [15, 19, 3, 20].

The COMDAT testbed processes data from a subset of available ownship sensors (medium-range and long-range radars, IFF, ESM) and remote data sources (Link-11 and Global Command and Control System – Maritime (GCCS-M)). A representative set of data attributes that are input from these sensors and sources is summarized in Table 1. The attributes Type, Ship Class, and Country are recognition attributes, with Type encompassing recognition levels from a general category (e.g., air, surface) to a more specific platform type (e.g., bomber, destroyer) and Ship Class (e.g., HALIFAX, AEGIS) being at the recognition level of the frame of discernment. Allegiance is an identification attribute (e.g., Friend, Neutral), with IFF and ESM sensors providing values based on contact measurements and Link-11 and GCCS-M providing off-board identification assessments. Frequency and Speed differ from other attributes in that they are continuous-valued, rather than having discrete values. The Speed attribute is derived from MSDF track kinematic estimates and the source for this attribute is indicated accordingly as 'MSDF' in Table 1. The COMDAT testbed, as a concept demonstrator, is not designed to process all of the operational inputs to the recognition process, including, for example, relational attributes. Implications of this limitation are discussed in Sections 4 and 5.

Source	**Type**	**Ship Class**	**Allegiance**	**Country**	**Emitter**	**Freq**	**Speed**
Radar	X						
IFF			X				
ESM	X		X		X	X	
Link-11	X		X			X	
GCCS-M	X	X	X	X		X	
MSDF							X

Table 1. Summary of attribute data processed by COMDAT's MSDF subsystem.

The three primary functions that make up COMDAT's automated recognition capability are identity alignment, identity update, and identity output. *Identity alignment* converts raw attribute data from input reports into D-S proposition sets, termed *input proposition sets*. For an input attribute with discrete values, a typical input proposition set has two elements: a proposition corresponding to the attribute's value, with a mass based on a general assessment of the accuracy and reliability of the sensor; and an ignorance proposition that is assigned the remainder of the evidence. As an example, the attribute value Type = Air is translated into an 'Air' proposition which represents the set of PDB platforms which are of type 'Air'. The Air proposition might then be assigned a mass 0.8, and the remaining mass of 0.2 is assigned to ignorance. The result is the input proposition set {(Air, 0.8), (Θ, 0.2)}. Continuous-valued attributes, such as speed and frequency, are first converted to discrete values using a fuzzy set-based approach.

Identity update updates an MSDF track's proposition set with evidence provided by input proposition sets, using Dempster's rule of combination. Given information from two sources of evidence, characterized by BPAs m_1 and m_2, Dempster's rule defines the new BPA m_3

$$m_3(P_k) = \frac{1}{1-\kappa} \sum_{P_i \cap P_j = P_k} m_1(P_i) m_2(P_j)$$

where the conflict κ, which quantifies the degree to which two BPAs are contradictory, is defined by

$$\kappa = \sum_{P_i \cap P_j = \varnothing} m_1(P_i) m_2(P_j),$$

The factor 1/(1-κ) in Dempster's rule re-normalizes the BPA by redistributing the conflicting belief over the non-null propositions. Identity update may result in the creation of new propositions. Consequently, truncation rules are applied after each execution of Dempster's rule to control the size of the proposition set.

Identity output generates two categories of output data for use by the operator: *generic propositions*, which are analogous to the terms in a recognition hierarchy; and *specific propositions*, which are the elements of the proposition set. COMDAT generates generic propositions at two recognition levels: Category, from the set {Air, Surface, Subsurface}, and Sub-Category, from the set {Missile, Helicopter, Fighter, Bomber, Commercial Air, Carrier, Cruiser, Destroyer, Frigate, Fishing Boat}. The level of confidence in each generic proposition is calculated based on the D-S concept of Belief, defined as the sum of the mass assigned directly to a proposition and to all of its subsets. Several alternative display concepts used for displaying generic propositions are possible including: displaying all members of the proposition set along with associated belief values; displaying only the proposition having the highest belief, and its associated belief value; displaying only the proposition having the highest belief, and then only if its belief exceeds a threshold value. The latter concept is currently implemented in COMDAT.

The above functions are performed in conjunction with MSDF subsystem functions for data association and kinematic fusion. It is important to note that identity update assumes the correct assignment of input reports to MSDF tracks, as determined by the data association function.

One of the key challenges in implementing an automated recognition capability has been to provide accurate and reliable recognition estimates in spite of possibly conflicting and dependent input attribute data. Sources of conflicting data in the recognition process include spurious input reports that are inserted into the recognition process; inconsistencies between the attribute data contained in input reports and the content of the PDB; and basic errors in the attribute data received from input sources. An extreme example of conflicting data would be one sensor report that asserts a contact is an aircraft and a second sensor report that asserts the same contact is a surface ship. Dependent data violate one of the key assumptions of D-S theory, that each input proposition represents an independent piece of evidence. If dependencies are ignored, the accuracy of the beliefs assigned to competing recognition propositions may be affected. Independence assumptions can be violated by input reports that contain multiple attributes, but which may or may not be derived from independent measurements, and by multiple input reports, having different timestamps, but containing one or more of the same attribute values.

Errors in the data association process can also negatively impact recognition by inserting spurious and potentially conflicting attribute data - originating from contacts other than the contact undergoing recognition - into the identity update function. Conversely, the presence of conflicting recognition data can be used as an indication that information items do not originate from the same source (based on the assumption that if different input reports do originate from the same real-world object, and if that object's recognition is static, then the input reports' recognition-related attributes should be consistent). In COMDAT, the conflict κ is used as one of three constraints within the data association process (the other two constraints being a traditional kinematics-based likelihood function and a constraint based on track number). Using D-S conflict as an association criterion has proven to be beneficial in improving data association performance, particularly in situations such as closely spaced or crossing contacts, where kinematic constraints alone may be insufficient to make high-confidence association decisions.

