

Agents in Tank Battle Simulations

Jeremy Baxter and Richard Hepplewhite

Networks of computers can be used to produce a digital virtual environment (DVE) where multiple participants can interact. This technology is extremely attractive to the military to provide training simulations. By the use of mock-up vehicles and high-fidelity visual systems, trainees get a window onto a virtual world populated by simulated vehicles interacting over a realistic terrain surface. Some of these vehicles are controlled by human trainees, others by computers. It is essential that trainees find the behavior of the computer-controlled vehicles realistic. Currently most computer forces are semiautomated using finite state machines or rule bases to govern their behavior, but requiring constant supervision by a human controller [1]. However, the vehicles can become increasingly autonomous as AI and agent techniques develop, thus reducing the number of human controllers as well as the hefty manpower bill associated with running big training simulations [2, 3].

We are concentrating on developing agents to control tanks within ground battle simulations. Here, tactical behavior is governed by two main factors—the terrain over which the tanks are moving and their beliefs about the enemy. In trying to produce battlefield behavior that mimics a human tactician, it is advantageous to model the command structure used by the army. This helps with the gathering of knowledge from subject matter experts and enables a hierarchical decomposition of the problems. The figure appearing in this sidebar shows the hierarchy of agents—high-level commanders are given objectives that are used to produce lower-level objectives for their subordinates.

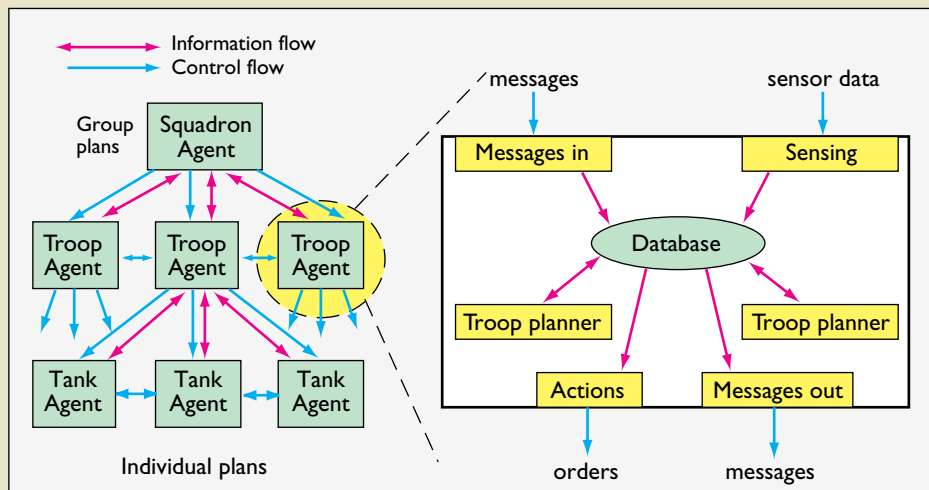
Information flows both up and down the command chain and agents need to cooperate with their peers to achieve the overall goal set by their commander. This natural decomposition of the problem allows higher-level agents to work on long-term plans while the individual tank agents carry out orders designed to achieve more immediate objectives.

The Agent Toolkit. To provide the framework within which agents operate, we use the SIM_AGENT toolkit (see the article by Sloman and Logan in this section). It allows multiple agents to be run and controls their communication with each other and with the physical simulation of the battlefield. Internally, these agents run a number of processes which share data held in a central database; as shown in the figure. The processes are scheduled to run for a few steps at a time and each performs a different task, for example, assessing incoming sensor data, monitoring the progress of a plan, or communicating with other agents. This allows agents to pursue many mental tasks simultaneously. Scheduling ensures each agent and process gets a fair share of the available processing power and enforces real-time operation.

Agents need to incorporate fast reactions to their environment to cope with unexpected events while at the same time perform complex reasoning about the terrain. Our agents, therefore, combine the use of anytime planning techniques with reactive plan execution systems designed to operate in a real time environment [2]. These are implemented as separate processes within the agents allowing the combination of reactive and deliberative behavior.

