Modelling Social Agents: Towards Deliberate Communication

Frank Dignum
Faculty of Mathematics & Computer Science, Eindhoven University of Technology
P.O. Box 513, 5600 MB Eindhoven, The Netherlands
dignum@win.tue.nl
Bernd van Linder
ABN AMRO
P.O. Box 283 Amsterdam, The Netherlands
bernd.van.linder@nl.abnamro.com

July 22, 1999

Abstract

In this paper we present a formal framework for social agents. The social agents consist of four components: the information component (containing knowledge and belief), the action component, the motivational component (where goals, intentions, etc. play a role) and the social component (containing aspects of speech acts and relations between agents). The main aim of this work was to describe all components in a uniform way, such that it is possible to verify each component separately but also formally describe the interactions between the different components. This is of prime importance in order to integrate communicative actions into the plan and goal formation of the agent.

1 Introduction

The formalization of rational agents is a topic of continuing interest in AI. Research on this subject has held the limelight ever since the pioneering work of Moore [18] in which knowledge and actions are considered. Over the years contributions have been made on both *informational* attitudes like knowledge and belief [17] and *motivational* attitudes like intentions and commitments [2, 7].

In our basic framework [9, 15] we modelled the informational attitudes of agents as well as various aspects of action by means of a theory about the *knowledge*, *belief* and *abilities* of agents, as well as the *opportunities* for, and the *results* of their actions. In this framework it can for instance be modelled that an agent knows that it is able to perform an action and that it knows that it is correct to perform that action to bring about some result.

In [7, 14] we dealt with the motivational attitudes of agents. In these papers we defined the concepts of *wishes*, *goals*, *intentions*, and *commitments* or *obligations*. By combining this formalization with the basic framework it is for instance possible to model the fact that an agent prefers some situation to hold and it also knows that it is able to achieve that situation by performing a sequence of actions. Furthermore it can be modelled that after an agent commits itself to achieve a goal it is obliged to perform those actions that bring about its goal.

Finally, in [6, 13] we formalized communication between agents. In this theory we can show both the communication itself as well as the consequences of communication. For instance, if some authorised agent gives orders to another agent to perform a certain action, the latter agent will be obliged to perform the action. Also if an authorised agent asserts a fact to another agent, the latter agent will believe this fact to be true. The examples show that the communication can change both the informational attitudes of agents as well as the behaviour of agents.

In this paper we intend to bring the different fragments of the framework together in one all-embracing formal system. That is, we will define a model for the following concepts: *belief, knowledge, action, wish, goal, decision, intention, commitment, obligation* and *communication*. Following [7, 14] we base this model on dynamic logic [8], which is extended with epistemic, doxastic, temporal and deontic (motivational) operators. The semantics will be based on Kripke structures with a variety of relations imposed on the states. We characterise this integrated framework by pointing out some differences with other work.

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, 10, 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. How do these communication protocols fit into the goal directed behaviour of the agents? Of course the protocol will fit in case the agents have a joint goal to start with, but what happens when an agent gets an unexpected message? In order to react appropriately to such a message the agent has to be able to reason about the intended effect of it and it should be able to figure out a response that is in line with its own goals. This is what we would like to call *deliberate communication*.

Note that we do not argue that agents do not have to deliberate when they follow a protocol. The counterproposal to a proposal message can be made through very intricate reasoning about the possible strategy of the other party! However, no deliberation is needed about what type of message is sent (only an accept, reject or counterproposal is usually possible).

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. 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.

In this paper we will describe the different components and their relationships informally. Even though we will give a precise definition of our language and the models used to interpret this language, we only have room to sketch the actual semantics.

2 The Concepts

The concepts that we formalize can roughly be situated at four different levels: the informational level, the action level, the motivational level and the social level. We will introduce the concepts of each of these levels in the following subsections.

2.1 The Informational Level

At the informational level we consider both knowledge and belief. Many formalizations 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. Both concepts are interpreted in a Kripke-style semantics, where each of the operators is interpreted by a relation between a possible world and a set of possible worlds determining the formulas that the agent knows respectively believes. We demand knowledge to obey an S5 axiomatisation, belief to validate a KD45 axiomatisation, and agents to believe all the things that they know. Of course we are aware that the above axiomatization of knowledge and beliefs can lead to some problems about logical omiscience on the one hand and to problems with conflicting beliefs on the other hand. However, we chose to use the most simple axiomatization here in order not to complicate the semantics too much. It is possible, without altering the other levels, to choose for more complex and/or realistic axiomatizations of knowledge and belief.

