

Interchanging Agents and Humans in Military Simulation

Clinton Heinze, Simon Goss, Torgny Josefsson, Kerry Bennett, Sam Waugh, Ian Lloyd, Graeme Murray, and John Oldfield

■ The innovative reapplication of a multiagent 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 defense analysts with the tools required to advise and assist the Australian Defense Force in the conduct of maritime surveillance and patrol. 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.

When a modern Air Force takes delivery of a major new capability, such as an aircraft or a significant upgrade to a sensor or weapon system, a significant amount of preparatory work will have already been undertaken. To more quickly use the capabilities of the system on introduction into service, it is essential that the aircrews have substantial procedures in place for deploying and operating the aircraft. The task of developing these operating procedures often falls to the squadron that will operate the new aircraft. Existing experience with similar capabilities can assist the transition but often the task of filling gaps in the operational knowledge falls to military analysts and modeling and simulation.

A recent major upgrade to the radar of the *AP-3C Orion* maritime patrol aircraft (figure 1) required a major rethinking of the procedures that govern the operational employment of

the aircraft. The Defense Science and Technology Organization (DSTO) was tasked with assisting the flight crews of the Maritime Patrol Group with examining certain aspects of the tactical operation of the aircraft. It was believed that by exploiting the upgraded capabilities of the radar, significant improvements in operational performance could be achieved together with improvements in fuel and airframe fatigue life use.

The modification and development of existing constructive, multiagent military systems by the DSTO to provide the Royal Australian Air Force (RAAF) with a human-in-the-loop (HIL) capability is reported in this article. It extends applications developed and described earlier (Tidhar et al. 1999; Tidhar, Heinze, and Selvestrel 1998) and begins the process of integrating intelligent agent developments (McIlroy and Heinze 1996) with HIL systems research (McIlroy et al. 1997). The application described here differs in purpose from most other deployed HIL simulations in that it is used for exploration, evaluation, and development of tactics and procedures rather than training (Tidhar, Murray, and Steuart 1995; Tambe, Schwamb, and Rosenbloom 1995).

The innovative use of an existing multiagent 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 (Heinze, Papasimeon, and Goss 2000). By adopting methodologies that do not distinguish between humans and agents, the construction of systems in which humans and agents can be interchanged is simplified. Two systems using the same base




Figure 1. The Various Sensors and Human Operators That Are Modeled in This Work Superimposed on a Photograph of an AP-3C Orion Maritime Patrol Aircraft.

Each of these components (including the crew) must faithfully be represented within the simulation.

architecture are used to support operations research.

First is the original system that uses intelligent agents for modeling all the military personnel within a scenario. This system is conceptually identical with the systems reported by Heinze, Smith, and Cross (1998) and Tidhar et. al. (1999), although applied to different aircraft and missions.

Second is the new system that removes the intelligent agents for a particular aircraft of interest and provides user interfaces that allow the actual crew of this aircraft to fly simulated missions to validate and develop tactics.

The use of intelligent agents in 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 using an interface (Tambe et al. 1994). 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 recognize the specific requirements that agents will place on the system (Jones and Laird 1997; Tambe,

Schwamb, and Rosenbloom 1995). These problems can be alleviated with the careful design of new systems or a costly remediation of existing systems.

This article 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 modeled them. 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.

The following section details the domain for this technology, that of maritime patrol and surveillance by the RAAF. Operational analysis that incorporates both constructive and HIL (or interactive) simulation can offer significant savings to the Air Force. Savings are realized 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 realized in extending the life of a type of aircraft through smarter operation.

The Process

Air Operations Division's (AOD) experience with agent-based simulation in support of Air Force operations research is extensive. With a background in the fighter world with the simulation of air combat tactics (Heinze, Smith, and Cross 1998) and the modeling and simulation in support of Project Wedgetail, the airborne early warning and control acquisition, AOD was well placed to undertake the required analysis.

Much of AOD's previous experience was the development of tactics for aircraft already in service. This experience provided the opportunity for the tactics that emerged from the analysis process to be tested, validated, and refined on the actual aircraft and within the operational environment of the squadrons. All the experience of the Operational Analysis Group was with the development and use of constructive simulation. That is, simulation that has no human interaction. This Monte Carlo simulation is characterized by the exploration of large parameter space over many thousands, tens of thousands, or even hundreds of thousands of runs. By simulating the important aspects of the domain, randomizing appropriate parameters, and collecting

