

Coordinating Plans for Agents Performing AAW Hardkill and Softkill for Frigates

D. Blodgett*, S. Paquet, P. Plamondon**, B. Chaib-draa** and P. Kropf*****

* Lockheed Martin Canada, Building S-82, 2649 Gottingen St., Halifax, NS, Canada, B3K 3C0 Email: dale.blodgett@lmco.com

** Computer Science Dept., Laval University, Ste-Foy, PQ, Canada, G1K 7P4 Email: chaib@ift.ulaval.ca

*** Computer Science Dept., University of Montréal, Montréal, PQ, Canada, H3C 3J7 Email: kropf@iro.umontreal.ca

Abstract

The coordination of anti-air warfare (AAW) hardkill (HK) and softkill (SK) weapon systems is an important aspect of command and control for the HALIFAX Class Frigate. This led to the development of a rapid prototyping environment, described here, which supports the investigation of methods to coordinate the plans produced by AAW HK and SK agents. The HK and SK planning agents are described. An overview of agent coordination methods is provided, with a focus on our initial approach to HK and SK coordination via a Central Coordinator. This approach was successfully implemented, and proved effective in mitigating interference between HK and SK actions, and improved the overall survivability of the Frigate. Finally, future directions of this research are presented.

Keywords

Agents, coordination, planning, resource management.

Introduction

The Combat System of the Canadian HALIFAX Class Frigate is composed of weapon systems, sensor systems, navigation systems, information systems, support systems, and the command and control system (CCS). The CCS lies at the heart of the Combat System. It constitutes an integrated system involving a combination of people, procedures, hardware, and software, which is used to enhance the ability of the personnel in performing Command and Control (C2). Technological advances in threat technology, the increasing speed and diversity of open-ocean and littoral scenarios, and the volume and imperfect nature of data to be processed under time-critical conditions pose significant challenges for current and future shipboard C2 systems as well as the operators who must use them. For several years now, Lockheed Martin Canada (LM Canada) has been working on Multi-Sensor Data Fusion (MSDF), Situation and Threat Assessment (STA), and Resource Management (RM) to support C2 tasks for the HALIFAX Class Frigate.

To advance this work, LM Canada, Laval University and the Defense Research Establishment Valcartier (DREV) have been collaborating since March 2000 to develop a framework for the design and implementation of RM decision aids, based on intelligent agent technology and techniques for multi-agent coordination. The collaboration makes use of a rapid prototyping environment to explore and develop tools and techniques for intelligent agents and multi-agent coordination, as well as approaches for resource planning and scheduling algorithms. This paper will focus on research done in the rapid prototyping environment on coordinating plans of anti-air warfare (AAW) RM agents for “hardkill” (HK) and “softkill” (SK) on a naval platform like the HALIFAX Class Frigate.

AAW HK and SK Systems for Frigates

The AAW HK weapons are weapons that are directed to intercept the target and actively destroy it through direct impact or explosive detonation in the proximity of the target. The range of different types of HK weapons varies, and the effectiveness of the weapon depends on a variety of factors like range, type of target, target speed, environment, etc. The AAW HK weapons for the HALIFAX Class Frigate include surface-to air missiles (SAMs) that have the greatest range, an intermediate range gun, and a Close-In Weapons System (CIWS) that is a short-range, rapid-fire gun. Closely allied to these weapons are two Separate Tracking and Illuminating Radars (STIRs) that are used to guide a SAM to a target, and to point the intermediate range gun. This effectively provides two concurrent fire channels for the AAW HK weapons. The CIWS has its own pointing radar.

The AAW SK weapons use techniques to deceive or disorient a target to cause the target to destroy itself, or at least lose its fix on its intended victim. Again, the range and effectiveness of these weapons varies considerably. The AAW SK weapons for the HALIFAX Class Frigate include chaff and jamming systems. The chaff system launches a shell that produces a burst at a designated position. The resultant chaff cloud has a significant radar cross section that can be used to screen the ship or produce an alternate target on which a radar-guided threat can fix. The jamming system uses electromagnetic emissions to

confuse the threat's sensors to cause the threat to either lose its fix on its intended target, or to improperly assess the position of its target.

