The Utility of Communication in Coordinating Intelligent Agents

Piotr J. Gmytrasiewicz*, Edmund H. Durfee†, and David K. Wehe*
*Department of Nuclear Engineering
†Department of Electrical Engineering and Computer Science
University of Michigan
Ann Arbor, Michigan 48109

Abstract

When intelligent agents who have different knowledge and capabilities must work together, they must communicate the right information to coordinate their actions. Developing techniques for deciding what to communicate, however, is problematic, because it requires an agent to have a model of a message recipient and to infer the impact of a message on the recipient based on that model. We have developed a method by which agents build recursive models of each other, where the models are probabilistic and decision-theoretic. In this paper, we show how an agent can compute the impact of a message in terms of how it increases (or decreases) its expected utility. By treating the imperfect communication channel probabilistically, our method allows agents to account for risk in committing to nonintuitive courses of action, and to compute the utility of acknowledging messages.

Introduction

When operating in multiagent environments, intelligent agents must generally coordinate their activities to avoid interfering with each other and to cooperate when they can mutually benefit. A crucial hindrance to effective coordination, however, is that intelligent agents might not know enough about each other's intentions, abilities, and perspectives to anticipate interactions. Unless they are somehow designed initially with rich models of each other, the intelligent agents must rely on communication among themselves to share the knowledge that is vital for coordination.

Communication is, however, a two-edged sword. If agents communicate too little, they risk interfering with each other and missing opportunities for achieving goals more efficiently. If agents communicate too much, they risk overwhelming each other with unimportant information, which impedes timely decision making. The challenge in designing algorithms for communication decision making, therefore, is in providing those algorithms with the ability to estimate the impact—or utility—of a message on both the sender and receiver. This ability is especially critical when communication bandwidth is extremely restricted, so that agents must be very selective about what messages are worth sending.

An agent that is considering sending a message should thus base its decision on an estimate of whether the message's recursive impact on the sender's and receiver's beliefs will improve the expected outcome of its decisions. We have developed a rigorous approach for modeling the utility of communication based on decision and game theoretic methods. In our approach, an agent begins with a recursively elaborated set of models about another agent. Using the probabilistic nature of these models, the agent can compute the expected utilities for the other agents' alternative decisions in the situation. It can then model how an exchange of information will influence the probabilities, and thus affect the other agent's decisions' expected utilities. This in turn will impact the initial agent's probability distribution about the other agent's activities, which can increase the expected utility of the initial agent's action.

As we describe in the next section, previous work on intelligent communication has emphasized static policies for deciding what messages are important to send, or treated communication as a tool for making deals among the agents. The contribution of the work we report in the subsequent sections is that it uses a recursive modeling technique to explicitly compute the expected utility of a message as the expected utility of the decision given the message minus the expected utility of the best decision prior to the message. This allows an agent to quantitatively estimate the importance of a message. Moreover, by working backward from a desired recursive model, it can guide the search for appropriate messages to transform the current recursive model into the desired one. We also show how our method can be applied to cases in which the communication is transmitted via an imperfect channel.

[0]This research was supported, in part, by the Department of Energy under contract DG-FG-86NE37969, and by the National Science Foundation under grant IRI-9015423.

Related Work

The intuition that communication is an essential tool for coordination has been demonstrated using game- and decision-theoretic frameworks, and is evident even in simple two-person games such as the "Battle of the Sexes" [Luce and Raiffa, 1957]. These frameworks provide analytical tools for describing the consequences of communication, and have been extended by Rosenschein and his colleagues, who have been developing a unified negotiation protocol [Rosenschein and Genesereth, 1985; Zlotkin and Rosenschein, 1989; Zlotkin and Rosenschein, 1990b]. In their work, communication is primarily a tool to get agents to converge on a joint plan, or deal, that guarantees them payoffs higher than they expect to get if they do not make a deal. They have also examined how agents can exchange joint plans or the information needed to converge on joint plans [Rosenschein and Genesereth, 1987], and our results similarly show the importance of both types of messages.