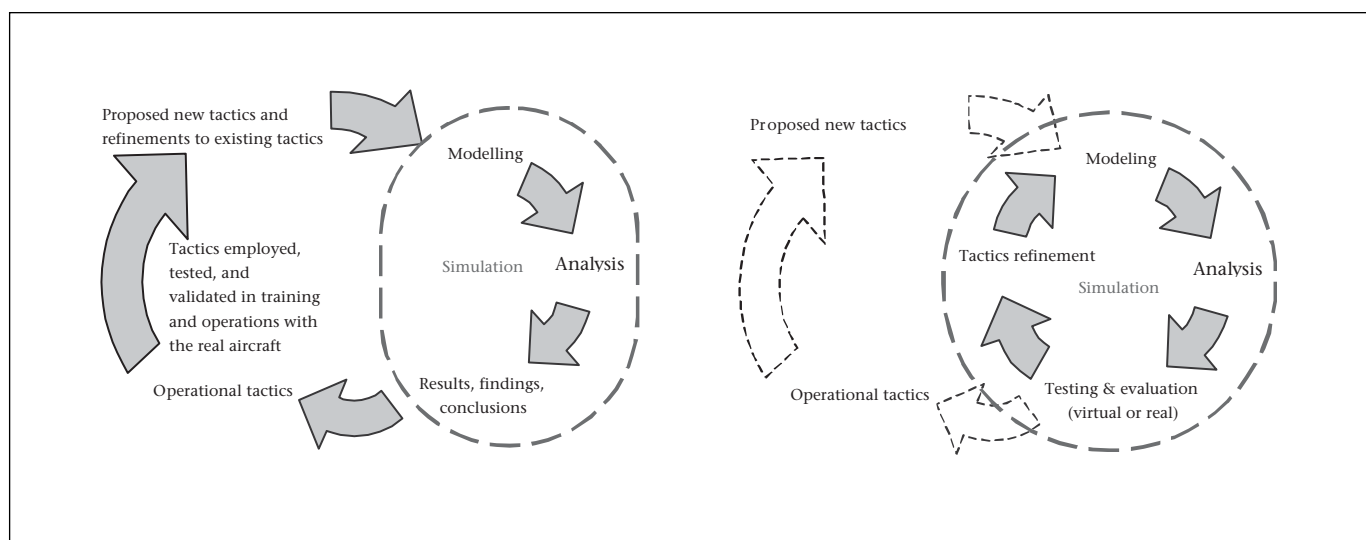


Figure 2. The Business Process Surrounding the Development of Standard Operating Procedures (SOPs).

A. The existing SOP development process. B. The AP-3C SOP development process. To speed up the development of tactics, the loop was tightened through the provision of an HIL simulator.

and analyzing the resulting data, measures of performance and effectiveness of various tactics can be produced and compared. Typically, these simulations, by necessity, run many times faster than real time.

The AP-3C tactics development task posed two new and significant challenges. First, the upgrade was expected to radically alter the manner in which the aircraft was flown—almost to the point of presenting a completely new aircraft. Second, the impending arrival of the upgraded aircraft placed short timelines on the analysis.

To address these challenges required some changes to the normally adopted business processes. It was clear that a closer than normal interaction with the flight crews would be necessary to focus the analyses on particularly important tactics and to test, evaluate, and validate the developed tactics. To gain a first-cut understanding of the performance of the aircraft, it was considered desirable to expose the Air Force flight crews to a simulation of the performance of the upgraded AP-3C. This exposure would provide the crews with the chance to gain some exposure to the new aircraft and to explore, in a structured way, the performance characteristics of the radar, and it would provide AOD analysts with the opportunity to observe, document, and record the tactical decision making of the crews for model development purposes (figure 2). With the business process determined, the challenge was to develop two systems.

First is a constructive simulation, similar to those that had been built in the past by the group but tailored to the physical systems and

tactics and decision making of the Maritime Patrol domain. This simulation would model the tactics using intelligent agents to represent the crew and would require the construction or acquisition of models of all the physical systems of the AP-3C aircraft, the environment in which it operated, and the other players within the scenarios. We wanted a constructive operations research–focused system that could be used to process many thousands of runs of many scenarios to carefully evaluate tactics and procedures.

Second is an interactive HIL system similar to the first but capable of providing actual flight crews the opportunity to review, replay, interact with, and participate in a small number of scenarios. To produce a completely new simulator, particularly one that differed radically from any simulator that had previously been developed within the group was a risk. It was a risk that it turned out could be mitigated by the design methodologies that had grown up around the engineering of intelligent agent systems and had resulted in the development of constructive simulations that were ideally suited to transformation into HIL simulations. Furthermore, the models of the physical systems, the radars, aircraft performance, ships, weather, and so on, could be reused. Thus, there was an expressed requirement for two systems.

Before describing the two simulators, a description of typical maritime patrol and surveillance tactics is provided, setting the scene for the types of modeling required and offering insights into the nature of the domain.

Maritime Patrol and Surveillance Tactics

AOD of the DSTO supports the RAAF's maritime patrol group in developing new tactics and concepts of operation for the upgraded *AP-3C Orion* maritime patrol aircraft (figure 1). The *Orions* are used by the RAAF in peacetime for maritime search, surveillance, and operations in and around Australian territorial waters.

