Holonic Control of Distributed Military Sensors for Littoral Surveillance

Abderrezak Benaskeur
Defense R&D Canada – Valcartier
2459 Pie XI North
Quebec, Quebec G3J 1X5
CANADA

Tel. +1 (418) 844-4000 ext. 4396

Abderrezak.Benaskeur@drdc-rddc.gc.ca

Peter McGuire, Geoff Liggins, Peter Wojcik
C-CORE, Bartlett Building
Morrissey Road
St. John's, Newfoundland A1B 3X5
CANADA

Tel. +1 (709) 737-3103

Peter.Wojcik@c-core.ca

Robert Brennan
School of Engineering
University of Calgary
2500 University Dr. N.W.
Calgary, Alberta T2N 1N4
CANADA

rbrennan@ucalgary.ca

ABSTRACT

This paper outlines a holonic control design, simulation and results for sensor management in a military setting. It details the issues involved in military sensor and holonic control, and demonstrates a holonic control solution for sensor management in littoral surveillance applications. An overview of the problem scope is presented, followed by a proposed holonic control architecture design. The proposed design is used to develop a software simulation using a military scenario in which the holonic control system is employed in the sensor management role. The results of this simulation are then presented.

1.0 INTRODUCTION

The military typically operate in dynamic, semi-structured and large-scale environments. This reality makes it difficult to detect and track all targets within the Volume of Interest (VOI), thus increasing the risk of late detection of threatening objects [1]. This can be very critical to own-force survival in high target-density operations, such as in littoral environments. A key challenge facing the military operators, in these contexts, is the focus of attention and effort, that is, how to make the most effective use of the available, but scarce, resources to gather the most relevant information from a dynamic environment.

Military platforms are generally outfitted with a set of sensors that provide a wealth of data when properly managed. Efficient Sensor Management (SM) can aid the processes of information gathering by automatically

allocating, controlling, and coordinating sensing resources to meet mission requirements. This paper presents results of a project that aims to assess the applicability of holonic control paradigm to SM in military surveillance operations. The holonic approach, originally devised for use in manufacturing environments [2], is deemed well suited to SM applications, in part, due to its ability to use autonomous agents approach while maintaining a strict command hierarchy.

A holonic control for SM was designed and implemented in the Matlab environment. To evaluate the performance of the latter, a scenario has been developed. Therein, a group of military platforms, located off the coast of Canada, is tasked with conducting surveillance operations for force protection in the port of Victoria, B.C. A control strategy is employed that maintains high quality tracks for targets that pose an actual threat, while lowering quality for all other tracks. Very encouraging results were generated.

The paper describes briefly the SM problem in Section 2.0, discusses the developed holonic structure, and explains the operating methodology in Section 3.0. An overview the surveillance application and the related scenario are discussed in Section 4.0. The simulation software and the obtained results are presented in Section 5.0.

2.0 SENSOR MANAGEMENT

The objective of any surveillance mission is to gather information about the presence and activity of all objects within the VOI. The information gathered is used to build a representation of the situation of interest. SM aids the surveillance process by directing sensing resources in a manner to acquire data that is the most relevant to mission objectives. The military organization by its very nature has a hierarchical command structure to manage its resources, men, and equipment. SM, a subset of this command structure, is not only hierarchical but also recursive. An example of such a recursive hierarchy is shown in Figure 1 for a typical naval Task Force configuration.

