

Fusion of Imaging and Non-imaging Data for Surveillance Aircraft[†]

Elisa Shahbazian, Langis Gagnon, Jean-Rémi Duquet, Maciej Macieszczak, Pierre Valin

Lockheed Martin Canada Inc., 6111 Royalmount Ave.,

Montréal, Québec, H4P 1K6, Canada

ABSTRACT

This paper describes a phased incremental integration approach for application of *image analysis* and *data fusion* technologies to provide automated intelligent target tracking and identification for Airborne Surveillance on board an Aurora Maritime Patrol Aircraft. The sensor suite of the Aurora consists of a radar, an Identification Friend or Foe (IFF) system, an Electronic Support Measures (ESM) system, a Spotlight Synthetic Aperture Radar (SSAR), a Forward Looking Infra-Red (FLIR) sensor and a Link-11 tactical datalink system. Lockheed Martin Canada (LMCan) is developing a testbed, which will be used to analyze and evaluate approaches for combining the data provided by the existing sensors, which were initially not designed to feed a fusion system. Three concurrent research proof-of-concept activities provide techniques, algorithms and methodology into three sequential phases of integration of this testbed. These activities are: (a) analysis of the fusion architecture (track/contact/hybrid) most appropriate for the type of data available, (b) extraction and fusion of simple features from the imaging data into the fusion system performing automatic target identification, and (c) development of a unique software architecture which will permit integration and independent evolution, enhancement and optimization of various decision aid capabilities, such as Multi-Sensor Data Fusion (MSDF), Situation and Threat Assessment (STA) and Resource Management (RM).

Keywords: Intelligent Decision Aids, Data Fusion, Image Analysis, Automatic Target Recognition, Real-Time Adaptive Planning, Situation and Threat Assessment, Resource Management.

1. INTRODUCTION

The Intelligent Decision Aid capabilities within a CCS are developed using *image analysis*, *data fusion*, *real-time adaptive planning* and other technologies. As a technology, *data fusion* is actually the integration and application of many traditional disciplines and new areas of engineering, such as communication and decision theory, epistemology and uncertainty management, estimation theory, image processing, digital signal processing, computer science and artificial intelligence to provide object kinematics and identity estimation, situation and threat assessment, and process refinement/sensor management. The recent arrival of sophisticated imaging sensors, such as the Spotlight Synthetic Aperture Radar (SSAR), to complement existing non-imaging sensors, provides the opportunity for a new type of fusion that synergetically makes use of contextual information, positional data, information extracted from images, using various *image analyses* techniques (e.g. Image Processing, Image Segmentation, Feature Extraction, Image-based target classification), and data from other sources. *Real-time adaptive planning techniques* provide decision aid in Resource Management (RM).

There are many technical challenges for providing a real-time integrated and automated CCS employing data fusion and other decision aid capabilities. These technological challenges include:

1. Dissimilar data from various types of sensors (2D, 3D, imaging, etc.),
2. Fusion of incomplete data, and data of uncertain track quality,
3. Alignment (registration) of geographically distributed sensors,
4. Processor resource requirements and concurrent implementation strategies to satisfy the real-time performance requirement,

[†] SPIE Proc. #3067, conference "Sensor Fusion: Architecture, Algorithms and Applications", Orlando 1997

5. Availability of experts and methodology for building the situation and threat assessment knowledge agents and knowledge bases,
6. Availability of realistic testbeds and test data to demonstrate and measure the CCS performance.

To analyze, develop and demonstrate solutions to these issues, LMCan is building a testbed called Adaptable Data Fusion Testbed (ADFT) that has a modular design and is incrementally implementing an intelligent decision aid capability for a typical remote sensing application where information is provided both by imaging sensors and non-imaging sensors.

The ADFT development is being done in three incremental phases that tackle successively more complex sensor inputs and more sophisticated algorithms for decision aid, thus minimizing the overall technical risks, yet ensuring a useful product at each completed step. Three concurrent research proof-of-concept activities provide techniques, algorithms and methodology into each phase of integration. These activities are: (a) analysis of the fusion architecture (track/contact/hybrid) most appropriate for the type of data available, (b) extraction and fusion of simple features from the imaging data into the fusion system performing automatic target identification, and (c) development of a unique software architecture which will permit integration and independent evolution, enhancement and optimization of various decision aid capabilities, such as Multi-Sensor Data Fusion (MSDF), Situation and Threat Assessment (STA) and Resource Management (RM). This paper describes some research results and methodology developed in these proof-of-concept activities and the phased incremental integration approach to gradually build an automated intelligent target tracking and identification capability for Airborne Surveillance on board an Aurora Maritime Patrol Aircraft.

2. ANALYSIS OF THE FUSION ARCHITECTURE

The Intelligent Decision Aid capabilities are becoming more and more critical for both Naval and Airborne platforms. These platforms are currently equipped with a sensor suite which may or may not be upgraded before they have to be integrated into the CCS through advanced decision aid capabilities. Furthermore, the current trend is to make the sensor systems more and more "intelligent", and not much is published about any analysis on how to optimize tracking and identification by intelligently distributing the various operations of processing sensor data between the sensor systems and the combat system. Since very few of the suppliers of these "intelligent" sensors discuss with the CCS integrator what are their requirements for combining the sensor data with the other data within the CCS, in the first phase of the ADFT development LMCan has decided to focus on techniques for fusing "what is available" (or easily derivable), rather than "what is optimal to fuse". In the subsequent phases, based on the results of parallel proof-of-concept analyses, incrementally more optimal data fusion algorithms and more complex image analysis methods for target classification will be incorporated.