The *Orions* are in an extensive upgrade program that includes new sensors and avionics that significantly improve the capability of the aircraft. Because the Maritime Patrol Group has 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 affect mission effectiveness. To meet these requirements, aspects of the domain are modeled: (1) a detailed physical model of the aircraft, its sensors, including flying characteristics, fuel consumption, and sensor performance; (2) the tactical decision-making processes on board the aircraft representing the human operators, crew workload (including the type and amount of information available at any given time), the sensor data-fusion process, and chain of command; (3) the environment, including weather and sea state; and (4) several types of ships. An example mission with some of the factors that need to be considered in tactical decision making is outlined in figure 3.

Constructive Simulation

The first system to be implemented was the constructive simulation. Models of the aircraft, ships and related subsystems, the weather, and the environment were engineered. Intelligent agents were used to model all the human players in the system. From a design and implementation perspective, the simulation system was similar to previously developed intelligent agent systems (Tidhar et al. 1999). The acquisition of the knowledge required to construct these systems involves familiarization with *AP-3C* tactical procedures documentation, debriefing of flight crew, and the regular participation

of DSTO staff in operational missions on board the 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 the Maritime Patrol Group 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 need arose through the lack of a simulation facility within the Maritime Patrol Group 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.

In this work, the *BATTLEMODEL* is used both for faster than real-time constructive Monte Carlo simulations of missions, with the intelligent agents making the tactical decisions, as well as with the interactive or HIL version, where tactical decisions are made by actual crew members. The constructive mode is used to gather information on several hundred simulated missions for statistical analysis and robustness tests of various tactics.

System Design

The design of *BATTLEMODEL* and all subsequent AOD simulations have built on 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 and the agent acting in the simulation. In a software engineering sense, an approach that abstracts away the differences between agents and actors tends to merge the business domain model, the use cases, and the system architectural design (Heinze, Papisimon, and Goss 2000). This approach was taken because explicit knowledge representation with agents closely matches the knowledge acquired from Air Force 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 (Heinze, Smith, and Cross 1998).

Battle Model

In light of previous experience with constructive simulation, the *BATTLEMODEL* simulation framework was chosen as the primary modeling environment for this work. *BATTLEMODEL* is

An AP-3C maritime patrol aircraft is tasked with finding and monitoring 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 residing in. 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 infrared optical sensors to try and locate and classify all possible ships in the region 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 the desire of the aircraft to advertise its position with regard 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 might 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 minimizing the amount of unnecessary flying distance to complete the mission as soon as possible (in effect solving traveling 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 behavior that might single out one unknown contact as the target. In effect, *suspicious behavior* 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 example, an ESM contact identified as a powerful but unknown navigation radar on a ship might be lost just prior to the same contact being picked up on the radar sensor and classified as a relatively small ship.

Figure 3. Typical AP-3C Scenario.

a simulation framework developed by the DSTO. It is currently used to support a program of work, including fighter combat, strike missions in hostile environments, airborne early warning and control, and the maritime version described here. 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.)

More recently, Air Force personnel have requested interactive tools that provide them with the capacity to simulate Air Force operations for the purpose of developing, testing, and validating concepts of operation and tactical procedures. This request in part has been driven by exposure to the tools used by analysts in validation sessions for air-mission models.

Intelligent Agents

The intelligent agents for the AP-3C simulation, as with all the AOD-developed simulators, model the tactical reasoning of the aircrews. The fine motor controls, the operation of consoles or control systems on the aircraft, are handled by modules external to the agent. The intelligent

agent takes in sensory data from the physical system models; processes it; and responds with high-level, tactical commands that are then implemented by particular modules. Typical examples, might include “fly a racetrack orbit,” “fly to the next waypoint,” “descend to 300 feet,” and “set the radar to mode X.”

This type of model is particularly useful when the task is the development of standard operating procedures.

The tactical decision-making component was modeled using individual intelligent agents (implemented in the dMARS language [d’Inverno et al. 1997]) for each crew member that has a significant role in the tactical decision-making process (figure 4). 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 making assessments of the intent of other entities.

Plans are graphic representations of actions to perform and goals to adopt. With attention to design, plans can be understood by subject matter experts, bridging the gap between the

