F. Dignum. Social interactions of autonomous agents; private and global views on communication. In J.-J. Meyer and P-Y. Schobbens, editors, F*ormal models of agents (LNCS-1760)*, Springer-Verlag, 1999, pages 103--122.

Social Interactions of Autonomous Agents; Private and Global Views on Communication

F. Dignum

Fac. of Maths. & Comp. Sc., Eindhoven University of Technology
P.O. Box 513, 5600 MB Eindhoven, The Netherlands
e-mail: dignum@win.tue.nl
phone: +31-402473705, fax: +31-402463992

Abstract. In describing the interactions between agents we can take either a global view, where the set of all agents is seen as one big system, or a private view, where the system is identified with a single agent and the other agents form a part of the environment. Often a global view is taken to fix some protocols (like contract net) for all the possible social interactions between agents within the system. Privately the agents then have fixed *reaction* rules to respond to changes in the environment. In a sense the agents are no longer autonomous in that they always respond in a fixed way and their behaviour can be completely determined by other agents. In this paper we investigate the case where there might not be a (or one) fixed protocol for the social interaction and where the agents do not necessarily react in the same way to each message from other agents. We distinguish between the agents perception of the world and the "real" state of the world and show how these views can be related.

Keywords: Multi-Agent Systems, Multi-Modal logic, Communication, Speech acts.

1 Introduction

In the area of Multi-Agent Systems much research is devoted to the coordination of the agents. Many papers have been written about protocols (like contract net) that allow agents to negotiate and cooperate (e.g. [19, 4]). Most of the cooperation between agents is based on the assumption that they have some joint goal or intention. Such a joint goal enforces some type of cooperative behaviour on all agents (see e.g. [3, 13, 23]). The conventions according to which the agents coordinate their behaviour is hard-wired into the protocols that the agents use to react to the behaviour (cq. messages) of other agents.

This raises several issues. The first issue is that, although agents are said to be autonomous, they always react in a predictable way to each message. Namely their response will follow the protocol that was built-in. The question then arises how autonomous these agents actually are. It seems that they react always in standard ways to some stimulus from other agents, that can therefore determine their behaviour.

Besides autonomy, an important characteristic of agents is that they can react to a changing environment. However, if the protocols that they use to react to (at least some part of) the environment are fixed, they have no ways to respond to changes. For instance, if an agent notices that another agent is cheating it cannot switch to another protocol to protect itself. (At least this is not very common). In general it is difficult (if not impossible) for agents to react to violations of the conventions by other agents.

As was also argued in [21], autonomous agents need a richer communication protocol than contract net (or similar protocols) to be able to retain their autonomy. A greater autonomy of the agent places a higher burden on the communication. An autonomous agent might negotiate over every request it gets. In this paper we will describe a mechanism to avoid excessive communication. It is similar to the one employed in [21], but defined more formally and still more generally applicable.

Negotiation between autonomous agents is only necessary if the agents do not have complete knowledge of the state of the world. If they did have complete knowledge (including knowledge about the state of minds of the other agents) they could calculate the optimum deal for both agents and agree in one step. This fact makes it important to distinguish between private and global views of the state of the world. And even more important the private and global view of actions and communication. We argue that agents do not only have limited knowledge of the world, but that they also can only *acquire* limited knowledge about the world. This holds especially for knowledge about the state of mind of other agents. In general it is not efficient for each agent to be able to "test" the truth of any statement about the world. This would require that all agents use the same language and have access to all facts about the world. However, one reason to introduce agents is to split up the work in manageable packets that can be handled by different agents. Each agents only reasons about its own part of the data. I.e. one agent for managing the weather reports and another agent to handle stock prices.

The same principle holds for the reasoning about actions. An agent cannot take into account all possible actions of other agents and possible events occurring in the environment. If an agent could do this, no unforeseen circumstances could arise and the goals would always be reached. Therefore, we assume that agents can only reason about a limited set of influences on their actions.

However, in order to describe the "actual" effect of actions (and especially communication) we need to use a global (agent independent) view. In this case the set of all agents is seen as one system. Using a global view on communication we can describe properties of communication protocols and proof their termination, fairness, etc.

In this paper we show how to describe the formal effects of communication both in the private view as well as in the global view. This gives rise to an integrated formal framework for communicating agents. An important aspect in the description of the effects of the communication is the use of deontic concepts.

This enables us to describe commitments resulting from communication without destroying the autonomy of the agents.

The second important point in this paper is the distinction between the private and the global view of the world in a formal framework and more specifically what are the consequences for the communication between agents. We describe a formal framework for communication that can be used to model all types of protocols. Instead of fixing some protocol the framework indicates possible meaningful sequences of messages for certain situations and goals of the agents. For instance, after a proposal is received a counterproposal can be given. However, it does not make sense to follow up a proposal with an identical counterproposal. The ultimate goal is to formally describe communication rules for autonomous agents. With these rules the effects of communication protocols (like contract net) can be calculated and more flexible ways of dealing with communication protocols can be devised.

In the next section we describe the four components that we use to describe autonomous communicating agents. In section 3 we show how communication can be formally described using our formalism, using the communication primitives for negotiating agents in the ADEPT system ([21]) as example. In section 4 we describe the differences between the local and global view on communication. In section 5 we give a sketch of a formalisation of the framework given in the previous sections. We give some conclusions in section 6.

2 Communicating Agents

The definition of the agents is based on the framework developed in [8, 9]. However, we added a private view on the actions. The concepts that we formalise can roughly be divided over four different components: the informational component, the action component, the motivational component and the social component. For readability we will mention all the concepts (including the ones described in previous publications) of each of these components in the following subsections. However, we will only go into the details of those concepts that are new for this paper.

2.1 The informational component

At the informational level we consider both knowledge and belief. Many formalisations have been given of these concepts and we will follow the more common approach in epistemic and doxastic logic: the formula $K_i\phi$ denotes the fact that agent i knows ϕ and $B_i\phi$ that agent i believes ϕ. We demand knowledge to obey an S5 axiomatisation, belief to validate a KD45 axiomatisation, and agents to believe all the things that they know.

2.2 The action component

In the action component we consider both dynamic and temporal notions. The main dynamic notion that we consider is that of actions, which we interpret as functions that map some some state of affairs into another one. Following [12, 26] we use parameterised actions to describe the event consisting of a particular agent's execution of an action. We let $\alpha(i)$ indicate that agent i performs the action α.