Figure 2 summarizes the different SM problems independently of the level they belong to. These problems are described in more details below. Note that this list is not intended to be exhaustive.
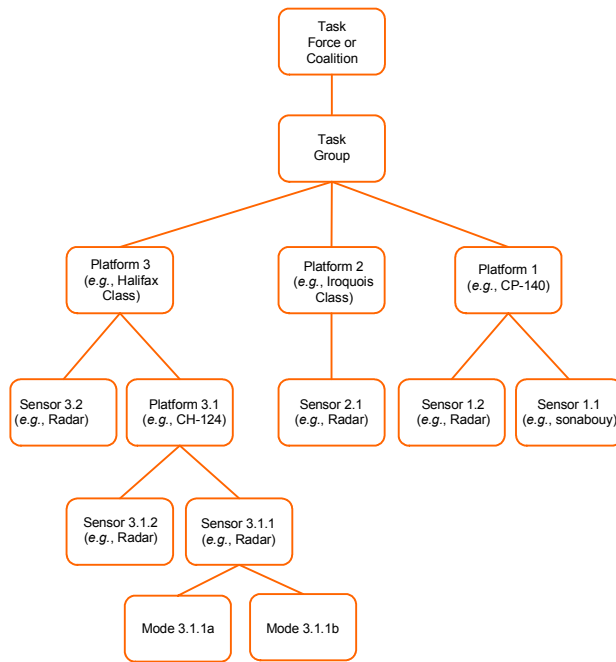
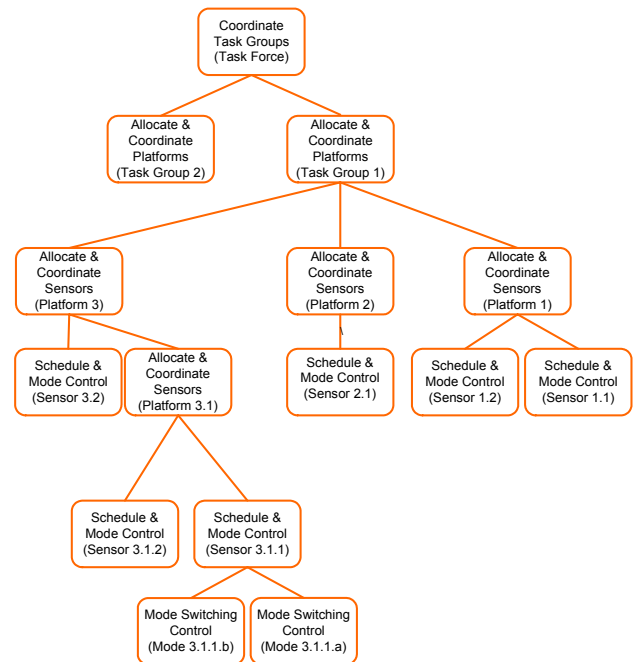**Figure 1: Hierarchy of Naval Sensing Resources**

Figure 2: Hierarchy of SM Problems

- **Allocation** – This is concerned with the determination of which sensing resource, or set of resources, to use to achieve the sensing objectives. By breaking these objectives down into a series of tasks, the SM needs to determine the most suitable resource to allocate to each task.

- **Coordination/Cooperation** - If a sensing resource when in operation is in conflict with other resources then the SM must determine which resource is more important and prevent the others from operating or must allow for some schedule to allow one resource to operate for a period of time and then the other. This defines the coordination, or conflict resolving, problem. Dual to this problem is the cooperation, where synergy among complementary resources is maximized by the SM module.

- **Scheduling** - Scheduling is the designation of time segments to specific tasks or activities, the nature of which is defined during the allocation or coordination stages. Scheduling typically uses time as its base variable; tasks are expected to start at a specified time and to execute for a fixed time interval.

- **Mode Control** – In case of sensors offering multiple modes, the SM should make use of the most optimal mode for the tasks being done provided that there is no other overriding reason not to.

- **Mode Switching Control** - While changing sensor modes, the data stream may be halted during the transition. The SM must address whether it is more important to maintain operation in possibly a sub-optimal mode while maintaining a live data stream or to change to a more optimal mode.

- **Others Problems** - Other potential issues in tactical surveillance, for which strategies within SM are required at all levels, would include: emission control: (SM system must trade off the gathering of more complete information using active sensor over self-security); failure recovery (SM must alter the sensing allocation and schedule in case of disabled or diminished sensing capability); and contingency handling (SM must address when and how to make the necessary changes If situation/objectives change).

SM requires a control architecture that matches the underlying command structure (see Figure 1). At the highest level, decision-making is based on very high-level information. As one descends the tiers in the hierarchy, the decision-making becomes more focused. As data moves up the hierarchy, it is transformed into information necessary for high-level decision-makers. The selection of the appropriate control architecture, to address the SM problems, is discussed the next section.

3.0   HOLONIC CONTROL

In distributed control schemes, each node in the architecture has a controller that allows it to works collectively with its neighbours to achieve some overall goal. There are several organizational structures that allow nodes and their controllers to work together. Given the hierarchical and recursive nature of the SM problems, the desirable characteristics of a control architecture to address them are : 1) hierarchy to account for a clear chain of command; 2) adaptability to the current situation; 3) sufficient autonomy of each node to perform its function without being encumbered from actions taken at the top level; 4) sufficient robustness to maintain operations even if (elements of) the network are incapacitated; and 5) recursiveness: where each node could be composed of one or more nodes of a lower abstraction level. The following presents a list of candidate control architectures to address SM problems (see Figure 3).

- **Centralized** - not suitable because it requires that a central node be kept intact at all times; this is a significant risk in the military context. It has the advantage of relatively simple control; however, if the situation for which it is configured changes then it requires a massive effort to reconfigure it. Centralized architectures are characterized by high communications requirements, a high computational burden at the central node, and a lack of general robustness and flexibility.

- **Decentralized** - (or heterarchical) is also unsuitable for the problem at hand because of its lack of structure. This makes the heterarchical architecture relatively robust because there is very little to break, but also makes it very difficult to control, which can lead to an undesirable chaotic behaviour.