These implementations of different fusion architectures within ADFT will help find answers to questions such as:

- How can one supplement the incomplete data provided by the sensors?
- What type of fusion architecture (Track Level or Contact Level) is appropriate for the type of data provided by the sensors?
- How does one take into account Track-to-Track cross-correlation in the Track Level Fusion?
- For each type of sensor data, how much pre-processing should be done before providing it to the fusion system?

Before describing the specifics of each of these phases, the Aurora aircraft's sensors and the data provided by them is discussed in the next section.

2.1 Input Data into the Fusion System

The enhanced combat system of the Aurora is expected to fuse data from:

- a. attribute measurement oriented sensors (Electronic Support Measures (ESM), Identification Friend or Foe (IFF), etc.);
- b. imaging sensors (Forward Looking Infra Read (FLIR), Spotlight Synthetic Aperture Radar (SSAR), etc.);
- c. tracking sensors (radar, acoustics, etc.);
- d. data from remote platforms (data links); and
- e. non-sensor data (intelligence reports, environmental data, visual sightings, encyclopaedic data, etc.).

Currently the radar and the Forward Looking Infra-Read (FLIR) sensor data are not digitized, and the Spotlight Synthetic Aperture Radar (SSAR) is not yet installed, therefore the tracking and estimation of identity is done manually. To be able to introduce automation into the aircraft target processing, two programs are currently in progress: (1) the development of the Digital Scan Converter (DSC); and (2) the implementation of the Advanced Development Model (ADM) of the SSAR.

In Table 1 below, the inputs to data fusion that are available from the sensors in its **current** version are listed in **bold**, while the others can be made available through further processing, either in an improved version of the sensor, or in the data fusion function itself (provided digital data is given to it by the sensor). In the last column a list of sensors that could be fused (but not necessarily should be) with the given sensor is shown. The main objective of the ADFT project is to find and propose the extent and the algorithms for sensor data processing (mostly imaging) and for fusion that best support the Aurora's missions. The capability of post-flight analysis and subsequent update of the MSDF databases provided by the FLIR and camera exists but is not shown in the table. Also not shown is the possibility of an input to MSDF that is made by an operator visually identifying a platform and providing this information to the fusion function via a keyboard entry.

Sensors	Input to MSDF	Fusion with ...
1. AN/APS-506 Search Radar	range, bearing, speed, acceleration	2, 3, 4, 5, 7, 8, 9
2. SSAR	Range, bearing, elevation, platform ID, mapping, video recording for post-flight analysis	1, 3, 4, 5, 7, 8, 9
3. AN/ALR-502 ESM	Bearing, emitter ID, threat number, AOP, platform ID	1, 2, 4, 5, 7, 8, 9
4. AN/APX-502 IFF	Range, bearing, allegiance	1, 2, 3, 5, 8
5. OR-5008/AA FLIR	Bearing, target attitude, platform ID	1, 2, 3, 4, 7, 8, 9
6. KA-501A camera	Platform ID (for database update)	N/A
7. AN/ASQ-502 MAD	Target position vertically below flight path	1, 2, 3, 5, 8, 9
8. Data Link (Link-11)	Other PU's tactical picture (position, velocity, ID, etc.)	1, 2, 3, 4, 5, 7, 9
9. OL-5004/AYS ADP	Position, velocity and ID	1, 2, 3, 7, 8
10. AN/ARS-501 SRS	Positions of up to 31 sonobuoys	N/A

Table 1- Present suite of sensors and their inputs to a data fusion function

Onboard the Aurora there is also non-sensor data which can be used in various ways to aid the fusion system as discussed below. The ADFT architecture is designed to be able to introduce such information when appropriate.

The intelligence reports can come from many sources, such as satellite information, tracking of targets over several days by foreign government agencies or from Sound Surveillance System (SOSUS) information. None of these sources are relevant for real-time fusion on-board the aircraft but they can aid considerably in the performance of the fusion function by allowing it to focus on the relevant database.

The analyses for the selection of algorithms and architectures also include considerations for the fusion function to allow for environmental data affecting both the positional uncertainty in the reported contact and/or track and the relative quality associated with a given sensor, depending on its sensitivity to environmental conditions. The local environment along the flight path will be used to set the priorities and/or weights given to the different sensor reports.

Visual sightings, particularly after a low-altitude pass, serve mainly to confirm the platform identity, particularly when the offending platforms have to be later brought to justice. Under proper operation of the fusion function, the platform ID should have been obtained considerably before such a fly over, at least in as far as platform type and threat level are concerned. Information reported by visual sightings coming from other sources is inherently hard to quantize in a form similar to other sensor data and may be best entered by an operator's keyboard set after he has himself correlated the relevant information to his best knowledge.

2.2 Fusion Architecture Selection

Figure 1 shows the diagram of the first phase of the fusion system for the airborne application. Here Track Fusion is performed, as in this phase the contact data will not be accessible from the sensors, and also this is a lower risk solution if the performance proves to be satisfactory.

