"Exhibitionists" and "Voyeurs" do it better: A Shared Environment for Flexible Coordination with Tacit Messages

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
{tummoli,castel}@ip.rm.cnr.it
[2] DEIS, Università degli Studi di Bologna
via Venezia 52, 47023 Cesena, Italy
{aricci,mviroli,aomicini}@deis.unibo.it

Abstract. Coordination between multiple autonomous agents is a major issue for open multi-agent systems. This paper proposes the notion of *Behavioural Implicit Communication* (BIC) originally devised in human and animal societies as a new 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 their actions. In this paper we abstract from the agents' architecture to focus on the interaction mediated by the environment. Observability of the environment – and in particular of agents' actions – is crucial for implementing BIC-based form of coordination in artificial societies. Accordingly in this paper we introduce an abstract model of environment providing services to enhance observation power of agents, enabling BIC and other form of observation-based coordination. Also, we describe a typology of environments and examples of observation based coordination with and without implicit communication.

1 Introduction

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

BIC is a critical coordination mechanism that 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 decentralised coordination for Multi-Agent Systems [4], and it is usually characterised 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 basic properties / services which can be used to instrument an agent (working) environment so as to support BIC and other observation-based forms of coordination. In other words we are looking for an abstract model of environment which could provide a systematic support to BIC, which could be exploited by MAS for their heterogeneous cooperative activities. In this paper a formal specification of such a model will be provided, so as to (i) making unambiguous the description of observation-form of coordination, (ii) making easier its engineering on top of MAS infrastructure, (iii) enabling form of automated reasoning on the dynamic behaviour of the MAS exploiting the services provided by the environment.

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. For instance, agents 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? On the other hand, objective approaches are mainly concerned with the viewpoint of an observer that is external to the agents. According to this interpretation, coordination is instilled in multi-agent systems by means of *ad hoc* abstractions, often termed as *coordination artifacts* [10], that mediate agent interactions. Coordination artifacts globally affect the behaviour of a multi-agent system, and are typically provided by agent coordination infrastructures [11, 12] that shape the environment where agents live and interact.

The remainder of the paper is structured as follows. In Section 2 we summarise what Behavioural Implicit Communication is, and why it is relevant for coordination in a multi-agent system. Section 3 focuses on the role of the environmental properties that can enable BIC, in particular the capacity of the environment to affect the observability of agents' actions: we advance a notion of *shared environment* and formalise a first typology. In Section 4 we provide a formal characterisation for multi-agent systems based on the notion of shared environment and BIC, and in Section 5 we show its usefulness in modelling properties of BIC and other scenarios. Section 6 describes how forms of observation-based coordination can be realised by exploiting the observability features provided by shared environments. Finally, Section 7 concludes trying to identify a path toward a future implementation of the shared environment.

2 Behavioural 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 behaviour is communication" [13] is true: in artificial multi-agent system, interaction with other agents or with the environment is usually implemented in terms of a message passing protocol, typically "wrapping" non-agent environmental resources to shape them as agents. Even the only widespread standards for agent technologies, provided by FIPA, currently account for speech acts only, neglecting in practice physical acts of any form [14].

However, interaction via messages is not the only viable solution to achieve coordination. As a more powerful framework, indirect interaction has been proposed [15] as a way to implement *stigmergy* for MAS societies. Decentralised coordination would be achieved thanks to interaction via persistent observable state changes. Indirect interaction is modelled 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 [16]. However from a functional perspective, even a pheromone is a message, like one written on a blackboard. Everyone autonomously accessing the blackboard can read the message and act upon it.

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 only 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 a *predefined* conventional language, and this capacity should be also instilled into artificial agents.

In order 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 at* informing Y. X's behaviour has the goal or the function of informing Y. X is executing a certain action "in order" to have other agents receiving a message and updating their beliefs or epistemic state. Communication is an intentional or functional notion in the sense that it is always goal oriented such that a behaviour is selected also for its communicative effect.[3]

While we agree with [17] that coordination can be seen as a causal process of correlation between agents' actions always involving an information flow between an agent and its environment, we do not consider always this flow as a process of communication. Consider a case where an hostile agent, whose actions are "observable", is entering a MAS. If another agent becomes aware of his presence and 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?

