Rational Communicative Behavior in Anti-Air Defense*

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

Our research addresses rational decision-making among anti-air units whose mission is to defend a specified territory from a number of attacking missiles. In this paper, we present a decision-theoretic method the automated agents can use to select optimal communicative acts in this domain, given the characteristics of the threat situation. Since the communication bandwidth is usually limited in a battlefield environment and disclosure of any information to hostile agents should be avoided, it is critical for a defending agent to be selective as to what messages should, or should not, be sent to other friendly agents. This work builds on our earlier work that uses the Recursive Modeling Method (RMM) for coordination, and applies RMM to rational communication in the anti-air defense domain. We show how a defending automated agent can rank the values of alternative messages that could be communicated in a given defense situation. Further, we compare this ranking to communicative choices of human subjects under the same circumstances of the coordinated anti-air defense.

Content Areas: Optimal Communication, Decision Theory, Communication and Coordination, Coordinated Anti-Air Defense.

Word Count: 4103

*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.

1 Introduction

This paper describes the *decision-theoretic message selection* of an automated agent engaged in an anti-air defense with other defense agents. The goal of defending agents is to minimize damages to their territory [10]. 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, and disclosure of any information to hostile agents should be avoided, it is critical for a defending agent to be selective as to what messages should be sent to other agents.

In our work, the automated agents make the communicative decisions so as to maximize the expected utility. Thus, they compute the expected utility of alternative communicative behaviors, and execute the one with the highest value. Following the principle of maximizing the expected utility [13], our agents are intended to be rational in their communicative behavior; they will send only the most valuable messages, never send messages that are damaging in any way, and be able to determine under what circumstances they should abstain from communicating all together.

To achieve the *coordination by communication* between agents, most of the previous approaches rely on predetermined protocols [12, 1]. Quintero [12] specified a multi-agent protocol language which can be used to model a cooperative environment. Their rule-based agents (called LALO agents) execute rules if messages are matched to the rules in the message rule set. Barbuceanu, in [1], postulates that conversation plans and rules be available to the agents for performing the coordinated behavior. In their method, rewards denote the utility of finding the optimal action given a state, and are assumed to be known before run time.

Tambe [15] suggested decision theoretic communication selectivity to establish mutual belief among agents in a team. This approach is similar to ours in that the focus is on whether or not an agent should transmit a given message (fact) to others. Their approach to selective communication uses a decision tree containing reward and penalty values that are domain knowledge. However, obtaining these values when the environment is dynamic, not fully observable or when the status of another agent is not fully predictable, may be problematic. Our framework, in contrast, does not rely on pre-determined rewards and penalties, but computes them based on the, possibly incomplete and uncertain, models of the other agents. Tambe's work is further closely related to prior work by Cohen [2] and Grosz [8].

As we mentioned, we investigate decision-theoretic message selection using Recursive Modeling Method (RMM) [6, 4]. RMM endows an agent with a representation of the other agents' beliefs, abilities, and intentions. As such, it allows the agent to predict the message's pragmatics, i.e., how a particular message will change the state of knowledge of the other agent(s), and how are the other agents likely to react to it. We propose that modeling the other agents while communicating is crucial; clearly without a model of the other agents' state of beliefs it would be impossible to properly assess the impact of the communicative act. Based on the message's pragmatics, our method quantifies the gain obtained due to the message as the increase in the expected utility obtained as the result of the interaction.

This paper applies our RMM-based approach to a specific environment of the anti-air defense. We will show how the expected utility of various messages is computed in four different example settings. Further, we will compare our decision-theoretic results with communicative decisions made by human subjects in the same settings. It turns out that the raking of messages obtained by our method closely mirrors the choices of the human subjects, thus demonstrating the psychological plausibility of our approach.

