The Utility of Embedded Communications: Toward the Emergence of Protocols*

Edmund H. Durfee
EECS Department
University of Michigan
Ann Arbor, MI 48109
durfee@umich.edu

Piotr Gmytrasiewicz
Computer Science Dept.
University of California-Riverside
Riverside, CA
piotr@ucrengr.ucr.edu

Jeffrey S. Rosenschein
Computer Science Dept.
Hebrew University
Jerusalem, ISRAEL
jeff@cs.huji.ac.il

Abstract

A fundamental feature of effective distributed systems is that the entities comprising the system have some set of guidelines—some plan to follow—that leads them into making good decisions about what to communicate and when. Traditionally, these *protocols* for communication have been given to the entities at the time that they are designed. For example, knowledge-based entities (agents) have been designed with protocols that allow them to make deals, allocate tasks, negotiate over solutions, and so on. Such distributed systems, however, will be brittle if the agents ever need to go beyond the pre-existing protocol. To constitute a robust system, the agents would benefit from the ability to discover new ways of communicating, and to generalize these into new protocols.

This paper extends the recursive modeling method to address issues of embedded communications—communications occurring in a larger context of other physical and/or communicative activities, and describes how behaviors like question-answering and order-following could emerge as rational consequences of agents' decisionmaking. These types of embedded communicative acts can form the building blocks of more complex protocols, given that agents can not only derive these embedded communicative acts but can generalize and reuse them appropriately.

Introduction

Agents typically exist in ongoing multiagent environments, and so, interactions among agents are embedded within a larger multiagent context. In particular, verbal and nonverbal communication among agents is embedded in an evolving multiagent context. Planning a communicative action in this context is a complex task, since considering all of the features of the context, and all of the possible messages that could be transmitted (through information channels or through changes to the physical environment), can be prohibitively time-consuming.

To reduce the combinatorics of communication planning, intelligent agents often rely on retrieving and following scripted interactions that delineate appropriate communicative acts and responses for typical multiagent contexts. Designers of artificial agents generally build in such "protocols" for interaction so as to constrain the interactions among agents to achieve predictable and efficient performance. For example, when building multiagent systems, designers have devised protocols for:

- consistent information distribution/replication (as in the distributed database literature);
- establishing commitments over resource allocation (as in resource locking protocols in distributed systems);
- remotely executing commands and accessing services (as in remote procedure calls);
- querying processing in database and knowledge base systems (as in the KQML and SKTP methods [15]);
- contracting and negotiation among agents (as in the Contract Net [17], the Unified Negotiation Protocol [18], and partial global planning [5]).

When developing systems for accomplishing well-understood tasks in well-defined environments, it is reasonable and proper for a system designer to define and institute appropriate protocols. However, when the nature of agents' tasks might change, or their environment might undergo substantial changes, being locked into a particular interaction protocol might lead to ineffective action and interaction on the parts of the agents. Accomplishing tasks in such environments might require agents to invent new protocols based on experience and on expectations about how messages will affect each other.

Our work is particularly interested in the second means for inventing new protocols, and is interested in

*The authors are listed alphabetically. This work was supported, in part, by the University of Michigan Rackham Graduate School's International Partnerships Program, by the National Science Foundation under PYI award IRI-9158473, by the Golda Meir Fellowship and the Alfassa Fund administered by the Hebrew University of Jerusalem, by the Israeli Ministry of Science and Technology (Grant 032-8284), and by the S. A. Schonbrunn Research Endowment Fund (Hebrew University) (Grant 034-7272).

answering the following question: what first-principles knowledge about the pragmatics of communication can agents use to plan rational communication actions in the absence of protocols, such that new protocols can emerge? Our objectives in this paper are thus to articulate how such first-principles knowledge might be represented and used, and to explain how individual decisions about messages to send can prompt further physical and communicative actions, leading to common types of embedded interactions such as question answering, order following, and threatening. We thus extend the concept of isolated message utility into embedded message utility, and augment our previous work by introducing communication and deliberation costs into the framework to motivate why self-interested agents would nevertheless engage in extended interactions such as answering questions and following orders. This work thus lays the foundation for designing agents who can create their own interactive protocols based on their knowledge of embedded interactions and the costs of reasoning from first principles.

We begin by describing prior work in representing agents as rational utility-maximizers who model each other as utility maximizers. By developing nested models of each other, they can represent the pragmatic meaning of (verbal or nonverbal) communicative actions as how these actions change the recursive nesting, leading to changes in expected utility. Having established the pragmatics of individual messages, we then turn, in Section , to communicative actions embedded in ongoing physical activities, where the utility of a message is dependent on the physical actions preceding and succeeding the message. It is in this context that messages embodying threats can make sense. In Section , we investigate the utility of messages that are part of ongoing communicative activities—that is, that are part of a dialogue—and motivate how dialogues involving question answering and order taking can in fact emerge from decisions made by autonomous agents. We conclude with a discussion of how patterns of dialogues could lead to the identification of protocols, and outline our ongoing work in this and other directions.

Recursive Modeling and the Utility of Communication

In [10], a method called the Recursive Modeling Method (RMM) was used to represent and evaluate all of the knowledge relevant to the problem of choosing the best action, in the presence of other agents. RMM uses a recursive hierarchy of simple models of the agents' decision-making situations (payoff matrices) to represent the information an agent has about its physical environment, about how the other agents view this environment, about how the other agents could view the original agent, how they could view the original agent viewing them, and so on.

An Example Model

To put our description of RMM in concrete terms, we will consider a particular decision-making situation encountered by an autonomous outdoor robotic vehicle, called R_1, (see Figure 1), attempting to coordinate its actions with another robotic vehicle, R_2. Outdoor robotic vehicles have multiple uses, predominantly acting in environments that are too hostile or hazardous for human-controlled vehicles. Among these uses are information gathering—or reconnaissance—activities to assess, for example, the extent of a chemical spill (in the case of an industrial accident), the positions of opposing forces (on a battlefield), or the location of a submerged ship (in the case of an underwater vehicle).

For a ground-based vehicle, gathering large amounts of information depends on moving to vantage points that command a wide view, such as hilltops. Thus, we will assume that a robotic vehicle, whose mission is to gather as much information as it can while minimizing its cost (fuel and/or time consumed), will prefer to move to nearby locations with high elevation. From the perspective of robot R_1, whose point of view we will take in analyzing this situation, two possible vantage points P1 and P2 are worth considering. P2 has a higher elevation and would allow twice as much information to be gathered as P1, and so the robot is willing to incur greater cost to go to P2. Based on domain-specific knowledge, in this example R_1 expects that gathering information at P2 will be worth incurring a cost of 4 (or, put another way, the information gathered from P2 has an expected value of 4), while the observation from P1 will be worth 2.

R_1 thus has three possible courses of action: it can move to P1 and gather information there (a_1^1); it can move to P2 and gather information there (a_2^1); or it can do neither and just sit still (a_3^1).[1] The expected cost (time or energy) to R_1 of pursuing each of these courses of action is proportional to the distance traveled, yielding a cost of 1 for a_1^1, 2 for a_2^1, and 0 for a_3^1. We further assume in this example that each of the robots can make only one observation, and that each of them benefits from *all* information gathered (no matter by which robot), but incurs cost only based on its own actions.

Given the above, robot R_1 can build a payoff matrix that summarizes the information relevant to its decision-making situation. The relevant alternative behaviors of R_2 that matter will be labeled a_1^2 through a_3^2, and correspond to R_2's taking the observation from

[1] Of course, each of these high-level courses of action must be further elaborated by the robot. While all possible detailed plans for these high-level courses of action could be enumerated and represented in a payoff matrix, the abstraction over actions and plans permits evaluation of choices at a level of detail where the quality of the decision is maximized while the costs of making the decision are minimized [12].

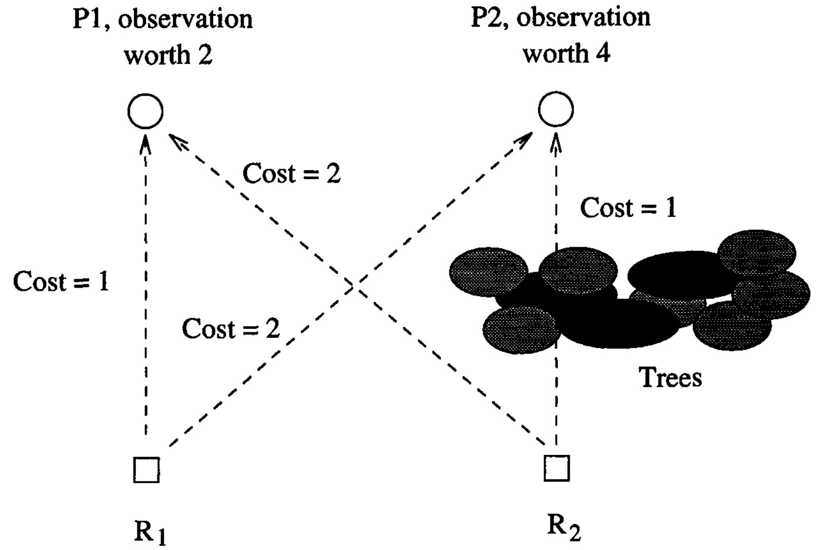

Figure 1: Example Scenario of Interacting Agents

point P1, P2, and staying put (or doing something else), respectively. Thus, the entry in the matrix corresponding to R_1's pursuing its option a_1^1 and R_2's pursuing a_2^2 is the payoff for R_1 computed as the total value of the information gathered by both robots from both P1 and P2, minus R_1's own cost: $(2+4)-1=5$. The payoff to R_1 corresponding to R_1's pursuing a_1^1 and R_2's pursuing a_1^2 is $(2+0)-1=1$, since the information gathered is worth 2, and redundant observations add no value. All of the payoffs can be assembled in the following payoff matrix:

$$\begin{array}{cccc} & & R_2 & \\ & a_1^2 & a_2^2 & a_3^2 \\ R_1 \quad a_1^1 & 1 & 5 & 1 \\ a_2^1 & 4 & 2 & 2 \\ a_3^1 & 2 & 4 & 0 \end{array}$$

Since R_1's payoff depends on what R_2 decides to do, R_1 has to model R_2's decision-making. If R_1 thinks that R_2 will attempt to maximize its own expected utility, then R_1 can adopt the intentional stance toward R_2 [4], and treat R_2 as rational. Thus, R_2's payoff matrix, if it knows about both observation points, is arrived at analogously to R_1's matrix above, and has the following form:

$$\begin{array}{cccc} & & R_1 & \\ & a_1^1 & a_2^1 & a_3^1 \\ R_2 \quad a_1^2 & 0 & 4 & 0 \\ a_2^2 & 5 & 3 & 3 \\ a_3^2 & 2 & 4 & 0 \end{array}$$

That is not all, though, because R_1 realizes that robot R_2 possibly does not know about the observation point P2 due to the trees located between R_2 and P2.[2] R_1, therefore, has to deal with another source of uncertainty: there are two alternative models of R_2's decision-making situation. If R_2 is unaware of P2, then it will not consider combinations of actions involving a_2^1 or a_2^2, and its payoff matrix is 2 × 2, as follows:

$$\begin{array}{ccc} & R_1 & \\ & a_1^1 & a_3^1 \\ R_2 \quad a_1^2 & 0 & 0 \\ a_3^2 & 2 & 0 \end{array}$$

A sensible way for R_1 to deal with its uncertainty as to which of the models of R_2 is correct is to represent its subjective belief as to their likelihood of being correct as a probability distribution. In this case, we assume that R_1, having knowledge about the sensors available to R_2 and assessing the density of the foliage between R_2 and P2, assigns a probability for R_2 seeing through the trees as 0.1.

Let us note that R_2's best choice of action, in each of the intentional models that R_1 has, also depends on what it, in turn, thinks that R_1 will do. Thus, R_1 should, in each of these models, represent how R_2 might model R_1. If it were to model R_1 as rational as well, the nesting of models would continue. For now, however, let us assume that R_1 has no knowledge about how it might be modeled by R_2. As we discuss more fully in [8], the proper representation of R_1's lack of knowledge about how it might be modeled by R_2 is to use a uniform probability distribution over R_1's set of actions on this level. Thus, if R_1's model of R_2 indicates that R_2 thinks R_1 could undertake any one of alternative actions a_1^1, a_2^1, and a_3^1, the uniform distribution is $[\frac{1}{3}, \frac{1}{3}, \frac{1}{3}]$. This distribution, according to the principle of entropy maximization [14], precisely

[2] R_2 not knowing about P2 assumes that R_2 does not have a complete and accurate map of the terrain, or if it does then it cannot locate its own position on that map.

represents the lack of information on the part of the robot R_1 in this case, since its informational content is zero.

In the above example, robot R_1's knowledge as to the decision-making situation that it faces can be represented graphically as depicted in Figure 2, which we will call the recursive model structure. The top level of this structure is how R_1 sees its own decision-making situation, represented as R_1's payoff matrix. On the second level are the alternative models R_1 could form of R_2, with the alternative branches labeled with the probabilities R_1 assigns to each of the models being correct. The third level is occupied by uniform probability distributions representing R_1's lack of knowledge of how it is being modeled by R_2.

Solving the Example

As has been summarized elsewhere [9], the recursive model generated by RMM can be solved so as to determine a rational action for the agent represented at the root of the tree, based on expected rational actions taken on the part of agents modeled deeper in the tree. Thus, the solution method proceeds from the leaves upward. In the example interaction show in Figure 2, consider the left branch. R_1 knows nothing about R_2's model of R_1, and thus can only model R_2 as believing R_1 is equally likely to take either of the 2 actions R_2 knows about (recall, this branch corresponds to the case where R_2 does not know about P2). R_1 will believe that R_2, under these circumstances, would take action a_3^2 (staying still), with an expected payoff of 1, which is better than the expected payoff of 0. On the other hand, if R_2 were aware of P2 (right branch), R_1 will believe it will prefer action a_2^2 (going to P2) as having the highest expected payoff ($\frac{11}{3}$). Since R_1 believes there is only a probability of .1 that R_2 knows about P2, it attributes an expected strategy of (0 .1 .9) to R_2. Using this at the root of the recursive tree, R_1 will prefer action a_2^1: it will move to P2 (with an expected payoff of 2).

Utility of Communication

Since any nontrivial message is bound to transform the knowledge encoded in the hierarchy of nested models—the transformation corresponding to the *pragmatic meaning* of the message, it is natural to define the utility of a message for a sending agent as the difference in its expected utility resulting from the state of knowledge before and after the message is sent:

$$U(M) = U_{p_M}(Y) - U_p(X) \qquad (1)$$

where $U_{p_M}(Y)$ is the expected utility of the agent's best action Y after sending message M, and $U_p(X)$ is the expected utility of the agent's best action X before. p_M encapsulates the knowledge of other agents' behavior after message M has been sent.

For example, consider what would happen in our original scenario in Figure 1 if R_1 were to send a message M_1 stating "There is an observation point P2, twice as high as P1, behind the trees". If we assume that communication channels never lose messages (but see [10]) and messages are always believed (see [7], then R_1 can be sure that R_2 will know about the point P2 as a result of the message having been sent. Thus, the recursive model structure will change due to the pragmatic meaning of M_1, as depicted in Figure 3.

Before the message was sent, R_1's best option was a_2^1, that is, to observe from point P2 and expect a payoff of 2. The new structure that reflects how R_1's knowledge would look after sending M_1 can be solved easily, but now R_1 would be sure that R_2 would observe from point P2, taking action a_2^2: $p_{R_2}^{M_1;R_1} = [0,1,0]$. The best alternative for R_1 now is to make an observation from P1, but the expected payoff has increased to 5! Thus, by sending the message M_1 to R_2, R_1 was able to increase the expected utility it gets from the interaction from 2 to 5. As defined in Equation 1, the utility of sending the message M_1 is $U(M_1) = 5 - 2 = 3$.

Communicative Actions Embedded among Physical Interactions

The previous section reviewed how communicative actions can increase utility because of how they are expected to change the physical actions of other agents. Thus, the utility of a communicative action is only realized assuming the subsequent physical actions are carried out. More generally, physical and communicative interactions are interleaved to much greater extents, allowing physical actions to be conditional upon other physical actions, and verbal commitments to physical actions.

Consider, for example, the case of a threat [11]. An important issue involved in the treatment of threats is that they change the character of our prototypical interactions in a way not considered before. The threatening agent must *wait* for the action of the threatened agent, and take its own action accordingly. Thus, unlike in the previous case above (and the other cases considered in [9, 10]), the actions of the players cannot be considered as simultaneous.

Threats will be assumed to have the form "If you do A, then I will do B," where A is an option of an opponent and B is an option available to the threatening agent. One of the subtleties of threats, as discussed for example in [16], is that they often seem to involve a degree of irrationality on the part of the threatening side. Thus, typically, action B that an agent is threatening to perform is *not* the optimal response to his opponent's action A (if it were, the "threatened action B" would simply be the rational response that the opponent would expect, even without having been threatened).

For example, consider a minor variation of the scenario of Figure 1, where the only difference is that P2 now has a worth of 4 only to R_1 (it is worth 0 to R_2). In this case, and assuming R_2 is made aware of P2, R_1

			R_2	
		a^2_1	a^2_2	a^2_3
	a^1_1	1	5	1
R_1	a^1_2	4	2	2
	a^1_3	2	4	0

$p = 0.9$ $\qquad$ $p = 0.1$

			R_1
		a^1_1	a^1_3
R_2	a^2_1	0	0
	a^2_3	2	0

[0.5, 0.5]

			R_1	
		a^1_1	a^1_2	a^1_3
	a^2_1	0	4	0
R_2	a^2_2	5	3	3
	a^2_3	2	4	0

[1/3, 1/3, 1/3]

Figure 2: Recursive Model Structure depicting R_1's Knowledge in Example 1.

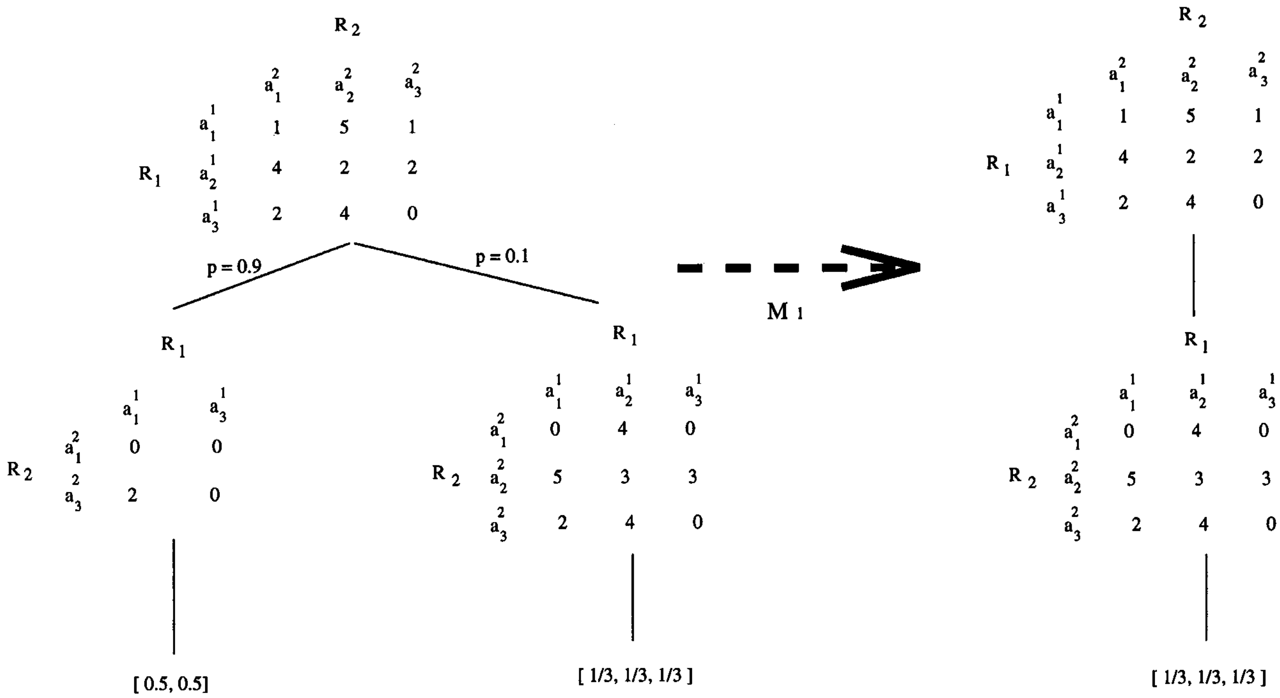

Figure 3: Expected Change of R_1's Knowledge due to Message M_1.

still wants R_2 to go to P2 while it goes to P1 for a net payoff of 5. However, if R_2 stays put, it expects to do better (2) than if it incurs the cost of going to P2. But if R_2 stays put, then R_1 expects 1 if it goes to P1, or 2 if it goes to P2, neither of which is very good.

R_1 thus needs to use communication to provide incentive to R_2 to go to P2. If payoffs are transferable, it could make a deal, hinging a physical action of transfering some utility to the physical action of R_2 going to P2. R_1 could transfer 1 unit of utility to R_2 so that R_2 would come out even, and R_1 would still be better off. Of course, R_2 could hold out for more. This kind of negotiation has been looked at most recently by Kraus [13].

If utility is not transferable, or if R_1 is more greedy (and willing to take some risk), then R_1 could instead threaten R_2 with not going to P1 unless R_2 goes to P2. This could be an effective strategy, but only if R_2 takes the threat seriously. Issues of how threats can be made more believable have been covered elsewhere [11]. For the purposes of this paper, the important consideration is that the utility of a threatening message is no longer based solely on the immediate transformation of the RMM hierarchy, but instead generally involves physical actions to make threats believable and carry them out, embedding communication within ongoing physical activity.

Communicative Actions Embedded within Dialogues

The issues of computing the utilities of questions and imperatives are more difficult than those of message types considered above, since they also touch on the agents' autonomy. The basic question here is: why should a fully autonomous being pay any attention to other agents' orders (in the case of imperatives) or requests for information (in the case of questions). As we have shown above, the computation of the utilities of messages and actions is performed exclusively from the point of view of a given agent, and the fact that another agent would value a certain physical or communicative action does not enter into these calculations. To understand question-answering and order-taking behavior on the part of an agent, therefore, we have to understand why it would be in the agent's self-interest to do so.

Questions as Mitigating the Cost of Communication

To develop this understanding in the context of questions, we must first recognize that, so far, we have treated communicative actions as being without cost. Once these actions have costs, then an agent must balance the benefits of communicating with the costs. For example, let us consider the scenario depicted in the first example interaction , but let us modify it slightly such that now the stand of trees between R_2 and P2 consists of only a single tree. Assuming that R_1 now expects R_2 to be extremely likely to see P2 (say, with probability .99), it computes R_2's intentional probabilities as:

$$.01 \times [0,0,1] + .99 \times [0,1,0] = [0,.99,.01].$$

In turn, this gives R_1 expected utilities of 4.96 for a_1^1, 2 for a_2^1, and 3.96 for a_3^1.

As detailed in [8], if R_1 were to send R_2 a message to ensure that R_2 knew about P2, then R_1 could increase its expected utility to 5 (assuming correct message transmission). Thus, in this case, the message would increase the expected utility by $5.0 - 4.96 = 0.04$. Assuming that it costs nothing to send a message, sending this message would be rational. But let us assume that sending a message costs something, say 0.1. Now the utility of the message minus the cost of the communicative action is $0.04 - 0.1 = -0.06$, so sending the message would be irrational.

However, now imagine that R_1 receives from R_2 a question message, M_7, asking "Are there any observation points in the region on the other side of the trees from me?" The immediate pragmatics of this new knowledge is to cause R_1 to transform its recursive model structure, leading it to now believe that R_2 only knows about P1 and will have intentional probabilities [0, 0, 1], meaning that R_2 will sit still. In turn, now R_1 expects utilities of 1 for a_1^1, 2 for a_2^1, and 0 for a_3^1. Its best action now is a_2^1, with an expected payoff of 2; the message it received caused it to revise its expected utility downward, from 4.96 to 2!

But now R_1 should reconsider sending the message about P2 to R_2, in effect answering the question. As before, successfully sending the message leads R_1 to expect a payoff of 5, which is much better than its current expectation of 2. The utility of the message minus the cost of sending it is $3 - 0.1 = 2.9$, and R_1, being rational, will answer R_2's question.

Now that we have established how answering a question can be rational even for an autonomous agent, let us turn to asking the question in the first place. Looking at the situation from R_2's perspective, all it sees in its environment are R_1, P1, and trees. Based on prior knowledge (for example, that observation points commonly come in pairs), R_2 might hypothesize that with probability 0.4 there is another observation point hidden behind the trees. If it assumes that this observation point will be worth 2 like P1,[3] and that it will cost 1 to get there, the expected utility for R_2 to go toward the hoped-for observation point is the probability the point is there times the worth of the point, minus the cost of going there: $(0.4\times 2)-1 = -0.2$. This negative expected utility means that it would be irrational for R_2 to act on the hunch that another observation point might be behind the tree.

But rather than taking physical action, R_2 can consider communicative action. When R_2 considers sending R_1 question M_7, "Are there any other observation points on the other side of the trees from me?" it computes the expected utility as follows. It believes that there is a 0.4 probability that it will receive an affirmative answer (using the prior knowledge above), and in that case it goes to the point to gain a payoff of 1 (since the expected worth is 2 and the expected cost of going to the point is 1). With 0.6 probability, it will receive a negative answer, and will stay still, gaining no additional payoff from its actions. Since it is currently expecting to gain nothing from its actions, the expected utility of asking the question is the expected improvement to its payoff (0.4×1) minus the cost of sending the message (0.1). Asking the question thus has a utility of 0.3, and it is rational for R_2 to ask.

Note that this calculation looked beyond the immediate pragmatics of the message (how it would transform the hearer's model of the situation) to consider how the hearer's possible subsequent (communicative) action might transform the original speaker's model of the situation. Thus, while both the question and the answer separately have immediate pragmatic effects like those of simple assertions about the situation (about what an agent knows or about observation points in the world), it is the longer sequence of them that induces the ultimately desired effect (increasing the questioner's awareness). This leads to considering the pragmatics, and utility, of messages in terms of the possible sequences of communicative and physical actions surrounding them. Section is a step toward capturing this notion.

Of course, the analysis done above has R_2 assume that R_1 will correctly interpret and truthfully respond to the message. Before sending the question, R_2 can first model R_1 to model the possible ways that the message might transform R_1's recursive model structure, and then decide from those whether it can expect R_1 to truthfully answer. In a manner similar to the above analysis about answering questions, R_2 will conclude that, if there is another worthwhile observation point behind the tree, it will be rational for R_1 to respond truthfully. If there is no observation point behind the tree, saying this will not benefit R_1, and in fact will cost it 0.1 (the cost of sending a message). Thus, if there is no observation point, R_1 will never respond. With this analysis, R_2 will conclude that it is rational to ask the question, to believe R_1 if it responds, and to stay still (which is what it would have done anyway) otherwise.[4]

[3] Note that, in reality, this is an underestimate.

[4] Of course, if R_1 were to value sending R_2 on a wild goose chase, and R_2 did not know about this propensity, then R_1 could successfully lie to R_2. For further investigation into honesty and trust among rational agents, the reader is referred to [7, 19, 20].

Imperatives as Mitigating the Costs of Computation

To understand imperatives requires a similar analysis, except that now we also have to take into account the costs (time, effort) of decision-making. In a nutshell, it is rational for an agent to obey an order from another if the default utility of obeying is greater than the (possibly larger) utility of independent decision-making minus the cost of that decision-making. At an intuitive level, when I hear someone shout "Duck!" it might be better for me to rapidly follow this instruction than respond more slowly (too late) after assessing the situation and deciding for myself on the proper action.

In its simplest form, an imperative transforms an agent's recursive nesting of models from which it derives its strategy into simply a strategy to follow. That is, the agent discards the deeper machinery and simply obeys the command. Because the decision to do so involves a tradeoff of costs and benefits of using the machinery, an imperative causes reasoning at the meta-level. Since RMM has not yet addressed meta-level reasoning issues to any great extent, the decision as to when to follow the command cannot at this time be reduced to operational terms. Clearly, the deciding factors will involve comparing the expected payoff of following the command (what does the command imply about the likely decisions of others, and what kind of payoffs are likely to be received based on these implications) against the expected payoff of deeper reasoning (what better decisions might be made, how much better might they be, and how much will it cost to find them). Our ongoing research is delving into meta-level issues, so as to eventually capture such reasoning about imperatives.

Of course, an agent can, upon receiving a command, consider what actions others will take if they assume the command will be obeyed, and then the agent can decide what it actually should do (which might or might not be the same as the commanded actions). Of course, if it thinks that other agents might believe that it might go through this reasoning, then it should model those other agents as being uncertain about whether it will follow the command or not, and the modeling gets much more complicated. This kind of modeling resembles agents reasoning about whether to believe what they hear [7].

There is another aspect of issuing orders, and obeying them, among autonomous agents that is closely related to mitigating the cost of independent decision-making. This is the notion that the agent issuing the command has more knowledge about the environment in the first place. The modeling of other agents' knowledge using a recursive knowledge-theoretic formulation has been investigated in [6]. There, the agents' own knowledge about the world is represented as a Kripke structure, with possible worlds and the agent's inability to tell them apart corresponding to the agents' uncertainty. Further, we have postulated that this representation be made recursive, so that the agent can model the state of knowledge of another agent in a way similar to RMM. If the agent models another using a structure with fewer possible worlds, this indicates that the agent believes that the other is less uncertain about the current state of affairs, and so might believe that the other agent is more likely correct in its orders than this agent is in its own choices. This particular aspect of rational behavior with respect to imperatives lets us observe that an interesting paradigm "knowledge is power" tends to naturally emerge from our formalism.

Arguments Based on Iteration

Finally, questions and imperatives might also arise not only due to specific interactions, but based on the game-theoretic paradigm of cooperation examined by Aumann in [1, 2], where it is shown how repetitive interactions allow cooperation to emerge as a rational choice of selfish interacting agents. We think it is safe to extend these conclusions to cases of questions and imperatives in communication, since obeying imperatives and answering questions can be seen as a form of cooperative behavior. Thus, even when it might be true that obeying imperatives and answering questions may be irrational in a one-time interaction, they may be beneficial overall if one adheres to the "you scratch my back now, I will scratch yours later" maxim, which is a generalized version of a "tit–for–tat" paradigm showed successful for the repeated Prisoner's Dilemma [3].

Towards the Emergence of Protocols

As the RMM approach is extended to anticipate even more prolonged dialogues among agents, agents will be able to search through possible courses of communicative exchanges to identify the most promising candidates. However, as the number of possible messages grows (increasing the branching factor) and the number of exchanges considered is enlarged (increasing the depth), the combinatorics of the search will make the process prohibitively expensive, at least if employed on most occasions. If, once again, the costs of computation are factored into the utility computations, it might be better sometimes to engage in a suboptimal dialogue than to identify an optimal one. And, moreover, once a good dialogue has been found, the computational efforts should be reused as much as possible by either retrieving and revising an appropriate dialogue from memory when similar situations arise in the future, or by generalizing a dialogue into a more generic pattern of communication.

This gives rise, naturally, to the establishment of protocols. The combinatorics of the search can be greatly reduced by restricting the types of messages considered at each stage (reducing the branching factor) and by clustering collections of communications together into "standard" interactions (question-answer, or announce-bid-award, for example). As clusters grow and are reapplied to different (but similar) situations, they become stored plans of communication, which means that they begin to resemble protocols. Thus, while substantial research is needed to adequately operationalize this process of going from computing utilities of individual actions based on transforming a nested modeling structure all the way up to representing generalized protocols for guiding dialogues, the research path looks both important (to build systems that can survive and establish new protocols when known protocols fail) and promising (since we have, in this paper, described how to derive primitives for such protocols, such as query-response, within this framework).

References

[1] Robert J. Aumann and Sergiu Hart. Repetition as a paradigm for cooperation in games of incomplete information. Technical report, Hebrew University, 1981.

[2] Robert J. Aumann and Sylvian Sorin. Cooperation and bounded recall. *Games and Economic Behavior*, 1:5–39, 1989.

[3] Robert Axelrod. *The Evolution of Cooperation.* Basic Books, 1984.

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

[5] Edmund H. Durfee and Victor R. Lesser. Partial global planning: A coordination framework for distributed hypothesis formation. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(5):1167–1183, September 1991. (Special Issue on Distributed Sensor Networks).

[6] Piotr J. Gmytrasiewicz and Edmund H. Durfee. A logic of knowledge and belief for recursive modeling: A preliminary report. In *Proceedings of the Tenth National Conference on Artificial Intelligence*, July 1992.

[7] Piotr J. Gmytrasiewicz and Edmund H. Durfee. Toward a theory of honesty and trust among communicating autonomous agents. *Group Decision and Negotiation*, 2:237–258, 1993.

[8] Piotr J. Gmytrasiewicz and Edmund H. Durfee. Rational coordination and communication in multiagent environments through recursive modeling. *Submitted for Publication*, 1994.

[9] 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.

[10] Piotr J. Gmytrasiewicz, Edmund H. Durfee, and David 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.

[11] Piotr J. Gmytrasiewicz and Jeffrey S. Rosenschein. The utility of embedded knowledge-oriented actions. In *Proceedings of the 1993 Distributed AI Workshop*, Hidden Valley, PA, May 1993.

[12] Craig A. Knoblock. *Generating Abstraction Hierarchies: An automated approach to reducing search in planning.* Kluwer Academic Publishers, 1993.

[13] Sarit Kraus and Jonathan Wilkenfeld. The function of time in cooperative negotiations. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, August 1991.

[14] Richard E. Neapolitan. *Probabilistic Reasoning in Expert Systems.* John Wiley and Sons, 1990.

[15] Ramesh S. Patil, Richard E. Fikes, Peter F. Patel-Schneider, Don Mckay, Tim Finin, Thomas Gruber, and Robert Neches. The DARPA knowledge sharing effort: Progress report. In *Proceedings of the Eighth IEEE Conference on AI Applications*, March 1992.

[16] Thomas C. Schelling. *The Strategy of Conflict.* Havard University Press, 1960.

[17] Reid G. Smith. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers*, C-29(12):1104–1113, December 1980.

[18] Gilad Zlotkin and Jeffrey S. Rosenschein. Cooperation and conflict resolution via negotiation among autonomous agents in non-cooperative domains. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(6), December 1991. (Special Issue on Distributed AI).

[19] Gilad Zlotkin and Jeffrey S. Rosenschein. Incomplete information and deception in multi-agent negotiation. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, pages 225–231, August 1991.

[20] Gilad Zlotkin and Jeffrey S. Rosenschein. A domain theory for task oriented negotiation. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, pages 416–422, August 1993.