

Interchanging Agents and Humans in Military Simulation

Clinton Heinze¹, Simon Goss¹, Torgny Josefsson¹, Kerry Bennett¹,
Sam Waugh¹, Ian Lloyd¹, Graeme Murray¹ & John Oldfield²

¹ Defence, Science, and Technology Organisation
506 Lorimer Street, Fishermen's Bend,
Victoria, Australia
firstname.lastname@dsto.defence.gov.au

² Royal Australian Air Force
Evaluation Analysis & Standards Flight, 92 Wing,
RAAF Base Edinburgh, South Australia, Australia
john.oldfield@defence.gov.au

Abstract

The innovative reapplication of a multi-agent system for human-in-the-loop (HIL) simulation was a consequence of appropriate agent oriented design. The use of intelligent agents for simulating human decision making offers the potential for analysis and design methodologies that do not distinguish between *agent* and *human* until implementation. With this as a driver in the design process the construction of systems in which humans and agents can be interchanged is simplified. Two systems have been constructed and deployed to provide defence analysts with the tools required to advise and assist the Australian Defence Force in the conduct of maritime surveillance and patrol. The systems both simulate maritime air operations. One utilises intelligent agents to provide models of tactical decision making, the other provides the same environment but provides a set of user interfaces to allow air force flight crews to participate in human in the loop simulation by replacing the intelligent agents. The experiences gained from this process indicate that it is simpler, both in design and implementation, to add humans to a system designed for intelligent agents than it is to add intelligent agents to a system designed for humans.

Introduction

The modification and development of existing constructive¹, multi-agent military systems by the Australian Defence Science and Technology Organisation (DSTO) to provide the Royal Australian Air Force (RAAF) with human-in-the-loop (HIL) capability is reported in this paper. It extends applications developed and described earlier [1,2], and commences the process of integrating intelligent agent developments [8] with HIL systems research [3]. The applications described here differ in purpose from most other deployed HIL simulations in that they are used for exploration, evaluation and development of tactics and procedures rather than for training [15].

The innovative use of an existing multi-agent system for HIL simulation was a consequence of appropriate agent oriented design. The use of intelligent agents for simulating human decision making offers the potential for analysis and design methodologies that do not distinguish between agent

and human until implementation [14]. With this as a driver in the design process the construction of systems in which humans and agents can be interchanged is simplified. Two systems using the same base architecture are used to support operations research:

- The original system that utilises intelligent agents for modelling all of the military personnel within a scenario. This is conceptually identical with the systems reported by Tidhar et. al. [1] although applied to different aircraft and missions.
- The new system that removes the intelligent agent for a particular aircraft of interest and provides user interfaces that allow the actual crew of that aircraft to fly simulated missions for the purpose of validating and developing tactics.²

The use of intelligent agents to military simulation is maturing. For several years intelligent agents have been applied to constructive military simulation. Architectures, methodologies and programming patterns in support of this development are improving.

The incorporation of intelligent agents into HIL simulation is generally post-hoc engineering of large legacy systems or the injection of entities into a large distributed simulation via an interface [10]. Agents have requirements on systems that are not apparent in mainstream HIL simulations. Difficulties associated with the successful incorporation of intelligent agents into extant systems are often associated with a failure to recognise the specific requirements that agents will place on the system [9,16]. These problems can be alleviated by careful design of new systems or by costly remediation of existing systems.

This paper provides a case study of a deployed HIL simulation used for development of tactical procedures for a maritime surveillance aircraft. An existing constructive simulation that used intelligent agents to model all human components was modified. The modifications provide user interfaces that allow air force personnel to replace the agents that previously modelled them. *The experiences gained from this process indicate that it is simpler, both in*

¹ The simulation community uses the term *constructive* to designate those systems that contain no human interaction.

² There is a phase of significant duration in an acquisition or mid-life refit programme where military units are in a work-up mode in anticipation of delivery of platforms with enhanced capability that makes them effectively 'first of kind' and requires development of new or significantly altered tactics, operating procedures, and doctrine.

design and implementation, to add humans to a system designed for intelligent agents than it is to add intelligent agents to a system designed for humans.

The following section details the domain of application of this technology, that of maritime patrol and surveillance by the RAAF. Operational analysis that incorporates both constructive and HIL simulation can offer significant savings to the Air Force. Savings are realised both in mission performance and in support costs with respect to fuel used and time to complete a mission. By far the biggest savings are realised in extending the life of type of an aircraft through smarter operation.

Maritime Patrol and Surveillance Tactics

Air Operations Division (AOD) of the DSTO supports the RAAF's Maritime Patrol Group (MPG) in developing new tactics and 'concepts of operation' for the upgraded AP-3C Orion Maritime Patrol Aircraft (see Fig. 1). The Orions are used by the RAAF in peacetime for maritime search, surveillance and operations in and around Australian territorial waters.

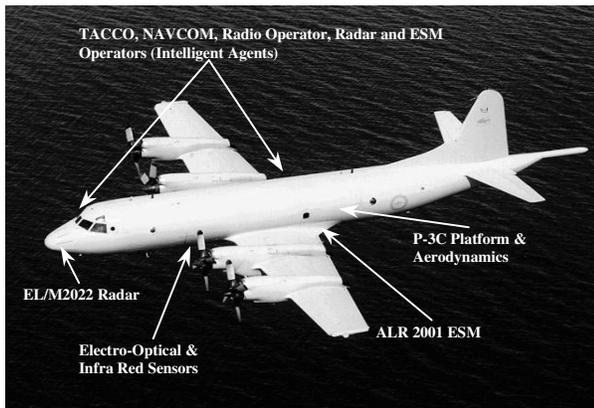


Figure 1. The various sensors and human operators that are modelled in this work superimposed on a photograph of an AP-3C Orion maritime patrol aircraft.

The Orions are in an extensive upgrade program that includes new sensors and avionics that significantly improve the capability of the aircraft. Because MPG have no previous operational experience with some of these new sensors, AOD's operational analysts work closely with them to baseline the expected mission performance of the aircraft in typical mission profiles and scenarios, and to develop new, integrated flying and sensor employment policies that allow the aircraft to function at its full mission potential.

The requirement from the RAAF was for AOD to investigate the effectiveness of flying tactics and sensor employment policies as a function of weather conditions, geography and other factors that impact on mission effectiveness. To meet these requirements aspects of the domain are modelled:

- A detailed physical model of the aircraft, it's sensors, including flying characteristics, fuel consumption and sensor performance.

- The tactical decision making processes on board the aircraft representing the human operators, and crew workload (including the type and amount of information available at any given time), the sensor data-fusion process and chain of command.
- The environment, including weather and sea state.
- Several types of ships.

An example mission with some of the factors that need to be considered in tactical decision making is outlined below:

Typical AP-3C Scenario

An AP-3C Maritime Patrol Aircraft is tasked to find and monitor the movements of a target whose location is not precisely known. In this situation the aircraft will fly a predefined search pattern over the region of ocean that the target is suspected of being located within. The aircraft will use its various sensors such as the radar, the ESM (Electronic Support Measures: for detecting the radar transmissions of other ships and aircraft) and infra-red optical sensors to try and locate and classify all possible ships in the region in order to find the target. The radar operator and the ESM operator perform their duties trying to detect and classify various 'contacts' on their sensors. Typically these two operators have hundreds of 'contacts' on their sensors at any given time. The protocol for using the radar and ESM sensors (they have many different modes) depends on a number of factors such as weather conditions, altitude, target type, the presence or otherwise of other ships and aircraft and whether the aircraft wants to advertise its position to the target it is looking for. Contacts that cannot be eliminated at this stage are passed up to the Sensor Employment Manager (SEM) who performs data-fusion duties, such as associating data from the two sensors if they are deemed to be from the same source, and who directs the use of different sensor classification techniques. The SEM passes on information about contacts that may possibly be the target to the Tactical Coordinator (TACCO). The TACCO decides which contacts need further investigation and in which order, either by flying closer or changing aspect for different sensor classification techniques. The TACCO must balance many competing factors, including minimising the amount of unnecessary flying distance in order to complete the mission as soon as possible (in effect solving 'travelling-salesman' type problems) and not concentrating on identifying one suspicious contact at the expense of others.

The TACCO and SEM are always on the alert for 'suspicious' behaviour that may single out one unknown contact as the target. In effect 'suspicious' behaviour means a reaction or response from the contact that is consistent with a particular goal specific only to the target (such as remaining covert). For instance, an ESM contact identified as a powerful but unknown navigation radar on a ship may be lost just prior to the same contact being picked up on the radar sensor and classified as a relatively small ship. This may be due to a target (with a low radar signature) switching off its own surveillance radar upon detecting the radar transmission of the AP-3C aircraft.

A system capable of simulating scenarios of this type was required. Decisions about the nature of the constructive simulation were guided by previous experience with fighter combat, strike missions, and airborne early warning and control.