Due to their different mechanisms, the HK and SK weapons have historically led independent existences in terms of design and operational deployment. Generally, the HK and SK weapons are supervised by separate control personnel. Thus, the complex task of *optimally combining* the two weapon types falls squarely on the shoulders of the person responsible for overall air defense. The inherent differences between HK and SK weapons, and the nature of their deployment history on the HALIFAX Class Frigate, leads naturally to a representation of HK and SK as two software agents which each determine an anytime plan for their resources and which coordinate plans between them.

The exact nature of the specifications and capabilities of the various AAW HK and SK weapons on the HALIFAX Class Frigate is obviously very complex, and much of that information is classified. In order to maintain emphasis on the coordination techniques and not be burdened by the complexity and fidelity of the representation of HK and SK, a considerably simplified model of the relevant AAW HK and SK weapons was used. This model is a simple, non-classified version of AAW HK and SK for the HALIFAX Class Frigate, but does preserve the fundamental features of these weapons. The details of the model for HK can be found in (Blodgett et al. 1998). The model for SK is described in more detail in the section of SK Systems Planning later in this paper.

Our main emphasis concerns the interaction between weapon systems and controllers on board a single ship, and how to coordinate them in an efficient way. Interactions between such systems on board a ship have always occurred. Sometimes these are *negative* and take the form of undesired interference. The number of potential interactions to be checked, when there are n systems on board, is $n(n-1)$. This number is exhaustive and becomes quite large with increasing n . Fortunately, only a small proportion of the combinations really exist and need to be further studied. Consequently, we should coordinate the HK system with the SK system, because the actions of one system can have a negative impact on the other system. For example, if the SK system launches chaff, it will make the HK system unable to fire a missile in that direction. So, if we do not coordinate these systems they will negatively interact, thus rendering the Frigate more vulnerable.

In our case, each Frigate should manage 3 HK resources (SAM, intermediate range gun, and CIWS) and 2 SK resources (Chaff and Jammer). Each of these resources has a probability of success, a weapon range and a list of interactions with other resources. Such interactions are used for detecting conflicts between resources during the coordination process. In addition to these resources, we also have two STIRs that are not directly managed by our agents. Information issued from the STIRs are however crucial for agents, particularly for guiding missiles.

Coordination Between Agents: A Brief Overview

Today, most would agree that coordination is a central issue in the domain of intelligent agents. Without coordination, agents can waste their efforts and squander resources or fail to accomplish objectives that require collective effort (Durfee 2000). Generally, coordination can be characterized as being the act of managing interdependencies between agents' activities (Lizotte and Chaib-draa 1997, Malone and Crowston 1994, Martial 1991).

Most early work on coordination between agents was guided by cooperation strategies, i.e. strategies that allow agents to improve their collective performance (Cammarata, MacArthur, and Steeb 1983, Durfee 1988, Lesser 1991). Thus, early work on distributed planning took the approach of *complete* planning before action. To produce a coherent plan, the agents must be able to recognize positive and negative interactions and either avoid them or else resolve them. For instance, (Georgeff 1983) included a synchronizer agent to recognize and resolve such interactions. Other agents send their plans to this synchronizer, which examines the plans for critical regions in which, for instance, contention for resources could cause them to fail. The synchronizer then inserts synchronization primitives to ensure mutual exclusion. (Kabanza 1995) has taken the same road by proposing a formal approach based on Metric Temporal Logic (MTL).

Early work on air traffic control also studied coordination strategies for resolving conflicts among plans for aircraft routes where the aircraft are considered as agents (Cammarata, MacArthur, and Steeb 1983). In this work, aircraft are in conflict if they are very close according to their current flight paths. To solve this kind of conflict, agents choose the most-informed agent for elaborating a new flight path and the least-constrained agent to execute this new flight path. The authors carried out experimental evaluations to compare new plans issued from this choice.

Another important approach, which specifically addresses the sub-problem of interdependencies, is the "Functionally Accurate Model (FA/C)" (Lesser 1991). In this model, agents do not need to have all the necessary information available locally in order to solve their sub-problems, but instead interact through their partial results. Starting with the FA/C model, a series of sophisticated distributed control schemes for agent coordination were developed, such as the use of static meta-level information specified by an organizational structure and the use of dynamic meta-level information developed in Partial Global Planning (PGP) (Decker and Lesser 1995, Durfee 1988).