For example, when considering how to place forces

to block enemy movement through an area, the squadron commander has to consider a number of factors. Positions must be identified that give protection to the defending forces, but also provide a good view of the potential enemy approach routes and are close enough to other



Hierarchy of agents and internal structure of an agent

groups to offer mutual support. The squadron defensive planner identifies candidate positions for the defending troops by analyzing the protection afforded by the terrain. Combinations of these positions are ranked in terms of the overall breadth and depth of the engagement area which can be seen (and fired upon) from them. During this optimization process, the best deployment identified so far is cached, so that it can be executed if the time for planning runs out.

The battlefield is so dynamic that detailed individual tank plans are unlikely to remain valid for long, so tanks operate by selecting actions from a recipe book covering general situations. These short-term actions combine the agent's goals with reactions to the enemy and the terrain. Plans are developed as sequences of these actions and are assessed by carrying out an internal simulation of their probable effects to identify how well they would perform in the present situation. The use of an internal simulation also allows assumptions about the future state of the world to be incorporated into the plan. During execution, the agent can identify cases where the assumptions turn out to be false and a new plan is required.

Simulations using these agents have been shown to military experts who have confirmed their terrain-related behavior is more realistic than that produced by simple, finite state machine-based approaches. The agent-based approach to this problem has several advantages, including decomposition of the problem, natural distribution between machines, and permitting clear comparisons with reality by human experts. Future work will focus on the deficiencies in overall group behavior and a need to plan to gather information. **C**

REFERENCES

1. Courtemanche, A.J., and Ceranowicz, A. ModSAF development status. In *Proceedings of the 5th Conference on Computer Generated Forces and Behavioural Representation*. (May 9–11, Orlando Fla.). Institute for Simulation and Training, 1995, 3–13.
2. Hepplewhite, R.T., and Baxter, J.W. Broad agents for intelligent simulation. In *Proceedings of the 6th Conference on Computer Generated Forces and Behavioural Representation*. (July 23–25, Orlando, Fla.) Institute of Simulation and Training, 1996, 319–327.
3. Hill, R., Chen, J., Gratch, J., Rosenbloom, P., and Tambe, M. Intelligent agents for the synthetic battlefield: A company of rotary wing aircraft. In *Proceedings of AAAI97 and IAAI97* (July 27–31, Providence, R.I.). AAAI Press, 1997, 1006–1012.

JEREMY BAXTER (jbaxter@signal.dera.gov.uk) is a senior scientist at the UK Defence Evaluation and Research Agency (DERA), Malvern, UK.

RICHARD HEPPLEWHITE (rth@signal.dera.gov.uk) is a senior scientist at the UK Defence Evaluation and Research Agency (DERA), Malvern, UK.

human workers and economy of plant operation. Simulated battlefield commanders or simulated antiterrorist strategists may have to detect and handle conflicts between protecting civilians and capturing opponents.

In these contexts an agent may benefit from an additional architectural layer with the ability to monitor, evaluate, and possibly modify internal processes of various kinds (see Figure 3). This could be based on a mixture of internal reactive and deliberative processes directed at the agent's internal sensory buffers, its goals and preferences, deliberative strategies, records of recent decisions, and the like. Such a system might discover conflicts of motivation and devise a strategy for dealing with them. Or it may notice that certain problem solving processes are taking too long, so that a less cautious but more speedy strategy is required.

We expect that some synthetic agents, including agents working in groups or teams, will require a similar variety of internal processing. For instance, metamanagement is required to decide when to stop trying to solve a problem alone and ask for help.

In humans, this "metamanagement" capability seems to be used for a variety of purposes including social control via inculcation of ethical and other standards of self-assessment, and also some kinds of learning in which deliberative processes are evaluated, found wanting, and improved. It also seems to become impaired during times of stress or in certain emotional disturbances.

Some of the internal processes, including both deliberative processes and reactive processes in which chains of associations are needed to solve a problem, may be too slow for an environment where opportunities and dangers requiring immediate action can turn up unexpectedly. A *global alarm system* using fast pattern-recognition mechanisms able to trigger stereotyped internal and external responses could help in dealing with such situations. The need for rapid reactions could make such a mechanism sometimes produce erroneous responses. This trade-off between speed and intelligence can be found in some human emotions and may also occur in synthetic agents. A trainable alarm system might reduce the frequency of mistakes.

Multiprocessing Within Agents

Sophisticated agents will include many concurrent processes performing various kinds of tasks, such as: