Multi-Agent Resource Sensitive Communication

Pamela W. Jordan
University of Pittsburgh ISP
901CL Pittsburgh, PA 15260
jordan@isp.pitt.edu

Marilyn A. Walker
Mitsubishi Electric Research Laboratories
201 Broadway Cambridge, MA 02139
walker@merl.com

Abstract

In order for a model of rational agency to be used for multi-agent social interaction, the model must be extended to account for communicative behaviors that are inefficient for one agent, but which increase the efficiency of the interaction of two agents in a dialogue. In this paper, we argue that naturally occurring communicative behaviors are often directed at the resource limits of the other agent and that communicative behaviors that are effective for resource unbounded agents are not suited for resource limited agents. We describe a set of experiments on communicative behaviors in a multi-agent simulation testbed and explore the efficacy of various information sources on which agents can base their decisions about what and when to communicate with other agents. In particular we argue that a model of what the other agent is attending to makes a significant contribution to communicative efficiency.

Introduction

A theory of multi-agent interaction requires a model of communication that guides an agent's decisions about what information to communicate to a hearer and how and when to communicate it. Consider the following short natural dialogue, part of a discussion about which route to take:

(1) a. Let's walk down Walnut St.

b. It's shorter.

Agent A says (1b) in order to motivate the hearer, agent B, to accept the proposal in (1a); we call (1b) a WARRANT for the proposal. In the context of this dialogue, A said (1b) despite the fact that it was already known to B. Why didn't A simply say (1a)?

We claim that A included (1b) because she believed that although B **knew** the warrant, B was not **attending** to the warrant. We will argue that a model of multi-agent interaction through communication should include a model of attentional state. We shall use the term SALIENT to describe the subset of an agent's beliefs and intentions that are currently attended to.[1]

Contrary to what we see in (1), a common assumption of theories of rational communicative behavior is that A should not say information that B already knows (Allen 1983; Grice 1967). We call this the REDUNDANCY CONSTRAINT. An analysis of naturally occurring dialogues showed that this constraint is frequently violated (Walker 1993). In our view the problem with the REDUNDANCY CONSTRAINT is that it does not distinguish between what B knows and what is salient for B, and many processes that underlie the interpretation of communicative acts, such as inference and deliberation, operate most efficiently and most deterministically on beliefs that are salient (Baddeley 1986; Walker 1993). Thus, even when it is **possible** for agent B to make an inference from, or deliberate using, a belief that is not salient, it may be more efficient for agent A to state the belief, guaranteeing its salience, and thus making it **easier** for B to carry out these processes (Norman & Bobrow 1975).

We call the communicative strategy in which A makes salient a warrant that is already known, the Explicit-Warrant strategy. In previous work, we showed that the Explicit-Warrant strategy is an efficient strategy for certain types of resource limited agents. In particular, the Explicit-Warrant strategy is more efficient when the effort involved in retrieval from memory dominates other processing effort, but it is not efficient when communication effort dominates, except when the task requires agents to be coordinated on warrants underlying intended actions (Walker 1994). In addition, *pace* economic models of information value, the Explicit-Warrant strategy may be **detrimental** for agents that are attention limited because it changes their attentional state, and this can detrimentally affect their performance on a task (Walker 1994).

In the remainder of the paper, we describe our extension of the IRMA agent architecture to agents that communicate with one another under assumptions of limited attentional capacity. We test our extended agent model by implementing it in the Design-World

[1] Salient is a cognitive term, but it is also possible to view facts that are currently stored in a cache as salient facts. The critical assumption is that the salient set is a subset of what is known.

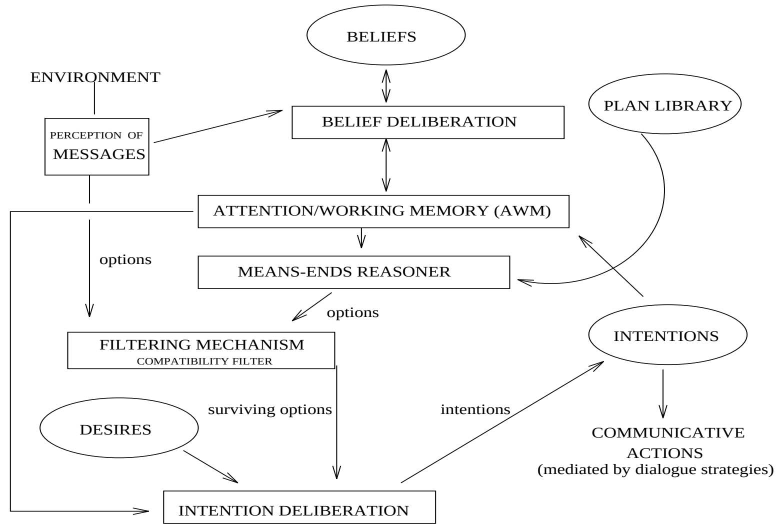

Figure 1: Design-World version of the IRMA Agent Architecture for Resource-Bounded Agents with Limited Attention (AWM)