We can reason about the results of actions on both a private level and a global level. The global level reasoning is the "standard" one using dynamic logic as described by Harel in [11]. We use $[\alpha(i)]\phi$ to indicate that *if* agent i performs the action indicated by α the result will be ϕ. I.e. no matter what happens, if agent i performs α the system will change to a state where ϕ holds. Note that this is a very strong statement! No unforeseen action can disturb the execution of α by i.

We also introduce a private level of reasoning about actions in this paper. We use $[\alpha(i)]_j\phi$ to indicate that agent j concludes that ϕ will hold *if* agent i performs the action indicated by α. Each agent j will only consider a subset of all possible actions that might intervene with α. For instance, it might be that $[\mathbf{read-record}]_\mathbf{j}\mathbf{K_j(correct\ number\ of\ computers\ sold\ this\ year)}$. But if j did not consider that agent i could just update the sales database at the same time we also have (globally) $\neg[\mathbf{read-record}]\mathbf{K_j(correct\ number\ of\ computers\ sold\ this\ year)}$.

Besides these formulas that indicate the results of actions we also would like to express that an agent has the reliable opportunity to perform an action. This is done through the predicate OPP: $OPP(\alpha(i))$ indicates that agent i has the opportunity to do α, i.e. the event $\alpha(i)$ will possibly take place.

Besides the OPP operator, which already has a temporal flavour to it, we introduce two genuinely temporal operators: $PREV$, denoting the events that actually just took place, and the "standard" temporal operator $NEXT$, which indicates, in our case, which event will actually take place next. We also define a more traditional $NEXT$ operator on formulas in terms of the $NEXT$ operator on events.

$$NEXT(\phi) \text{ iff } NEXT(\alpha(i)) \wedge [\alpha(i)]\phi$$

This means that the formula ϕ is true in all next states iff an action $\alpha(i)$ is performed next and the formula ϕ is true after the performance of $\alpha(i)$.

In this paper we introduce two special action types. These are the *test* action and the *Reveal* action. Both actions have an epistemic character. Although the test action is already introduced in standard dynamic logic, we give it an epistemic flavour conform [17]. I.e. after i tests the truth of a formula i knows whether the formula is true or not. The test action on formula ϕ is written as $\phi?$. So, more formally we have:

$$\phi \rightarrow [\phi?(i)]K_i(\phi)$$
$$\neg\phi \rightarrow [\phi?(i)]K_i(\neg\phi)$$

As we argued before, an agent cannot test every possible formula. Every agent has a restricted domain on which it can perform tests. However, an agent i can reveal certain information to an agent j by using the reveal action. The result of this action is that agent j can test the truth of that formula himself. Formally:

$$[Reveal(i, j, \phi)]OPP(\phi?(j))$$

The reveal action is especially useful to function as grounding mechanism for discussions about the validity of some formula. It is equivalent to the physical action of showing some evidence as support to your claim.

2.3 The motivational component

In the motivational component we consider a variety of concepts, ranging from preferences, goals and decisions to intentions and commitments. The most fundamental of these notions is that of conditional preferences.(See also [1, 16]). Formally, (conditional) preferences are defined as the combination of implicit and explicit preferences. A formula ϕ is preferred by an agent i in situation ψ, denoted by $Pref_i(\phi|\psi)$, iff ϕ is true in all the states that the agent considers desirable when ψ is true, and ϕ is an element of a predefined set of (explicitly preferred) formulas. We assume a (total) ordering between the explicit preferences of each agent in each world. (The ordering may vary between worlds because the preferences are conditional upon some statement to hold true.) The use of conditional preferences, instead of the traditional "desires", makes it possible to use the qualitative decision theory developed in [1, 16] and also to make a connection with game theoretic work used for negotiations between agents (see e.g. [22]).

Goals are not primitive in our framework, but instead defined in terms of preferences. Informally, a preference of agent i constitutes one of i's goals iff i knows that the preference does not hold yet, but is *achievable*. Formally:

$$Achiev_i\phi \equiv \exists\beta : [\beta(i)]_i\phi \wedge OPP(\beta(i))$$

Note that we use $[\beta(i)]_i\phi$ to indicate that agent i privately concludes that ϕ holds after performing β. In most cases it will hold that (globally) $\neg[\beta(i)]\phi$ or even $[\beta(i)]\neg\phi$.
A goal is now formally defined as a preference that does not hold but is achievable:

$$Goal_i(\phi|\psi) \equiv Pref_i(\phi|\psi) \wedge \neg\phi \wedge Achiev_i\phi$$

Note that our definition implies that there are three ways for an agent to drop one of its goals: since it no longer considers achieving the goal to be desirable, since the preference now holds, or since it is no longer certain that it can achieve the goal. This shows that our framework complies to the standard notions of goals given in e.g. [2].

Goals can either be known or unconscious goals of an agent. Most goals will be known, but we will later on see that goals can also arise from commitments and these goals might not be known explicitly.

Intentions are divided in two categories, viz. the intention to perform an action and the intention to bring about a proposition. The latter category of

intentions is seen as goals in our framework.
We define the intention of an agent i to perform a certain action α as primitive, denoted by $INT_i\alpha$. An intention to perform an action is based on the decision to try to reach a goal. The agent can only make a decision to try to achieve the goal that has the highest preference (the utility principle). Because the order of the preferences may differ in each world, this does not mean that once a goal has been fixed the agent will always keep on trying to reach that goal (at least not straight away). The above is described formally by

$$\gamma \rightarrow OPP(DEC(i,\alpha)) \text{ iff } \exists\phi : Goal_i(\phi|\gamma)\wedge$$
$$\gamma \rightarrow [\alpha;\beta(i)]\phi \wedge \neg\exists\psi(Pref_i(\psi|\gamma) \wedge \phi <_i \psi)$$

$$OPP(DEC(i,\alpha)) \rightarrow [DEC(i,\alpha)]INT_i\alpha$$

There is no direct relation between the intention to perform an action and the action that is actually performed next. We do, however, establish an indirect relation between the two through a binary *implementation* predicate, ranging over pairs of actions. The idea is that the formula $IMP_i(\alpha_1,\alpha_2)$ expresses that for agent i executing α_2 is a reasonable attempt at executing α_1.