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:

[3] An agent's behaviour can be goal oriented for different reasons. An intentional agent (i.e. a BDI agent) is a goal governed agent (the goal is internally represented) which instantiates a communicative plan to reach the goal that another agent is informed about something. However also simple reactive agents (i.e. insect-like) can act purposively (hence can communicate) if their behaviour has been shaped by natural or artificial selection, by reinforcement learning or by design (in the interest of the agent itself). In these latter cases the behaviour has the *function* of communicating in the sense that it has been selected *because of* a certain communicative effect.

the agent's and the designer's [6]. Relative to the agents' world, the designer acts as Natural Selection or God does on our world. Even in the case that an agent's perception of the action of another agent is actually implemented as an information passed from a sender to a receiver, this should not be necessarily considered as a form of "communication", and correspondingly the information passed should not be necessarily labelled 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.

2.2 Communication is not always Explicit

Communication is normally conceived as implemented through specialised actions such as those defined in the FIPA ACL protocol [18]. 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 others are interpreting his action? This sort of communication without a codified action but with a communicative intention is what we intend for Behavioural 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) the *observed* executed plan also as a sign, the plan is used as a message but it is not shaped, selected, or designed to be a message.

An agent B has the same goal but observing the other's action he decides to clean the other 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, in the same way that 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 behaviour. The most basic message will be that the agent is

doing the action α. 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.

Behavioural 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 categorise or react to his behaviour.

A general definition for BIC is:

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

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 a decentralised coordination mechanism. Besides, being *the production of a certain behaviour as a consequence of the effects produced in the local environment by previous behaviour* or *indirect communication through the environment* [4], 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 Subsection 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 for audio signals).

As in BIC, real stigmergic communication does not exploit any *specialised 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. In this perspective, stigmergy to a special form of BIC where the addressee does not perceive the *behaviour* (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 organisation with no need for intelligent behaviour. However there also exist interesting form of stigmergic communication at the intentional level. Consider a sergeant who – while crossing a mined field – says "walk on my prints!" to his soldiers. From that very moment any print is a mere consequence of a step, plus a stigmergic (descriptive "here I put my foot" and prescriptive "put your foot here!") message to the followers.

2.4 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 be non social as when an agent coordinate with a moving object. For instance, it can be *unilateral*, *bilateral* and *reciprocal* (see Fig. 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.

For the goals of this paper, we distinguish five different forms of coordination:

Unilateral — X intends to coordinate with Y by observing Y's actions.

Bilateral — In this case we have the unilateral form of coordination for both agents, so: X intends to coordinate with Y by observing Y's actions, and viceversa: Y intends to coordinate with X by observing X's actions.

Unilateral-AW — In this case we have a unilateral form of coordination, but with a first form of awareness: X intends to coordinate with Y by observing Y's actions, and Y is aware of it (i.e. knows to be observed).

Reciprocal — In this case the we have both a bilateral form of observation based coordination and awareness by both the agents: X intends to coordinate with Y by observing Y's actions, Y is aware of it, Y intends to coordinate with X by observing X's actions and X is aware of it.

Mutual — This case extends the reciprocal form by introducing the explicit awareness of each other intention to coordinate: X intends to coordinate with Y by observing Y's actions, Y is aware of it, Y intends to coordinate with X by observing X's actions, X is aware of it, X is aware of Y's intention to coordinate, and Y is aware of X's intention to coordinate.

Behavioural implicit communication is necessary for mutual coordination while it is possible and useful in the other kinds of observation-based coordination.

3 Toward a Shared Environment: Objective & Intentional Observability

The goal of this paper is to instrument an agent working environment with properties and services that can enable the observation-based forms of coordination discussed above, BIC in particular. In other words, we aim at defining an abstract model of an agent environment, which could be engineered on top of a