In the next phase Contact Level Fusion is planned with more low level imaging information for automatic target recognition. This architecture is given in Figure 2

In Phase 1, the Track Level MSDF system combines Radar, ESM, IFF, Link and Long Range (LR) target features provided by the SSAR. All data pre-processing including Radar and IFF Contact Level fusion is done as part of the sensor subsystems. LMCAN has already implemented such a data fusion system for the Canadian Patrol Frigate¹⁻³ (CPF). The differences for airborne applications include types of targets and background, timing requirements and types of data to be fused. The similarities include the fact that the pre-processing will remove significant amount of background clutter and a Nearest Neighbour or Jonker-Volgenant-Castanon (JVC) association algorithms are appropriate, and that Track-to-Track cross-correlation must be accounted for. The target identification will be estimated using the LMCAN's variety of the Dempster-Shafer algorithm, the Truncated Dempster-Shafer^{4,5} combining attributes available from the sensors in this phase.

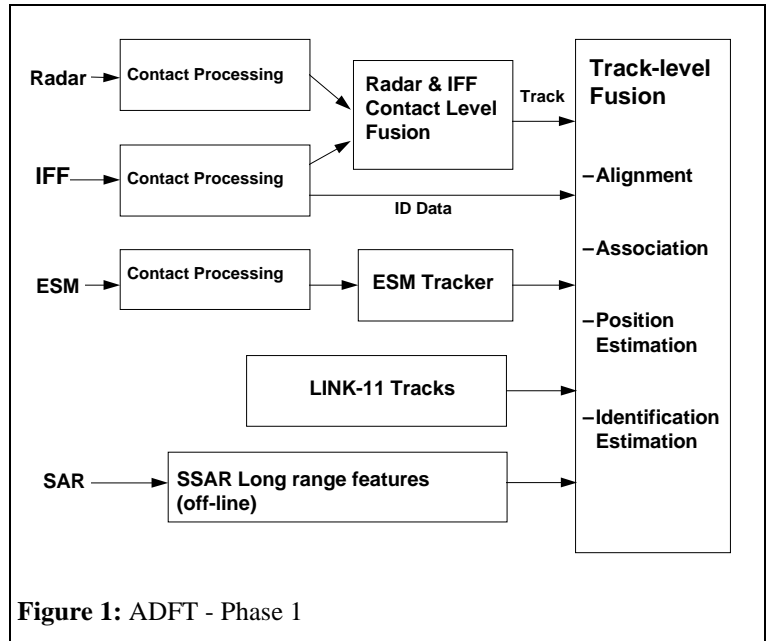


Figure 1: ADFT - Phase 1

In Phase 2, a Contact Level MSDF system will combine Radar, ESM, IFF, and Link data with SAR and FLIR image features. Here the SAR and FLIR feature extraction is not done to the point of object identification. By extracting simple features, and combining these with information from other sensors, the target identification can be more easily performed. This is shown in the following section.

The fusion architecture shown in Figure 2 for Phase 2 is a Hybrid architecture. There are two fusion centres: one performing a Contact Level fusion for position estimation from onboard sensors, and the other performing a Track Level fusion for kinematic fusion of ESM and Link data and for Target Identification fusing attribute data from all sensors. The Contact Level fusion centre will use Multi-Hypothesis Tracking (MHT) algorithm, while the kinematic fusion algorithms from Phase 1 system can be used in the Track Level fusion centre. Here too there will be Track-to-Track cross-correlation issues that must be accounted for.

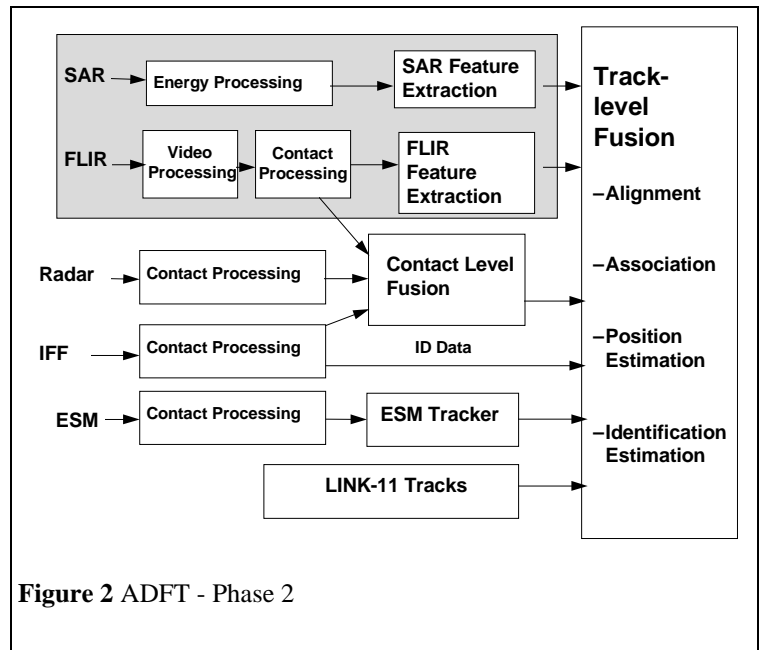


Figure 2 ADFT - Phase 2

There is an alternate fusion architecture that could be selected for this phase as shown in Figure 3. Here all the positional information is fused in one Contact Level fusion centre using an MHT algorithm. The main issues to be resolved in this case are due to the fact that: (a) not all the data coming into this algorithm are contacts (e.g. the Link data); and (b) not all covariances are available or straightforward to estimate. However in this case there will not be any Track-to-Track cross-correlation issues.

The analyses for specifics of these two approaches are currently in progress and will continue in parallel with the ADFT Phase 1 implementation and evaluation.

The Truncated Dempster-Shafer algorithm will again be used for target identity estimation in Phase 2. However in this phase there will be more complex features to combine to derive target identity.

