

Organizations as Socially Constructed Agents in the Agent Oriented Paradigm

Guido Boella¹ and Leendert van der Torre²

¹Dipartimento di Informatica - Università di Torino - Italy

²CWI - Amsterdam - The Netherlands

Abstract. In this paper we propose a new role for the agent metaphor in the definition of the organizational structure of multiagent systems. The agent metaphor is extended to consider as agents also social entities like organizations, groups and normative systems, so that mental attitudes can be attributed to them - beliefs, desires and goals - and also an autonomous and proactive behavior. We show how the metaphor can be applied also to structure organizations in functional areas and roles, which are described as agents too. Thus, the agent metaphor can play a role similar to the object oriented metaphor which allows structuring objects in component objects. Finally, we discuss how the agent metaphor addresses the problems of control and communication in such structured organizations.

1 Introduction

The role of software engineering is to provide models and techniques that make it easier to handle the complexity arising from the large number of interactions in a software system [12]. Models and techniques allow expressing knowledge and to support the analysis and reasoning about a system to be developed. As the context and needs of software change, advances are needed to respond to changes. For example, today's systems and their environments are more varied and dynamic, and accommodate more local freedom and initiative [23].

For these reasons, agent orientation emerged as a new paradigm for designing and constructing software systems [12, 23]. The agent oriented approach advocates decomposing problems in terms of autonomous agents that can engage in flexible, high-level interactions. In particular, this is a natural representation for complex systems that are - as many real systems are - invariably distributed [12]. Compared to the still dominant software paradigm, namely object orientation, agent orientation offers a higher level of abstraction for thinking about the characteristics and behaviors of software systems. It can be seen as part of an ongoing trend towards greater interactivity in conceptions of programming and software system design and construction. Much like the concepts of activity and object that have played pivotal roles in earlier modelling paradigms - Yu [23] argues - the agent concept can be instrumental in bringing about a shift to a much richer, socially-oriented ontology that is needed to characterize and analyze today's systems and environments.

The shift from the object oriented perspective to the agent oriented one is not, however, without losses. Booch [7] identifies three tools which allow coping with complexity: “1) Decomposition: the most basic technique for tackling any large problem is to divide it into smaller, more manageable chunks each of which can then be dealt with in relative isolation. 2) Abstraction: the process of defining a simplified model of the system that emphasises some of the details or properties. 3) Organisation: the process of identifying and managing interrelationships between various problem solving components.”

In the agent oriented approach, however, decomposition, abstraction and organization are not yet addressed with the same efficacy as in the object oriented approach, where an object can be composed of other objects, which can be ignored in the analysis at a certain level of abstraction. The agent metaphor is sometimes proposed as a specialization of the object metaphor [1]: agents do not only have - like objects - a behavior which can be invoked by the other agents, but they also autonomously act and react to changes in the environment following their own goals and beliefs. In contrast, the component view of objects in the object metaphor could to be lost. The property of agents, i.e., sociality, closest to the property allowing the aggregation of objects to form more complex objects is not enough to overcome the gap. In particular, multiagent systems offer as aggregation methods the notion of group or of organization. According to Zambonelli *et al.* [24] “a multiagent system can be conceived in terms of an organized society of individuals in which each agent plays specific roles and interacts with other agents”. At the same time, they claim that “an organization is more than simply a collection of roles (as most methodologies assume) [...] further organization-oriented abstractions need to be devised and placed in the context of a methodology [...] As soon as the complexity increases, modularity and encapsulation principles suggest dividing the system into different suborganizations”. According to Jennings [12], however, most current approaches “possess insufficient mechanisms for dealing with organisational structure”. Moreover, what is the semantic principle which allows decomposing organizations into suborganizations must be still made precise.

The research question of this paper, thus, is: how can the agent oriented paradigm be extended with a decomposition structure isomorphic to the one proposed by the object oriented paradigm? How can a multiagent system be designed and constructed as an organization using this structure?

The methodology we use in this paper is a normative multiagent framework we proposed in [2, 3, 6, 5]. The basic idea of this framework is: agents attribute mental attitudes, like beliefs, desires and goals, to the other agents they interact with and also to social entities like groups, normative systems, and organizations. Thus these social entities can be described as agents too, and at the same time, the components of organizations, namely, functional areas and roles, can be described as agents, as in the ontology we present in [3]. We call them *socially constructed agents*.

This paper is organized as follows. In Section 2 we discuss the progress from object orientation to agents and socially constructed agents. In Section 3 we present the formal model and in Section 4 we discuss the issue of control and communication in an multiagent system structured as an organization. A summary closes the paper.