Having defined the binary IMP predicate, we may now relate intended actions to the actions that are actually performed. We demand the action that is actually performed by an agent to be an attempt to perform one of its intentions. Formally, this amounts to the formula

$$(INT_i(\alpha_1(i)) \wedge NEXT(\alpha_2(i))) \rightarrow IMP_i(\alpha_1,\alpha_2)$$

The last concept that we consider at the motivational level is that of commitment. Many interpretations have been given to the concept of commitment (see e.g. [2, 13, 15]). We chose a deontic interpretation of commitment. That is, a commitment of an agent to reach a goal is expressed as an obligation of the agent towards itself to reach the goal. Although the obligation does not ensure the actual performance of the action by the agent, it does have two practical consequences. If an agent commits itself to an action and afterwards does not perform the action a *violation* condition is registered, i.e. the state is not ideal (anymore).

The second consequence of registering a commitment as an obligation is, as we argued in [6], that obligations lead to (conditional) preferences which are ordered. From this it follows that an agent will be very committed to a goal if the preference following from a commitment has a very high ranking. In the other hand the commitment of an agent towards a goal is low if the generated preferences get a low ranking.
The relation between obligations and preferences is formally described as follows:

$$\forall i,j,\phi Pref_i(\phi|O_{ij}(\phi))$$

and for actions:

$$\forall i,j,\alpha Pref_i(PREV(\alpha(i))|O_{ij}(\alpha(i)))$$

Note that the latter is sufficient to create a goal if i has the opportunity to perform α, because $PREV(\alpha(i))$ does not hold presently (the action is not performed yet when the obligation arises) and it is achievable (by performing the action $\alpha(i)$.

The above connection between commitments and preferences (and thus goals) makes our agents sincere. Whenever an agent commits itself there automatically arises a preference to fulfil the commitment. Whether the commitment is kept depends on the priority of the resulting preference and the achievability of it. This is especially important if the commitment is made towards other agents. In that case the commitment forms a part of the social component. We will say more about the social component in the next section.

2.4 The social component

The $COMMIT$ described in the previous section is one of the four types of *speech acts* [24] that play a role in the social component. Speech acts are used to communicate between agents. The result of a speech act is a change in the doxastic or deontic state of an agent, or in some cases a change in the state of the world. The speech acts are the main actions for which synchronization between agents is essential. A speech act always involves at least two agents; a speaker and a hearer. If an agent sends a message to another agent but that agent does not "listen" (does not receive the message) the speech act is not successful. We will describe the speech acts first on the global level to indicate the interaction between the agents. Then we will show the private views of the agents on the speech acts.
The most important feature in which our framework for speech acts differs from other frameworks for speech acts (based on the work of Searle) is that a speech act in our framework is not just the sending of message by an agent but is the composition of sending and receiving of a message by two (or more) agents!

We distinguish the following speech act types: *commitments, directions, declarations* and *assertions*. The idea underlying a direction is that of giving orders, i.e. an utterance like 'Pay the bill before next week'. A typical example of a declaration is the utterance 'Herewith you are granted permission to access the database', and a typical assertion is 'I tell you that the earth is flat'. Each type of speech act should be interpreted within the background of the relationship between the speaker and the hearer of the speech act. In particular for directions and declarations the agent uttering the statement should have some kind of basis of authority for the speech act to have any effect.

We distinguish three types of relations between agents: *peer* relation, *power* relation and *authorization* relation. The first two relations are similar to the ones used in the ADEPT system [21, 14]. The power relation is used to model hierarchical relations between agents. We assume that these relations are fixed during the lifecycle of the agents. Within such a relation less negotiation is possible about requests and demands. This reduces the amount of communication and therefore increases the efficiency of the agents.

The peer relation exists between all agents that have no prior contract or obligations towards each other (with respect to the present communication). This relation permits extensive negotiations to allow a maximum of autonomy for the agents.
The last relation between agents is the authorization relation which is a type of temporary power relation that can be build up by the agents themselves.

The power relation is formalized as a partial ordering between the agents, which is expressed as follows: $i \ll j$ means that j has a higher rank than i.
The authority relation is formalized through a binary predicate $auth$; $auth(i, \alpha)$ means that agent i is authorised to perform α. It seems that this specifies a property of one agent, however, the other agent is usually part of the specification of α. Therefore the authorization to perform an action implicitly determines an authorization relation between the agents involved in that action as well.

One way to create the authorisation relations is by agent j giving an implicit authorisation to i to give him some directives. For example, when agent i orders a product from agent j it implicitly gives the authorisation to agent j for demanding payment from i for the product (after delivery). We will see later that most communicative actions have also implicit components and effects that are usually determined by the context and conventions within which the communication takes place.

Besides the implicit way to create authorizations, they can also be created explicitly by a separate speech act which is formally a declaration that the authorization is true.

The speech acts themselves are formalised as meta-actions (based on earlier work [5]):

- $DIR(x, i, j, \alpha)$ formalises that agent i directs agent j to perform α on the basis of x, where x can be either *peer*, *power* or *authority*.
- $DECL(i, f)$ models the declaration of i that f holds.
- $ASS(x, i, j, f)$ formalises the assertion of i to agent j that f holds.
- $COMMIT(i, j, \alpha)$ describes that i commits itself towards j to perform α.

Note that the commit and the declarative do not take a relation parameter. This is basically because the effect of a commit is the same irrespective of the relation between the agents, while the declarative does only involve one agent.

A directive from agent i to agent j to perform α results in an obligation of j towards i to perform that action *if* agent i was either in a power relation towards j or was authorized to give the order. In a similar way the assertion of proposition f by i to j results in the fact that j will believe f *if* I had authority over j. Creating the authorizations is an important part of the negotiation between agents when they are establishing some type of contract. On the basis of the authorizations that are created during the negotiation some protocol for the transactions between the agents can be followed quick and efficiently. (See [25] for more details on contracts between agents).

Formally, the following formulas hold for the effects of commitments, orders and declaratives:

- $[COMMIT(i,j,\alpha)][DECL(j,P_{ij}(\alpha(i)))]O_{ij}\alpha$
- $auth(i,DIR(authority,i,j,\alpha)) \rightarrow [DIR(authority,i,j,\alpha)]O_{ji}\alpha$
- $j \ll i \rightarrow [DIR(authority,i,j,\alpha)]O_{ji}\alpha$
- $[DIR(peer,i,j,\alpha)]K_j INT_i \alpha(j)$
- $auth(i,DECL(i,f)) \rightarrow [DECL(i,f)]f$
- $[DECL(i,f)]Pref_i(f|true)$
- $[ASS(peer,i,j,f)]K_j B_i f$
- $auth(i,ASS(authority,i,j,f)) \rightarrow [ASS(authority,i,j,f)]B_j f$
- $j \ll i \rightarrow [ASS(power,i,j,f)]B_j f$

A commitment always results in a kind of conditional obligation. The obligation is conditional on the permission of the agent towards which the commitment is made. (This is very close to the ACCEPT action in other frameworks). The giving of permission is formally described by $[DECL(j,P_{ij}(\alpha(i)))]$, where $P_{ij}(\alpha(i)) \equiv \neg O_{ij}(\overline{\alpha(i)})$. I.e. the permission to perform α is equivalent to the fact that there is no obligation to perform the negation of α.
The permission of j is necessary because j might play a (passive) role in the action α initiated by i. Of course j must be willing to play its part. It signifies this by giving the permission to i. In contrast to the other speech acts no precondition has to hold for a commitment to obtain its desired result.