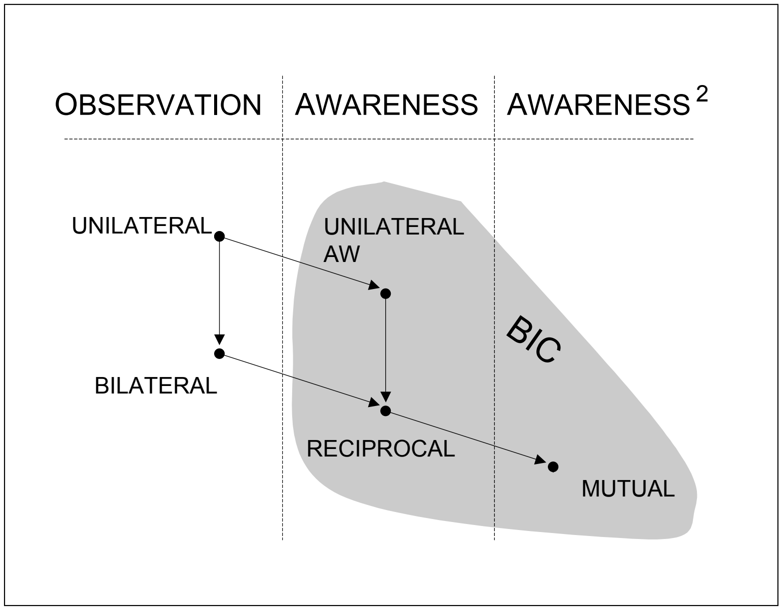

Fig. 1. Forms of coordination in relation to observation capability and awareness.

MAS infrastructure, so as to be exploited by agents living in the MAS to exploit BIC for their cooperation. Given this objective, it will be fundamental to identify a formal model of this support, in order to ease its engineering on top of existing MAS infrastructures: accordingly in this section the abstract model of the shared environment enabling observability is introduced and in Section 4 its formal semantics is described.

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 that can ease or hamper the individual agents' goals. An example is a ground that 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*), that 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 this observability. These features will be shown to support (unilateral, bilateral, reciprocal, mutual) coordination.

3.1 Observability in Shared Environments

Each s-env 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-env* 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.

The level of observability of an s-env is formalised by a relation $Pow : A \times A \times Act$, where A is the set of agents and Act is the set of usual practical actions. When $\langle x, y, \alpha \rangle \in Pow$, also written $Pow(x, y, \alpha)$, it means that action $\alpha \in Act$ 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. More generally, $Pow(x, B, \alpha)$ holds for the set $B \subseteq A$ of agents which x has the power to observe through action α, and similarly, $Pow(B, y, \alpha)$ holds for the set $B \subseteq A$ of agents that have the power to observe executions of α by agent y.

Pow relation can be then conceived as rules that 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 is 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. So, as the relation *Pow* is introduced to statically describe the set of opportunities and constraints related to agents' observability, a relation *Obs* (a subset of *Pow*) has to be introduced to characterise the state of the s-env at a given time, so that $Obs(x, y, \alpha)$ means that agent x is actually observing executions of action α by agent y. That is, $Obs(x, y, \alpha)$ means that an execution of action α by agent y will be perceived by 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. The ES of an agent x includes its *environmental knowledge* which is then given by information *(i)* on the agents he is observing, *(ii)* on the agents that are observing him, and *(iii)* on the action executions that he is observing. We generalise, and write $B_z obs(x, y, \alpha)$ for agent z believing that x is observing executions of action α by z, and $B_x done(y, \alpha)$ for x believing that y has executed action α.

3.2 Observation is Interaction with the Environment via Epistemic Actions

The epistemic state of an agent evolves through *epistemic actions*, which are actions aimed at acquiring knowledge from the environment [20]. In our framework epistemic actions are formalised as a class of *interactions* with the environment. Typically, an epistemic action is fired by an agent intention, by which the s-env reacts updating the epistemic state of the agent. To model agent's intention, we introduce the concept of motivational state: besides the epistemic state, an agent is characterised by a *motivational state* (MS).

A first case of epistemic action is used by the agent which is willing to know whether he is observing another agent, whether another agent is observing him, or generally, whether an agent x is observing an agent y. So, suppose the MS of z includes intention $I_z check(x, y, \alpha)$, which means that agent z intends to know whether x observes executions of α by y. At a given time, an epistemic action is executed by which the ES of agent z will include the belief about whether $Obs(x, y, \alpha)$ holds or not.