2 Decision-Theoretic Communication

We identify a communicative act with its *decision-theoretic (DT) pragmatics*, defined as the transformation of the state of knowledge about the decision-making situation the act brings about. We model DT pragmatics using the RMM representation to investigate the utility of the communicative act [6, 5].[1] The transformation in the agent's decision-making situation, as represented by RMM's recursive model structure, may change the expected utilities of alternative actions. It is natural to identify the change of the expected utility brought about by a communicative action as the expected utility of this action itself. By evaluating the alternative communicative acts in this way, the agent can select and send the highest utility message—the message that causes the greatest gain in the expected utility of the agent's action.[2]

Formally, the expected utility of the communicative act M, executed by an agent, is the difference between the payoff the agent expects before and after executing the act:

$$U(M) = U_{p^M}(Y) - U_p(X). \tag{1}$$

where $U_p(X)$ is the utility of the best action, X, expected before sending the message, and $U_{p^M}(Y)$ is the utility of the best action, Y, expected if the message were to be sent. Further details and discussion is contained in [6, 5]. We now go on to describing our application domain.

3 The Anti-Air Defense Domain

Our model the anti-air domain consists of a number of attacking missiles, in the simplest case two as depicted in Figure 1(a), and a number of defending units, also two on the bottom of

[1]Since our formalism is probabilistic, it naturally handles cases when the meaning of a message is itself uncertain.

[2]The notion of the utility of a message we use here differs from the notion of the value of information considered in decision theory [11, 13]. The latter expresses the value of information to its recipient. We, on the other hand, consider the value of a message to its sender, since, of course, it is the sender that makes the decision of if, and what, to communicate. The two notions coincide in two special cases, when the preferences of the agents perfectly coincide, and when a speaker requests information from the hearer by asking a question.

Figure 1(a). In this paper we assume that, for simplicity, each of the batteries has only one interceptor, but see [10] for more complex examples involving multiple salvos and more attacking missiles. As we mentioned, the mission of the defense units is to attempt to intercept the attacking missiles so as to minimize damages to own territory. Let us note that coordination is necessary; if the defense batteries miscoordinate and attempt to intercept the same threat the other threat will reach its destination and cause damage proportional to the warhead size.

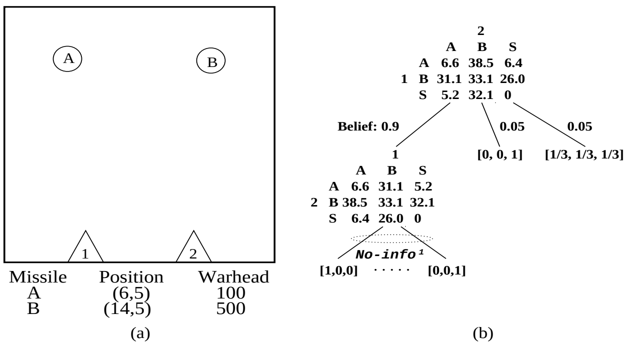

Figure 1: Anti-air defense (a) an example scenario (b) a recursive model structure.

To analyze the defense situation in decision-theoretic terms we form attributes quantifying the quality of alternative actions. First, each attacking missile presents a threat of a certain value. We compute the missile's threat by dividing the size of its warhead by the altitude of the missile. Thus, the defense units should prefer to intercept threats that are larger and closer to the defended territory. Further, the defense units should consider the probability with which their interceptors would be effective against each of the hostile missiles. The interception probability, $P(H_{ij})$, is dependent on the angle, γ, between the missile j's direction of motion and the line of sight between the battery i and the missile, as follows:

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

where μ is an interceptor-specific constant (assumed here to be 0.01).

Given these factors, the expected benefit of intercepting threats can be quantified as a product of the sizes of the threats and the hit probabilities. For example, in Figure 1(a), the combined benefit of Battery1's shooting at the threat A and Battery2's shooting at the threat B amounts to 38.5. This value is entered in the payoff matrix, such as one on top in Figure 1(b). In this payoff matrix the rows represent the alternative actions of Battery1 of shooting at A, B, or not shooting, respectively, and the columns represent the alternative actions of Battery2.

The top payoff matrix in Figure 1(b) represents the decision-making situation that Battery1 is facing. To solve it, Battery1 needs to predict what Battery2 is likely to do; for example if

Battery2 were to shoot at A it would be best for Battery1 to shoot at B, but if Battery2 is expected to shoot at B then Battery1 should definitely target A.

