Executing plans in a tank battle simulation.

Jeremy W. Baxter
Defence Research Agency
St. Andrews Road, Malvern, Worcestershire WR14 3JP
UNITED KINGDOM
Email: jbaxter@signal.dra.hmg.gb

Abstract

This paper briefly describes some of the plan execution problems prevalent in the simulated battlefield domain and how they can be addressed. The system described uses generalised actions which are expanded to low level actions at the time they are executed, ensuring their relevance to the prevailing situation. The continuing value of plans is ensured by monitoring the assumptions and constraints used during planning and calling for re-planning when a violation is detected. An example of the operation of the planner is given and potential future enhancements briefly discussed.

1 Introduction

As part of work on battlefield simulations our project is investigating AI techniques for providing plausible automated behaviour for tank commanders. From a planning perspective the battlefield domain is extremely complex. The problem is continuous in space and time, and can contain many other battlefield agents, both friendly and hostile. The future, and possibly the present, state of the battlefield is uncertain. A planning system [1], based on a POPLOG toolset [2], has been developed which generates plans by combining a series of pre-determined potential actions on the basis of their costs in distance, time and risk. To enable our tank command agents to estimate the risk involved in performing an action they carry out a simplified simulation of the actions and use this cost as a basis for selecting the most appropriate actions. This technique is similar in its underlying philosophy to that of Lee and Fishwick [3]. Execution of these plans takes place in an uncertain and highly dynamic world and so execution has to be tightly linked with up to date information about the world and re-planning activated when failure is detected.

2 Domain description

The battlefield domain can be characterised primarily as being highly dynamic with numerous semi-independent agents altering the world. In addition to this the information available about the world is restricted by the capabilities of a vehicle's sensors and may be inaccurate. The success of an action cannot be guaranteed because the agent may be destroyed or damaged during execution. Other features such as changing objectives and interruption of actions also feature in the problem because the dynamism of the world combined with uncertain and incomplete information mean that it is extremely rare that a plan can be executed in its entirety. This leads to a requirement that the plan execution system should be fully aware of the assumptions used when the plan was made and be able to either adapt or reject an action by comparing the present situation with the situation envisaged when the action was incorporated into the plan.

3 Planning to fail

In such a fluid situation the plan execution system needs means of rapidly assessing plan failure. In our application this means that the plan execution system has to monitor several features. First, and most obviously the potential achievement of the present goal has to be monitored, second the system must check if any of the constraints on the plan have been violated and third it must consider whether any of the assumptions used in the planning process no longer apply.

3.1 Monitoring Assumptions and Constraints

An assumption represents some feature of the environment which the planner has identified as crucial to the success of the action being performed. For instance in our tank battle simulations a plan implies an assumption about the acceptable level of risk. While the agent is considering actions for inclusion within a plan it predicts the threat which each known opposing agent would pose to it during that action, opposing agents are assumed to keep on a constant heading and velocity. The threat level predicted for the most dangerous opponent to an action is recorded and used as an assumption of the largest individual threat level which will be perceived during the action.

If during execution of that action within the plan an entity is encountered which poses a greater threat than was foreseen, either due to a known entity moving to a more dangerous position or a previously unknown entity being detected, then re-planning takes place.

A constraint represents some state of the world the planner was attempting to avoid. Its violation requires an adjustment to the plan. Examples are that an agent should be in a loose formation with respect to other friendly agents, or should be within a specified corridor or area. Violation of a constraint triggers a re-assessment of the plan so that actions which restore the desired state can be considered. In an ideal system the particular constraint or assumption violated would be used to guide the planner in a plan repair process rather than requiring a full re-plan as is currently the case in our system.

Care has to be taken to prevent transient violations of constraints from causing unnecessary re-planning episodes. For example in the transition period between actions the maximum threat assumption has to be set to be the greater of the levels associated with both actions to allow time for any reduction in threat predicted by the execution of the new action to be detected.

3.2 Using Generalised Actions

One of the ways we have attempted to prolong the life of plans in dynamic situations is to use generalised actions which, while having been evaluated in detail by the planner, are re-evaluated when they are executed to conform to the situation as it is at the moment of execution rather than the situation envisaged by the planner. This allows minor corrections to be made as an action is executed and for the detection of inappropriate actions if the situation has changed drastically. This is similar to the Reactive Action Packages used by Firby [4]. As an example one of the actions available in our planner is 'face threat' where a tank turns to face the direction of hostile fire to ensure any hits it receives will be on its stronger frontal armour. The heading it faces may well be different from the heading which would have been generated when the plan was made since the opposing tank may have moved to a different position. The action is flexible enough to cope with such minor differences. If the opposing tank has been destroyed or retreated out of sight the action is no longer appropriate and would therefore trigger a re-planning episode.