A directive from agent i results in an obligation of agent j (towards i) if agent i was authorised to give the order or i has a power relation towards j.
If i has no authority or power over j then the directive is actually a request. It results in the fact j knows that i wants him to perform α. If j does not mind to perform it can commit himself to perform α and create an obligation.

Assertions can be used to transfer beliefs from one agent to another. Note that agent j does not automatically believe what agent i tells him. We do assume that agents are sincere and thus we have the following axiom:

$$OPP(ASS(x,i,j,f)) \rightarrow B_i f$$

That is, an agent can only assert facts that it believes itself.
The only way to directly transfer a belief is when agent i is authorised to make a statement. Usually this situation arises when agent j first requested some information from i. Such a request for information (modelled by a directive without authorisation) gives an implicit authorisation on the assertions that form the answer to the request.

A declaration can change the state of the world if the agent making the declaration is authorised to do so. (This is the only speech act that has a direct effect on the states other than a change of the mental attitudes of the agents!). If agent i has no authority to declare the fact, then the only result of the speech act is that i establishes a preference for itself. It prefers the fact to be true.

Although we do not attempt to give a (complete) axiomatization, we want to mention the following axioms for the declaratives, because they are very fundamental for creating relationships between agents.

$$[DECL(i,auth(j,DIR(authority,j,i,\alpha(i))))]auth(j,DIR(authority,j,i,\alpha(i)))$$

which states that an agent i can create authorisations for an agent j concerning actions that i has to perform.
The following axiom is important for the acceptance of offers:

$$[DECL(i, P_{ji}(\alpha(i)))]P_{ji}(\alpha(i))$$

which states that an agent can always give permission to another agent to perform some action.
Note that it may very well be that another agent forbids j to perform α! The permission is only with respect to i!

3 Formal Communication

In the previous section we gave a brief overview of the basic messages that agents can use in our framework. To show the power of our framework and to show the relation with other work on communication between agents we show how the basic illocutions that are used for the negotiating agents in the ADEPT system (and that also form the heart of many other negotiation systems) can be modelled within our framework. We only show this for the negotiation because it forms an important part of the communication between agents. In a later paper we will show how the communication in the stages after the negotiation (the performance and satisfaction stages) can also be formally modelled in our framework.

The negotiating agents in the ADEPT system use the four illocutions: PROPOSE, COUNTERPROPOSE, ACCEPT and REJECT. These four illocutions also form the basic elements of many other negotiation systems.

The PROPOSE is directly translated into a COMMIT. The obligation that follows from a proposal depends on the acceptance of the receiving party. However, the ACCEPT that is used as primitive in ADEPT and most other systems involves more than the giving of permission that we already indicated above.

The ACCEPT message has three components. That is, we consider the ACCEPT to be the simultaneous expression of three illocutions.

1. Giving permission to perform the action
2. Commitment to perform those actions that are necessary to make the proposal succeed
3. Giving (implicit) authority for subsequent actions (linked to the proposal by convention)

For example if agent i sends the following message to j:

> PROPOSE,i,j,
> I will deliver 20 computers (pentium, 32M, etc.) to you for $1000,- per computer

then the ACCEPT message of j to i:

ACCEPT,j,i,
You will deliver 20 computers (pentium, 32M, etc.) to me for $1000,- per computer

means:

1. You are permitted to deliver the computers: $DECL(j, P_{ij}(deliver))$
2. I will receive the computers (sign a receipt): $COMMIT(j, i, receive)$
3. I give you authority to ask for payment after delivery: $DECL(j, [deliver]auth(i, DIR(authority, i, j, pay)))$

It is important to notice that only the first component of the meaning of the ACCEPT message is fixed. The other two components depend on the action involved and the conventions (contracts) under which the transaction is negotiated.

The REJECT message is the denegation of the ACCEPT message. It means that the agent is either not giving permission for the action, not committing itself to its part of the action or not willing to give authority to subsequent actions. Formally this is expressed as the disjunction of the negation of these three parts. Due to space limitations we will not work this out any further.

The COUNTERPROPOSE is a composition of a REJECT and PROPOSE message. Formally it can thus be expressed as the parallel execution of these two primitives.

Besides the formal representation of the illocution of the message we can also give some preconditions on the basic message types. Only the PROPOSE message type does not have preconditions. This is as expected because the PROPOSE is used to start the negotiation. The other types of messages are all used as answer to a PROPOSE (or COUNTERPROPOSE) message. We can formally describe the precondition that these message types can only be used after a PROPOSE or COUNTERPROPOSE as follows:

- $OPP(ACCEPT(j, i, \alpha)) \leftrightarrow (PREV(PROPOSE(i, j, \alpha)) \vee PREV(COUNTERPROPOSE(i, j, \alpha)))$
- $OPP(REJECT(j, i, \alpha)) \leftrightarrow (PREV(PROPOSE(i, j, \alpha)) \vee PREV(COUNTERPROPOSE(i, j, \alpha)))$
- $OPP(COUNTERPROPOSE(j, i, \beta)) \leftrightarrow \beta \neq \alpha \wedge (PREV(PROPOSE(i, j, \alpha)) \vee PREV(COUNTERPROPOSE(i, j, \alpha)))$

In the precondition of the COUNTERPROPOSE we included the fact that a counterproposal should differ from the proposal that it counters. (Although not mentioned in this paper, the semantics of actions does give an equivalence relation between actions). More elaborate conversation rules are needed to describe long term dependencies within protocols. E.g. one cannot repeat the same proposal later on if it already has been rejected. These rules should be incorporated within the protocols that the agents are using.