The actual behavior of Battery2, however, depends on a number of factors that Battery1 may be uncertain about. For example, if Battery2 has been hit and incapacitated it will not be able to launch any interceptors. If it is not incapacitated then its own decision-making situation can be represented as another payoff matrix. These two models, of Battery2 being operational, and incapacitated, are depicted as the second level models in Figure 1(b), with their associated probabilities, respectively. The third model, assigning equal chance to every action of Battery2, is also included to represent the case in which nothing is known about Battery2's behavior. This is called a No-Information model. The probabilities associated with the alternative models can be arrived at using Bayesian learning, as we describe in [7].

Under the assumption that the nesting of models is finite, as in the structure in Figure 1(b), Battery1's decision-making problem can be solved using dynamic programming. The solution proceeds bottom-up and results in expected utility values being assigned to Battery1's alternative behaviors, best on which can then be executed. For a more detailed discussion, see [9].

In considering the set of messages available to an automated agent, we classify the messages into two groups: intentional messages and modeling messages [5]. Intentional messages describe the intended actions of the agent, and modeling messages describe the environment or the status of agent relevant to the achievement of their goal. Both types of messages are intended to remove some of the uncertainty present in the absence of communication. The agents are allowed to use two primitives: *Attempt* and *Assert* [14]. When they want to send intentional messages, they use *Attempt*. In case of modeling messages, they use *Assert*. For simplicity we assume that the agents can mutually understand communication primitives [1].

We will consider the optimal communicative behavior of the Battery1 only, and, for simplicity, assume that Battery2 is silent but can receive messages. In all of the anti-air defense scenarios considered below Battery1 have a choice of six communicative behaviors:

$No\ Comm.$: No communication
$M1$: I'll intercept Missile A.
$M2$: I'll intercept Missile B.
$M3$: I have both long and short range interceptors.
$M4$: There is a missile A, whose position and warhead size are P_A and W_A, respectively.
$M5$: There is a missile B, whose position and warhead size are P_B and W_B, respectively.

In our implementation, the anti-air defense simulator with communication was written in Common LISP and built on top of the MICE simulator [3]. Our experiment was to investigate how RMM agents rank the messages among a list of them above, and whether there is an agreement between the communicative behavior advocated by RMM and human communicative behavior. As human subjects we used 32 CSE and EE graduate students. Each of them was presented with the scenario, and was given a description of that was known and what was uncertain in each case. The students were then asked to indicate which of the six messages was the most appropriate in each case, and which one was the second choice. The results are summarized next to each of the scenarios described below.

4 Air Defense Scenarios

In this section we consider four different anti-air defense scenarios. In each case we present the decision-theoretic evaluation of the alternative messages performed by our agents, and compare results to communicative behaviors chosen by the human subjects.

4.1 Scenario 1

Let us consider the scenario depicted on the left in Figure 2. Here, the defense batteries face an attack by missiles A and B, with the size of A larger than that of B, and the altitude of missile B is lower than that of missile A. The state of Battery1's state of knowledge before communication is summarized as a two-level recursive model structure on the left in Figure 3. Battery1 assigns the probability of 0.9 to Battery2's being fully operational and having both long and short range interceptors, probability 0.05 to Battery2's being incapacitated. The remaining probability of 0.05 is assigned to the No-Information model representing all of the possible unknown cases. In this scenario Battery1 is assumed to have no more information. In particular, Battery1 does not know what action Battery2 expects of Battery1. This is represented by another No-Information model below in Figure 2, as discussed at length in [5].

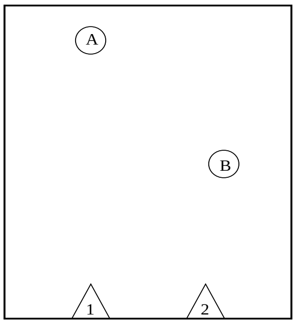

Missile	Position	Warhead
A	(7,3)	500
B	(14,10)	100

Messages	RMM U(M)	Human 1st	Human 2nd
No Communication	0	4	8
M1: I'll shoot down A.	7.31	24	5
M2: I'll shoot down B.	-1.26	3	10
M3: I have both armors.	0	0	3
M4: A - (7,3), 500	0	1	1
M5: B - (14,10), 100	0	0	5

Figure 2: Scenario 1, and Summary of Results.