One of the dangers of this approach is that a series of minor corrections to actions can accumulate to give an overall outcome radically different from that envisaged by the planner. In our experience these differences are usually detected by the violation of a constraint or assumption. Our plans are also fairly short, typically consisting of two to five actions each lasting ten to fifteen seconds which reduces the possibilities for differences to accumulate.

If the system is to operate in real time temporary replacements for inappropriate actions are necessary to fill in until a full new plan is available. These replacements need to be linked to the way in which the action failed. In our example if a 'face threat' action fails due to lack of a threat it is reasonable to replace it with a 'go-aim-goal' action which causes the tank to head for its present goal position. Actions which fail due to increased threat levels are replaced by 'face threat' actions as a stop-gap measure until a plan accounting for the new or increased threat is ready.

4 Example

Figure 1 shows an example of the sort of situation in which re-planning is called for by an agent. The agent in the bottom right of the figure is advancing to a goal location in the top left. It has a plan consisting of a series of 'go-aim-goal' actions which cause the agent to head direct for the goal position. When it reaches point A it spots an opposing agent which has just crested a hill 2km away at point P. As part of its situation assessment the agent identifies the potential threat has increased from zero, the expected level for its present action to a much higher level, representing the risk posed by the newly spotted opposing agent. The violation of this threat assumption causes a re-planning episode to begin.

While the agent considers its options the now inappropriate 'go-aim-goal' action is replaced by a 'face-

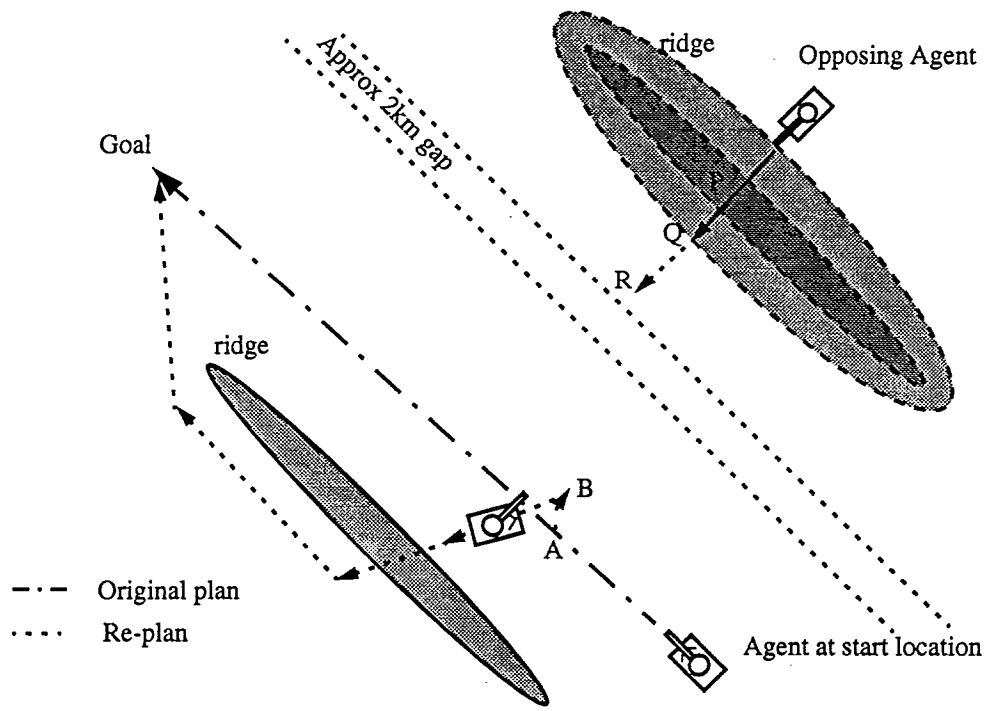

FIGURE 1. A Re-plan forced by the appearance of an opposing agent

threat' action which causes the agent to steer towards the threat and begin to slow down. This takes it to point B, at which point the planner has come up with a new plan consisting of two 'retreat-from-threat' actions which make the agent reverse over a nearby small ridge, a 'follow-ridge' action which causes it to parallel the ridge and a 'go-aim-goal' action to complete the plan. The agent recognises that there will be a continuing, but reduced, threat as it reverses into cover, followed by a zero threat level while it is masked by the ridge.

As the agent executes the plan the situation diverges from the predictions which the plan was based upon. As shown in Figure 2 the agent assumed that the opposing agent would continue on its previous path and would be at point R when the agent finished the first retreat action at C. In reality the other agent made a sharp turn and by the time the agent finishes the first action in the plan it is at point S. This has not violated the threat assumption made for the plan but when the second 'retreat-from-threat' action is executed its heading is adjusted to take account of the actual position of the threat, thus altering the path taken by the agent. This illustrates how minor adjustments to the plan can be made during execution so long as the real situation does not differ significantly from that envisaged during planning.