Several components were candidates for reuse from

previous developments. Additionally the aircrew of MPG expressed interest in having the ability to interact with a simulation as it ran, and using the simulation as a tool for exploring tactics. This was in part driven by exposure to the tools used by analysts in validation sessions for air mission models. This need arose through the lack of a simulation facility within the MPG to account for the upgraded system. Opportunities arose through insights gained from the development of tactical development environments for other projects currently being undertaken at AOD. Thus there was an expressed requirement for two systems:

- a constructive operations research focussed system that could be used to process many thousands of runs of many scenarios to carefully evaluate tactics and procedures; and
- an interactive HIL system capable of providing the flight crews the opportunity to review, replay, interact with and participate in a small number of scenarios.

Simulation Framework

The components common to both systems are detailed here. The details of the intelligent agents and the HIL user interfaces discussed in the section following this.

BattleModel

In light of previous experience with constructive simulation, the 'BattleModel' simulation framework was chosen as the primary modelling environment for this work. BattleModel is a simulation framework developed by the DSTO. The details of the framework and its suitability for agent based modelling of military decision-making have been documented [13]. BattleModel was designed several years ago to conduct constructive simulations in support of operations research. A strong design requirement was that BattleModel support the integration of intelligent agents. (The support of HIL simulation was not considered a priority.)

AP-3C Orion

The aircraft itself is modelled to the extent that it has the same manoeuvring characteristics of the AP-3C including the same fuel consumption, climb and descent rates and cruise performance as a function of weight and altitude. The sensor suites modelled in this scenario include a high fidelity radar model originally built by the Surveillance Systems Division of DSTO [12]. This model includes all the radar tracking and radar imaging modes available on the AP-3C Elta radar. Further systems modelled on the aircraft include the ESM electronics, visual model and Electro-Optical sensor.

Weather and the Environment

The weather and environmental conditions in the area significantly influence mission effectiveness. The presence of certain cloud formation and thunderstorm activity

severely constrains where the aircraft can fly and affects sensor performance – particularly visual detection and classification range. Strong winds and rainfall have some effect on the radar tracking capability but it mainly affects visual classification ranges and the performance of the various classification modes on the radar. The sea state and the size and direction of the swell also affects the capability of the radar to identify contacts. Sufficiently detailed models of the weather and the sea and land environment provide the inputs into the sensor performance models and into the tactical decision making of the agents.

Ships

The types of military and commercial ships found in the Australian region are modelled. These provide the surface radar contacts for which the Orion will search. The ships are fitted with suitable radio transmitters and radars providing the ability to model the ability of the Orion to detect the ships by radar, or to detect their electronic emissions by ESM.

Intelligent Agents and System Design

The first system to be implemented was the constructive simulation. Models of the aircraft, ships and related subsystems, the weather and environment were engineered. Intelligent agents were used to model all of the human players in the system. From a design and implementation perspective this was similar to previously developed intelligent agent systems [1]. The acquisition of the knowledge required to construct these systems involves familiarisation with AP-3C tactical procedures documentation, debriefing of flight crew and the regular participation of DSTO staff in operational missions on board the aircraft.

The second system was the interactive or HIL variant. This reused most of the components of the first systems but replaced the agents used to model the AP-3C crew and replaced them with user interfaces and a HIL capability.

System Design

The design of BattleModel and all subsequent AOD simulations have built upon experiences with agent oriented systems. This experience has led to a view of system development that does not distinguish between the human acting in the world or the agent acting in the simulation. In a software engineering sense this tends to merge the business domain model, the use cases and the system architectural design [14]. This approach was taken because explicit knowledge representation with agents closely matches the knowledge acquired from RAAF personnel. Constructing agents that, at a knowledge level, closely resemble the humans that they model reduces the software engineering effort by closing the gap between knowledge engineer and software engineer [7].

Agent Oriented Design

The development described in this paper incorporated the modelling of three related, but different, application systems. Figure 2 illustrates the modelling similarity when each activity is represented with the symbology of the Unified Modeling Language (UML) and use case symbology. The lowest row gives a standard application of UML as we used in the development of Human-in-the-Loop simulations. The top row corresponds to domain or business modelling of real-world defence operations for the documentation of concepts of defence operations. The process of applying and testing techniques and methodologies from the knowledge acquisition community (such as cognitive systems engineering and cognitive work analysis) to software engineering, is ongoing. The middle row is an extension to the UML to include the modelling of intelligent agents as actors within the design scope of the system (see also Heinze, Papisimeon and Goss [14]).