Figure 3 depicts the change of the decision-making situation expected due to Battery1's sending a message $M1$, "I will intercept missile A." As we mentioned this change is the $M1$'s decision-theoretic (DT) pragmatics. It illustrates that, as a result, Battery1 expects Battery2 to know that Battery1 will select missile A, which is represented as a [1, 0, 0] probability distribution Battery2 would use to describe Battery1's actions.

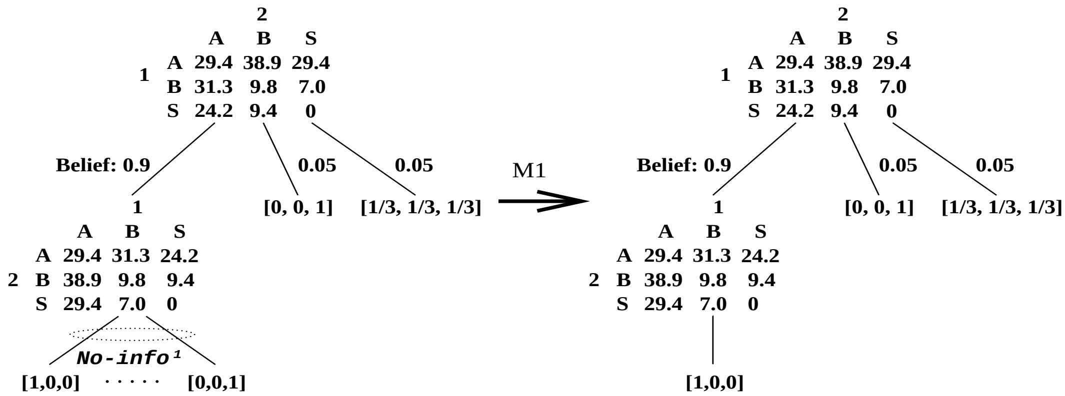

Figure 3: The DT Pragmatics of message $M1$ in Scenario 1.

To compute the value of communication according to Eq. 1, one has to solve both modeling structures in Figure 3 and compare results. Before communication (the left part of Figure 3), Battery1 computes that if Battery2 is operational then the probability distribution over Battery2's actions A, B, and S is $[0.85, 0.15, 0.0]$. This result is obtained using logic sampling, as described in [10]. Using dynamic programming one can now easily compute that Battery1's best option is to shoot at missile A, with the expected utility, $U_p(A)$, of 30.83 $(= 0.78 \times 29.4 + 0.15 \times 38.9 + 0.07 \times 29.4)$.

After sending the message $M1$ (the right part of Figure 3), the probability distribution over Battery1's actions at Level 2 will be [1,0,0]. Thus, if Battery2 is fully operational it will choose to shoot at missile B, i.e., the probability distribution over Battery2's actions becomes [0,1,0]. This probability distribution is combined with the model of Battery2 being incapacitated and with the third *No-info*[1] model: $(0.9 \times [0, 1, 0] + 0.05 \times [0, 0, 1] + 0.05 \times [1/3, 1/3, 1/3]) = [0.02, 0.92, 0.06]$. The resulting distribution is Battery1's overall expectation of Battery2's actions, given all of the remaining uncertainty. Propagating these results into Level 1, the combined probability distribution describing Battery2's actions is used to compute the expected utility of Battery1's action of shooting A. We have: $U_{p^{M1}}(A) = 0.02 \times 29.4 + 0.92 \times 38.9 + 0.06 \times 29.4 = 38.14$. According to the Equation 1, the expected utility of the communicative act $M1$, $U(M1)$, is $7.31 (= 38.14 - 30.83)$.

