Demonstration of Rational Communicative Behavior in Coordinated Defense

Sanguk Noh and **Piotr J. Gmytrasiewicz**
Department of Computer Science and Engineering
University of Texas at Arlington
{noh, piotr}@cse.uta.edu

Introduction

The primary goal of our demonstration is to show results we obtained on communication among artificial and human agents interacting in a simulated air defense domain. For artificial agents, we advocate a decision-theoretic message selection mechanism which maximizes the expected utility of the communicative actions. Thus, the agents compute the expected utility of alternative communicative behaviors, and execute the one with the highest value (Noh & Gmytrasiewicz 1999). Following the principle of maximizing the expected utility, our agents are intended to be rational in their communicative behavior; they send only the most valuable messages, and never send messages they consider to be damaging given the circumstances at hand.

In the anti-air defense the agents, human or artificial, are coordinating their defense actions, in this case interception of threats, with other defense agents. The goal of defending agents is to minimize damages to their territory. To fulfill their mission, the agents need to coordinate and, sometimes, to communicate with other agents. However, since the communication bandwidth is usually limited in a battlefield environment, it is more valuable for a defending agent to be selective as to what messages should be sent to other agents. Endowing the agents with a decision-theoretic method to choose their own communicative behavior given a situation at hand frees our agents from depending on communication protocols, frequently advocated in other work. We feel that relying on protocols drawn up beforehand could lock the agents into suboptimal behavior in unpredictable domain like the battlefield, in which situations that were not foreseen by the designer are likely to occur.

Our approach uses the Recursive Modeling Method proposed before in (Gmytrasiewicz & Durfee 1995). RMM endows an agent with a compact specialized representation of other agents' beliefs, abilities, and intentions. As such, it allows the agent to predict the message's decision-theoretic (DT) pragmatics, i.e., how a particular message will change the decision-making situation of the agents, and how the other agents are likely to react to it. We propose that modeling other agents while communicating is crucial; clearly, without a model of the other agents' states of beliefs it would be impossible to properly assess the impact of the communicative act on these beliefs (Cohen & Levesque 1990). Based on the message's DT pragmatics, our method quantifies the gain obtained due to the message as the increase in expected utility obtained as a result of the interaction.

Our demonstration consists of our RMM and human agents interacting in two different air defense scenarios in cases when communication is, and is not, available. We will show how communication can benefit the agents in coordination tasks, and compare performance of RMM and human agents.

Demonstration Settings

In our implementation, the anti-air defense simulator with communication was written in Common LISP and built on top of the MICE simulator. Our demonstration is intended to compare the performance achieved by RMM team with that of human team, with and without communication, in two different scenarios.

In the anti-air defense domain, two defense units are faced with an attack by seven incoming missiles, as depicted in Figure 1. The warhead sizes of missiles are 470, 410, 350, 370, 420, 450, and 430 unit for missiles A through G, respectively. The positions of defense units are fixed and those of missiles are randomly generated. Each of two defense units is assumed to be equipped with three interceptors, if they are not incapacitated. Thus, they can launch one interceptor at a given state, and do it three times during a course of one defense episode.

For all settings, each defense unit is initially assumed to have the following uncertainties (beliefs) in its knowledge base:

- The other battery is fully functional and has both long- and short-range interceptors with probability 60%;
- The other battery is operational and has only long-range interceptors with probability 20% (In this case,