In identity alignment, the translation of input attribute data to input propositions is achieved using the *a priori* knowledge of the expected attribute values for each platform, which is stored in the PDB. However, this process will generate conflict, in the form of null input propositions, if there are no PDB platforms that match an input report's attribute values. The underlying cause for this type of mismatch can lie with the input source (for example, if there are errors in the input data), with the translation between attribute values used in input reports and the attribute values in the PDB, or with the content of the PDB (for example, if the PDB is incomplete or contains incorrect data values). The translation between input report and PDB attribute values can be particularly complex, because it must be performed independently for each input source, and then can be complicated further by factors such as attribute values with alternate names (an example being the HALIFAX Class, which has the alternate name, Canadian Patrol Frigate).

The identity alignment process is further complicated in the case of input reports containing multiple attributes. If attributes are assumed to be dependent, as can be the case for input reports from remote data sources, it is appropriate to generate a single proposition set, based on the intersection of the individual attributes' values. This approach will generate null propositions if two or more of the input attributes are in complete conflict and will have the unwanted effect of discarding the attribute data contained in the input report.

D-S has been demonstrated to provide some tolerance to the occasional piece of conflicting input data. Through the use of Dempster's rule, conflicting data can be accommodated via the ignorance proposition Θ: as long as an input proposition set has a non-zero mass assigned to Θ, it will never be entirely in conflict with an existing recognition proposition. However, if the conflict is large, significant changes in proposition masses can result from the re-normalization process. This can result in undesirable effects, such as oscillations in the belief between competing propositions [20]. The truncated D-S algorithm, which maintains a minimum level of ignorance in the recognition proposition set, bounds the possible level of conflict and is able to recover more quickly from the effects of a conflicting data input than the conventional D-S algorithm [3].

In order to better understand options for improving the identity alignment process, a study [4] was conducted to compare the effect of different independence assumptions and evidence combination methods on the identity alignment of both conflicting and non-conflicting multi-attribute input reports. In addition to Dempster's rule, other evidence combination methods included: five rules based on Dempster-Shafer Theory (Yager's [22], Smets' [18], Dubois and Prade's [6], Murphy's [11] and the Disjunctive [5] rule); two rules based on Dezert-Smarandache Theory (DSmT) [17]; the Taxonomy method [16]; and a method, termed Best Intersection [4], proposed during the implementation of COMDAT. The other Dempster-Shafer based rules differ from Dempster's rule in their approach to how conflicting evidence is handled. For example, Yager's rule transfers conflict to the ignorance proposition, whereas Dubois and Prade's rule transfers conflict to the union of the conflicting propositions. In the case of independent attributes, the study indicated that the standard combination of data using Dempster's rule is adequate. In the case of dependent attributes, the study indicated that:

- treating dependent data as independent (combining attributes using Dempster's rule) results in over-confidence in the resulting beliefs;
- the decision as to which data to combine is more important then the choice of evidence combination method. The study proposed a heuristic that the minimal set of independent and conflicting data be selected for combination;
- the choice of evidence combination rule is not significant when conflict is small;
- the Taxonomy, Dubois and Prade's, and DSmT rules provide more accurate results in cases of highly conflicting evidence; but,
- taking into account efficiency and ease of use, Dempster's rule is recommended for overall use, based on the use of the data selection heuristic.

Various methods have been devised to modify the identity update process to handle dependent data resulting from repeated attribute measurements. As a starting point, each MSDF track can store a record of the input source for each reported attribute and then use this information to filter out duplicate attribute measurements. However, filtering out all duplicate attribute measurements has the disadvantage of making the recognition process vulnerable to random errors in received attribute values. A second approach, described in [2], allows multiple attribute measurements from the same sensor, but only if they are well enough separated in time to be considered independent. COMDAT implements a variant of this approach, termed the *consecutiveness criterion*. A third approach, also implemented in COMDAT, allows multiple attribute measurements from the same source but only while the conflict exceeds a user-defined threshold.

4. Assessment of the potential role of an automated evidential reasoning capability to support recognition tasks

A significant outcome of the various analyses described in Section 2 was the development of a Concept of Operations (CONOPS) and Operator-Machine interface (OMI) for an automated decision support capability based on incorporating D-S algorithms (see Section 5). In COMDAT, the analyses helped to focus design attention on the work operators face in performing their recognition tasks. One specific way in which this was done was to look more generally at the role of evidential reasoning in the D-R processes that operators currently engage in when processing contacts.

By evidential reasoning we mean the process of examining individual and multiple pieces of data to determine what information content they collectively provide to make as complete a recognition of a contact as possible. In general, as attribute information is acquired on contacts, operations room team personnel make a number of recognition decisions, such as: "Is the contact an air or surface contact?", "Where is the contact?", "What is its IFF, altitude, etc.?", "Have I seen this before?", "Do I have information to expect something?" D-S functionality in COMDAT (as described in Section 3) represents a particular scheme for effecting automated evidential reasoning.

While the analysis of operators' recognition processes provided evidence of evidential reasoning, comparative assessments of the nature of this processing with automated reasoning algorithms suggested a limited number of operational situations where such functionality could be especially beneficial. Such situations included circumstances around a saturated air attack where the ability of the team to identify and recognise contacts of interest could be compromised, such that data queues start to form. In such cases, it is possible that an automated evidential reasoning capability could usefully enhance the performance of the team by automatically processing the data available to make recommendations about recognition characteristics.