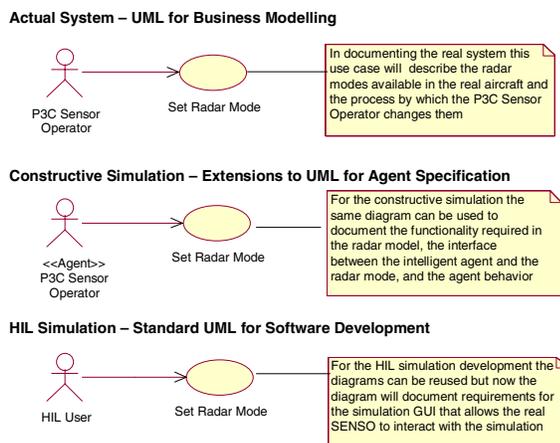


Figure 2. An example use case diagram showing the actual system, the agent based simulation and the HIL simulation. This diagram emphasises the reusability of system documentation and design that is available with attention to modelling choices.

The fundamental similarity between the actual system operated by an actual crew, a simulated system operated by simulated crew, or a simulated system operated by an actual crew enabled selection of modelling and documentation techniques that could be reused across this spectrum. Use of comparable tools and languages, together with the ability to re-use modelling and documentation over a number of development projects, reduces development time and allows for efficient focusing of resources.

A fundamental premise of this article is that it is easier to add humans to a system designed for intelligent agents (ie, to an environment designed for constructive simulations), than it is to add intelligent agents to a system designed for humans (ie, to an environment designed for HIL simulation). This premise has been formulated following experiences of attempting to inject intelligent agents (computer generated forces [CGFs]) into simulations that incorporate a human flying an aircraft (stationary cockpit platform with realistic consoles, controls and communications) within a virtual world. Though there are

many engineering challenges in immersing intelligent agents into such HIL simulators, an indicative problem area is the construction of things and spaces that can be understood and acted upon by agents. Much HIL simulation development is focused on constructing an information environment of relevance to the human perceptual and action system. Projected visualisations (rendered scenes provided by graphics engines) are quite realistic with near-real-time changes of the environment through which the pilot navigates, and with identifiable symbology of ‘things’ in this environment (eg, other planes). Communications are representative of human voices and human languages. In summary, the back-end graphical and voice databases of these simulators do not provide a structured representation for ready access by reasoning agents. There are thus significant issues associated with ensuring that intelligent agents that must interact with terrain and objects in the scenario also have the ability to take in relevant information and act in an appropriate manner accordingly. One manner of achieving this is to program agents to perceive and act in a human-like manner. Another is to construct a parallel system that gives an agent-compatible environment. In either case, considerable development work is required because the application programming interface (API) for humans cannot be used by intelligent agents.

In contrast, relatively lower costs, resources and development time are required to allow a human immersion within a system that has been primarily designed for agent play. At the implementation level, humans and intelligent agents connect to the rest of the simulation through exactly the same API. This allows a “plug-and-play” philosophy with respect to interchanging agents and humans. The visual (a graphical user interface) and other databases of relevance to human perceptual and action requirements are readily constructed for ‘plug-in’ together with the human. Such a system also shows potential for some members of a AP-3C crew to be real, and some to be intelligent agents. Work is currently being directed to investigate this concept.

Intelligent Agents

The tactical decision making component was modelled using individual intelligent agents (implemented in the dMARS language [5]) for each crew-member that has a significant role in the tactical decision making process. intelligent agents were chosen because of the requirement to model decision-making based on a degree of awareness of the environment or tactical situation the aircraft finds itself in. Maritime surveillance tactics, as with almost all tactical decision making, rely on making an assessment of the current situation based on fusing data from different sensors and also on making assessments of the intent of other entities.

The construction of the models of the AP-3C crew was in and of itself a considerable challenge. The task of modelling military decision-making in a manner that allows for the explicit representation of tactical options to support evaluation and development is made more difficult because of the nature of the specific domain. The AP-3C intelligent

agent development differed from previous systems (see Tidhar et. al. [1] for a description of these) because it involved the modelling of an aircraft with several interacting crew. Issues of teamwork, communication, and cooperation are more significant.

The AP-3C crew agents are based upon an existing agent design that characterises decision making as a four stage process of situation awareness, situation assessment, tactics selection, and tactics implementation. The basic design was extended to provide a flexible way of modelling different tactical options so that different tactics could be simply tried and evaluated. The dMARS -Agent formalism is particularly suited to modelling this type of situation awareness based behaviour. Additionally, the plan language which expresses procedural knowledge affords knowledge acquisition and validation advantages [6] (See Fig. 3).

