COOPERATION WITHOUT COMMUNICATION

Michael R. Genesereth, Matthew L. Ginsberg, and Jeffrey S. Rosenschein*

Logic Group, Knowledge Systems Laboratory,
Computer Science Department, Stanford University,
Stanford, California 94305

ABSTRACT

Intelligent agents must be able to interact even without the benefit of communication. In this paper we examine various constraints on the actions of agents in such situations and discuss the effects of these constraints on their derived utility. In particular, we define and analyze *basic rationality*; we consider various assumptions about independence; and we demonstrate the advantages of extending the definition of rationality from individual actions to decision procedures.

I Introduction

The affairs of individual intelligent agents can seldom be treated in isolation. Their actions often interact, sometimes for better, sometimes for worse. In this paper we discuss ways in which cooperation can take place in the face of such interaction.

A. Previous work in Distributed AI

In recent years, a sub-area of artificial intelligence called distributed artificial intelligence (DAI) has arisen. Researchers have attempted to address the problems of interacting agents so as to increase efficiency (by harnessing multiple reasoners to solve problems in parallel [29]) or as necessitated by the distributed nature of the problem domain (e.g., distributed air traffic control [30]).

Smith and Davis' work on the contract net [6] produced a tentative approach to cooperation using a contract-bid metaphor to model the assignment of tasks to processors. Lesser and Corkill have made empirical analyses of distributed computation, trying to discover cooperation strategies that lead to efficient problem solutions for a network of nodes [3,4,7,21]. Georgeff has attacked the problem of assuring noninterference among distinct agents' plans [12,13]; he has made use of operating system techniques to identify and protect critical regions within plans, and has developed a general theory of action for these plans. Lansky has adapted her work on a formal, behavioral model of concurrent action towards the problems of planning in multi-agent domains [20].

These DAI efforts have made some headway in constructing cooperating systems; the field as a whole has also benefited from research into the formalisms necessary for one agent to reason about another's knowledge and beliefs. Of note are the efforts of Appelt [1], Moore [24], Konolige [19,18], Levesque [22], Halpern and Moses [8,16].

B. Their assumptions

Previous DAI work has assumed for the most part that agents are mutually cooperative through their designer's fiat; there is built-in "agent benevolence." Work has focused on how agents can cooperatively achieve their goals when there are no conflicts of interest. The agents have identical or compatible goals and freely help one another. Issues to be addressed include those of synchronization, efficient communication, and (inadvertent) destructive interference.

C. Overview of this paper

1. True conflicts of interest

The research that this paper describes discards the benevolent agent assumption. We no longer assume that there is a single designer for all of the interacting agents, nor that they will necessarily help one another. Rather, we examine the question of how high-level, autonomous, independently-motivated agents ought to interact with each other so as to achieve their goals. In a world in which we get to design only our own intelligent agent, how should it interact with other intelligent agents?

*This research has been supported by the Office of Naval Research under grant number N00014-81-K-0004 and by DARPA under grant numbers N00039-83-C-0136 and N00039-86-C-0033.

There are a number of domains in which autonomous, independently-motivated agents may be expected to interact. Two examples are resource management applications (such as an automated secretary [15]), and military applications (such as an autonomous land vehicle). These agents must represent the desires of their designers in an environment that includes other intelligent agents with potentially conflicting goals.

Our model of agent interaction thus allows for true conflicts of interest. As special cases, it includes pure conflict (i.e., zero sum) and conflict-free (i.e., common goal) encounters. By allowing conflict of interest interactions, we can address the question of why rational agents would choose to cooperate with one another, and how they might coordinate their actions so as to bring about mutually preferred outcomes.

2. No communication

Although communication is a powerful instrument for accomodating interaction (and has been examined in previous work [28]), in our analysis here we consider only situations in which communication between the agents is impossible. While this might seem overly restrictive, such situations do occur, e.g., as a result of commmunications equipment failure or in interactions between agents without a common communications protocol. Furthermore, the results are valuable in the analysis of cooperation with communication [28,27].

Despite the lack of communication, we make the strong assumption that sufficient sensory information is available for the agents to deduce at least partial information about each other's goals and rationality. For example, an autonomous land vehicle in the battlefield may perceive the actions of another autonomous land vehicle and use plan recognition techniques [9] to deduce its destination or target, even in the absence of communication.

