Stanford Heuristic Programming Project
Report No. HPP-84-36

September 1984

Cooperation without Communication

Michael R. Genesereth
Matthew L. Ginsberg
Jeffrey S. Rosenschein

COMPUTER SCIENCE DEPARTMENT
Stanford University
Stanford, California 94305

Cooperation without Communication

Abstract

Work in distributed artificial intelligence has until now addressed the issues of multi-agent interaction only under the strict assumption of "agent benevolence," where each participant in an interaction freely aids all others. The approach taken in game theory, however, assumes "agent selfishness," where each participant acts solely to ensure its own well-being. This paper attempts to clarify the differences between these two approaches by developing a model general enough to describe both of them, and by investigating the consequences of this model in some specific instances, such as the problems of coordination of actions among non-communicating agents and the prisoner's dilemma. In particular, the expanded model includes multi-agent interactions where the participants have distinct or conflicting goals. The mechanism presented ensures harmonious interactions without communication under many circumstances.

§1. The Need for Cooperation

1.1 Ubiquity of interaction and conflict

The world functions through interacting agents. Each person pursues his own goals through encounters with other people or machines. Items are purchased, schedules coordinated, arrangements made so that the original goals can be satisfied. Only in extremely limited situations can an individual's goals be pursued without interaction; when such circumstances do arise, it is only because of careful preparation (and this preparation has invariably been accomplished through interactions).

When people deal with one another, they often bring to the encounter differing goals, and the interaction process takes this conflict into account. People promise, threaten, and together find compromises that will satisfy all parties, without such outcomes necessarily being ideal for all. The process of negotiation takes place in both formal and informal contexts; it is so much a part of daily life that it often passes unnoticed. Deciding on a meeting time or place, for example, often involves suboptimal satisfaction of conflicting interests.

It would seem that such negotiation only takes place among humans, and that human-machine interactions occur without such processes having a role. In a sense this is correct.

Man-machine interactions currently take place in very restricted, formal environments due to the imperatives of unintelligent machines.

The rigidity of the protocols governing these interactions is not so much a consequence of their non-communicative nature as of general assumptions regarding the inflexibility of the machines involved. In this paper, we will discuss formalisms that allow machines to interact flexibly with other agents without communicating directly with them.

1.2 Extreme cases

There are two extreme cases in the study of interaction. The first is when participants are engaged in a "pure" conflict: any advantage gained by one is exactly balanced by a loss to another. This case is often called "zero sum," and has been examined in great detail by game theorists [16]. The other extreme is the absence of conflict, when two participants have identical goals. Under these circumstances, the agents are mainly concerned with sound and efficient methods of helping each other with their shared tasks [22].

"Pure" conflicts in the real world are extremely rare, and it is difficult to generate compelling examples of their occurrence. War, or battles within a war, are often taken as canonical examples of pure conflict; in fact, the options available in most military situations allow for outcomes that are mutually preferred by both parties (e.g., retreat of one side, rather than annihilation of that side at a great cost to the victor). Thus, although the game theory work on zero-sum conflicts is impressive in its elegance, it has achieved this elegance by abstracting away significant features of conflict.

Conflict-free interactions are also extremely rare. Overlapping of interests notwithstanding, there is almost always enough variation among the participants' desires to cause them to prefer slightly different outcomes. Two agents may want the same event to occur but may differ as to which should exert itself in order to bring it about, for example.

The vast majority of real-world interactions lie between the two extremes of total conflict and absence of conflict, and it is these partial conflict interactions that will be the primary concern of this paper.

1.3 Real world examples

Current AI programs, when they deal with human-machine or machine-machine interactions at all, assume a formal and rigid protocol for those interactions. This approach is inadequate for many tasks. The reasons for this invariably stem from the need for flexibility, even with machine-machine interactions. If a machine is going to act as a surrogate for a human, and is to accomplish tasks at a level comparable to that of a human, the machine must not operate under restrictive, crippling assumptions. What follows are three scenarios where intelligent agents would need sophisticated and flexible interaction capabilities in order to perform satisfactorily.

