# On Two Families of Paradigms of Group-Solvability*

Alessandro Agostini
ITC-IRST Trento

November 7, 2002

**Abstract**

We advance and compare two families of **coalitional paradigms** of solvability.† A coalitional paradigm is distinguished from a "noncoalitional" paradigm primarily by its focus on what *groups* of agents can achieve, rather than on what *individual* agents can do—even if cooperating. As a criterion of group formation, our models engage a kind of pairwise, context-dependent coordination between knowledge-based "learning agents," eventually able to communicate the *complete & local* meaning of expressions taken from the literals of a common first-order language.

**Keywords:** Formalisms and logics for agents and MAS; coalition formation; coordination, groups, teams, and group-dynamics.

## 1 Introduction

Following [JFL+01], automated negotiation research can be considered to deal with three broad topics, namely: (a) *Negotiation Protocols*: the set of "rules" that govern the interaction among the agents in play. (b) *Negotiation Objects*: the range of issues, or attributes, over which agreement must be reached. (c) *Agents' Decision Making Models*: the decision making system the agents employ to act in line with the negotiation protocol in order to achieve their objectives. Either negotiation protocols togheter with negotiation objects or agents' decision making models is the dominant concern; it depends on the negotiation scenario.

The scope of this article falls primarily into the domain of decision making models, see for instance [JFL+01, Syc89, ZS97] and the references cited there. More precisely, decision making models for coalition formation and teamwork [SLA+99, SK95, Ago00b]. Our main contribute is to provide a framework to study the important topic of group-learning, where pairwise coordination for group formation is explicitely stated. We do this in the tradition of formal learning theory, say

*Accepted as Poster Paper at the 2nd International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-03). Electronic correspondence to: agostini@irst.itc.it.

†In this paper we prefer to use "solvability" instead of "learning" under the assumption that "learnable" denotes a wider class of objects than "solvable" does. Apart from this somewhat philosophical remark, in this paper we shall always mean the two terms as synonymous.

[OdJMW97, MO98, MO99], that descends from the pioneering studies on inductive inference developed by Solomonoff, Putman, Gold, L. Blum and M. Blum among others. (See for instance [JORS99] for a survey.)

First, we define a communication protocol among a set of agents that always move simultaneously and all relevant moves are made by the agents (agents' moving components are functions in mathematical sense; no randomness ever intervenes). The communication protocol we consider is pairwise, so that the kind of models we discuss is suitable for modeling group situations where communication is not with the whole group, as in an auction, but indeed pairwise, as happens for instance in commercial transactions. Second, we provide some examples of coalition formation and show how *coalitions by coordination* may be eventually used to solve certain classes of structures, in a sense made precise within a model-theoretic paradigm—we call this paradigm of **group-solvability**. A key of success for coordination of a set of agents is, as we will see, the existence of some common knowledge that states who has to be the leader in each pairwise interaction of the agents in the set. Third, we rework some of our models in order to make computability explicit. Some introduced notions (e.g., "coordination sentence," "group-solvability") are proved useful to specify the benefits and the limits of our approach for multiagent systems. On the one hand, the technique for group formation we provide suggests new ways of structuring goal selection and, expecially, goal decomposition in multiagent systems. On the other, we exhibit a problem solvable by a single learner that a computable learner fails to solve.

## 1.1 What is Coordination?

For the purposes of this paper, we consider coordination as a particular process of negotiation. In other words, negotiation underpins attempts to coordinate and, by starting from coordination, to cooperate. We hope the following equational *slogan* to be of intuitive and immediate help for the reader.

(1) Cooperation = coordination + solvability.

In other words, two agents cooperate whenever they coordinate in solving some goal-problem. Of course, we all have an intuitive sense of what the word "coordination" means. When say Frank and Gertrude pass one another in the hall every day (*The Passing Game*, [Ago01, Ch. 1]), when "Sam and Sally meet daily in the park, pretending each time that it's yet another chance encounter, walking side by side in shy silence" [MO99, p. 363], when we watch the Italian volleyball team winning the 2002 Woman's World Championship or, by a counterexample, when we spend hours waiting our best friend in the wrong place or a plane on the "Malpensa" runway because the airline cannot find a gate for it.

For many purposes, this intuitive meaning is sufficient. However, in trying to characterize in a formal way the behavior of certain computable agents willing to "coordinate," it is sometimes important to have a more precise idea of what we mean by "coordination." [Dur88, MC94, MO99, RZ94] and, more recently, [CJ02, Sessions 11A and 9B] list a number of definitions that have been suggested for this term. For our purposes here, however, it is useful to begin with the following simple definition.

(2) *Coordination is managing dependencies between activities in a given context.*[1]

As the definition suggests, we believe that it is helpful to use *coordination* in a fairly specialized sense of "coordination in." In this paper we will explain exactly how and why. Moreover, the definition is consistent with the intuition that, neither if there is no interdependance nor if there is no context, there is nothing to coordinate.

This paper is organized by following the dimensions on which our vision of cooperation (1) in based on, namely: absence of coordination, absence of solvability, presence of both. The first dimension is introduced in Section 2. The second dimension is advanced in Section 3 and extended to "groups" of agents along the third dimension in Sections 4 and 5. In Section 6 we explain some consequences of our paradigms and results for learnability and cooperation in multiagent systems and present some related and future work. In Section 7 we draw the conclusions.

## 2 Individual Solvability

The paradigm of solvability on presentation in this section is a model-theoretic paradigm. Our approach follows a first-order perspective, wherein "sequences" and "language" are basic ingredients.

### 2.1 A First-Order Framework

*Sequences.* We denote the set $\{0, 1, 2, ...\}$ of natural numbers by $N$. We denote the usual linearly ordered structure with domain $N$ by $\omega$. Let $\eta$ be an infinite sequence.[2] For $i \in N$, we write $\eta|_i$ for the proper initial sequence of length $i$ in $\eta$. We write $length(\sigma)$ for the length of a finite sequence, $\emptyset$ for the finite sequence of length zero, $\sigma_i$ or also $(\sigma)_i$ for the $i$th element of $\sigma$, $0 \leq i < length(\sigma)$, and $last(\sigma)$ for the last element in $\sigma$. The set of elements in a (finite, infinite) sequence $\tau$ is denoted by $range(\tau)$.

*Language.* We write $\mathcal{L}_{form}$ to denote a first-order language with equality built up from a (countable, decidable) vocabulary $\mathcal{L}$ consisting of predicates and function symbols of various arities, along with constants symbols and countably many variables $Var = \{v_i \,|\, i \in N\}$. We write $\mathcal{L}_{sen}$ to denote the subset of $\mathcal{L}_{form}$ containing no free variables (**sentences**). We write $\mathcal{L}_{basic}$ for the set of literals of $\mathcal{L}_{form}$.

Our semantic notions are standard. Let $\mathcal{S}$ be a structure that interprets $\mathcal{L}$. Let $\models$ denote the model theoretic concept of truth in a structure. Then $\mathcal{S}$ is a **model** of $\Gamma \subseteq \mathcal{L}_{form}$, and $\Gamma$ is said to be **satisfiable in** $\mathcal{S}$, if there is a mapping $h$ from

[1] This definition is an extension by using the concept of a "context" (see below) to that of [MC94]. We refer the interested reader to the work by Malone and Crowston for further reference that connect our definition to organization theory and coordination science.

[2] By "infinite sequence" we shall always mean an $\omega$-sequence, or a total function defined on $N$.

*Var* to dom($S$) with $\mathcal{S} \models \Gamma[h]$. In this case, we call such mapping an **assignment to** $\mathcal{S}$. $\Gamma$ is **satisfiable** if it is satisfiable in some structure. An assignment $h$ to $\mathcal{S}$ is **complete** if $h$ is a mapping **onto** dom($S$). The **basic diagram** of $\mathcal{S}$ **under** complete assignment $h$ is the subset of $\mathcal{L}_{basic}$ made true in $\mathcal{S}$ via $h$.

## 2.2 Components

The paradigm is now to be introduced in detail by stepping through its basic components. The following example may help.

(3) EXAMPLE: [OdJMW97, p. 740] "First, a class of possible realities is specified in advance; the class is known to both players of the game. Nature is conceived as choosing one member from the class to be the 'actual world'; her choice is initially unknown to the scientist. Nature then provides a series of clues about this reality. These clues constitute the data upon which the scientist will base his hypotheses. Each time Nature provides a new clue, the scientist may produce a new hypothesis. The scientist wins the game if there is sufficient guarantee that his successive conjectures will stabilize to an accurate hypothesis about the reality Nature has chosen."

Observe that the game is asymmetric. In other words, the kind of agents concerned is not uniformly the same. In contrast, our successive models of coordination and group learning concern with symmetric interactions, since the agents in these models have identical structure and eventually differ only on the available actions and beliefs (respectively $\Psi$ and $\mathbf{A}$, see subsection 3.1.1 below).

We now state the components of the foregoing illustrative picture of scientific discovery. Let *SEQ* denote the collection of all the *finite* sequences over $\mathcal{L}_{basic}$.

### 2.2.1 Learners.

A **learner** (or "formal scientist") is any mapping from *SEQ* to $\mathcal{L}_{sen}$. A learner thus examines finite pieces of elementary data and—on the basis of such data, advances sentential statements. Given $\sigma \in SEQ$, learner $\Psi$ then applies to the data-stream formalized by $\sigma$ and, if $\Psi(\sigma)$ is defined, outputs the statement $\Psi(\sigma) \in \mathcal{L}_{sen}$. Intuitively, $\Psi(\sigma)$ might be thought to be the "belief" of $\Psi$ faced with $\sigma$. $\Psi(\sigma)$ might also be interpreted as $\Psi$'s hypothesis on the structure for $\mathcal{L}$, if there exists, that satisfies the formula $\sigma_1 \wedge \sigma_2 \wedge \cdots \wedge \sigma_n$, $n = \mathit{length}(\sigma)$, at position $n$ of the play history.

Learners can be partial or total, computable or uncomputable. We consider computable learners in Section 5.

### 2.2.2 Worlds.

As the first-order perspective we adopted suggests, the underlying reality or "possible worlds" a learner is concerned about are all the structures that interpret $\mathcal{L}$.

### 2.2.3 Environments.

We define an **environment** to be any infinite sequence over $\mathcal{L}_{basic}$. Thus, for all $\sigma \in SEQ$, there is an environment $e$ such that $\sigma = e|_n$ with $n = length(\sigma)$. To consider *consistent* data-streams, we need to relate them to a structure. We do it in the next definition.

(4) DEFINITION: Let structure $\mathcal{S}$ and complete assignment $h$ to $\mathcal{S}$ be given. An environment $e$ is **for** $\mathcal{S}$ **via** $h$ just in case $range(e) = \{\beta \in \mathcal{L}_{basic} \,|\, \mathcal{S} \models \beta[h]\}$. An environment $e$ is **for** $\mathcal{S}$ just in case $e$ is an environment for $\mathcal{S}$ via some complete assignment.

In other words, an environment for a structure $\mathcal{S}$ via complete assignment $h$ lists the basic diagram of $\mathcal{S}$ under $h$.

### 2.2.4 Contexts.

By a (learning) **context** we mean a set of sentences, and precisely any subset of $\mathcal{L}_{sen}$. Intuitively, given a learner, a context expresses the set of "intelligible hypotheses" $\pi$ for the learner's discovery of the underlying reality eventually made actual by Nature. Because of the relationship of the reality and a learning context, we sometimes bring them together in "problems." Precisely, a (learning) **problem** is a pair $(\mathbf{W}, \pi)$, where $\mathbf{W}$ is a collection of structures, and $\pi$ is a learning context. Given a problem, a natural question is: "Is the problem solvable?"

### 2.2.5 Success.

We say that learner $\Psi$ **solves** problem $(\mathbf{W}, \pi)$ just in case for all $\mathcal{S} \in \mathbf{W}$ and for every environment $e$ for $\mathcal{S}$, there is $\theta \in \pi$ such that:

($a$) for all but finitely many $n \in N$, $\Psi(e|_n) = \theta$ [that is, $\Psi$ **stabilizes on** $e$ to $\theta$], and

($b$) $\mathcal{S} \models \theta$ [that is, $\theta$ is **right**].

If some learner solves $(\mathbf{W}, \pi)$, then $(\mathbf{W}, \pi)$ is called **solvable**, otherwise **unsolvable**. We write "$\mathbf{W}$ is $\mathbf{Ex}^{\pi}$-solvable" in place of "$(\mathbf{W}, \pi)$ is solvable." We also write "$\mathbf{W} \in \mathbf{Ex}^{\pi}(\Psi)$" in place of "$\Psi$ solves $(\mathbf{W}, \pi)$." We define $\mathbf{Ex}^{\pi} = \{\mathbf{W} \,|\, \mathbf{W} \text{ is } \mathbf{Ex}^{\pi}\text{-solvable}\}$. This completes the formalization of the basic components of the paradigm.

In the perspective of a Multi-Agent Theory, an important question is how information-processing limitations of the learners affect the desirability of different paradigms of solvability. We provide an answer to this question by Theorem (13) in Section 5.

# 3 Coordination

We now advance a paradigm of "contextual coordination" suitable to represent and analyze coordination processes. In contextual coordination the component of learning is less evident than in the paradigm $\mathbf{Ex}^{\pi}$. Moreover, contextual coordination is a symmetric process, while learning in the paradigm $\mathbf{Ex}^{\pi}$ is asymmetric, in the sense that there is both a functional and a structural difference between the learner and the "teacher" (Nature) that there is not in contextual coordination. To coordinate, an agent (in the role of a formal scientist) has not only to guess which structure $\mathcal{S}$ the environment played by another agent (in the role of Nature) is for, but she has also to play—by interacting with the agent, an environment for a structure "similar" to $\mathcal{S}$ with respect to a given set of goals or "coordination sentences." The structure $\mathcal{S}$ is interpreted as the learner's beliefs on the environment she inputs. If either the environments played by the two agents differ or the learning agent does not guess the right sentence, each agent is faced with two alternative choices: (a) start changing his/her environment; (b) try to convince the other agent to change her/his environment. In both cases, coordination is relevant to the agents' choices.

An example of a process wherein coordination subsumes a convention over leadership is the following.

(5) EXAMPLE: (*A Game of Leadership*) Two players interact extensively by perfect knowledge of the game. Perfect knowledge is required to the agents in order to deal with a "rule of leadership," which is fixed by convention at the beginning of the game. A "leader" chooses an action and a "follower," informed of the leader's choice takes an action as well. In the first part of the game, the players' moves are aimed at making leadership emerge. Then, each player plays according to his role, either of "the leader" or "the follower." Their moves depend on the criterion of coordination and characterize the game.

A primary vehicle to extend individual solvability to a cooperative, multi-agent setting is to identify and study the basic processes involved in coordination. Are there fundamental coordination processes ("strategies") that occur in all coordinated systems? If so, how can we represent and analyze these processes? One of the advantages of the definition we have used for coordination is that it suggests a direction for addressing these and related questions.

### 3.1 Components

In order to analyze a situation in terms of coordination, it is sometimes important to explicitely identify the components of coordination in that situation. According to our informal definition of coordination (2), coordination means "managing dependencies between activities in a given context." Therefore, since activities must, in some sense that will be made clearer below, be performed by "actors," the definition implies that

(6) all instances of coordination include *agents* performing *activities* ("actions') that are *interdependent* in a *context.*

Agents are thus the first leading concept of our new paradigm. Once we have defined what is an "agent," all other components will follow in a natural way.

### 3.1.1 Agents.

An agent in a contextual coordination paradigm, or **learning agent**, is a pair $\langle \Psi , \mathbf{A} \rangle$, where $\Psi$ is any mapping from *SEQ* to $\mathcal{L}_{basic} \times \mathcal{L}_{sen}$, and $\mathbf{A}$ is a nonempty class of structures for $\mathcal{L}$. Thus, $\Psi(\sigma) = \langle (\Psi(\sigma))_0 , (\Psi(\sigma))_1 \rangle$ for all $\sigma \in SEQ$.[3] We say that $(\Psi(\sigma))_0$ is $\langle \Psi , \mathbf{A} \rangle$'s **action** (on $\sigma$), and that $(\Psi(\sigma))_1$ is $\langle \Psi , \mathbf{A} \rangle$'s **guess** (on $\sigma$). Intuitively, faced with $\sigma \in SEQ$, $\langle \Psi , \mathbf{A} \rangle$ believes $\Psi(\sigma)$ if whenever $\Psi(\sigma)$ is defined, there are some $\mathcal{A} \in \mathbf{A}$ and assignment $h$ to $\mathcal{A}$ such that $\mathcal{A} \models (\Psi(\sigma))_0[h]$ and $\mathcal{A} \models (\Psi(\sigma))_1$. Of the two components of a learning agent, the first is called **communication function** or "ability," while the second component is called **background world**. We say that $\langle \Psi , \mathbf{A} \rangle$ is **based on A**.

(7) *Remark:* As one might observe, $\Psi$ is rather general: any mapping of the right form is allowed. It would be possible to make agents more "intelligent," for example by letting them have a default (partial) 'revision' mapping on $\mathbf{A}$. In this case, our approach can be comparable to work in update semantics or belief revision, where the "beliefs" of an agent are represented by his background world. Since we take any class of first-order structures as an agent's background word, it follows that beliefs may be inconsistent in principle. In other words, we do not assume that $\mathbf{A}$ is the class of models of any "belief set" [Lev80, Gär88]. To consider belief sets is a particular constaint one might give to the paradigm, in order to study beliefs as sets of sentences. We discuss further this point in Section 6 (6.0.2).

### 3.1.2 Protocols.

We now define a protocol of communication between two learning agents. We restrict attention to dynamics based on simultaneous moves, that is, the agents make decisions at the same time, to pairwise communication only, that is, interactions involve just two agents, and to hidden guesses, that is, the agents' guesses are not explicitly communicated to the other agents. Our presentation may be generalized to a number of different protocols. We discuss further this point in Section 6 (6.0.1).

Let finite or infinite sequence $\zeta$ of length $n > k$, $k \in N$ be given. $\zeta[k]$ denotes the finite sequence $\langle \zeta_0 \cdots \zeta_k \rangle$ and ${}_k|\zeta$ denotes the sequence obtained from $\zeta$ by deleting its first $k+1$ elements.

(8) Definition: Let learning agents $\langle \Psi , \mathbf{A} \rangle$ and $\langle \Phi , \mathbf{B} \rangle$ be given.

(*a*) The **interaction sequence** (or "play" of $\langle \Psi , \mathbf{A} \rangle$ and $\langle \Phi , \mathbf{B} \rangle$ is the infinite sequence

$$D_{\Psi,\Phi} = (\langle \overline{\Psi}_i , \overline{\Phi}_i \rangle \; : \; i \in N),$$

where $\overline{\Psi}_i$ is the $i$th move of $\Psi$ and $\overline{\Phi}_i$ is the $i$th move of $\Phi$, defined by induction as follows.

[3] The notation $()_0$ and $()_1$ here and in similar circumstances should not be confused with the usage of $(\sigma)_i$ we make to denote the $i$th element of a sequence $\sigma$ of appropriate length.

i. $\overline{\Psi}_{00} = (\Psi(\emptyset))_0$ and $\overline{\Psi}_{10} = (\Psi(\emptyset))_1$;
$\overline{\Phi}_{00} = (\Phi(\emptyset))_0$ and $\overline{\Phi}_{10} = (\Phi(\emptyset))_1$.

ii. $\overline{\Psi}_{0n+1} = (\Psi(\overline{\Phi}_0[n]))_0$ and $\overline{\Psi}_{1n+1} = (\Psi(\overline{\Phi}_0[n]))_1$;
$\overline{\Phi}_{0n+1} = (\Phi(\overline{\Psi}_0[n]))_0$ and $\overline{\Phi}_{1n+1} = (\Phi(\overline{\Psi}_0[n]))_1$.

(*b*) Let $k \in N$ be given. The **interaction sequence** of $\langle \Psi , \mathbf{A} \rangle$ and $\langle \Phi , \mathbf{B} \rangle$ **starting at** $k$ is the infinite sequence

$$^{k}D_{\Psi,\Phi} = (\langle\, {}_{k}|\overline{\Psi}_i \, , {}_{k}|\overline{\Phi}_i \,\rangle \; : \; i \in N),$$

where ${}_{k}|\overline{\Psi}_i$ is the $i$th element in ${}_{k}|\overline{\Psi}$ and ${}_{k}|\overline{\Phi}_i$ is the $i$th element in ${}_{k}|\overline{\Phi}$.

The **response sequence**

$$R(\Psi, \Phi) = (\overline{\Psi}_i \; : \; i \in N)$$

is the (finite or infinite) sequence of moves by learning agent with ability $\Psi$ in response to learning agent with ability $\Phi$, and the **response sequence**

$$R(\Phi, \Psi) = (\overline{\Phi}_i \; : \; i \in N)$$

is the sequence of moves by $\Phi$ in response to $\Psi$. We write $R(\Psi_0, \Phi)$, $R(\Psi_1, \Phi)$, ${}_{k}|R(\Psi_0, \Phi)$ and ${}_{k}|R(\Psi_1, \Phi)$ to denote, respectively, $\overline{\Psi}_0$, $\overline{\Psi}_1$, ${}_{k}|\overline{\Psi}_0$ and ${}_{k}|\overline{\Psi}_1$.

### 3.1.3 Contexts.

Learning contexts in 2.2.4 are called "coordination contexts" henceforth, since they are at the service of coordination only. Coordination contexts determine what agents or designers think to be an "interesting" set of goals for coordination. We will see that these goals are "coordination sentences" in the language $\mathcal{L}_{sen}$. Given a context $\pi$, a question is always, roughly: "Is there a class of learning agents to coordinate with in $\pi$?" We provide a class of such learning agents and $\pi$ in the paradigm of group learning below.

### 3.1.4 Success.

Learning agents have to describe by induction a structure representing a part of their beliefs, eventually after a finite number of failures and a finite sequence of moves. Each description is an appropriate response to the behavior of an agent, who acts by playing on the basis of the "behavioral" language $\mathcal{L}_{basic}$. On the other hand, the agents' moves are concerned with success, in particular, with the "coordination sentence" chosen by each agent at any step of the process of coordination. According to the next definition, the agents can restart their mutual interaction finitely many often, but after the last disagreement they must eventually coordinate.

(9) Definition: Let $\pi \subseteq \mathcal{L}_{sen}$ and learning agents $\langle \Psi , \mathbf{A} \rangle$, $\langle \Phi , \mathbf{B} \rangle$ be given. We say that $\langle \Psi , \mathbf{A} \rangle$ **$\pi$-coordinates with** $\langle \Phi , \mathbf{B} \rangle$ (written: $\Psi \rightleftharpoons_{\pi} \Phi$) just in case for some $s, t \in N$, there is $\theta \in \pi$ such that:

($a$) $_s|\overline{\Psi}_0$ is an environment for some $\mathcal{A} \in \mathbf{A}$;
($b$) $_t|\overline{\Phi}_0$ is an environment for some $\mathcal{B} \in \mathbf{B}$;
($c$) for all but finitely many $n \in N$, $(_s|\overline{\Psi}_1)_n = (_t|\overline{\Phi}_1)_n = \theta$;
($d$) $\mathcal{A} \models \theta$ and $\mathcal{B} \models \theta$.

In this case, $\theta$ is said to be a **coordination sentence**. If ($c$) holds, we say that $\langle \Psi, \mathbf{A} \rangle$ ($\langle \Phi, \mathbf{B} \rangle$) **guesses** $\theta$.

Observe that $\rightleftharpoons_\pi$ is symmetric, that is, for all learning agents $\langle \Psi, \mathbf{A} \rangle$, $\langle \Phi, \mathbf{B} \rangle$, if $\Psi \rightleftharpoons_\pi \Phi$ then $\Phi \rightleftharpoons_\pi \Psi$. This seems to be a natural property of coordination. However, $\rightleftharpoons_\pi$ is neither reflexive nor transitive, as it is easy to verify.

According to Definition (9), given a coordination context $\pi$ learning agents must coordinate twice. First, the agents eventually stabilize on an environment for a structure $\mathcal{S}$ in their own background world. This step is done by the agents through their external, behavioral ability component. Second, the agents eventually stabilize to a "coordination sentence" in $\pi$ true in $\mathcal{S}$. This second step is done by the agents through their internal, hidden ("cognitive") ability. Both steps are processes in the limit, which eventually start after a finite number of disagreements.

In order to analyze a situation in terms of coordination, it is often useful to identify *evaluation criteria* for judging how well the dependencies between activities in a given context are being "managed" by the agents. For instance, some coordination processes may be faster or more accurate than others. We don't study efficiency in this paper, however. Rather, it is important to realize that accuracy depends of what situation we want to model, and that there is no single "right" way to identify the criteria of coordination. Our choice of the criterion of coordination is motivated by the next paradigm. We said that one of the advantages of the definition we have used for coordination is that it suggests a direction to represent and analyze the fundamental coordination processes and strategies. If coordination is defined as *managing dependencies in a context*, then further progress should be possible by characterizing different kinds of dependencies and identifying the winning strategies that can be used by a *group* of agents to manage them. The next section suggests the beginnings of such an analysis.

## 4 Group Solvability

There is a sense in which some overall evaluation criterion is necessarily implied by the definition of coordination. The most commonly analyzed case of managing dependency in a context occurs when an individual or group decides to pursue a goal from a set of possible goals (in fact, from a context) and then decomposes this goal into activities, or subgoals, which together will achieve the original goal. In this case, we call the process of chosing the goal **goal selection** and the process of chosing the activities **goal decomposition**.

(10) EXAMPLE: (*Top-Down Goal Decomposition*) A common kind of dependency

among activities is that a group of activities[4] are all "subtasks" for achieving some overall goal. The strategic-planning process in human organizations may be viewed as involving this kind of goal selection and goal decomposition process. In computer systems, we usually think of the goals as being predetermined, but an important problem involves how to break these goals into activities that can be performed separately. In a sense, for example, the essence of distributed information systems and multiagent systems is to decompose a "goal" into elementary versus autonomous activities. Planning in artificial intelligence (see for instance [AHT90, GHT02] and the references cited there) is another example of goal decomposition in multiagent systems. A "plan," intuitively, is what we called an "environment" in this paper.

In what follows, we advance a criterion of coordination where: (a) goal selection from a set $\pi$ of possibilities corresponds to convergence to a coordination sentence $\theta$ (the "selected goal") in $\pi$, that is, to $\pi$-coordinate; and (b) goal decomposition into sub-goals corresponds to group formation, where for each member of the group that may eventually solve the goal, there is another member that plays a sub-goal with the aim to help towards a complete solution. Goal selection and goal decomposition are instances of the following definition.

(11) DEFINITION: Let $\pi \subseteq \mathcal{L}_{sen}$, nonempty class $\mathbf{A}$ of structures, and set $\Sigma$ of learning agents based on $\mathbf{A}$ be given.

(*a*) $\Sigma$ **$\mathbf{Gr}^{\pi}$-solves $\mathbf{A}$** (written: $\mathbf{A} \in \mathbf{Gr}^{\pi}(\Sigma)$) just in case:

i. for all $\langle \Psi, \mathbf{A} \rangle, \langle \Phi, \mathbf{A} \rangle \in \Sigma$, $\langle \Psi, \mathbf{A} \rangle$ $\pi$-coordinates with $\langle \Phi, \mathbf{A} \rangle$;
ii. for all $\langle \Phi, \mathbf{A} \rangle \in \Sigma$ and for all $\mathcal{A} \in \mathbf{A}$, there are $\langle \Psi, \mathbf{A} \rangle \in \Sigma$ and $k \in N$ such that ${}_k|R(\Psi_0, \Phi)$ is an environment for $\mathcal{A}$.

In this case, $\mathbf{A}$ is said to be **$\mathbf{Gr}^{\pi}$-solvable**.

(*b*) $\mathbf{Gr}^{\pi} = \{\mathbf{A} \mid \mathbf{A} \text{ is } \mathbf{Gr}^{\pi}\text{-solvable}\}$.

The definition identifies the key components of a paradigm of **group-learning** (or learning by groups).[5] We call a **group** ("based on $\mathbf{A}$") the set $\Sigma$ of learning agents that satisfies clause (*a*).ii of the definition. Goal selection is modeled by clause (*a*).i. Goal decomposition follows from a composition of clauses (*a*).i and (*a*).ii. According to this interpretation of goal selection and goal decomposition within the paradigm $\mathbf{Gr}^{\pi}$, an example of group formation and, as a consequence, of goal selection and goal decomposition is provided by the proof of the following theorem.

(12) THEOREM: [AM01, Ago01] Let $\pi \subseteq \mathcal{L}_{sen}$ and let countable class $\mathbf{A}$ of structures be such that for every $\mathcal{A} \in \mathbf{A}$ there is $\theta \in \pi$ such that $\mathcal{A} \models \theta$. Then $\mathbf{A}$ is $\mathbf{Gr}^{\pi}$-solvable.

[4] A group of agents, by argument (6).

[5] In fact, it identifies a *family* of paradigms over the parameter $\pi$. This remark partially justifies the title of this article.

The proof of the theorem provides a procedure of coordination between two agents sharing a "convention of leadership." We refer to [AM01] for the formal proof. Here is a sketch.

Let $\pi = \{\theta_i \in \mathcal{L}_{sen} \,|\, i \in N\}$ and let $\{\mathcal{A}_j \,|\, j < \omega\}$ be an infinite-repetition enumeration of the countably many members of $\mathbf{A}$. Let $e^{\mathcal{A}_n}$ be an environment for $\mathcal{A}_n \in \mathbf{A}$ and let $\theta_m \in \pi$ be such that $\mathcal{A}_n \models \theta_m$. We describe a limiting game of a learning agent $\langle \Psi_{<n,m>}, \mathbf{A} \rangle$ who interacts with a learning agent $\langle \Psi_{<k,j>}, \mathbf{A} \rangle$ in order to $\pi$-coordinate. To make $\pi$-coordination possible, we assume that the agents stipulate a convention regulating leadership. This is assumed to happen before the beginning of the game. The convention we refer to in this proof is *lexicographic order on the agents indexing.*[6] According to this convention, the leader is who among the agents has the lower index. So, in particular, $\langle \Psi_{<n,m>}, \mathbf{A} \rangle$ is the leader iff $<n,m>$ is lower than $<k,j>$ according to lexicographic order. The assessment of the leadership is done by $\langle \Psi_{<n,m>}, \mathbf{A} \rangle$ soon after the first two stages of the agents' interaction are played. Now the game begins. By using the communication function, $\langle \Psi_{<n,m>}, \mathbf{A} \rangle$ starts moving $v_n \doteq v_n$ to communicate to $\langle \Psi_{<k,j>}, \mathbf{A} \rangle$ that "if she will be the leader then she completely describe $\mathcal{A}_n$." This is done by enumerating environment $e^{\mathcal{A}_n}$. At the second interaction stage and, again, by the first component of communication function (since guesses are "hidden," communicating them to the partner is useless), $\langle \Psi_{<n,m>}, \mathbf{A} \rangle$ moves $v_m \doteq v_m$ to communicate to $\langle \Psi_{<k,j>}, \mathbf{A} \rangle$ that "if she will be the leader then she guesses (interacting with $\langle \Psi_{<k,j>}, \mathbf{A} \rangle$) $\theta_m$ forever." From the third stage of the interaction onwards, both the agents will know who is the leader. If $\langle \Psi_{<n,m>}, \mathbf{A} \rangle$ is not the leader, namely, $\langle \Psi_{<n,m>}, \mathbf{A} \rangle$ is "the follower" so that $\langle \Psi_{<k,j>}, \mathbf{A} \rangle$ results to be the leader, then $\Psi_{<n,m>}$ simply copies the moves of $\Psi_{<k,j>}$, precisely those moves $\Psi_{<k,j>}$ advances by the first component of his communication function. Also, $\Psi_{<n,m>}$ outputs the guess $\theta_j$ in $\pi$ produced by $\Psi_{<k,j>}$. Because guesses are "hidden," in particular to $\langle \Psi_{<n,m>}, \mathbf{A} \rangle$—in fact $\Psi_{<n,m>}$'s input domain is *SEQ*, $\theta_j$ is communicated by $\langle \Psi_{<k,j>}, \mathbf{A} \rangle$ to $\langle \Psi_{<n,m>}, \mathbf{A} \rangle$ via a literal of the form $v_j \doteq v_j$. On the other hand, if $\langle \Psi_{<n,m>}, \mathbf{A} \rangle$ is the leader, then she is expected to act consistently to what announced in advance by her first two moves, namely, she starts enumerating an environment for $\mathcal{A}_n$ (specifically: $e^{\mathcal{A}_n}$), and starts guessing sentence $\theta_m$ forever after. The game ends in the limit.

Set $\Sigma = \{\langle \Psi_{<n,m>}, \mathbf{A} \rangle \,|\, \mathcal{A}_n \models \theta_m\}$. Our claim that $\Sigma$ witnesses $\mathbf{A} \in \mathbf{Gr}^{\pi}(\Sigma)$ is then easy to verify by the following argument. Observe that $\Sigma$ is nonempty. For all $\mathcal{A} \in \mathbf{A}$ and for all $\langle \Psi_{<n,m>}, \mathbf{A} \rangle \in \Sigma$, let $k > n$ be such that $\mathcal{A}_k = \mathcal{A}$. Since $\{\mathcal{A}_j \,|\, j < \omega\}$ is an infinite-repetition of the structures in $\mathbf{A}$, such $k$ exists. Then, $\langle \Psi_{<k,m>}, \mathbf{A} \rangle \in \Sigma$ is the leader who enumerates an environment for $\mathcal{A}$. Moreover, for all but finitely many $t \in N$, $(R(\Psi_{<k,m>1}, \Psi_{<n,m>}))_t = \theta_m$. On the other hand, it is easy to verify that the convention on leadership stipulated by any two

[6]One can identify several different types of conventions, or "agreement conditions," which may be used in different coordination scenarios. It is always assumed, however, that the agents will settle on the convention to be used *before* the actual coordination process proper begins. The selection of a piece of knowledge as shared by the agents is thus an important but *meta-coordination* issue, which falls outside the scope of the present work.

agents in $\Sigma$ ensures that each pair of learning agents in $\Sigma$ $\pi$-coordinate. Hence, $\mathbf{A} \in \mathbf{Gr}^{\pi}(\Sigma)$. This completes the sketched proof.

The technique for group formation provided by the proof of the theorem suggests new ways of structuring goal selection and, expecially, goal decomposition in multiagent systems. In particular, the technique suggests that in human organizations as well as in computer systems and networks, distribution of activities might sometimes be better off not as strict hierarchies but as multilayered systems in which any agent at one level—namely, an agent who is not the leader in interacting with a group of agents, could direct the activities of any other agent within a larger group of agents at the next level down. In other words, the technique for group formation as depiced by the proof above suggests, we believe, an innovative way to intend group formation through goal decomposition. It also gives an intuition of how we intend distribution of agents' activities within a system of "peers." In fact, our paradigm reflects the "peer-to-peer" vision [PSW01].

# 5 Computable Solvability

How do information-processing limitations of a learner affect the desirability of different paradigms of solvability? Are some methods of coordination appropriate for coordinating people that would not be appropriate for coordinating computable agents? For example, the method we used to prove Theorem (12) is appropriate for coordinating people, but it is not appropriate for coordinating computable agents, because the undecidable choice of satisfiable goals. In this section, we present an appropriate method for coordinating computable agents. Moreover, we provide an answer to the first question above with regard to the paradigm $\mathbf{Ex}^{\pi}$. For doing this, we address the study of individual learning and of coordination by narrowing the interpretation of learners and learning agents to computable objects.

## 5.1 Computable Learners

We now prove that there are $\mathbf{Ex}^{\pi}$-solvable problems that a computable learner fails to solve. Our result shows the limits of individual solvability in a computable system, and fixes a cut-off in the competence of computable learners relative to a particular context.

That for some $\pi \subseteq \mathcal{L}_{sen}$ there is a $\mathbf{Ex}^{\pi}$-solvable class of structures that cannot be $\mathbf{Ex}^{\pi}$-solved by any computable learner is a corollary of the following theorem.

(13) THEOREM: [AM01, Ago01] Suppose that $\mathcal{L}$ is limited to a binary predicate symbol, a unary function symbol, and two constants symbols. For every countable class $\Sigma$ of learners there is a problem $(\mathbf{K}, \pi)$ with the following properties.[7]

- $(a)$ $\mathbf{K}$ is $\mathbf{Ex}^{\pi}$-solvable.
- $(b)$ No member of $\Sigma$ $\mathbf{Ex}^{\pi}$-solves $\mathbf{K}$.

[7] Recall that we say that class $\mathbf{W}$ of structures is $\mathbf{Ex}^{\pi}$-solvable iff problem $(\mathbf{W}, \pi)$ is solvable.

The proof of this theorem can be found in [Ago01, Thm. 4.(20)].

(14) COROLLARY: Let $\mathcal{L}$ be as in Theorem (13). Then there are $\pi \subseteq \mathcal{L}_{sen}$ and $\mathbf{Ex}^{\pi}$-solvable class $\mathbf{K}$ of structures such that no computable learner $\mathbf{Ex}^{\pi}$-solves $\mathbf{K}$.

*Proof:* Observe that the class of computable learners form a countable set. Then apply Theorem (13). ■

## 5.2 Computable Agents

In the light of group formation among computable agents, we now present a result that overcomes the limitations of applicability of the procedure of coordination developed in the proof of Theorem (12). To state and prove the result, we rely on the following technical definition.

(15) DEFINITION: Let $(\mathcal{A}_n : n \in N)$ be a countable class of structures. We say that $(\mathcal{A}_n : n \in N)$ is **uniformly recursive** just in case there is a total computable function $E$ from $N \times N$ to the set of infinite sequences over $\mathcal{L}_{basic}$ such that for every $n \in N$, $\lambda i \,.\, E(n,i)$ is a environment for $\mathcal{A}_n$.

To illustrate, one might think to structure $\mathcal{A}_n$ as a classification structure, for example a tree, a graph, eventually a "concept hierarchy" [BBH98]. $\mathcal{A}_n$ may be used by an agent with background world included in $(\mathcal{A}_n : n \in N)$ to classify $n$ documents, or also a set of documents according to $n$ related concepts. In the latter case, a suitable classification theory $\Gamma_c \subseteq \mathcal{L}_{form}$ should have $\mathcal{A}_n$ as a model. Observe:

(16) LEMMA: Let countable structure $\mathcal{S}$ and environment $e$ be given. Suppose that $e$ is for $\mathcal{S}$. If $e$ is recursive, then $\mathcal{S}$ is computable.

The lemma motivates the following definition of a "computable agent."

(17) DEFINITION: A learning agent $\langle \Psi \,, \mathbf{A} \rangle$ is **computable** just in case $\Psi$ is computable and $\mathbf{A}$ is uniformly recursive.

We are now ready to set the computable (recursive) version of the paradigm of group-learning $\mathbf{Gr}^{\pi}$.

(18) DEFINITION: Let $\pi \subseteq \mathcal{L}_{sen}$ be given. $[\mathbf{Gr}^{\pi}]^{rec} = \{\mathbf{K} \,|\, \mathbf{K} \in \mathbf{Gr}^{\pi}(\Sigma)\}$, where $\Sigma$ is a set of computable learning agents based on $\mathbf{K}$.

Mimicking the proof of Theorem (12) and taking $\lambda i \,.\, E(n,i)$ to be a computable environment for $\mathcal{A}_n$, we have:

(19) THEOREM: [AM01, Ago01] Suppose that $(\mathcal{A}_n : n \in N)$ is a uniformly recursive class of structures with infinite repetitions. Let $\pi = \{\theta_i \in \mathcal{L}_{sen} \,|\, i \in N\}$ be such that there is a total computable function $f : N \longrightarrow N$ such that for every $i \in N$, $\mathcal{A}_i \models \theta_{f(i)}$. Then $(\mathcal{A}_n : n \in N)$ is $[\mathbf{Gr}^{\pi}]^{rec}$-solvable.

The proof of the theorem provides a procedure of coalition formation similar to that of Theorem (12), but computable.

# 6 Discussion

Even though our paradigms omit many important aspects of human coordination and distributed computer systems, they help illuminate a wide range of phenomena. For instance, the models are consistent with a number of previous work about individual learning [OdJMW97, MO98], coordination [Ago00a, MO99] and coalition/group formation [APM00, RS98, Syc89, Ago00b]. These models also help analyze design alternative for distributed systems of "peers," and they suggest ways of analyzing the structural changes associated with introducing the inductive approach to coordination into multiagent systems and organizations' dynamics. The proofs of Theorem (12) and Theorem (19) are examples of limiting games, where there is a winning strategy for pairwise coordination on a set of agents with common knowledge. If two agents agree upon a convention of leadership, that is, they eventually answer in the same way to a question like "who should be the leader now?" and, moreover, they use that strategy, then coordination and group formation are guaranteed to exists.

In addition to the processes described above for managing dependencies within a context, two other processes deserve specific attention: *communication* and *belief revision*. So far we have analyzed these processes as implicit ways of managing specific dependencies. For instance, belief revision can be viewed as a way the learning agents have of managing the structures in their background word in order to coordinate. However, because of the importance of these two processes in almost all instances of coordination and teamwork, we describe them separately here.

### 6.0.1 Communication

One obvious way of generating new coordination procedures is by considering alternative forms of communication for all the places in a coordination process where information, or also "knowledge," needs to be transferred.[8] We restricted attention to communication functions where actions are "basic" (i.e., formulas in $\mathcal{L}_{basic}$) and where the discovery machinery runs over sentences (i.e., formulas in $\mathcal{L}_{sen}$). Our presentation may be generalized to a number of different protocols—examples are sequential moves, $n$-person communication, explicit guesses—, with regard to what requires the particular scenario or application. Moreover, actions available to the agents can be limited to just atomic formulas (no negations), or enriched to include universal or other kinds of formulas. Many alternative definitions of an "environment" are possible. We chosen to consider "complete" environments in this paper as it seems the more simple solution. It is worth noting

[8] "*Knowledge* is information transformed into capabilities for effective action. In effect, knowledge is action."(P. C. Murray) Retrieved August 8, 2001 from the World Wide Web: http://www.ktic.com/topic6/13-term0.htm.

that "real" multiagent environments may suffer omissions, erroneous intrusions, or both omissions and intrusions.

Our coordination framework may also highlight new aspects of communication. For instance, think to the problem of "meaning negotiation" [Ago02, BW02]. In any system of autonomous and distributed agents willing to coordinate, autonomy is the core condition for agents to make independent assignments of meaning to world objects. The problem of meaning negotiation arises from such assumption, which is also related in our framework to the assumption that the agents' guesses are hidden. For example, when we view communication as a way of managing a seeker/provider relationship, we may be concerned about how to make meaning relative to the seeker's needs "usable" by the provider in order to fulfill the seeker's requests, that is, to coordinate. How, for instance, can the agents establish a common language interpretation over a shared context, that allows them to communicate meaning? Moreover, how do such a shared context eventually emerge by meaning negotiation? The procedures of coordination we have advanced by the proofs of Theorems (12) and (19) at least provide a first answer to these questions. They also address the solution of a problem of *knowledge management*, as knowledge management is "the process whereby knowledge seekers are linked with knowledge sources, and knowledge is transferred." (A.J. Murray)[9]

### 6.0.2 Belief revision

At the borderline of our approach is the problem of belief changes—how an agent should revise her beliefs to coordinate upon learning new information. The arrival of data $\sigma$, however, modifies an agent's choice $\mathcal{A}$ of a structure in the agent's background word according to some fixed, but implicit scheme of belief revision. The resulting new beliefs are denoted $\mathcal{A} \dotplus \sigma$ and signifies the impact of $\sigma$ on $\mathcal{A}$. Belief revision is so far mostly defined for agents whose background world is a set of sentences, possibly a "theory" (see for instance [Lev80, Gär88, MO98]). On the other hand, we have seen that our paradigms consider an agent's beliefs to be represented by a class of structures. New paradigms of coordination among 'revision-based' rather than 'knowledge-based' agents may then be advanced from our work by answering the following question, among others: Is there a natural (justifiable on intuitive grounds), semantic generalization of belief revision in the context of inductive inference, in which revision applies directly to classes of structures? If the answer is "yes", how to extend the mathematics of inductive coordination as presented in this paper so as to be able to frame appropriate conditions on the revision operator so that its use has a "rational" quality to it? We leave these important questions to future work

[9]Retrieved August 8, 2001 from the World Wide Web: http://www.3-cities.com/~bonewman/what-is.htm.

# 7 Conclusions

The key point of the foregoing discussion, and indeed of much of this paper, is that the concepts of a coordination theory as that we have presented can help identify similarities among concepts in different disciplines, among other we have mentioned planning, knowledge management, and of course coordination, cooperation and (meaning) negotiation in a multiagent setting. These similarities, in turn, suggest how ideas can be transported back and forth across disciplinary boundaries and where opportunities exist to develop even deeper analyzes.

To summarise, in this paper we have provided a framework to study the important topic of group-learning, where pairwise coordination is modeled over sets of belief-based agents, eventually able to communicate the complete (see "environments") and local (that is, of each agent) meaning of their beliefs by using the literals of a common first-order language. We have defined a communication protocol for agents that always move simultaneously, and all relevant moves they do are application of functions in mathematical sense; no randomness ever intervenes. Our protocol is illustrative and it is suitable for modeling group situations as happens for instance in commercial transactions. We have provided some examples of coalition formation, and have showed how *coalitions by coordination* may be used to solve some classes of structures, in a precise sense we have made actual within the model-theoretic tradition of formal learning theory. As an important consequence for real applications to multiagent systems, we have showed how certain classes of structures may be restricted to a computable setting, thus providing a suitable paradigm to analyze coordination and coordination mechanisms, for example "by leadership," among groups of agents with an internal "schema," such as a classification, beliefs, a concept hierarchy or any other type of structured or semi-structured background knowledge.

# References


[Ago00a] A. Agostini. Notes on formalizing coordination. In E. Lamma and P. Mello, editors, *AI*IA 99: Advanced in Artificial Intelligence*, pages 285–296. Springer-Verlag LNAI 1792, 2000.

[Ago00b] A. Agostini. Team-solvability: A model-theoretic perspective. In W. Horn, editor, *Proceedings of the Fourteenth European Conference on Artificial Intelligence (ECAI-00)*, pages 333–337, Amsterdam, 2000. IOS Press.

[Ago01] A. Agostini. *Paradigms of Coordination and Solvability.* PhD thesis, Università di Siena, 2001.

[Ago02] A. Agostini. Coordination through inductive meaning negotiation. Technical report, ITC-irst, Trento, Italy, November 2002.

[AHT90] J. F. Allen, J. Hendler, and A. Tate, editors. *Readings in Planning*, San Mateo, CA, 1990. Morgan Kaufmann Publishers.

[AM01] A. Agostini and F. Montagna. Group-solvability. Technical report, Università di Siena, Siena, Italy, October 2001. Electronically available at
`http://sra.itc.it/people/agostini/`.

[APM00] L. Amgoud, S. Parsons, and N. Maudet. Arguments, dialogue, and negotiation. In W. Horn, editor, *Proceedings of the Fourteenth European Conference on Artificial Intelligence (ECAI-00)*, pages 338–342, Amsterdam, 2000. IOS Press.

[BBH98] A. Büchner, D. Bell, and J. Hughes. A contextualised object data model based on semantic values. In *Proceedings of the Eleventh International Conference on Parallel and Distributed Computing Systems (PDCS-98)*, pages 171–176, 1998.

[BW02] P. Bouquet and M. Warglien, editors. *Meaning Negotiation - Papers from the AAAI Workshop*, Technical Report WS-02-09, Menlo Park, CA, 2002. AAAI Press.

[CJ02] C. Castelfranchi and W. Lewis Johnson, editors. *Proceedings the First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-02)*, New York, NY, 2002. ACM Press.

[Dur88] E. H. Durfee. *Coordination of Distributed Problem Solvers.* Kluwer Academic Publishers, Boston, MA, 1988.

[Gär88] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States.* The MIT Press, Cambridge, MA, 1988.

[GHT02] M. Ghallab, J. Hertzberg, and P. Traverso, editors. *Proceedings of the Sixth International Conference on AI Planning and Scheduling (AIPS-02)*, Menlo Park, CA, 2002. AAAI Press.

[JFL+01] N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge. Automated Negotiation: Prospects, methods and challenges. *International Journal of Group Decision and Negotiation*, 10(2):199–215, 2001.

[JORS99] S. Jain, D. Osherson, J. Royer, and A. Sharma. *Systems That Learn - An Introduction to Learning Theory, 2nd edition.* The MIT Series in Learning, Development, and Conceptual Change, v. 22. The MIT Press, Cambridge, MA, 1999.

[Lev80] I. Levi. *The Enterprise of Knowledge.* The MIT Press, Cambridge, MA, 1980.

[MC94] T. W. Malone and K. Crowston. The Interdisciplinary Study of Coordination. *ACM Computing Surveys*, 26(1):87–119, 1994.

[MO98] E. Martin and D. Osherson. *Elements of Scientific Inquiry.* The MIT Press, Cambridge, MA, 1998.

[MO99] F. Montagna and D. Osherson. Learning to coordinate: A recursion theoretic perspective. *Synthese*, 118(3):363–382, 1999.

[OdJMW97] D. Osherson, D. de Jongh, E. Martin, and S. Weinstein. Formal Learning Theory. In J. van Benthem and A. ter Meulen, editors, *Handbook of Logic and Language*, pages 737–775. Elsevier Science B.V., 1997.

[PSW01] M. Parameswaran, A. Susarla, and A.B. Whinston. P2P networking: An information-sharing alternative. *IEEE Computer*, 34(7):31–38, 2001.

[RS98] S. K. Rustogi and M. P. Singh. The bases of effective coordination in decentralized multi-agent systems. In J. P. Müller, M. P. Singh, and A. S. Rao, editors, *Intelligent Agents V - Agent Theories, Architectures, and Languages*, pages 149–161. Springer-Verlag LNAI 1555, 1998.

[RZ94] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers.* The MIT Press, Cambridge, MA, 1994.

[SK95] O. Shehory and S. Kraus. Task allocation via coalition formation among autonomous agents. In C. S. Mellish, editor, *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-95)*, pages 655–661, San Mateo, CA, 1995. Morgan Kaufmann.

[SLA+99] T. W. Sandholm, K. Larson, M. Andersson, O. Shehory, and F. Tohmé. Coalition structure generation with worst case guarantees. *Artificial Intelligence*, 111(1-2):209–238, 1999.

[Syc89] K. P. Sycara. Argumentation: Planning other agents' plans. In N. S. Sridharan, editor, *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI-89)*, pages 517–523, San Mateo, CA, 1989. Morgan Kaufmann.

[ZS97] D. Zeng and K. Sycara. How can an agent learn to negotiate? In J. P. Müller, M. J. Wooldridge, and N. R. Jennings, editors, *Intelligent Agents III - Agent Theories, Architectures, and Languages*, pages 233–244. Springer-Verlag LNAI 1193, 1997.