The details of image analysis techniques and methods used to provide data into the Identification Estimation process from the imaging sensors are discussed in the following section.

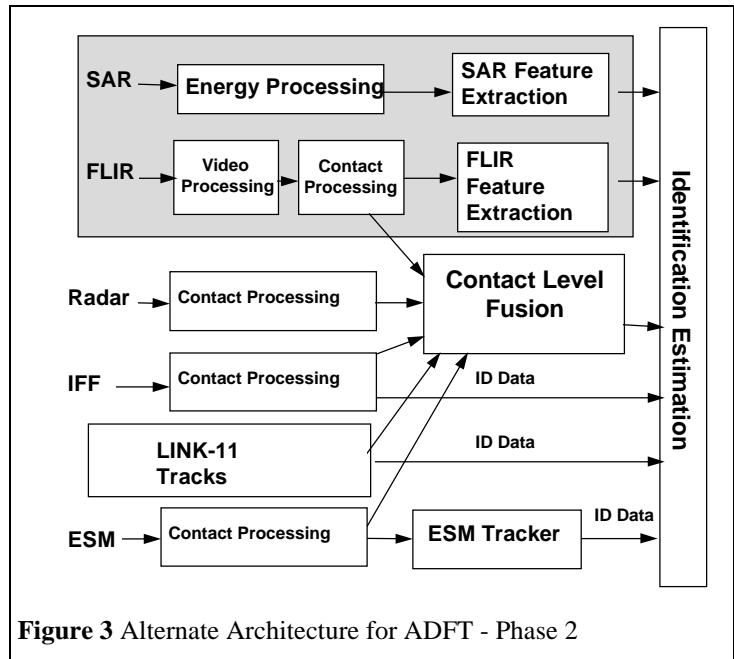


Figure 3 Alternate Architecture for ADFT - Phase 2

During Phase 2, the capture of real-time requirements will be performed on LMCAN’s Simulated Real-Time Environment (SRTE), which is currently designed to realize the same functionality for the naval CCS of the future Canadian Patrol Frigate.

In Phase 3, validation of the algorithms will be performed using real data from NATO exercises and ADM imagery as it becomes available, concentrating on the Spotlight modes (usually adaptive for naval targets but also, to a certain extent, non-adaptive for land targets), and sometimes using StripMap land imagery and Range Doppler Profiling (RDP).

3. IMAGE ANALYSES AND FEATURE EXTRACTION

Imaging sensors (SAR and FLIR) onboard the Aurora aircraft are planned to provide assistance in a variety of surveillance and sovereignty operations over the Canadian landmass and coastal areas [IEEE Canadian Review]. These include patrols over the remote northern areas, support in monitoring and enforcing the 200 mile economic zone, pollution monitoring and counter drug patrols. Such activities involve a great deal of digital image enhancement, segmentation, image analyses, target detection and classification. LMCAN has engaged a few R&D activities in these areas over the last 2 years, in collaboration with University partners⁶⁻⁹. The following paragraphs describe one of these activities that is specific to the feature extraction and classification of ships from airborne SAR images and fusion with non-imaging sensors.

The image analysis “box” in the ADFT system is designed to provide SAR-based ship target classification capability, through fusion (evidential reasoning) with other non-imaging sensors. Figure 4 illustrates the system LMCAN is studying for ship feature extraction and recognition from SSAR or ISAR images. Both target features and identification are sent to the fusion system, following a four-step hierarchical procedure and through small dedicated modular classifiers. This approach has numerous advantages: (1) low technical risk, (2) easier system design refinement and upgrade, (3) incremental implementation in ADFT and (4) easier and more efficient retraining procedures than for large classifiers.

The preliminary target detection step and the subsequent SAR antenna orientation are assumed to be performed with the help of the range-bearing information provided by the conventional radar mode. At the operator request, a potential ship target is imaged at high resolution using the SSAR or ISAR modes. For now, we assume that the SAR image generator provides adequate platform and target motion compensations to assure minimal blurring effects in the image (otherwise, we recognize that this can impose a severe limit in the classification performance). The resulting target image, along with the target range (slant-range), the platform altitude and the target aspect angle (which could be acquired from a previous long tracking time), constitute the input data in Figure 4.

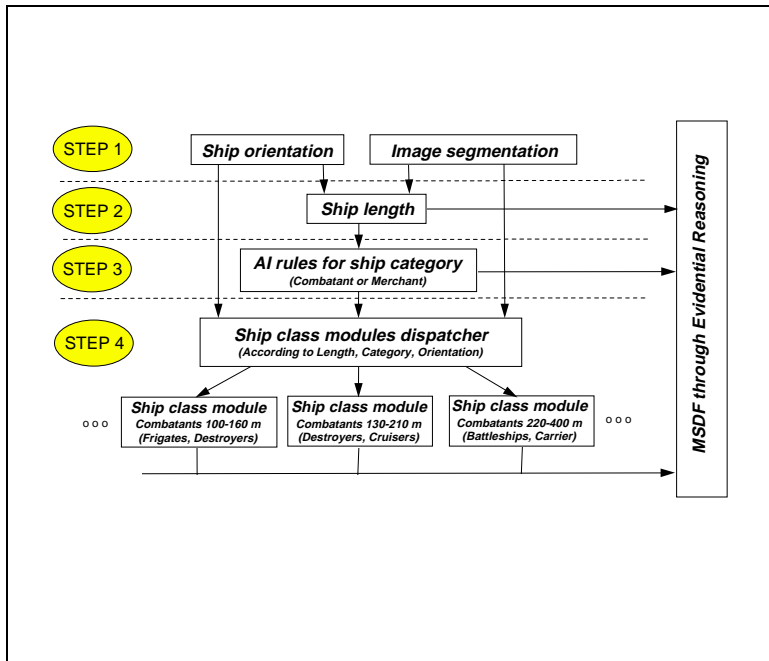


Figure 4: Hierarchical ship features extraction and classification from SAR images