Resource management applications

People manage a variety of resources in their daily lives. The most conspicuous of these is time; humans continually interact with one another in order to allocate their time satisfactorily. One task for an automated personal secretary would certainly be to manage its owner's time effectively. It would have to be able to negotiate with other similar entities (or human secretaries) and arrive at, for example, compromise meeting schedules.

The typical human would be thoroughly unsatisfied with an automated secretary that performed less well at this task than its human counterpart. The full complexities of this scenario (commitments, threats, quid pro quo offers, hiding of information) require fully flexible interacting agents, beyond the capabilities currently addressed in AI systems. Management of other resources (such as space) would require similarly capable systems.

Space applications

Intelligent autonomous agents can be expected to operate under extreme circumstances; indeed, it is precisely under extremes that are difficult for humans that automated capabilities become attractive. Construction of vehicles or stations in space is an example of a scenario where automated builders could usefully be employed. Such construction agents can be expected to have certain intelligent characteristics, since the location and fitting of free-floating segments requires considerable skill.

In addition to the robotics and vision capabilities required, however, are abilities of coordination and synchronization. Although there may be rigid solutions to the problem of construction (and there does exist a globally accepted set of goals), it is much more likely that flexible capabilities, similar to those that would be employed by human builders, will be more appropriate for the job. Uncertainties over such things as the order in which pieces will be located (or the fuel that an agent will have expended at a particular point in the process), and the negotiation needs of synchronization argue for a flexible approach to interaction.

Military applications

Much attention has been paid recently to the possibility of autonomous land vehicles (ALV's), intelligent automated agents that would be capable of independently carrying out military missions. While well-specified, rigid missions might be performed without sophisticated interaction skills, such rigid missions are exceedingly rare.

As pointed out above, even conflicts arising in war are not of the zero-sum variety; there are cases where mutually preferred outcomes can be tacitly or explicitly agreed upon. In encounters with allies this is especially important. An ALV from one NATO country may be allied with an ALV from another NATO country, but in an encounter their goal structures could be very different (this is true even of two ALV's that have been sent out on different missions by the same commander). They must have some way of reconciling their conflicting goals. A sophisticated method of interaction is especially necessary in cases (such as this one) where encounters among the agents will occur haphazardly, the domain of interaction is complicated and unpredictable, and the possibilities of central arbitration of disputes are minimal.

1.4 Payoff matrices as examples

The game theory literature has adopted a simple and elegant notation to represent multi-agent interactions: the payoff matrix. Here is an example:

	A	B
A	3/1	2
B	2/5	0/1

Our notation here is that the first player selects a move labeling one of the two rows and the second selects one of the two columns. A single number indicates an identical payoff for both players, while '3/1', for example represents a payoff of 3 for the first player and 1 for the second.

We will use the same mechanism to illustrate several types of interactions that can occur among agents. These are only a few of the many possibilities, and will be discussed in further detail below. For a full discussion of the types of interactions that can occur, even within the simplified world of abstract matrices, see [21] (which includes a 78 member taxonomy of two player games).

Best plan

One kind of interaction is that in which there is a mutually agreed-upon "best plan." Although the participants might disagree among themselves as to what the second best or third best alternatives are, they do agree (and might be expected to converge without communication) on a unique first choice. Rapoport and Guyer [21] call these "no-conflict games," and variations on the basic theme occupy positions 1-6 and 58-63 in their taxonomy. An example of a best plan game is the following:

	A	B
A	7	4
B	5	6

In this case the joint move (A,A) is preferred by both players.

Prisoner's Dilemma

One of the best known examples of a two-person interaction is the so-called "prisoner's dilemma." Its motivating scenario involves two prisoners being held by the police. Each is questioned individually and is offered some incentive to implicate the other, but they will be better off maintaining collective silence than if both talk. Thus each would really prefer to confess while the other remains silent, failing that, to mutually remain silent, and least of all, to mutually confess. The payoff looks like this:

	C	D
C	3	0/5
D	5/0	1

The dilemma revolves around the fact that each player sees that it is better for him to defect regardless of the other's action. Yet if both defect, they end up with an outcome that is worse *for both* than if they both had cooperated. This "temptation" to defect is the hallmark of the prisoner's dilemma (Rapoport and Guyer characterize the matrix, number 12 in their taxonomy, as having a "single strongly stable deficient equilibrium"; it is the only matrix in their taxonomy having this characteristic). Attempts to surmount the unpleasant conclusion that both players will defect have preoccupied much game theory and philosophy literature [2,5,11,20,24]. Below, we present an approach that solves the prisoner's dilemma under a particular set of assumptions.