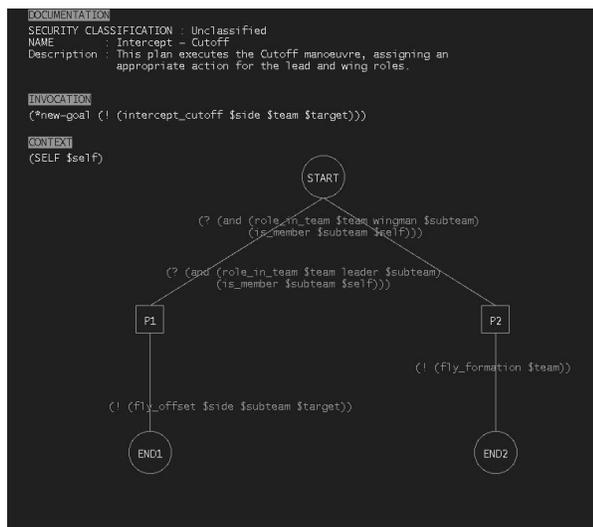


Figure 3. A dMARS plan. Plans are graphical representations of actions to perform and goals to adopt. With attention to design, plans can be understood by subject matter experts. This bridges the gap between the domain expert and the software engineer [11].

In terms of the tactical decision making process, six of the crew members on the aircraft are modelled (using Intelligent Agents technology) to the extent that the type of information, the amount of information and the communications and chain of command on the aircraft are accurately reproduced [4,17].

The advantages of using Intelligent Agents for this are two-fold. Firstly, the roles and area of responsibility of each individual crew member can be incorporated into each agent individually which facilitates modification of tactics and monitoring of workloads³. Secondly, the execution of the tactical model can be displayed graphically and understood by a non computer programmer. This allows actual aircrew to validate the existing tactical model and also to determine not just what decisions were made but also why they were

³ This approach enables crewing options to be explored by simulation in AOD work for other platforms.

made.

The agents receive information from the simulation, reason about it and make tactical decisions that alter the aircraft states – or send message to other agents. The agents receive only that subset of data that would be available to the real crew member that they model and efforts are to use knowledge representations within the agent that match to those used by air-crew.

HIL Interface

For the interactive variant the agents that modelled the crew of the AP-3C were removed and interfaces that allowed for HIL participation were added. Figure 4 shows the high level plan view that is presented to the AP-3C crew. Added to this are interfaces that allow them to control the AP-3C by changing its course, altitude, or speed, and to control the radar.

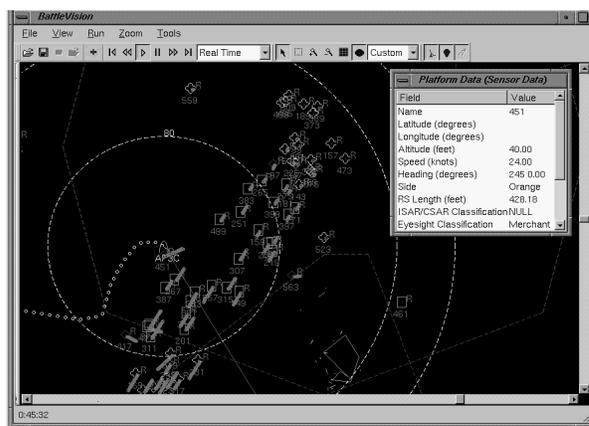


Figure 4. A screen capture of the main plan view of the battle-space. The squares surround ships that are potential Targets of Interest for the Orion. This view does not reflect the radar or situation displays that might be found on a real aircraft but provides the necessary level of information for evaluating tactical options.

The commands available to the crew using the interactive AP-3C are identical to the set available to the agents. This allows for reuse of all of the other components of the system. The interactive simulation is not intended to replicate the AP-3C’s on-board displays as it is designed to provide the crew with a display for considering tactics and not practicing procedures⁴.

Studies and “CONOPS” Development

In this work the ‘BattleModel’ is used both for faster than real-time constructive simulations of missions, with the Intelligent-Agents making the tactical decisions, and also interactively or “crew-in-loop”, where tactical decisions are made by actual crew members. The constructive mode is

⁴ A suitable lack of fidelity is in fact highly desirable to keep the simulation and exploration of tactical options at the function level and separate from the specific interfaces of the deployed aircraft.

used to gather information on several hundred simulated missions for statistical analysis and robustness tests of various tactics⁵.

Figure 5 illustrates the tactical development cycle used by the RAAF and where BattleModel fits into the process. Initial tactics, mission profile and measures of effectiveness are supplied to AOD for initial modelling and analysis at the top of the cycle. Constructive simulations of hundreds of missions are then performed using BattleModel for a given tactic or environmental condition. The outcomes of these missions are then analysed, both statistically and individually to determine the drivers of mission effectiveness and identify worst-case-scenarios. These results are then presented to the RAAF for evaluation and analysis. The initial tactics are then refined further using the HIL mode based on the modelling results and subsequently input into another iteration or round of modelling. Once the tactics are deemed suitably mature they become operational.