Whereas Rosenschein and Genesereth have developed communication strategies for logic-based agents, other researchers have developed strategies for other types of agents. For example, Durfee and his colleagues have employed heuristic communication policies to balance the relevance, timeliness, and completeness of messages [Durfee *et al.*, 1987]. These heuristics guide cooperating problem solvers into selectively exchanging partial solutions in a more effective manner.

Speech act theory [Cohen and Levesque, 1990; Perrault, 1990] is also concerned with the impact of a communication act on a recipient agent. Thus, agents need to model each other, and to model how others model them, and so on. While our approach also exploits the agents' recursive nesting of beliefs, our emphasis is not so much on developing a logical formalism for modeling interagent communication as it is on quantifying the expected utility of communication. Halpern and Moses [Halpern and Moses, 1984] have considered how the recursive nesting of beliefs leads to difficulties in converging on "common knowledge." Our investigations confirm these difficulties, but emphasize that decisions are possible in some cases without common knowledge.

The Recursive Modeling Method (RMM)

In multiagent worlds, the utility of an agent's action can depend on the concurrent actions of other agents, so an agent should attempt to predict the intended actions of others when deciding on its own. Because the other agents are likely to be modeling it as well, the agent must recursively build models of others and itself. We have developed a Recursive Modeling Method (RMM) to create these models explicitly [Gmytrasiewicz *et al.*, 1991a]. By employing RMM, an agent exploits any information about others that it has,

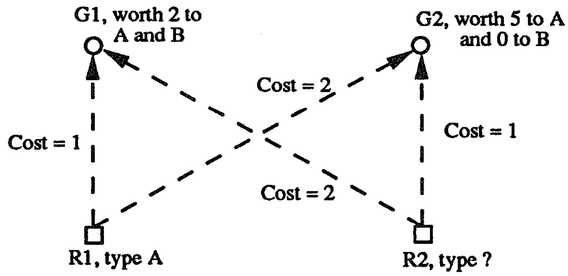

Figure 1: Example Scenario of Interacting Agents

and summarizes uncertainties as probability distributions. Furthermore, an agent can use RMM to model how other agents model it, and so on into deeper levels of modeling.

Before we introduce the general form of RMM, consider this example (Figure 1). We assume that the environment can be populated by type A agents and type B agents. Type A agents can perceive all of the goals in the environment, can perform actions to achieve them all, and know about both types of agents. Type B agents can only see goals of type G1, cannot perform actions to achieve type G2 goals, and know only about type B agents. The utility of achieving goals of type G1 is 2 for both types of agents, and for achieving goals of type G2 is 5 for type A agents and 0 for type B agents. The cost of attempting to achieve the farther goal is 2, and the closer goal 1, irrespective of an agent's type. For simplicity, each agent can only achieve one goal.

Agent R1, a type A agent, has 3 options: pursue G1, pursue G2, or do something (including nothing) else (G1, G2 and S for short). R1 computes its payoffs as the difference between the sum of the worth it assigns to *all* the achieved goals (whether or not it personally performed them all) and the costs it personally incurred in achieving its own goal. These payoffs are represented as a matrix (top of the hierarchy in Figure 2). The expected utilities of R1's options naturally depend on what R2 chooses to do concurrently. R1 can assume that R2 will maximize its own payoff [Dennett, 1986], but R1 does not know whether R2 is of type A or B, so R1 does not know whether R2 will pursue G2. The payoff matrices of R2, as modeled by R1, depend on R2 being type A or B, with probabilities $p^{R1}_{R2(A)}$ and $p^{R1}_{R2(B)}$, respectively (where $p^{R1}_{R2(A)} + p^{R1}_{R2(B)} = 1$). The alternative views are depicted on the second level of the hierarchy in Figure 2.

Furthermore, if R2 is of type B, then it will model R1 as type B as well (because it does not know of type A agents). If R2 is of type A, then its model of R1 would account for R1 being either type A or B. Thus, R1 will model R2, and R2's model of R1, and R2's model of R1's model of R2, and so on, as depicted in Figure 2.

