Co-ordinating Agents in the Simulated Battlefield

Graham S. Horn and Jeremy W. Baxter
Defence Evaluation and Research Agency
DERA Malvern
St Andrews Road, Malvern
WR14 3PS, U.K.
{ghorn, jbaxter}@signal.dera.gov.uk

ABSTRACT

Co-ordinating groups to carry out a complex task is difficult in the real world and equally challenging in the simulated battlefield. This paper outlines the core of a framework, based on theories from multi-agent systems, which is used to support group behaviours. It makes use of military institutions such as the command structure, and the associated norms and conventions.

Keywords

Multi-agent systems, simulation, co-ordination, behavioural frameworks.

1. 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. The need for such a framework became apparent in previous work [1] when more complicated group behaviours were put together using sets of existing simple behaviours. The framework should enable such compositions of behaviours to be straight forward, and allow the co-ordination to be flexible. To this end we are making use of military institutions such as the command structure, and the norms and conventions that come with it.

A typical problem in this domain that requires successful co-ordination is a squadron assault. Two sub-groups are needed – one to give supporting fire while the other moves to a concealed form up position before assaulting the enemy position. It is essential that the two groups act in a co-ordinated manner. Supporting fire should commence before or at the same time as the assault group begins its attack, and should cease as the assault group closes with the enemy (to prevent fratricide). A number of failure cases must be accounted for, including: destruction of one of the sub-groups; discovery of other enemy positions; loss of individual vehicles. These might lead to re-planning of the assault, or its abandonment.

2. MILITARY COMMAND HIERARCHY

The agents are organised into a command hierarchy based upon the military command structure. This is an institution that has a number of norms and conventions associated with it. For example, communication of orders flows down the hierarchy, whereas intelligence should be shared between peers and sent to superiors (See figure 1). The responsibilities of the agents are clearly defined by this structure.

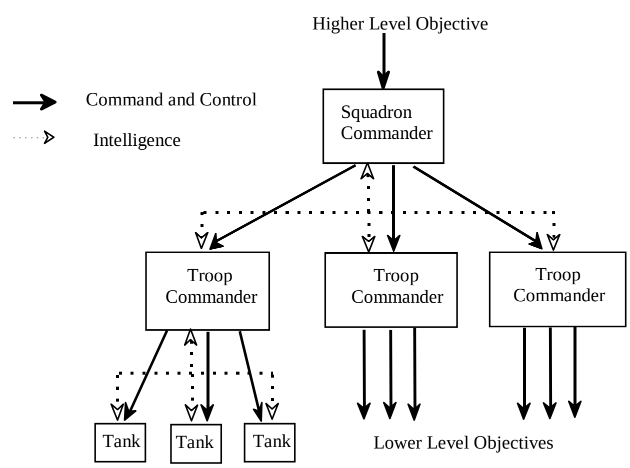

Figure 1. Military Command Hierarchy

The hierarchy enables the planning of complex group orders to be divided up into several smaller problems. A commander agent will produce a plan containing tasks for different sub-groups, each of which is assigned to one of his subordinates who will in turn produce more detailed plans for their tasks. Commander agents are responsible for monitoring the progress of their group and sub-groups towards achieving the goals assigned to them. When success or failure has been detected the usual convention is to notify the agent's superior. For example, when a Troop Commander detects failure a message should be sent to its Squadron Commander. Different co-ordination norms include carrying out actions at times specified relative to a common clock or using messages to trigger the start of co-ordinated actions.

3. 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 [2], 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 [3] 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 [4] 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.

Ideas from Joint Intentions have been successfully applied by Tambe to improve the team behaviour of attack helicopters [5]. We have used Joint Intentions to determine when agents need to send messages and when to wait for other agents to complete their task(s) before continuing with the next task.

4. THE CO-ORDINATION MATRIX

We have designed an object called a Co-ordination Matrix (CM) that is used to represent the orders received or produced by a commander agent. These two forms are known as Instructed (ICM) or Constructed (CCM). The ICM represents the tasks a group must achieve and how they relate to the super-group and its other sub-groups. The CCM contains tasks for sub-groups belonging to the commander agent that produces it and is converted to an ICM when it is sent as an order. The CCM is used to monitor and control the progress of the sub-groups towards completing the tasks in this order.

The Co-ordination Matrix consists of group states and task sequences for the sub-groups. Each task or state has an associated completion condition, constraint and convention. The condition and constraint may be a conjunction or disjunction of conditions and constraints. When the completion condition is satisfied the state or task is advanced and the convention is used to determine any co-ordination actions, for example whether to communicate

[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.

