Issues affecting the Control and Co-ordination of Agent Teams in the Simulated Battlefield

Graham S. Horn and Jeremy W. Baxter

Defence Evaluation and Research Agency
St. Andrew's Road
Malvern
Worcestershire
WR14 3PS
United Kingdom
{ghorn, jbaxter}@signal.dera.gov.uk

Abstract

The control and co-ordination of groups of agents in battlefield simulations is difficult. Some of the important issues include the production of plans, the assignment of agents to roles and the monitoring of plan execution. A framework, based on theories from multi-agent systems, has been implemented to support co-ordinated group behaviours. This has been used for plan execution, monitoring and repair. Current work is focused on mixed-initiative planning, making use of existing terrain-based planners.

Introduction

We are researching AI techniques for providing intelligent opposing forces within battlefield simulations. In particular we are working on an agent-based system for the planning and execution of actions of small groups of tanks. The main emphasis of the work is to produce a framework to ensure co-ordination between groups engaged in a common task. We also plan to improve the control of the system by enabling interaction by the human operator, who might wish to modify the behaviour of the agents during execution. The main difficulty here is preventing the user interaction from damaging the co-ordination of the agent teams.

Issues

The main issues that are of particular relevance to our work and the workshop themes are described below.

Mixed-Initiative Planning and Control

Previous work focused on terrain reasoning and planning in an agent-based system (Hepplewhite and Baxter 1998). We plan to extend this to provide some level of mixed-initiative planning, with the user being able to modify both routes and assignment of agents to sub-teams. Typical mixed-initiative planning issues will apply. Also it is essential to represent the commitments and obligations between agents in the plans produced for agent teams. These commitments must be made explicit so that the user does not break them when altering plans.

Resource management

The agents in a team are normally only capable of carrying out one role at any given time. The assignment of agents to sub-teams will depend on a number of factors, such as the capabilities possessed and physical location. This assignment will normally be automated, but the user may wish to contribute to the selection process (e.g. add the requirement that "the troop on the left should operate on the left flank, and the troop on the right should be fire support").

The user may wish to take control of some sub-teams during execution, but intend for them to be returned to computer control at a later time and to maintain the commitments between sub-teams throughout. For example, the user may wish to navigate a sub-team through a narrow passage that is not part of the planned route, since it is a short cut, before relinquishing control. If the commander of that sub-team is not aware of the user's intention the overall plan might fail.

Alternatively the user may wish to use the sub-team for a different purpose, and not remain committed to the group plan. This means that the group commander may need to re-plan and should not wait for the units now under user control.

Execution, Monitoring and Repair

Most of our work of the past year has concentrated on execution and monitoring issues (Baxter and Horn 2000). We have developed a framework that uses a device called a Co-ordination Matrix (CM) to co-ordinate groups of agents. Each agent has a CM that contains its orders, and another (produced internally) which contains orders for its sub-groups. A CM contains a sequence of group states and a number of sequences of tasks for sub-groups. Each task contains completion conditions, constraints and conventions on how to co-ordinate, as well as information specific to the type of task.

The user of the system may want to look at the internal state of an agent in order to understand its actions. At present the user of our system is able to look at the two current CMs for each agent. The current tasks and states are highlighted and information on any task or state can be accessed. However, as the number of agents in the system grows it is not feasible to display information for each agent. It might be useful to have an additional agent that chooses which information should be displayed, in order to best inform the user of the overall state of the system.

Co-ordination Techniques

The multi-agent systems literature has been studied to identify how theories of group action can be applied to this domain. Joint Intentions, as described by Cohen and Levesque (1990), is a formal model of co-operation that provides a set of definitions of what it means for a group of agents to hold a joint goal and to jointly intend to do something. This model was developed from the BDI agent model (Rao et al., 1993) which ascribes beliefs, desires and intentions to agents and uses this to formulate and predict their behaviour. An agent may have a desire to achieve something (a goal) but is described as intending to do it only if it is executing a plan which it believes will achieve that goal. Cohen and Levesque extend this notion to groups using the example of two drivers, one of which is following the other until he knows his way home. They identify that in order for co-operative behaviour to be robust it is not sufficient for both agents to mutually believe[1] that they are carrying out an action together, they need to be committed to keeping the other members of the group informed about their beliefs and intentions. The key to avoiding problems is that this commitment persists even when an individual agent believes the action has been completed, is impossible or irrelevant.

Joint Intentions theory therefore provides a framework within which groups of agents can hold mutual goals and execute mutual actions and describes the cases where some form of communication is necessary to ensure that agents acting together maintain a coherent state. Jennings (1993) describes the important role of commitments and conventions in co-ordination schemes. Commitments are made by agents to carry out certain actions and they agree on a set of conventions on how to monitor their commitments.

[1] Mutual belief means I know that you know that I know… and so on to an infinite level of nesting. In practice it is a requirement that agents are confident that they all believe the same thing.

Ideas from Joint Intentions have been successfully applied by Tambe to improve the team behaviour of attack helicopters (Tambe 1998). The main difference between our framework and the STEAM rules produced by Tambe is the presence of an agent representing the group as a whole that is responsible for instructing and co-ordinating all group members.