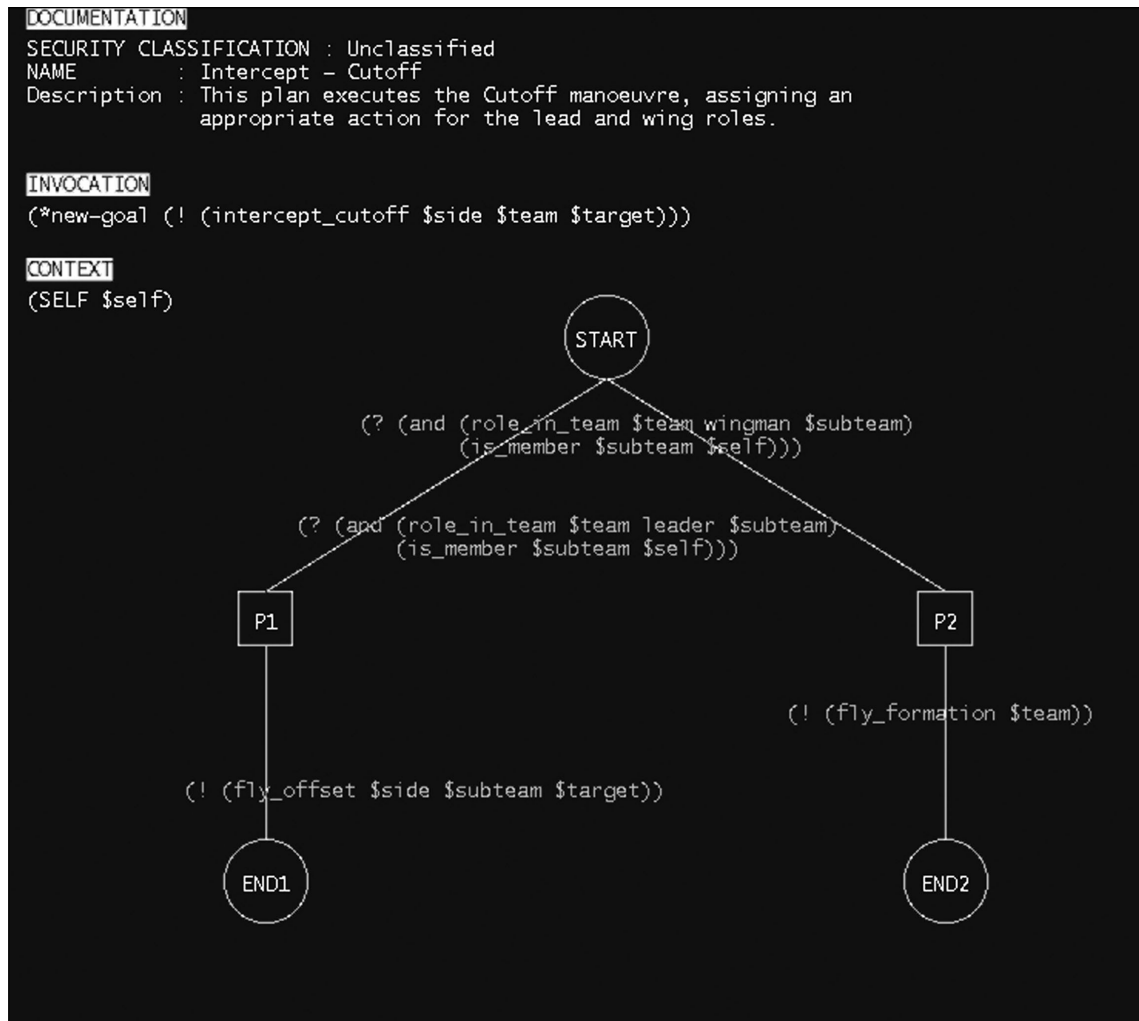


Figure 4. A dMARS Plan.

domain expert and the software engineer (Heinze, Smith, and Cross 1998; Rao 1997).

The dMARS agent formalism is particularly suited to modeling this type of situation awareness-based behavior. Additionally, the plan language in which procedural knowledge is expressed affords knowledge-acquisition and validation advantages (Georgeff and Lansky 1986) (figure 2).

In terms of the tactical decision-making process, six of the crew members on the aircraft are modeled (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 (Wooldridge and Jennings 1998). The six crew members are the pilot, the navigation and communications officer (NAVCOM), the TACCO, the SEM, and the radar and ESM operators.

The advantages of using intelligent agents for

this are twofold. First, 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 work loads. Second, the execution of the tactical model can be displayed graphically and understood by a non-computer programmer, allowing actual aircrew to validate the existing tactical model and determine not just what decisions are made but also why they are made.

The agents receive information from the simulation, reason about it, and make tactical decisions that alter the aircraft states or send a message to other agents. The agents receive only the 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 those used by the aircrew.

AP-3C ORION

The aircraft itself is modeled to the extent that it has the same maneuvering 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 modeled in this scenario include a high-fidelity radar model originally built by the Surveillance Systems Division of DSTO (Antipov, Reid, and Baldwinson 1999). This model includes all the radar-tracking and radar-imaging modes available on the AP-3C Elta radar. Further systems modeled on the aircraft include the ESM electronics, the visual model, and the electrooptical sensor.

Weather and the Environment

The weather and environmental conditions in the area significantly influence mission effectiveness. Certain cloud formations and thunderstorm activity severely constrain where the aircraft can fly and affect 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 affect the capability of the radar to identify contacts. Sufficiently detailed models of the weather and the sea and land environment provide the input into the sensor performance models and the tactical decision making of the agents.