3. Study of constraints

In this paper we examine various constraints on the actions of agents in such situations and discuss the effects of these constraints on the utility derived by agents in an interaction. For example, we show that it can be beneficial for one agent to exploit information about the rationality of another agent with which it is interacting. We show that it can also be beneficial for an agent to exploit the similarity between itself and other agents, except in certain symmetric situations where such similarity leads to indeterminate or nonoptimal action.

The study of such constraints and their consequences is important for the design of intelligent, independently motivated agents expected to interact with other agents in unforeseeable circumstances. Without such an analysis, a designer might overlook powerful principles of coooperation or might unwittingly build in interaction techniques that are nonoptimal or even inconsistent.

Section 2 of this paper provides the basic framework for our analysis. The subsequent sections analyze progressively more complicated assumptions about interactions between agents. Section 3 discusses the consequences of acting rationally and exploiting the rationality of other agents in an interaction; section 4 analyzes dependence and independence in decision making; and section 5 explores the consequences of rationality across situations. The concluding section discusses the coverage of our analysis.

II Framework

Throughout the paper we make the assumption that there are exactly two agents per interaction and exactly two actions available to each agent. This assumption substantially simplifies our analysis, while retaining the key aspects of the general case. Except where indicated to the contrary, all results hold in general [10,11,14,27].

The essence of interaction is the dependence of one agent's utility on the actions of another. We can characterize this dependence by defining the payoff for each agent i in an interaction s as a function p_i^s that maps every joint action into a real number designating the resulting utility for i. Assuming that M and N are the sets of possible moves for the two agents (respectively), we have

$$p_i^s : M \times N \rightarrow \mathbb{R}.$$

In describing specific interactions, we present the values of this function in the form of payoff matrices [23], like the one shown in figure 1. The number in the lower left hand corner of each box denotes the payoff to agent J if the agents perform the corresponding actions, and the number in the upper right hand corner denotes the payoff to K. For example, if agent J performs action a in this situation and agent K performs action c, the result will be 4 units of utility for J and 1 unit for K. Each agent is interested in maximizing its own utility.

K

		c	d
J	a	1 4	4 2
	b	3 1	2 3

Figure 1: A payoff matrix

Although the utilities present in a payoff matrix can generally take on any value, we will only need the *ordering* of outcomes in our analysis. Therefore, we will only be using the numbers 1 through 4 to denote the utility of outcomes.

An agent's job in such a situation is to decide which action to perform. We characterize the *decision procedure* for agent i as a function W_i from situations (i.e., particular interactions) to actions. If S is the set of possible interactions, we have

$$W_i : S \rightarrow M.$$

In the remainder of the paper we take the viewpoint of agent J.

III Basic Rationality

We begin our analysis by considering the consequences of constraining agent J so that it will not perform an action that is *basically irrational.* Let R_i^s denote a unary predicate over moves that is true if and only if its argument is rational for agent i in situation s. Then agent J is basically rational if its decision procedure does not generate irrational moves, i.e.,

$$\neg R_J^s(m) \Rightarrow W_J(s) \neq m.$$

W_J here is a function that designates the action performed by J in each situation, as described above. In order to use this definition to judge which actions are rational, however, we need to further define the rationality predicate R_J^s.

An action m' *dominates* an action m for agent J in situation s (written $D_J^s(m', m)$) if and only if the payoff to J of performing action m' is greater than the payoff of performing action m (the definition for agent K is analogous). The difficulty in selecting an action stems from lack of information about what the other agent will do. If such information were available, the agent could easily decide what action to perform. Let the term $A_K^s(m)$ denote the action that agent K will perform in situation s if agent J performs action m:

$$W_K(s) = A_K^s(W_J(s)).$$

In what follows we call A_K^s the *reaction function* for K. Then the formal definition of dominance is

$$D_J^s(m', m) \iff p_J^s(m', A_K^s(m')) > p_J^s(m, A_K^s(m)).$$

We can now define the rationality predicate. An action is *basically irrational* if there is another action that dominates it.

$$(\exists m' \; D_J^s(m', m)) \Rightarrow \neg R_i^s(m)$$

Even if J knows nothing about K's decision procedure, this constraint guarantees the optimality of a decision rule known as *dominance analysis.* According to this rule, an action is forbidden if there is another action that yields a higher payoff for every action of the other agent, i.e.,

$$(\exists m' \forall n \forall n' \; p_J^s(m, n) < p_J^s(m', n')) \Rightarrow W_J(s) \neq m.$$

Theorem Basic rationality implies dominance analysis.