A further situation where this functionality could improve upon the speed and rate of processing, is when the team as a whole is focused on a specific contact and a large number of other contacts build in the interim in the wider area. Automation could then have a role in processing these contacts in the background to facilitate the team in rebuilding the wider picture. However, the team's immediate requirement in regaining the larger picture is to identify the most pressing threats, and existing operations room functionality already provides some guidance and information in this respect. In both of the cases described above, further investigation is required to quantify the benefit to the recognition process of adding an automated reasoning capability.

Finally, it is important to note that the analyses compiled significant evidence of the differences between recognition processes as currently conducted in a HALIFAX Class operations room and the processing of the D-S recognition algorithms that were implemented in COMDAT. Some of these differences relate to the range and types of information and *a priori* knowledge that operators employ. For example, not surprisingly, operators use a large amount of contextual knowledge in making their recognition assessments which are not currently accounted for in D-S processing. This suggests: the need for careful consideration of operator intervention and interaction with automated D-S functionality; the need to incorporate mechanisms to make D-S processing sufficiently observable and directable by operators; as well as the potential extension of the processing capability of D-S recognition algorithms to include contextual information, where possible.

5. Concept of operations and user interface concepts for operator involvement with automated recognition processes

5.1 Concept of operations

At the time the COMDAT project was initiated, it was recognized that the necessary human factors research to support the implementation of MSDF did not exist. Thus, a significant component of the overall project was devoted to the development of relevant human factors knowledge [8]. A major goal of this research was to develop a CONOPS and prototype OMI for the provision of a semi-automated and automated decision support capability based upon the D-S algorithms to assist the processes of contact recognition and identification in the HALIFAX Class frigate.

The basis for the CONOPS and the OMI were the results of the information flow and decision process analyses and the capabilities of the D-S algorithms discussed in the previous sections. This information was used to identify the relevant decision processes that would be impacted by the algorithms, to develop an implementation concept, and finally to develop OMI concepts.

An overriding assumption in the development of the CONOPS was that the human operator would continue to make assessments of contact data in parallel with the algorithms. Based on the results of the analysis, there was always the possibility that operators would have access to information that was not part of the data analysed by the algorithms including contact behaviours (such as changes in previous patterns of activity, unexpected changes in course or speed) and the geographical location of contacts. Moreover, it was envisaged that this approach would lead to greater operator acceptance. This decision meant that the CONOPS would need to fully integrate the output of ongoing human analysis processes in detection and recognition, so that the D-S algorithms could leverage the information provided by humans to amplify and modify the data provided by the sources available to them.

The process followed in developing the CONOPS considered the following steps:

- Determine if the implementation will be evolutionary or revolutionary in terms of existing systems and personnel organization,
- Resolve the respective roles of the human and the decision support system in making decisions,
- Establish the appropriate level of support for human decisions, and
- Integrate decision support functionality into the operational context.

The design and implementation of a revolutionary system generally assumes a clean slate. Hence, it is rarely possible to implement a revolutionary system. Even if a system is installed in a new class of ship, it is unlikely that the personnel operating the system would have different roles and responsibilities. Since MSDF must support current tasks (i.e. recognition and identification), it is even more critical that the implementation be evolutionary. Thus, a critical constraint on the implementation was that it be evolutionary.

The second step was to consider roles and responsibilities the human and MSDF would have in the new system. In the current system, the assignment of Standard Identification (ID) is ultimately the responsibility of the Commander of either the ship or the task group. That authority is usually delegated under well defined conditions. For example, the recognition and identification of contacts traveling in commercial air lanes, where the contact data are typically consistent with those of commercial aircraft, may be assigned to a junior operator. To maintain this capability, a requirement of the system was that it should allow the human to specify the responsibilities that MSDF can assume within a specified context.

The third step investigated the level of automation or decision support to be implemented. Table 2 [14] shows the levels of support that a decision support system could provide ranging from manual to complete automation.

With simple systems, a high level of automation may be appropriate. However, when the decision support system is replacing some of the operator’s decision making responsibility, a more conservative level of automation may be preferable. In naval operations, there are critical Rules of Engagement (ROE) to adhere to that provide the operators the legal framework under which to take potentially lethal action. Moreover, operators want to be in control or feel they are in control, especially over circumstances that are potentially life threatening. A more conservative approach to decision support is less likely to lead to conflicts in authority between the human and the system, and the rejection of the decision support system.

On the other hand, a too conservative automation level is likely to result in sub-optimum human-system performance. It is important that the decision support operates at a high enough level to reduce the workload that would otherwise be required to gather and comprehend the information needed to make an informed decision. By providing operators with support levels that should provide tangible benefits, there will be more immediate initial acceptance and a likely change in subsequent attitudes that may allow higher levels of support to be implemented in the future.

Level	Description
10	The system acts autonomously without human intervention
9	The system informs the user after executing the action only if the system decides it is necessary
8	The system informs the user after executing the action only upon user request
7	The system executes an action and then informs the user
6	The system allows the user a limited time to veto before executing an action
5	The system executes an action upon user approval
4	The system suggests one alternative
3	The system narrows the selection down to a few
2	The system offers a complete set of action alternatives
1	The system offers no assistance

Table 2: Hierarchy of levels of decision support.