Similarly, an agent may have the intention $I_x obs(x, y, \alpha)$ in exploiting the observability power of the environment to observe y's actions α. The intention activates the observation service provided by the s-env, causing: *(i)* the $B_x obs(x, y, \alpha)$ knowledge to be added to agent's epistemic state (i.e. agent x knows that he is observing actions by agent y); (ii) the element $Obs(x, y, \alpha)$ to be added to the set defining *Obs* relation (meaning that the s-env enables the observation for agent x of actions α executed 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 $I_x drop(x, y, \alpha)$ appears in the agent motivational state, the effects of $obs(x, y, \alpha)$ are reversed, i.e. no longer the agent continues to observe action α in the future.

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 execution of an action α by agent y (denoted as $done(y, \alpha)$) causes the creation of a new belief $B_x done(y, \alpha)$ in the epistemic state of all the agents x of the environment such that $Obs(x, y, \alpha)$ holds.

4 Formal Model

In the following, we provide a syntax and an operational semantics for modelling MAS according to the conceptual framework defined in previous sections. This formalisation has the primary goal of been a precise description of the concepts described in previous section, and of their impact on the dynamic evolution of a MAS. Then, given the operational character of the model, it can be used as an abstract reference implementation for an infrastructure supporting s-envs, as well as to pave the way towards the application of some analysis tool.

We let metavariables x, y, z range over agent identifiers, and α, β over usual practical actions.

The syntax of MAS configurations is as follows:

$$
\begin{array}{lll}
S ::= 0 \mid A \mid E \mid S \,||\, S & & \text{MAS Configuration} \\
A ::= 0 & & \text{Agent Configuration} \\
\quad \mid \; B_x\phi & & \text{Belief of } x \\
\quad \mid \; I_x\phi & & \text{Intention of } x \\
\quad \mid \; A \,||\, A & & \text{Composition} \\
E ::= 0 & & \text{Environment Configuration} \\
\quad \mid \; Pow(x, y, \alpha) & & x \text{ has the power to observe } y\text{'s } \alpha \\
\quad \mid \; Obs(x, y, \alpha) & & x \text{ is observing } y\text{'s } \alpha \\
\quad \mid \; E \,||\, E & & \text{Composition} \\
\phi ::= & & \text{Formulas} \\
\qquad obs(x, y, \alpha) & & x \text{ is observing } y\text{'s } \alpha \\
\quad \mid \; coord(x, y, \alpha) & & x \text{ coordinates with } y \text{ through } \alpha \\
\quad \mid \; check(x, y, \alpha) & & \text{check whether } x \text{ is observing } y\text{'s } \alpha \\
\quad \mid \; drop(x, y, \alpha) & & \text{prevent } x \text{ from observing } y\text{'s } \alpha \\
\quad \mid \; done(x, \alpha) & & x \text{ executes actions } \alpha \\
\quad \mid \; \neg\phi \mid I_x\phi \mid B_x\phi & & \text{Structured formulas}
\end{array}
$$

The operator for parallel composition is assumed to be commutative, associative, and to absorb the empty configuration 0.

The metavariable S ranges over configurations of the MAS, which at our abstraction level are a parallel composition of agent configurations and environment configurations. Environment configurations are parallel composition of terms, each denoting either the power of agent x to observe action α executed by agent y ($Pow(x, y, \alpha)$), or the fact that the environment is currently supporting the fact that x is observing action α executed by agent y ($Obs(x, y, \alpha)$). Agent configurations are parallel compositions of mental properties, namely beliefs (B) and intentions (I) qualified by the agent x, and about a formula ϕ. Notice that we model a MAS configuration as a multiset of either agent and environment properties, without a separation, by simply following the abstraction process induced by the formalism adopted.

A formula ϕ can be believed and/or intended by an agent. Atomic formulas are: *(i)* $obs(x, y, \alpha)$, used to express that x is observing executions of α by y, *(ii)* $coord(x, y, \alpha)$, used to express that x coordinates its behaviour with y by observing executions of α, *(iii)* $check(x, y, \alpha)$, used to check if x is observing executions of α by y, *(iv)* $drop(x, y, \alpha)$, used to prevent x from observing executions of α by y, and *(v)* $done(x, \alpha)$, used to express an that x executes/has executed α. Moreover, formulas can be structured ones: $\neg\phi$ expresses negation of ϕ, $I_x\phi$ and $B_x\phi$ that agent x intends/believe ϕ, respectively. A number of assumptions on such formulas are clearly to be made as usual, e.g. that $\neg\neg\phi \equiv \phi$ or $B_x\phi \equiv B_xB_x\phi$, but we abstract away from this aspect for it plays no significant role in this paper.