Proof: A straightforward application of the definition of rationality. □

As an example of dominance analysis, consider the payoff matrix in figure 2. In this case, it is clearly best for J to perform action a, no matter what K does (since 4 and 3 are both better than 2 and 1). There is no way that J can get a better payoff by performing action b.

K

		c	d
J	a	4 4	3 3
	b	2 2	1 1

Figure 2: Row Dominance Problem

Of course, dominance analysis does not always apply. As an example, consider the payoff matrix in figure 3. In this situation, an intelligent agent J would probably select action a. However, the rationale for this decision requires an assumption about the rationality of the other agent in the interaction.

		K	
		c	d
J	a	4, 3	2, 1
	b	1, 4	3, 2

Figure 3: Column Dominance Problem

In dealing with another agent it is often reasonable to assume that the agent is also basically rational. The formalization of this assumption of *mutual rationality* is analogous to that for basic rationality.

$$\neg R^s_K(m) \;\Rightarrow\; W_K(s) \neq m$$

Using this assumption one can prove the optimality of a technique called *iterated dominance analysis.*

Theorem Basic rationality implies iterated dominance analysis.

Proof: For this proof, and those of several following theorems, see [10] and [11]. □

Iterated dominance analysis handles the column dominance problem in figure 3. Using the basic rationality of K, we can show that action d is irrational for K. Therefore, neither ad nor bd is a possible outcome, and J need not consider them.[2] Of the remaining two possible outcomes, ac dominates bc (from J's perspective), so action b is irrational for J.

IV Action Dependence

Unfortunately, there are situations that cannot be handled by the basic rationality assumptions alone. Their weakness is that they in no way account for dependencies between the actions of interacting agents. This section offers several different, but inconsistent, approaches to dealing with this deficiency.

The simplest case is complete independence. The independence assumption states that each agent's choice of action is independent of the other. In other words, each agent's reaction function yields the same value for every one of the other agent's actions. For all m, m', n, and n', we would then have

$$\begin{aligned} A^s_K(m) &= A^s_K(m') \\ A^s_J(n) &= A^s_J(n'). \end{aligned}$$

[2] We write mn to describe the situation where J has chosen action m, and K has chosen action n; we call this a *joint action.*

The main consequence of independence is a decision rule commonly known as *case analysis.* If for every "fixed move" of K, one of J's actions is superior to another, then the latter action is forbidden. The difference between case analysis and dominance analysis is that it allows J to compare two possible actions for each action by K without considering any "cross terms."

As an example, consider the payoff matrix in figure 4. Given independence of actions, a utility-maximizing agent J should perform action a: if K performs action c, then J gets 4 units of utility rather than 3, and if K performs action d, then J gets 2 units of utility rather than 1. Dominance analysis does not apply in this case, since the payoff (for J) of the outcome ad is less than the payoff of bc.

		K	
		c	d
J	a	4, 2	2, 4
	b	3, 3	1, 1

Figure 4: Case Analysis Problem

Theorem Basic rationality and independence imply case analysis.

By combining the independence assumption with mutual rationality, we can also show the correctness of an iterated version of case analysis.

Theorem Mutual rationality and independence imply iterated case analysis.

As an example of iterated case analysis, consider the situation in figure 5. J cannot use dominance analysis, iterated dominance analysis, nor case analysis to select an action. However, using case analysis K can exclude action c. With this information and mutual rationality, J can exclude action a.

		K	
		c	d
J	a	3, 3	2, 4
	b	1, 1	4, 2

Figure 5: Iterated Case Analysis Problem

Note that, if two decision procedures are not independent, the independence assumption can lead to nonoptimal results. As an example, consider the following well-known "paradox." An alien approaches you with two envelopes, one marked "\$" and the other marked "¢". The first envelope contains some number of dollars, and the other contains the same number of cents. The alien is prepared to

		Alien: 1000	Alien: 1
J	\$	x 1000	x 1
J	¢	x 10	x .01

Figure 6: Omniscient Alien Problem

give you the contents of either envelope. The catch is that the alien, who is omniscient, is aware of the choice you will make. In an attempt to discourage greed on your part, he has decided to put one unit of currency in the envelopes if you pick the envelope marked \$ but one thousand units if you pick the envelope marked ¢. Bearing in mind that the alien has decided on the contents of the envelope before you pick one, which envelope should you select?