5 Discussion and Conclusions

Our approach to plan execution has been to link it in tightly with the assumptions about the world made by the planner and to monitor these assumptions as well as checking for achievement of the goal. By allowing actions to be interpreted as they are executed our system can respond to minor changes in the situation and detect when actions are no longer applicable. This enables the system to operate in a highly dynamic world. The difficulties lie in identifying what assumptions and constraints should be monitored and how much actions should be permitted to differ from the effect envisaged during planning. At present the system concentrates almost exclusively on short term plans to ensure its survival. We hope to integrate longer term planning to guide groups of tanks in the form of a commander agent or agents.

Another aspect which needs further consideration is how pro-active agents should be. It is possible for agents to gather information which affects the viability of future stages of their plans. For instance an agent may have a plan which involves crossing a ridge line. While executing the plan it receives a message informing it that there are enemy forces on the other side of the ridge. With our present system the agent will not be able to react to this until it notices the increased threat level as it crosses the

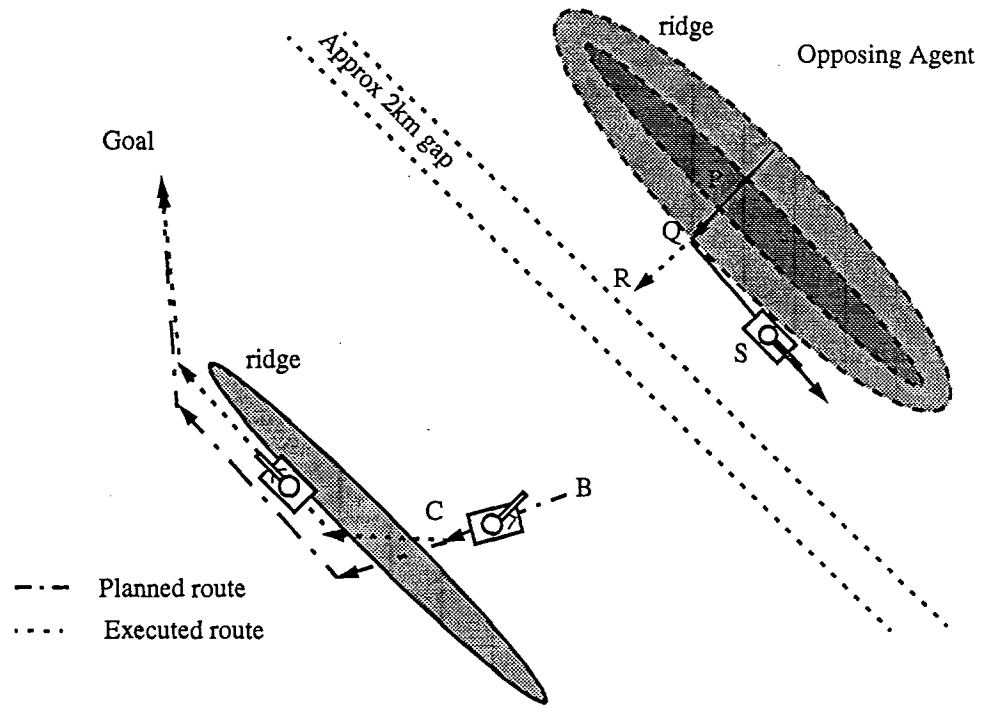

FIGURE 2. Adjustment to the plan during execution

ridge. A more pro-active agent might be able to predict the increase in threat level and immediately reassess its plan. This would require an ability to recognise and deal with potential future threats. It is a difficult problem to decide how and when an agent should consider revising or abandoning a plan due to a potential rather than a real violation of its assumptions.

6 References

1. Baxter J W and Hepplewhite R T: 'Intelligent Agents for Battlefield Simulation', *AISB Quarterly* Autumn 1995 no.93 pp 26-33
2. Sloman, A and Poli R:'SIM_AGENT: A toolkit for exploring agent designs', in *ATAL-95, Workshop on Agent Theories, Architectures, and Languages*, IJCAI-95, Montreal, August 1995. To be published 1996 by Springer-Verlag Also Cognitive Science technical report: CSRP-95-3, The University of Birmingham.
3. Lee J J and Fishwick P A: 'Incorporating Simulation-Based Models into Planning Systems', *Proceedings of the 5th Conference on AI, Simulation, and Planning in High Autonomy Systems*, Gainsville Florida December 7-9 1994: IEEE Computer Society Press Los Alamitos, California pp 113-119
4. Firby J: 'Adaptive Execution in Complex Dynamic Worlds.' PhD diss., Dept. of Computer Science, Yale University. 1989