dialogue testbed (Bratman, Israel, & Pollack 1988; Pollack & Ringuette 1990; Walker 1994). We then explain how we used this model to simulate a multi-agent interaction to test the effectiveness of a variety of communicative decision algorithms.

Since we claim the decision algorithm needs an attentional model of B in order to determine whether a warrant is salient for B, we vary the source for the model of B's attentional state as well as the decision algorithm to arrive at different communication strategies. In particular, we propose a new strategy, Self, in which A's own attentional state is used as a model for B's attentional state. We will compare the Self strategy with four other strategies based on other ways of approximating B's attentional state.

The Design-World Testbed

Design-World is a testbed that supports experiments on the interaction of agents' resource limits, their communicative strategies and the complexity of the task. The agent architecture used in the Design-World testbed is based on the IRMA architecture for resource-bounded agents, shown in figure 1 (Bratman, Israel, & Pollack 1988; Pollack & Ringuette 1990).

Because attentional state is an important factor in communicative effectiveness, the architecture has been extended with a model of limited Attention/Working Memory (AWM) with properties that model human cognition: there is a higher probability that beliefs stored in memory are salient and thus accessible when (1) they have been accessed recently (2) they have been accessed frequently (Landauer 1975). In the presentation, we will show how Landauer's model, when appropriately parameterized, fits the psychological data on human memory and learning.

Furthermore, processes such as reasoning and deliberation are limited to the subset of SALIENT beliefs and intentions currently in AWM as shown in figure 1[2]. The size of the subset is determined by a parameter to be described below. Figure 1 also shows that in addition to options generated by reasoning from salient beliefs, salient options can also be generated by communication.

The Design-World task requires two agents to carry out a dialogue in order to negotiate an agreement on the design of the floor plan of a two room house (Whittaker, Geelhoed, & Robinson 1993). At the beginning of the simulation, both agents know the structure of the DESIGN-HOUSE plan[3]. Each agent has 12 items of furniture that can be used in the plan, and items used in the plan are assigned a utility ranging from 10 to 56. Assigning utility serves two functions: (1) it is the basis for the agent's deliberation about which plan options are better; and (2) it provides the basis for an objective performance evaluation metric of the agents' communicative behaviors. Both the 12 allocated furniture items and the utility information for all 24 furniture items are stored as the agent's beliefs. Due to the way AWM is modelled, the salience of this information varies according to the recency and frequency with which it is accessed.

The agents' goal is to select 4 pieces of furniture for each room so that the maximum utility is achieved. Subgoals such as color matches within or across rooms can be introduced where meeting these subgoals adds

[2] AWM is a processing mechanism as well as a data structure in that it selects which subset of an agent's beliefs and intentions are salient

[3] The plan library bypasses AWM so this knowledge is always accessible.

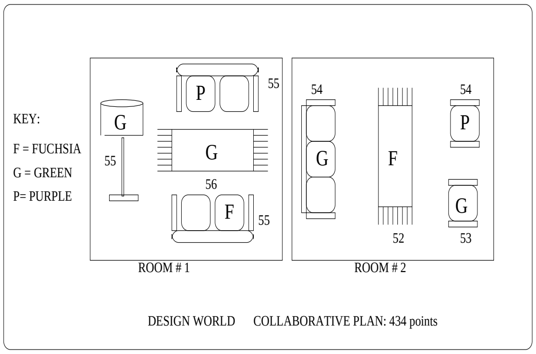

Figure 2: One Final State for Design-World Standard Task: Represents the Collaborative Plan Achieved by the Dialogue, 434 points

extra utility to the design. However, for our testing of the decision algorithm, no subgoals were added to the task. The simulation supports parameters that reflect different assumptions about agents' resource limits when carrying out this task so that different communicative strategies have differential effects on their performance.