- **Hierarchical** - is a top-down decomposition of tasks and division of labour approach. This structure is good in that it forces an expected behaviour; but it is inflexible and branches can become uncontrollable if an intermediate element is incapacitated. The degree of autonomy of an element in a hierarchy is quite limited. The top-down approach is convenient for planning purposes and the dissemination of instructions and goals. However, if the situation changes significantly then a new entire plan must be derived, which can be a significant computational and communications burden.

- **Federated** - is a compromise between the hierarchical and the heterarchical structures. Like the heterarchical approach the nodes have a high degree of autonomy but form a structure through communications and the use of specialized middle nodes. This approach has improved robustness and flexibility over the other architectures but does not allow for the dynamic restructuring that is an integral part of holonic architectures, described below.

- **Holonic** – is a hybrid architecture that takes the best of different architectures and avoids many of their pitfalls. The holonic architecture takes advantage of the distributed capabilities from classical Multi-Agent Systems (MAS) while incorporating the benefits of the hierarchical command structure that allows for strong goal orientation. This architecture is discussed in more details in the next section.
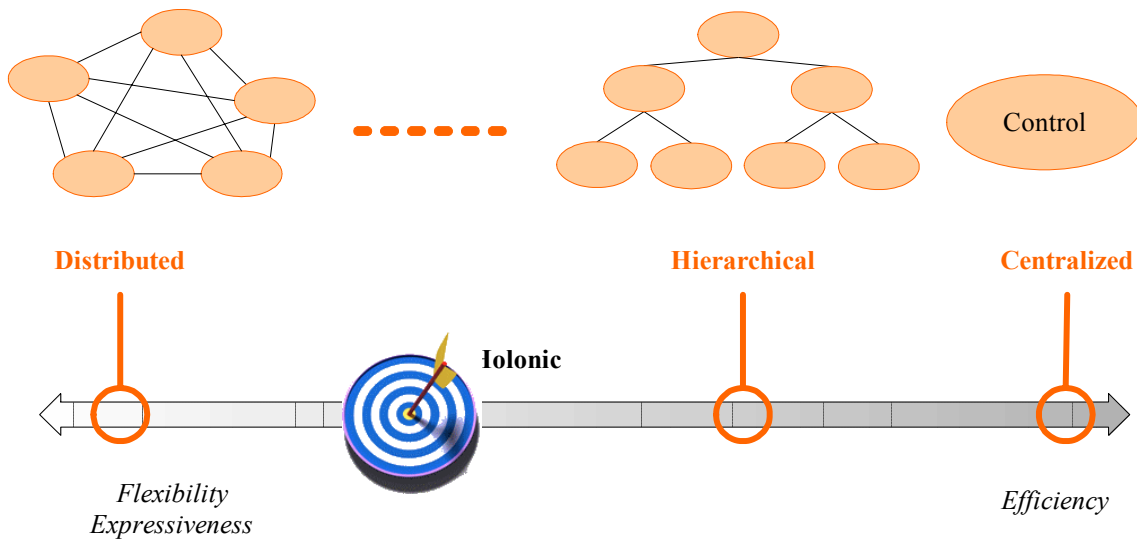
Figure 3: Architectures Spectrum

3.1 Holonic Architecture

To explain complex biological and social systems, Arthur Koestler [3] made two key observations: 1) these systems evolve and grow to satisfy increasingly complex and changing needs by creating stable intermediate forms which are self-reliant and more capable than the initial systems; and 2) in living and organizational systems it is generally difficult to distinguish between *wholes* and *parts*: almost every distinguishable element is simultaneously a whole (an essentially autonomous body) and a part (an integrated section of a larger, more capable body). To explain this concept, Koestler suggested a new term: *holon*, from the Greek *holos* meaning whole and the suffix *on* implying particle as in *proton* or *neutron*.

3.1.1 Characteristics of Holons

A Holonic System (HS) *consists of autonomous, self-reliant units, called holons* that co-operate to achieve the overall system objectives [4]. Some key properties of a HS developed from Koestler's model are [2]: 1) *autonomy* - the capability of a holon to create and control the execution of its own plans and/or strategies (and to maintain its own functions); 2) *cooperation* - the process whereby a set of holons develop mutually acceptable plans and execute them; 3) *self-organization* - the ability of holons to collect and arrange themselves in order to achieve an overall system goal; and 4) *Reconfigurability* - the ability of the function of a holon to be simply altered in a timely and effective manner. Another important holonic concept is the notion of functional decomposition. The complexity of dynamic systems can be dealt with by decomposing the systems into smaller parts. A consequence of this is the idea that holons can contain other holons (*i.e.*, they are recursive). Problems are solved by *holarchies* (hierarchies of holons), or groups of autonomous and co-operative basic holons and/or recursive holons that are themselves holarchies. The recursive and hierarchical structure of holonic architecture and its ability to generate dynamic linkages to form an impromptu command structure to perform a task make it well suited to the above described military SM problems.