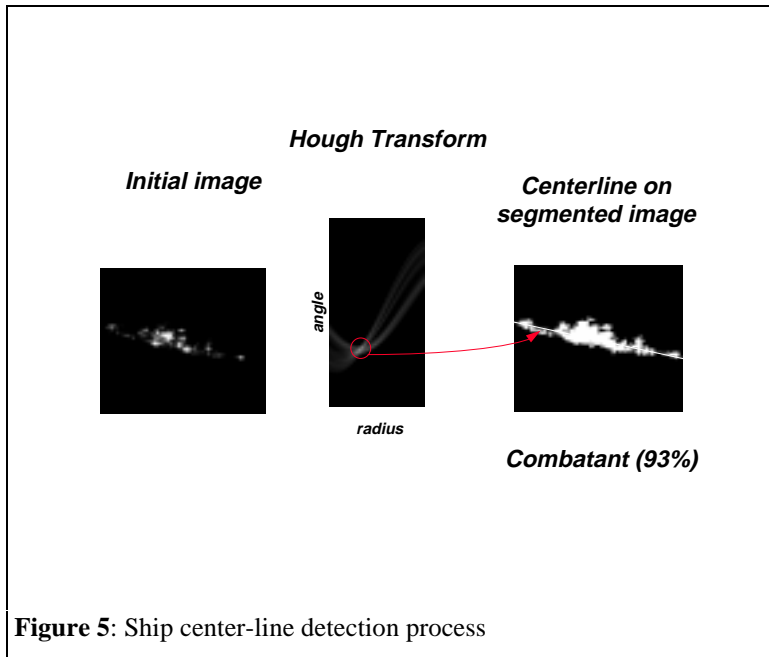


Figure 5: Ship center-line detection process

Step 1 in Figure 4 consists in segmenting out the target pixels from the ocean clutter. Two algorithms have been implemented and tested for that purpose. The first is based on a neural clustering scheme, called Probabilistic Winner-Take-All (PWTA), which aims to provide a neural approximation to the optimal Bayesian clustering⁷. As opposed to the standard k-means algorithms, the PWTA adapts the means *and the variances* of the clutters, which provides a better representation of the input distribution. The second segmentation algorithm is based on a combination of directional pixel intensity thresholding (to help removing linear strikes and artifacts caused respectively by rotating ship antenna and SAR image generator) followed by morphological operations that eliminate small non-connected regions and merge large ones.

Step 2 aims to calculate the ship length, which is a simple but very discriminating feature in ship classification. This feature is calculated only on image frames for which target eccentricity is high, which is the case for elongated objects like ships (from broadside or planar views). Image frames with low target eccentricity are discarded assuming they result from (1) a fail in the previous segmentation step or (2) a lack of target profile or plan information in ISAR images due to small pitch or yaw motions of the ship target. Ship length is calculated by identifying the target's end-points and using their respective coordinate information in the SSAR or ISAR image-plane geometry formula. Ship end-points are obtained by estimating the ship center-line from the maximum peak of the Hough transform of the segmented image¹⁰. If the ship is classified as "large ship" (lets say, ship length above 70 m),

then the original image is sent to the first target classification step (Step 3); otherwise it is labeled as small ship and not analyzed further (our current ship data-base does not contain small ships yet).

Step 3 is the first target classification step in the hierarchy. It has been designed to provide a high-level identification declaration of the ship category, that is military ship or merchant ship. This classification is based on the spatial distribution

of radar scattering along 9 sections of the ship target¹¹; military (more precisely military combatant ships such as frigate, destroyer, cruiser) have large structures, thus radar scattering, concentrated in the middle part of the ship while merchant ships (e.g. cargo, tanker) radar scattering comes mainly from the ship end regions. Because of a limitation in the number of real data, a set of 9-dimensional scatterer distributions are simulated and used to train and validate a 3-layer back-propagation Neural Net (NN) with an expert-system supervisor. The test set is made of real data only. The NN has 3 outputs corresponding to combatant, merchant or unknown. The outputs are normalized to provide a confidence level associated to the declaration (the sum of all confidence levels is 100%).

Ship length and ship category, along with the associated confidence levels, are sent to the fusion module where all sensor declarations and features are processed using the Truncated Dempster-Shafer evidential reasoning algorithm.

If a more accurate target identification is required, like ship type (e.g. frigate, destroyer, cruiser, cargo, tanker) or even ship class (e.g. Mackenzie-class frigate, Belknap-class cruisers), and if the confidence level on the ship category is sufficiently high (e.g. 85%), then the second step (Step 4) in the ship recognition process is engaged. This necessitates much more sophisticated classifiers¹² in order to (1) encode small differences in ship features and (2) deal with a large database (a typical database for ship surveillance may contain up to 1000 ships imaged at 180 aspect angles and 10 depression angles¹¹). Rather than training a large classifier that would most probably lead to convergence problems, we are considering a modular approach consisting of many small dedicated classifiers, each of them being specialized to recognize a subset of the ship database under a small viewing angle range. This approach has also the advantage of avoiding the training of the whole classifier if a new ship is entered in the database. During the recognition process, only modules corresponding to the radar viewing angle and the most probable ship classes are fired. This information may be obtained from the LMCAN's fusion system which has the capability of providing a list of the most probable target identifications. Two proof-of-concept