The expected utilities of the other messages are computed analogously, and the results are depicted in Figure 2. As expected, some of the messages have no value in this situation, and their computed expected utility is zero, since they do not convey anything useful and novel. Let us note that the message $M2$ is damaging, with a negative expected utility. Thus, it is a bad idea for Battery1 to commit itself to shooting at missile B in this scenario.

The results of human choices are also summarized in Figure 2. Twenty four, out of thirty two, subjects chose message $M1$ as the best in this situation, while five subjects judged it as a second best. We elaborate on the experimental results further later.

4.2 Scenario 2

In this scenario, in Figure 4, the warhead size of missile B is larger than that of missile A. However, due to bad weather (rain or cloud), Battery1 thinks that Battery2 is unlikely to be able to detect Missile B, and assigns to it the probability of only 0.1. Again, we allow for further uncertainty: Battery2 can be incapacitated by enemy fire with probability 0.05, and all of the other possibilities about which Battery1 has no information about Battery2 are grouped into a No-Information model with probability 0.05.

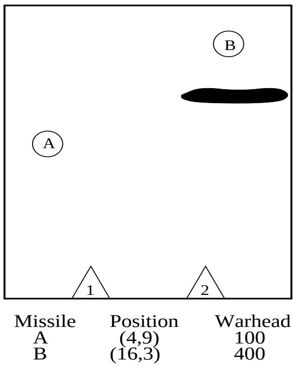

Messages	RMM U(M)	Human 1st	Human 2nd
No Communication	0	0	3
M1: I'll shoot down A.	0	1	7
M2: I'll shoot down B.	0.09	9	11
M3: I have both armors.	0	0	3
M4: A - (4,9), 100	0	0	1
M5: B - (16,3), 400	3.64	22	7

Figure 4: Scenario 2.

In this case, the most interesting is the message $M5$, "There is a missile B, whose position and warhead size are P_B and W_B, respectively." The result of this message would be to eliminate the possibility that Battery2 does not know about Missile B, and its DT pragmatics is depicted in Figure 5.

The recursive model structures in Figure 5 can be solved using dynamic programming and logic sampling as described above. Before communication, Battery1 would assign a distribution of $[0.82, 0.12, 0.06]$ to battery2's actions, and Battery1 could get the expected utility of 23.42. After sending the message $M5$, the probability that Battery2 will shoot at B increase to 0.9, and the new distribution is $[0.04, 0.90, 0.06]$, and the expected utility of Battery1, $U_{p^{M5}}(A)$, is: $U_{p^{M5}}(A) = 0.04 \times 8.7 + 0.90 \times 29.1 + 0.06 \times 7.8 = 27.06$. Therefore, the expected utility of the communicative act $M5$, $U(M5)$, is $3.64 (= 27.06 - 23.42)$.

The expected utilities of the other communicative alternatives are computed similarly and the results are in Figure 4. The results of experiments with human subjects is also summarized. This time, 22 out of 32 subjects chose message $M5$ as the best one. Interestingly, quite a number of subjects picked message $M2$ as the best, which the RMM agent rated as second best in this case.

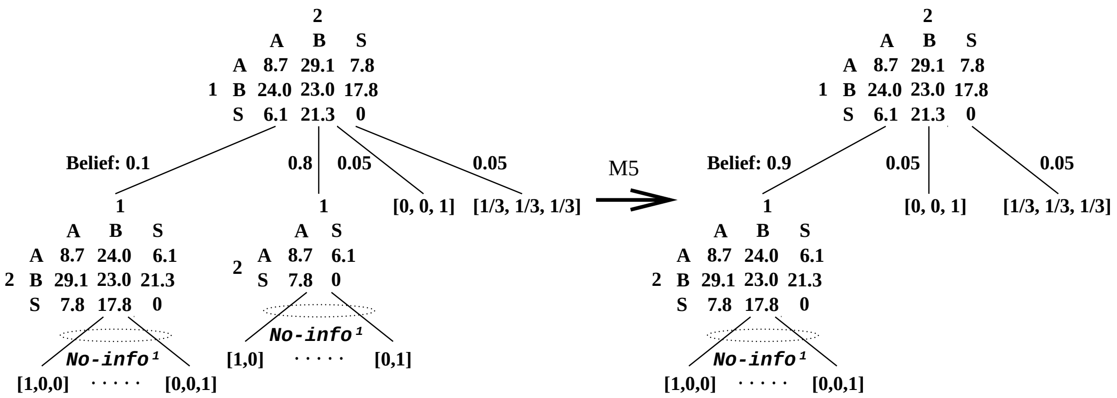

Figure 5: DT pragmatics of $M5$ in scenario 2.