3.1.2   Holons and Agents

The common thread that runs throughout the work on HSs is the close link between MAS and HSs. Given this link, various co-operation, communication and organizational techniques from the MAS world can be used to implement autonomous, co-operative and recursive agents (*i.e., holons*). HS can be considered a general paradigm for distributed intelligent control, whereas MAS are regarded as software technologies that can be used to implement HSs. There are similarities and differences between holons and agents. One of the major difference concerns the recursiveness. A holon may be composed of other holons, while there is no recursive architecture as such in MAS.

3.2   Sensor Management Using Holonic Control

The holonic control paradigm can be applied to the challenge of implementing a SM system. This approach will maintain a command structure similar to the one used by the military organizations. The primary benefit of holonic control is its ability to form a localized structure, or *holarchy*, to address needs as they arise. Holons have a certain degree of autonomy that allows them to make decisions of limited scope. The flexibility displayed by HSs is the result of the combined behaviour of the holarchy and not the actions of individual holons.

The concept of applying holonic control methodology to the task of SM is illustrated in Figure 4. The proposed holonic control architecture is decomposed into three main levels: sensor, platform, and group. The levels are related to each other in a recursive hierarchical manner typical of HSs. The sensors represent the lowest level of the hierarchy. Each platform coordinates the sensors that are located aboard it, but do not control the sensors aboard other platforms. Likewise, the group level manager coordinates sensing activities between platforms but does not directly manage the sensors aboard those platforms. Figure 4 represents a single SM system that may be located at any level in the control hierarchy. The following paragraphs describe, briefly, the main elements (holons) of the holonic SM system.

- **Service Interface & Command Holon (SICH)** – is the most important holon that acts as facilitator. It interfaces with the requesting holons to define the constraints of their request and then spawns a task holon. The SICH then releases this task holon into the network. When the task holon returns, the resulting information is presented to the requesting holon and the corresponding task holon that gathered the information is then killed off. A fixed communications hierarchy is defined with respect to the SICHs. SICHs require constant two-way communications with each other within the hierarchy in order to keep track of resource availability and to transfer information up the chain of command. This communications hierarchy is defined according to functional significance, *i.e.* group-level SICH connects to platform holons while each platform SICH connects to its own internal resources.

- **Task Holon (TH)** - is a holon that once created autonomously makes use of the resources that are either allocated to it or that it negotiates for itself. This holon differs from other holons in the network as it only exists for the duration of the request and is in effect a roaming client in a network of services. The main purpose of this class of holons is to utilize resource holons to fulfill its mandate. THs establish communication links with the resources directly below it and is also in communication with the SICH above it.
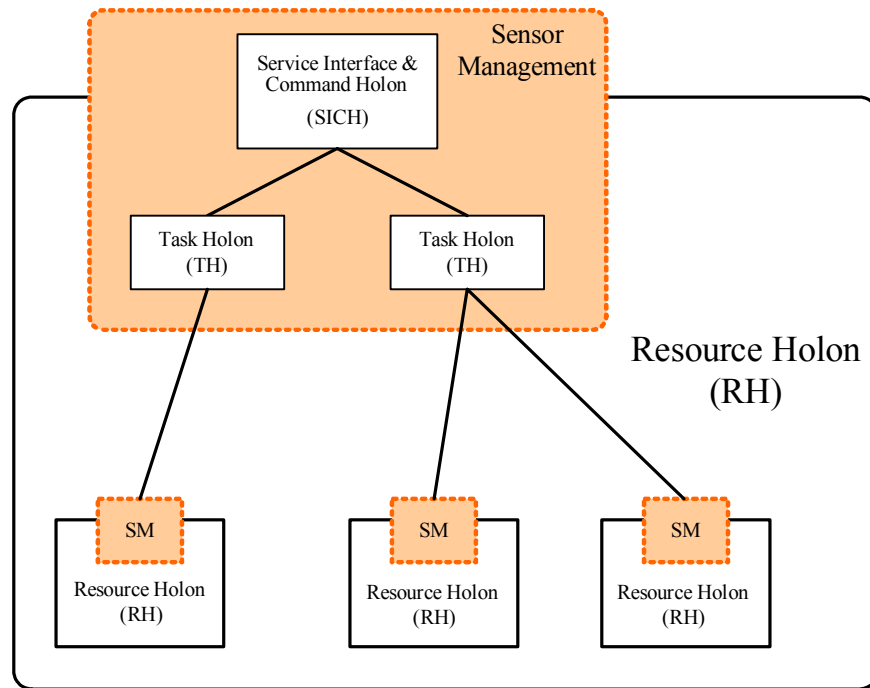
Figure 4: Recursive Holonic Structure for SM

