The Intentions of Teams:

Team Structure, Endodeixis, and Exodeixis

Munindar P. Singh[1]

Abstract. Teams arise in a number of important multiagent applications. Several theories of intentions for teams have been proposed. By and large, these theories tend to model team intentions exclusively on the basis of mental concepts, and fail to acknowledge the internal structure of teams.

We present a formal theory of intentions for teams that considers the structure of teams explicitly. In this context, we distinguish between exodeictic and endodeictic intentions, which are conceptualized as pointing outward or inward from a team. These concepts are formalized in a framework that models the structure of teams in terms of their members' commitments and coordination requirements. In this way, our approach combines mental and social concepts in a principled manner. We describe some postulates concerning intentions and structure, and give technical results establishing or falsifying these postulates with different definitions.

1 Introduction

Intentions have drawn much attention in multiagent systems research. With few exceptions, previous theoretical work has generally considered only the usual mental primitives of traditional AI. However, multiagent systems are inherently social entities. We restrict attention to *teams*, which are multiagent systems that are viewed as having different members playing specific roles and usually cooperating to achieve some higher end. Although mental concepts such as intentions apply on teams, they must be properly related to social concepts. To do so is the main objective of this paper.

This exercise is of theoretical and practical importance, because not only is this issue of centrality to DAI, but implemented systems involving the intentions of teams now exist, e.g., STEAM [19] and ARCHON [12]. Because of the limitations of the present theories, existing systems were forced to invent additional representations to effectively capture the social dimension. By including this in our theory, we can hope to offer a more accurate and implementable theory, which will facilitate the designer's task while providing rigor and flexibility. We motivate a set of definitions of intentions of teams that combine aspects of previous work on intentions, coordination, social commitments, and structure.

Social Stance. Two powerful and well-known ways of looking at agency are the *intentional stance* [6, 14], and the *knowledge level* [15]. These approaches legitimize the ascription of intentions to complex physical systems. We implicitly adopt these approaches in defining the intentions of teams.

[1] Department of Computer Science, North Carolina State University, Raleigh, NC 27695-7534, USA, singh@ncsu.edu

ECAI 98. *13th European Conference on Artificial Intelligence*
Edited by Henri Prade
Published in 1998 by John Wiley & Sons, Ltd.

Traditionally, the intentional stance supervenes of the *design* stance, and the knowledge level on the *symbol* level. However, in the case of teams, this hierarchy is less clear. Although a team may be viewed as a single agent, the members of a team exist independently and are themselves intentional. Intuitively, when a team is opened up with the design stance, we find not mere mechanisms, but other agents, some of which may be teams. For this reason, we introduce the *social* stance or level, which plays a central role in our approach.

Just as the intentional stance justifies mental constructs, the social stance enables social ones, such as social commitments. For simplicity, we take the social stance as including organizational aspects [2].

We take the notion of individual agents as given, and present a recursive definition of teams, which are also treated as agents..

Mutual Beliefs. Traditional theories, such as [13, 11], involve the notion of *mutual belief* (essentially the same as *common knowledge* for our purposes). Roughly, a set of agents mutually believe p iff each of them believes p, and each of them believes that each of them believes p, and so on, *ad infinitum* [9]. Traditional approaches require mutual beliefs among the team members essentially to achieve the effect that can be more simply attained through social commitments. In fact, it is known that in settings with asynchronous, unreliable or unboundedly delayable communication, mutual beliefs cannot be created. They exist only if present from the start [4]. Thus, mutual beliefs are used primarily to establish impossibility results for distributed computing protocols. It is puzzling that the basis for impossibility results would form a cornerstone of theories that seek application in real environments!

Exodeixis and Endodeixis. The presence of structure in teams causes a systematic variation in meaning, which has led to much confusion in some previous work. In one sense, the intentions of a team apply outside of the team. These are the intentions of the team as viewed by others, who effectively view the team as a single monolithic autonomous entity. This is the sense that is suggested in statements such as "North Korea intends to invade South Korea." In another sense, the intentions of a team apply within the team. These are the intentions of the team as viewed by the team itself or by others with an interest in the team's composition. Here the team is thought of as having structure. Further, the members of the team are committed to the intentions of the team and may cooperate according to their roles. This is the sense that is suggested in statements such as "The North Koreans intend to invade South Korea."

