What I See is What You Say: Coordination in a Shared Environment with Behavioral Implicit Communication

Luca Tummolini[1], Cristiano Castelfranchi[1], Alessandro Ricci[2], Mirko Viroli[2] and Andrea Omicini[2]

[1] Institute of Cognitive Sciences and Technologies – CNR, Viale Marx, 15, 00137, Roma, Italy
{castel,tummoli}@ip.rm.cnr.it
[2] DEIS, Universita' degli Studi di Bologna a Cesena, Via Venezia 52, 47023, Cesena, Italy
{aomicini,aricci,mviroli}@deis.unibo.it

Abstract. Coordination between multiple autonomous agents is a major issue for open multi-agent systems. This paper proposes the notion of Behavioral Implicit Communication (BIC) originally devised in human and animal societies as a different and critical coordination mechanism also for artificial agents. BIC is a parasitical form of communication that exploits both some environmental properties and the agents' capacity to interpret each other's actions. In this paper we abstract from the agents' architecture to focus on the interaction mediated by the environment. To implement BIC in artificial societies two environmental properties are necessary: the "observability" of the software agents' actions and the "traceability" of the environment . The goal of this paper is to address the first property defining a model of observation mediated by the environment. From the viewpoint of the agents, a crucial distinction is proposed between the environment and the artifacts they can use. Both the environment and the artifacts can affect the visibility of agents in several respects and facilitate different forms of implicit communication between them. A typology of environments and examples of observation based coordination with and without implicit communication are described.

1 Introduction

In this paper we advance the notion of Behavioral Implicit Communication (BIC) as a kind of communication which does not involve specific codified actions aimed only to communicate [1]. We have BIC when usual practical actions are *contextually used* as messages for communicating. We argue that providing agents with an environment eases coordination achievement [2] *also* because it can enable a more flexible form of communication between agents.

BIC is a critical coordination mechanism which is mainly responsible for the overall social order of human societies. A sub-category of BIC, commonly known as stigmergy [3], is shared also with animal societies and is widely considered as a necessary means to achieve coordination without a central control. Stigmergy has been proposed also as a model of decentralized coordination for Multi-Agent Systems [4], and it is usually characterized as a form of communication mediated by the environment which simply needs ant-like agents. BIC is proposed as a general framework able to provide a more comprehensive theory that covers also intentional BDI agents. This paper focuses on the environmental properties that can enable BIC.

Approaches to coordination have been recently classified in two main categories: subjective and objective coordination [5][6]. Subjective approaches rely on the viewpoint of the individual agent that can "perceive" and understand the actions of its peers and can agree on a coordinated plan thanks to explicit communication [7] or plan recognition[8][9].

However what does it mean in this approaches that an agent can "perceive" or "observe" another agent? Do perception and observation always imply a form of communication between the two agents?

Differently, objective approaches are mainly concerned with the viewpoint of an external observer and provide the agents with an infrastructure to mediate their interactions [10][11]. Objective approaches are concerned with modeling the environment in which agents interact using coordination artifacts to influence the global behavior of the MAS.

In these approaches however environment and artifacts seems to collapse in a single category. On the contrary we will show that, for the sake of implementing BIC, agents societies should be provided both with an environment and a set of artifacts as different abstractions to interact with.

The remainder of the paper is structured as follows: in Section 2 we summarize what is Behavioral Implicit Communication and why it is relevant for coordination in a Multi-Agent System; in Section 3 we focus on the role of the environmental properties that can enable it, in particular the capacity of the environment to affect the observability of the agents, we advance a notion of *shared environment* and formalize a first typology; in Section 4 a description of how forms of observation-based coordination can be realized exploiting the observability features provided by shared environments, focusing in particular on the BIC approach. Finally, Section 5 concludes identifying a future implementation of the shared environment.