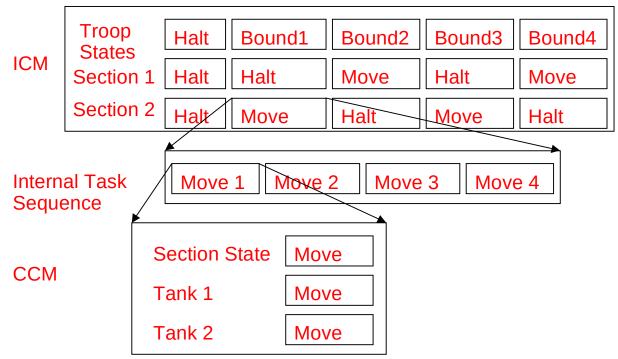

Figure 2. Coordination Matrix Expansion

this change to the superior or subordinate agent(s). These completions can be based on observations, such as a vehicle crossing a phase line, or be related to the group state or position in the task sequence. Constraints are typically based on observations, such as the number of enemy vehicles visible.

When a commander agent receives an ICM he identifies which task sequence is assigned to the group that he commands. For each task one or more CCMs will be produced and sent as orders to the commanders of the sub-groups. Figure 2 shows the expansion of an overwatch move where a group has been split into two sections, one of which halts and watches over the other while it moves forward. The ICM shown belongs to Section 2 whose commander has broken the first move order into a sequence of intermediate moves. The CCM is for the first of these intermediate moves and contains move orders for the two tanks belonging to the section.

The commander agent is able to re-organise his assets for each CCM, increasing or decreasing the number of sub-groups as appropriate to the task in hand. In the overwatch example the commander of the Troop (typically 3 or 4 tanks) split the group into two Sections when producing the order represented by the ICM in figure 2. The commander of Section 2 split his group into two sub-groups, each consisting of a single tank.

Each agent has a set of rules that describe how to use the Co-ordination Matrix to co-ordinate with its peers (for an ICM) or subordinates (for a CCM). These rules are derived from Joint Intentions theory and are designed to ensure the group as a whole maintains a shared group state and therefore remains co-ordinated. The convention associated with a task/state indicates what to do on completion of that task/state. Other rules deal with what to do on failure. Completion conditions based on task/group states are used to co-ordinate the actions of the sub-groups in carrying out an order. Using the overwatch example again, suppose that the current group state is Bound 1 and Sections 1 and 2 are carrying out the tasks Halt and Move, respectively. Section 2 finishes its Move and it transitions to its next task, which is a Halt. The completion condition on the group state Bound 1 is that Section 2 is in this Halt state. So the group state can now advance to Bound

2. Section 1 can then begin its Move because the completion condition on its Halt task was based on the group state being Bound 2. Completions based on group states are therefore used to advance the position in a task sequence and those based on task position are used to advance the group state.

In our current implementation a Co-ordination Matrix contains two representations of the current group state denoting local and mutual belief. Completions based on group state can use either of these. Work to date has only used the mutual group state which can only be advanced by subordinates upon receipt of a message from the commanding agent. Therefore in the current implementation only the superior agent (who is monitoring the CCM) uses the completions based on the position in the task sequences. In the example above this means that the Troop commander uses the completion to advance the mutual group state in his copy of the order (the CCM) and sends messages to the commanders of the two Sections which enable them to advance the mutual group state in their copies (the ICMs).

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.

5. SUMMARY AND FUTURE WORK

We have described a device (the Co-ordination Matrix) that can be used to co-ordinate the execution of group tasks. It allows for different conventions to be applied to the constituent tasks and group states. Future work will investigate the use of different conventions, and will also investigate changing them during execution. Currently the only convention implemented is to send a message when the group state or task changes, but we might wish to enforce radio silence – in which case the agents will need to reason about when messages should be sent. Further details of our current implementation can be found in [6].

The military command structure is the institution that provides norms for the agents such as obligations to superiors, ability to delegate to subordinates, etc. and gives context for the use of the Co-ordination Matrix. We plan to extend this structure so that the agents also have responsibilities to the human operator of the system, so that s/he can interact with the agents in order to modify their behaviour without damaging the co-ordination of the groups.

6. ACKNOWLEDGMENTS

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

7. REFERENCES

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

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

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

[4] N. R. Jennings. Commitments and conventions: the foundations of coordination in multi agent systems. Knowledge Engineering Review, vol. 8 no. 3, pages 223—250, 1993.

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

[6] J.W. Baxter and G. S. Horn. Representing and Executing Group Tasks. In Proceedings of the 9th conference on Computer Generated Forces and Behavioral Representation, Orlando, Florida, May 2000 (to appear).