Multiagent Coordination in Antiair Defense: A Case Study*

Sanguk Noh and **Piotr J. Gmytrasiewicz**

Department of Computer Science and Engineering
University of Texas at Arlington
Arlington, TX 76019, Box 19015
{noh, piotr}@cse.uta.edu
Office: (817)272-3399, 272-3334, Fax: (817)272-3784

Abstract. This research addresses rational decision-making and coordination among antiair units whose mission is to defend a specified territory from a number of attacking missiles. The automated units have to decide which missiles to attempt to intercept, given the characteristics of the threat, and given the other units' anticipated actions, in their attempt to minimize the expected overall damages to the defended territory. Thus, an automated defense unit needs to model the other agents, either human or automated, that control the other defense batteries. For the purpose of this case study, we assume that the units cannot communicate among themselves, say, due to an imposed radio silence. We use the Recursive Modeling Method (RMM), which enables an agent to select his rational action by examining the expected utility of his alternative behaviors, and to coordinate with other agents by modeling their decision-making in a distributed multiagent environment. We describe how decision-making using RMM is applied to the antiair defense domain and show experimental results that compare the performance of coordinating teams consisting of RMM agents, human agents, and mixed RMM and human teams.

1 Introduction

This paper describes rational decision-making and rational coordination among the antiair defense units facing a missile attack. The task of automated defense units is to defend a specified territory and to coordinate their attempts to intercept the attacking missiles, given the characteristics of the threat, and given what they can expect of the other defense units.

Our approach to this coordinated decision-making problem is based on the assumption that the task of each of the defense units is to minimize the overall damages to the attacked territory. Under a realistic threat situation, friendly defense units cannot expect to have an advanced knowledge of the character of the incoming attack. It is, therefore, crucial that each of the defense units make

* This research has been sponsored by the Office of Naval Research Artificial Intelligence Program under contract N00014-95-1-0775, and by a research initiation grant from the CSE Department of the University of Texas at Arlington.

a decision as to which incoming threat to intercept with an available interceptor by analyzing all potential threats acquired from the radar. In such cases, *coordination* requires an agent to recognize the current status, and to model the actions of the other agents to decide on his own next behavior. Since in any realistic combat situation the integrity of the defense team cannot be guaranteed, and even the very existence of the other friendly units cannot be counted on, relying on pre-established coordination protocols can be suboptimal or even dangerous. Therefore, our approach is that each unit is to independently decide on and execute his action, and that coordination among the units is to emerge on-the-fly as the result of the units' individual rational actions.

We begin by formulating the antiair defense as a decision-theoretic problem from the point of view of an individual defense unit. As we mentioned, the objective is to *minimize damages*. Since the classical notion of a symbolic goal doesn't provide a sufficient basis for choice of action in uncertain situations [13], we need the attributes quantifying the quality of choices in the design of decision-making procedures. First, each attacking missile has its threat value. We compute the missile's threat considering such attributes as the altitude of the missile and the size of its warhead. Further, the defense units should consider the probability with which their interceptors would be effective against each of the hostile missiles. Based on these attributes combined, each unit has to determine the optimal action from his probabilistic decision model.

For the purpose of coordinated decision-making in a multiagent environment, our research uses the Recursive Modeling Method (RMM), previously reported in [2, 3]. RMM enables an agent to model the other agents and to rationally coordinate with them even if no protocol or overall plan can be established explicitly in advance. Using RMM as a decision-making tool, an agent rationally selects his action under uncertainty guided by the principle of *expected utility maximization*. We expect RMM to be appropriate to the antiair defense domain, since the coordinating units have to be able to react to threats in previously unforeseen circumstances. These can include the changing nature of the attack, other defense units being shot at and possibly incapacitated, communication lines broken down, sudden need for complete radio silence, and so on. In these unexpected conditions, relying on globally consistent view of the situation, achieved by communication or by pre-established protocols, is unrealistic or likely to lock the agents into suboptimal forms of behavior. Further, by being competent decision makers and able to rationally model the action of other agents, RMM agents can effectively interact with human-controlled units, in spite of a lack of a predetermined protocol that a human would have to learn and follow.