2 From objects to socially constructed agents

The trend in software and requirements engineering and in programming languages paradigms has been from elements that represent abstract computations towards elements that represent the real world: from procedural to structured programming, from objects to agents. Agent systems have no central control authority, instead each agent is an independent locus of control, and the agent's task drives the control. Delegating control to autonomous components can be considered as an additional dimension of modularity and encapsulation. Intentional concepts such as goals, beliefs, abilities, commitments, *etc.*, provide a higher-level characterization of behavior. One can characterize an agent in terms of its intentional properties without having to know its specific actions in terms of processes and steps. Explicit representation of goals allows motivations and rationales to be expressed. The agent concept provides a local scope, for reconciling and making tradeoffs among competing intentionality, such as conflicting goals and inconsistent beliefs. By adopting intentional modelling, the networks of dependencies among the agents can be modelled and reasoned about at a high level of abstraction. Moreover, cooperation among agents cannot be taken for granted. Because agents are autonomous, the likelihood of successful cooperation is contingent upon many factors. However, an agent that exists within a social network of expectations and obligations has behaviors that are confined by them. The agent can still violate them, but will suffer the consequences. The behavior of a socially situated agent is therefore largely predictable, although not in a precise way.

Given that agents are nowadays conceived as useful abstractions for modelling and engineering large complex systems, the need for a disciplined organizational principle for agent systems emerges clearly in the same way as the formalization of the object decomposition principle does in the case of object oriented systems.

One of the main features of the object perspective is that objects are composed by other objects and that objects can be replaced by other objects with the same properties (e.g., the same interface). This is not entirely true for agents. According to Jennings [12], "the agent oriented approach advocates decomposing problems in terms of autonomous agents", but no further decomposition seems possible. To overcome this flatness limitation, the organization metaphor has been proposed, e.g., by [10, 24]. Organizations are modelled as collections of agents, gathered in groups [10], playing roles [12, 17] or regulated by organizational rules [24]. What is lacking is a notion of organization as a first class abstraction which allows decomposing into subproblems the problem which a system wants to solve, using a recursive mechanism (as the object decomposition is) until autonomous agents composing a multiagent system are reached.

The desired solution is required to model at least simple examples taken from organizational theory in Economics as the following one. Consider a simple enterprise which is composed by a direction area and a production area. The direction area is composed by the CEO and the board. The board is composed by a set of administrators. The production area is composed by two production units; each production unit by a set of workers. The direction area, the board, the production area and the production units are *functional areas*. In particular, the direction area and the production areas belong to the organization, the board to the direction area, *etc.* The CEO, the administrators and

the members of the production units are *roles*, each one belonging to a functional area, e.g., the CEO is part of the direction area.

This recursive decomposition terminates with roles: roles, unlike organizations and functional areas, are not composed by further social entities. Rather, roles are played by other agents, real agents (human or software) who have to act as expected by their role.

The object metaphor is not adequate to deal with such a structure, because each entity can be better described in terms of belief, desires and goals, and of its autonomous behavior. We talk, e.g., about the decisions of the CEO, or about the organization's goal to propose a deal, about the belief of the production area that the inventory is finished, *etc.* Hence, at first sight, these entities can be described as autonomous agents. But this is not sufficient, since the agent metaphor does not account for the decomposition structure of an organization relating it with its functional areas and roles. Moreover, organizations, functional areas and roles do not exist in the same sense as (human or software) agents do. Thus, if we want to follow this intuition, the agent metaphor must be extended. Inspired by Searle [21]'s analysis of social reality we define organizations, functional areas and roles as *socially constructed agents*. These agents do not exist in the usual sense of the term, but they are abstractions which other agents describe as if they were agents, with their own beliefs, desires and goals, and with their own autonomous behavior. The argument goes as follows:

1. agents can attribute to other (human or software) agents mental attitudes and an autonomous behavior to explain how they work, regardless of the fact that they really have any mental attitudes (the *intentional stance* of Dennett [8]);
2. according to Searle [21], agents create new social entities like institutions - e.g., money and private property - by means of collectively attributing to existing entities - e.g., paper bills - a new functional status - e.g., money - and new qualities.
3. if the new functional status is composed by mental attitudes and autonomous behavior, the new entities are described as agents: *socially constructed agents*.
4. hence, socially constructed agents, *qua* agents, can create new socially constructed agents by attributing mental attitudes to them, in turn.