	Average of the highest network output	Average of the highest normalized network output
Class 1	0.97	0.92
Class 2	0.92	0.83
Class 3	0.93	0.86
Class 4	0.97	0.93
Confuser	0.90	0.51
Clutter	0.12	0.55

Table 2: Highest network output trained on 4 ship classes

classifiers have been tested by our University collaborators: low-resolution back-propagation NN and Principal Component Analysis (PCA) classifiers. The input to the NN is a small 10x10 image chip obtained by a 5:1 resolution reduction of the original radar image. As shown on Table 2, NN has demonstrated high robustness on the classification of simulated ISAR responses of 4 ship classes, one confuser ship and one clutter, for an aspect angle range of 35 degrees⁸. The ships of the first four classes are always assigned correctly with high reliability (second column). However, the confuser rather than the clutter was usually recognized as possible ship target (first column).

Although these results were encouraging, further experiments with a higher number of ship classes and the introduction of simulated speckle noise lead us to refine the approach. In order to increase the classifier robustness, our collaborators are now investigating a multi-resolution NN architecture. In this case, classification is made at the lower resolution and, if necessary, higher resolution images are processed.

At this time, the PCA classifier has been tested mainly for its robustness to systematic deformation in ship range profiles with the aim of clarifying the encoding process. It turns out that the most important ship features actually encoded with PCA is the ship length and the major ship structures. We are currently investigating whether this information is sufficient to discriminate between a set of 20 similar ship classes.

More specific test results about multi-resolution NN and PCA classifiers are under way and will be released soon⁹.

4. TESTBED ARCHITECTURE

LMCan has developed a unique software architecture which will permit integration and independent evolution, enhancement and optimization of various decision aid capabilities, such as Multi-Sensor Data Fusion (MSDF), Situation and Threat Assessment (STA) and Resource Management (RM). This architecture uses a Knowledge Based System (KBS) shell based on a BlackBoard (BB)-based problem-solving engine. This architecture is shown in Figure 6.

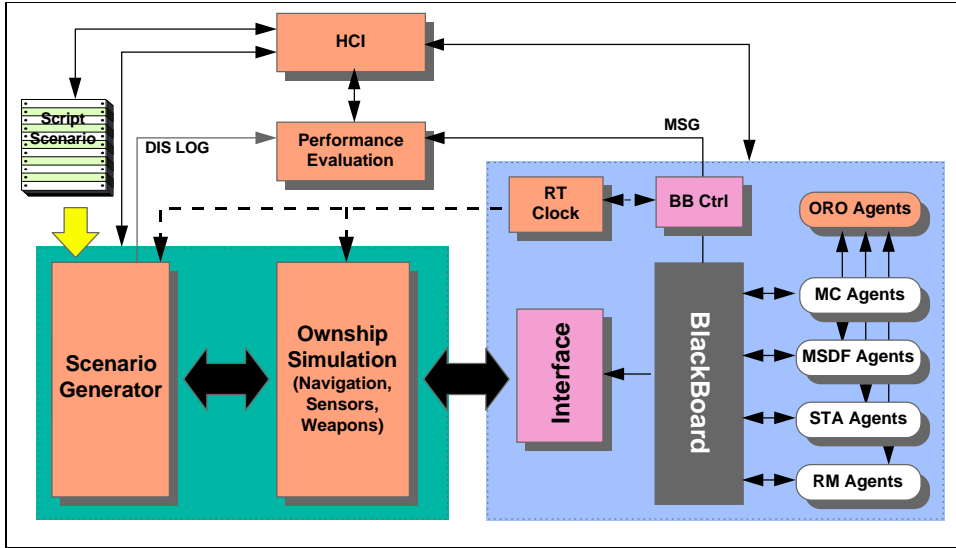


Figure 6: Testbed Architecture using a BB based KBS shell

The ADFT is being developed to function within this architecture. The development of the STA and RM agents is outside the scope of the ADFT's three phases. The ADFT includes:

- a) Infrastructure comprising simulation, performance evaluation and HCI, and
- b) MSDF applications.

This architecture will permit incremental developments and enhancements of the infrastructure including the input data simulation components for a selected Data Fusion implementation, as well as each