2.2 The Action Level

At the action level 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 [9, 29] 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. It is also possible to compose actions into an action expression using operators for the parallel execution of actions "&", the choice between actions "+", the sequential composition of actions ";" and the negation of an action "-". Action expressions do not play a major role in this paper where we want to lay out the basic concepts, but they can be used at any place where actions are used.

The results of actions are modelled using dynamic logic as described by Harel in [8]. We use $[\alpha(i)]\phi$ to indicate that *if* agent i performs the action indicated by α the result will be ϕ. Note that it does not state anything about whether the action will actually be performed. So, it might for instance be used to model a statement like: 'If I jump over 2.5m high I will be the world record holder'.

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. In this paper we say that an agent i has the opportunity to perform α if there is some way for i to perform α. The agent not necessarily has the opportunity to perform α in any possible way! From this observation it follows that:

$$OPP(\alpha_1) \rightarrow OPP(\alpha_1 + \alpha_2)$$

but also that

$$OPP(\alpha_1) \not\rightarrow OPP(\alpha_1 \& \alpha_2)$$

In [9] we also used an *ability* predicate. This predicate was used to indicate that an agent has the inherent possibility to perform a certain action (at a certain place and time). So, whether an agent can perform an action depends on both its inherent ability and the external opportunity to perform the action. In the present paper the abilities of agents do not play a major role. Therefore we left them out of the formalization and assume (for the present) that the abilities of an agent are incorporated in the opportunities for an agent to perform an action.

Besides the OPP operator, which already has a temporal flavor 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. Note that the standard dynamic logic operator "<>" can only be used to indicate the *possible* next action. That is, we can use $< \alpha > true$ to indicate from the present state a next state can be reached by performing α. However, to denote the *actual* next action a new operator is needed! See [5] for a more in depth discussion of this issue. A similar argument holds for the $PREV$ operator for actions. We also define a more traditional temporal logic $NEXT$ 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 when an action $\alpha(i)$ is performed next the formula ϕ is true after the performance of $\alpha(i)$.

For the $NEXT$ and $PREV$ predicates we assume the following relations to hold:

1. $NEXT(\alpha_1 + \alpha_2) = NEXT(\alpha_1) \vee NEXT(\alpha_2)$
2. $NEXT(\alpha_1 \& \alpha_2) = NEXT(\alpha_1) \wedge NEXT(\alpha_2)$
3. $NEXT(\overline{\alpha}) = \neg NEXT(\alpha)$

4. $NEXT(\alpha_1;\alpha_2) = NEXT(\alpha_1) \wedge [\alpha_1]NEXT(\alpha_2)$

1. $PREV(\alpha_1 + \alpha_2) = PREV(\alpha_1) \vee PREV(\alpha_2)$

2. $PREV(\alpha_1 \& \alpha_2) = PREV(\alpha_1) \wedge PREV(\alpha_2)$

3. $PREV(\overline{\alpha}) = \neg PREV(\alpha)$

4. $PREV(\alpha_1;\alpha_2) = PREV(\alpha_2) \wedge Before(\alpha_2)PREV(\alpha_1)$

We will not make the operator $Before(\alpha)$ explicit in this paper due to space and time constraints. However, it can be expressed quite naturally using standard temporal operators.

2.3 The Motivational Level

At the motivational level we consider a variety of concepts, ranging from wishes, goals and decisions to intentions and commitments. The most fundamental of these notions is that of wishes. Formally, wishes are defined as the combination of implicit and explicit wishes, which allows us to avoid all kinds of problems that plague other formalizations of motivational attitudes. A formula ϕ is wished for by an agent i, denoted by $W_i\phi$, iff ϕ is true in all the states that the agent considers desirable, and ϕ is an element of a predefined set of (explicitly wished) formulas.