The payoff matrix for this situation is shown in figure 6. Since the payoff for \$ is greater that that for ¢ for **either of the alien's options, case analysis dictates choosing the \$ envelope.** Assuming that the alien's omniscience is accurate, this lead to a payoff of \$1.00. While selecting the envelope marked ¢ violates case analysis, it leads to a payoff of \$10.00.

We can easily solve this problem by describing the alien's reaction function and abandoning the independence constraint. The appropriate axioms are $A^s_K(\$) = 1$ and $A^s_K(¢) = 1000$.

These constraints limit J's attention to the lower left hand corner and the upper right hand corner of the matrix. Since the payoff for selecting the ¢ envelope is better than the payoff for selecting the \$ envelope, a rational agent will choose ¢. Although the example given here is whimsical, there are real-world encounters where the assumption of independence is unwarranted, and where the effect illustrated above must enter the rational agent's analysis.

Another interesting example of action dependence is *common behavior*. The definition requires that we consider not only the current situation s but also the permuted situation s' in which the positions of the interacting agents are reversed. An agent J and an agent K have *common behavior* if and only if the action of K in situation s is the same as that of agent J in the permuted situation s'.[3]

$$\forall s\ W_K(s) = W_J(s')$$

Common behavior is a strong constraint. While it may be insupportable in general, it is reasonable for interaction among artificial agents, especially those built from the same design. Unfortunately, it is not as strong as we would like, except when combined with general rationality.

[3] This constraint is similar to the *similar bargainers* assumption in [28,27].

V General Rationality

General rationality is a stronger version of basic rationality, the primary difference being that general rationality applies to decision procedures rather than to individual actions. We introduce a new set of relations and functions to define general rationality. Let $\mathcal{R}_i$ denote a unary predicate over procedures that is true if and only if its argument is rational for agent i. A *generally rational* agent can use a procedure only if it is rational.

$$\neg\mathcal{R}_J(P) \Rightarrow \exists s\ (W_J(s) \neq P(s)).$$

Recall that W_J here is a function that designates the action performed by J in each situation, as described above. In order to use this definition to judge which actions are rational, we of course need to define further the rationality predicate $\mathcal{R}_J$.

The definition of rationality for a decision procedure is analogous to that for individual actions. A procedure is irrational if there is another procedure that *dominates* it:

$$(\exists P'\ \mathcal{D}_J(P', P)) \Rightarrow \neg\mathcal{R}_J(P).$$

One procedure dominates another if and only if it yields as good a payoff in every game and a better payoff in at least one game. Let the term $\mathcal{A}^s_K(P)$ denote the action that agent K will take in situation s if agent J uses procedure P. Then the formal definition of dominance is:

$$\begin{gathered}\mathcal{D}_J(P', P) \iff \\ \forall s\ p^s_J(P'(s), \mathcal{A}^s_K(P')) \geq p^s_J(P(s), \mathcal{A}^s_K(P)) \\ \wedge\ \exists s\ p^s_J(P'(s), \mathcal{A}^s_K(P')) > p^s_J(P(s), \mathcal{A}^s_K(P)).\end{gathered}$$

The advantage of general rationality is that, together with common behavior (defined in the last section), it allows us to eliminate joint actions that are dominated by other joint actions for all agents, a technique called *dominated case elimination.*

Theorem General rationality and common behavior imply dominated case elimination.

Proof: Let s be a situation with joint actions uv and xy such that $p^s_J(u,v) > p^s_J(x,y)$ and $p^s_K(u,v) > p^s_K(x,y)$, and let P be a decision procedure such that $P(s) = x$ and $P(s') = y$ (where s' is the permuted situation, where J and K's positions have been reversed). Let Q be a decision procedure that is identical to P except that $Q(s) = u$ and $Q(s') = v$. Under the common behavior assumption, Q dominates P for both J and K and, therefore, P is generally irrational. □

In other words, if a joint action is disadvantageous for both agents in an interaction, at least one will perform a different action. This conclusion has an analog in the informal arguments of [5] and [17].

No-conflict situations are handled as a special case of this result. The *best plan* rule states that, if there is a joint action that maximizes the payoff to all agents in an interaction, then it should be selected.

Corollary General rationality and common behavior imply best plan.

Proof: Apply dominated case elimination to each of the alternatives. □

As an example of best plan, consider the situation pictured in figure 7. None of the preceding techniques (e.g., dominance analysis, case analysis, iterated case analysis) applies. However, *ac* dominates all of the other joint actions, and so J will perform action a and K will perform action c (under the assumptions of general rationality and common behavior).