R1 builds and uses this hierarchy to help guess what

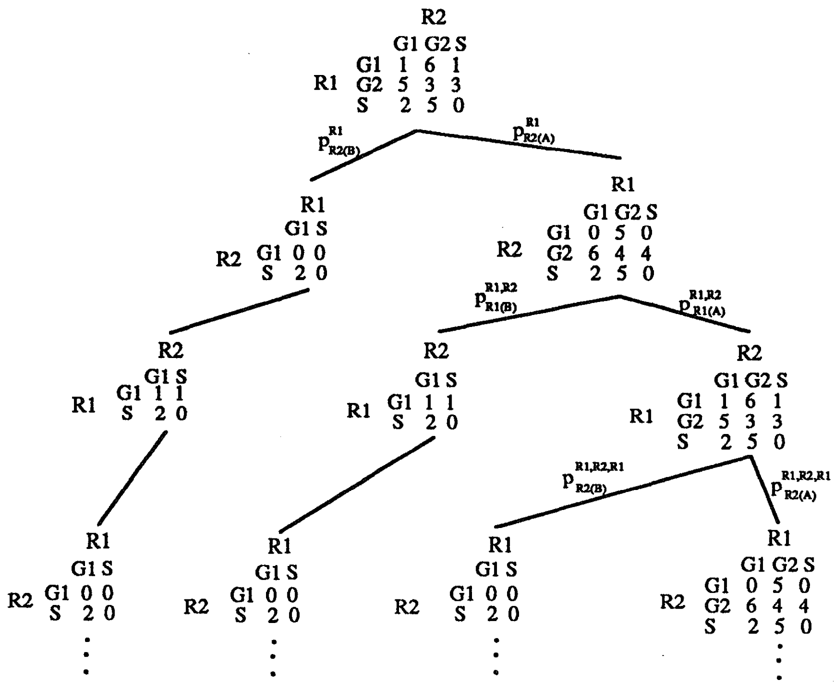

Figure 2: Example of the Recursive Hierarchy

R2 will decide to do. It summarizes its guesses as a probability distribution $p^{R1}_{R2} = (p_{G1}, p_{G2}, p_S)$ indicating the probability that R2 will intend to pursue G1, G2, and S. We call this representation an *intentional probability distribution*, because it specifies R1's view of R2's intentions, given what R1 knows about R2. Assume that $p^{R1}_{R2(A)} = p^{R1,R2}_{R1(A)} = \cdots = 0.5$, meaning that R1 believes that R2 is equally likely to be type A or B, and that R2, if type A, will believe that R1 is equally likely to be type A or B, and that R2, if type A, will believe that R1 will believe that R2 is equally likely to be type A or B, and so on. Treating all other uncertainties equiprobabilistically as well, R1 uses the hierarchy to compute the intentional probability distribution over R2's options as $p^{R1}_{R2} = (0, 0.5, 0.5)$. Intuitively, if R2 is type B, then it will expect R1 to pursue G1 and S with equal probability, and so R2 would be better off, on average, pursuing S. If R2 is type A, then it will see R1 as either type A (and so expect it to pursue G2) or type B (and so expect it to pursue G1 or S), and R2 would thus do best pursuing G2 (the details of the computation are given elsewhere [Gmytrasiewicz *et al.*, 1991a]). The p^{R1}_{R2} distribution gives expected utilities of 3.5, 3, and 2.5 for R1's options G1, G2, and S, respectively. R1 should thus pursue G1.

Temporarily stepping back from this example, the more general formulation of RMM assumes that R1 is interacting with $(N-1)$ other agents, R2–RN. The utility of R1's m-th alternative option can be evaluated as:

$$u^{R1}_m = \sum_k \cdots \sum_l \{p^{R1}_{R2-k} \cdots p^{R1}_{RN-l} u^{R1}_{m,k,\ldots,l}\} \qquad (1)$$

where p^{R1}_{Ri-k} represents the probability R1 assigns to Ri's intending to act on the k-th element of Ri's set of options; as mentioned before, we will refer to these as intentional probabilities. $u^{R1}_{m,k,\ldots,l}$ is R1's payoff (utility) as an element of the N-dimensional game matrix.

R1 can estimate the intentional probabilities p^{R1}_{Ri-k} by guessing how the game looks from Ri's point of view. R1 models each Ri using probability distributions p^{R1}_{Pi}, p^{R1}_{Ai}, and p^{R1}_{Wi}, which we call modeling probabilities. p^{R1}_{Pi} summarizes R1's knowledge about Ri's preferences (goals it will value). p^{R1}_{Ai} summarizes R1's knowledge about Ri's abilities (goals it can pursue), given its preferences. p^{R1}_{Wi} summarizes R1's knowledge about Ri's world model, given its abilities.[1] In every case of Ri having various preferences, abilities and world models, R1 assumes that Ri is rational and considers the probability that the k-th element of Ri's set of options is of the highest utility to Ri. The modeling probabilities can then be used to compute the intentional probabilities p^{R1}_{Ri-k}, as the following probabilistic mixture:

[1] In the example shown in Figure 2, the modeling probabilities were combined into probabilities over agent types. This simplifies the example because agent types serve to encapsulate preferences, abilities, and world models.

$$p^{R1}_{Ri-k} = \sum_{Pi}\sum_{Ai}\sum_{Wi}\{p^{R1}_{Pi}p^{R1}_{Ai}p^{R1}_{Wi} \times Prob(Max_{k'}(u^{R1,Ri}_{k'}) = u^{R1,Ri}_{k})\} \quad (2)$$

where $u^{R1,Ri}_{k'}$ is the utility R1 estimates that Ri will assign to its option k', and is computed as

$$u^{R1,Ri}_{k'} = \sum_{r} \cdots \sum_{s} \{p^{R1,Ri}_{R1-r} \cdots p^{R1,Ri}_{RN-s} u^{R1,Ri}_{k',r,\ldots,s}\} \quad (3)$$

The $u^{R1,Ri}_{k',r,\ldots,s}$ is how R1 sees Ri's payoffs in the N-dimensional game matrix. The probabilities R1 thinks Ri assigns to agent Rn acting on its o-th option $p^{R1,Ri}_{Rn-o}$, can in turn be expressed in terms of $p^{R1,Ri,Rn}_{Rv-w}$ and $u^{R1,Ri,Rn}_{o',w,\ldots}$ and so on.

As we detail elsewhere [Gmytrasiewicz *et al.*, 1991a], it is possible to determine the convergence of RMM without going very deep into the recursive levels (usually 4-th or 5-th level). In the cases when RMM does not converge on a unique probability distribution over options of the other agents, we compute the expected intentional distribution as a probabilistic mixture of the distributions between which RMM cannot decide and use it to uniquely determine the agent's best option (see [Gmytrasiewicz *et al.*, 1991a] for an example).

The Utility of Communication

We treat decisions about communication just like decisions about any other actions, and thus employ decision-theoretic techniques to select the action with the highest expected utility [Gmytrasiewicz *et al.*, 1991b]. For communication actions, the agents use their nested models to predict the impact of messages on the expected utilities of alternative actions, and then send the highest utility message—the message that causes the greatest gain in the expected utility of the agent's action.

For example, using a hierarchy such as that in Figure 2, an agent computes an intentional probability distribution p (ignoring the superscript and subscript for now), over the other agents' options. The initial agent can thus compute its best choice, which we denote as X, as the action with the highest expected utility $U_p(X)$.

If the initial agent sends a message M to the other agent, the message causes the receiving agent to modify its hierarchy, and thus causes the intentional probability distribution over its options to change to p_M. This new distribution in turn can affect the expected utilities of the initial agents actions, such that the action Y that it will now take (which may or may not be the same as X) has an expected utility of $U_{p_M}(Y)$. The utility of a message, M, is defined as the difference between the expected utility of the preferred action before M was sent and the expected utility of the agent's chosen action after the message was sent:

$$U(M) = U_{p_M}(Y) - U_p(X). \quad (4)$$

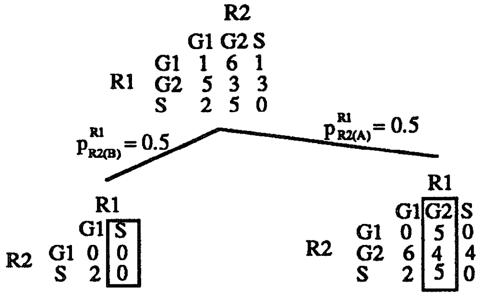

Figure 3: Recursive Hierarchy After a Received Intentional Message

We broadly classify messages into types, depending on how they will impact a recipient and sender. In this paper we investigate two types—intentional and modeling messages.

Intentional Messages

An intentional message corresponds to an agent committing to a choice of action, and informing other agents about it, i.e. it contains information about the intentional probabilities $p^{R1,Ri}_{R1-r}$ in equation (3). If we assume that agents must meet their commitments, then a recipient can use this message to predict exactly what the sender will do. In modeling the recipient, therefore, the sender can truncate the recursion because it knows exactly how it will be modeled by the recipient.

For example, consider the scenario in Figure 1 and the hierarchy in Figure 2. As discussed before, R1's best option before communication is to pursue G1, with its expected utility of 3.5. On inspecting the hierarchy (Figure 2), however, note that R1 has, on the average, better payoffs if it pursues G2. The question is, can it change R2's preferences to take advantage of these payoffs?

The answer, in this case, is yes. Suppose that R1 considers transmitting an intentional message M_1 to R2 declaring its intention to pursue G2. R1 can thus truncate the hierarchy (Figure 3). If R2 is type B and receives M_1, it models R1 as pursuing S, and so, for R2, the options G1 and S are equally good and equally likely. If R2 is type A and receives M_1, it also sees options G1 and S as equally likely. Thus, the new probability distribution over R2's options is $p_{M_1} = (0.5, 0, 0.5)$. R1 has committed itself to G2, but now R1 computes the expected utility of G2 as 4. According to equation (4), therefore, the utility of the message M_1 is $U(M_1) = U_{p_{M_1}}(G2) - U_p(G1) = 0.5$.

The above analysis assumes that R2 is guaranteed to receive M_1. Unfortunately, communication channels are seldom so reliable.[2] Because M_1 commits R1

[2] While low-level networking protocols can make commu-

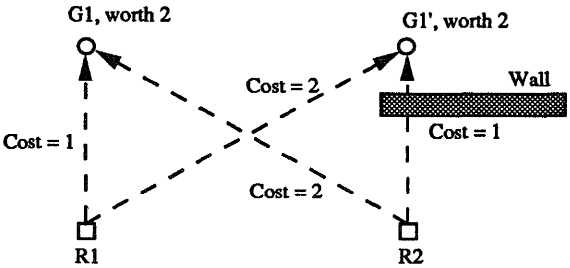

Figure 4: Variation of the Example Scenario

to pursuing G2, which is not what R2 will expect it to do, the failure of the message to arrive at R2 might disadvantage R1. We can formalize R1's risk by assuming that the communication channel will correctly deliver a message with probability p_c, where $0 \leq p_c \leq 1$. From R1's perspective: with probability p_c, M_1 will be received, in which case the probability distribution over R2's options is (0.5, 0, 0.5) as we just derived; and with probability $(1 - p_c)$, M_1 will not be received, so the probability distribution over R2's options is the same as with no communication at all: (0, 0.5, 0.5). Combining these we get:

$$p_c(0.5, 0, 0.5) + (1 - p_c)(0, 0.5, 0.5) = (0.5p_c, 0.5 - 0.5p_c, 0.5).$$

Because R1 is committed to pursuing G2, it computes the expected utility of this option to be $U_{p_{M_1}}(G2) = 3 + p_c$, so $U(M_1) = p_c - 0.5$. In other words, R1 should only send M_1 when $p_c > 0.5$. When $p_c < 0.5$, in fact, the communication is ill-advised: R1 is taking too big a risk that it will commit to an option that R2 will not know to support.

Modeling Messages

Modeling messages contain information about the modeling probabilities p_{Pi}^{R1}, p_{Ai}^{R1}, and p_{Wi}^{R1} in equation (2) and update the hearer's and the speaker's model of the multiagent world. For example, consider what would happen in a variation of our original scenario. In this variation (Figure 4), both agents are of type A and, instead of G2, we have G1′, so that both agents will regard both goals as equally valuable. Also, there is a wall that probably obstructs R2's view of G1′.

The recursive hierarchy R1 will build for this scenario is depicted in Figure 5, where $p_{W2(G1')}^{R1}$ represents the probability that R1 assigns to R2 having G1′ in its world model, which we assume to be low (0.01) because of the wall. R1 also assumes that, if R2 does see G1′, it knows that R1 sees it too. As the progressively deeper levels of the left branch are analyzed, the solution quickly converges on R2's best option being S. The analysis of the right branch shows that the best option of R2, if it sees G1′, is to pursue it. The resulting probability distribution over R2's moves is then $p = (0, 0.01, 0.99)$, which results in G1 being R1's best choice, with its expected utility of 1.02.

Intentional messages will not help in this case: R1's committing itself to G1 results in the same expected utility (because that is what R1 believes R2 will expect R1 to do, anyway); R1's commitment to G1′ gives $p_{M_{G1'}} = (0.5, 0, 0.5)$ over R2's options, which results in an expected utility of G1′ of 1; and R1's commitment to S results in an expected utility of 1.01. Thus, none of these messages can better the expected utility of 1.02 gained without communication.

However, on inspecting R1's payoff matrix, it is clear that R1's option of G1 will be better if R1 can increase the chances of R2 pursuing G1′. This requires that R1 increase the probability of the right branch of the model. A simple way to do this is for R1 to send R2 message M_2 stating that goal G1′ is behind the wall. If, for the time being, we assume that communication channels never lose messages, then R1 models R2's response by changing $p_{W2(G1')}^{R1}$ to 1, and $p_{W2(G1')}^{R1,R2,R1}$ to 1, and so on. Due to message M_2, the hierarchy has only one branch all the way down. Computing utilities, R1's best option is still G1, but now it expects R2 to pursue G1′. R1's option thus now has an expected utility of 3 rather than 1.02, so $U(M_2) = 1.98$.

Considering that communication channels are not perfect, let us see how the probability of communication p_c factors into these calculations. Combining the weighted intentional probabilities for R2, we get:

$$\begin{aligned} p_{M_2} &= p_c(0, 1, 0) + (1 - p_c)(0, 0.01, 0.99) \\ &= (0, 0.01 + 0.99p_c, 0.99 - 0.99p_c) \end{aligned}$$

which gives an expected utility for G1 (still R1's best choice) as $U_{p_{M_2}}(G1) = 1.02 + 1.98p_c$, for a message utility of $U(M_2) = 1.98p_c$. In other words, sending the message is always to R1's benefit, unless $p_c = 0$. And if $p_c = 1$, the message allows R1 to maximize its expected utility.

Finally, consider how R1's model changes if R2 acknowledges receiving M_2, as depicted in Figure 6. Even after an acknowledgement, the model still includes uncertainty associated with R2 not knowing whether R1 received the acknowledgement. Because R1 now knows that R2 knows about G1′, p_c no longer enters into the probability mixture R1 has of R2's intentions (because R2's only rational option given that it now knows about G1′ is to pursue G1′). G1 is still R1's best choice, but now has an expected utility of 3, meaning that the utility of the acknowledgement message is equal to $1.98(1 - p_c)$. Additional acknowledgement messages have no influence on R1's expected utility in this case because, once it knows that R2 has received the message, it knows that R2 should pursue G1′ regardless of deeper uncertainties about knowledge about knowledge. Thus, in this case, the agents con-

nication more reliable, they still cannot ensure that message exchange can lead to common knowledge [Halpern and Moses, 1984].

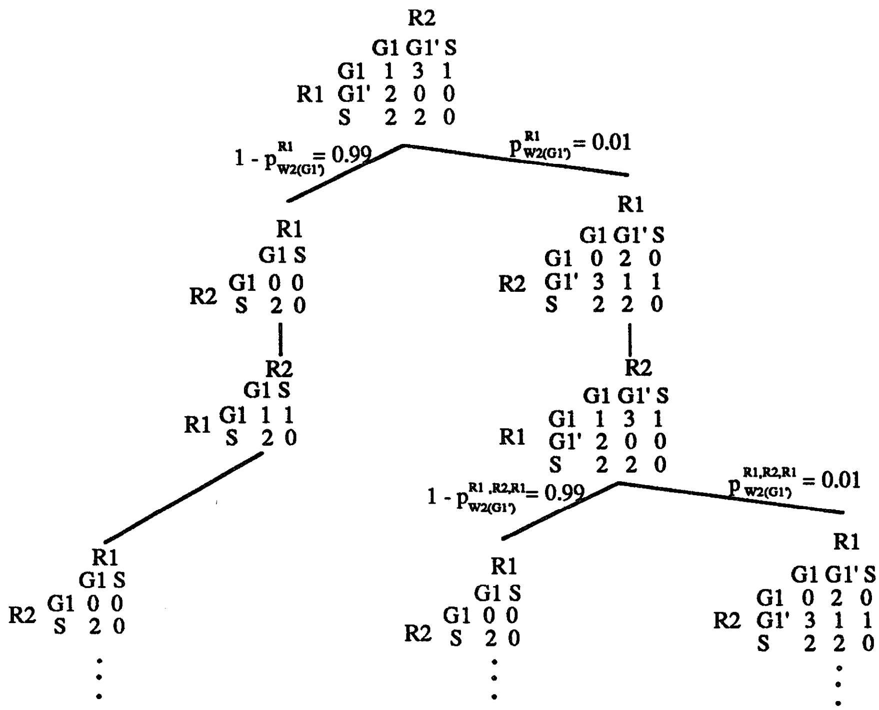

Figure 5: Recursive Hierarchy for Scenario Variation

verge on appropriate actions without "common knowledge" [Halpern and Moses, 1984]. More generally, our analysis of cases where the choices of actions depend on deeper levels of knowledge indicate that the utility of successive acknowledgement messages decreases due to higher polynomial terms in p_c. By deciding upon a threshold for message utility, for instance equal to the price of message transmission, the agents can truncate this infinite regress of acknowledgements.

Discussion

It is intuitive that, when deciding whether or not to send a message to another agent, a sending agent should consider the expected benefits it will accrue due to that message. Given that it must limit communications to the most important messages (due to limited communication and/or processing resources), computing the utilities of messages and using this measure to guide communication decisions makes sense. What we have developed in this paper is a formalization of these intuitions, in which agents recursively model each other and can assess the expected impact of messages based on these models. These results advance the state of the art, which generally employs static strategies for making communication decisions rather than evaluating the expected utility of each potential message. The tradeoff, however, is that the recursive modeling can be computationally costly, and this overhead must be weighed against the costs and benefits of a simpler but

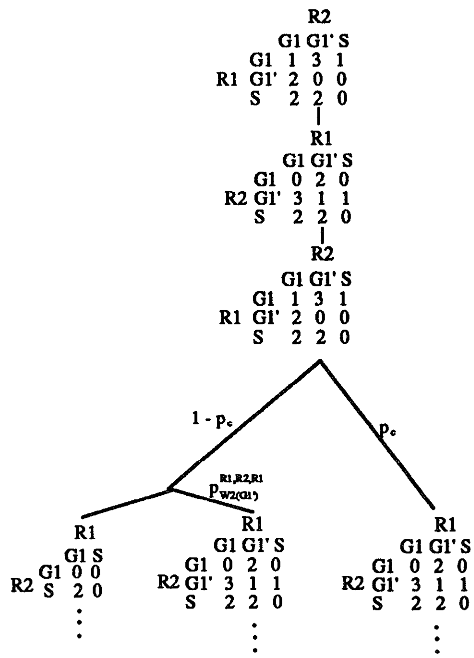

Figure 6: Recursive Hierarchy After M_2 was Acknowledged

cruder approach.

We are now extending our investigation in several ways. The reader will likely have noticed in our analyses that the generation of messages was only vaguely described in terms of an agent trying to prune or truncate the probabilistic models. While we have identified some criteria for generating potential messages, we need to define an algorithm for this process. Moreover, we would like to account for other message types, such as requests for information and imperatives [Cohen and Levesque, 1990]. We would also like to extend the modeling and messaging to allow agents to potentially lie to each other, and to model each other as potential liars [Zlotkin and Rosenschein, 1990a]. Finally, we are exploring the practical implications of using this approach in the context of a robotic application for nuclear power plant environments, which are rigorously designed such that developing probabilistic models is especially feasible.

References

[Cohen and Levesque, 1990] 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.

[Dennett, 1986] D. Dennett. Intentional systems. In *Brainstorms.* MIT Press, 1986.

[Durfee *et al.*, 1987] Edmund H. Durfee, Victor R. Lesser, and Daniel D. Corkill. Cooperation through communication in a distributed problem solving network. In Michael N. Huhns, editor, *Distributed Artificial Intelligence*, Research Notes in Artificial Intelligence, chapter 2, pages 29–58. Pitman, 1987. (Also in S. Robertson, W. Zachary, and J. Black (eds.), *Cognition, Computing, and Cooperation*, Ablex 1990.).

[Gmytrasiewicz *et al.*, 1991a] Piotr J. Gmytrasiewicz, Edmund H. Durfee, and David K. Wehe. A decision-theoretic approach to coordinating multiagent interactions. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, August 1991.

[Gmytrasiewicz *et al.*, 1991b] Piotr J. Gmytrasiewicz, Edmund H. Durfee, and David K. Wehe. Combining decision theory and hierarchical planning for a time-dependent robotic application. In *Proceedings of the Seventh IEEE Conference on AI Applications*, pages 282–288, February 1991.

[Halpern and Moses, 1984] Joseph Y. Halpern and Yoram Moses. Knowledge and common knowledge in a distributed environment. In *Third ACM Conference on Principles of Distributed Computing*, 1984.

[Luce and Raiffa, 1957] R. D. Luce and H. Raiffa. *Games and Decisions.* John Wiley and Sons, 1957.

[Perrault, 1990] C. R. Perrault. An application of default logic to speech act theory. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in Communication.* MIT Press, 1990.

[Rosenschein and Genesereth, 1985] Jeffrey S. Rosenschein and Michael R. Genesereth. Deals among rational agents. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 91–99, Los Angeles, California, August 1985. (Also published in *Readings in Distributed Artificial Intelligence*, Alan H. Bond and Les Gasser, editors, pages 227–234, Morgan Kaufmann, 1988.).

[Rosenschein and Genesereth, 1987] Jeffrey S. Rosenschein and Michael R. Genesereth. Communication and cooperation among logic-based agents. In *Proceedings of the Sixth Phoenix Conference on Computers and Communications*, pages 594–600, Scottsdale, AZ, February 1987.

[Zlotkin and Rosenschein, 1989] Gilad Zlotkin and Jeffrey S. Rosenschein. Negotiation and task sharing among autonomous agents in cooperative domains. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 912–917, August 1989.

[Zlotkin and Rosenschein, 1990a] Gilad Zlotkin and Jeffrey S. Rosenschein. Blocks, lies, and postal freight: Nature of deception in negotiation. In *Proceedings of the 1990 Distributed AI Workshop*, October 1990.

[Zlotkin and Rosenschein, 1990b] Gilad Zlotkin and Jeffrey S. Rosenschein. Negotiation and conflict resolution in non-cooperative domains. In *Proceedings of the National Conference on Artificial Intelligence*, pages 100–105, July 1990.