(Tambe 1996, Tambe 1997) has also contributed to the coordination of a teamwork. His model is called STEAM (Shell for TEAM work) and is based on enhancements to

the Soar architecture. The basic building block of a teamwork in STEAM is *Joint Intentions* as suggested by (Cohen and Levesque 1991). More precisely, a teamwork in STEAM is based on agents building up a (partial) hierarchy of joint intentions. Based on the teamwork operationalized in STEAM, three teams have been implemented, two that operate in a commercially available simulation for military training and a third in ROBOCUP.

Military applications have also been investigated by researchers from SRI (Stanford Research Institute) under the multi-agent planning aspect. In (Wilkins and Desimone 1994), the SRI authors describe a prototype system for quickly developing a joint military actions course of action. The system, SOCAP (System for Operations Crisis Action Planning), combines Artificial Intelligence planning, SIPE-2 (System or Interactive Planning and Execution), and a color map display and applies this technology to military operations planning. This approach is extended in (Wilkins and Myers 1995) by a language (called ACT) for representing the knowledge required to support both the generation of complex plans and reactive execution of those plans in dynamic environments. In 1998, the SRI authors introduced both theoretical and practical issues relevant to reasoning about locations (Myers and Wilkins 1998). These techniques were developed during application of the location theories to several large-scale planning tasks within the SIPE-2 planning framework applied to military operations.

HK and SK Planning

As stated before, we have two agents, one for the HK system and the other for the SK system. When they face one or several threats, these two specific agents plan the use of weapon resources of the Frigate for countering the threat(s). Planning weapon resources in this context means allocating and scheduling the deployment of the Frigate's weapon resources against threats with a precise order on the intervention time. The HK and SK planning agents were implemented using the simplified model of HK and SK for the HALIFAX Class Frigate, as discussed above.

HK Systems Planning

This type of planning is accomplished by the HK agent. To do that, this agent should manage three types of resources: SAMs, an intermediate range gun, and CIWS. It has two sorts of algorithms for managing these resources: reactive planning and deliberative planning.

Reactive Planning. Generally, reactive planning uses very low-level reasoning techniques for a simple response to a situation to give a very short reaction time. This is very important in our context because defending Frigates brings a very hard and usually very short time constraint.

To construct a reactive plan, the HK agent maintains a list of threats coming on the Frigate. This list is sorted according to some threat evaluation (i. e., the list is sorted

from the most to the least dangerous threat). Then, it applies some predefined rules for allocating the resources. These predefined rules are: (1) allocating a SAM and a gun to the most dangerous threat; (2) allocating a SAM to the second most dangerous threat; (3) allocating the CIWS to all threats that enter into the CIWS's range.

Though these rules are simple, they allow using all available resources in an efficient way. Unfortunately, the available resources are only allocated to the two most dangerous threats, and all others in the list (if any) are not considered in the reactive plan. In the case where a kill assessment indicates that a hostile threat has been destroyed, the resources that have been allocated to this threat become available for the next most dangerous threat in the list.

Deliberative Planning. Deliberative planning uses complex, high-level reasoning techniques, often over an extended time horizon. Consequently, these plans take more time to construct than the reactive plans. In exchange, they offer more flexibility than the reactive plans since they allow taking into consideration a great number of threats.

In deliberative planning, a decision tree is first produced that explicitly considers, in a probabilistic manner, all possible outcomes of a particular action. Such a tree reflects in fact a plan with different conditional branches. The conditional branches allow us to take into account results of actions. For instance, during the plan execution, one should follow one branch or another depending on the result of an engagement to some threat x . If this engagement has succeeded, then one continues the plan by following a branch where one does not consider the threat x anymore. If the engagement has failed, then one pursues a branch where other engagements are planned for x . All these conditional branches reflect in fact contingent plans and are very important in the sense that engagements to threats are uncertain. Notice that without conditional branches, the time horizon of the plan would be very limited, and we would need to re-plan each time that an engagement fails. The latter can take a long time, thus causing problems for the subsequent threat engagements.

The initial tree is then improved by a tabu search (Blodgett et al. 1998) through the removal or addition of defense actions, followed by update operations aimed at maintaining the consistency of the plan. In recent years, tabu search has been applied with a high degree of success to a variety of problems. It is based on an iterative neighborhood search method where modifications to the current solution that degrade the solution value are admissible. The latter moves allow the method to escape from bad local optima (as opposed to a pure local search approach). To avoid cycling, a short-term memory, known as the tabu list, stores previously visited solutions or components of previously visited solutions. It is then forbidden or tabu to come back to these solutions for a certain number of iterations. Our tabu search may be summarized as follows:

- a) Generate an initial solution s using the construction heuristic;
- b) $s^* = s$;
- c) While stopping criteria of tabu search is not met do:
 1. Generate a neighborhood of s through non tabu moves (or tabu moves that lead to solutions that improve s^*) and select the best solution s' ;
 2. If s' is better than s^* then $s^* = s'$;
 3. $s = s'$;
- d) Output s^* .