Agents create organizations by collectively attributing them mental attitudes; organizations, as socially constructed agents, can create new social entities like functional areas and roles which are the components of the organization. Functional areas, as agents, can in turn apply the agent metaphor to create subareas and further roles, and so on. Roles are descriptions of the behavior which is expected by agents who, with their own mental attitudes, play these roles: the role's expected behavior is described in terms of mental attitudes, since roles are considered socially constructed agents. Modelling roles by attributing them mental attitudes allows a more expressive way to describe the expected behavior with respect, e.g., the scripts proposed by Activity Theory [20]. In this manner, we have a way to structure an organization in components with an homogeneous character - since they are all agents - in the same way as the object orientation allows structuring objects by means of objects. An advantage of this way of structuring an organization is that its components can be described as agents with beliefs, desires and goals. Hence, the same decomposition approach advocated by [12] is used for structuring an organization: it is decomposed in a set of autonomous agents: not only real ones, but socially constructed agents like functional areas and roles; socially

constructed agents do not exist, but they are only used as abstractions in the design analysis to structure an organization. At the end of the process there are only human or software agents which, to coordinate their behavior, behave as if they all attribute the same beliefs, desires and goals to the organization. This is a subjective approach to coordination [20].

Another reason why organizations, functional areas and roles should be all considered as agents - and not simply groups - is that they have private properties and agents who are employed in them; so a department can possess a building and machines, employ people, *etc.* Moreover they are the addressees of obligations (e.g., to pay the employees), permissions (e.g., a role can use a certain machine) and powers (e.g., the role of CEO can take decisions). This is what is also meant by the law when such social entities are defined as “legal persons”: they are considered persons with obligations and rights [19]. Finally, organizations and functional areas, as legal institutions, are normative agents themselves: they are agents who can pose obligations on the roles and on the employees, e.g., by giving orders to them, or endow them with permissions and powers.

There is a difference with the decompositional view of the object oriented perspective which must be noticed. The parts of an object exist by themselves and the object itself exists only as long as its (essential) parts exist. In contrast, in an organization the perspective is reversed: the “components” of the organization exist only as long as the organization exists, while the organization itself can exist even without its components. The role of CEO does not have sense if the organization which the role belongs to does not exist anymore. The reason is that an organization as a social entity has no physical realization. The organization exists because of the attribution of mental attitudes by the agents of a society. In turn, functional areas and roles exist only as long as the organization attributes mental attitudes to them. An important consequence of this view is that an organization can restructure itself while continuing to exist.

As [9, 10] claim, a multiagent system should not make any assumption about the implementation of the agents. As Yu [23] notices, the agent perspective does not mean necessary that entities should be implemented with mental attitudes:

Agent intentionality is externally attributed by the modeller. From a modelling point of view, intentionality may be attributed to some entity if the modeller feels that the intentional characterization offers a useful way for describing and analyzing that entity. For example, some entity that is treated as an agent during modelling may end up being implemented in software that has no explicit representation and manipulation of goals, *etc.*

Socially constructed agents defined in terms of beliefs, desires and goals are only an abstraction for designing the system. Moreover, the behavior of roles is described by mental attitudes, but this does not require that the agents playing roles in the organizations are endowed with beliefs and motivations: it is sufficient that their behavior conforms to that of the role they are playing.

In Figure 1, we summarize the approach: the multiagent system in the oval is composed of three real agents (boxes) who collectively attribute beliefs (B), desires (D) and goals (G) to the organization (parallelogram). The organization, in turn, attributes mental attitudes to two functional areas and functional areas to three roles. The organization and the functional areas are attributed also norms (V), facts (f), institutional facts (i) and decisions (the triangle d).




Fig. 1. The attribution of mental attitudes.

3 The conceptual model

We introduce the conceptual model necessary to cope with socially constructed agents: first the multiagent system with the attribution of mental attitudes to agents, then the normative system.

First of all, the structural concepts and their relations. We describe the different aspects of the world and the relationships among them by introducing a set of propositional variables X and extending it to consider also negative states of affairs: $L(X) = X \cup \{\neg x \mid x \in X\}$. The relations between the propositional variables are given by means of conditional rules written as $R(X) = 2^{L(X)} \times L(X)$: the set of pairs of a set of literals built from X and a literal built from X , written as $l_1 \wedge \dots \wedge l_n \rightarrow l$ or, when $n = 0$, $\top \rightarrow l$. The rules are used to represent the relations among propositional variables existing in beliefs, desires and goal of the agents.

Then there are the different sorts of agents A we consider. Besides real agents RA (either human or software) we consider as agents in the model also socially constructed agents like organizations OA , functional areas FA , and roles RO . The different sorts of agents are disjoint and are all subsets of the set of agents A : $RA \cup OA \cup FA \cup RO \subseteq A$. All these agents have mental attitudes; by mental attitudes we mean beliefs B , desires D and goals G .