The Co-ordination Matrix

The purpose of the Co-ordination Matrix is to provide a single repository for the information required by a group to carry out a co-ordinated activity. The subordinate members of the group do not have the authority to modify the tasks that it contains, but are free to re-plan internally generated tasks. Group members must ensure that beliefs that affect the Co-ordination Matrix are mutually held. Moreover, information that does not affect the Co-ordination Matrix does not need to be shared.

Current Work

Many of the issues described above have yet to be addressed in our work. We are currently implementing a mixed-initiative planning system which will be used to plan a Squadron (US Company) sized assault. The planners will be based on existing fully automated planners (Hepplewhite and Baxter 1998) that make extensive use of terrain reasoning.

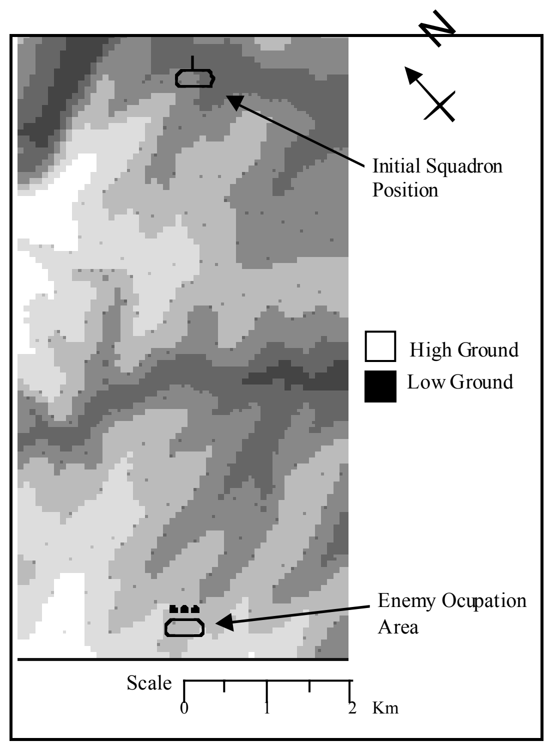

Figure 1: Terrain Profile

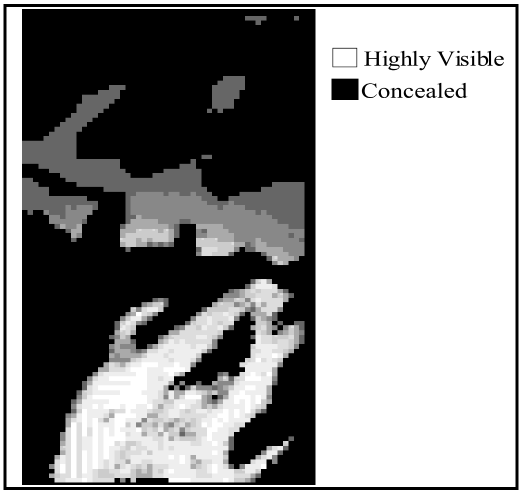

Figure 2: Visibility from the enemy

The area under consideration is represented in the planners using grids of costs. These costs will be displayed as overlays on the map. This aids the understanding of the user because the options available and their impact can be displayed in a format s/he is familiar with. Using a geographically related grid of costs ensures that the machine reasoning components make decisions clearly related to the map features accessed by the user.

For example, figure 1 shows the terrain profile of the area of responsibility for the assaulting squadron. Figure 2 shows an example cost map of the visibility from the enemy area. The scale represents how well each position can be observed, lighter points represent positions which can be seen better and from more positions in the enemy area. Other costs that will be considered include distance, gradient and visibility of the objective.

Summary

In order to improve the control and co-ordination of battlefield simulation agents a number of planning and execution issues need to be addressed. A framework, based on co-ordination theories from the multi-agent systems literature, has been implemented. This is currently being used to investigate the execution, monitoring, and repair of plans represented by Co-ordination Matrices. The production of a mixed-initiative planning system is in the preliminary stages. User interaction at runtime has yet to be addressed.

Acknowledgements

This work was funded by Technology Group 5 (Human Sciences and Synthetic Environments) of the MOD Corporate Research Programme.

References

Hepplewhite, R. T. and Baxter, J. W. 1998. Planning and Search Techniques to Produce Terrain Dependent Behaviours. In *Proceedings of the 7th conference on Computer Generated Forces and Behavioral Representation*, Orlando, Florida.

Baxter, J. W. and Horn, G. S. 2000. Representing and Executing Group Tasks. In *Proceedings of the 9th conference on Computer Generated Forces and Behavioral Representation*, Orlando, Florida. Forthcoming.

Levesque, H., Cohen, P. and Nunes, J. 1990. On Acting Together. In Proceedings of the Eighth National Conference on Artificial Intelligence, 94—99.

Rao, A., Lucas, A., Morley, D., Selvestrel, M. and Murray, G. 1993. Agent-oriented architecture for air-combat simulation. Technical Note 42, Australian Artificial Intelligence Institute.

Jennings, N. R. 1993. Commitments and conventions: the foundations of coordination in multi agent systems. *Knowledge Engineering Review*, 8(3):223—250.

Tambe, M. 1998. Implementing Agent Teams in Dynamic Multi-agent Environments. Applied Artificial Intelligence, vol. 12.