SK Systems Planning

This type of planning is accomplished by the SK agent. This agent manages two types of resources: jammers and chaff. In our case, we have two jammers and four chaff launchers. Jammers can act on two threats each. Starting from these considerations, the SK agent elaborates a reactive plan. To do that, it starts from the list of threats attacking the ship (sorted by order of importance, from the most to the least dangerous) and then applies a simple rule which consists of allocating a jammer and a chaff in order to the four most dangerous threats.

During an attack, jammers and chaff must act concurrently and in a complementary way. First, the jammer is used to break the missile threat's radar lock on ownship. Once the missile has lost its target, the jammer creates a false target position on the missile's radar. Then chaff is deployed at a position consistent with the false one provided by the jammer. In this way, the missile's radar locks onto the chaff cloud as its new target.

Planning Coordination Between HK and SK

Methods of Coordination

There are many ways to coordinate the two agents HK and SK. For instance, we can use a Central Coordinator which, after receiving the two plans, one from each agent, will merge them. If there are some negative interactions between the planned actions, it will modify the plans to eliminate those negative interactions, or if not possible, it will try to reduce their effects.

We can also use a direct method where agents communicate with each other and try to coordinate their actions. In this case, communications can be used for commitments and convention as suggested by (Jennings 1994) and they can be used for synchronizing plans and conflict solving.

A third method might be a kind of whiteboard (a common data space) in which the two agents HK and SK will construct a coordinated plan by some successive

refinements. In this case, the coordination will be implicit because they will work on the same plan.

Similar to the whiteboard is the mediator, which in fact plays the role of a Central Coordinator with the possibility of communication and negotiation with SK and HK agents on synchronizing plans and conflict resolution.

The method that uses communications for commitments and conventions, the whiteboard method and mediator approach all seem to be time consuming, and consequently they can probably decrease the ultimate success of the plan for our time critical application. That is why we have opted to initially investigate a Central Coordinator which does not use communication between agents, and for which the coordination process is only based on some simple rules.

Coordination by a Central Agent

As specified in Figure 1, faced with threats, a Situation Assessment Agent proposes a tactical situation (e.g., threat list, kinematic data, etc.) to SK and HK agents as well as to the Central Coordinator. Starting from this tactical situation, HK and SK agents elaborate their plans and send them to the Central Coordinator. Then, the coordinator tries to come up with a coordinating global plan that it proposes to the human operator (the commander or someone else).

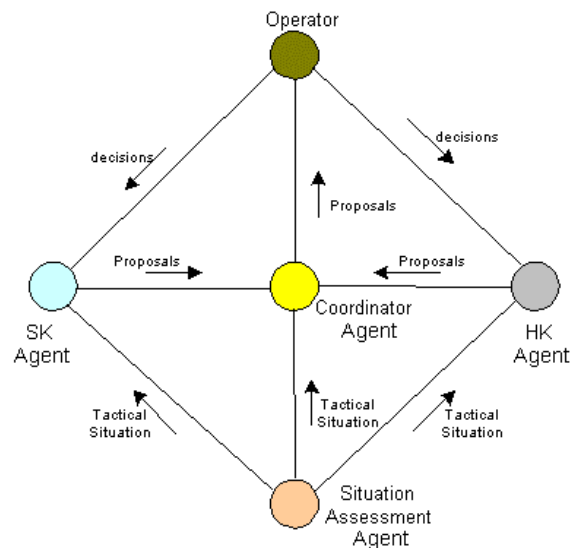


Figure 1. Coordination between HK/SK agents

We assume here that given a situation, SK and HK agents can always come up with plans. When the coordinator has received the two plans, it tries to merge them in order to obtain a coherent global plan. To achieve that, it tries to juxtapose the different portions of both plans taking into account the time of engagement. After that, the coordinator verifies if there are conflicts in the global plan. Notice that most of these conflicts are due to the chaff deployment since they can decrease the range of STIR (radars) drastically. Consequently, we assume we cannot use SAMs, the intermediate range gun and CIWS against

threats blocked by chaff. We also assume that (i) a jammer has very few impacts on HK weapons, (ii) HK weapons have no impact on SK weapons. In this case, the coordinator should only focus on the management of interactions induced by chaff. According to these considerations the coordinator acts as follows:

- a) It checks if there are some HK resources waiting to be launched in the same direction of a chaff already deployed.
 1. If this is the case, the agent verifies if it can delay the deployment of the HK weapon until the chaff effect has been completely dissipated
 2. If not, the agent retracts the deployment of the HK weapon from the global plan.
- b) It checks if there are some chaffs to be launched and if these chaffs are in conflict with HK weapons
 1. If this is the case, the agent gives priority to HK weapons (by removing chaffs from its global plan or by delaying them). The reason to do that is that HK weapons have higher probabilities of success than SK weapons.
 2. If not, the agent tries to see what is the best way of merging chaff and HK according to the situation.

We plan in the future to add to this coordination mechanism a second complementary option, which consists of slightly moving the chaff deployment so that it will not be in the same angle as the HK weapon(s) that we want to launch next. We can do this because the chaff deployment can be moved slightly without losing its efficiency.

As we see, the Central Coordinator uses rules that are relatively simple for the coordination between the two agents. This of course leads to a coordination process that can be very fast.

Preliminary Tests

Test Environment

Even though complete results about the performance of this approach are not yet available, it is worth presenting here an overview of its potential as a promising approach for the Frigate's C2.

The rapid prototyping environment has been implemented with an agent development tool called Jack Intelligent Agents (<http://www.agent-software.com.au/>). Precisely, Jack is an Agent Oriented development environment built on top of and fully integrated with the Java programming language. It includes all components of

the Java development environment as well as offering specific extensions to implement agent-oriented concepts: Agents, Capabilities, Events, Plans, Agent Knowledge Bases (Databases). The agents used in JACK are intelligent agents. They model reasoning behavior according to the theoretical Belief Desire Intention (BDI) model of artificial intelligence. Following the BDI model, JACK intelligent agents are autonomous software components that have explicit goals to achieve or events to handle (desires). To describe how they should go about achieving these desires, these agents are programmed with a set of plans. Each plan describes how to achieve a goal under varying circumstances. Set to work, the agent pursues its given goals (desires), adopting the appropriate plans (intentions) according to its current set of data (beliefs) about the state of the world. This combination of desires and beliefs initiating context-sensitive intended behavior is part of what characterizes a BDI agent.

The basic version of the system as described here comprises two Frigates, for which the number of threats varies between 1 and 20. The threats are all missiles of an identical type. The time at which these threats appear and their initial coordinates are generated randomly. In addition to that, we have assumed that all generated threats move in a straight-line path in the direction of one of the Frigates.

Currently, there is no coordination between the two Frigates. This type of coordination is in fact a very complex process and is left for future work. In this case, each Frigate *only* responds to threats that move in its direction, and it tries as a complex agent to coordinate its HK and SK weapons as proposed previously.

Preliminary Results on Planning Coordination

With regards to the distinction between reactive planning and deliberative planning in the case of HK, the preliminary results show that the deliberative plans are generally more effective than the reactive plans (as indicated in Figure 2). These results also show that the effectiveness of the deliberative plans degrades more quickly than the reactive plans when the number of threats increases. These results are in line with our intuition, because it is normal that at some point the time of deliberation becomes too high and consequently, the agents do not have enough time to build good deliberative plans.

The results on coordination show that chaff may conflict with the HK weapons. This becomes more prevalent as the number of threats or the number of chaff clouds increases. This can be explained by the fact that the chaff clouds are relatively large and they remain in the air a long time before being dissipated, and may block threat(s) from the STIRs used to direct HK weapons. Furthermore, unlike the actual HALIFAX Class Frigate, our simulation and SK agent does not yet consider navigational maneuvers of the ship to reduce these conflicts. Consequently, in our model, and due to the fact that the HK weapons are privileged, the chaff is often not deployed.




Figure 2. Reactive versus deliberative HK planning

Figure 3 shows the relative merits of HK and SK in missile defense, when deliberative plans are used. As we can see, the softkill weapons are less effective than the hardkill weapons. Another important result is that the combination of the hardkill and softkill weapons is more effective than using only one kind of weapon. This last result demonstrates the importance of using both types of weapons and making sure that two are well coordinated.