In the following sections, this paper explores the overall objective of the antiair defense scenario. To compute the expected utilities of alternative action, we elaborate on a set of attributes describing the relevant features of the problem domain. Then, we explain how RMM leads an agent to the subjective decision-theoretic optimality with a concrete example of a coordination in our domain, and we discuss the experimental results. In conclusion, we summarize our results and further research issues.

2 Antiair Defense Environment

Let us consider a situation depicted in Fig. 1. This scenario has six incoming missiles and two defending units in a 20 by 20 grid world. Each of the defense units independently decides to launch interceptors against the incoming missiles in the absence of communication. The incoming missiles keep going straight top-down and attack the overall ground site on which the units are located.

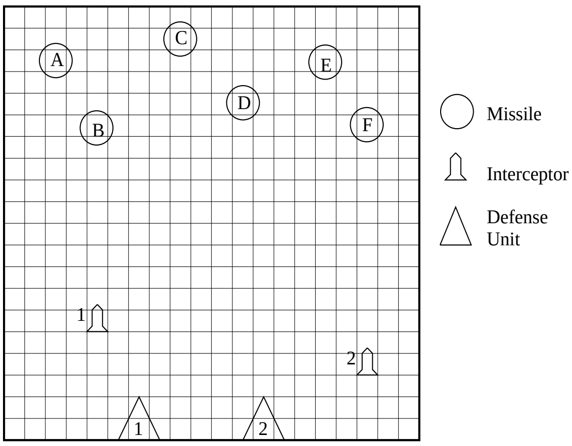

Fig. 1. The antiair defense scenario.

2.1 Attribute Analysis

One of the major problems encountered in decision making process is how to model preferences and utilities. For the purpose of formalizing the decision-making problem of minimizing damages, we first consider the attributes that influence the units' decision-making. Each missile has its intrinsic threat value. In our model, threat is evaluated by considering the altitude of a missile and its warhead size. Intuitively, a defense battery should give priority to the missile which is closer to the ground and bigger than the others. An explosion of a missile close to the ground results in a more powerful the blast, which increases the damages. Further, the measure of damage is proportional to the size of a missile's warhead. We calculate the missile threat, MT, using the following formula:

$$MT_n = W_n \times 1/A_n \tag{1}$$

where

• W_n: the warhead size of missile n

• A_n: the altitude of missile n

A factor that does not enter into the missile threat calculation, but one that should be considered when an intercept decision is made, is the probability that an interceptor would be successful against a given missile. This probability is assumed to depend on the angle between a missile's direction of motion and the battery's line-of-sight. This probability is maximized when this angle is 0, as follows:

$$P(HIT_{ij}) = e^{-\mu\gamma_{ij}} \tag{2}$$

where
• γ_{ij}: the angle between battery i's line-of-sight and missile j's direction of motion. Thus, $\gamma_{ij} = tan^{-1}\alpha - tan^{-1}\beta$ such that $0 \leq \gamma_{ij} \leq 90$
• α: the slope of the missile j's direction of motion
• β: the slope of the line-of-sight with which battery i aims at missile j
• μ: an interceptor-specific constant (assumed here as 0.01)

We will use the values of missile threat and the hit probability to develop the decision-theoretic assessment of the agent's alternative plans of action in an antiair defense environment.

2.2 Planning and Execution Cycle

In the domain in Fig. 1, the defense units, {Battery1, Battery2}, is a set of planning and executing agents. The targets, {MissileA, MissileB, MissileC, MissileD, MissileE, MissileF}, have no plans of their own, and are assumed not to make any decisions.

The ordered actions (Scan-Area, Select-Target, Launch-Interceptor) available to agents are repeatedly used to achieve the subgoal of intercepting one of the attacking missiles during the overall plan-action cycle. As we mentioned, there is no notion of a symbolic goal in this planning. Instead, the goal of minimizing the damages is represented as a quality measure assigned to plans, which is then used to coordinate plans among multiple agents. In this case study, our work addresses the rationally coordinated target selection, i.e., the Select-Target step, providing the best defense strategy that can be implemented by independent defense units.