§2. Rationality

Although any individual agent can assume the rationality of its own behavior, it is important for that agent to be able to remain flexible with regard to the rationality of the others. We will therefore investigate the behavior of an individual rational agent under a variety of assumptions regarding the rationality of the agents with which it is interacting. By "rational" here we mean an agent capable of deriving the logical consequences of the precise definition of rationality we will present shortly.

We will also investigate only that small set of interactions to which we referred in the last section—specifically, situations where there is no communication between the players and where the goal of each participant is to maximize his return as measured by some sort of payoff function.

To formalize these notions, let P be a set of players and, for each player $i \in P$, let M_i be a set of possible moves for i. For $S \subset P$, we denote $P - S$ by $\bar{S}$; we will also write i instead of $\{i\}$ where no confusion is possible. Thus $\bar{\imath} = P - \{i\}$. We also write M_S for $\prod_{i \in S} M_i$.

We will denote by m_S an element of M_S; this is a collective move for the players in S. To $m_S \in M_S$ and $m_{\bar{S}} \in M_{\bar{S}}$ correspond an element $\vec{m}$ of M_P. By a *game* we will now mean a real-valued function

$$p : P \times M_P \to \mathbb{R}; \tag{1}$$

$p(i, \vec{m})$ is the payoff for player i if move $\vec{m}$ is made. Note that we think of a game not as a sequence of moves or interactions but as a single situation, and we identify the game with the function which describes the outcomes of the actions available to the various players.

Rather than analyze specific moves in specific games, it will be more appropriate for us to discuss the manner in which the players decide to make these moves. Each player i is thus assumed to have some decision procedure D_i by which he decides what move to make in any given game g. Thus if we denote by g_i the set of possible moves for i in g, and write G for the set of all games and G_i for the set of all moves legal for i in *some* game, we can define a *decision procedure* to be a function

$$D_i : G \to G_i$$

such that $D_i(g) \in g_i$ for all games g. A decision procedure therefore assigns to any game a specific move which is legal in that game; intuitively, each player uses his decision procedure to select from his legal alternatives in any given game. For a fixed game g and set S of players, we will denote $\prod_{i \in S} [D_i(g)]$ by $D_S(g)$. D_S can be thought of as a "collective" decision procedure by which a group S of players selects a joint move.

Suppose now that player i uses his decision procedure D_i to make a move m_i in a game g with payoff function p. The other players in the game use their collective decision procedure $D_{\bar{\imath}}$ to make a move $m_{\bar{\imath}}$ and the resulting payoff to i is therefore

$$pay(i, D_i, g) = p(i, \vec{m}). \tag{2}$$

The function *pay* here gives the payoff to i in game g if he uses decision procedure D_i; note that the appearance of $\vec{m}$ in (2) means that *pay* is implicitly a function of all of the D_j's, and not just D_i.

We are now in a position to define what it means for a decision procedure to be rational. Specifically, we will say that a decision procedure D_i is *irrational* if there exists a decision procedure C_i such that

$$pay(i, D_i, g) \leq pay(i, C_i, g) \tag{3}$$

for all games g, with the inequality being strict in at least one case. In other words, a decision procedure is irrational if there is another decision procedure which is better in some specific game, and no worse in all others.

The power of the characterization (3) depends on the decision procedure $D_{\bar{i}}$ which provides us with $m_{\bar{i}}$ in the right hand side of (2). This in turn depends on the individual decision procedures D_j for $j \neq i$. If we had complete knowledge of these other decision procedures, we would effectively be moving with the knowledge of what the other players' moves would be, and *our* decision procedure would be easily determined. In practice, of course, this will not be the case, and gaps in our knowledge of $D_{\bar{i}}$ will make it impossible for us to eliminate all of the strategies which are deemed to be irrational by (3). Our interest will be in finding constraints on our own decision procedure given various assumptions about the others.