The operational semantics is defined by the following rewrite rules for system configurations.

$$\frac{-}{I_z check(x,y,\alpha) \,||\, Obs(x,y,\alpha) \,||\, S \rightarrow B_z obs(x,y,\alpha) \,||\, Obs(x,y,\alpha) \,||\, S} \quad \text{[CHK]}$$

$$\frac{Obs(x,y,\alpha) \notin S}{I_z check(x,y,\alpha) \,||\, S \rightarrow B_z obs(x,y,\alpha) \,||\, S} \quad \text{[NCHK]}$$

$$\frac{-}{I_z drop(x,y,\alpha) \,||\, B_z obs(x,y,\alpha) \,||\, Obs(x,y,\alpha) \,||\, S \rightarrow B_z \neg Obs(x,y,\alpha) \,||\, S} \quad \text{[YDRP]}$$

$$\frac{Obs(x,y,\alpha) \notin S}{I_z drop(x,y,\alpha) \,||\, B_z obs(x,y,\alpha) \,||\, S \rightarrow B_z \neg obs(x,y,\alpha) \,||\, S} \quad \text{[NDRP]}$$

$$\frac{-}{I_z obs(x,y,\alpha) \,||\, Pow(x,y,\alpha) \,||\, S \rightarrow B_z obs(x,y,\alpha) \,||\, Pow(x,y,\alpha) \,||\, Obs(x,y,\alpha) \,||\, S} \quad \text{[ASK]}$$

$$\frac{I_x done(x,\alpha) \,||\, S \rightarrow I_x done(x,\alpha) \,||\, S'}{I_x done(x,\alpha) \,||\, Obs(y,x,\alpha) \,||\, S \rightarrow I_x done(x,\alpha) \,||\, Obs(y,x,\alpha) \,||\, B_y done(x,\alpha) \,||\, S'} \quad \text{[OBS1]}$$

$$\frac{Obs(y,x,\alpha) \notin S}{I_x done(x,\alpha) \,||\, S \rightarrow B_x done(x,\alpha) \,||\, S} \quad \text{[OBS2]}$$

$$\frac{-}{A \,||\, S \rightarrow A' \,||\, S} \quad \text{[AG]}$$

Rule [CHK] says that if agent z intends to check/know if x is observing y's action α and this is the case, then such an intention will be turned into a belief. Dually, rule [NCHK] deals with the case where this is not the case ($Obs(x, y, \alpha)$ does not occur in the system configuration), so that z will believe that $obs(x, y, \alpha)$ does not hold.

Rule [YDRP] says that if agent z know that x is observing y's action α (which is the case) and wants to stop him, term $Obs(x, y, \alpha)$ is dropped from the environment and z's belief is updated correspondingly. By rule [NDRP] we deal where the similar case, but supposing the agent had a wrong belief (x was not actually observing y's actions α), which is dealt with trivially.

Rule [ASK] is about agent z willing that x observes y's actions α: if this is allowed ($Pow(x, y, \alpha)$), x's beliefs will be updated as well as the environment state.

Rule [OBS1] and [OBS2] recursively define how the environment broadcasts information about an action to all the observers. When agent x wants to execute α, each observer y (rule [OBS1]) will be recursively added the belief $B_y done(x, \alpha)$: when none needs to be managed, x intention can simply become a fact, that is, he will belief the action to be executed.

The final, trivial rule [AG] is used to represent the fact that at any given time some agent configuration can change autonomously, thus modelling any belief revision or intention scheduling.

Notice that formulas $B_z coord(x, y, \alpha)$ or $I_z coord(x, y, \alpha)$ never appear in this semantics. This is because the fact that an agent coordinates its behaviour with another is not an aspect influencing/influenced by the environment: it is rather a mental property characterising the forms of observation-based coordination an agent participates to thanks to the s-env support.

5 Applications of the model

The formal model described above serves multiple purposes:

- clearly and rigorously identifying basic primitives / general-purpose mechanisms which can be composed to specify various type of observation-based coordination patterns;
- helping the engineering of the approach on top of MAS infrastructures. The operational semantics provides a rigorous description of the observation features we aim at supporting at the infrastructure level. So it is a fundamental guide for designer and developers of MAS infrastructures which want to support BIC;
- supporting agent reasoning. By formally defining the observability rules characterising environment configuration, we promote their inspection and formal reasoning by intelligent agents, so as to automate the analysis of the dynamic behaviour of the MAS: recognising failures, providing suggestions, and so on.

A formal semantics makes it possible to establish some rigorous properties about observed events and observation rules of the environment, which necessarily impact on the reasoning process of observing / observed agents. In other words, the environment (infrastructure) provides some guarantees which can be taken as assumptions by agents exploiting the services, so as to alleviate their reasoning process: for instance, the environment can guarantee agents to observe *all* the actions executed by a certain other agent in the right order.

In this section we deepen this issue showing how concepts and applications related to the s-env notion can be formally addressed by our model.

5.1 Specifying Observation-based Coordination

The formal framework can be adopted to specify rigorously the forms of coordination devised in Section 2. Given two agents x and y, an action α, and the system configuration S we introduce the following predicates:

- Unilateral

$$Uni(x, y, \alpha, S) \triangleq I_x coord(x, y, \alpha) \in S \ \wedge \ Obs(x, y, \alpha) \in S$$

– Unilateral with Awareness

$$UniAW(x, y, \alpha, S) \triangleq Uni(x, y, \alpha, S) \ \wedge \ B_y obs(x, y, \alpha) \in S$$

– Bilateral

$$Bi(x, y, \alpha, S) \triangleq Uni(x, y, \alpha, S) \ \wedge \ Uni(y, x, \alpha, S)$$

– Reciprocal

$$Rec(x, y, \alpha, S) \triangleq UniAW(x, y, \alpha, S) \ \wedge \ UniAW(y, x, \alpha, S)$$

– Mutual

$$Mut(x, y, \alpha, S) \triangleq Rec(x, y, \alpha, S) \ \wedge \ B_x I_y coord(y, x, \alpha) \ \wedge \ B_y I_x coord(x, y, \alpha)$$

So forms of unilateral coordination are obtained by instrumenting the environment configuration with the simple rule $Pow(x, y, \alpha)$ and with agent x manifesting the intention $I_x(obs(x, y, \alpha))$, causing the instrumentation of the environment with the rule $Obs(x, y, \alpha)$.

Bilateral coordination can be obtained by extending previous approach to include also y observation of x's actions, instrumenting the environment with the rules $Pow(y, x, \alpha)$ and $Obs(y, x, \alpha)$, the latter instantiated by the intention of the agent y $I_y obs(y, x, \alpha)$.

The unilateral and bilateral forms of coordination can be extended then with forms of awareness, by agents intention $I_y check(x, y, \alpha)$ enabling y awareness of the observability of his actions to x – obtaining the unilateral-aw form – and $I_x check(y, x, \alpha)$, enabling also x awareness of the observability of his actions to y – obtaining the reciprocal form of coordination.

5.2 From Overhearing to Oversensing

As an example scenario possibly enjoying the features of s-envs we consider overhearing. This has been introduced in MAS as a technique / architecture to realise forms of collaboration and coordination non-preplanned, typically in unstructured and unpredictable environments, based on unobstrusive observation and unsolicited suggestion [21]. Roughly speaking, overhearing consists in one agent – the overhearer – sniffing messages exchanged by two or more agents. The overhearer collects the messages and makes them available to suggester agents through a sort of publish / subscribe service: suggesters subscribe their interest to be notified by the overhearer when a certain type of event occurs concerning the communication among observed agents.

Overhearing has been used for supporting group formation in open environments [22], monitoring the interactive behaviour of organization – in particular implicit organisation [23], to enable awareness among agents [24], plan and conversation dynamic recognition [25].

Overhearing can be suitably implemented on top of the abstract model of environment described in this paper. In our case the overhearer agent disappears, since its functionalities are directly provided by the s-env environment (infrastructure), which is responsible to enable interactions (communications) among agents. The publish / subscribe service among the overhearer agent and suggester agents is mapped onto the s-env services.