4.3 Scenario 3

If we allow more uncertainty in anti-air defense, the decision-theoretic message selection will be more complicated. For scenario 3, missile A is coming toward Battery2 and missile B toward Battery1. The altitude of missile B is lower than that of missile A and the warhead sizes of missiles are comparable. (See Figure 6) In this case, we assumed more uncertainty: Battery1 is uncertain whether Battery2 has any short range interceptors left. If Battery2 has only long range interceptors, it will be unable to attack missile B, and can only attempt to shoot down missile A.

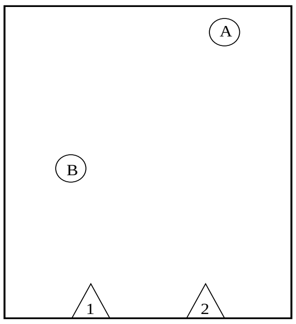

Missile	Position	Warhead
A	(14,2)	120
B	(6,10)	100

Messages	RMM U(M)	Human 1st	Human 2nd
No Communication	0	1	3
M1: I'll shoot down A.	-1.47	0	3
M2: I'll shoot down B.	2.63	27	3
M3: I have both armors.	2.02	3	19
M4: A - (14,2), 120	0	1	4
M5: B - (6,10), 100	0	0	0

Figure 6: Scenario 3.

From Battery1's point of view, therefore, Battery2's decision making situation could be modeled as one of four cases: Battery2 has both short and long range interceptors; Battery2 has only long range interceptors; Battery2 has been damaged or incapacitated; or there is no information as to what action Battery2 would undertake.

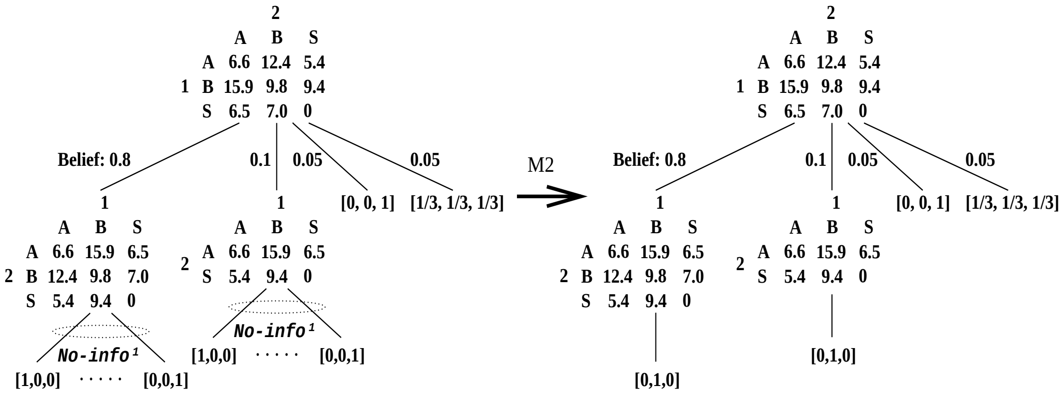

Figure 7: DT pragmatics of $M2$ in scenario 3.

If Battery1 broadcasts $M2$, "I will intercept missile B.", the resulting transformation is as depicted in Figure 7. However, another viable message is message $M3$, with the DT pragmatics as depicted in Figure 8. This transformation shows that if Battery1 sends the message "I have both long and short range interceptors.", Battery2 will include the fact that Battery1 has both ammunitions in its modeling of Battery1, and will solve these models to arrive at Battery1's target selection.