At this point we have made no such assumptions at all. Specifically, we have *not* assumed that the other decision procedures are independent of our decision procedure D_i. In fact, the ability to drop this independence assumption is the reason we are working not with the moves themselves, but with the analyses which lead to them. By working in this fashion, it is possible to assume that the players have knowledge of each others' strategies without being led to the circular arguments that otherwise pervade this sort of analysis.

We will in fact deal with a hierarchy of assumptions regarding the decision procedures of the other players. Here are the definitions we will adopt:

1. **Minimal rationality:** In this case, we make no assumptions at all about the other players, although any particular agent still assumes its own rationality. The effect of this is for us to make our decisions as if the other players were moving randomly. Perhaps this would be a suitable assumption for robots to make when dealing with humans.
2. **Separate rationality:** The decision procedures D_i of all players are assumed to be rational. A reasonable approach for robots dealing with other robots.
3. **Unique rationality:** D_i and D_j are functionally independent for $i \neq j$. The functional independence means that each player's decision procedure does not influence the others. This assumption is generally made in the game theory literature, and amounts to assuming that the moves of the other players are fixed in advance (since their decision procedures are), although the exact nature of those moves is unknown.
4. **Common rationality:** $D_i = D_j$ for all i and j. This assumption (which cannot be described in the conventional game-theoretic formalism) is peculiarly appropriate to an artificial intelligence setting; it would be straightforward to equip potentially interacting agents (such as those described in the first section of this paper) with identical decision procedures.

It is clear that minimal rationality is entailed by any of the other conditions, in that if a decision procedure can be shown to be irrational assuming minimal rationality, it will also be irrational under any of the other definitions. Separate and unique rationality are independent but consistent; there are decision procedures which are rational under one assumption but irrational under the other. We will refer to the combination of both separate and unique rationality as **individual rationality.**

Common and unique rationality are inconsistent. This is intuitively clear from the definitions; we will show that there is no decision procedure that is rational under both definitions, since they generate conflicting strategies for the prisoner's dilemma. Finally, we have:

Common rationality implies separate rationality if at least one of the players has a rational *decision procedure.*

If the common decision procedure D_i were irrational for some player i, with C_i being an improvement, it would follow from the completely symmetric definition of common rationality that C_i would be an improvement for all of the other players as well.

Note that it is not true in general that a decision procedure that is rational for one player is rational for all others. A third player whose actions are known to be biased in favor of one of the others can introduce asymmetry of this sort, since we approach problems differently if there are additional agents present upon whose support and cooperation we can rely.

§3. Results

In this section we will investigate the consequences of our definitions of rationality. We will assume throughout that g is some fixed game with payoff function p, and that D_i is a rational decision procedure for player i.

Theorem 1 (Case analysis). *Assuming unique rationality, if for some moves c_i and d_i, for all $m_{\bar{\imath}}$,*

$$p(i, d_i \times m_{\bar{\imath}}) < p(i, c_i \times m_{\bar{\imath}}),$$

then $D_i(g) \neq d_i$.

Proof. Suppose that $D_i(g) = d_i$, and let C_i be the decision procedure given by

$$C_i(g') = \begin{cases} D_i(g'), & \text{if } g' \neq g; \\ c_i, & \text{if } g' = g, \end{cases}$$

and set $m_{\bar{\imath}} = D_{\bar{\imath}}(g)$. Now

$$pay(i, D_i, g) = p(i, d_i \times m_{\bar{\imath}}) < p(i, c_i \times m_{\bar{\imath}}) = pay(i, C_i, g),$$

while $pay(i, C_i, g') = pay(i, D_i, g')$ for $g' \neq g$, so that the decision procedure D_i is irrational. □

It is important to realize that this proof depends critically upon the independence of the decision procedures of the various players. If this were not the case, we would not be able to show that $pay(i, C_i, g') = pay(i, D_i, g')$ for $g' \neq g$, since we could not be sure that the other players would not change their strategies for g' in reaction to i's changing his strategy for g.