Goals are not primitive in our framework, but instead defined in terms of wishes. Informally, a wish of agent i constitutes one of i's goals iff i knows the wish not to be brought about yet, but implementable, i.e. i has the opportunity to achieve the goal. In terms of [2] our agents only have achievement goals. To formalize this notion, we first introduce the operator $Achiev$; $Achiev_i\phi$ means that agent i has the opportunity to perform some action which leads to ϕ, i.e.

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

A goal is now formally defined as a wish that does not hold but is achievable:

$$Goal_i\phi \equiv W_i\phi \wedge \neg\phi \wedge Achiev_i\phi$$

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 wish now holds, or since it is no longer certain that it can achieve the goal. This implies in particular that our agents will not indefinitely pursue impossible goals.
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.
Because the goals also depend on a set of explicit wishes we can avoid the problem of having all the consequences of a goal as goal as well.

Intentions are divided in two categories, viz. the intention to perform an action and the intention to bring about a proposition. We define the intention of an agent to perform a certain action as primitive. We relate intentions and goals in two ways. Firstly , the intention to reach a certain state is defined as the goal to reach that state. The second way is through *decisions*. An intention to perform an action is based on the decision to try to bring about a certain proposition. We assume a (total) ordering between the explicit wishes of each agent in each world. On the basis of this ordering the agent can make a decision to try to achieve the goal that has the highest preference. Because the order of the wishes 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). As the result of deciding to do α, denoted by $DEC(i,\alpha)$, the agent has the intention to do α, denoted by $INT_i\alpha$. I.e. $[DEC(i,\alpha)]INT_i\alpha$.
The precondition above is described formally by

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

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. For example, if I intend to jump over 1.5m and I jump over 1.4m it can be said that I tried to fulfil my intention, i.e. the latter action is within the intention of performing the first action. However, if instead of jumping over 1.5m I killed a referee it cannot be said anymore that I performed that action with the intention of jumping over 1.5m.

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

$$NEXT(\alpha_2(i)) \rightarrow (\exists \alpha_2 : INT(\alpha_1(i)) \wedge IMP_i(\alpha_1, \alpha_2))$$

being valid.

The last concept that we consider at the motivational level is that of commitment. This concept is also part of the social level if the commitment is made towards another agent. As the result of i performing a $COMMIT(i, j, \alpha)$ action the formula $O_{ij}\alpha$ becomes true. (See [6] for more details.) I.e. by committing itself to an action, an agent i obliges itself towards j to perform the action α. The commitment can be a private one if j is the same as i. In that case the result is an obligation of the agent towards itself to perform the action. 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 registration of the violation is done through the introduction of a deontic relation between the worlds. This relation connects each world with the set of ideal worlds with respect to that world. More details about the formal semantics of this deontic operator can be found in [7].

Secondly, an obligation to perform an action leads to the goal of having performed the action. Formally this is achieved with the following formula:

$$O_{ij}\alpha \rightarrow W_i(PREV(\alpha(i)))$$

Note that this is sufficient to create a goal, because $PREV(\alpha(i))$ does not hold currently (the action is not performed yet when the obligation arises) and it is achievable (by performing the action $\alpha(i)$. The above formalisation is actually also a shortcut on the more cognitive correct assumption that the agent actually wants to avoid the violation of the obligation. I.e. $W_i(\neg Violation(O_{ij}\alpha))$. Depending on how severe are the consequences of this violation for the agent this wish will get a higher priority.

We only consider sincere agents and therefore we assume that an agent can only commit itself to actions that it intends to do eventually, i.e. intention provides a precondition for commitment.

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. 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.

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 the 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, 11]. 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 built 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 [6]):

- $DIR(x, i, j, \alpha)$ formalises that agent i directs agent j to perform α on the basis of x, where x can be either $peer$, $power$ or $authority$.
- $DECL(i, j, f)$ models the declaration of i in the "presence" of j that f holds.
- $ASS(x, i, j, f)$ formalises the assertion of i to agent j that f holds. on the basis of x, where x can be either $peer$, $power$ or $authority$.
- $COMMIT(i, j, \alpha)$ 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 has to be authorized not by the other party but by some third party.