- **Resource holons (RH)** - are the systems and devices that are being coordinated at each level in the hierarchy (see Figure 1). Complex resources, such as platforms, would be subdivided into holons that represent functional systems on board that must function together but are well defined in their own right. Simple resources, such as a single sensor, are at their most basic level already and would not benefit from further subdivision. RHs make use of the SICH that act as an interface to the external holarchy through which THs make their request for access to resources.

To explain how the structure of Figure 4 operates, a SM holon at task group level is considered. When a task request is made to the group holon, its SICH acts as a mediator and negotiates with the requesting RH whether it can service that task. If the task is accepted by the holon then a TH is spawned. This new TH negotiates with the RHs (*i.e.*, the platforms is this case) that it is aware of in order to complete its task. The TH can be killed off by the SICH if it has been supplanted by a newer TH. Information is distributed through the group holon by the TH and the SICH. If a TH cannot complete its objective because a RH has recently become unavailable, then the TH will look for alternate solutions by negotiating with other RHs (*i.e.*, platforms) for a replacement or completing its task without the resource. If no solution is found, the TH will report to the SICH above it that it will find an alternate solution either by creating an additional TH to aid the first one or by terminating the TH and creating a new one. In this way, problems are addressed locally and information about the problem is propagated up through the holarchy to the level where it is needed.

The SICHs are connected to one another through fixed links, while THs are responsible for making the dynamic hierarchical links and operating within these formed links. A TH making a request of a SICH negotiates access to available RHs. The result of a successful negotiation allows direct access to the allocated portion of the RHs. The RHs comprise the next level down in the holarchy. These holons have the same internal general structure as the group holon (see Figure 4).

This HS consists of a loosely defined architecture corresponding to the relative scales of the problems to be addressed. The system is mostly a distributed one, but a hierarchy is imposed to demark logical division of effort and resources. Decomposing the problem is done according to the level of abstraction. For example, at the higher levels, the holons are concerned with "big picture" building problems. At the lower levels they are concerned with more detailed (and arguably simpler) problems.

4.0 SURVEILLANCE APPLICATION & SCENARIO

This section summarizes the operation of the proposed holonic control system for SM in military settings. A force protection scenario is assumed in this description. This section discusses the role of SM in acquiring and maintaining target tracks across multiple platforms. In the presented scenario, the platforms perform most of the sensing activities, including: detecting targets, tracking targets, and modifying sensor configurations based on analysis of the local situation. The group level SM role is to selectively acquire data from the platforms and modify platform sensing-operations based on a group level situation analysis.

4.1   Initialization

At the beginning, each platform will first be required to configure their sensors for searching. This proceeds with the platform level SICHs creating configuration holons that control their sensors. In parallel, the platform level SICH creates a number of search holons in order to detect any targets in the sensing domains. Configuration holons, when created, have priority access to the sensors.

4.2   Target Tracking

When a search holon reports a detection, the platform level SICH issues a tracking holon that maintains the target track. The platform SICH then creates a message holon in order to notify the group level SICH. The latter, in response, creates a monitor holon that intermittently updates the group level track database. There are two levels of Situation Analysis (SA) that are implemented; group level and platform level. The platform level SA is used to balance resource usage amongst the tracking tasks. At the group level, the communication limitations change the nature of the SM. Situation analysis provides therefore a means for balancing the limited communications resource.

4.3   Cue/Handoff Events

There are two main types of cue/handoff events that can occur: those due to targets transiting between sensor domains aboard a single platform, and those due to targets transiting from one platform sensing domain to another. Within a single platform, the tracking holon responsible for a transiting target usually can handle the handoff by simply acquiring service on a different sensor as the situation warrants. Between platforms, cue/handoff requires the coordination at the group level. In this case, it is up to the group SA to recognize the impending event and cue a search or handoff the track to the receiving platform. If the track will be lost or degraded due to too much time passing before the target reaches the receiving platform, the group level SICH will issue a group level search holon that will delegate a search task to the receiving platform. The platform level SICH will, in turn, create a platform level search holon that will attempt to reacquire the target track.

4.4   Load Balancing

In times of high load, tracking unimportant objects with more than one platform can be unnecessary. Releasing the platforms from performing redundant tasks makes these platforms more accessible to other

tasks that may require their attention. This is referred to as load balancing. The load-balancing approach employed here uses the list of track holons that are currently running to find objects that are being tracked by more than one platform. If the duplicate tracks are not involved in a cue/handoff and the objects do not pose a threat to own-force, then the duplicate tracks are removed.