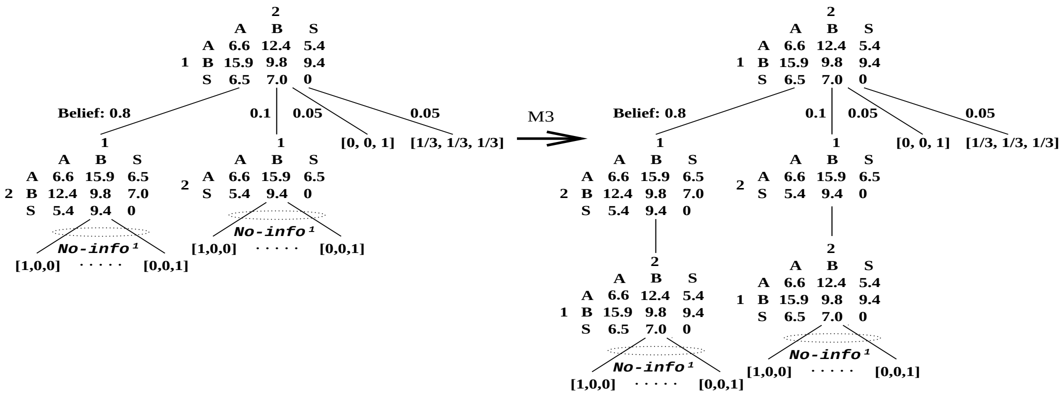

Figure 8: DT pragmatics of $M3$ in scenario 3.

Before communication, Battery1's best option is to shoot down missile B, with its expected utility of 12.76. If $M1$ were to be sent Battery1 could expect the utility of 11.29. After sending $M2$, Battery1 could expect 15.39 ($= 0.92 \times 15.9 + 0.02 \times 9.8 + 0.06 \times 9.4$), while $M3$, would result in the expected utility of $14.78 (= 0.82 \times 15.9 + 0.11 \times 9.8 + 0.07 \times 9.4)$.

These results, together with selections of the human subjects, are summarized in Figure 6. 27 out of our 32 subjects agreed with the RMM calculation and picked $M2$ as the best message in this case. Further, 19 subjects rated $M3$ as their second choice.

4.4 Scenario 4

Our fourth scenario was intended as a test of a case when communication is not necessary. In general, since the communication may be expensive, a rational agent should not communicate with others if there is no gain from communication. We set the warhead size of missile B as bigger than that of missile A. The altitudes of missiles are the same, as depicted in Figure 9. Also Battery1 assigned the probability 0.05 to Battery2 being incapacitated by enemy fire, with the No-Information model having the weight of another 0.05. Thus, Battery1 thinks that Battery2 has both interceptors and can see Missile A and B with probability 0.9.

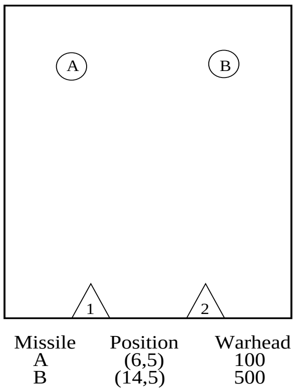

Messages	RMM U(M)	Human 1st	Human 2nd
No Communication	0	11	10
M1: I'll shoot down A.	0	19	9
M2: I'll shoot down B.	-3.31	2	5
M3: I have both armors.	0	0	1
M4: A - (6,5), 100	0	0	2
M5: B - (14,5), 500	0	0	5

Figure 9: Scenario 4.

For this scenario, the expected utilities of intentional messages are less than or equal to the expected utility of no communication, as depicted in Figure 9. The modeling messages are useless as well, since they do not change the decision-making situation in any beneficial way. The results of experimens with the human subjects suggests that the majority of them were indifferent between communicating and not communicating in this case. There is, however, a slight preference humans seem to exhibit for communicating useless messages. This can be accounted for by two factors; first it was probably difficult for our subjects to compute that the values of messages that, under usual conditions make perfect sense, are useless in this particular situation. Second, human communicative behavior may be driven not strictly by the values of messages we have been computing here. People may want to talk just to stay in touch and maintain the social bond (babble), even when the situation does not demand it.