Figure 3. Relative merits of HK and SK in missile defense (with deliberative HK plans).

We also note the effectiveness of a Central Coordinator agent for the coordination of plans for our type of environment. When it applies a set of simple rules, the coordinator makes it possible to merge plans and to

manage conflicts very quickly, while keeping the number of interactions as low as possible.

Future Plans

Enhancements

The rapid prototyping environment developed in this study has already proven useful for the implementation of HK and SK planning agents, and for the investigation of the Central Coordinator method of coordinating HK and SK planning.

Enhancements to this work are already underway or are planned for the future. These include: (1) improving the realism and increasing the complexity of the HK and SK models (e.g., more realistic chaff clouds and chaff deployment), (2) increasing the complexity of the threat scenarios (e.g., more realistic threat trajectories), (3) adding more elements to the coordination process (e.g., navigational maneuvers and changes to ship orientation to reduce conflict between HK and SK), (4) extending the local HK and SK coordination to a more complex coordination between multiple platforms, and (5) investigating other coordination methods (i.e., involving communication strategies).

Technology Demonstration Test Bed

Initial investigations and proof-of-concept evaluations of the agents and coordination techniques were performed using a rapid prototyping environment. However, the agents and coordination techniques will eventually be transitioned to a more powerful Technology Demonstration Test Bed, which incorporates a real-time knowledge based system infrastructure. Here, development and implementation issues can be more fully and realistically explored. A schematic of the Technology Demonstration Test Bed is shown in Figure 4.

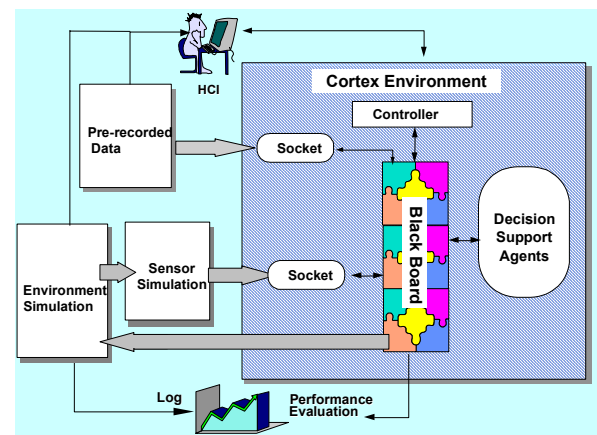


Figure 4. Technology Demonstration Test Bed

The Test Bed was designed to investigate naval (and airborne) C2 information system applications. Cortex, a software tool developed by LM Canada, is a real-time knowledge-based system using a blackboard architecture and communications layer (Lockheed Martin Canada 2001) that is used to build decision support agents. The Test Bed can process pre-recorded data. However, it also provides simulation of the environment (including entities like ownship and threats) and realistic simulation of sensors' (e.g., for the HALIFAX Class Frigate) perceptions of the environment. It also permits the agents to act upon and affect the environment simulation. The Test Bed accommodates operator interaction with the agents, provides monitoring and control of the Test Bed, and permits performance evaluation of Test Bed applications.

Agents for HK and SK developed in the rapid prototyping environment would be ported to Cortex, where they could benefit from potentially great increases in speed and performance. Using Cortex may also make it possible to explore methods of HK and SK coordination that were not viable in the rapid prototyping environment. The Test Bed provides more flexible and realistic simulation capabilities for the HALIFAX Class Frigate than exist in the current rapid prototyping environment. Combined with the opportunities for performance evaluation, the Technology Demonstration Test Bed makes possible a more robust, comprehensive and realistic representation and evaluation of HK and SK coordination methods.

Conclusions

As a consequence of this work, a useful rapid prototyping environment has been developed. The infrastructure incorporates agent tools and a closed-loop simulation of environment, scenarios and resources for the target application (HK and SK for the HALIFAX Class Frigate). Rudimentary HK and SK agents for a naval platform like the HALIFAX Class Frigate were implemented. Due to the hard and very short time constraints for the availability of HK and SK plans, the most suitable initial approach for coordination between HK and SK was via a central coordination agent. This method was successfully implemented, and proved effective in mitigating interference between HK and SK actions and improved the overall survivability of the platform. The simplified nature of the weapon specifications and capabilities, as well as the simple threat scenarios, means that care should be taken not to interpret the results of this study as precise for the real world. Nonetheless, the general results certainly validate intuition in many ways, and are valuable as a guide to what is and isn't a viable strategy for the more complex real-world situation with the HALIFAX Class Frigate. The rapid prototyping environment developed here will be enhanced to provide a more realistic test environment, and will be used to investigate other methods for coordination between HK and SK.