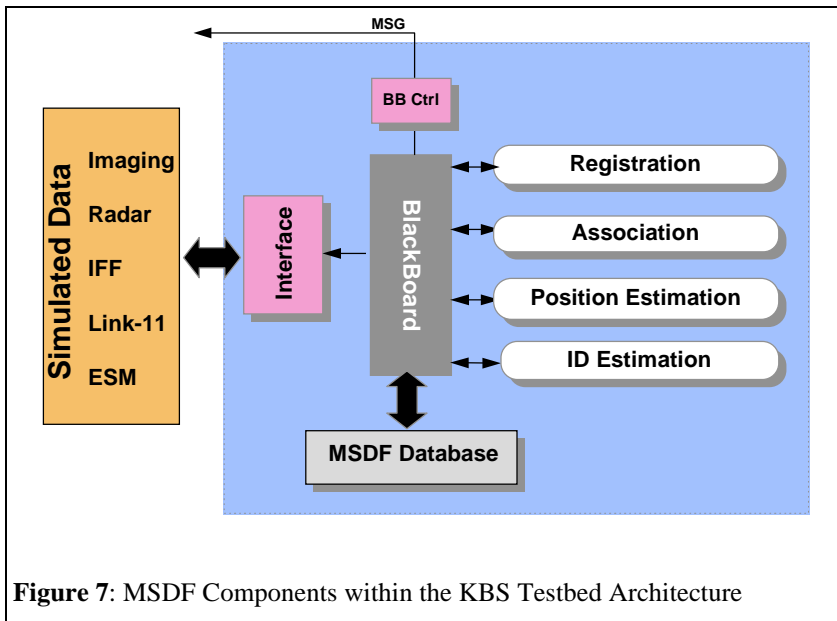


Figure 7: MSDF Components within the KBS Testbed Architecture

MSDF component independently. Figure 7 shows the input data simulation and MSDF components within this architecture, that will incrementally evolve and be optimized as described in section 2 of this paper.

The choice of this specific KBS architecture as the generic skeleton for an integrated MSDF/STA/RM System was made to satisfy a number of requirements such as:

- support both numeric and Artificial Intelligence (AI) techniques
- real-time efficient
- distributed (multiprocessor) environment
- design flexibility
- guaranteed real-time execution for MSDF/STA/RM components

The survey of the open literature concluded that at this point in time, there is no

commercially available KBS that could allow us to fulfill this set of requirements, hence we developed this BB-based KBS shell to become the generic architecture for demonstrating and evaluating incorporation of data fusion and other intelligent decision aids into all C⁴I systems that we become involved in (e.g. CPF), and to be able to leverage work done on different programs. An overview of this BB-based KBS shell is given below.

Generally speaking, a BlackBoard system, like any expert system computer program, is a problem-solving engine whose architecture is based on three main components¹³:

- the knowledge sources,
- a global data base (in this case, the BlackBoard itself),
- a control structure.

The **Knowledge Sources** (also referred to as Agents) are the parts of the system that contain the algorithms, facts, logical rules or heuristics representation needed to perform the ultimate task of the engine, namely to find a solution to a given problem. The goal of each knowledge source is to contribute pieces of information by modifying some control or data structure that is present on the blackboard, that eventually leads to a solution to the problem.

The **BlackBoard** itself is a global structure that is available to all knowledge sources in the system. It contains problem-specific data objects needed by, or produced by, the knowledge sources, and that are part of the solution space to the problem. These data objects can then be connected to others through named links or relations, or organized into more refined structures or hierarchies, for example by partitioning or introducing “layers” on the blackboard itself.

The **Control Structure** is arguably the most delicate part of the system, since any action can create new conditions that in turn require drastically different problem solving strategies. It is well known that the solution reached by a complex system generally depends on the sequence of intervention of the knowledge sources. A lot of control mechanisms are available to the developer; typically, they involve choosing a “focus of attention” that is triggered in priority, according to a ranking based on pre-defined evaluation criteria.

A BlackBoard environment such as the one described above can be viewed as a high-level set of routines and instructions, similar to a programming language, than can be used by a developer to implement knowledge rules in order to solve a specific problem. Currently the core of our BB-based KBS shell has been developed and integrated with some basic MSDF/STA/RM applications for the CPF. This system is totally generic, and as such could be used just as well to implement a Command and Control System or to play chess. It has been implemented in C++ rather than in a higher-level language (such as LISP, Smalltalk, ...) to satisfy the real-time requirement.

A very detailed discussion of the specific implementation of our BB-based KBS system will be part of a separate publication. Figure 6 above illustrates all the basic components of our system: a set of agents, a blackboard and a control mechanism.

The BlackBoard can be viewed as a global database accessible by all agents. It contains all the active processing data (interfaced, generated and modified by the agents). The agents performing knowledge transformation on the blackboard are independent from each other and are “self-activating” or “data-driven”, in the sense that their activation is triggered by new pieces of data being put on the blackboard. The agents “examine” the state of the BlackBoard, create new data structures and, when it is necessary, modify them by taking a data structure of one type and producing a data structure of another type. In the case of an application to a CCS, the classical model had to be modified for more flexibility, and new data types definitions with labels and attributes were introduced; this allows the derivation of new data types (so-called “derived” agents) through cloning or grouping.

Self-activating agents typically cause an explosive demand on the system resources. In the current system, this is being taken care of by adding activation constraints to the agents, and by providing the controllers with a negotiation mechanism in order to ensure timely response from the agents. These controllers were designed so the system can be implemented on a distributed architecture environment. The BlackBoard itself can also be partitioned into context-dependent “areas of interest” (dimensions), giving rise to the possibility of efficient associative data processing.

This system can be quite complex and obscure for an external user that would want to implement his own set of rules on the system. In order for our KBS system to be easier to use, its complexity has to be hidden from the end user by the addition of *user libraries*, which we have started to develop. These libraries will lie on top of the *KBS shell*, and will eventually hide most of the low-level details by providing a higher-level interface to the system. This is a dynamic process, with the system starting with a “naked” shell and growing in an “onion-like” structure, incorporating cumbersome or frequently used functionality into more generic and user-friendly functions.

One must keep in mind that with more layers, the system becomes more user-friendly but less generic. For instance, the outer layers of the user libraries should allow the user to easily implement new agents, but only in the specific context of a Command and Control System. However, the fact that these libraries follow a concentric shells structure allows the KBS system to remain fully reusable, since the outer shells may or may not be used for a given application.