3 Decision-Theoretic Agent

To be rational in decision-theoretic sense, the agents follow the principle of maximum expected utility (PMEU) [10]. In this section, we will show how PMEU can be implemented in this case study using the Recursive Modeling Method (RMM). RMM [2, 3] will be used to model the other agent, and to select the most appropriate missile to intercept by a given defense battery.

3.1 An Example Scenario

Our approach is to take the agent-oriented perspective. In the examples scenario (Fig. 1), we view the decision-making through the eyes of an individual defense unit, Battery1, and his radar-acquired data.[1] Fig. 2 depicts the information acquired in the example scenario by Battery1 for the missiles A through F. In Fig. 1, the left top corner of the screen is (0,0), x is pointing right, and y is pointing down. Applying formula (1) and (2) to the acquired data, Battery1 can compute the relevant attributes of altitude, warhead size, and the angle γ that determines the hit probability. Battery1 also generates the expected hit probabilities for Battery2, assuming his hypothetical intercepting actions. The results are summarized in Fig. 2.

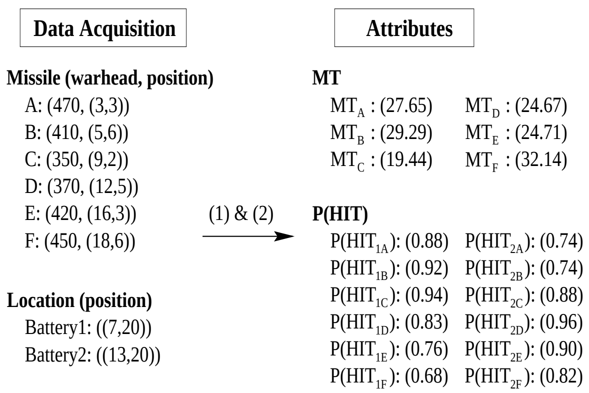

Fig. 2. Radar data acquired by Battery1 for the missiles A through F.

3.2 Generation of the Payoff Matrix

Commonly, a decision problem is represented by a version of belief network [6]. Poh and Horvitz [7] use an influence diagram that includes a decision variable, with values ranging over the possible decisions, chance variables, that represent the uncertainty of the domain, and a utility node. In our work, we rely on the payoff matrix representation, used in game theory. Payoff matrices, while different from belief nets, can be seen to faithfully summarize the information contained in belief nets by listing the expected payoffs (obtained from the utility node) of possible decisions, depending on the parameters describing the domain (chance nodes).

The expected payoffs, corresponding to batteries' attempting to intercept the respective missiles, can be expressed as a combination of the threat of the

[1] Battery2 acquires the information about environment from his point of view. Battery1 and Battery2 independently maintain their knowledge bases.

missiles and the probability of their interception. For example, if the Battery1 is faced with n missiles at some state, and he targets a missile j, the resulting threat will be reduced by the missile threat MT_j multiplied by the probability of successful interception $P(HIT_{1j})$. If both batteries target missiles at the same time, the reduction of threat, and therefore the total payoff, is equal to the sum of the threats that each of them removes.

3.3 Modeling Other Agent – Recursive Model Structure

In order to solve his decision-making situation, described by the payoff matrix above, Battery1 needs to hypothesize the likely actions of the Battery2. In the Recursive Modeling Method, the actions of the other rational agents are anticipated using a model of their decision-making situation. If it were known that the Battery2 is a rational decision-maker, then Battery2's decision-making situation could be modeled as a payoff matrix as well. In our case study, we considered a more realistic situation in which it is not known for sure that the Battery2 is rationally maximizing the payoff as well. For example, it could be that Battery2 has been damaged or otherwise incapacitated, in which case there is no information as to what action he would undertake. Thus, there are two alternative models that Battery1 can use to model Battery2; one has the form of the payoff matrix, and the other one contains no information about Battery2's action. We call the latter model the *No-info* model. In RMM, each of the alternative models is assigned a probability indicating the likelihood of its correctness.

Further, in case the Battery2 is not incapacitated, it is likely that Battery2 is modeling Battery1 as well, in his own attempt to coordinate with Battery1. This leads to the nesting of models in RMM. Since, in our scenario, Battery2 also may be uncertain whether Battery1 is not damaged, there are two alternative models on this level of modeling as well. One of the models Battery2 can have of Battery1 is that of a rational maximizer, while the other one is unknown and again labeled as a *No-info* model.