---

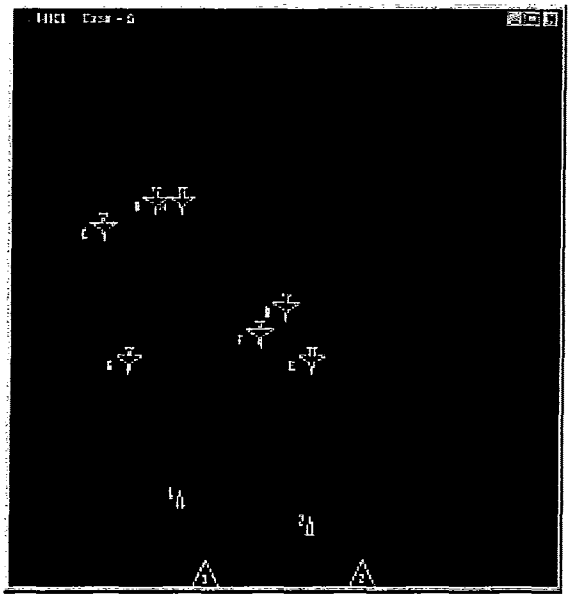

Figure 1: A complex air defense scenario.

it can shoot down only distant missiles, which are higher than a specific altitude.);

- The other battery has been incapacitated by enemy fire with probability 10%;
- The other battery is unknown with probability 10%.

In each demonstrated scenario we allow for one-way communication at a time between defense units. Thus, if both agents want to send messages, the speaker is randomly picked in the RMM team, and the human team flips a coin to determine who will be allowed to talk. The listener is silent and can only receive messages. Each of human subjects is presented with the scenarios, and is given a description of what is known and what is uncertain in each case. They are then asked to indicate which of the 11 messages is the most appropriate in each case. In all of the anti-air defense scenarios, each battery is assumed to have a choice of the following communicative behaviors:

- "No communication."
- "I'll intercept missile A."

...

- "I'll intercept missile G."
- "I have both long- and short-range interceptors."
- "I have only long-range interceptors."
- "I'm incapacitated."

Given the uncertainties and the communicative behaviors, we set up two different scenarios. For each scenario, RMM and human agents intercept incoming targets with and without communication, respectively. We demonstrate their target selection sequences in all settings by retrieving them from http://dali.uta.edu.

To evaluate the quality of the agents' performance, we express the results in terms of (1) the number of selected targets, i.e., targets the defense units attempted to intercept, and (2) the total expected damage to friendly forces after all six interceptors were launched. The total expected damage is defined as a sum of the residual warhead sizes of the attacking missiles. Thus, if a missile was targeted for interception, then it contributed $\{(1 - probability\_of\_hit) \times warhead\_size\}$ to the total damage. If a missile was not targeted, it contributed all of its warhead size value to the damage.

Conclusion

Our demonstration presents the implementation and evaluation of the decision-theoretic message selection used by automated agents coordinating in an anti-air defense domain. When the communication requires valuable bandwidth resources, our message prioritization method can be useful for agents coordinating in a military situation by allowing them to determine the expected utility value for each possible communicative act.

We measure the increase in performance achieved by rational communicative behavior in the RMM team, and compare it to the performance of the human-controlled defense batteries. The results are intuitive: as expected, communication improves the coordinated performance achieved by the teams. An interesting aspect of the demonstration is that it shows the differences between the communicative behaviors exhibited by RMM and human agents. While human communicative behaviors are often similar to those selected by the RMM agents, there are telling differences that, in our experimental runs, allow the RMM team to achieve a slightly better performance. It may be that the differences in processing of probabilistic information about the uncertainties involved explain why decision making achieved by artificial agents tends to be somewhat superior to that of human agents.

References

Cohen, P. R., and Levesque, H. J. 1990. Rational interaction as the basis for communication. In Cohen, P. R.; Morgan, J.; and Pollack, M. E., eds., *Intentions in Communication*. MIT Press.

Gmytrasiewicz, P. J., and Durfee, E. H. 1995. A rigorous, operational formalization of recursive modeling. In *Proceedings of the First International Conference on Multi-Agent Systems*, 125–132. Menlo Park: AAAI Press/The MIT Press.

Noh, S., and Gmytrasiewicz, P. J. 1999. Implementation and evaluation of rational communicative behavior in coordinated defense. To appear in *Proceedings of the Third International Conference on Autonomous Agents*, Seattle, Washington.