In order to explore a range of resource capacities, both the processing effort to access AWM and the size of the subset representing what is salient within AWM are parameters of the model. In our tests, the AWM subset size parameter varies between 1 and 16, where agents with a subset size of 1 are severely limited in attention[4] and a size of 16 means that everything known is salient. While we have not tested the performance of humans on this task, we would expect their performance to be somewhere in the middle of this scale. Since the AWM parameter limits what an agent is attending to, an agent may not be able to recall its highest utility pieces at any given time. Figure 2 shows a potential final design plan negotiated via a (simulated) dialogue.

An Example of Multi-Agent Interaction

In Design-World, negotiating an agreement between two agents consists of a cycle in which: (1) individual agents perform means-end reasoning about OPTIONS to use various furniture pieces that they have; (2) individual agents deliberate about which options are preferable based on the utility of using the item in the plan. See figure 1. Then the agents realize their intentions through DISCOURSE ACTS such as PROPOSALS, ACCEPTANCES, REJECTIONS and CLARIFICATIONS: (3) agents use the preferred options to make PROPOSALS to other agents to PUT a piece of furniture into a room; (4) then the other agents ACCEPT, REJECT, or request CLARIFICATION of these proposals.

[4]An agent with AWM size parameter of 1 can access about 7 beliefs and intentions.

This is illustrated by the dialogue excerpt in 2. Both the artificial language that the agents communicate with and a gloss generated from that language in *italics* are shown:[5]

(2) 1: BILL: *Let's put the green rug in the study.*
(propose agent-bill agent-kim option-10: put-act (agent-bill green rug room-1))

2: KIM: *Then, let's put the green lamp there.*
(propose agent-kim agent-bill option-33: put-act (agent-kim green lamp room-1))

3: BILL: *Next, let's put the green couch in the study.*
(propose agent-bill agent-kim option-45: put-act (agent-bill green couch room-1))

4: KIM: *No, instead let's put the purple couch in the study.*
(reject agent-kim agent-bill option-56: put-act (agent-kim purple couch room-1))

In this dialogue, when agent Kim receives agent Bill's proposal, she evaluates it and decides whether to accept or reject it. During this evaluation, she will attempt to retrieve the utility propositions stored earlier in memory. Thus the propositions about the utility of furniture items are WARRANTS for supporting deliberation similar to the naturally occurring (1b). Proposals, such as (2-1) and (2-2), are inferred to be implicitly ACCEPTED because they are not rejected(Walker & Whittaker 1990). If a proposal is ACCEPTED, either implicitly or explicitly, then the option contained in the proposal becomes a mutual intention that contributes to the final design plan(Power 1984; Sidner 1994). Agents REJECT a proposal if deliberation produces a better option, as with (2-4).

A discourse strategy is a particular way of achieving a DISCOURSE ACT such as a PROPOSAL. In the experiments we will discuss in our presentation, agents are parametrized for different discourse strategies by placing different expansions of discourse plans in their plan libraries. We examine five strategies which vary the expansions of PROPOSALS: (1) All-Implicit (2) Explicit-Warrant (3) Self (4) Oracle (5) Random.

The All-Implicit strategy is an expansion of a discourse plan to make a PROPOSAL, and decomposes trivially to the communicative act PROPOSE. In dialogue (2), both Design-World agents communicate using the All-Implicit strategy, and proposals expand to the PROPOSE communicative acts shown in (2-1), (2-2), and (2-3). The All-Implicit strategy **never** includes warrants in proposals, leaving it up to the other agent to retrieve them from memory. An agent utilizing this strategy acts as though it assumes that everything the other agent knows is salient.

The Explicit-Warrant strategy expands the PROPOSAL discourse act to be a WARRANT followed by a PROPOSE utterance (Suthers 1993). For example (3-1a) is Ted's WARRANT for his proposal in (3-1b):

[5]The generation of the gloss was not a focus of this study and was done automatically but by adhoc methods.

(3) 1a: TED: **Putting in the green rug is worth 56.**
1b: TED: *So, let's put the green rug in the study.*
2a: BEN: **Putting in the green lamp is worth 55.**
2b: BEN: *So, let's put the green lamp in the study.*

An agent utilizing this strategy **always** includes the warrant. It acts as though it assumes that nothing is salient for the other agent. The other three strategies are all variations of the Explicit-Warrant strategy, that vary what information is used as a model of B's attentional state. All of these strategies model the naturally occurring example in (1), but vary the conditions under which an agent makes salient a warrant that is already known to the other agent.