The resulting hierarchy of models, which we call the recursive model structure, terminates with a *No-info* model when the agent (in this case Battery2) runs out of modeling information. Fig. 3 is the Battery1's model structure of depth three for the example scenario in Fig. 2.

To summarize, level 1 in the recursive model structure represents the way that Battery1 observes the situation to make his own decision, shown as Battery1's payoff matrix. Level 2 depicts the models Battery1 has of Battery2's situation, and level 3 contains the models that Battery1 anticipates Battery2 may be using to model Battery1. The recursive modeling could continue into deeper levels, but in this case we assumed that the batteries have no further information. In other words, we are examining the reasoning of Battery1 in the particular case, when equipped with a finite amount of information about Battery2, and nested to the third level of modeling.[2]

[2] We do not elaborate on the important issues of learning and belief revision here. Thus, we analyze decision-making given a pre-existing state of knowledge, but we do

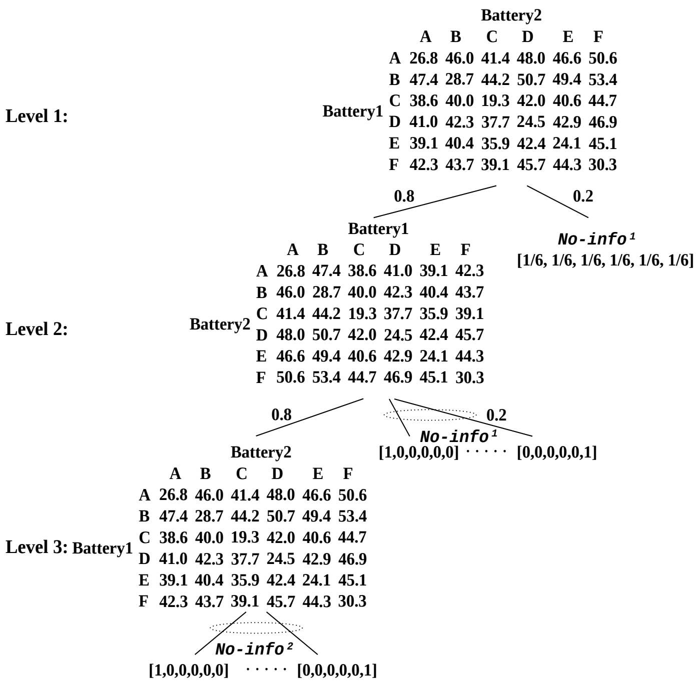

Fig. 3. The recursive model structure for Battery1.

In general, the *No-info* models have to differentiate among knowledge limitations of the different agents involved. For example, the situation in which Battery1 does not have any information about how he is modeled by Battery2 is different from the situation in which Battery1 knows that Battery2 has no information about Battery1. *No-info*^{1} on level 2 in Fig. 3 represents the fact that Battery1 has no information on how Battery2 models Battery1's actions if Battery1 is incapacitated. *No-info*^{2} on level 3 represents Battery1's belief that Battery2 has no information about how Battery1 models Battery2's actions, if Battery1 is not incapacitated.

Since the modeling structures, such as one in Fig. 3, express the optimal choices in the agents' decision-making situations recursively depending on the choices of the other agents, one can use dynamic programming to solve them in a bottom-up fashion. We elaborate on the solution procedure in more detail in [2], in which we have also outlined an analytical method of dealing with the *No-info* models. For the purpose of this case study, however, we have designed and implemented a numerical method of solving the *No-info* models using logical sampling [10], according to the algorithm below.

Let $A_i = \{a_i^1, a_i^2, ..., a_i^n\}$ be the set of actions available to agent i, and $P_i =$

not detail how the models of the other agents were obtained. For our current work on learning and model formation see [4].