4.5 Threat Evaluation

The holonic architecture and the associated threat assessment feedback aids in the refinement of sensing resource utilization. In particular, the feedback is used to rank the sensing tasks (*i.e.*, set priority) and establish sensing objectives (*i.e.*, track quality). When sensor loads are near maximum, these task attributes play a significant role in insuring the most threatening tracks are maintained regardless of other activity in the sensing region. For the scenario under consideration, the threat evaluation (relative to the protected forces, assumed in the port of Victoria, B.C.) uses primarily the closest point of approach and time to closest point of approach to compute the level of threatiness. Also are used the intelligence information and behaviour of the objects with respect to commercial transport lanes while approaching Victoria.

4.6 Scenario

In order to demonstrate the holonic control, the scenario "Surveillance/Control of Canadian Territory and Approaches", from the Departmental Force Planning Scenarios[1], was identified as a suitable candidate. The scenario has been modified to better illustrate elements of the SM design and force protection problematic.

The area of interest is limited to the port of Victoria and surrounding sea approaches. The surveillance region of operation is limited to the triangular section of sea bounded by the coastline and the predefined limits of responsibility as depicted in Figure 5. This area will be monitored with a group of platforms and a single ground station located at Victoria. The platforms perform a search and cue the ground station upon detection. In order to reduce its signature, the ground station is set not to perform target search itself, but rather rely on the referral (cue/handoff) from the platforms.

One of the main difficulties with this mission is the presence of a large number of spurious objects. Objects are constantly entering and leaving the sensing domain of the platforms, and it may become difficult or impossible to track all of them simultaneously while maintaining a search capability for new objects. The proposed SM will aid this process by tailoring the use of the sensing resources based on an assessment of the likelihood of the target being a threat. Without this feedback, all targets would be given equal attention by the sensing resources. With feedback, those targets that are deemed unimportant to the force protection mission will be tracked less closely providing greater sensing resources for the important ones. This provides significant advantage as the number of targets to be tracked increases.

The mission will employ four platforms, each of which is equipped with a single Electronically Scanned Aperture (ESA)-type sensor. Figure 5 illustrates the placement of the platforms within the covered area, beyond the sensing range of the base station. The sensing ranges of these platforms do not overlap. The platforms are positioned such that any object that approaches Victoria within the area of responsibility will be likely to cross their detection region. This arrangement of platforms will provide early detection of potential threats, without presenting a large radar signature from the ground station. The ground station will only track the most significant objects, while the platforms will track all objects within their sensing range. To reduce emissions, these platforms will use only 20% of their sensing capacity to perform their tasks.

---

[1] http://www.vcds.forces.gc.ca/dgsp/pubs/rep-pub/dda/scen/intro_e.asp
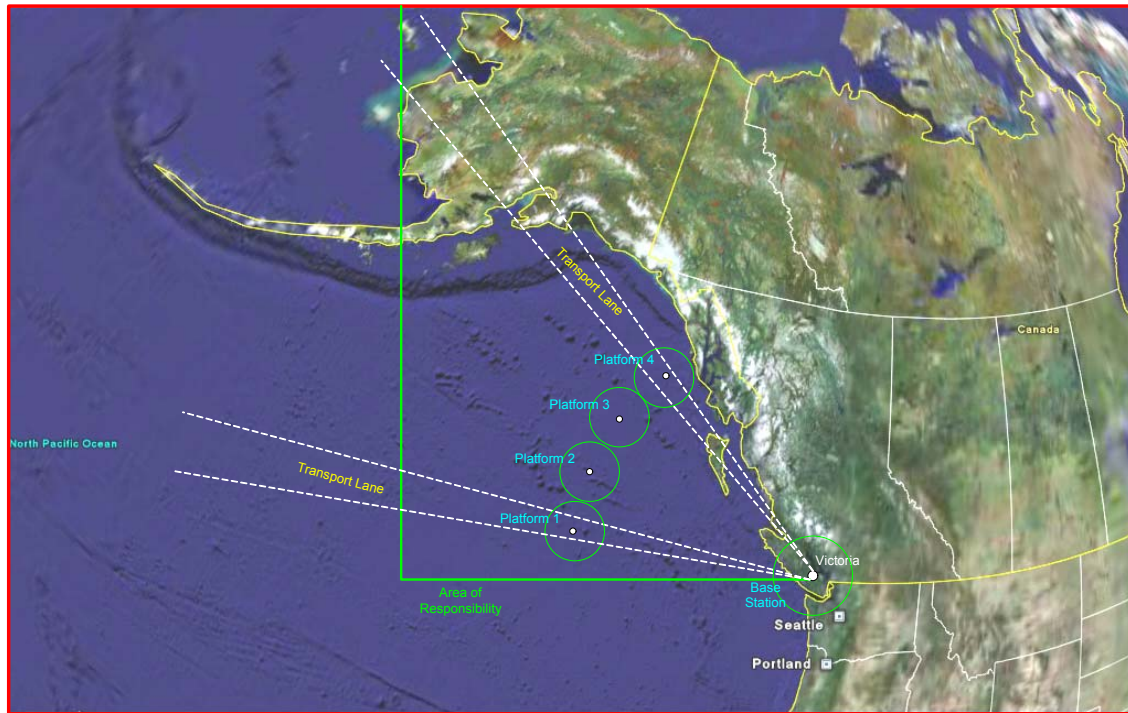
Figure 5: Area of Responsibility