		K: *c*	K: *d*
J	*a*	4, 4	2, 2
J	*b*	1, 1	3, 3

Figure 7: Best Plan

General rationality and common behavior also handle difficult situations like the prisoner's dilemma [2,5,25] pictured in figure 8. Since the situation is symmetric (i.e., $s = s'$, using our earlier notation), common behavior requires that they both perform the same action; general rationality eliminates the joint action *bd* since it is dominated by *ac*. The agents perform actions a and c respectively, and each receives 3 units of utility. By contrast, case analysis dictates that the agents perform actions b and d, leading to a payoff of only 2 units for each.

		K: *c*	K: *d*
J	*a*	3, 3	1, 4
J	*b*	4, 1	2, 2

Figure 8: Prisoner's Dilemma

Unfortunately, general rationality and common behavior are not always consistent. As an example, consider the battle of the sexes problem in figure 9. Again the situation is symmetric, and common behavior dictates that both agents perform the same action. However, both joint actions on the *ac*/*bd* diagonal are forbidden by dominated case elimination.

This inconsistency can be eliminated by occasionally restricting the simultaneous use of general rationality and

		K: *c*	K: *d*
J	*a*	1, 1	3, 4
J	*b*	4, 3	2, 2

Figure 9: Battle of the Sexes

common behavior to non-symmetric situations. Nevertheless, in a no-communication situation, the resolution of a conflict such as that in figure 9 remains undetermined by the constraints we have introduced.

VI Conclusions

A. Coverage of this approach

There are 144 distinct interactions between two agents with two moves and no duplicated payoffs. Of these, the techniques presented here cover 117. The solutions to the remaining 27 cases are unclear, e.g., the situation in fig. 10.

		K: *c*	K: *d*
J	*a*	4, 2	1, 4
J	*b*	3, 3	2, 1

Figure 10: Anomalous situation

For a discussion of a variety of other techniques that can be used to handle these situations, as well as a discussion contrasting all of these approachs with those used in game theory, see [27].

B. Suitability of this approach

This paper's analysis of interactions presupposes a variety of strong assumptions. First, the agents are assumed to have common knowledge of the interaction matrix, including choices of actions and their outcomes. Second, there is no incompleteness in the matrix (i.e., there are no missing utilities). Third, the interaction is viewed in isolation (i.e., no consideration is given to future interactions and the effects current choices might have on them). Fourth, there must be effective simultaneity in the agents' actions (otherwise, there are issues concerning which agent moves first, and the new situation that then confronts the second agent).

Admittedly, these are serious assumptions, but there are some situations where they are satisfied. Consider as an example two ALVs approaching opposite ends of a narrow tunnel, each having the choice of using the tunnel or trying one of several alternate routes. It is not unreasonable to assume that they have common knowledge of one another's approach (e.g., through reconnaisance). Nor is it unreasonable to assume that the agents have some models of one another's utility functions. Finally, in the domain of route navigation, the choices are often few and well-defined. There might be no concern (in this case) over future encounters, and the decisions are effectively simultaneous. The types of analysis in this paper are an appropriate tool to use in deciding what action to take.

For most domains, of course, the assumptions listed above are far too limiting, and clearly more work needs to be done in developing this approach so that each of the most restrictive assumptions can be removed in turn. The work in [28,27] represents steps in that direction. Currently, research on the question of incomplete matrices is being pursued, so that the type of conflict analysis presented in this paper can be applied to interactions with incomplete information [26]. Future work will focus on issues arising from multiple encounters, such as retaliation and future compensation for present loss.

Intelligent agents will inevitably need to interact flexibly with other entities. The existence of conflicting goals will need to be handled by these automated agents, just as it is routinely handled by humans. The results in this paper and their extensions should be of use in the design of intelligent agents able to function successfully in the face of such conflict.

REFERENCES

[1] D. E. Appelt. *Planning Natural Language Utterances to Satisfy Multiple Goals.* PhD thesis, Stanford Univ., 1981.

[2] Robert Axelrod. *The Evolution of Cooperation.* Basic Books, Inc., New York, 1984.

[3] D. Corkill. *A Framework for Organizational Self-Design in Distributed Problem-Solving Networks.* PhD thesis, University of Massachusetts, Amherst, MA, 1982.