$(p_i^1, p_i^2, ..., p_i^n)$ be the probability distribution over the available actions, i.e., the conjecture as to agent i's action. The conjecture, P_i, is calculated based on frequencies, $F_i = (f_i^1, f_i^2, ..., f_i^n)$, with which each action turns out to be optimal, as follows:

Procedure No-info-PDF
Input the payoff matrix M_i of agent i.
Output the probability distribution, P_i, over the actions of agent i.
begin
  $F_i \leftarrow (0, 0, ..., 0)$
  $N \leftarrow 0$ // the total number of probability distributions.
  for each probability distribution, P_{NI}, in the set of sampled probability distributions comprising the *No-Info* model,
    Multiply a probability distribution P_{NI} by M_i.
    Select an action a_i^k which has the maximum expected utility.
    $f_i^k \leftarrow f_i^k + 1$
    $N \leftarrow N + 1$
  end for
  $P_i \leftarrow (f_i^1/N, f_i^2/N, ..., f_i^n/N)$
  return P_i
end

Using the sampling algorithm, with the sampling density of 0.1, on the third level of the recursive model structure, reveals that the probability distribution over Battery1's actions becomes [0.24, 0.71, 0.0, 0.01, 0.0, 0.04] for interception of missiles A through F, respectively. This probability distribution summarizes Battery1's knowledge as to Battery2's expectations, if Battery2 is not incapacitated, about the actions of Battery1, if Battery2 thinks that Battery1 is not incapacitated.

The dynamic programming bottom-up solution then proceeds to level 2. The above distribution over the actions of Battery1 is combined with each individual probability distribution of $No{-}info^1$ on level 2, after each distribution is multiplied by weight 0.8 and 0.2, respectively. The sampling No-info-PDF procedure is invoked again, resulting in [0.0, 0.0, 0.0, 0.03, 0.0, 0.97] over Battery2's actions of intercepting the missiles A through F, respectively. Thus, in spite of uncertainties on the third level of modeling, Battery1's knowledge is enough to expect that, if Battery2 is rational, it will most likely attempt to intercept MissileF.

Propagating these results into the level 1, the probability distribution describing Battery2's actions is obtained as a combination (0.8 × [0.0, 0.0, 0.0, 0.03, 0.0, 0.97] + 0.2 × [1/6, 1/6, 1/6, 1/6, 1/6, 1/6]), which results in [0.03, 0.03, 0.03, 0.06, 0.03, 0.82]. Now, we can compute the expected utilities of Battery1's alternative behaviors as follows:

$U_A : 0.03 \times 26.8 + 0.03 \times 46.0 + 0.03 \times 41.4 + 0.06 \times 48.0 + 0.03 \times 46.6 + 0.82 \times 50.6 = 49.20$
$U_B : 0.03 \times 47.4 + 0.03 \times 28.7 + 0.03 \times 44.2 + 0.06 \times 50.7 + 0.03 \times 49.4 + 0.82 \times 53.4 = 51.92$
$U_C : 0.03 \times 38.6 + 0.03 \times 40.0 + 0.03 \times 19.3 + 0.06 \times 42.0 + 0.03 \times 40.6 + 0.82 \times 44.7 = 43.33$
$U_D : 0.03 \times 41.0 + 0.03 \times 42.3 + 0.03 \times 37.7 + 0.06 \times 24.5 + 0.03 \times 42.9 + 0.82 \times 46.9 = 44.85$

$U_E : 0.03 \times 39.1 + 0.03 \times 40.4 + 0.03 \times 35.9 + 0.06 \times 42.4 + 0.03 \times 24.1 + 0.82 \times 45.1 = 43.71$
$U_F : 0.03 \times 42.3 + 0.03 \times 43.7 + 0.03 \times 39.1 + 0.06 \times 45.7 + 0.03 \times 44.3 + 0.82 \times 30.3 = 32.67$

Thus, if Battery1 is rational, he will attempt to maximize his own expected utility and prefer to intercept MissileB first, given that he expects it most likely that Battery2 will select MissileF simultaneously. According to the planning and execution cycle, the interceptor will be launched toward the selected missile, and the decision-making process will be repeated until there are no missile threats, or until Battery1 runs out of interceptors.