The simulation is restricted to a single target type, namely aircraft. Targets are grouped into two categories, incoming (headed toward Victoria) and outgoing (headed away from Victoria). Incoming targets not in the transport lanes will be rated highly likely to be threats, and will need to be tracked closely by the platforms. Outgoing targets and targets in the transport lanes will have their threat level set to zero. They serve as a complicating factor in the mission as they must also be tracked though they are not as important to the mission. They will increase the sensor loading, thereby reducing the amount of sensor time available for searching. An additional important incoming target designated "the actual threat" is simulated with a heading directed towards Victoria.

Each target is given a heading that sends it directly towards or away from Victoria; however, each heading is also subject to a random perturbation of up to ±10 degrees. The simulation will be run for a time that is sufficient for the targets to pass completely through the platform sensing domains. This simulation will provide a rapid increase in tracking load as the targets enter the sensing domains. The general approach employed is to simulate relatively few incoming targets and gradually increase the number of outgoing targets in each experiment. Different simulations, with varying incoming/outgoing threat ratios (5/10, 5/21, 5/39, 5/75, 5/95, and 5/195) have been run. This setup provides suitable configurations for evaluating mission performance as the number of objects is increased.

5.0   RESULTS

Several differing methods were identified in order to evaluate the holonic control. Comparison of the holonic architecture in operation with, and without, threat assessment feedback further illustrates the utility of such closed-loop control. In addition, this comparison demonstrates the capability of the holonic architecture to

utilize high-level situation analysis for resource management purposes. The performance of the HS is assessed using different metrics, the application of which is discussed below.

5.1   Resource usage (Platform Load)

Resource usage is measured in terms of the fraction of the available sensing resources available at any given time. In the results reported here, platform load is averaged over several platforms to provide a composite measure. Figure 6 shows the platform load, averaged over the lifetime of the simulation, as a function of the number of outgoing targets. The feedback configuration (in red) maintains a lower average load and its load increases less quickly than in the non-feedback configuration (in blue).

5.2   Tracking performance

Track holons are created with a desired track quality attribute $Q_{desired}$. Their performance Q_{track} is computed according to the fractional difference in the desired and actual quality. For comparison purposes a sum, across all tracks (N) maintained by the group, is used.

$$TP(t) = \sum_{i=1}^{N} TP(t,i) = \sum_{i=1}^{N}\left(1 + \min\left[\frac{Q_{track} - Q_{desired}}{Q_{desired}}, 0\right]\right)$$

From Figure 7, it can be observed that the non-feedback configuration (in red) performs slightly worse than the feedback configuration (in blue). The difference is most pronounced during the highest loading periods in the first half of the simulation. During this period, the sensors in the non-feedback configuration cannot adequately service all of the track holons, resulting in some decay in performance.   The feedback configuration, on the other hand, is better able to service all of its tracks.

5.3   Searching performance

On all platforms an ongoing search task is established in order to detect new targets. The search holons will secure any sensor time not used up by tracking tasks. By subtracting the instantaneous platform load from the total sensing capacity of the platform, an instantaneous measure of search performance $RC(t,p)$ is obtained. The latter is therefore inversely proportional to platform loading (see Section 5.1). The relative increase in search capacity in the feedback case results in targets being detected sooner than in the non-feedback case.

5.4   Global surveillance performance

A global surveillance performance, that combines the capacity to conduct target search $RC(t,p)$ by a group of platforms with the tracking performance $TP(t)$, is defined as follows:

$$S(t) = \frac{1}{N}\sum_{p=1}^{N} S(t,p) = \frac{1}{N}\sum_{p=1}^{N}\left(RC(t,p)\bullet\sum_{i=1}^{M}\left[\left(1 + \min\left[\frac{Q_{track}(i) - Q_{desired}(i)}{Q_{desired}(i)}, 0\right]\bullet Z_b(t,p,i)\right)\right]\right)$$

Here $Z_b(t,p,i)$ is the assessed threat level for track i aboard platform p at time t. The sum is conducted over the M tracks maintained by platform p at time t. An average over N platforms defines the instantaneous surveillance metric.