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 [28] for more details on contracts between agents). Formally, the following formulas hold for the effects of commitments, orders and declaratives:

- $[COMMIT(i, j, \alpha)]K_j([DECL(j, i, P_{ij}(\alpha(i)))]O_{ij}\alpha)$
- $auth(i, DIR(authority, i, j, \alpha)) \rightarrow [DIR(authority, i, j, \alpha)]K_j O_{ji}\alpha$
- $j \ll i \rightarrow [DIR(power, i, j, \alpha)]K_j O_{ji}\alpha$
- $[DIR(x, i, j, \alpha)]K_j INT_i\alpha(j)$
- $auth(i, DECL(i, j, f)) \rightarrow [DECL(i, j, f)]K_j f$
- $[DECL(i, j, f)]K_j W_i f$

- $[ASS(x,i,j,f)]K_jB_if$
- $auth(i,ASS(authority,i,j,f)) \rightarrow [ASS(authority,i,j,f)]B_jf$
- $j \ll i \rightarrow [ASS(power,i,j,f)]B_jf$

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). It is important that j knows that giving this permission results in an obligation for i. It means that j can use this in the plan (and/or goal) formation! Because we assume that $K_i\phi \rightarrow \phi$ it is also true in the world. The giving of permission is formally described by $[DECL(j,i,P_{ij}(\alpha(i)))]$, 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. Also in this case it is important that j knows that it has this obligation!
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 α 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:

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

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! Although this might not be obvious it follows from the fact that j knows f afterwards, which can only be the case if f is actually true).
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. Because j is ”listening” to this declaration it also knows of this preference.

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.

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

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:

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

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!

Before we give a sketch of a formalization of all the above concepts, in the next section we will first illustrate how the communicative acts as defined above can be used to model the basic messages in other frameworks.

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.

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. e.g. accepting a delivery means that you authorize the other party to demand payment)

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 could be expressed as the disjunction of the negation of these three parts. However, usually it is seen as not giving permission for the action. If the permission is not given then the other two parts are of no consequence any more.

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:
  $[PROPOSE(i,j,\alpha)(i)][ACCEPT(j,i,\alpha)(j)][\alpha(i)]auth(i,\beta(i))$
- $[PROPOSE(i,j,\alpha)(i)][REJECT(j,i,\alpha)(j)]\neg O_{ij}(\alpha(i)) \wedge \neg P_{ji}(\alpha(i))$ (reject)
- $[PROPOSE(i,j,\alpha)(i)][COUNTERPROPOSE(j,i,\beta)(j)]$
  $\neg O_{ij}(\alpha(i)) \wedge \neg P_{ji}(\alpha(i)) \wedge OPP(ACCEPT(i,j,\beta)(i) + REJECT(i,j,\beta)(i))$ (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 and there is also no permission to perform the action. 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.

Although in this section we have shown how some messages as used in ADEPT and their effects can be formally described in our framework this is only the start of deliberative communication. The next interesting step that we intend to take is to show why a certain reaction is given to a message from another agent given the situation and the goals of the agent. In principle it is possible to check in the above example whether a proposal should be accepted, rejected or countered given the goals of the agent and the effects of these answers. We leave a detailed exploration of this very interesting area for future work.

4 A Sketch of a Formalization

In this section we precisely define the language that we use to formally represent the concepts described in the previous section, 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, pair-wise 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 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 *The set Act of action expressions is defined to be the smallest set closed under:*

1. $At \cup \{\mathbf{any}, \mathbf{fail}\} \subseteq Act$
2. $\alpha_1, \alpha_2 \in Act \Longrightarrow \alpha_1;\alpha_2, \alpha_1 + \alpha_2, \alpha_1 \& \alpha_2, \overline{\alpha_1} \in Act$

The set $MAct$ of general action expressions contains the regular actions and all of the special meta-actions informally described in the previous section. For these meta-actions it is not always clear whether they can be performed in parallel or what the result is of taking the negation of a meta-action. This area needs a more thorough study in the future. For simplicity, we restrict ourselves in this paper to closing the set $MAct$ under sequential composition.