Based on the above analysis, the following factors were considered in choosing the appropriate level of support: (a) the importance of the decision and consequences for error, (b) the rate at which the system (including the human operator) must process data and (c) the operator's ongoing workload. In the case of COMDAT, it was decided that it would be appropriate for a human to be the final arbiter of major decisions but that MSDF would work in the background to bring all the relevant pieces of information to the attention of the operator to support decision making. This role was particularly critical with respect to factors (b) and (c) since under some tactical situations data rates could be very high and the operator would typically only have sufficient moment to moment capacity to attend to a specific situation that was most time sensitive. Thus, after a period in which the operator was focused on a critical tactical problem, the decision support system should allow the operator to readily re-acquire wide area picture situation awareness by processing relevant data in the background and presenting the operator with a summary of what new information had been acquired relevant to the operator's pre-defined information requirements.

The fourth step in the process was deciding how to integrate MSDF into the operational context. If a decision support system is fully developed, as was the case with COMDAT, its functionality is fixed, potentially restricting the designer's flexibility in implementing an optimal level of support based on operator requirements. As a result, it may be difficult to successfully integrate the system into either the existing technology or the culture of the organization. However, it is not impossible. Moreover, the technology is always evolving. By designing the interface to meet the operator requirements, one can hopefully guide the evolution of the underlying decision support algorithms in a direction that will lead to greater operator acceptance and consequently a more effective system.

Thus, the proposed implementation was based on the analyses of operator requirements discussed above, but took into account the potential limitations of MSDF as implemented in the COMDAT project. As mentioned previously, even with the potentially useful integration, or fusion, of numerous data sources to arrive at an assessment of recognition or ID, MSDF under COMDAT did not have the capacity to integrate all the available data sources used by human operators in compiling the picture, such as intelligence and contextual information. Contextual information may reduce the number of possible recognition options presented by MSDF about a specific contact. This possibility influenced the decision to position MSDF as a collaborator in rather than a replacement for the current picture compilation processes carried out by the operator. In order for MSDF to be an effective collaborator, it was important that it be able to accept operator input. Most existing approaches to the implementation of decision support algorithms do not necessarily provide for this capability.

In the case of the current CONOPS, this was done by building a capability for operators to define different operating modes including: the ability to weigh or discount various sources, to define thresholds for various actions to be taken, and to interact with the various recommendations presented by MSDF.

5.2 Operator machine interface concepts

The user interface that was developed as a result of the CONOPS had four main components:

1. A *setup page* for configuring the system,
2. Three types of *data pages* through which the operator interacts with the system while it is operational,
3. *Alerts and indicators* which are operator-configured to provide improved situation awareness under prescribed conditions,
4. A *contact symbol and data block* in the tactical situation display.

The setup page allowed the operator to configure the level of autonomy or operating mode that MSDF would have and the conditions under which the system was to alert the operator. For example, the operator might request that the system issue an alert whenever a contact going faster than speed "x" was within distance "y" of own ship. Communication between the MSDF algorithms and the operators occurred through the data pages. The most important of these was the overview page. It provided the operator with a summary of all the contacts being monitored by MSDF as well as any alerts associated with them. The operator would use the overview page to regain situation awareness about the overall tactical picture and to select the next contact requiring evaluation. Since an evolutionary implementation was envisaged, these pages would be used in conjunction with, rather than replacing, the current tactical situation display. The tactical display provides the geospatial picture while the overview page provides MSDF's overall assessment of the contacts it is monitoring. A more detailed description of some of these OMI concepts can be found in [9].

The most relevant concept in support of the recognition process was the Operator As Input (OAI) function. This function was designed to allow the operator to influence the recommendations provided by an MSDF system that normally would not accept operator input. For each contact the operator could bring up a single contact proposition page (Figure 3). The "organic MSDF" column on that page shows the current recommendation for that contact (based on the D-S analyses described earlier). If the operator disagreed with one of these probabilities, they could enter an alternative recommendation in the "operator input" column including the reason for the recommendation. This change would not affect the MSDF-only analysis, but would be processed in the parallel OAI MSDF analysis. For example, MSDF might indicate that there was a 72% probability that a contact is a merchant vessel based on its speed. However, the operator might have had a visual sighting that the contact was a warship and enters that there is a 95% probability that the contact is a warship. The MSDF system would continue to collect data from its sources to try to refine its recommendation concerning the type of merchant ship and presents this in the "organic MSDF" column. A parallel analysis, using the same data plus the new OAI information, would be carried out to further refine the recommendation that the contact was a warship (e.g., the type and class of warship). The results of that analysis would appear in the "operator input" column.