The increase in search capacity in the feedback configuration, leads to significantly better global surveillance score in the sensing domain of the platforms, as is illustrated in Figure 8. It is clear that the feedback configuration provides better surveillance, especially during periods of high platform load.

5.5   Actual Threat Tracking Performance

Since each simulation includes one target designated as "the actual threat", it is useful to examine the behaviour of the system with respect to this target. Figure 9 shows the track quality for "the threat" in both feedback and non-feedback configurations. Note, that while both configurations maintain a target track, the "threat" is detected much sooner in the feedback configuration. In has also been noticed that the penetration of "the threat" into the area of responsibility before detection increases as the number of spurious targets does.

6.0   CONCLUSION

The paper reviewed the role of SM in the context of military operations. SM in a military setting is hierarchical and recursive in nature and can be decomposed into a series of smaller control problems at each level in the hierarchy. Several control architectures were considered; and the holonic architecture appeared to be the most suitable. The conceptual design for SM using a holonic control approach was presented. Three levels of SM are considered: sensor, platform and group. A force protection scenario was presented that details the situation to which the SM design will be applied. The threat assessment of individual target tracks is used as the basis upon which the SM system allocates sensor resources. Threatening targets are assigned more sensor time, and are therefore tracked more closely, than non-threatening targets. This threat-based feedback approach helps conserve resource usage and focus resources on the most important targets.

7.0   REFERENCES

[1]   Benaskeur, A.R. "Sensor management in command and control", In: Proceedings of the 7th International Command and Control Research and Technology Symposium, September 2002.

[2]   McFarlane, D.C. and Bussmann, S. (2000) "Developments in holonic production planning and control", Production Planning and Control, 11(6), 522-536.

[3]   Koestler, A. (1967): The Ghost in the Machine. Arkana.

[4]   Bussmann, S. (1998): An agent-oriented architecture for holonic manufacturing control. First Open Workshop on Intelligent Manufacturing Systems Europe, 1-12.
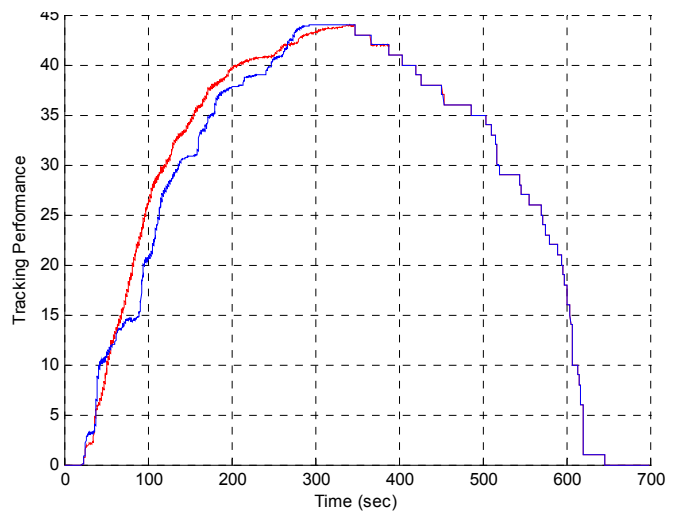
Figure 6: Platform Loading

Figure 7: Tracking Performance
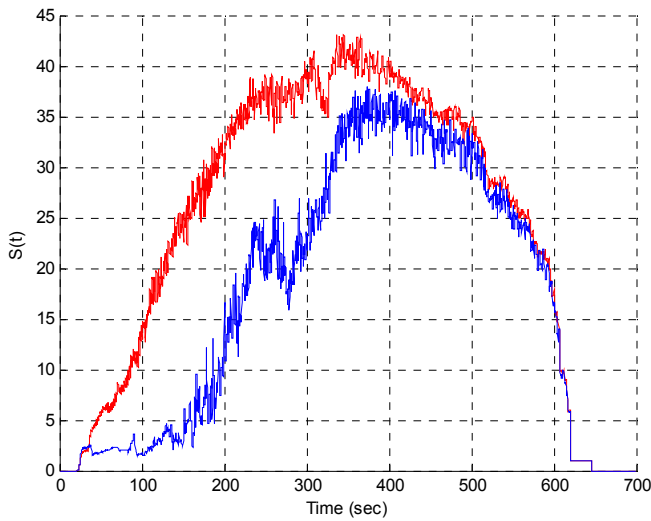
Figure 8: Scene Global Surveillance
Performance

Figure 9: Threat Detection & Tracking
Performance