Ships

The types of military and commercial ships found in the Australian region are modeled. These ships 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 detect their electronic emissions by ESM.

Interactive (Crew-in-the-Loop) Simulator

The second system was the interactive or HIL variant. This system 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 an HIL capability.

Leveraged Development

The many models of the physical elements within the simulation were reused in the devel-

opment of the interactive simulator, including the aerodynamics and aircraft performance models; the radars and other sensors; the weather, atmosphere, and the sea models; and the various ships, aircraft, and ground vehicles and their subsystems. Also reused was the simulation kernel itself. The BATTLEMODEL had been fitted for real-time simulation during its early development, and although this feature had never been used, it proved simple to modify.

The BATTLEMODEL, although well suited to simulations with intelligent agents, does not require intelligent agents, and it has been used for many nonagent studies. During the development of the agent-based simulation, the knowledge-elicitation process had resulted in a domain model that was geared for the agent-based modeling.

The knowledge-elicitation process extracted and recorded information about the decision making, the actions, the reasoning, and the behavior of each of the crewmembers, together with their interactions, commands, and communications. The documentation of this knowledge, although it was intended for the development of the intelligent agents and becomes a requirements specification for their functions, can equally be interpreted as a domain model and, because of the nature of the analysis and design process, is also easily translated into a specification for the interfaces of the interactive simulator.

Thus, the domain model that describes the tactical decision making of the actual flight crews, the agent use case model that describes the functional requirement for the intelligent agents, and the standard use case model that describes the requirements for the aircrews operating the interactive simulator are effectively identical.

Furthermore, the resultant system uses precisely the same interfaces for the intelligent agents as for the crew-in-the-loop interfaces. Thus, the systems present a core simulation kernel that allows the seamless interchanging of agents and humans.

HIL Interface

For the interactive variant, the agents that modeled the crew of the AP-3C were removed, and interfaces that allowed for HIL participation were added. Figure 5 shows the high-level plan view that is presented to the AP-3C crew. These interfaces allow the crew to control the AP-3C by changing its course, altitude, and speed and controlling its radar.

The squares surround ships that are potential targets of interest for the *Orion*. This view

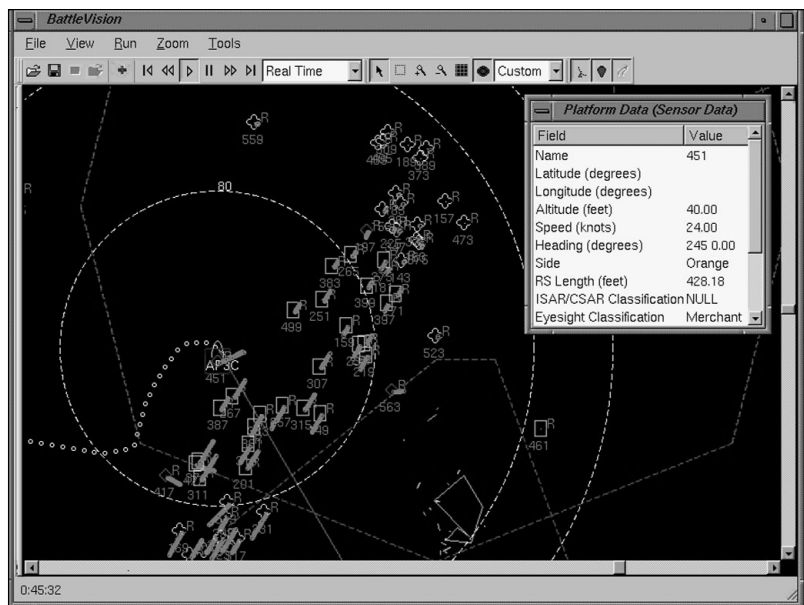


Figure 5. A Screen Capture of the Main Plan View of the Battle Space.

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, allowing for reuse of all the other components of the system. The interactive simulation is not intended to replicate the *AP-3C*'s on-board displays because it is designed to provide the crew with a display for considering tactics, not practicing procedures.

Application Deployment

The operations research and analysis of tactics is carried out by a process of specifying scenarios, tactics, and measures of effectiveness and conducting extensive Monte Carlo simulation to explore the space created by the systematic variation of important parameters and randomness introduced into the simulation. Typically existing tactics are base lined, and suggested improvements or variations are 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.

Scenarios run much faster than real time, allowing many more instances to be looked at than is possible with an HIL facility. Results are obtained, analyses conducted, and the findings reported to the operational squadrons or the Maritime Patrol Group (figure 6).

These types of graphs detail time of flight,

fuel used, altitudes maintained, and other measures of performance that are used to compare tactics and determine standard operating procedures.

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 characterizing 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 preanalysis role in defining the types of scenarios that might be of interest and the range of tactics that might be explored and then a postanalysis in validating the usefulness of tactics and the behavior of the agents.