We do not want to give the formalisation of complete protocols at this place due to space limitations. However, we can indicate quite easily the results of the most common pairs of messages where agent i first proposes something to agent j after which agent j can accept it, reject it or counterpropose it. These moves are formally described as follows:

- $[PROPOSE(i,j,\alpha)(i)][ACCEPT(j,i,\alpha)(j)]O_{ij}(\alpha(i)) \wedge P_{ji}(\alpha(i))$ (accept)
  Furthermore, if the success of $\alpha(i)$ depends on the performance of $\beta(j)$ by j:
  $[PROPOSE(i,j,\alpha)(i)][ACCEPT(j,i,\alpha)(j)]O_{ji}(\beta(j))$
  And if conventions determine that i can perform $\beta(i)$ after acceptance of the proposal then:
  $[PROPOSE(i,j,\alpha)(i)][ACCEPT(j,i,\alpha)(j)][\alpha(i)]auth(i,\beta(i))$
- $[PROPOSE(i,j,\alpha)(i)][REJECT(j,i,\alpha)(j)]\neg O_{ij}(\alpha(i))$ (reject)
- $[PROPOSE(i,j,\alpha)(i)][COUNTERPROPOSE(j,i,\beta)(j)]\neg O_{ij}(\alpha(i))$ (counter)

Note that the counterproposal has no effect of itself yet. Only the reject component of the counterproposal has immediate effect. The proposal component of the counterproposal only takes effect after an appropriate answer of i.
For the reject we only indicated that the obligation does not arise. The rest of the effect depends on the context and is usually not of prime interest.

The formalisation of the basic messages in the ADEPT system shows two things.
First, that our framework is powerful enough to formally describe the negotiation in the ADEPT system including the effects of the communication.
Secondly, that seemingly simple message types, like ACCEPT, have complicated meanings that partly depend on the context in which they are used.

4 Private and global views on communication

In the previous sections we gave a formal description of communication between agents. This description was given from a global viewpoint. That is, the communication was seen as actions that change the complete system of agents from one state to another state. This is quite natural when considering material actions like database updates. If an agent changes a database, the system will be in a different state where some values in the database are changed. No other agents are necessarily (directly) involved in this action. However, communicative actions (except for the declaratives) always require the participation of two agents: the speaker and the hearer.

In this section we will give a private view on communication based on the global view defined in the previous sections. In a private view of the system we try to ascribe each action, that takes place in the system, to an agent that has control over that action. Also we try to make clear which part of the system can be "seen" by each of the agents. I.e. which formulas can be checked by the agents.
To explain the private description of the communication between agents we will use only one type of message. All remarks hold mutatis mutandis for the other

types of messages.
In a global view we have the following axiom for directives:

$$auth(i, DIR(authority, i, j, \alpha)) \rightarrow [DIR(authority, i, j, \alpha)]O_{ji}\alpha$$

I.e. after an authorized directive an obligation arises.

In the private view the following features of communication can be better described:

1. Communication consists of speaking and listening.
2. Speaker and hearer might not share the same language.
3. Not all pre-conditions and effects of communications can be (directly) checked by both speaker and hearer.

Ad.1. The first and most important step that should be taken to privatize the view on this communication is to split up this action into a speaker and hearer part. Agent i can never perform the complete directive by itself. It can only send the message and hope that agent j receives the message. So, although agent i initiates the action it does not have complete control over it. It cannot assure that the action completes successfully. Because there is not a single entity that has control over the communicative actions we will split up the communicative actions into a send and receive action to get a private view on them. $DIR(authority, i, j, \alpha) \equiv$

$$send(DIR(authority, i, j, \alpha))(i) \& receive(DIR(authority, i, j, \alpha))(j)$$

The parallel decomposition of the directive should be read as a synchronization between the agents. In an actual implementation the actions might be serialized.

Although in the global view we cannot assume that an obligation holds after the sending of (an authorized) directive by agent i, agent i can privately conclude this if we assume the following axiom:

$$auth(i, DIR(authority, i, j, \alpha)) \rightarrow [send(DIR(authority, i, j, \alpha))]_i O_{ji}\alpha$$

This means that agent i assumes that agent j will always receive the messages that agent i sends.

In the same way we have of course (and with more right probably):

$$auth(i, DIR(authority, i, j, \alpha)) \rightarrow [receive(DIR(authority, i, j, \alpha))]_j O_{ji}\alpha$$

That is, if agent j receives an authorized directive it will conclude that it now has an obligation towards i.

Ad.2. Because the communication is now split up into a send and receive part it is also possible to indicate whether the receiver can "understand" the message that was send. I.e. whether the receiving agent talks the same language in terms of formulas that it incorporates in its private language. It is possible to incorporate some general *translation rules* in the system that indicate how terms can be translated from one agent's language to another's. In this paper we will assume that all agents use the same language in order not to complicate the formalisation to much. See [20] for an example how an agent system can be described in which agents can use different languages.

Ad.3. The last part that plays a role in the privatization of communication is the checking of the pre-conditions and effects of communication. If agent j does not know that agent i is authorized to give him an order it might not accept the consequent obligation. Often agent j can also not check the authority directly. Therefore, we think that in each protocol it should be possible for j to question the authority of i if j cannot check this authority himself. This is conform the theory from Habermas about communication protocols [10] where this is classified as an attack on the validity claims. Agent j can attack the validity of the authority of i by directing agent i to make the authority available for inspection of agent j. We get the following possibilities:

1. $(auth(i, DIR(authority, i, j, \alpha)) \wedge$
   $OPP(auth(i, DIR(authority, i, j, \alpha))?(j))\ ) \rightarrow$
   $[DIR(authority, i, j, \alpha)]O_{ji}\alpha$
   I.e. if agent j has the opportunity to check the authority of agent i then the authoritative direction of i to j to perform α results in an obligation.
2. $(auth(i, DIR(authority, i, j, \alpha)) \wedge$
   $\neg OPP(auth(i, DIR(authority, i, j, \alpha))?(j))\ ) \rightarrow$
   $[DIR(authority, i, j, \alpha)]$
   $auth(j, DIR(auth., j, i, Reveal(i, j, (auth(i, DIR(auth., i, j, \alpha)))))$
   If agent j does not have the opportunity to check the authority of i then the direction of i only results in the authority of j to direct i to reveal the status of his authority to j. We admit that this formula is not very readable, but it is of course very easy to find some suitable abbreviations for these standard formulas.

The establishment of the truth of the authority of i does not have to be the end of the discussion, because, according to Habermas, agent j might now question the reason for this authority. For instance, it is based on law, on a previous agreement, on a contract, etc. We will not go further into this at this place.