We define a single move m_i to be *rational* if there is a rational decision procedure D_i which produces it, and now have:

Corollary 2 (Iterated case analysis). *Assuming individual rationality, if for some moves c_i and d_i, for all rational $m_{\bar{i}}$,*

$$p(i, d_i \times m_{\bar{i}}) < p(i, c_i \times m_{\bar{i}}),$$

then $D_i(g) \neq d_i$.

Proof. Clear. □

If this result allows us to conclude that $D_i(g) = c_i$ for some specific c_i, this is known as the *solution in the complete weak sense* in the game theory literature.

The consequences of case analysis are quite well known. Consider the example used to introduce our payoff matrix notation:

	A	B
A	3/1	2
B	2/5	0/1

Independent of the second player's choice, the first will be better if he makes move A, since his payoff will be 3 as opposed to 2 if the other player chooses move A, and 2 as opposed to 0 if B is chosen. Case analysis therefore implies that A is the only rational move for the first player in this game.

There is no such implication for the second player. If we assume individual rationality, however, the second player will realize that the first can be counted on to make move A, and will therefore respond with move B (receiving a payoff of 2 instead of 1).

Alternatively, consider the prisoner's dilemma, presented earlier:

	C	D
C	3	0/5
D	5/0	1

(4)

Since the two strategies are independent, each player can reason that whatever the other does, his best move is D (defection). The paradox lies in the fact that this move produces payoffs for each player which are less than if they both choose move C (cooperation).

Another difficult problem for individual rationality is presented in the best plan game:

	A	B
A	7	4
B	5	6

(5)

It is fairly clear that both players should choose move A, since this will result in their obtaining the highest payoff available to them. It is impossible to reach this conclusion assuming individual rationality, however, since it is not enough to assume merely that the other player in this game is rational. One must also assume that he knows *you* are rational, that he knows you know he is rational, that he knows you know he knows you are rational, and so on. (This is a consequence of what is referred to in [10] as *common knowledge*.) The usual game-theoretic approach is to address this problem directly by finding characterizations of the joint move (A,A) which can then be incorporated into a definition of group rationality. No attempt is made to derive this sort of "coordinated action" from the definition of rationality for an individual. We find this unsatisfactory.

Finally, we should mention that case analysis (and therefore individual rationality generally) has difficulty dealing with altruism. However much of a hurry we may be in to get our laundry, it seems reasonable to stop and help someone who has just been hit by a truck, even though our personal utility does not go up by doing so. The conventional solution is to build some sort of "altruism factor" into the payoff matrices, but this is unnecessarily *ad hoc*.

We should point out here that just as we object to a formalism that forces one *away* from altruistic behavior, we would have similar difficulties with one which forced the players toward it. What is needed is a theory that is uncommitted on this point, leaving it in some sense as an implementation detail. We will see presently that common rationality is such a theory. Before doing so, however, we return to the game (5):

Theorem 3 (Coordination). *Assuming common rationality, if there exists a collective move $\vec{c}$ such that*

$$p(i,\vec{c}) > p(i,\vec{d})$$

for all i and $\vec{d}$, then $D_i(g) = c_i$.

Proof. Let G' be the set of all games for which the conditions of the theorem are satisfied, i.e., for which there is a single move which maximizes the payoffs to all of the players. For a game g in G', let $\vec{m}(g)$ be this uniformly best move, and define a new decision procedure C_i given by:

$$C_i(g') = \begin{cases} m_i(g'), & \text{if } g' \in G' \\ D_i(g'), & \text{otherwise.} \end{cases}$$

Since the description of the games in G' is symmetric under interchange of the players involved, we have $pay(i, C_i, g') = pay(i, D_i, g')$ for $g' \notin G'$, and $pay(i, C_i, g') \geq pay(i, D_i, g')$ for $g' \in G'$. If $D_i(g) \neq c_i$, the inequality will be strict for g and D_i will be irrational. □

The game (5) is of course handled as a special case of this result. To deal with the prisoner's dilemma and related games, we have:

Theorem 4 (Cooperation). *Assume common rationality, and suppose that g has moves $\vec{c}$ and $\vec{d}$ such that*

$$p(j,\vec{d}) \leq p(j,\vec{c})$$

for all j, with the inequality being strict in at least one case. Then there exists an i such that $d_i \neq D_i(g)$.

In other words, if mutual defection is to everyone's disadvantage, then at least one player will cooperate. For a game such as (4) which is symmetric (in that the payoff function is), it follows that all of the players cooperate. This conclusion has an analog in the informal arguments of [5] and [11].

Proof. The technique is the same as that of the earlier proof. We let G' be the set of games which are obtained from g by permuting the players in g in some (possibly trivial) fashion, and consider the decision procedure which agrees with D except on the games in G', where the move $\vec{c}$ is made. This produces unchanged payoffs for all of the games not in G', and payoffs which are no worse in G'. Since at least one player actually benefits from the change in the original game g, it follows that there is some game in G' which benefits any particular player i. The conclusion follows. $\square$

It is important to realize both that the cooperation theorem does *not* imply that all of the players cooperate, and the reason for this. The possibility of one player cooperating while the others defect (undoubtedly beneficial to the defectors and disastrous for the cooperator) is typical of the "altruism" allowed under common rationality. Consider the following non-result:

Non-theorem 5 (Restricted case analysis). *Assuming minimal rationality, if for some c_i and d_i, for all $c_{\bar{i}}$ and $d_{\bar{i}}$,*

$$p(i,\vec{d}) < p(i,\vec{c}),$$

then $d_i \neq D_i(g)$.

Proof? If $d_i = D_i(g)$, consider the decision procedure given by

$$C_i(g) = \begin{cases} D_i(g'), & \text{if } g' \neq g; \\ c_i, & \text{otherwise.} \end{cases}$$

Whatever the potential dependence of the other decision procedures on D_i, we will still have $pay(i, C_i, g) > pay(i, D_i, g)$. $\square$

The problem is that we cannot show that

$$pay(i, C_i, g') \geq pay(i, D_i, g') \tag{6}$$

for games g' other than g without additional assumptions, such as unique rationality (since restricted case analysis is indeed a special case of case analysis) or the assumptions in the coordination theorem. In fact, assuming common rationality, if the "better" move c_i forces a reduction in the payoff for some other player, then (6) will specifically *not* hold for some permutation of g. This of course does not mean that C_i is irrational; it merely means that C_i cannot be used to prove that D_i is.

Restricted case analysis is in fact independent of the assumption of common rationality, although it is consistent with it. This is the precise version of our earlier comment that common rationality is ambivalent with regard to altruism.

§4. Previous Work

The problem of interacting agents has been examined in several areas; in this section, we briefly review related work in distributed artificial intelligence (DAI) and in game theory, noting where our work differs from previous approaches.

4.1 AI background

DAI has attempted to address the problems of interacting agents through an artificial intelligence perspective. The single unifying assumption in this work is that one or more of the interacting agents will be using artificial intelligence techniques to guide their actions (including, of course, their communication actions). We call this the "intelligent-agent paradigm." Within this broad categorization, the many individual efforts to give AI systems the capability to interact with other rational systems are seen as potentially increasing efficiency (by harnessing multiple reasoners to solve problems in parallel) or as necessitated by the distributed nature of the problem (e.g., distributed air traffic control [3]).

Smith and Davis' work on the contract net [7] produced a tentative approach to cooperation using a contract-bid metaphor to model the assignment of tasks to processors. One agent, upon receiving a task, decomposes it into smaller subtasks; the subtasks are then announced, other agents bid for the right to perform them, and the original agent, after examining the bids, assigns each subtask. A similar method, using a more explicitly economic model, is employed by others in the enterprise system [17]. Davis has also investigated cooperative problem-solving strategies that lead to easier interactions among agents [6].

Lesser and Corkill have made empirical analyses of distributed computation, trying (through use of a testbed) to discover cooperation strategies that lead to efficient problem solutions for a network of nodes [4]. Georgeff has attacked the problem of assuring non-interference among distinct agents' plans [8,9]. 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.

These DAI efforts have made some headway in constructing cooperating systems; the field as a whole, however, has lacked the formal foundation that might speed progress. Over the past few years, several researchers have begun to construct the theoretical underpinnings. In particular, Appelt [1], Moore [18,19] and Konolige [12,13,14,15] have worked on the formalisms necessary for one agent to reason about another's knowledge and beliefs.

Earlier assumptions

Previous DAI work has assumed 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. Problems to be overcome include those of synchronization, communication, and (inadvertent) destructive interference.

Occasionally, this cooperation arises naturally out of the sense of shared goals the agents are assumed to have. Since, for example, in Lesser and Corkill's systems the agents are always assumed to be working on the same problem overall, it makes little sense to ask *why* they are helping one another: they help each other because they have been designed that way, and they have been designed that way because, in some sense, they are all part of a single problem-solving entity.

Other systems, however, have agents with potentially distinct goals. Nevertheless, these systems have avoided dealing with the issue of non-cooperative agents. Thus, for example, Georgeff's theory treats agents as having distinct plans that can potentially interfere with one another (e.g., through uncoordinated access to the identical resource). Still, he does not deal with the common occurrence of truly conflicting plans: what if agents cannot simply avoid conflict, but must deal with it directly? What principles guide activity under such circumstances? Previous AI work has not addressed these questions.

Our model of agent interaction 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 (even without communication) so as to bring about mutually preferred outcomes.

4.2 Game theory background

Game theory has examined the issues of rational agent interaction for over thirty years, and there is much to be gained from their theoretical treatment of the subject. Here we touch only briefly on some of the ideas that are related to our work; a general survey of game theory can be found in [16].

Related work

Game theory has treated the zero-sum interaction in impressive detail, effectively "solving" it in the general case (the generated solution is conditionally normative in that it specifies what action should be taken for a particular result, such as getting no less than some specified payoff). Communication plays no useful role in the zero-sum theory; since interests of the parties are diametrically opposed, there can be no chance for cooperation or purpose to collusion.

The non-zero-sum interaction, the specific type of encounter in which we are most interested, has not been "solved" in the same way. There are, to be sure, special cases where particular actions are clearly warranted and lead to specific results, but the complete formal results to be found in zero-sum theory have eluded game theorists working on the more general case.

Nevertheless, much energy has been expended on the non-zero-sum interaction both with and without communication (the former is termed "cooperative" by game theorists). When communication is present, it is possible for participants to arrive at compromises that reconcile their conflicting interests. Nash's "solution" to the bargaining problem, for example, satisfies certain axiomatic assumptions as to what characteristics we might want a solution to have; yet it is in no sense completely satisfactory, since the axioms are not necessarily accepted by the agents posited in the theory. In contrast, we need only assume utility maximization in the zero-sum theory to arrive at its normative conclusions.

Their assumptions

Agents in the game theory literature are assumed to be utility maximizers, in the sense that they have a consistent (i.e., transitive) pattern of preferences, and strive to achieve more preferred outcomes over less preferred ones. Game theory assumes that each player is also fully aware of the well-defined set of alternative actions for itself and for other players, and the preference pattern that every agent has over the potential outcomes. In a sense, all agents are assumed to have common knowledge [10] of the payoff matrix (however, they cannot be sure that the other agents see the matrix in the same way—there may be a different encoding of the agents and their choices; this assumption that the agents and choices are "unlabeled" removes certain contextual clues that could lead to solutions, and is severely criticized in [23]).

The common terminology for these assumptions is that the agents are "rational," but as Luce and Raiffa [16] point out

> Though it is not apparent from some writings, the term "rational" is far from precise, and it certainly means different things in the different theories that have been developed. Loosely, it seems to include any assumption one makes about the players maximizing something, and any about complete knowledge on the part of the player in a very complex situation... [*Games and Decisions*, p. 5]

The assumption of "rationality," never made quite precise, has restricted some of the kinds of results that game theorists are able to derive. Within the setting of artificial intelligence, however, it makes sense to define the term in considerably greater detail. Once the definition of rationality has been made concrete (in a variety of alternative formulations), more definite progress can be made on the issues of cooperative interactions.

§5. Conclusion

Intelligent agents will inevitably need to interact flexibly with other entities. While previous AI research has been concerned with the development of protocols for use in conflict-free interactions, game theorists have considered instances of total conflict, the zero-sum games. Our interest has been in situations where there may be only partial conflict between the goals of the agents involved.