4 Experiments

The antiair defense simulator is written in Common LISP and built on top of MICE simulator [1], running on a LINUX platform. In the experiments we ran, each of two defense units could launch three interceptors, and were faced with an attack by six incoming missiles. We put all of the trials under the following conditions. First, the initial positions of missiles were randomly generated and it was assumed that each missile must occupy a distinct position. Second, the performance assessments of agents with different policies were compared using the same threat situation. Third, the size of each missile was constant during all of experimentations and the warhead sizes were comparable. In this experiment, the warhead sizes were 470, 410, 350, 370, 420, 450 for missiles A through F, respectively. Further, each interceptor could intercept only one missile, and it was moving twice as fast as the incoming missile. Finally, although there was no communication between agents, each agent could see which threat was shot at by the other agent, and then used this information to make the next decision.

Our experiments were aimed at determining the quality of modeling and coordination achieved by the RMM agents in a team, and when paired with human agents, when compared to other strategies. To evaluate the quality of the agents' performance, the results were expressed as (1) the number of 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 $(1 - P(HIT)) \times Warhead\_Size$ to the total damage. If a missile was not targeted, it contributes all of its warhead size value to the damage.

The target selection strategies are as follows:

- Random: selection randomly generated.
- Independent, no modeling: selection of $arg\ max_j\{P(HIT_{ij}) \times MT_j\}$ for agent i.
- Human[3] : selection by human.

[3] We should remark that our human subjects were simply CSE and EE graduate students who were informed about the criteria for target selection. We would expect that antiair specialists, equipped with a modern defense doctrine, could perform substantially better than our subjects. However, the defense doctrine remains classified, and was not available to us at this point.

• RMM: selection by RMM.

The random agents were included to provide the worst base-line case of the system performance in our experiments. The performance achieved by independent agents was included to show what coordination can be expected when agents maximize but do not model each other. It turned out that the ways that human agents choose a missile were different and sometimes quite arbitrary. For example some of our human subjects shot the only 3 missiles coming to their own side by dividing the grid world into left and right sides. They sometimes had difficulties in splitting the screen when the missiles were clustered at the center area, which led to much duplicated effort. Others tended to choose missiles with the largest missile size. Still others tried to consider the multiplication of the missile size and the hit probability, but did not model the other agent appropriately.

4.1 Performance assessments

We experimented with the above policies to understand the agent interactions in two groups: the agent teams with the same policy and the mixed agent teams among different policies.

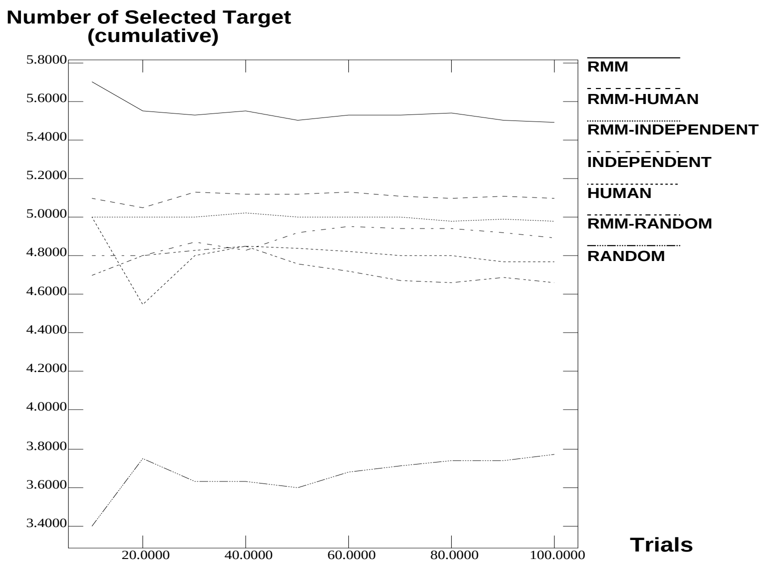

Fig. 4. The number of selected targets (averages over 10 runs).

As shown by Fig. 4 and 5, we found that the all-RMM team outperformed the human and independent teams. The average number of selected targets by RMM after 100 trials is 5.5, compared to 4.8 for the independent team, and 4.7 for the all-human team. Further, the RMM-controlled coordinated defense resulted in the total expected damage of 488, which is nearly half that of independent team (732) and that of the all-human team (772).