The above points indicate that the private view on communication between agents reveals new aspects of the communication that are not visible in the global view. Especially the difference in awareness about actions and facts by different agents leads to new communicative acts that did not seem necessary in the global view.

5 A sketch of a formalisation

In this section we precisely define the language that we use to formally represent the concepts described in the previous sections, and the models that are used to interpret this language. We will not go into too much detail with regard to the actual semantics, but try to provide the reader with an intuitive grasp for the formal details without actually mentioning them.

The language that we use is a multi-modal, propositional language, based on three denumerable, pairwise disjoint sets: Π, representing the propositional symbols, Ag representing agents, and At containing atomic action expressions. The language $FORM$ is defined in four stages. Starting with a set of propositional formulas ($PFORM$), we define the action- and meta-action expressions, after which $FORM$ can be defined.

The set Act of regular action expressions is built up from the set At of atomic (parameterised) action expressions (denoted by $\underline{a}$...) using the operators ; (sequential composition), + (nondeterministic composition), & (parallel composition), and $^-$ (action negation). The constant actions **any** and **fail** denote 'don't care what happens' and 'failure' respectively.

Definition 1. *Let $\underline{a} \in At$ then the set Act of action expressions is given by the following BNF:*

$$\alpha :: -\underline{a}|\mathbf{any}|\mathbf{fail}|\alpha_1 + \alpha_2|\alpha_1 \& \alpha_2|\overline{\alpha}$$

The set $MAct$ of general action expressions contains the regular actions and all of the special meta-actions informally described in section 2. For simplicity, we restrict ourselves in this paper to closing the set $MAct$ under sequential composition.

Definition 2. *Let $\alpha \in Act$, $i, j \in Ag$ and $x \in \{peer, authority, power\}$ then the set $MAct$ of general action expressions is given by the following BNF:*

$$\gamma\alpha :: -\alpha|DEC(i,\alpha)|COMMIT(i,j,\alpha)|DIR(x,i,j,\alpha)|\gamma\alpha_1;\gamma\alpha_2$$

Not all actions can be defined at this level, because some actions like $DECL$ contain formulas from $FORM$ as parameters. These actions will be defined in the next stage.

The complete language $FORM$ is now defined to contain all the constructs informally described in the previous section. That is, there are operators representing informational attitudes, motivational attitudes, aspects of actions, and the social traffic between agents.

Definition 3. *Let $\psi \in PFORM$, $\gamma\alpha \in Mact$, $\alpha, \alpha_1, \alpha_2 \in Act$, $i, j, k \in Ag$ and $x \in \{peer, authority, power\}$ then the language $FORM$ of formulas is given by the following BNF:*

$$\begin{aligned}\phi :: -\ &\psi|\neg\phi|\phi_1 \wedge \phi_2|K_i\phi|B_i\phi|[\gamma\alpha]\phi|[\gamma\alpha]_i\\ &[DECL(i,\psi)]\phi|[ASS(x,i,j,\psi)]\phi|[Reveal(i,j,\psi)]\phi|[\psi?(i)]\phi\\ &[DECL(i,\psi)]_k\phi|[ASS(x,i,j,\psi)]_k\phi|[Reveal(i,j,\psi)]_k\phi|[\psi?(i)]_k\phi\\ &[\gamma\alpha;\gamma\beta]\theta|[\gamma\alpha;\gamma\beta]_i\theta|PREV(\alpha)|OPP(\alpha)|NEXT(\phi)\\ &Pref_i(\phi|\psi)|\psi <_i \phi|i \ll j|INT_i\alpha|IMP_i(\alpha_1|\alpha_2)|O_{ij}(\alpha)|auth(i,\alpha)\end{aligned}$$

Note that the ASS, $DECL$, $Reveal$ and $test$ action are introduced in $FORM$ at this stage. The postcondition ϕ does not have any meaning except as a placeholder in these formulas.

The models used to interpret $FORM$ are based on Kripke-style possible worlds models. That is, the backbone of these models is given by a set Σ of states, and a valuation π on propositional symbols relative to a state. Various relations and functions on these states are used to interpret the various (modal) operators. These relations and functions can roughly be classified in four parts, dealing with the informational component, the action component, the motivational component and the social component, respectively. We assume tt and ff to denote the truth values 'true' and 'false', respectively.

Definition 4. *A model Mo for $FORM$ from the set CMo is a structure $(\Sigma, \pi, I, A, M, S)$ where*

1. *Σ is a non-empty set of states and $\pi : \Sigma \times \Pi \to \{tt, ff\}$.*
2. *$I = (Rk, Rb)$ with $Rk : Ag \to \wp(\Sigma \times \Sigma)$ denoting the epistemic alternatives of agents and $Rb : Ag \times \Sigma \to \wp(\Sigma)$ denoting the doxastic alternatives.*
3. *$A = (Sf, Mf, Sfa, Mfa, Ropp, Rprev, Rnext)$ with $Sf : Ag \times Act \times \Sigma \to \wp(\Sigma)$ yielding the global interpretation of regular actions, $Mf : Ag \times MAct \times (CMo \times \Sigma) \to (CMo \times \Sigma)$ yielding the global interpretation of meta-actions, $Sfa : Ag \times Ag \times Act \times \Sigma \to \ wp(\Sigma)$ yielding the private interpretation of of regular actions, $Mfa : Ag \times Ag \times MAct \times (CMo \times \Sigma) \to (CMo \times \Sigma)$ yielding the private interpretation of meta-actions, $Ropp : Ag \times \Sigma \to \wp(Act)$ denoting opportunities, $Rprev : Ag \times \Sigma \to Act$ yielding the action that has been performed last and $Rnext : Ag \times \Sigma \to Act$ yielding the action that will be performed next.*
4. *$M = (Rp, Rep, <, Ri, Ria, Ro)$ with $Rp : Ag \times \Sigma \to \wp(\Sigma)$ denoting implicit preferences, $Rep : Ag \times \Sigma \to \wp(FORM)$ yielding explicit preferences, $<\subseteq Ag \times \Sigma \to FORM \times FORM$ which is a preference relation on preferences, $Ri : Ag \times \Sigma \to \wp(Act)$ denoting intended actions, $Ria : Ag \times \Sigma \to \wp(Act) \times \wp(Act)$ denoting implementation relations between actions and $Ro : Ag \times Ag \to \wp(\Sigma \times \Sigma)$ denoting obligations.*
5. *$S = (Auth, \prec)$ with $Auth : Ag \times \wp(MAct) \to \{tt, ff\}$ yielding authorisations and $\prec : Ag \times Ag \to \{tt, ff\}$ yielding hierarchical relations between agents.*

such that the following constraints are validated:

1. *$Rk(i)$ is an equivalence relation for all i, and $Rb(i,s) \neq \emptyset$, $Rb(i,s) \subseteq \{s' \mid (s,s') \in Rk(i)\}$ and $(s,s') \in Rk(i) \Longrightarrow Rb(i,s) = Rb(i,s')$, which ensures that knowledge validates an S5 axiomatisation and belief obeys a KD45 axiomatisation, while agents indeed believe all things they know.*
2. *Sf yields the global state-transition interpretation for regular actions. This function satisfies the usual constraints ensuring an adequate interpretation of composite actions in terms of their constituents. Sfa satisfies the same constraints as Sf but also should satisfy that $Sfa(i,j,\alpha,s) \subseteq Sf(j,\alpha,s)$. I.e. the private interpretation of an action is more limited than the global one. The function Mf models the global model-transforming interpretation of meta-action. Because we do not allow the composition of meta-actions with other actions yet, we require for the moment that $Mf \equiv Mfa$. Below we elaborate on the definition of Mf for the meta-actions introduced in the previous section.*

3. $Rnext(i,s) \in Ropp(i,s) \subseteq \{\alpha \mid Sf(i,\alpha,s) \neq \emptyset\}$*, which ensures that opportunities are a subset of the actions that are possible by virtue of the circumstances and that the next action performed is an opportunity. Furthermore,* $Rprev(i,s) = \alpha$ *iff* $\alpha \in Ropp(i,s')$ *for some* s' *with* $s \in Sf(i,\alpha,s')$*, which relates previously executed actions to past opportunities.*
4. $Ri(i,s) \subseteq \{\alpha \mid Sf(i,\alpha,s) \neq \emptyset\}$ *and for all* $s \in \Sigma$ *some* $s' \in \Sigma$ *exists with* $(s,s') \in Ro$.

The complete semantics contains an algebraic semantics of action expresses, based on the action semantics of Meyer [18]. In this paper we will abstract from the algebraic interpretation of actions and instead interpret actions as functions on states of affairs. For the meta-actions the state-transition interpretation is not adequate, because meta-actions do not change states but they change relations between states. For instance, in the case of an assertion, the effect is to change the doxastic state of the receiving agent, and nothing else. To formalise this behaviour, we interpret meta-actions as model-transforming functions. In the case of an assertion, the resulting model will differ from the starting model in the doxastic accessibility relation of the receiving agent.

Definition 5. *The binary relation* $\models$ *between an element of* $FORM$ *and a pair consisting of a model* Mo *in* CMo *and a state* s *in* Mo *is for propositional symbols, conjunctions and negations defined as usual. Epistemic formulas* $K_i\phi$ *and doxastic formulas* $B_i\phi$ *are interpreted as necessity operators over* Rk *and* Rb *respectively. For the other formulas* $\models$ *is defined as follows:*

$$
\begin{array}{lcl}
Mo,s \models [\alpha(i)]\phi & \Longleftrightarrow & Mo,s' \models \phi \text{ for all } s' \in Sf(i,\alpha,s)\\
Mo,s \models [\alpha(i)]_j\phi & \Longleftrightarrow & Mo,s' \models \phi \text{ for all } s' \in Sfa(j,i,\alpha,s)\\
Mo,s \models [\gamma\alpha(i)]\phi & \Longleftrightarrow & Mo',s' \models \phi \text{ for all } Mo',s' \in Mf(i,\alpha,Mo,s)\\
Mo,s \models [\gamma\alpha(i)]_j\phi & \Longleftrightarrow & Mo',s' \models \phi \text{ for all } Mo',s' \in Mfa(j,i,\alpha,Mo,s)\\
Mo,s \models PREV(\alpha(i)) & \Longleftrightarrow & \alpha \in Rprev(i,s)\\
Mo,s \models OPP(\alpha(i)) & \Longleftrightarrow & \alpha \in Ropp(i,s)\\
Mo,s \models NEXT(\alpha(i)) & \Longleftrightarrow & \alpha(i) \in Rnext(i,s)\\
Mo,s \models Pref_i(\phi|\psi) & \Longleftrightarrow & \text{If } Mo,s \models \psi \text{ then}\\
 & & Mo,s' \models \phi \text{ for all } s' \in Rp(i,s) \text{ and } \phi \in Rep(i,s)\\
Mo,s \models \psi <_i \phi & \Longleftrightarrow & (\psi,\phi) \in <(i,s)\\
Mo,s \models i \ll j & \Longleftrightarrow & i \prec j\\
Mo,s \models INT_i\alpha & \Longleftrightarrow & \alpha \in Ri(i,s)\\
Mo,s \models IMP_i(\alpha_1,\alpha_2) & \Longleftrightarrow & (\alpha_1,\alpha_2) \in Ria(i,s)\\
Mo,s \models O_{ij}(\phi) & \Longleftrightarrow & Mo,s' \models \phi \text{ for all } s' \text{ with } (s,s') \in Ro(i,j)\\
Mo,s \models O_{ij}(\alpha) & \Longleftrightarrow & Mo,s \models [\mathbf{any}(i)]O_{ij}(PREV(\alpha(i)))\\
Mos, \models auth(i,\alpha) & \Longleftrightarrow & Auth(i,\alpha,s) = tt
\end{array}
$$

The functions interpreting the special meta-actions (**?**, *Reveal, DEC, COMMIT, DIR, DECL* and *ASS*) can be described in terms of the preconditions and the postconditions for execution of the actions. Due to space limitations we leave them out here. See [9] for more details.

6 Conclusions

In this paper we have shown that it is possible to formally describe communicating agents. The emphasis in this paper was on the formal description of the communication between agents. A very important aspect of this formalism is that it is possible to (formally) describe the effects of the communication. Therefore it is possible to check what is the resulting situation after a communication protocol has been followed. We can analyze a protocol and find out what are reasonable moves at any point in the protocol. We have shown how the message types of the ADEPT system can be described in our primitives. This revealed that a seemingly simple primitive like ACCEPT contains a lot of hidden meanings.