Definition 2 *The set $MAct$ of general action expressions is defined to be the smallest set closed under:*

1. $Act \subseteq MAct$
2. $\alpha \in Act, i, j \in Ag, x \in \{peer, authority, power\}$
$\Longrightarrow DEC(i, \alpha), COMMIT(i, j, \alpha), DIR(x, i, j, \alpha) \in MAct$
3. $\gamma\alpha_1, \gamma\alpha_2 \in MAct \Longrightarrow \gamma\alpha_1; \gamma\alpha_2 \in MAct$

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 *The language $FORM$ of formulas is defined to be the smallest set closed under:*

1. $PFORM \subseteq FORM$
2. $\phi, \phi_1, \phi_2 \in FORM \Longrightarrow \neg\phi, \phi_1 \wedge \phi_2 \in FORM$
3. $\phi \in FORM, i \in Ag \Longrightarrow K_i\phi, B_i\phi \in FORM$
4. $\gamma\alpha \in MAct, \phi \in FORM \Longrightarrow [\gamma\alpha]\phi \in FORM$
5. $\psi, \phi \in FORM, i, j \in Ag, x \in \{peer, authority, power\} \Longrightarrow$
$[DECL(i, j, \psi)]\phi, [ASS(x, i, j, \psi)]\phi \in FORM$
6. $[\gamma\alpha]\phi, [\gamma\beta]\psi, \theta \in FORM \Rightarrow [\gamma\alpha; \gamma\beta]\theta \in FORM$
7. $\alpha \in Act, \phi \in FORM \Longrightarrow PREV(\alpha), OPP(\alpha), NEXT(\phi) \in FORM$
8. $\phi, \psi \in FORM, i, j \in Ag, \alpha, \alpha_1, \alpha_2 \in Act \Longrightarrow W_i\phi, \psi <_i \phi, INT_i\alpha, i \ll j$
$IMP_i(\alpha_1, \alpha_2), O_{ij}(\alpha), auth(i, \alpha) \in FORM$

Because the assert and declare speech acts range over formulas they can only be introduced at this stage!

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 level, the action level, the motivational level and the social level, 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. *$A = (Sf, Mf, Ropp, Rprev, Rnext)$ with $Sf : Ag \times Act \times \Sigma \to \wp(\Sigma)$ yielding the interpretation of regular actions, $Mf : Ag \times MAct \times (CMo \times \Sigma) \to (CMo \times \Sigma)$ yielding the interpretation of meta-actions, $Ropp : Ag \times \Sigma \to \wp(Act)$ denoting opportunities, $Rprev : \Sigma \to Act$ yielding the action that has been performed last and $Rnext : \Sigma \to Act$ 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 wishes, $Rep : Ag \times \Sigma \to \wp(FORM)$ yielding explicit wishes, $< \subseteq Ag \times \Sigma \to FORM \times FORM$ which is a preference relation on wishes, $Ri : Ag \times \Sigma \to \wp(Act)$ 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 state-transition interpretation for regular actions. This function satisfies the usual constraints ensuring an adequate interpretation of composite actions in terms of their constituents. The function Mf models the model-transforming interpretation of meta-action. Below we elaborate on the definition of Mf for the meta-actions introduced in the previous section.*
3. *$Rnext(s) \in Ropp(i, s) \subseteq \{\alpha \mid Sf(i, \alpha, s) \neq \emptyset\}$, 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, $Rprev(s) = \alpha$ 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 [16]. 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 formalize 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:*

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

The functions interpreting the special meta-actions are described below in terms of the preconditions and the postconditions for execution of the actions. The precondition describes on which models the model-transforming function has the desired effect and the postcondition describes the model yielded by the application of the meta-action.

DEC The precondition for execution of $DEC(i,\alpha)$ is that for some $\phi \in FORM$, $Goal_i\phi \wedge [\alpha(i);\beta(i)]\phi$ holds, for some $\alpha(i),\beta(i) \in Act$ and furthermore no ψ exists such that $P_i\psi$ and $\phi < \psi$ hold. Thus agents may only decide to intend to do those actions that fulfil some most preferred goal. As the result of execution of $DEC(i,\alpha)$ the model is changed in such a way that $INT_i\alpha$ holds in the resulting model.

$COMMIT$ There are no preconditions for execution of $COMMIT(i,j,\alpha)$ by agent i. The effect of the commitment is that the model is changed in such a way that $O_{ij}(\alpha)$ holds afterwards.

DIR The preconditions for execution of $DIR(authority,i,j,\alpha)$ by i are given by $auth(i,DIR(i,j,\alpha))$. This implies that agent i should have the authority over j before it can order it around. The effect of such an action is that j is obligated to i to perform α, which is implemented in a way similar to the implementation of the $COMMIT$ action. The same holds for a directive on the basis of power.