5. CONCLUSION

Many issues must be addressed and resolved before data fusion and other technologies can become part of Combat Systems of existing and future platforms. These issues involve sensor pre-processing and sensor data quality/completeness, overall fusion architecture to provide the desired system performance and overall system implementation architecture to guarantee real-time execution. LMCAN's efforts for resolving these issues for airborne surveillance onboard an Aurora Maritime Patrol Aircraft have been described. An evolutionary incremental approach for demonstrating and evaluating solutions of increasing sophistication has been selected, where the solutions in each phase could be used if proven satisfactory.

Due to the unique BB-based architecture of the testbed, the implementation is very modular and can be re-used for other programs both in terms of the data fusion algorithms, as well as the infrastructure.

ACKNOWLEDGEMENTS

We acknowledge our University collaborators (J.-Y. Potvin, M. Gendron, J. Patera, F. Drissi Smaili and V. Gouaillier from Université de Montréal, S.D. Blostein and H. Osman from Queens University,) for their help in these R&D activities through partial funding from the Natural Sciences and Engineering Research Council (NSERC) of Canada.

We would like to thank the Defence Research Establishment Valcartier (DREV) collaborating team (E. Bossé, R. Carling, B. Chalmers and J. Roy) on the MSDF/STA/RM implementations within the BB-based KBS shell for the CPF, and the Aerospace Radar and Navigation Section at the Defence Research Establishment Ottawa (DREO) for providing real image data and in particular R. Klepko for his great help in understanding practical issues in automatic classification of ships from SAR images.

REFERENCES

1. F. Bégin, E. Boily, T. Mignacca, E. Shahbazian, and P. Valin, "Architecture and Implementation of a multi-sensor Data Fusion Demonstration Model within the real-time Combat System of the Canadian Patrol Frigate", *Proc. of AGARD symposium on Guidance and Control for Future Air-Defence Systems*, Copenhagen, **AGARD-CP-555**, pp. 28.1-28.8, 1994.
2. J. Couture, E. Boily, and M.-A. Simard, "Sensor Data Fusion of Radar, ESM, IFF and Datalink of the Canadian Patrol Frigate and the Data Alignment Issues", *Proc. of SPIE's Aerospace Sensing / Orlando'96, Signal and Data Processing of Small Targets*, Oliver E. Drummond, ed., **SPIE Vol. 2759**, pp. 361-372, 1996.
3. P. Valin, J. Couture, and M.-A. Simard "Position and Attribute Fusion of Radar, ESM, IFF and Datalink for AAW missions of the Canadian Patrol Frigate", *Proc. of IEEE / SICE / RSJ International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI'96)*, Washington D.C., pp. 63-71, 1996.
4. M.-A. Simard, P. Valin, and E. Shahbazian, "Fusion of ESM, Radar, IFF and other attribute information for target identity estimation and a potential application to the Canadian Patrol Frigate", *Proc. of AGARD 66th symposium on Challenge of Future EW System Design*, Ankara, **AGARD-CP-546**, pp. 14.1-14.18, 1993.
5. M.-A. Simard, J. Couture, and E. Bossé, "Data Fusion of Multiple Sensors Attribute Information for Target Identity Estimation using a Dempster-Shafer Evidential Combination Algorithm", *Proc. of SPIE's Aerospace Sensing / Orlando'96, Signal and Data Processing of Small Targets*, Oliver E. Drummond, ed., **SPIE Vol. 2759**, pp. 577-588, 1996.
6. L. Gagnon, and F. Drissi Smaili, "Speckle Noise Reduction of Airborne SAR Images with Symmetric Daubechies Wavelets", *Proc. of SPIE's Aerospace Sensing / Orlando'96, Signal and Data Processing of Small Targets*, **SPIE Vol. 2759**, pp. 14-24 (1996); L. Gagnon, and A. Jouan, "Speckle Filtering of SAR Images - A Comparative Study Between Complex-Wavelet-Based and Standard Filters" (submitted SPIE Annual Meeting, San Diego, 1997).

7. H. Osman, and S.D. Blostein, "SAR Imagery Segmentation Using Probabilistic Winner-Take-All Clustering", *Proc. of SPIE's Aerospace Sensing / Orlando'96, Algorithms for Synthetic Aperture Radar Imagery III*, **SPIE Vol. 2757**, pp. 217-226, 1996.
8. H. Osman, and S.D. Blostein, "SAR Image Processing Using Probabilistic Winner-Take-All Learning and Artificial Neural Networks", *Proc. of IEEE ICIP*, **Vol. II**, pp. 613-616, 1996.
9. L. Gagnon, and V. Gouaillier, "Ship Silhouette Using Principal Component Analysis"; H. Osman, S.D. Blostein, and L. Gagnon, "Classification of Ships in Airborne SAR Imagery Using Backpropagation Neural Networks" (submitted SPIE Annual Meeting, San Diego, 1997).
10. S. Musman, D. Kerr, and C. Bachmann, "Automatic Recognition of ISAR Images", *IEEE Trans. Aerospace Electron. Systems*, **Vol. 32**, pp. 1392-1403, 1996.
11. R. Klepko, "Automatic Pattern Classification of Airborne SAR Images of Ships", *DREO Report #1283*, (SECRET)
12. M. Menon, "An Automatic Ship Classification System for ISAR Imagery", *Proc. of Applications and Science of Artificial Neural Networks*, **SPIE Vol. 2492**, pp.373-388, 1995
13. H. Penny Nii, "Blackboard Systems: The Blackboard Model of Problem Solving and the Evolution of Blackboard Architectures", *The AI Magazine*, **7** (2), pp. 38-53.