5 Conclusions and Future Work

This paper presented work that applied decision-theoretic message selection to a domain of coordination in anti-air defense domain. We chose four different scenarios, and our decision-theoretic computation showed that each demands a different communicative behavior. In some cases the intentional messages, that obligate the agent and committ it to some course of action, are the most appropriate. Sometimes, however, the modelling messages, intended to share information about the environment ot the agent's state, are the best. In some other cases, either messages could be used, while sometimes communication has no immediate value at all.

Our experiments with human subjects were intended to compare our normative decision-theoretic calculations with the communicative behaviors that humans consider most appropriate under various circumstances. We found that in cases when the calculations show one message that is clearly the best among all of the others, our human subjects were more likely to choose this message themselves. Thus, in scenarios 1, 2, and 3, we had 75%, 69%, and 84% agreement with the human selection. In the scenario 3, in which the next best message had the value close to that of the optimal message, 60% of the subjects picked it as their second choice. In the fourth scenario, the results of human selection are inconclusive, which mirrors the result of our calculation that placed the best message as equally valuable to no communication at all.

The results above show that the decision-theoretic message selection exhibits psychological plausibility: In simple situations above humans tend to communicate in the same way. However, we expect that our approach to rational communicative behavior will be even more valuable in more complex situations. For example, we can easily incorporate factors such as unreliability of the communication channels, and the probability of the messages being intercepted by the enemy, into our decision-theoretic calculations. In these situations, humans, accustomed to communicative traits that usually do not include such factors, may exhibit communicative behaviors that is further from the optimal. we will test the above conjecture in our future work.

References

[1] M. Barbuceanu and M. S. Fox. Integrating communicative action, conversations and decision theory to coordinate agents. In *Proceedings of Agents '97*, Marina del Rey, Feb. 1997.

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

[3] 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.

[4] 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, Menlo Park, 1995. AAAI Press/The MIT Press.

[5] P. J. Gmytrasiewicz and E. H. Durfee. Rational interaction in multiagent environments: Communication. *Submitted for publication*, 1997. Available in postscript from http://www-cse.uta.edu/ piotr/piotr.html.

[6] P. J. Gmytrasiewicz, E. H. Durfee, and D. K. Wehe. The utility of communication in coordinating intelligent agents. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, pages 166–172, July 1991.

[7] P. J. Gmytrasiewicz, S. Noh, and T. Kellogg. Bayesian update of agent models: Foundations and results. *Submitted for publication*, 1997. Available in postscript from http://www-cse.uta.edu/ piotr/piotr.html.

[8] B. J. Grosz and C. Sidner. Plans for discourse. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in Communication*. MIT Press, 1990.

[9] S. Noh and P. J. Gmytrasiewicz. Agent modeling in antiair defense. In A. Jameson, C. Paris, and C. Tasso, editors, *Proceedings of the Sixth International Conference on User Modeling*, pages 389–400, New York, June 1997. Springer-Verlag.

[10] S. Noh and P. J. Gmytrasiewicz. Multiagent coordination in antiair defense: A case study. In M. Boman and W. V. de Velde, editors, *Multi-Agent Rationality – MAAMAW'97 Workshop, Lecture Notes in Artificial Intelligence Vol. 1237*, pages 4–16. Springer, New York, May 1997.

[11] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufman, San Mateo, CA, 1988.

[12] A. Quintero. Multi-agent system protocol language specification. Technical report, Universidad de los Andes, Colombia, 1996.

[13] S. J. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach.* Prentice-Hall, Englewood Cliffs, New Jersey, 1995.

[14] I. A. Smith and P. R. Cohen. Toward a semantics for an agent communications language based on speech-acts. In *Proceedings of the National Conference on Artificial Intelligence*, pages 24–31, Portland, Oregon, Aug. 1996.

[15] M. Tambe. Agent architectures for flexible, practical teamwork. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, pages 22–28, Providence, Rhode Island, 1997. AAAI Press/MIT Press.