[4] Daniel D. Corkill and Victor R. Lesser. The use of meta-level control for coordination in a distributed problem solving network. IJCAI-83, pp. 748–756.

[5] L. Davis. Prisoners, paradox and rationality. *American Philosophical Quarterly*, 14, 1977.

[6] Randall Davis and Reid G. Smith. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, 20(1):63–109, 1983.

[7] Edmund H. Durfee, Victor R. Lesser, and Daniel D. Corkill. Increasing coherence in a distributed problem solving network. IJCAI-85, pp. 1025–1030.

[8] R. Fagin and J. Y. Halpern. Belief, awareness, and limited reasoning: preliminary report. IJCAI-85, pp. 491–501.

[9] Michael R. Genesereth. The role of plans in automated consultation. IJCAI-79, pp. 311–319.

[10] Michael R. Genesereth, Matthew L. Ginsberg, and Jeffrey S. Rosenschein. *Cooperation without Communication.* HPP Report 84-36, Heuristic Programming Project, Computer Science Department, Stanford University, September 1984.

[11] Michael R. Genesereth, Matthew L. Ginsberg, and Jeffrey S. Rosenschein. *Solving the Prisoner's Dilemma.* Report No. STAN-CS-84-1032 (HPP-84-41), Computer Science Department, Stanford University, November 1984.

[12] Michael Georgeff. Communication and interaction in multi-agent planning. AAAI-83, pp. 125–129.

[13] Michael Georgeff. A theory of action for multi-agent planning. AAAI-84, pp. 121–125.

[14] M. L. Ginsberg. Decision procedures. In *Proceedings of the Distributed Artificial Intelligence Workshop*, pages 43–65, AAAI, Sea Ranch, CA, December 1985.

[15] Ira P. Goldstein. *Bargaining Between Goals.* A. I. Working Paper 102, Massachusetts Institute of Technology Artificial Intelligence Laboratory, 1975.

[16] Joseph Y. Halpern and Yoram Moses. *Knowledge and Common Knowledge in a Distributed Environment.* Research Report IBM RJ 4421, IBM Research Laboratory, San Jose, California, October 1984.

[17] D. R. Hofstadter. Metamagical themas—computer tournaments of the prisoner's dilemma suggest how cooperation evolves. *Scientific American*, 248(5):16–26, May 1983.

[18] Kurt Konolige. A computational theory of belief introspection. IJCAI-85, pp. 502–508.

[19] Kurt Konolige. *A Deduction Model of Belief and its Logics.* PhD thesis, Stanford University, 1984.

[20] A. L. Lansky. *Behavioral Specification and Planning for Multiagent Domains.* Technical Note 360, SRI International, Menlo Park, California, November 1985.

[21] V. Lesser and D. Corkill. The distributed vehicle monitoring testbed: a tool for investigating distributed problem solving networks. *AI Magazine*, 4(3):15–33, Fall 1983.

[22] Hector J. Levesque. A logic of implicit and explicit belief. AAAI-84, pp. 198–202.

[23] R. Duncan Luce and Howard Raiffa. *Games and Decisions, Introduction and Critical Survey.* John Wiley and Sons, New York, 1957.

[24] R. Moore. A formal theory of knowledge and action. In J. R. Hobbs and R. C. Moore, editors, *Formal Theories of the Commonsense World*, Ablex Publishing Co., 1985.

[25] D. Parfit. *Reasons and Persons.* Clarendon Press, Oxford, 1984.

[26] Jeffrey S. Rosenschein. *Cooperation in the Presence of Incomplete Information.* Technical Report, Knowledge Systems Laboratory, Computer Science Dept., Stanford Univ., 1986. In preparation.

[27] Jeffrey S. Rosenschein. *Rational Interaction: Cooperation Among Intelligent Agents.* PhD thesis, Stanford University, 1986. Also published as STAN-CS-85-1081 (KSL-85-40), Department of Computer Science, Stanford University, October 1985.

[28] Jeffrey S. Rosenschein and Michael R. Genesereth. Deals among rational agents. IJCAI-85, pp. 91–99.

[29] R. G. Smith. *A Framework for Problem Solving in a Distributed Processing Environment.* PhD thesis, Stanford University, 1978.

[30] R. Steeb, S. Cammarata, F. Hayes-Roth, and R. Wesson. *Distributed intelligence for air fleet control.* Technical Report WD-839-ARPA, The Rand Corporation, Dec. 1980.