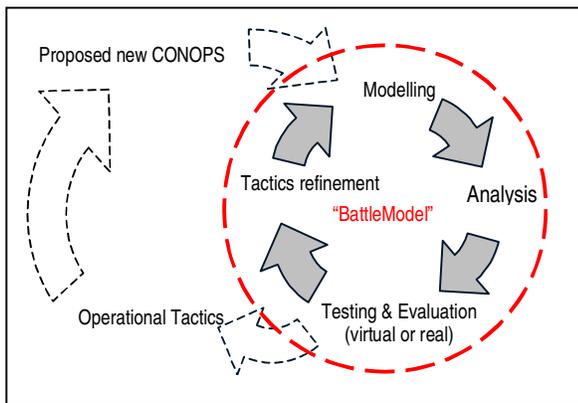


Figure 5. The tactical development cycle used by the RAAF illustrating where BattleModel fits into the process.

In the context of Figure 4 the AP-3C BattleModel is used to provide data for the modelling phase and the interactive simulation is used to test and evaluate the tactics and to refine them for further modelling.

Application Deployment

The deployment of this system within the tactical development program outlined above, currently involves three DSTO analysts and several RAAF officers on a full-time basis. Fifteen DSTO scientists, commercial software engineering contractors and several RAAF 92WG AP-3C aircrew are also involved to various degrees on a casual basis. The initial development and deployment of this system required approximately eleven staff-years of DSTO scientist and analyst effort over a three year period and approximately AU\$380K for software engineering support and associated hardware. Approximately three staff-years of effort was devoted to the design, development and testing of the intelligent agent module.

⁵ Faster than real time studies are essential: varying 8 parameters across three variables within typical scenarios requires three years to complete if constrained to real time performance.

With the system currently deployed at DSTO Melbourne, regular three-day workshops are held on a half-yearly basis with the schedule to be increased considerably with a further deployment at RAAF Edinburgh in 2001. The HIL applications of the system are required to execute in real-time and hence require several hours to complete a single mission. Although complete missions are not always required this typically limits the system to three or four missions per day.

The agent based applications are not limited to execute at real-time and hence are only hardware limited. These applications execute at approximately 5 X real-time on SGI Origin 200, 4 processor, 360MHz R12K machines with 2GB memory. Approximately 40% of this CPU time is directly attributable to the agent's process. Typical support studies for the tactical development program requires statistical evaluation of approximately a dozen agent controlled scenarios with approximately 100 missions simulated in each scenario. Total CPU time (per Origin 200) is of the order of 500 hrs for a given study.

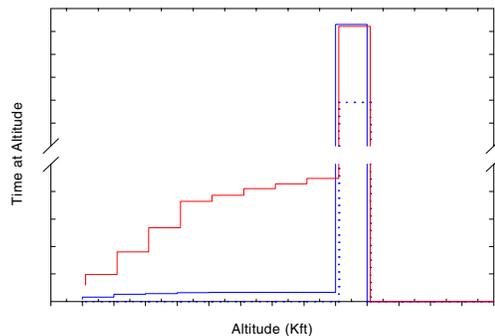
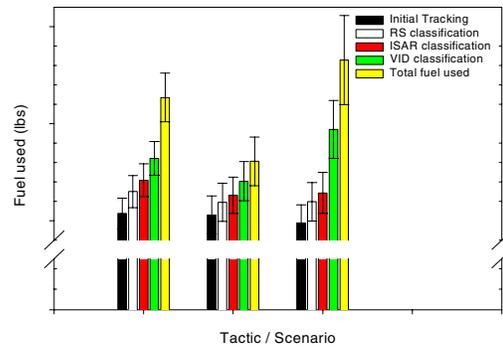


Figure 6. Typical results from "AP-3C BattleModel", the constructive simulation. These types of graphs detail time of flight, fuel used, altitudes maintained and other measures of performance that are used to compare tactics and to determine standard operating procedures.

The operations research and analysis of tactics is carried out by a process of specifying tactics, measures of effectiveness, scenarios, and conducting extensive simulation. Typically existing tactics are base-lined and suggested improvements or variations evaluated relative to the baseline. In this way, measures such as time of flight, fuel used, or time at a specific altitude can be used to evaluate tactics.

By parameterising the tactics search of a range of tactics is possible. Typical scenarios run much faster than real time allowing many more instances to be looked at than is possible with a HIL facility. Results are obtained, analyses conducted and the findings reported to the operational squadrons or the MPG (see Fig. 6). The system deployment results from the tactical development program are national security classified and cannot be discussed further.