Recall that an agent using the Self strategy uses its own attentional state as a model for the other agent. Self is plausible because agents in conversation appear to expend a great deal of effort to stay coordinated(Thomason 1990), and because A is not always in a position to evaluate what B is currently attending to. The Self strategy has the added advantage for resource bounded agents of not requiring the processing involved with maintaining a separate model of B's attentional state. While in our simulation the attentional states of Agents A and B will **never** be identical due to their different initial information and due to the probabilistic nature of AWM, we will show that A's attentional state is still a plausible model for B's attentional state.

In the Oracle strategy, A has access to B's actual attentional state. Oracle provides data on how a perfect attentional model of the other agent would affect the performance of the decision algorithm.

Finally the Random strategy randomly says the warrant (probability = 0.5). This strategy is included as a baseline comparison in order to test whether the extra effort of consulting an attentional model is better than random.

Performance Evaluation

To compare strategies, we need to evaluate PERFORMANCE. We assume that agents are working as a team and are attempting to minimize COLLABORATIVE EFFORT. We assume that collaborative effort consists of the combined effort involved in sending and understanding messages, in accessing facts in memory, and in reasoning and deliberation (inferencing). In order to make the model independent of the implementation, we require parameters of: (1) COMMCOST: cost of sending a message; (2) INFCOST: cost of inference; and (3) RETCOST: cost of retrieval from memory. Then

COLLABORATIVE EFFORT =
(COMMCOST × total messages for both agents)
+ (INFCOST × total inferences for both agents)
+ (RETCOST × total retrievals for both agents)

Although in the results presented here we will assume that communication cost dominates all other processing costs and that retrieval cost is not free, our results hold across a range of settings.

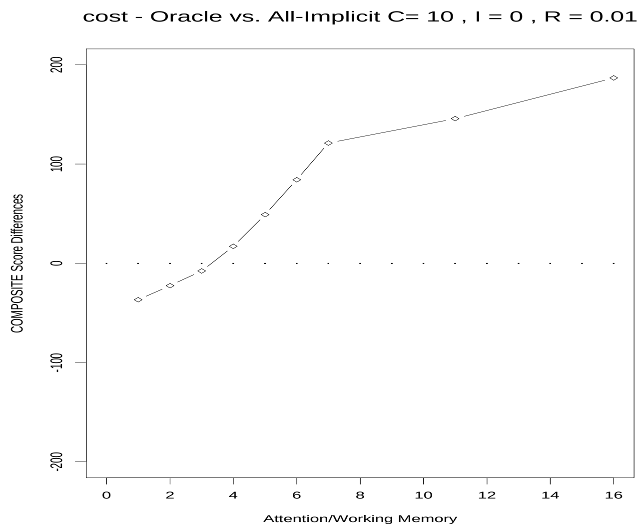

Figure 3: The Oracle communication strategy (a perfect attentional model) is Better than All-Implicit under assumptions that communication costs dominate other costs

Performance is the difference between the RAW SCORE for the task and COLLABORATIVE EFFORT. RAW SCORE is task specific: here we simply sum the utility of the furniture pieces in each PUT-ACT in the final design.

PERFORMANCE = Task Defined RAW SCORE
− COLLABORATIVE EFFORT.

An experiment consists of simulating 100 dialogues at each parameter setting for each strategy, and comparing the resulting performance distributions using the Kolmogorov-Smirnov (KS) two sample test (Siegel 1956). When comparing two strategies for a set of fixed parameter settings, a strategy is BENEFICIAL if the difference in distributions using the KS two sample test is significant at $p < .05$, in the positive direction, for two or more AWM settings.

We present our results in plots of the **differences** in performance of agents using different communicative strategies, such as the plot in figure 3. **Differences** in performance between two strategies are plotted on the Y-axis against the complete range of AWM size parameter settings on the X-axis. Each point in the plot represents the difference in the means of 100 runs of each strategy at a particular AWM size setting. These plots summarize the information from 18 performance distributions (1800 simulated dialogues).