$DECL$ The action $DECL(i,j,f)$ has as precondition that i is authorised to declare f. (Only some civil servants can declare people to be married in the Netherlands). Execution of an action $DECL(i,j,f)$ in a certain state of a model will be a modification of the valuation π such that f is true in all the resulting states of the resulting models. Note that whenever f is inconsistent no model results.

ASS The precondition for $ASS(x,i,j,f)$ is that i is demanded to believe f, i.e. B_if should hold. This implies in particular that agents are not allowed to gossip, i.e. spread around rumors that they themselves do not even believe. As the result of executing $ASS(authority,i,j,f)$ by i in some state s, two cases arise. If $auth(i,ASS(authority,i,j,f))$ holds then the model under consideration is modified such that $Rb(j,s)$ contains only states in which f is true, which indeed implies that B_jf holds in s in the resulting model. Regardless whether i is authorised to make the assertion the model under consideration is modified such that $Rk(j,s)$ contains only states s' such that $Rb(i,s')$ contains only states in which f is true. This implies that K_jB_if holds in the resulting model.

5 Related Approaches

In this section we very briefly indicate the main differences between our approach and three other approaches to model rational agents, viz. the framework proposed by Cohen & Levesque [2], the BDI-framework of Rao & Georgeff [22] and the theoretical framework for multi-agent systems of M. Singh [25]. After that we will shortly compare our approach to some other work on communicating agents.

The main difference between our approach and the one of Cohen & Levesque is that they define intentions in terms of goals and beliefs. We agree with this approach when it concerns intentions on propositions. However, we do not take a goal to be a primitive notion as is the case in the approach of Cohen & Levesque. Because they take a goal to be a primitive notion they have to define different types of goals in order to define the persistence of a goal, the achievability of a goal, etc. All these properties are direct consequences of our definition of a goal in terms of preferences and achievabilities. The distinction between goals and preferences allows for a bigger flexibility than possible in the approach of Cohen & Levesque. Furthermore, whereas we define the intention to perform an action as primitive, Cohen & Levesque define the intention to perform an action as the goal to reach a state where that action has been performed. Although both types of intentions of Cohen & Levesque are based on the notions of goals and beliefs it is not clear what is the relation between the intention to reach a certain state and the intention to perform an action; in fact these notions seem to be unrelated. However, it seems desirable that the intention to reach a certain goal induces the intention to perform an action which is needed to reach that goal. In our approach this relation is established through the notion of decisions. A goal can induce a decision. The decision then induces the intention to perform an action. Another relation that remains unclear in the theory of Cohen & Levesque is that between an intended action and the action that is actually performed. The only relation they give is that the intended action should be the same as the action whose goal it is to be performed. Which in itself does not mean anything for the actual course of events. In our approach we introduce the notion of an implementation relation between actions, which introduces a loose coupling between intended actions and the actions that are actually performed.

The last point also shows one of the main differences between our framework and that of Rao & Georgeff. In their framework it holds that if an agent intends to perform an action it will also actually perform the action. They weaken this assumption in [12], where they investigate different strategies with respect to commitment to a goal. An agent can be single minded when it always performs its intended actions and open minded when it discards its goals on the basis of new information it receives. This work, however, is of a very practical nature and does not have a theoretical counterpart. Therefore it remains unclear how this would be reflected in the BDI framework.

In the BDI framework they avoided making the intention operator into a temporal operator by introducing the notion of a successful performance of an action and a failed performance of an action. However, the relation between a successful performed event and an event that failed to be performed is unclear. Can this be any other event? Can it include the event itself? At present the best one can say if an event has been performed (either successful or failed) is that *some* event has been performed.

A last, rather important, point of difference between our framework and the other two is the fact that we also include the social level, which we consider essential, but is only briefly mentioned by Cohen & Levesque, and not considered at all by Rao & Georgeff.

The social level is treated in the theory of Singh, in the form of communication between agents. The framework of Singh is based on CTL^*, a branching-time logic. The main difference with our framework is that he does not incorporate deontic relations between agents. It means that the success of a speech act is not defined in the same way as we do. A directive is successful if it results in the hearer intending to perform the action that it was ordered to perform. This is much stronger than in our framework. Of course, an obligation can lead to a goal, but only if the hearer places the directive above the plan it is currently performing.