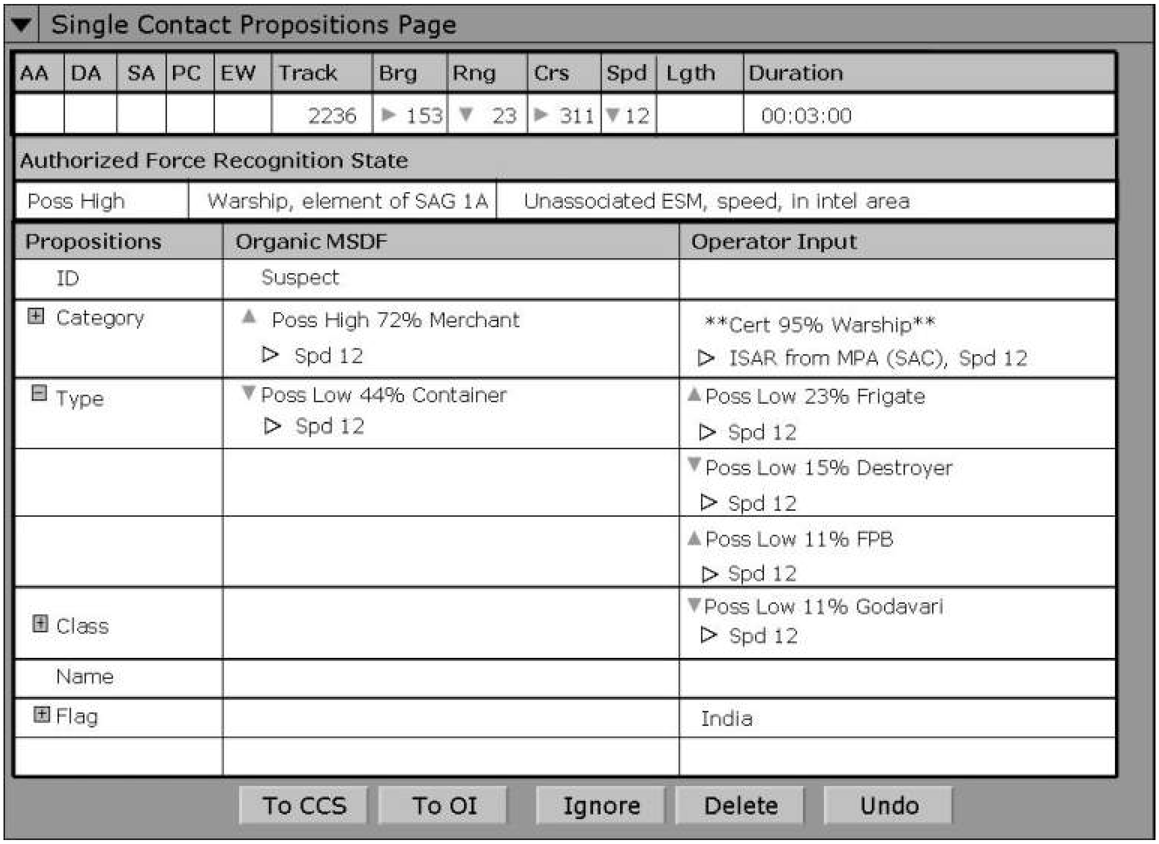

AA	DA	SA	PC	EW	Track	Brg	Rng	Crs	Spd	Lgth	Duration
					2236	▶ 153	▼ 23	▶ 311	▼12		00:03:00

Authorized Force Recognition State		
Poss High	Warship, element of SAG 1A	Unassociated ESM, speed, in intel area

Propositions	Organic MSDF	Operator Input
ID	Suspect	
⊞ Category	▲ Poss High 72% Merchant ▷ Spd 12	**Cert 95% Warship** ▷ ISAR from MPA (SAC), Spd 12
⊟ Type	▼ Poss Low 44% Container ▷ Spd 12	▲ Poss Low 23% Frigate ▷ Spd 12
		▼ Poss Low 15% Destroyer ▷ Spd 12
		▲ Poss Low 11% FPB ▷ Spd 12
⊞ Class		▼ Poss Low 11% Godavari ▷ Spd 12
Name		
⊞ Flag		India

Figure 3: Proposed OMI concept for displaying the output of MSDF on a single contact and allowing operator input to the MSDF system.

This approach has several advantages. It allow operators to use the decision support system to test out their own hypotheses and also provides a second check on those hypotheses through a comparison of the recommendations from the MSDF-only analysis and those from the OAI analysis. If the MSDF algorithms simply cannot support the operator input assessment, then an alert is generated, allowing the operator to reassess the underlying rationale for their assessment.

Although a mock-up of the OMI concepts was evaluated by Navy personnel, it was not possible to produce a working version of the OMI. Thus, the actual utility of the OAI concept remains to be evaluated.

6. Measures and methods for assessing operator and system performance in carrying out recognition tasks

6.1 Operator measures of performance

With the introduction of any new system aimed at improving operator and system performance, a key question is "Does it improve performance and, if so, by how much?" At best, the evaluation of a new system involves subjective assessments by experienced operators whose opinion is often biased by their attitude towards the new technology. There are no objective evaluations to determine if operators can conduct their tasks more efficiently or effectively using the new system. A major hindrance to a more objective evaluation is the lack of standard Measures of Performance (MOPs) and estimates of performance with the current system. To overcome this limitation, one component of COMDAT was the development of relevant MOPs for the recognition process and the collection of baseline data.

The initial development of MOPs was carried out in conjunction with a cognitive analysis of the Operations Room Officer (ORO) position. For each goal identified during the cognitive analysis, corresponding criteria, measures, and methods were developed. The criteria identified a dimension of evaluation interest; while the measure and method operationalized the criteria. For example a key goal in the recognition process is building and maintaining an awareness of the air picture. Some criteria associated with this goal are to:

- acquire and maintain awareness of significant issues with the air picture;
- detect pertinent changes in key elements of the air picture promptly and accurately;
- identify hostile contacts.

Measures that can be used to evaluate these criteria include:

- percent of current air contacts processed,
- response time to identify hostile contact, and
- number of contact changes responded to.

Methods for assessing these measures could include

- freeze probe,
- embedded probes, and
- SME review of real time or video of the scenario execution.