The human performance is very similar to the performance of independent agents. The most obvious reason for this is that humans tend to depend on their intuitive strategies for coordination, and, in this case, found it hard to engage in deeper, normative, decision-theoretic reasoning. The common reason for the lower score of human teams were choices of the greatest threat simultaneously – they made redundant efforts to attack the same missile. This, again, suggests that human agents attempted to minimize their expected damage but did not model the decision-making of the other human very well, while RMM agents were rationally optimizing given what they expected of the other agent. The performance of the RMM team is not perfect, however, since they were equipped with limited and uncertain knowledge of each other.

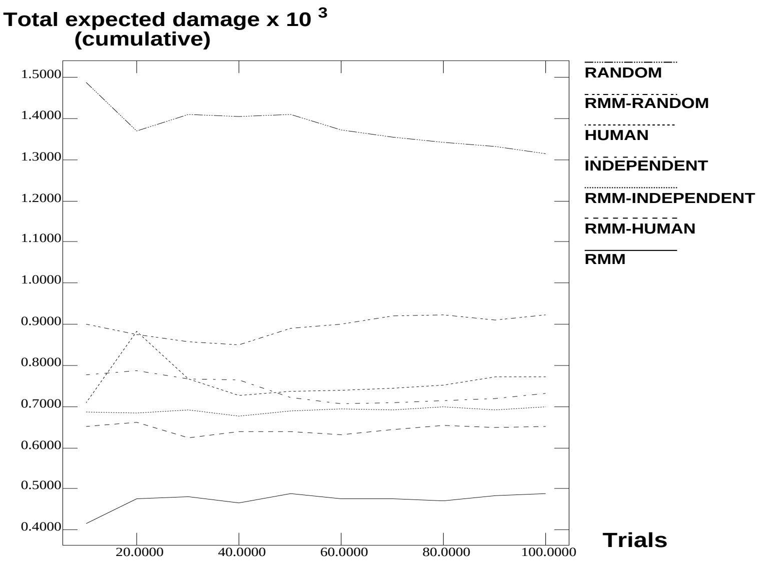

Fig. 5. The total expected damage (averages over 10 runs).

The performance of the heterogenous teams again suggests the favorable quality of coordination of RMM agents. Comparing a heterogenous team with a homogeneous team, the average number of selected targets for the RMM-Human team is 5.1, and for the all-human team it is 4.7; that of the RMM-Independent team is 4.98 vs. 4.8 for the Independent team; that of the RMM-Random team is 4.66, and 3.77 for the all-Random team. As shown by the above comparison, the recursive modeling allows the RMM agent to improve the performance of target selection when paired with any other agent. A particularly promising facet of these results is that the recursive model structure (Fig. 3.) seems to be a robust mechanism for modeling and coordination with the human subjects, and not only among RMM agents. (For the demonstration of the antiair defense domain, refer to the Web page at $http://dali.uta.edu/Air.html$.)

5 Related Work

The approaches related to the modeling and coordination in the multiagent environment are in the area of multiagent plan recognition and plan coordination. In work on plan recognition [9] and [12], the objective is to enable an agent to model or recognize the other agents through observation to anticipate the other agents' future action, given prior knowledge about these agents. In these approaches, an agent usually compares a pre-calculated plan, or a protocol, with the on-going situation, and then chooses his next action accordingly. The process of mental-state recognition [9] assumes the correct and complete knowledge of the plans of the other agents, and it does not represent uncertainty that might be present in real-world domains. As Tambe et al. [12] pointed out, ambiguities may persist when an agent must infer unobserved actions and intentions.

The multiagent coordination without communication has been dealt with in [11] and [5]. Sen, in [11], used a particular reinforcement learning methodology and concentrated on the learning classifier systems where agents share no problem-solving knowledge. However, the reinforcement learning method needs a super-agent for the external reward and requires numerous training cases to converge on the desirable entries. Mor et al. [5] handled the opponent's modeling in the complexity of polynomial time, based on a game theoretic techniques. They represented a strategy of using a deterministic finite automaton and the learning process as a means to reach equilibrium.