The interactive, or crew-in-loop, 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, and so on) superimposed on a geographic map of the region. This approach allows the crew to focus on developing and evaluating higher-level tactical procedures (figure 7) rather than on establishing low-level interactions with individual controls.

The simulations are housed within a facility at DSTO (figure 7). Within this facility, *AP-3C* crews can simulate missions and explore tactical options. The interactive simulation can record, replay, and rerun 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 because of 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 organizational links between AOD and the Maritime

Patrol Group and has fostered strong cooperation.

A proposal to transition the technology from the defense scientist to the operational flight crew is currently being considered. If accepted, the technology will be 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 simulation systems are for the development of tactical procedures and the testing and evaluation 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 developments 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 conceptualizations 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 modeling teams and command and control and the mainstreaming of the technology. The former should allow for easier modeling of socially complex scenarios, and the latter will allow agent technologies to be deployed directly into operational Air Force units.

Concluding Remarks and Lessons Learned

Defense organizations 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 defense systems and the expertise of the humans operating them. Modeling and simulation of these operations is a priority for defense organiza-

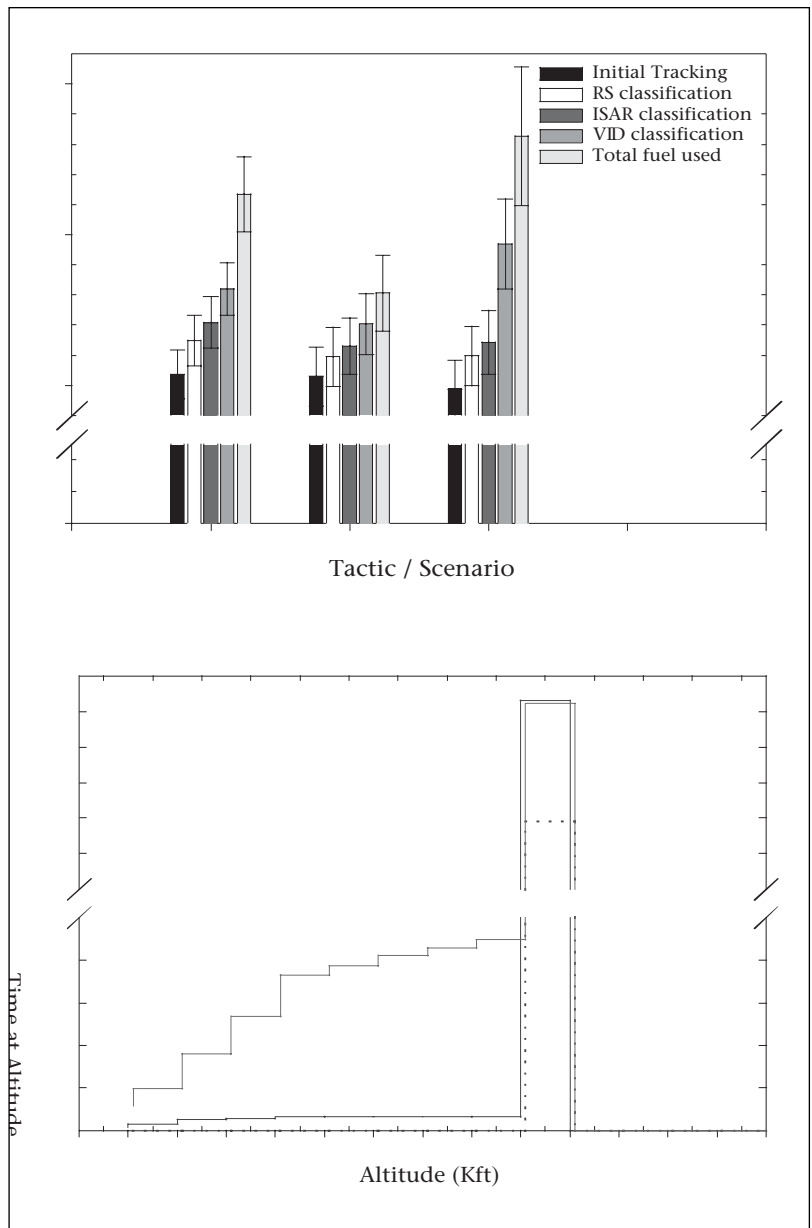


Figure 6. Typical Results from AP-3C BATTLEMODEL, the Constructive Simulation.

tions in evaluating existing and proposed defense systems. Modeling the human operators is critical to conducting such evaluations.

By providing a combination of constructive and interactive technologies, DSTO has been able to supply advice about the tactical operational performance of the AP-3C Orions to the Maritime Patrol Group of the RAAF. 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 Monte Carlo simulations that



Figure 7. Six Royal Australian Air Force AP-3C Crew Members from Maritime Patrol Group Forty-Six Minutes into a Simulated Mission Using the Interactive AP-3C Simulation with its Crew-in-Loop Mode to Evaluate Some Maritime Search and Classification Tactics.

cover thousands of scenarios. This method has caused the crew to reflect on procedures and develop insights about their own performance not otherwise available to them, promoting double-loop organizational learning (Senge 1998).