Unusual, suspect, promising, or otherwise interesting combinations of tactics and scenarios can be examined in detail in the HIL facility. Furthermore, the HIL facility can be used to reduce the search space by characterising and constraining the types of tactics that need to be considered by exploratory investigation by experienced AP-3C crews. Thus the human in the facility performs a valuable *pre-analysis* role in defining the types of scenarios that might be of interest and the range of tactics that might be explored and then *post-analysis* in validating the usefulness of tactics and the behaviour of the agents.

The interactive or HIL mode is used to test and evaluate new tactics in a realistic environment and to refine existing tactics based on statistical analysis of constructive simulation results. In this mode, the tactical picture is projected onto a large screen showing the current sensor information (radar, ESM contacts, aircraft location and state etc) superimposed on to a geographical map of the region. This allows the crew to focus on developing and evaluating higher-level tactical procedures (see Figure 7) rather than on low-level interactions with individual controls.



Figure 7. Six RAAF AP-3C crew members from Maritime Patrol Group forty six minutes into a simulated mission using the interactive AP-3C simulation with its HIL mode to evaluate some maritime search and classification tactics.

The simulations are housed within a facility at DSTO (see Figure 7). Within this facility AP-3C crews can simulate missions and explore tactical options. The interactive simulation can record, replay, and re-run allowing specific missions to be studied, reviewed, and alternate tactics explored and evaluated. Typically two crews will alternate missions lasting many hours over a period of several days. During this time proposed tactics can be reviewed, checked against a variety of scenarios, and otherwise evaluated.

Maintenance and Future Development

The current implementations of the constructive and the Interactive AP-3C simulations are maintained and run by DSTO directly with strong infrastructure and financial support from the RAAF due to the priority nature of this work. AP-3C crews spend time at AOD on a regular basis – refining their concepts and following the results of operational studies. The interactive nature of the process has tightened organisational links between AOD and the MPG and has fostered strong cooperation.

A proposal to transition the technology from the defence scientist to the operational flight crew is currently being considered. If accepted this would see the technology transferred from DSTO into the squadron where it could be used regularly by operational crews for tactical development. It is important to distinguish this system from training simulators that exercise procedures and skill based reactions of operators. These simulations systems are for the development of tactical procedures and the trialing of concepts and hence do not have the expensive development and maintenance costs associated with the high-end graphics of training simulators.

Future developments in agent languages are expected to feed the AP-3C project. These include technological improvements in agent languages; methodological improvements in the software engineering, development and maintenance of these systems, and knowledge engineering aspects that provide techniques for closing the gap between the conceptualisations of the domain held by the domain experts and the explicit representations within the computational system. Currently DSTO is undertaking research and development in support of the existing agent developments such as that described here and the systems that are expected to enter service throughout the next decade. It has been possible to rapidly develop and commission new operational systems that are at the leading edge of agent development but where the risk is mitigated by maintaining a strong coupling between research and development and the operational systems.

The next two advances being pursued by the development group lie in the areas of modelling teams and command and control and the mainstreaming of the technology. The former should allow for easier modelling of socially complex scenarios whilst the latter will allow agent technologies to be deployed directly into operational air force units.

Concluding Remarks and Lessons Learned

Defence organisations are primarily concerned with developing and maintaining the capability to conduct successful and efficient military operations. These operations rely heavily on the performance of the available defence systems and the expertise of the humans operating them. Modelling and simulation of these operations is a priority for defence organisations in evaluating existing and proposed defence systems. Modelling the human operators

is critical to conducting such evaluations.

A combination of constructive and interactive technologies has allowed DSTO to supply advice about the tactical operational performance of the AP-3C Orions to MPG. The HIL system allows the AP-3C crews to gain familiarity with the system and to explore, prototype and workshop tactics that can then be studied in-depth using intelligent agents, as substitutes for the crew, in constructive simulations that cover thousands of scenarios. This method has caused crew to reflect on procedures and come to insights about their own performance not otherwise available to them. This promotes 'double loop' organisational learning [18].

The ability to *plug-and-play* intelligent agents and humans within the same basic system has dramatically improved the ability of DSTO to obtain valuable input from the air force. The system has provided a clear means of validating the behaviour of agents and has value in knowledge acquisition.

Significant savings in dollars and aircraft life can be obtained if tactics can be evaluated and refined with modelling and simulation. By maintaining systems that explicitly model flight crew and their tactical decision making with intelligent agents it has been possible to rapidly develop HIL equivalents. These systems provide valuable advice to the operators of military aircraft and provide mechanisms for validation, exploration, and evaluation of tactical procedures. By including the operational crews in the development of simulation improvements in knowledge acquisition and validation of the intelligent agents have been realised.