How does our work compare with the other current work on communicating agents?
Bretier & Sadek [1] formally define the effects of communicative actions. However, this theory is limited to effects on the beliefs and intentions of agents. We have shown that the concept of obligation also forms an important ingredient of the effects of some speech acts. That is, directives and commissives always result in obligations. It must be remarked that the application for which the theory in [1] is used does not need these illocutions.

Another difference between our work and that of [1] is that our theory is geared to communication between agents while their theory is geared towards the management of human-agent dialogue. The same holds for the theory of Traum [26]. Although many aspects are compatible, there are also some differences. One of the main differences is that in the communication between humans and agents the detection of the illocution of messages is not trivial and important to steer the dialogue in a natural way. Agents will use a formal language in which the illocution of each message is clear.

Very interesting in the work of Traum is that he also recognizes the importance of obligations between the participants in a dialogue. Besides mutual knowledge and believe this is an important relation, because it indicates the expectations that the participants have of each other.

Finally the work of Noriega & Sierra [20] comes very close in spirit to our work. They also try to give one formal framework for the different components of communicating agents. Their theory is more flexible than ours in that they allow different agents to use different languages and even the different components might use different inference rules. The relation between the components is given by so-called bridge rules. We assume a uniform language for all agents and components. This provides for an automatic integration of the components. The main point in which the theory in [20] is lacking compared to ours is the semantics of the speech acts themselves and their effects.

6 Conclusions

In this paper we presented an informal overview and a sketchy formalization of the concepts that we consider essential to model rational agents. In our very flexible and highly expressive framework we propose a variety of concepts, which are roughly situated at four different levels: the informational level, where knowledge and belief are considered, the action level, where we consider various aspects of action, the motivational level, where we dealt with preferences, goals, intentions, etc., and the social level, which is concerned with the social traffic between agents.

The resulting multi-modal logic is quite complex. However, we want to make two remarks 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. It is very important that this is realized when the communication protocols are automatized. (As is the aim in communication between agents). If the communication is automatic it becomes very important to know the exact effects of the messages. What is the knowledge of each agent and what are its obligations (resp. expectations).

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!

In subsequent work we want to show how communication protocols that are used in more practical work like [19, 21] can be given a formal semantics in our framework. For the basic illocutions used in the ADEPT system [21] this has been done in this paper and has already led to the discovery that the seemingly simple ACCEPT message has unexpected results.

Also we want to define an agent architecture for communicating agents that adhere to our theory. Some groundwork in this respect has been done in [27]. What needs to be done in this respect is to show how communicative actions can form part of plans to reach a goal of the agent and also how goals of an agent can be adjusted under the influence of communication from other agents. We believe both parts can be done using the formalism described in this paper.

References

[1] P. Bretier and M. D. Sadek. A rational agent as the kernel of a cooperative spoken dialogue system: Implementing a logical theory of interaction. In J. P. Müller, M. J. Wooldridge, and N. R. Jennings, editors, *Intelligent Agents III — Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages (ATAL-96)*, Lecture Notes in Artificial Intelligence 1193, pages 189-204, Springer-Verlag, Heidelberg, 1996.

[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. Using Transactions in Integrity Constraints: Looking forward or backwards, what is the difference? In *First International Workshop on Applied Logic: Logic at Work*, UVA, Amsterdam, 1992.

[6] F. Dignum and H. Weigand. Modelling communication between cooperative systems In J. Iivari et al., *Advanced information systems engineering,* pages 140-153, Springer-Verlag, Heidelberg, 1995.

[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. Brown and J. Carmo (eds.),*Deontic logic, agency and normative systems,* Workshops in Computing, Springer-Verlag, pages 80-97, 1996.

[8] D. Harel. First Order Dynamic Logic. LNCS 68 Springer-Verlag, Heidelberg, 1979.

[9] 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, Springer-Verlag, Heidelberg, pages 366-378, 1994.

[10] N. Jennings. Commitments and Conventions: The foundation of coordination in Multi-Agent systems. *Knowledge Engineering Review*, vol. 8(3), pages 223-250, 1993.

[11] 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-A gent Technology*, pages 345-360, The practical application company, London, 1996.