Each suggester agent manifests its intention to overhear certain communication events concerning the interaction of agents $a_1, \ldots a_n$ by formulating the intentions:

$$I_s obs(s, a_1, \textit{communication-event})$$
$$\ldots$$
$$I_s obs(s, a_n, \textit{communication-event})$$

Specific examples can be:

$$I_s obs(s, customer, ask(customer, service\_provider, price(X, Y)))$$
$$I_s obs(s, service\_provider, inform(service\_provider, \_, \_))$$

Here, all the requests made by the customer agent to a service provider about some good prices are observed, along with all the information provided by the service-provider. Actually, our model supports an extension of overhearing toward oversensing, i.e. applying the principle of overhearing to a general model of action / interaction, which includes – but is not limited to – communication.

6 Other Examples of BIC Coordination

Mutual coordination is at the basis of BIC, requiring not only observation based coordination and forms of awareness, but agents awareness of each other intention to coordinate.

Actually, 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 behaviour 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.

6.1 Information on the other 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 sub-team. This intention tracking does not need a complete plan recognition but simply that the individual or the sub-team 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 whether passengers 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.

6.2 Joint persistent goals achievement: "I have done it".

Joint intention theory [7, 26, 27] 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 mis-coordination can outweigh it [28]. Behavioural 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 [28] 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.

7 Conclusion

In this paper we have proposed a model of a shared environment for observation based coordination which can enable behavioural 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 *pow* and *obs* rules, as the set of rules defining respectively 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 check for agent intentions, causing epistemic actions and then the updating of the *obs* rules of the environment;
– 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 characterising the shared environment.

The concept of observation artifact is strictly related to the *coordination artifact* abstraction [10], 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 [11]: accordingly suitable infrastructure can be devised to support effectively observation artifacts, as runtime entities enhancing the observation capabilities of agents.

Acknowledgements

The authors are grateful to the anonymous reviewers, whose work, constructive remarks and criticisms have greatly helped improving the quality of this paper.

This research is under the auspices of: MIUR (the Italian Ministry of Education, University and Research), COFIN 2003 project "Fiducia e diritto nella società dell'informazione", paper no. 3; and EC (the European Community), FP6 project "AgentLink III".

References

1. Castelfranchi, C.: When doing is saying – the theory of behavioral implicit communication. Draft. Available at `http://www.istc.cnr.it/doc/62a_716p_WhenDoingIsSaying.rtf` (2003)
2. Ciancarini, P., Omicini, A., Zambonelli, F.: Multiagent system engineering: the coordination viewpoint. In Jennings, N.R., Lespérance, Y., eds.: Intelligent Agents VI — Agent Theories, Architectures, and Languages. Volume 1767 of LNAI., Springer-Verlag (2000) 250–259
3. Theraulaz, G., Bonabeau, E.: A brief history of stigmergy. Artificial Life **5:2** (1999) 97–117
4. Beckers, R., Holland, O., Deneubourg, J.L.: From local actions to global tasks: Stigmergy in collective robotics. In Brooks, R.A., Maes, P., eds.: Artificial Life IV. MIT Press (1994)
5. Schumacher, M.: Objective Coordination in Multi-Agent System Engineering – Design and Implementation. Volume 2039 of LNAI. Springer-Verlag (2001)