Acknowledgments

The authors would like to thank Martin Cross, Tracy Truong, Arvind Chandran from DSTO and Gil Tidhar, Phil Dart, Mario Selvestrel, and Darryl Hunter, from Kesem International who assisted with the development of these systems. Without the strong and active support of the Orion flight crews of No. 92 Wing of Maritime Patrol Group this work would not have been possible. The authors would also like to acknowledge the financial support from MPG the support of David Glenny, the responsible Research Leader.

References

1. G. Tidhar, C. Heinze, S. Goss, G. Murray, D. Appla, and I. Lloyd. Using Intelligent Agents in Military Simulation or "Using Agents Intelligently". In Proceedings of Eleventh Innovative Applications of Artificial Intelligence Conference. Deployed Application Case Study Paper, Orlando, Florida, 1999.
2. G. Tidhar, C. Heinze, and M. Selvestrel. Flying Together: Modelling Air Mission Teams, Applied Intelligence, vol. 8, pp. 195-218, 1998.
3. D. McIlroy, C. Heinze, D. Appla, P. Busetta, G. Tidhar, and A. Rao. Towards Credible Computer Generated Forces. In Proceedings of Second International Simulation Technology and Training Conference, (SimTecT '97), Melbourne, Australia, 1997.
4. AP-3C Tactical development workshop, held at AOD, AMRL Fishermens Bend, Melbourne, May, 2000.
5. d' Inverno M. and Kinny, D. and Luck, M. and Wooldrige, M., 'A Formal Specification of dMARS', In proceedings of the Fourth International Workshop on Theories, Architectures and Languages, 1997.
6. Georgeff, M. P. and Lansky, A. L., 'Procedural Knowledge', Proceedings of the IEEE Special Issue on Knowledge Representation, volume 74, pages 1383-1398, 1986.
7. C. Heinze, B. Smith, and M. Cross. Thinking Quickly: Agents for Modeling Air Warfare. In Proceedings of Australian Joint Conference on Artificial Intelligence, AI '98 Brisbane, Australia, 1998.
8. McIlroy, D. and Heinze, C., 'Air Combat Tactics Implementation in the Smart Whole AiR Mission Model (SWARMM)', In Proceedings of the SimTecT Conference, 1996, Melbourne, Australia.
9. Jones, R. M. and Laird, J. E. Constraints on the design of a high-level model of cognition. In *Proceedings of Nineteenth Annual Conference of Cognitive Science*, 1997
10. Tambe, M. and Jones, R. M. and Laird, J. E. and Rosenbloom, P. S. and Schwamb, K., 1994. Building Believable Agents for Simulation Environments: Extended Abstract. In *Collected Papers of the SOAR/IFOR Project*, Information Sciences Institute, University of Southern California, pages 78-81. Marina del Ray, CA.
11. Rao, A. S., 'A Unified View of Plans as Recipes', Contemporary Action Theory, ed. Holmstrom-Hintikka, G. and Tuomela, R., Kluwer Academic Publishers}, The Netherlands, 1997.
12. Detection Performance Prediction Model for the EL/M2022 Maritime Surveillance Radar System (U)Antipov, B. Reid, J. Baldwinson; DSTO -TR-0870 (CONFIDENTIAL), 1999
13. C. Heinze, D. Appla, and I. Lloyd. "The BattleModel". In Proceedings of Special Interest Group on Artificial Intelligence (SIGAI' 97), Fishermens Bend, Australia, 1997.
14. Heinze, C., M. Papisimeon, and S. Goss. Specifying Agent Behaviour With Use Cases. In Proceedings of Pacific Rim Workshop on Multi-Agents, PRIMA 2000.
15. Tidhar, G., Murray, G., and Steuart S., 'Computer-Generated Forces and Agent-Oriented Technology', In Proceedings of the Australian Joint Conference on Artificial Intelligence Workshop on AI in Defence, 1995, Canberra, Australia.
16. Wooldrige, M. J. and Jennings, N. R., 'Pitfalls of Agent-Oriented Development', In 'Proceedings of the 2nd International Conference on Autonomous Agents (Agents-98)', Minneapolis, USA, 1998.
17. Transcript of 92 Wing P-3C Tactical Walk-Through. T. W. Josefsson and S. Goss (editors), DSTO-IPIC-0027, 1999.
18. Senge, P., *The Fifth Discipline: The Art and Practice of the Learning Organisation*. Random House Publishers, Australia, 1998.