[12] D. Kinny and M. Georgeff. Commitment and Effectiveness of Situated Agents. In *Proceedings International Joint Conference on Artificial Intelligence*, Sydney, Australia, pages 82-88, Morgan Kaufman Publishers, 1991.

[13] B. van Linder, W. van der Hoek and J.-J.Ch. Meyer. Communicating rational agents. In Nebel and Dreschler-Fisher, eds, *Proceedings of KI'94,* LNCS 861, pages 202–213, Springer-Verlag, Heidelberg, 1994.

[14] B. van Linder, W. van der Hoek and J.-J.Ch. Meyer. How to motivate your agents. On making promises that you can keep. In Wooldridge, Müller and Tambe, eds, *Intelligent Agents II,* LNCS 1037, pages 17–32, Springer-Verlag, Heidelberg, 1996.

[15] B. van Linder *Modal Logics for Rational Agents,* PhD Thesis, Utrecht University, 1996.

[16] J.-J.Ch. Meyer. A different approach to deontic logic. In *Notre Dame Journal of Formal Logic,* vol.29, pages 109–136, 1988.

[17] J.-J.Ch. Meyer and W. van der Hoek. *Epistemic Logic for AI and computer science,* Cambridge University Press, Cambridge, 1995.

[18] R. Moore. A formal theory of knowledge and action. In J. Hobbs and R. Moore, eds, *Formal theories of the commonsense world,* pages 319-358, Ablex Publ. Comp., 1985.

[19] J. P. Müller. A cooperation model for autonomous agents. In J. P. Müller, M. J. Wooldridge, and N. R. Jennings, editors, *Intelligent Agents III — Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages (ATAL-96),* Lecture Notes in Artificial Intelligence 1193, pages 245-260. Springer-Verlag, Heidelberg, 1996.

[20] P. Noriega and C. Sierra. Towards layered dialogical agents. In J. P. Müller, M. J. Wooldridge, and N. R. Jennings, editors, *Intelligent Agents III — Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages (ATAL-96),* Lecture Notes in Artificial Intelligence 1193, pages 173-188. Springer-Verlag, Heidelberg, 1996.

[21] T. J. Norman, N. R. Jennings, P. Faratin, and E. H. Mamdani. Designing and implementing a multi-agent architecture for business process management. In J. P. Müller, M. J. Wooldridge, and N. R. Jennings, editors, *Intelligent Agents III — Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages (ATAL-96)*, Lecture Notes in Artificial Intelligence 1193, pages 261-276. Springer-Verlag, Heidelberg, 1996.

[22] A.S. Rao and M.P. Georgeff. Modeling rational agents within a BDI-architecture. In J. Allen et al., eds, *Proceedings of KR'91*, pages 473-484, Morgan Kaufman Publishers, 1991.

[23] G. Sandu. Reasoning about collective goals. In J. P. Müller, M. J. Wooldridge, and N. R. Jennings, editors, *Intelligent Agents III — Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages (ATAL-96)*, Lecture Notes in Artificial Intelligence 1193, pages 131-140. Springer-Verlag, Heidelberg, 1996.

[24] J.R. Searle. Speech Acts. Cambridge University Press, Cambridge, 1969.

[25] M. Singh. Multiagent Systems. LNAI 799, Springer-Verlag, Heidelberg,1994.

[26] D. R. Traum. A reactive-deliberative model of dialogue agency. In J. P. Müller, M. J. Wooldridge, and N. R. Jennings, editors, *Intelligent Agents III — Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages (ATAL-96)*, Lecture Notes in Artificial Intelligence 1193, pages 157-172. Springer-Verlag, Heidelberg, 1996.

[27] E. Verhagen, F. Dignum and H. Weigand. A Language/Action Perspective on Cooperative Information Agents. In F. Dignum at al., eds, *Communication Modeling - The Language/Action Perspective (LAP-96)*, electronic Workshops in Computing, Springer-Verlag, London, 1996.

[28] 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 Infor mation Systems Engineering (LNCS 1080)*, pages 193-209, Springer-Verlag, 1996.

[29] R. Wieringa, J.-J.Ch. Meyer and H. Weigand. Specifying dynamic and deontic integrity constraints. *Data & knowledge engineering*, vol.4, pages 157-189, 1989.