6 Conclusion and Further Research

This paper presented a case study in modeling and coordination in a multiagent distributed environment of the antiair defense. This investigation implies that RMM can be applied to high-level tactical decision making for minimizing the overall damages. It turned out that it was relatively simple to quantify the missile threats and the hit probability to generate fairly realistic payoff matrices. Based on threat evaluation, the RMM agents we implemented to model the other agents and to follow the principle of maximum expected utility. We anticipate that the approach presented will also be applicable to other decision problems that require modeling and coordination.

The most promising conclusion of our case study seems to be the high quality of the coordinated decision-making achieved by the RMM agents. This, in turn, suggests that modeling other agents, including humans, as rational decision-makers is a promising and viable approach to achieving plan recognition and coordination. Also, the relatively high quality achieved by the mixed human/RMM team shows how a nested decision-theoretic modeling can enable human-machine coordination without pre-established protocols that the human would have to learn and follow during the interaction.

In other experiments, we also implemented RMM in the predator-prey pursuit environment. There, the payoff matrices the agents used are 5-dimensional. The quality of coordination among the RMM agents for this domain is slightly inferior to that of human agents ($http://dali.uta.edu/Pursuit.html$). Further, we applied Bayesian learning for belief update to the antiair defense domain, as well as the pursuit domain. In our future research we will investigate the communication method between agents using KQML and probe the value of decision-model refinement [8].

References

1. E. H. Durfee and T. A. Montgomery. MICE: A flexible testbed for intelligent coordination experiments. In *Proceedings of the 1989 Distributed AI Workshop*, pages 25–40, Sept. 1989.
2. P. J. Gmytrasiewicz. On reasoning about other agents. In *Intelligent Agents II: Agent Theories, Architectures, and Languages*, pages 143–155. Springer, 1996.
3. P. J. Gmytrasiewicz and E. H. Durfee. A rigorous, operational formalization of recursive modeling. In *Proceedings of the First International Conference on Multi-Agent Systems*, pages 125–132. AAAI Press/The MIT Press, 1995.
4. T. Kellogg and P. J. Gmytrasiewicz. Bayesian belief update in multi-agent systems. *In preparation*, 1996.
5. Y. Mor, C. V. Goldman, and J. S. Rosenschein. Learn your opponent's strategy (in polynomial time)! In G. WeiB and S. Sen, editors, *Adaptation and Learning in Multi-Agent Systems - IJCAI'95 Workshop, Lecture Notes in Artificial Intelligence*, pages 164–176. Springer, New York, 1996.
6. J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufman, 1988.
7. K. L. Poh and E. J. Horvitz. A graph-theoretic analysis of information value. In *Proceedings of the Twelfth Conference on Uncertainty in Artificial Intelligence*, 1996.
8. K. L. Poh and E. J. Horvitz. Reasoning about the value of decision-model refinement: Methods and application. In *Proceedings of the Ninth Conference on Uncertainty in Artificial Intelligence*, pages 174–182. Morgan Kaufmann Publishers, Inc., July 1993.
9. A. S. Rao and G. Murray. Multi-agent mental-state recognition and its application to air-combat modelling. In *Proceedings of 13th International Distributed Artificial Intelligence Workshop*, pages 283–304, 1994.
10. S. J. Russell and P. Norvig. *Artificial Intelligence A Modern Approach.* Prentice-Hall, Englewood Cliffs, New Jersey, 1995.
11. S. Sen and M. Sekaran. Multiagent coordination with learning classifier systems. In G. WeiB and S. Sen, editors, *Adaptation and Learning in Multi-Agent Systems - IJCAI'95 Workshop, Lecture Notes in Artificial Intelligence*, pages 218–233. Springer, New York, 1996.
12. M. Tambe and P. S. Rosenbloom. Architectures for agents that track other agents in multi-agent worlds. In *Intelligent Agents II: Agent Theories, Architectures, and Languages*, pages 156–170. Springer, 1996.
13. M. P. Wellman and J. Doyle. Preferential semantics for goals. In *Proceedings of the National Conference on Artificial Intelligence*, pages 698–703, July 1991.