2 Behavioral Implicit Communication for coordination

2.1 Interaction is not always communication

There is a sense in which the famous claim of the Palo Alto psychotherapy school "any behavior is communication" [12] is true: in an artificial multi-agent system each

form of interaction with other agents or with the environment is implemented in form of a message passing protocol [13]. Direct interaction considered as interaction via messages has been criticized as the only viable solution to achieve coordination.
As a more powerful framework, indirect interaction has been proposed [14] as a way to implement *stigmergy* for MAS societies. Decentralized coordination would be achieved thanks to interaction via persistent observable state changes. Indirect interaction is modeled on the pheromone metaphor. To find the shortest way to reach food ants mark their trail with a pheromone that is attractive for other ants [15]. From a functional perspective, pheromone is like a message written on a blackboard. Everyone autonomously accessing the blackboard can read the message and act upon it. In this sense a blackboard is a powerful coordination artifact that extend the power of face to face personal interaction to asynchronous and impersonal forms (an analogous fundamental shift as the passage from oral communication to the written one).

While we will also argue for having persistence and observability of changes in the agents' environment as necessary requirements for having global coordination, we strive for a coordination mechanism which does not rely on explicit codified communication. In fact not all kinds of communication exploit codified (and hence rigid) actions. Our claim is that human and animals are able to communicate also without the need for a *predefined* conventional language and this capacity should be designed also for artificial agents.
To distinguish it from mere interaction, we define communication as a process where information arriving from Agent X (Sender) to Agent Y (Receiver) is *aimed to* informing Y. Agent X has the Goal of informing Agent Y. Agent X is executing a certain action "in order" to have other agents receiving a message and updating their beliefs. Communication is an intentional or functional notion in the sense that it is always goal oriented such that a behavior is selected also for its communicative effect.

When reasoning about agents we should be at the agents' level of explanation. There are at least two different viewpoints that need to be disentangled: the agent's and the designer's. Relative to the agents' world, the designer acts as Natural Selection or God does on our world. Even if an agent's perception of the action of another agent is necessary implemented as information transition from a sender to a receiver, this implementation of interaction should not be necessarily considered as "communication" and the passed information should not be always labeled as a "message".

From the external viewpoint of the designer a message passing of this sort is designed in order to inform the agent who is observing. However from the viewpoint of the agent a simple perception is not necessarily communication.

Consider a case where an hostile agent, whose actions are "observable" is entering a MAS. If another agent become aware of his presence, can observe him, should we say that the hostile agent is communicating his position? Or differently is the escaping prey communicating to the predator her movements?

While not being communication observable agents can be indeed very informative.

2.3 Communication is not always explicit

Communication is normally conceived as implemented through specialized actions such as those defined in the FIPA ACL protocol [16]. Such protocols are inspired by natural language or expressive signals where meaning is associated to a specific action by convention.

What about the case where the agent is aware of being observed (other agents believe that he is performing a given practical action) and he "intends that" [7] the other are interpreting his action? This sort of communication without a codified action but with a communicative intention is what we intend for Behavioral Implicit Communication [1]. What is relevant here is that the agent's execution plan is aimed to achieve a pragmatic goal as usual: i.e. an agent A is collecting trash to put it in a bin (as in [8]).

To implicitly communicate, the agent should be able to contextually "use" (or learn to use or evolve to use, see [17]) the *observed* executive plan also as a sign, the executive plan is used as a message but it is not shaped, selected, designed to be a message.

An agent B has the same goal but observing the other's action he decides to clean another side of the road. Since the agent A knows that an agent B is observing him, the practical action he is executing can be used *also* as a message to B such as "I am cleaning here". Such a possibility can lead agents to avoid a specific negotiation process for task allocation and can finally evolve in an implicit agreement in what to do.

There seems to be at least three different conditions to support such a form of communication.

- The first is relative to environmental properties. The "observability" of the practical actions and of their traces is a property of the environment where agents live, one environment can "enable" the visibility of the others while another can "constrain" it, like sunny or foggy days affect our perception. An environment could also enable an agent to make himself observable or on the contrary to hide his presence on purpose.
- The second is related to the capacity of agents to understand and interpret (or to learn an appropriate reaction to) a practical action. A usual practical action can be a message when an agent knows the way others will understand his behavior. The most basic message will be that the agent is doing the action a. More sophisticated form would imply the ability to derive pragmatic inference from it (what is the goal of doing? What can be implied?).
- The third condition is that the agent should be able to understand (and observe) the effect that his actions has on the others so that he can begin acting in the usual way *also* because the other understand it and react appropriately.

Behavioral Implicit Communication is in this sense a parasitical form of communication that exploits a given level of visibility and the capacity of the others to categorize or react to his behavior.

A general definition for BIC is:

the agent (source) is performing a usual practical action a, but he also knows and lets or makes the other agent (addressee) to observe and understand such a behavior, i.e. to capture some meaning m from that "message", because this is part of his (motivating or non motivating) goals in performing a.

2.3 BIC is not always Stigmergy

The need for an environment for a Multi Agent System is often associated with the goal of implementing stigmergy as decentralized coordination mechanism. Besides, being *the production of a certain behavior as a consequence of the effects produced in the local environment by previous behavior* or *indirect communication through the environment* [18], stigmergy seems very similar to the form of communication we are arguing for.

However these general accepted definitions makes the phenomenon too broad. It is too broad because it is unable to distinguish between the communication and the signification processes. As we have seen in 2.1 we do not want to consider the hostile agent's actions or the escaping prey as communicative actions notwithstanding that the effects of their actions elicit and influence the actions of other agents. Besides, every form of communication is mediated by the environment exploiting some environmental channel (i.e. air).

As in BIC, real stigmergic communication does not exploit any *specialized communicative* action but just usual practical actions (i.e. the nest building actions). In fact we consider stigmergy as a subcategory of BIC, being communication via long term *traces*, physical *practical* outcomes, *useful* environment modifications which preserve their practical end but acquire a communicative function. We restrict stigmergy to a special form of BIC where the addressee does not perceive the *behavior* (during its performance) but perceives other *post-hoc traces* and outcomes of it.

Usually stigmergy is advocated as a coordination mechanisms that can achieve very sophisticated forms of organization without the need for intelligent behavior. However there also exist interesting form of stigmergic communication at the intentional level. Consider a sergeant that – while crossing a mined ground –says to his soldiers: "walk on my prints!". From that very moment any print is a mere consequence of a step, plus a stigmergic (descriptive "here I put my foot" and precriptive "put your foot here!") message to the followers.

2.3 Coordination is not always cooperation

Coordination is that additional part or aspect of the activity of an Agent specifically devoted to deal and cope with the dynamic environmental interferences, either positive or negative, i.e. with opportunities and dangers/obstacles [19]. Coordination can either be non social as when an agent coordinate with a moving object. It can be *unilateral*, *bilateral* and *reciprocal* (see table 1) without being cooperative as when a leopard curves left and right and accelerates or decelerates on the basis of the *observed* path and moves of its escaping prey; but at the same time the gazelle jumps left or right and accelerates or not in order to avoid the leopard and on the basis of the *observed* moves of it. This is an observation based but not a communication/message based (BIC) *reciprocal* coordination.

We distinguish four different forms of coordination:

Table 1. Four kinds of observation-based coordination

	Unilateral	Bilateral	Reciprocal	Mutual
Ag X	X observes Y's behavior and is not observed	X observes Y's behavior and ignores to be observed	X observes Y and knows that Y is observing	X observes Y and knows that Y is observing and knows that Y wants to coordinate and wants Y to understand his intention to coordinate
Ag Y	Y ignores X's observation	Y observes X's and ignores to be observed behavior	Y observes X and knows that X is observing	Y observes X and knows that X is observing and knows that X wants to coordinate and wants X to understand his intention to coordinate

For the rest of the paper we will focus on the environmental properties that enable observation based coordination and we identify several examples of BIC based coordination.

3 Toward a Shared Environment: objective and intentional observability

Agents that live in a *common environment* (*c-env*) are agents whose actions and goals interfere (positively or negatively) and need coordination to manage this interference. In a pure *c-env*, actions and their traces are state transitions which can ease or hamper the individual agents' goals. An example is a ground which is common for different insects species but where no interspecies communication is possible. Agents can observe just the state of the environment and act on that basis without having access to the actions of their peers. Even a trace is seen as part of the environment and not as a product of other agents. A general property of a *c-env* is that it enables agents to modify its state and keep track of it.

We propose a notion of *shared environment* (*s-env*) which is a particular case of a *c-env* that enables (1) different forms of observability of each other action executions, as well as (2) awareness of such observability. These features will be shown to support (unilateral, bilateral, reciprocal, mutual) coordination.

3.2 Observability in shared environments

Each s-environment is defined by the level of observability that it can afford. The level of observability is the possibility for each agent to observe, *i.e.* to be informed about, another agent's actions or their traces.

The most general kind of s-envinronment can defined by the fact that each agent accessing it *can* observe all the others and is *observable* by them. A prototypical model of this sort of environment is the central 'square' of a town.

A level of observability is formalized by a set of relations $Power_{obs}$: A x A x Act, where A is the set of agents and Act is the set of usual practical actions. When $\langle x,y,\alpha\rangle \in Power_{obs}$, also written $Power_{obs}(x,y,\alpha)$, it means that action α executed by agent y is observable by agent x. In this case x has the role of observer agent and y that of observed agent. This means that in that *s-env*, it is possible for x to observe the actions of y.

$Power_{obs}$ relations can be then conceived as rules which define the set of 'opportunity and constraints' that afford and shape agents' observability within the environment. A specific rule is an opportunity or a constraint *for a specific agent* and in particular it so only relative to the agent's active goals while interacting with that environment.

A *public s-env* transfers to an agent a specific *observation power*: the power to be informed about others' actions. Given that A is the set of agents, $Power_{obs}\langle A, A, \alpha\rangle$ the observation relation that holds, $x \in A$ is the agent accessing the *s-env*, then $Power_{obs}\langle x, A, \alpha\rangle$ holds which means that x has the power (can) to access the practical actions α of all the agents in the environment.

When accessing a shared environment of this kind also the relation $Power_{obs}\langle A, x, \alpha\rangle$ holds which means that all the agents have the power to observe the action of x.
To take into account the agent's viewpoint over observation, we introduce the concept of agent *epistemic state (ES)*, representing the beliefs the agent has because of his observation role, namely, his occurrence in relations $Power_{obs}$ and his intention to observe other agents' actions.
The specific *beliefs* on the observation relations that hold in an *s-environment* constitute the *environmental knowledge* of the agent. The environmental knowledge of the accessing agent x in this kind of *s-env* is then given by:

1. $B_x Power_{obs}\langle x, A, \alpha\rangle$
2. $B_x Power_{obs}\langle A, x, \alpha\rangle$

The agent knows that it is possible for him to observe everybody and to be observed by everybody.

3.3 Observation is interaction with the environment via epistemic actions

Knowing that he can observe everybody and can be observed by everybody else, an agent can *check what* is doing a specific agent or can *check whether* somebody is observing him.

We define *epistemic actions* as any action aimed at acquiring knowledge from the environment [20]. In our framework epistemic actions are formalized as a class of *interactions* with the environment. While the environment specifies the opportunities and constraints by the set of $Power_{obs}$ rules, concrete observation acts, that exploit these powers, are a sort of epistemic actions.

In this paper we introduce only the most basic kind of epistemic action: the *check what/who* action which is an action aimed at acquiring a new belief and then is formalized as a request of information to the environment.

So, as the relation $Power_{obs}$ is introduced to statically describe the set of opportunities and constraints related to agents' observability, a relation *Obs* (a subset of $Power_{obs}$) has to be introduced to characterize the state of the s-env at a given time, so that obs(x,y,α) means that agent x is actually observing executions of action α by agent y. To model, agent's intention to observe other agents' actions, we introduce the concept of motivational state: besides the epistemic state, an agent is characterized by a *motivational state (MS)*, which represents agent's intentions in exploiting the observability power of the environment to observe other agents' actions. The motivational state is represented as a set of intentions of the kinds $I_X$check_what(act(y,α)), which means that the agent x intends to observe execution of actions α by agents y.

When such an intention appears in the MS of agent x, the s-env conceptually intercepts it and enacts the corresponding observations, that is, (i) the s-env adds $B_X$obs(y,α) to the agent's epistemic state (agent x knows that he is observing actions by agent y), and (ii) relation *Obs* is added the rule obs(x,y,α) (the s-env makes agent x observing actions α by agent y). In other words, we can think that the appearance of an intention in the motivation state of the agent causes the execution of an epistemic action toward the environment, enabling agent observations.

Similarly, an agent may want to stop observing actions. When the intention I_X-check_what(act(y,α)) disappears from the agent motivational state, its effects are reversed: $B_X$obs(y,α) is dropped from the agent's ES, and rule obs(x,y,α) is removed from the s-env. So, again dropping the intention from the MS causes the execution of an epistemic action on the environment, disabling agent observations.

Now we are ready to link the MS state of the agent, *Obs* rules and the ES state of the agent: according to the semantics of the actions, the occurrence of an event act(y,α,t) (meaning that the agent y executed actions a at time t) causes the creation of a new belief B_X act(y,α,t) in the epistemic state of all the agents x of the environment such that obs(x,y,α) holds.

4 Observation-based coordination

Hence different environments afford different levels/power of observability, which agents can exploit by issuing suitable epistemic actions to realize different forms of observation-based coordination on top of it, such as the BIC.

For instance, forms of unilateral coordination -- where x observes y's behavior and is not observed, y ignores x's observation – can be realized by defining in the rules of the s-env with a simple $Power_{obs}(x,y,\alpha)$ relationship, supposing that the coordination is based on the observation of actions α executed by Y. Then, by the intention I_X-check_what(act(y,α)), an agent X can coordinate with Y, without Y being aware of it. In particular, x can reason about y's actions dynamically observed, and then act accordingly.

Bilateral Coordination can be obtained by extending previous approach to include also the possibility of observing x's actions for agent y – including also $Power_{obs}(y,x,\beta)$ relation in the rules of the s-env – and by enabling y to observe x's actions, by means of a new intention $I_Y$check_what(act(x,β)), which causes a suitable epistemic action on the environment. In this way, both x and y are aware of the dynamic execution of each other actions, properly observed, and can adapt accordingly their behavior in order to reach the desired coordination.

4.1 Examples of BIC Coordination

Coordination is for the most part based on observation but in *mutual* coordination it is not simply observation that is necessary but BIC too. However tacit messages can be exchanged also in different other forms of coordination.

In coordination the most important message conveyed by BIC is not the fact that I intend to do (and keep my personal or social commitments – which is crucial in cooperation), or my reasons and motives for acting, or the fact that I'm able and skilled. It is more relevant communicating (informing) about when, how, where I'm doing my act/part in the shared environment, so that you can coordinate with my behavior while knowing time, location, shape, etc.

In what follows some examples of coordination with tacit messages are provided that are inspired mainly from the teamwork literature.

Information on the others members' activity: "I am ready". In [7] a trade off in the amount of information team members must maintain on each other intentions is discussed, particularly when a step involves only an individual or a subteam. This intention tracking does not need a complete plan recognition but simply that the individual or the subteam intend to execute that step. Consider as an example a sort of teamwork which is to drive an underground train. A coordination problem for the driver is to close the doors when all passengers are on board and this can be difficult when a station is overloaded. The driver is able to observe using a mirror the passengers rush in taking his train. Passengers usually don't know to be observed and they are not communicating their intentions. However usually before leaving a station the drivers make a first attempt to close the door which, although it is a practical action, is mainly used as a message like “The train is leaving”. The driver does not intend to really close the door. However either the passenger understand the message or simply infer the driver’s intention to leave they often go off the train and let the train leave safely the station. This is a case of bilateral coordination where only the drivers’ actions can be considered as messages.

Joint persistent goals achievement: “I have done it”. Joint intention theory [7,21,22] has been proposed as a framework for multi-agent coordination in a team. The team members are required to jointly commit to a joint persistent goal G. It also requires that when any team member acquires the belief that G has been achieved or turns out to be unachievable or irrelevant, a mutual belief about this event should be attained. Because of the domain is usually of partial observability, the team member is commonly designed to signal this fact to the other agent through *explicit communication*. However, in real world domains, explicit communication has a cost and sometimes the expected cost of miscoordination can outweigh it [23]. Behavioral implicit communication can be adopted in such cases even if it is possibly ambiguous because it can turn out to be good enough and better of not communicating at all. Drawing on [23] consider such scenario. Two helicopters with different abilities have a joint goal of reaching together a final destination but encounter a dangerous radar unit. Only one of them is capable of destroying the radar and should decide to communicate a message like “I destroyed the radar” to the other. However sending these message could be too expensive and risky (i.e. by being intercepted). If the destroyer believes that the other helicopter is following him and is observing him, by simply keeping on track to destination he can assume that the other will receive his silent message anyway and will keep the commitment to reach the final destination. This is a case of mutual coordination with tacit messages because also the follower’s action of keeping the track can be considered as a message.

5. Conclusion

In this paper we have proposed a model of a shared environment for observation based coordination which can enable behavioral implicit communication between the agents. The BIC approach and the related shared environment supporting framework can be suitably implemented in infrastructures supporting the MAS. In particular *governing infrastructures* – i.e. infrastructures providing abstractions and services also for governing / constraining agent interaction [6] – can be suitably adopted for the purpose, representing the s-env as a first class issue.

The requirement for a MAS infrastructure in order to support the observation-based coordination are:

- It must provide explicit abstractions storing, managing and enacting $Power_{obs}$ and *Obs*, as the set of rules defining the observability level of the environment and the set of rules defining actually what observations are taking place;

- It must have access to the motivational state of the agents, in order to dynamically update the *Obs* rules according to the new intention about observation found in MS;
- It must have access to the epistemic state of the agents, in order to dynamically update it according the action execution events and the *Obs* rules dynamically characterizing the shared environment;

The concept of observation artifact is strictly related to the *coordination artifact* abstraction [25], which represents first class runtime entities provided to agents to support their coordination. TuCSoN is a coordination infrastructure for MAS supporting the coordination artifact abstraction [10]: accordingly suitable infrastructure can be devised to support effectively observation artifacts, as runtime entities enhancing the observation capabilities of agents.

References

1. C. Castelfranchi. When doing is saying - the theory of behavioral implicit communication. Draft, Available at: http://www.istc.cnr.it/doc/62a 716p WhenDoingIsSaying.rtf, 2004
2. P. Ciancarini, A. Omicini, and F. Zambonelli. Multiagent system engineering: The coordination viewpoint. In N. R. Jennings and Y. Lesperance, editors, *Intelligent Agents VI. Agent Theories, Architectures, and Languages*, *volume 1757 of LNAI*, pages 250–259. Springer-Verlag, 2000.
3. G. Theraulaz, E. Bonabeau, "A Brief History of Stigmergy". *Artificial Life*, 5, 2, pp. 97-117, 1999.
4. R. Beckers, O. Holland and J.-L. Deneubourg. "From local actions to global tasks: stigmergy in collective robotics". In R. Brooks and P. Maes (Eds.) *Artificial Life 4*. MIT Press
5. M. Schumacher "Objective Coordination in Multi-Agent System Engineering–Design and Implementation",volume 2039 of LNAI .Springer-Verlag, April 2001.
6. A. Omicini, S. Ossowski. "Objective versus Subjective Coordination in the Engineering of Agent Systems" Intelligent Information Agents: The AgentLink Perspective. *LNAI 2586* (State-of-the-Art Survey). Springer-Verlag, 2003.
7. B.J Grosz, and S. Kraus, "Collaborative Plans for Complex Group Action" *Artificial Intelligence*, 86, (1996), 269-357.
8. A. S. Rao. "A unified view of plans as recipes". In G. Holmstrom-Hintikka and R. Tuomela, editors, *Contemporary Action Theory*. Kluwer Academic Publishers, The Netherlands, 1997.
9. M.J. Huber, E. H. Durfee "Deciding when to commit to action during observation based coordination". In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95)*, pages 163--170, Menlo Park, California, June 1995. AAAI Press.
10. A. Omicini, F. Zambonelli "Coordination for Internet Application Development" - Autonomous Agents and Multi-Agent Systems 2(3). Kluwer Academic Publishers, September 1999.
11. P. Noriega, C. Sierra "Electronic Institutions: Future Trends and Challenges. CIA 2002: 14-17.

12. P. Watzlavich, J. H. Beavin, D.D. Jackson, *Pragmatics of human communication: a study of interactional patterns, pathologies, and paradoxes*, W.W. Norton &Co., Inc., New York, 1967.
13. R. Milner. "Elements of Interaction" *Communications of the Association for Computing Machinery 36* (1), 78-89, 1993.
14. D. Keil, D. Goldin "Modeling Indirect Interaction in Open Computational Systems" *1st Int'l workshop on Theory and Practice of Open Computational systems (TAPOCS)* Linz, Austria June 2003, IEEE Computer Society Press
15. D. Helbing, J. Keltsch, P. Molnar "Modelling the evolution of human trail systems" *Nature*, 1997 ; 388(6637):47-50.
16. FIPA. FIPA communicative act library specification. http://www.fipa.org, 2000.
18. R.Beckers, O.E. Holland, J.L Deneuborg, "From Local to Global tasks Stigmergy and Collective Robotics", R. Beckers, O.E. Holland, J.L. Deneuborg, *Artificial Life IV*, Rodney A. Brooks and Pattie Maes (eds.), 1996, pag.181.
19. C. Castelfranchi, "Modelling Social Action for AI Agents", *Artificial Intelligence*, 103, pp. 157-182, 1998.
20. C. Castelfranchi, E. Lorini "Cognitive Anatomy and Functions of Expectations". Workshop on Cognitive Modeling of Agents and Multi-Agent Interactions. IJCAI Conference 2003.
21. P. Cohen, H. Levesque, "Teamwork", *Technical Report*, SRI-International, Menlo Park, Calif. USA. 1991.
22. N. R. Jennings, "Controlling cooperative problem solving in industrial multi-agent systems using joint intentions". *Artificial Intelligence* 75 (2) 195-240, 1995.
23. D. Pynadath, M. Tambe. "The communicative multiagent team decision problem: Analyzing teamwork theories and models". JAIR, 2002.
24. M. Tambe. "Towards flexible teamwork. Journal of Artficial Intelligence Research", 7:83--124, 1997.
25. Omicini, A. Ricci, A. Viroli, M., Castelfranchi C., Tummolini, L. Coordination Artifacts: Environment Based Coordination for Intelligent Agents. *International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'04)*, New York, USA