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 customer. The system has provided a clear means of validating the behavior of agents and has significant value in knowledge acquisition.

Significant savings in dollars and in aircraft life can be obtained if tactics can be evaluated and refined with modeling 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 HLL 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 realized.

Acknowledgments

The authors would like to thank Martin Cross, Tracy Truong, and Arvind Chandran from the DSTO development team 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 to per-

form some aspects of this work and the support of David Glenny as the research leader responsible for this work within AOD.

References

- Antipov, B.; Reid, J.; and Baldwinson, J. 1999. Detection Performance Prediction Model for the EL/M2022 Maritime Surveillance Radar System. Technical Report, DSTO-TR-0870, Defense Science and Technology Organisation, Melbourne, Australia.
- d'Inverno, M.; Kinny, D.; Luck, M.; and Wooldridge, M. 1997. A Formal Specification of dMARS. In *Intelligent Agents IV*, eds. M. P. Singh, A. Rao, and M. J. Wooldridge, 155–176. Lecture Notes in Artificial Intelligence 1365. Berlin: Springer-Verlag.
- Georgeff, M. P., and Lansky, A. L. 1986. Procedural Knowledge. In *Proceedings of the IEEE Special Issue on Knowledge Representation, Volume 74*, 1383–1398. Washington, D.C.: IEEE Computer Society.
- Heinze, C.; Papisimeon, M.; and Goss, S. 2000. Specifying Agent Behaviour with Use Cases. In *Design and Applications of Intelligent Agents, Proceedings of the Third Pacific Rim International Workshop on Multi-Agents, PRIMA 2000*, eds. C. Zhang and V. Soo, 128–142. Lecture Notes in Artificial Intelligence 1881. Berlin: Springer.
- Heinze, C.; Smith, B.; and Cross, M. 1998. Thinking Quickly: Agents for Modeling Air Warfare. Paper presented at the Australian Joint Conference on Artificial Intelligence, AI '98, 6–10 December, Brisbane, Australia.
- Jones, R. M., and Laird, J. E. 1997. Constraints on the Design of a High-Level Model of Cognition. Paper presented at the Nineteenth Annual Conference of Cognitive Science, 7–10 August, Stanford, California.
- McIlroy, D., and Heinze, C. 1996. Air Combat Tactics Implementation in the Smart Whole AiR Mission Model (SWARMM). Paper presented at the First Simulation Training and Technology Conference (SimTecT), 25–26 March, Melbourne, Australia.
- McIlroy, D.; Heinze, C.; Appla, D.; Busetta, P.; Tidhar, G.; and Rao, A. 1997. Toward Credible Computer-Generated Forces. In *Proceedings of the Second International Simulation Technology and Training Conference (SimTecT '97)*, eds. S. Sestito, P. Beckett, G. Tudor, and T. J. Triggs, 234–239. Melbourne, Australia: SimTecT.
- Rao, A. S. 1997. A Unified View of Plans as Recipes. *Contemporary Action Theory*, eds. G. Holmstrom-Hintikka and R. Tuomela, 309–331. Dordrecht, The Netherlands: Kluwer Academic.
- Senge, P. 1998. *The Fifth Discipline: The Art and Practice of the Learning Organization*. Scoresby, Australia: Random House.
- Tambe, M.; Schwamb, K.; and Rosenbloom, K. S. 1995. Building Intelligent Pilots for Simulated Rotary Wing Aircraft. Paper presented at the Fifth Conference on Computer-Generated Forces and Behavioral Representation, 9–11 May, Orlando, Florida.
- Tambe, M.; Jones, R. M.; Laird, J. E.; Rosenbloom, P. S.; and Schwamb, K. 1994. Building Believable Agents for Simulation Environments: Extended

Abstract. In *Collected Papers of the SOAR/IFOR Project*, 78–81. Marina del Ray, Calif.: University of Southern California Information Sciences Institute.

Tidhar, G.; Heinze, C.; and Selvestrel, M. 1998. Flying Together: Modeling Air Mission Teams. *Applied Intelligence* 8(3): 195–218.

Tidhar, G.; Murray, G.; and Steuart, S. 1995. Computer-Generated Forces and Agent-Oriented Technology. Paper presented at the Australian Joint Conference on Artificial Intelligence, Workshop on AI in Defense, 6–7 July, Canberra, Australia.

Tidhar, G.; Heinze, C.; Goss, S.; Murray, G.; Appla, D.; and Lloyd, I. 1999. Using Intelligent Agents in Military Simulation, or “Using Agents Intelligently.” In *Proceedings of the Eleventh Innovative Applications of Artificial Intelligence Conference*, 829–836. Menlo Park, Calif.: American Association for Artificial Intelligence.

Wooldrige, M. J., and Jennings, N. R. 1999. Software Engineering with Agents: Pitfalls and Pratfalls. *IEEE Internet Computing* 3(3): 20–27.

Clint Heinze is a cognitive scientist currently employed by the Defense Science and Technology Organization as a military analyst. Filling a role on the boundary between the software engineering of multiagent systems for the modeling of tactical decision making in military flight simulation and military operations research, he maintains an interest in software engineering, computational cognitive modeling, and intelligent agent technologies. From an undergraduate background in aerospace engineering, Heinze is currently completing a Ph.D. in the Department of Computer Science and Software Engineering at the University of Melbourne. His e-mail address is clinton.heinze@dsto.defence.gov.au.

Simon Goss holds a Ph.D. in physical chemistry from La Trobe University and Grad Dip (KBS) from RMIT. He is a senior research scientist in the Air Operations Division of DSTO. He leads the long-range research in Operations Research, where intelligent agents and multiagent systems are used for simulation and modeling. His interests lie in the interface of knowledge acquisition, software engineering, cognitive science, and operations research. He is a member of AAI, ACM, IEEE, AORS, and ACS and is a senior fellow in the Agent Lab at the University of Melbourne.

Torgny Josefsson is a principal research scientist with the Air Operations Division of DSTO. He graduated from Monash University with a BSc(Hons) (1989) and a Ph.D. (1993). After four years of post-doctoral work at the School of Physics, University of Melbourne, he joined Maritime Operations Division (MOD) Stirling in 1997. At MOD, his work involved operational analysis to support tactical development for Collins class submarines and studies into the enhancements of submarine capabilities. Since joining the Air Operations Division in 1999, he has primarily been involved in developing new concepts of operation for the AP-3C maritime patrol aircraft.


Kerry Bennett is currently deputy chief executive office of Neurosciences Victoria Limited and associate professor within the Centre for Neuroscience, The University of Melbourne, Victoria, Australia. Her research focus is on the neurophysiological mechanisms of human behavior, with particular emphasis on perception, attention, and cognition. With respect to intelligent agent technology, her interest lies in the application of perceptuomotor integration principles for motor control and learning to the design and management of multiagent systems.

Sam Waugh is a senior research scientist with the Air Operations Division of the Defence Science and Technology Organisation, Melbourne, Australia. He received a B.Sc. and a Ph.D. in computer science from the University of Tasmania, Australia, in 1991 and 1997, respectively. His research interests include machine learning and software engineering.

Ian Lloyd is head of mission and campaign analysis in the Air Operations Division of Defense Science and Technology Organization. He has extensive experience in military and civil aviation, including some time as an airline pilot. He has also worked on analysis of missile guidance and developed a manned simulator for rapid prototyping of cockpit displays and a terrain-referenced navigation system. For the last five years, he has been conducting operational analysis in support of the Airborne Early Warning and Control (AEW&C) acquisition. He is currently extending the methods and tools used for the AEW&C to studies supporting AP-3C and F-111 acquisitions and operations.

Graeme Murray has an academic background in physical chemistry and pure mathematics and holds a BSc, an MSc, and an MA. He joined the Operational Research Group at the Aeronautical Research Laboratories (now part of the Aeronautical and Maritime Research Laboratory) in 1969 as an analyst and has wide experience in military operational research and defense issues. He initiated the Aircraft Systems Division’s research program in the application of AI to defense problems in 1989 and led the Air Operations Division’s research contribution to the Cooperative Research Center for Intelligent Decision Systems from 1991 to 1996 as SWARMM project leader. His research interests include improvement of the representation of human decision making in computer simulation models. He is a member of the Australian Society for Operations Research.

John Oldfield is a squadron leader in the Royal Australian Air Force. He served in the 92 Wing of the Maritime Patrol Group on P-3C *Oriens* as an acoustic operator. He was the 92 WG Standards FLTCDR between 1998 and 2001. He is currently attending Joint Services Command and Staff College at Shrivenham in the United Kingdom.



Artificial Intelligence and Mobile Robots Case Studies of Successful Robot Systems

Edited by David Kortenkamp, R. Peter Bonasso, and Robin Murphy

The mobile robot systems described in this book were selected from among the best available implementations by leading universities and research laboratories. These are robots that have left the lab and been tested in natural and unknown environments. They perform many different tasks, from giving tours to collecting trash. Many have distinguished themselves (usually with first- or second-place finishes) at various indoor and outdoor mobile robot competitions.

Each case study is self-contained and includes detailed descriptions of important algorithms, including pseudocode. Thus, this volume serves as a recipe book for the design of successful mobile robot applications. Common themes include navigation and mapping, computer vision, and architecture.

6 x 9, 400 pp., ISBN 0-262-61137-6

**To order, call 800-405-1619 (<http://mitpress.mit.edu>).
Distributed by The MIT Press, 55 Hayward, Cambridge, MA 02142**