We have also shown in this paper that there exists an important difference between a private and global view on actions and in particular communications. The private view opens up new communication moves in the negotiation because the agents involved have different information!
The difference becomes of prime importance when we want to implement agents that have to follow the rules of our logical formalism. By using a private view of actions it becomes clear which agent has control over each action. This is important because in the implemented system each action has to be initiated by some agent.
The private view on actions also makes it possible to introduce unforeseen actions, which seems more realistic in a multi-agent system which usually has an open character. I.e. not all the actions of all agents can be checked all the time.

Two remarks should be made about the logical formalism. First, it is not our aim to build an automated theorem prover that can prove theorems in this very rich logic. The use of a logical formalism gives the opportunity to automatically generate the logical effects of a sequence of steps in a protocol. These could be subsequently implemented in a more efficient formalism. The logical description, however, can be used as a very general and precise specification of that implementation.
Secondly, the use of logic forces a very precise formal description of the communication. The use of logic led to the discovery that the primitive ACCEPT message has actually several components, some of which depend on the context within which the ACCEPT is used. It is very important that this is realized when the communication protocols are automatized. (As is the aim in communication between agents.)
We admit that the logical formulas get very complicated and are not very readable. However, it is easy to define suitable abbreviations for standard formulas. At least, working this way, it is clear what these abbreviations mean exactly!

References

1. C. Boutilier. Toward a Logic for Qualitative Decision Theory. In Jon Doyle, Erik Sandewall and Pietro Torasso (eds.), *Principles of Knowledge Representation and*

Reasoning, proceedings of the fourth international conference, pages 75-86, 1994, Morgan Kaufmann Publishers, San Francisco, California.

2. P. Cohen and H. Levesque. Intention is choice with commitment. *Artificial Intelligence*, vol.42, pages 213-261, 1990.
3. P. Cohen and H. Levesque. Teamwork *Nous*, vol.35, pages 487-512, 1991.
4. R. Davis and R. Smith. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, vol.20, pages 63-109, 1983.
5. F. Dignum and H. Weigand. Communication and deontic logic. In R. Wieringa and R. Feenstra, editors, *Information Systems, Correctness and Reusability*, pages 242–260. World Scientific, Singapore, 1995.
6. F. Dignum. Autonomous Agents and Social Norms. Submitted to ICMAS workshop on Norms, Obligations and Conventions.
7. F. Dignum, J.-J.Ch. Meyer, R. Wieringa and R. Kuiper. A modal approach to intentions, commitments and obligations: intention plus commitment yields obligation. In M.A. Brown and J. Carmo (eds.) *DEON'96 Workshop on deontic logic in computer science*, pages 174-193, Lisbon, Jan. 1996.
8. F. Dignum and B. van Linder. Modelling Rational Agents in a Dynamic Environment: Putting Humpty Dumpty Together Again. In J.L. Fiadeiro and P.-Y. Schobbens (eds.) *ModelAge-96*, pages 81-92,Sesimbra, Portugal, 1996.
9. F. Dignum and B. van Linder. Modeling Social Agents: Communication as Action In J. Mueller, M. Wooldridge and N. Jennings (eds.) *Intelligent Agents III - Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages (ATAL-96)*, pages 83-93, Budapest, Hungary, 1996.
10. J. Habermas. The Theory of Communicative Action: Reason and Rationalization of Society. Polity Press, Cambridge, 1984.
11. D. Harel. First Order Dynamic Logic. LNCS 68 Springer, 1979.
12. W. van der Hoek, B. van Linder and J.-J.Ch. Meyer. A logic of capabilities. In Nerode and Matiyasevich, eds, *Proceedings of LFCS'94*, LNCS 813, pages 366-378, 1994.
13. N. Jennings. Commitments and Conventions: The foundation of coordination in Multi-Agent systems. *Knowledge Engineering Review*, vol. 8(3), pages 223-250, 1993.
14. N. Jennings, P. Faratin, M. Johnson, P. O'Brien and M. Wiegand. Using Intelligent Agents to Manage Business Processes. In *Proceedings The Practical Application of Intelligent Agents and Multi-Agent Technology*, pages 345-360, London, 1996.
15. D. Kinny and M. Georgeff. Commitment and Effectiveness of Situated Agents. In *Proceedings Int. Joint Conf. on Artificial Intelligence*, pages 82-88, Sydney, Australia, 1991.
16. J. Lang. Conditional Desires and Utilities - an alternative logical approach to qualitative decision theory. In W. Wahlster, editor, *Proceedings of ECAI-96*, pages 318-327, Budapest, Hungary, 1996, John Wiley & Sons Ltd.
17. B. van Linder, W. van der Hoek and J.-J.Ch. Meyer. Tests as Epistemic Updates. Pursuit of Knowledge. Technical Report, UU-CS-1994-08, Utrecht University, 1994.
18. J.-J.Ch. Meyer. A different approach to deontic logic. In *Notre Dame Journal of Formal Logic*, vol.29, pages 109–136, 1988.
19. J. Muller. A cooperation model for autonomous agents. In J. Muller, M. Wooldridge and N. Jennings, eds, *Intelligent Agents III - Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages (ATAL-96)*, pages 135-147, Budapest, Hungary, 1996.

20. P. Noriega and C. Sierra. Towards layered Dialogical Agents In J. Muller, M. Wooldridge and N. Jennings, eds, *Intelligent Agents III - Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages (ATAL-96)*, pages 69-82, Budapest, Hungary, 1996.
21. T. Norman, N. Jennings, P. Faratin and E. Mamdani Designing and Implementing a Multi-Agent Architecture for business process management. In J. Mueller, M. Wooldridge and N. Jennings (eds.) *Intelligent Agents III - Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages (ATAL-96)*, pages 149-162, Budapest, Hungary, 1996.
22. J. Rosenschein and G. Zlotkin *Rules of Encounter* MIT Press, Cambridge Massachusetts, 1994.
23. G. Sandu. Reasoning about collective goals. In J. Muller, M. Wooldridge and N. Jennings, eds, *Intelligent Agents III - Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages (ATAL-96)*, pages 35-47, Budapest, Hungary, 1996.
24. J.R. Searle. Speech Acts. Cambridge University Press. 1969.
25. H. Weigand, E. Verharen and F. Dignum. Interoperable Transactions in Business Models: A Structured Approach. In P. Constantopoulos, J. Mylopoulos and Y. Vassiliou, eds, *Advanced Information Systems Engineering (LNCS 1080)*, pages 193-209, Springer, 1996.
26. R. Wieringa, J.-J.Ch. Meyer and H. Weigand. Specifying dynamic and deontic integrity constraints. *Data & knowledge engineering*, vol.4, pages 157-189, 1989.