Acknowledgments

The authors would like to thank Bob Mazerolle, Arthur Flores and Rick Pantano of Lockheed Martin Canada for their guidance concerning HK and SK weapon systems for the HALIFAX Class Frigate. We are also grateful for the financial support for this work from the Natural Sciences and Engineering Research Council of Canada (project 222802-98) and the Defence Research Establishment Valcartier.

References

- Blodgett, D. E., Gendreau, M., Guertin, F., Potvin, J-Y, and Séguin R. A tabu Search heuristic for resource management in naval warfare. Publication of Centre de Recherche du Transport, CRT-98-63, Montréal, 1998.
- Cammarata, S., MacArthur, D. and Steeb, R. Strategies of cooperation in distributed problem solving. In *Proc. of the 8th Int. Joint Conf on AI (IJCAI-83)*, Karlsruhe, Germany, 767-770, 1983.
- Cohen, P.R. and Levesque, H.J. Teamwork. *Nous*, 25(4), 487-512, 1991.
- Decker, K.S. and Lesser, V.R. Designing a family of coordination algorithms. in *Proc. of the First International Conf. on Multi-Agent Systems (ICMAS-95)*, AAAI Press: Menlo Park, CA, pp. 73--80, 1995.
- Durfee, E.H. *Coordination of distributed problem solvers*. Kluwer Academic Pub.: Boston, MA, 1988.
- Durfee, E.H. Scaling up agent coordination strategies, *IEEE Computer*, 34(7), pp. 39-46, 2000.
- Georgeff, M.P. Communication and interaction in multi-agent planning. In *Proc. Of the 3rd National Conf. on AI (AAAI-83)*, Washington, D.C., 125-129, 1983.
- Jennings N.R. *Cooperation in Industrial Multi-Agent Systems*. World Scientific Series in Computer Science, vol. 43, 1994.
- Kabanza, F. Synchronizing Multiagent Plans using Temporal Logic Specifications, *Proc. of the First International Conf. on Multi-Agent Systems*, AAAI Press: Menlo Park, CA, 217-224, 1995.
- Lesser, V.R. A retrospective view of FA/C distributed problem solving. *IEEE Trans, on Systems, Man and Cybernetics, Special Issue on Distributed AI*, 21(6), 1347-1362, 1991.
- Lesser, V.R. Reflections on the nature of multi-agent coordination and its implications for an agent architecture. In *Autonomous Agents and Multi-Agent Systems*, Kluwer Acad. Pub., vol 1(1), 89-111, 1998.

Lizotte, S. and Chaib-draa, B. Coordination in CE systems: An approach based on the management of dependencies between agents. *CERA: Concurrent Engineering: Research and Applications*, 5(4), 367--377, 1997

Lockheed Martin Canada. *Cortex: A Knowledge-Based System and Run-Time Engine*, Version 2.4. Lockheed Martin Canada: Montreal, Canada, February 7, 2001.

Malone, T.W. and Crowston, K. The interdisciplinary study of coordination. In *ACM Computing Survey*, vol 26(1), 87-120, 1994.

Martial F.V., Coordinating Plans of Autonomous Agents, *Lectures Notes in AI*, Springer-Verlag series, Berlin, 1991.

Myers, K.L. and Wilkins, D.E. Reasoning about locations in theory and practice. *Computational Intelligence*, 14(2), 1998.

Tambe, M, Teamwork in real-world, dynamic environments. In *Proceedings of the International Conference on Multi-agent systems (ICMAS96)*. AAAI Press, Menlo Park, 1996.

Tambe, M. Towards flexible teamwork. *Journal of Artificial Intelligence Research*, vol. 7, 83-124, 1997. AAAI Press: Menlo Park, CA, 1997.

Wilkins, D.E. and Desimone, R.V. Applying an AI planner to military operations planning. In M. Fox and M. Zweben (eds), *Intelligent Scheduling*, MK Pub., San Mateo, 685-709, 1994.

Wilkins, D.E. and Myers, K.L. A common knowledge representation for plan generation and reactive execution. *Journal of Logic and Computation*, vol. 5, 731-761, 1995.