The complete set of goals and criteria were analysed and the list refined to focus on the specific goals that MSDF might impact. A total of 67 MOPs were identified. The next step was to develop a detailed experimental plan to evaluate each MOP and review possible environments for conducting the planned studies. Potential environments identified included existing trainers, and onboard ship. Not surprisingly, operational and manpower constraints made it difficult, if not impossible, to utilize any of these environments to carry out the proposed experimental plan. The review did indicate a possible alternative. Training sessions in the Operations Room Team Trainer (a simulation/training facility used for the team training of HALIFAX Class operations room personnel) utilized a wide range of scenarios that were carried out by experienced naval personnel. Archives of these sessions were maintained including an audio/video recording of all team interactions. These recordings could be played back in parallel with the captured screen and keystroke information. A preliminary analysis indicated that baseline data on 64 of the MOPs defined earlier could be collected using archived training records. Although analysis of the training records appeared feasible, it proved extremely time intensive. A critical recommendation therefore was that the evaluators observe the actual training sessions in order to mark critical incidents or responses for later review.

The final step was to carry out a Baseline Study. Multiple anti-air and anti-surface warfare scenarios were analysed and means and standard deviations computed on 27 anti-air warfare and 17 anti-surface warfare MOPs. The resulting data set, while still relatively small, did show the potential of the approach and the type of data that could be collected. Relevant data were collected on task times and accuracy in initiating and resolving air contacts, task times and accuracy in detecting and recognizing surface contacts and task times and accuracy for first and subsequent standard ID for both air and surface contacts. In addition, errors that could be impacted by the introduction of MSDF were documented. These errors were categorized as procedural, situational awareness, or attribute recognition. It was concluded that the majority of performance errors could have been helped by MSDF-related support tools.

Although the method proved useful, there were limitations. The variance in the data collected was quite large. This is not surprising since each data point was based on a unique situation (i.e. different contacts, different sensors etc). In classical experimental studies, the data is collected using multiple instances of the same or very similar stimulus situations. Although the training sessions are more realistic, the probability of finding a statistically significant difference in performance between the baseline system and a prototype system such as COMDAT would be very low (at least using inferential statistics) unless very large quantities of data were collected. Given the workload-intensive nature of the data extraction process, this could be costly and time consuming. However, some methods for simplifying data collection were identified during the COMDAT study.

A further problem with this approach is correlating the baseline data with data collected using the new technology. Ideally, the prototype system would be installed in the Trainer. This was not feasible within the constraints of the COMDAT project. One recommendation is that future projects build this requirement into their project. A better approach to evaluation would be to build a simulation of the current system that was designed to easily incorporate new technology. This possibility is currently being pursued.

Despite these limitations, the resulting data does provide more objective estimates of current performance that developers can use as benchmarks in evaluating future systems. Moreover, the MOPs have value for use in more objectively evaluating team performance during training as well as the training regime itself. If the MOPs were used to evaluate training outcomes, data collection and analysis could probably be partially automated. Moreover, ongoing baseline data would be generated that would be of use not only for the training staff, but also for future development projects.

6.2 System measures of performance for automated recognition

Work was also conducted, as part of COMDAT, to develop an experimental process for the shore-based test and evaluation of the automated recognition capability implemented in the COMDAT testbed. The work included development of a test environment, methodology, scenarios, and system measures of performance. Eleven MOPs associated with the fusion of contact attribute information were considered which take into consideration the clarity, timeliness, accuracy, and completeness of the estimated recognition.

A test environment was established using the Navy's Combat System Test Center (CSTC) Mini-System, a test system used for the maintenance, development and testing of the HALIFAX Class combat system software. The CSTC Mini-system uses the Combat Systems Simulation (CSS) system, which simulates ownship sensors and weapons, as well as targets and tactical scenarios. The CSTC Mini-System was selected from other options, that included development and test facilities and operational trainers, on the basis of criteria that included cost, availability, security level, scenario development capabilities, supported data sources, data recording capabilities, and required infrastructure upgrades.

Six scenarios were developed to support areas of investigation that included fusion of multiple sources of attribute information, the impact of conflicting attribute data on recognition performance, and the impact of attribute data on data association. The approach used to develop the scenarios was to overlay a set of specified target manoeuvres on regions of differing sensor coverage. The topology of a sample scenario is shown in Figure 4. The scenario includes single targets as well as closely spaced target clusters. Scenario complexity is based on factors including the number of target clusters, the complexity of target manoeuvres, and by the degree of difference in the ground-truth recognitions of targets making up target clusters.

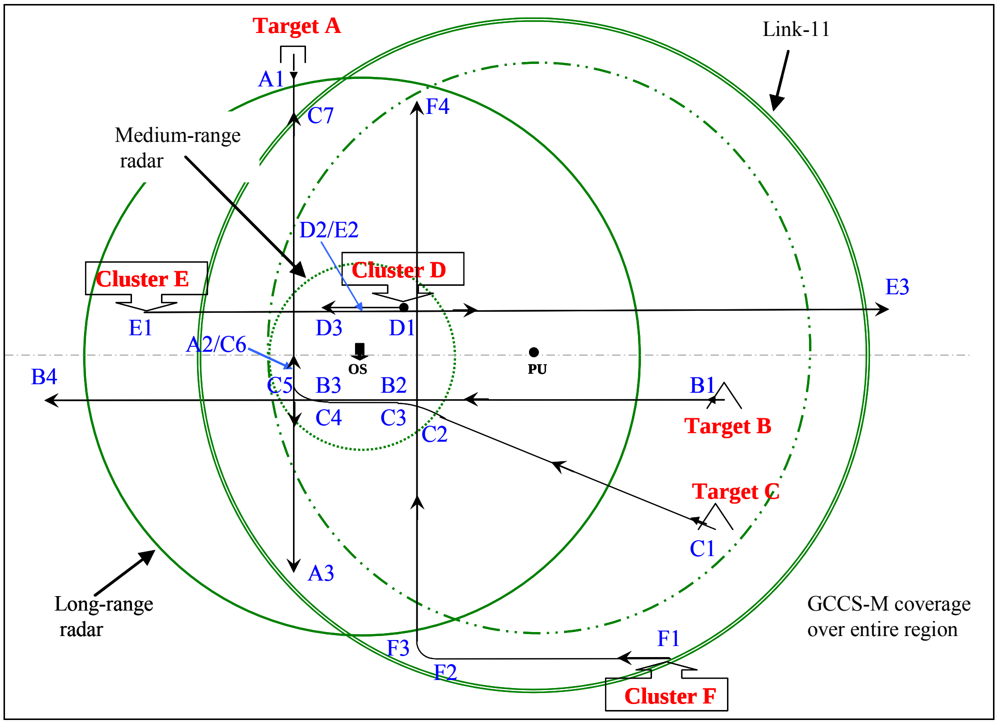

Figure 4. Sample scenario topology for shore-based trial. Included are: target trajectories (indicated by target symbols and arrows) and sensor coverage areas (indicated by green rings).

System MOPs for recognition were identified based on the following aspects of tactical picture quality: accuracy – the degree to which the assigned recognition matches the ground-truth recognition; completeness – the degree to which all real-world objects have been recognized; clarity – the stability of the recognition; and timeliness – the time required to achieve recognition to a specified level. Examples of MOPs considered and their relationship to picture quality include:

- fraction of recognition reports matching ground-truth (accuracy);
- fraction of recognition reports for which recognition is declared (completeness);

- time elapsed until recognition is declared (timeliness);
- number of switches in recognition value (clarity);
- degree of conflict between different recognition values (clarity).

Recognition conflict is based on the D-S conflict, κ, and measures the extent to which there are competing recognition propositions. In most cases, MOPs are calculated over individual tracks (or ground-truth contacts), and statistics gathered over individual scenarios, or over a set of scenarios.

One concept used in some of the MOPs calculations is the *identity confusion matrix*, an extension of the confusion matrix [10] used in the assessment of data association performance. With rows i corresponding to MSDF tracks t_i and columns j corresponding to ground-truth targets g_j, an identity confusion matrix entry $m(i,j)$ indicates the number of recognition updates for MSDF track t_i that were generated by input reports originating from a particular ground-truth target g_j and that match the recognition of g_j. The identity confusion matrix has the advantage that it can be used in cases where the a priori assignment of input data to ground-truth is not known, as is usually the case for experimental trials with real data.

Performance tests were carried out using three MOPs related to recognition accuracy and completeness – R^{ID}, the fraction of recognition reports for which a recognition is declared (the belief in the recognition value exceeds a threshold of 0.5); R^{GT}, the fraction of recognition reports which are correct (matches ground-truth); and $R^{ID\text{-}GT}$, the fraction of declared recognition reports which are correct. The recognition reports used for calculation of these MOPs were D-S generic propositions at the Category and Sub-Category level. Category was recognized accurately, as is consistent with the number of input sources which directly report Category via the Type attribute (see Table 1) and the relative degree of ease with which contacts having different Category values (Surface or Air) can be discriminated. In contrast, Sub-Category reports were generally correct once a recognition was declared but were more easily affected by errors in data association. The incorporation of data from attribute-rich sources, such as Link-11 and GCCS-M, noticeably increased the frequency of Sub-Category declarations. In addition, the presence of conflicting information was more likely to result in no declared recognition rather than an incorrect recognition value.

One challenge in conducting performance tests was to separate out the effects of tracking and data association on recognition performance. This was addressed partially through the design of the scenarios, by including a baseline of geographically well-separated tracks having minimal numbers of manoeuvres which will minimize the number of spurious input reports inserted into the tracks' data. In addition, some tests were conducted first using data with noise-free position reports (which had minimal tracking error) and then repeated with more realistic position data.

7. Conclusions

At the time it started, COMDAT was unique in its support of both a human factors component and decision-support technology research component that ran in parallel while providing input to each other. As a result of this approach, a broad range of knowledge and capabilities were provided to the Navy for potential integration into their capital acquisition programmes. One of the most important outputs of COMDAT was the detailed understanding of the recognition and decision processes used by operators. This knowledge was essential for determining where the D-S algorithms could be most effectively used and the development of a flexible CONOPS that would ensure the successful integration of the D-S algorithms into an operations room used to manual recognition and identification. The second important outcome was the development of MOPs and assessment methods to assess both system performance (hardware and software) and human-in-the-loop performance (hardware, software and human) that could be used across a wide range of environments including the shipboard environment itself. This outcome provided the

Navy with the capability to objectively assess new technologies and interfaces and to specify meaningful criteria that new systems must meet. Finally, significant extensions were made to the pre-existing D-S recognition algorithms based on the requirements to process realistic data and to consider new data sources such as GCCS-M.

Some of the important lessons learned from the various research efforts include:

- the importance of conducting cognitive analyses and work flow analyses of key decision making processes early in the project.
    - the development of algorithms preceded a detailed human factors analysis and so was based on technology-driven assumptions regarding the appropriate level and type of automation, and
    - the development of the OMI did not really occur until close to the end of the project. Thus, it was not possible to produce a working interface.
- the requirement for a capability that permits system evaluation at the earliest possible stage in the development process using a realistic representation of the operational environment
    - considerable effort was expended revising the algorithms in order to handle the characteristics of real input data.
- the inclusion of a plan to integrate the new technology (including the OMI) into a realistic representation of the existing environment.
    - the critical comparison of human-in-the-loop performance between the existing system and the D-S enhanced system was not possible.
    - because a single test environment that could accommodate system-level and human-in-the-loop (single operator and team) performance was not available time and resources were expended in developing and validating multiple test environments.

The recognition concepts developed under COMDAT can be extended in several areas. First, the OMI concept for user involvement with MSDF-based recognition processes has still to be implemented and tested. As part of such a development some adaptation of the truncated D-S algorithms is expected. Second, extensions to the existing D-S capability, such as the processing of selected relational attributes and the incorporation of some aspects of contextual information, are also possible. Third, implementation of the OMI concept should be followed by human-in-the-loop testing, using the measures and methods for assessing operator performance developed under COMDAT, and compared against baseline performance on the existing CCS. Finally, many of the broader issues raised in the modeling of the recognition process deserve further attention, a specific case being the further analysis of operational situations where automation, such as that based on D-S evidential reasoning, can have the greatest impact on the recognition process. The net result of this approach to development should be technology which is well-integrated with the work demands of the shipboard environment and which can provide maximum benefit in increasing the overall effectiveness of the Maritime Tactical Picture.

Acknowledgements

The authors would like to acknowledge the contributions made by Lockheed-Martin Canada and Human*systems* Inc. in carrying out aspects of the work described in this paper. The authors would also like to acknowledge the valuable assistance received from members of the Canadian Forces who participated in the various evaluations and data gathering sessions during the conduct of the COMDAT project.

References

1. Blackman, S., and Popoli, R. (1999). Design and Analysis of Modern Tracking Systems. Norwood, MA: Artech House.
2. Bogler, P. (1987), Dempster-Shafer Reasoning with Applications to Multi-sensor Target Identification, *IEEE Trans Systems Man and Cybernetics*, vol. 17, pp. 968-977.
3. Bossé, É. (2002), Identity Information Fusion Through Evidential Reasoning, *Proceedings NATO Advanced Research Workshop on Multisensor Data Fusion*, Pitlochry, Scotland, pp 457-480.
4. Couture, J. (2009), unpublished report.
5. Dubois, D., and Prade, H. (1986), A set-theoretic view of belief functions. *International J. General Systems*, vol. 12, pp. 193-226.
6. Dubois, D., and Prade, H. (1988), Representation and combination of uncertainty with belief functions and possibility measures. *Computational Intelligence*, vol. 4, pp. 244-264.
7. Klein, G.A., Calderwood, R. & MacGregor, D. (1989). Critical Decision Method for Eliciting Knowledge. *IEEE Transactions on Systems, Man and Cybernetics*, 19(3), 462-472.
8. Matthews, M. L., Rehak, L., and McFadden, S. (2007), Final report on human factors research conducted under the COMDAT Technology Demonstrator Project, DRDC-Toronto CR-2007-149, Human*systems* Inc., Guelph, ON.
9. Matthews, M. L., Keeble, A. R., and McFadden, S. M. (2009), Integrating multi source data fusion with the Navy operator *Proceedings of the Fourteenth International Ship Control Systems Symposium*, Ottawa ON: Canadian Association of Defence and Security Industries.
10. Mori, S., Chang, K.-C., Chong, C.-Y., and Dunn, K.P. (1989), Tracking Performance Evaluation: Prediction of Track Purity, SPIE Proceedings, vol. 1096.
11. Murphy, C. (2000), Combining belief functions when evidence conflicts. *Decision Support Systems*, vol. 29, pp. 1-9.
12. NATO (2002). Multinational Maritime Tactical Instructions and Procedures, MTP 1(D) Volume 1.
13. NATO RSG 14. (1999) Analysis Techniques for Man-Machine Systems Design. Technical Report AC/243 (Panel 8) TR/7.
14. Sheridan, T.B. and Verplank, W. (1978), Human and computer control of undersea operators, Man-Machine Systems Laboratory, Department of Mechanical Engineering, MIT, Cambridge, MA.
15. Simard, M.-A., Couture, J., Bossé, É. (1996), Data Fusion of Multiple Sensors Attribute Information for Target Identity Estimation using a Dempster-Shafer Evidential Combination Algorithm, SPIE Proceedings vol. 2759, Signal and Data Processing of Small Targets, pp. 577-588.
16. Simard, M.-A., Valin, P., Lesage, F. (2000), The ID Box: A Multisource Attribute Data Fusion Function for Target Identification, SPIE Proceedings vol. 4051, Sensor Fusion: Architectures, Algorithms, and Applications IV, pp. 180-190.
17. Smarandache, F., and Dezert, J. (eds.) (2004). Advances and Applications of DSmT for Information Fusion, collected works., Rehoboth, NM: American Research Press.

18. Smets, P. (1990). The combination of evidence in the transferable belief model. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 12(5), pp. 447-458.
19. Valin, P., Couture, J., and Simard, M. (1996), Position and Attribute Fusion of Radar, ESM, IFF and Datalink for AAW missions of the Canadian Patrol Frigate, Multisensor Fusion and Integration for Intelligent Systems (MFI 96), Washington, DC, Dec 8-11 1996, pp. 63-71.
20. Valin, P. (2002), Reasoning Frameworks, *Proceedings NATO Advanced Research Workshop on Multisensor Data Fusion*, Pitlochry, Scotland, pp 223-245.
21. Vicente, K. J. (1999). Cognitive work analysis: Toward safe, productive, and healthy computer-based work. Mahwah, NJ: Lawrence Erlbaum Associates.
22. Yager, R.R. (1987). On the Dempster-Shafer framework and new combination rules. *Information Sciences*, vol. 41, pp. 93-138.