We introduce the term *exodeixis* (from *exo-* "outward" and *deixis* pointing) for the first phenomenon, and *endodeixis* (from *endo-* "inward") for the second. The corresponding intentions are called exodeictic and endodeictic, respectively.

Organization. Section 2 introduces some background material. Section 3 formalizes team structure. Section 4 presents a formalization of intentions as motivated above. Section 5 presents several postulates concerning team intentions and structure, and show which of our formal definitions support which postulates. Section 6 concludes with a discussion.

2 Interaction-Oriented Programming

The present approach fits naturally in our ongoing research program of *interaction-oriented programming (IOP)*. IOP is about abstractions and techniques for programming interactions among agents. It has three layers, from lower to upper:

- *coordination*, which enables the agents to operate in a shared environment
- *commitment*, which reflects the social stance, capturing the agents' obligations, social structure, and norms
- *collaboration*, which combines the intentional and the social stances.

Some informal concepts, such as competition, may be classified into different layers: bidding in an auction requires no more than coordination among bidders and fairly rigid commitments between the bidders and the auctioneer, whereas commerce involves flexible commitments, and negotiation involves sophisticated protocols. Team intentions reside in the collaboration layer and build on top of coordination and commitment.

2.1 Coordination

Recognizing that the agents will be designed by different parties and will behave autonomously, we require that only limited knowledge of the agent's construction be used in coordination. This knowledge is the form of a compact *skeleton* for each agent, which includes its publicly visible *events* along with constraints on whether they can be reordered, triggered, or prevented. Coordination requirements are formulae expressing relationships among events. The formal language is simple, but can be processed automatically to ensure the occurrence or mutual ordering of events to satisfy the stated requirements.

Our specification language is propositional logic augmented with the *before* $(\cdot)$ temporal operator. *Before* is essentially a dual of the more conventional "until" operator. R3 suggests an enabling condition or a data flow from e to f. Table 1 presents some common examples. Relationships may involve multiple events. R8 captures requirements such as that if an agent does something (e), but another agent does not match it with something else (f), then a third agent can perform g. This is a typical pattern of coordination where an agent handles contingencies resulting from the other agents' actions. Additional details of semantics and processing are available elsewhere [17].

2.2 Commitments

Agents can commit to each other. The *debtor* commits to the *creditor* to bring about the *discharge condition*. Commitments are formed in a *context*, which is typically the enclosing team:

Definition 1 The formula $\mathsf{C}(x, y, p, \tau)$ denotes a commitment, where x is the debtor, y the creditor, τ the context, and p the discharge condition.

Operations on Commitments. These include:

O1. *Create* instantiates a commitment; it is typically performed as a consequence of an agent adopting a role or by exercising a social policy (explained below).
O2. *Discharge* satisfies the commitment; it is performed by the debtor concurrently with the (possibly physical) actions that lead to the given condition being satisfied.
O3. *Cancel* revokes the commitment. It can be performed by the debtor.
O4. *Release* essentially eliminates the commitment. This is distinguished from both *discharge* and *cancel*, because *release* does not mean success or failure of the given commitment, although it lets the debtor off the hook. The *release* action may be performed by the context or the creditor of the given commitment.
O5. *Delegate* shifts the role of debtor to another agent within the same context, and can be performed by the new debtor or the context.
O6. *Assign* transfers a commitment to another creditor within the same context, and can be performed by the present creditor (if authorized) or the context.

Through a minor abuse of notation, we write the above operations also as propositions, indicating their successful execution.

Social Policies. These are conditional expressions involving commitments and operations on commitments. Policies are useful, because they lead to a decoupling of the decision-making among the agents. These are essentially like operating procedures in general organizations. Each party can act without having to confirm every decision with others, unless the policies call for explicit confirmation. Agents can commit to social policies themselves, resulting in meta-commitments. Social policies are the basis for the commitments, including the commitments that are instantiated when a team is created. Ultimately, these policies must be rooted in the society from which the agents are drawn to compose a team.

3 Team Structure

One of the most important aspects of teams is their structure. Teams typically have several distinct roles. Both the exodeictic and endodeictic intentions of a team are constrained by its structure. This obvious point is worth emphasizing, because it has been ignored by previous approaches. Our approach admits two aspects of team structure, in terms of commitment and coordination, respectively. Both are represented via the interactions among team-members.

Committed Interactions. Certain high-level interactions among team-members occur at the level of their social commitments to each other. These interactions involve the operations on commitments as described in section 2.2. Most operations on commitments are carried out through illocutionary acts between agents [1]. These operations occur in a context where the agents' prior commitments include the applicable social policies.

Coordinating Interactions. Another subclass of interactions involves the establishment of various conditions in the world by some members that other members rely on. These could be because of constraints of the physical environment or because the members are falling into various habits that have not yet been raised to the level of explicit social commitments [20].

	Name	Description	Formal notation
R1	e is required by f	If f occurs, e must occur before or after f	$e \vee \overline{f}$
R2	e disables f	If e occurs, then f must occur before e	$\overline{e} \vee \overline{f} \vee f \cdot e$
R3	e feeds or enables f	f requires e to occur before	$e \cdot f \vee \overline{f}$
R4	e conditionally feeds f	If e occurs, it feeds f	$\overline{e} \vee e \cdot f \vee \overline{f}$
R5	guaranteeing e enables f	f can occur only if e has occurred or will occur	$e \wedge f \vee \overline{e} \wedge \overline{f}$
R6	e initiates f	f occurs iff e precedes it	$\overline{e} \wedge \overline{f} \vee e \cdot f$
R7	e and f jointly require g	If e and f occur in any order, then g must also occur	$\overline{e} \vee \overline{f} \vee g$
R8	g compensates for e failing f	if e happens and f does not, then perform g	$(\overline{e} \vee f \vee g) \wedge (\overline{g} \vee e) \wedge (\overline{g} \vee \overline{f})$

Table 1. Example Coordination Relationships

Teams Formalized. The following definition formalizes the intuition that the structure of a team is reflected in the constraints on the interactions of its members.

Definition 2 A team $\tau \triangleq \langle \{x_1, \ldots, x_n\}, S, R \rangle$, where x_i are agents, S is a set of social commitments among x_i, and R is the set of coordination relationships among x_i.

The x_i are the *members* of τ (notated, $x_i \in \tau^m$). To be more precise, we would define abstract teams in terms of their roles, and instantiate the roles with agents in order to create concrete teams. Here we go directly to concrete teams for expository ease, and to save space. Many of the commitments are realized through coordination actions, but coordination relationships can exist that are not explicitly part of the commitments. The implicit coordination relationships can be lifted into proper commitments when the team matures. Since the coordination component is used within the commitment component, R can be empty even in nontrivial cases, but S is empty only in trivial teams.

Example 1 A soccer team is a team with 11 players (in different roles). The players are committed to each other and the team, e.g.,

- the goal-keeper will obstruct any attempts at putting the ball in the goal he defends
- the center-left will obey the center-forward's signal to advance, i.e., the signal causes the *creation* of a commitment to advance
- each player will rush back to defend if the opposing team has a corner kick, i.e., the corner causes the *creation* of a commitment to return and a *release* from any prior commitment to advance.

There are coordination requirements underlying the commitments, e.g., the center-left will stay to the left and back of the center-forward while obeying his signal. Some requirements may not be explicitly committed to, e.g., the center-left rushes to assist the center-forward if the latter is surrounded, but that's just on his own accord. ▮

More complex teams, e.g., in business organizations, would use the other operations on commitments as well.

4 Team Intentions

We motivated the distinction between endodeictic and exodeictic intentions of teams in section 1. We now give formal definitions for these concepts, and compare them with regard to some technical properties that involve them and the structure of teams.

4.1 Formal Framework

Our formal framework is the usual formal language and model involving branching time and actions. We only have space to cover this informally; for a detailed exposition, please consult [18]. Briefly, the model, M, consists of *moments* arranged according to temporal precedence in a branching structure. Each moment represents a possible state of the world (and has associated intentions and commitments of each agent). Each branch or *path* (rooted at a moment) represents a possible execution of the multiagent system and its environment. The propositional operators are standard. For simplicity, we do not highlight the temporal operators of the language and assume they are included as needed in the propositions.

We define modal accessibility relations for intentions $\mathbf{I}$ and commitments $\mathbf{C}$. The former follows the development of [18]; the latter generalizes the construction in [8] to allow an explicit social context. $\mathbf{I}(x,t)$ gives the paths that are "intentional-alternatives" for agent x at moment t. This is assumed to be defined only for individuals. $\mathbf{C}(x,y,\tau,t)$ gives the "commitment-alternative" paths that reflect the commitments of x for y in context τ.

We stipulate that the notions of exodeictic and endodeictic intentions coincide for individuals. An individual intends p ($x|p$) iff p occurs on all paths that are intentional-alternatives for him at the given moment. An agent is committed to another agent in a given context for p iff p holds on all commitment-alternative paths for the given creditor, context, and moment.

SEM-1. $M \models_t x|p$ iff $(\forall P : P \in \mathbf{I}(x,t) \Rightarrow M \models_{P,t} p)$

SEM-2. $M \models_t \mathsf{C}(x,y,\tau,p)$ iff $(\forall P : P \in \mathbf{C}(x,y,\tau,t) \Rightarrow M \models_{P,t} p)$

We assume that the semantics for sets of commitments, and sets of coordination relationships is given.

Definition 3 $[\![R]\!] \triangleq$ the set of paths on which the coordination relationships in R are satisfied.

Definition 4 $[\![S]\!] \triangleq$ the set of paths on which the commitments in S are satisfied.

4.2 Intentions Formalized

So that we can make logical claims uniformly at all levels of a nested team, we seek to preserve the basic nature of Sem-1. Another reason for preserving symmetry is that often the change from viewing a system as an individual to viewing it as a team is a necessary step in designing or analyzing the given system. We surely would not want all of our claims about the intentions of the system to be automatically invalidated when that happens, although (as we described in section 1) some adjustment may be necessary to accommodate the structure of a team.

Consequently, to check whether a team τ intends condition p, we carry out the same basic steps as for a single agent, adding commitment and coordination requirements, of course:

- Identify the set of admissible paths that satisfy
  - the intentions of the members of τ
  - the coordination requirements on the members imposed by their membership in τ
  - the commitment requirements on the members imposed by their membership in τ
- Check whether the given condition p holds on each of these paths
- If it does, τ intends p; otherwise, τ does not intend p.

The above schematic is applied to each of the senses of the intentions, although some of the tests are modified. Because our definitions consider the necessary consequences of admissible paths, they are fairly strong and eliminate contingently satisfied conditions. Moreover, the definitions depend on the intentions, coordination requirements, and commitments of the members. There is no implicature that a team will succeed with its intentions.

Below, we define sets of paths $\mathbf{I}_e$, $\mathbf{I}_n$, and $\mathbf{I}_j$ for teams. Each definition recurses on the structure of teams; for each, the base case of individuals is set to $\mathbf{I}$. These sets of paths capture the different semantics of intentions.

Exodeictic Intentions. A team τ exodeictically intends p iff p holds on all paths that satisfy the exodeictic intentions of the members of τ, and satisfy the team structure requirements.

Definition 5 $\mathbf{I}_e(\tau, t) \triangleq \bigcap_{x \in \tau^m} \mathbf{I}_e(x, t) \cap [\![S]\!] \cap [\![R]\!]$

SEM-3. $M \models_t x\mathsf{I_e}p$ iff $(\forall P : P \in \mathbf{I}_e(x, t) \Rightarrow M \models_{P,t} p)$

A team may exodeictically intend conditions that arise from the combinations of the members' intentions (and the team structure). For example, if two members intend $a \vee b$ and $a \vee \neg b$, then their team may intend a, which neither had intended.

Endodeictic Intentions. A team τ endodeictically intends p iff p holds on all paths that satisfy the endodeictic intentions of the members of τ, satisfy the team structure requirements, and require that the members are committed to τ to bring about p.

Definition 6 $\mathbf{I}_n(\tau, t) \triangleq (\bigcup_{x \in \tau^m} \mathbf{I}_n(x, t) \cup \bigcup_{x \in \tau^m} \mathbf{C}(x, \tau, \tau, t)) \cap [\![S]\!] \cap [\![R]\!]$

SEM-4. $M \models_t x\mathsf{I_n}p$ iff $(\forall P : P \in \mathbf{I}_n(x, t) \Rightarrow M \models_{P,t} p)$

It is useful to distinguish another notion of endodeictic intentions that we term *joint* intentions. Here we additionally require that the members be committed to each other to bring about p.

Definition 7 $\mathbf{I}_j(\tau, t) \triangleq (\bigcup_{x \in \tau^m} \mathbf{I}_j(x, t) \cup \bigcup_{x \in \tau^m} \mathbf{C}(x, \tau, \tau, t) \cup \bigcup_{x,y \in \tau^m} \mathbf{C}(x, y, \tau, t)) \cap [\![S]\!] \cap [\![R]\!]$

SEM-5. $M \models_t x\mathsf{I_j}p$ iff $(\forall P : P \in \mathbf{I}_j(x, t) \Rightarrow M \models_{P,t} p)$

Notice that the above definition for the endodeictic intentions unions in sets of paths reflecting the members' commitments. According to our semantic definitions, this means that p must hold in a larger set of paths in order to be considered intended. Thus, we have that $x\mathsf{I_j}p$ entails $x\mathsf{I_n}p$, which entails $x\mathsf{I_e}p$.

5 Some Postulates

One of the ways in which theoretical approaches may be evaluated is by determining which postulates the given definitions support. In light of this, the formal consequences of our definitions are important. We next state some postulates using a generic intention operator I, which is replaced by a specific operator $\mathsf{I_e}$, $\mathsf{I_n}$, or $\mathsf{I_j}$ to determine if they satisfy the given postulate.

P1. **Singleton Teams (ST).** *Are the intentions of a single-member team identical to those of its sole member?*
$\models x\mathsf{I}p \leftrightarrow \langle\{x\}, \emptyset, \emptyset\rangle\mathsf{I}p$
For example, let Bill intend to go to Italy. Then the team formed out of Bill alone and with no restrictions also intends to go to Italy. Only $\mathsf{I_e}$ supports this in general, but with the weak additional assumption that the commitments of a team-member to the team are intended by him, $\mathsf{I_n}$ and $\mathsf{I_j}$ also support it.

P2. **Temporal Coherence (TC).** *Does a team intend the temporal combination of the intentions of its members?*
$\models (x\mathsf{I}p \wedge y\mathsf{I}q \wedge (x, y \in \tau^m)) \rightarrow \tau\mathsf{I}(p \wedge q)$
If Al intends to hang a picture and Bo intends to hang a mirror, then does their team intend to hang both objects? This holds for $\mathsf{I_e}$, but is invalid for $\mathsf{I_n}$ and $\mathsf{I_j}$, since it is possible that no member is committed to hanging both objects. This postulate is related to the intuition that a team might not internally intend a consequence of its intentions.

P3. **Strong Upward Closure (SUC).** *Does a team intend what any one of its members intends?*
$\models (\exists x : x \in \tau^m \wedge x\mathsf{I}p) \rightarrow \tau\mathsf{I}p$
An example of this postulate is the following. If a platoon of an army brigade intends to control a major bridge, then the army brigade intends to control that bridge. Only $\mathsf{I_e}$ supports this.

P4. **Weak Upward Closure (WUC).** *Does a team intend what all of its members intend?*
$\models (\forall x : x \in \tau^m \rightarrow x\mathsf{I}p) \rightarrow \tau\mathsf{I}p$
For example, if all players in a soccer team intend to win, then the team intends to win. Only $\mathsf{I_e}$ supports this in general, but if commitments to the team are intended, then $\mathsf{I_n}$ (but still not $\mathsf{I_j}$) supports it.

P5. **Weak Nonemergence (WN).** *Is each intention of the team directly reflected as an intention of at least one of its members?*
$\models \tau\mathsf{I}p \rightarrow (\exists x : x \in \tau^m \wedge x\mathsf{I}p)$
If WN fails, then in effect, the intentions of a team are *strongly* emergent, in that no member of a team has the same intentions as the team itself. WN is not valid for any definition of intentions, because the team's intention may be due to a combination of its members' intentions. For example, neither Al nor Bo above may intend to hang both the picture and the mirror. WN fails in general, because of S and R, but if $S = \emptyset$ and $R = \emptyset$, and commitments to the team are intended, then it holds for $\mathsf{I_n}$ (but not for $\mathsf{I_e}$ and $\mathsf{I_j}$).

P6. **Strong Nonemergence (SN).** *Is each intention of the team directly reflected as an intention of each of its members?*
$\models \tau\mathsf{I}p \rightarrow (\forall x : x \in \tau^m \rightarrow x\mathsf{I}p)$
If SN fails, then in effect, the intentions of a team are *weakly* emergent, in that they are not common to all members. SN is

invalid for the same reasons as WN—the team's intention may be a combination of its member's intentions as in, e.g., P2. If $S = \emptyset$ and $R = \emptyset$, and commitments to the team are intended, then it holds for $\mathsf{I_n}$ (but not for $\mathsf{I_e}$ and $\mathsf{I_j}$).

P7. **Consequential Closure (CC).** *Are the intentions of a team closed under logical consequence?*
$\models \tau \mathsf{I} p \wedge \Box(p \rightarrow q) \rightarrow \tau \mathsf{I} q$
$\Box(p \rightarrow q)$ means that p entails q; i.e., p implies q at every moment in M. CC holds here as for all other possible worlds based approaches. We include this as a kind of truth-in-advertising, because its presence makes the formal notion different from its informal counterpart. Note, however, that only intentions for logical consequences are entailed, not for material consequences.

	ST	TC	SUC	WUC	WN	SN	CC
$\mathsf{I_e}$	+	+	+	+	−	−	+
$\mathsf{I_n}$	±	−	−	±	±	±	+
$\mathsf{I_j}$	±	−	−	−	−	−	+

Table 2. Postulates Supported

The above postulates involve several important theorems and non-theorems of our theory. They help distinguish between the different senses of intentions and help relate our models to the reasoning that they support. Table 2 summarizes our results for the three variants of team intentions that we formalized. + indicates support, − indicates failure, and ± indicates support when we make the assumption that commitments to one's team are intended. For WN and SN, we also assume $S = \emptyset$ and $R = \emptyset$.

6 Discussion

The theory presented in this paper refines and formalizes some intuitions about team intention, especially as that concept may be used in modeling the behavior of complex teams. It grounds the intentions of a team in (1) the intentions of its members, and (2) their structure. It allows nested teams, and provides a framework in which different useful senses of intentions can be formalized.

It is interesting that the above definitions avoided the concept of mutual beliefs altogether, leaving open the possibility of realizing a team-based system in a wider range of environments. Also, the above approach can naturally accommodate the structure of heterogeneous teams, including large teams whose members may not be perfectly cooperative. We believe the notions of endodeixis and exodeixis are applicable in a broader range of situations than treated here, for which they would need a more careful analysis than we developed.

There are a number of important directions for further work. One is the relationship with group and individual rationality. The connection between rationality and social concepts remains especially under-studied, although some conceptual and theoretical advances have been made [3, 5].

A related issue is about how agents may form goals of mutual interest [7], or how the members of a team may collectively reason about their intentions. Some nice theories of argumentation and negotiation are being developed, e.g., [16]. A closer investigation of these forms of negotiation and the creation and maintenance of teams and team intentions remains to be made.

Lastly, there is a large body of work on emergent behavior in multiagent systems, e.g., [10], which gives primacy to the behavior of agents over mental concepts. We showed above how team intentions naturally depend on the coordination relationships of the teams: this provides a potentially powerful means to unite the two camps.

ACKNOWLEDGEMENTS

This research was supported by the National Science Foundation under grant IRI-9624425 (Career Award), IBM, and the NCSU College of Engineering. This paper benefited greatly from the comments of the anonymous reviewers.

REFERENCES

[1] John L. Austin, *How to Do Things with Words*, Clarendon Press, Oxford, 1962.
[2] Kathleen Carley and Les Gasser, 'Computational organization research', in *[21]*, chapter 7, (1998).
[3] Amedeo Cesta, Maria Miceli, and Paola Rizzo, 'Help under risky conditions: Robustness of the social attitude and system performance', in *Proceedings of the International Conference on Multiagent Systems*, pp. 18–25, (1996).
[4] K. M. Chandy and Jayadev Misra, 'How processes learn', *Distributed Computing*, **1**, 40–52, (1986).
[5] Rosaria Conte and Cristiano Castelfranchi, *Cognitive and Social Action*, UCL Press, London, 1995.
[6] Daniel C. Dennett, *The Intentional Stance*, MIT Press, Cambridge, MA, 1987.
[7] Frank Dignum and Rosaria Conte, 'Intentional agents and goal formation', in *Intelligent Agents IV: Agent Theories, Architectures, and Languages (ATAL-97)*, pp. 231–244. Springer-Verlag, (1998).
[8] Barbara Dunin-Kęplicz and Rineke Verbrugge, 'Collective commitments', in *Proceedings of the International Conference on Multiagent Systems*, pp. 56–63, (1996).
[9] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi, *Reasoning About Knowledge*, MIT Press, Cambridge, MA, 1995.
[10] Jacques Ferber and Jean-Pierre Müller, 'Influences and reaction: A model of situated multiagent sytems', in *Proceedings of the International Conference on Multiagent Systems*, pp. 72–79, (1996).
[11] Barbara J. Grosz and Sarit Kraus, 'Collaborative plans for complex group action', *Artificial Intelligence*, **86**(2), 269–357, (October 1996).
[12] Nick R. Jennings, E. H. Mamdani, Jose Manuel Corera, Inaki Laresgoiti, Fabien Perriollat, Paul Skarek, and Laszlo Zsolt Varga, 'Using Archon to develop real-world DAI applications, part 1', *IEEE Expert*, **11**(6), 64–70, (December 1996).
[13] H. J. Levesque, P. R. Cohen, and J. T. Nunes, 'On acting together', in *Proceedings of the National Conference on Artificial Intelligence*, pp. 94–99, (1990).
[14] John McCarthy, 'Ascribing mental qualities to machines', in *Philosophical Perspectives in Artificial Intelligence*, ed., Martin Ringle, Harvester Press, (1979).
[15] Allen Newell, 'The knowledge level', *Artificial Intelligence*, **18**(1), 87–127, (1982).
[16] Carles Sierra, Nick R. Jennings, Pablo Noriega, and Simon Parsons, 'A framework for argumentation-based negotiation', in *Intelligent Agents IV: Agent Theories, Architectures, and Languages (ATAL-97)*, pp. 177–192. Springer-Verlag, (1998).
[17] Munindar P. Singh, 'A customizable coordination service for autonomous agents', in *Intelligent Agents IV: Agent Theories, Architectures, and Languages (ATAL-97)*, pp. 93–106. Springer-Verlag, (1998).
[18] Munindar P. Singh, Anand S. Rao, and Michael P. Georgeff, 'Formal methods in DAI: Logic-based representation and reasoning', in *[21]*, chapter 8, (1998). www.csc.ncsu.edu/ faculty/ mpsingh/ papers/ mas/ formal-DAI.ps.
[19] Milind Tambe, 'Agent architectures for flexible, practical teamwork', in *Proceedings of the National Conference on Artificial Intelligence*, pp. 22–28, (1997).
[20] Karl E. Weick, *The Social Psychology of Organizing*, Addison-Wesley, Reading, MA, 2nd edn., 1979.
[21] *Introduction to Distributed Artificial Intelligence*, ed., Gerhard Weiß, MIT Press, Cambridge, MA, 1998.