6. Omicini, A., Ossowski, S.: Objective versus subjective coordination in the engineering of agent systems. In Klusch, M., Bergamaschi, S., Edwards, P., Petta, P., eds.: Intelligent Information Agents: An AgentLink Perspective. Volume 2586 of LNAI: State-of-the-Art Survey. Springer-Verlag (2003) 179–202
7. Grosz, B., Kraus, S.: Collaborative plans for complex group action. Artificial Intelligence **86** (1996) 269–357
8. Rao, A.S.: A unified view of plans as recipes. In Hölmstrom-Hintikka, G., Tuomela, R., eds.: Contemporary Action Theory. Volume 2: Social Action. Kluwer Academic Publishers (1997)
9. Huber, M., Durfee, E.: Deciding when to commit to action during observation based coordination. In: 1st International Conference on Multi-Agent Systems (ICMAS-95), Menlo Park, CA, USA, AAAI Press (1995) 163–170 Proceedings.
10. Omicini, A., Ricci, A., Viroli, M., Castelfranchi, C., Tummolini, L.: Coordination artifacts: Environment based coordination for intelligent agents. In: Autonomous Agents and Multi-Agent Systems. (2004) 3rd International Joint Conference (AAMAS 2004), New York, July 2004.
11. Omicini, A., Zambonelli, F.: Coordination for Internet application development. Autonomous Agents and Multi-Agent Systems **2:3** (1999) 251–269 Special Issue: Coordination Mechanisms for Web Agents.
12. Noriega, P., Sierra, C.: Electronic institutions: Future trends and challenges. In Klusch, M., Ossowski, S., Shehory, O., eds.: Cooperative Information Agents VI. LNAI. Springer-Verlag (2002) 14–17
13. Watzlavich, P., Beavin, J., Jackson, D.: Pragmatics of human communication: a study of interactional patterns, pathologies, and paradoxes. W.W. Norton &Co., Inc., New York (1967)
14. Omicini, A., Ricci, A., Viroli, M., Rimassa, G.: Integrating objective & subjective coordination in multiagent systems. In: 19th ACM Symposium on Applied Computing (SAC 2004), Nicosia, Cyprus, ACM (2004) 449–455 Special Track on Coordination Models, Languages and Applications.
15. Keil, D., Goldin, D.: Modeling indirect interaction in open computational systems. In: IEEE 12th International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WET ICE 2003), 1st International Workshop "Theory and Practice of Open Computational Systems" (TAPOCS 2003), Linz, Austria, IEEE CS (2003) Proceedings.
16. Helbing, D., Keltsch, J., Molnar, P.: Modelling the evolution of human trail systems. Nature **388()** (1997) 47–50
17. Parunak, H.V.D., Breuckner, S., Sauter, J., Odell, J.: A preliminary taxonomy of multi-agent interaction. In Giorgini, P., Müller, J., Odell, J., eds.: Agent-Oriented Software Engineering IV. LNCS. Springer-Verlag (2004) 4th International Workshop (AOSE 2003), Post-Proceedings.
18. FIPA: FIPA Communicative Act Library Specification. (2000) `http://www.fipa.org`.
19. Castelfranchi, C.: Modelling social action for AI agents. Artificial Intelligence **103** (1998) 157–182
20. Castelfranchi, C., Lorini, E.: Cognitive anatomy and functions of expectations. In: Cognitive Modeling of Agents and Multi-Agent Interactions. (2003) Workshop at IJCAI. Proceedings.
21. Aiello, M., Busetta, P., Dona', A., Serafini, L.: Ontological overhearing. In: Intelligent Agents VIII. Volume 2333 of LNAI. Springer-Verlag Berlin (2002) 175–189

22. Legras, F., Tessier, C.: Lotto: Group formation by overhearing in large teams. In Jennings, N.R., Sierra, C., Sonenberg, L., Tambe, M., eds.: 3rd international Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004), New York, USA, ACM (2004)
23. Rossi, S., Busetta, P.: Towards monitoring of group interactions and social roles via overhearing. In: CIA 2004. Volume 3191 of LNAI. Springer-Verlag Berlin (2004) 47–61
24. Novik, D., Ward, K.: Mutual beliefs of multiple conversants: a computational model of collaboration in air traffic control. In: AAAI-93, Washington, DC, USA (1993) 196–201
25. Gutnik, G., Kaminka, G.: Toward a formal approach to overhearing: algorithm for conversation identification. In Jennings, N.R., Sierra, C., Sonenberg, L., Tambe, M., eds.: 3rd international Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004), New York, USA, ACM (2004)
26. Cohen, P., Levesque, H.: Teamwork. Technical report, SRI-International, Menlo Park, CA, USA (1991)
27. Jenning, N.R.: Controlling cooperative problem solving in industrial multi-agent systems using joint intentions. Artificial Intelligence **75:2** (1995) 195–240
28. Pynadath, D., Tambe, M.: The communicative multiagent team decision problem: Analyzing teamwork theories and models. Journal of Artificial Intelligence Research (2002)