Figure 3 compares the Oracle strategy with the All-Implicit strategy. The Oracle strategy is superior to the terser All-Implicit strategy at AWM settings of 6 to 11, even when communication cost dominates other

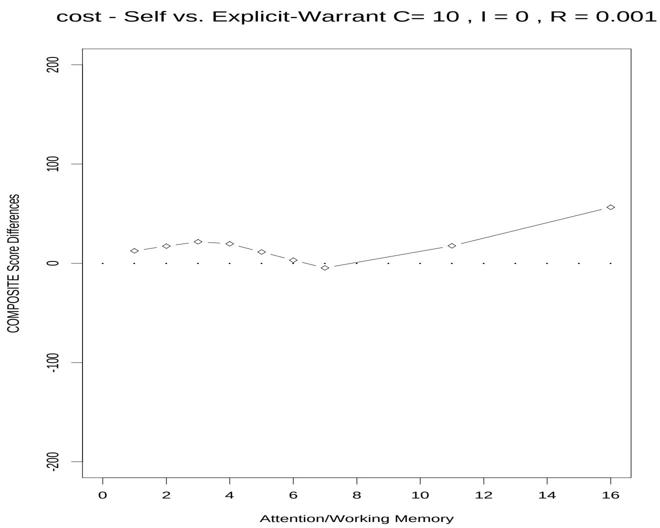

Figure 4: The Self communicative strategy is Better than Explicit-Warrant under assumptions that communication costs dominate other costs

processing costs. Figure 3 shows that communication efficiency improves if agent A has a perfect model of agent B's attentional state. This was also true when Oracle was compared to Explicit-Warrant.

Figure 4 plots the differences in performance of agents using the Self strategy vs. those using the Explicit-Warrant strategy. The figure shows that Self performs significantly better at AWM settings of 2–4 and 11–16, under the assumption that communication cost dominates other processing costs. Overall, our other results[6] show that Self performs better than the verbose Explicit-Warrant and no worse than the terse All-Implicit when communication costs dominate. As retrieval costs increase, Self performs better than All-Implicit and no worse than Explicit-Warrant. Thus these results show that the Self strategy is a plausible strategy that is beneficial for resource limited agents.

Although Self was not significantly better than Random, we are hypothesizing that Self would be better for tasks where more mutual knowledge is acquired through communication.

Conclusion

In this paper we have argued that models of rational agency must consider the effect of resource limits on processing communicative acts and that communicative strategies that are effective for resource unbounded agents, such as the All-Implicit strategy, are not effective for resource limited agents. We showed that a model of an agent's attentional state can make communication more effective. We presented a new strategy called Self in which a resource limited agent approximates another agent's attentional state by using its own attentional state, and showed that the Self strategy is superior to both the Explicit-Warrant and All-Implicit strategies under a wider range of processing conditions.

[6]Figures omitted due to space limitations.

References

Allen, J. F. 1983. Recognizing intentions from natural language utterances. In Brady, M., and Berwick, R., eds., *Computational Models of Discourse.* MIT Press.

Baddeley, A. 1986. *Working Memory.* Oxford University Press.

Bratman, M.; Israel, D.; and Pollack, M. 1988. Plans and resource bounded practical reasoning. *Computational Intelligence* 4:349–355.

Grice, H. P. 1967. *William James Lectures.*

Landauer, T. K. 1975. Memory without organization: Properties of a model with random storage and undirected retrieval. *Cognitive Psychology* 495–531.

Norman, D. A., and Bobrow, D. G. 1975. On data-limited and resource-limited processes. *Cognitive Psychology* 7(1):44–6.

Pollack, M., and Ringuette, M. 1990. Introducing the Tileworld: Experimentally Evaluating Agent Architectures. In *AAAI90*, 183–189.

Power, R. 1984. Mutual intention. *Journal for the Theory of Social Behaviour* 14.

Sidner, C. 1994. An artificial discourse language for collaborative negotiation. In *AAAI 94*, 814–820.

Siegel, S. 1956. *Nonparametric Statistics for the Behavioral Sciences.* McGraw Hill.

Suthers, D. D. 1993. Sequencing explanations to enhance communicative functionality. In *Proceedings of the 15th Annual Meeting of the Cognitive Science Society.*

Thomason, R. 1990. Propagating epistemic coordination through mutual defaults I. In Parikh, R., ed., *Proceedings of the Third Conference on Theoretical Aspects of Reasoning about Knowledge*, 29–39. Morgan Kaufmann.

Walker, M. A., and Whittaker, S. 1990. Mixed initiative in dialogue: An investigation into discourse segmentation. In *Proc. 28th Annual Meeting of the ACL*, 70–79.

Walker, M. A. 1993. *Informational Redundancy and Resource Bounds in Dialogue.* Ph.D. Dissertation, University of Pennsylvania.

Walker, M. A. 1994. Experimentally evaluating communicative strategies: The effect of the task. In *AAAI94.*

Whittaker, S.; Geelhoed, E.; and Robinson, E. 1993. Shared workspaces: How do they work and when are they useful? *IJMMS* 39:813–842.