Given the assumption of no communication, we have developed a formalism for dealing with these partial-conflict situations by considering the decision procedures of the various participants. This approach is flexible enough to subsume many of the earlier results; our notion of individual rationality, for example, shares many features with the game theorists' usual assumptions.

The decision procedure formalism does not force individual rationality on us, however, and we have also investigated the results of another choice which we referred to as common rationality. Common rationality allows for (but does not require) altruistic behavior on the part of the agents involved and also produces attractive features of coordination and cooperation. The former of these is apparent in its treatment of the best plan game (4), while the latter is demonstrated by its satisfactory handling of the prisoner's dilemma.

Acknowledgement

The authors have chosen to list themselves in alphabetical order. This research has been supported by the Office of Naval Research under grant number N00014-81-K-0004 and by DARPA under NAVELEX grant number N00039-83-C-0136.

References

[1] Appelt, D.E., Planning natural language utterances to satisfy multiple goals, Tech Note 259, SRI International, Menlo Park, California (1982).

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

[3] Cammarata, S., McArthur, D., and Steeb, R., Strategies of cooperation in distributed problem solving, *IJCAI-83*, Karlsruhe, West Germany (1983) 767–770.

[4] Corkill, D.D., and Lesser, V.R., The use of meta-level control for coordination in a distributed problem solving network, *IJCAI-83*, Karlsruhe, West Germany (1983) 748–756.

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

[6] Davis, R., A model for planning in a multi-agent environment: steps toward principles for teamwork, Working Paper 217, MIT AI Lab (1981).

[7] Davis, R., and Smith, R. G., Negotiation as a metaphor for distributed problem solving, *Artificial Intelligence* **20** (1983) 63–109.

[8] Georgeff, M., Communication and interaction in multi-agent planning, *AAAI-83*, Washington, D.C. (1983) 125–129.

[9] Georgeff, M., A Theory of action for multi-agent planning, *AAAI-84*, Austin, Texas (1984) 121–125.

[10] Halpern, J., and Moses, Y., Knowledge and common knowledge in a distributed environment, *Proceedings of the Third Annual ACM Conference on Principles of Distributed Computing*, Vancouver, British Columbia, Canada (1984).

[11] Hofstadter, D.R., Metamagical themas—the calculus of cooperation is tested through a lottery, *Scientific American* **248** (1983) 14–28.

[12] Konolige, K., A first-order formalization of knowledge and action for a multi-agent planning system, Tech Note 232, SRI International, Menlo Park, California (1980).

[13] Konolige, K., Circumscriptive ignorance, *AAAI-82*, Pittsburgh, Pennsylvania (1982) 202–204.

[14] Konolige, K., A deductive model of belief, *IJCAI-83*, Karlsruhe, West Germany (1983) 377–381.

[15] Konolige, K., *A Deduction Model of Belief*, Ph.D. Thesis, Stanford University (1984).

[16] Luce, R.D., and Raiffa, H., *Games and Decisions, Introduction and Critical Survey*, John Wiley and Sons, New York (1957).

[17] Malone, T.W., Fikes, R.E., and Howard, M.T., Enterprise: a market-like task scheduler for distributed computing environments, Working paper, Cognitive and Instructional Sciences Group, Xerox Palo Alto Research Center (1983).

[18] Moore, R.C., Reasoning about knowledge and action, Tech Note 191, SRI International, Menlo Park, California (1980).

[19] Moore, R.C., A formal theory of knowledge and action, Tech Note 320, SRI International, Menlo Park, California (1984). Also to appear in *Formal Theories of the Commonsense World*, Hobbs, J.R., and Moore, R.C. (Eds.), Ablex Publishing Co. (1984).

[20] Parfit, D., *Reasons and Persons*, Clarendon Press, Oxford (1984).

[21] Rapoport, A. and Guyer, M., A taxonomy of 2 x 2 games, *Yearbook of the Society for General Systems Research* **XI** (1966) 203–214.

[22] Rosenschein, J.S., and Genesereth, M.R., Communication and cooperation, HPP Report 84-5 (1984).

[23] Schelling, T.C., *The Strategy of Conflict*, Oxford University Press, New York (1963).

[24] Sowden, L., That there is a dilemma in the prisoners' dilemma, *Synthese* **55** (1983) 347–352.