Mental attitudes are represented by rules, even if they do not coincide with them: $MD : B \cup D \cup G \rightarrow R(X)$. When there is no risk of confusion we abuse the notation by identifying rules and mental states. To resolve conflicts among motivations we introduce a priority relation by means of $\geq : A \rightarrow 2^M \times 2^M$ a function from agents to a transitive and reflexive relation on the powerset of the motivations $M = D \cup G$ containing at least the subset relation. We write \geq_a for $\geq (a)$. Moreover, different mental attitudes are attributed to all the different sorts of agents by the agent description relation $AD : A \rightarrow 2^{B \cup D \cup G \cup A}$. We write $B_a = AD(a) \cap B$, $A_a = AD(a) \cap A$ for $a \in A$, etc.

Also agents are in the target of the agent description AD relation for the following reason: organizations, functional areas and roles exist only as profiles attributed by other agents. So they exist only as they are described as agents by other agents, according to the agent description relation. The AD relation specifies that an agent $b \in OA \cup FA \cup RO$ exists only as far as some other agents $\{a \in A \mid b \in A_a\}$ attribute to it mental attitudes. The set $(FA \cup RO) \cap A_o$ represents the immediate “components” of the organization or functional area $o \in OA \cup FA$. The decomposition structure of an organization ends with roles. Roles are described as agents, but they do not create further socially constructed agents; rather, roles are associated with agents playing them, $PL : RO \rightarrow RA$.

We introduce now concepts concerning informational aspects. First of all, the set of variables whose truth value is determined by an agent (decision variables) [14] are distinguished from those P which are not (the parameters). Besides, we need to represent also the so called “institutional facts” I . They are states of affairs which exist only inside normative systems and organizations: as Searle [21] suggests, money, private property, marriages, etc. exist only as part of social reality; since we model social reality by means of the attribution of mental attitudes to social entities, institutional facts can be modelled as the beliefs attributed to these agents, as done by [6]. Similarly, we need to represent the fact that social entities like normative systems and organizations are able to change their mental attitudes. The actions determining the changes are called creation actions C . Finally, inspired by Lee [15] we introduce the notion of documents DC : “we use the term ‘document’ since most information parcels in business practice are mapped on paper documents”.

As concerns the relations among these concepts, we have that parameters P are a subset of the propositional variables X . The complement of X and P represents the decision variables controlled by the different agents. Hence we associate with each agent a subset of $X \setminus P$ by extending again the agent description relation $AD : A \rightarrow 2^{B \cup D \cup G \cup A \cup (X \setminus P)}$. We write $X_a = AD(a) \cap X$.

Moreover, the institutional facts I are a subset of the parameters P : $I \subseteq P$. When a belief rule $Y \wedge c \rightarrow p \in B_a$ has an institutional fact $p \in I$ as consequent, we say that $c \in X$ counts as p in context Y - using Searle [21]’s terminology - for agent $a \in OA \cup FA \cup RO$.

The creation actions C are a subset of the institutional facts $C \subset I$. Since agents are attributed mental attitudes, we represent their modification by adding new mental attitudes expressed as rules. So the creation action relation $CR : \{b, d, g\} \times A \times R(X) \rightarrow C$ is a mapping from rules (for beliefs, desires and goals) to propositional variables, where $CR(b, a, r)$ stands for the creation of $m \in B_a$, $CR(d, a, r)$ stands for the creation of $m \in D_a$, and $CR(g, a, r)$ stands for the creation of $m \in G_a$, such that the mental attitude m is described by the rule $r \in R(X): r = MD(m)$.

Finally, the document creation relation $CD : DC \rightarrow X$ is a mapping from documents to decision variables representing their creation. We write $CD(d) \in X_a$ for the creation of document $d \in DC$.

We define a multiagent system as $MAS = \langle RA, OA, FA, RO, X, P, B, D, G, AD, MD, \geq, I, C, DC \rangle$.

We introduce obligations posed by organizations and functional areas by means of a normative multiagent system. Let the norms $\{n_1, \dots, n_m\} = N$ be a set. Let the norm description $V : OA \cup FA \rightarrow (N \times A \rightarrow X)$ be a function from agents to complete functions from the norms and agents to the decision variables: we write V_o for the function $V(o)$ and $V_o(n, a)$ for the decision variable of agent $o \in RA \cup OA \cup FA$ representing that it considers a violation of norm n by agent $a \in A$.

$NMAS = \langle RA, OA, FA, RO, X, P, D, G, AD, MD, PL, \geq, I, C, DC, N, V \rangle$ is a normative multiagent system .

Following [2], obligations are defined in terms of goals of the addressee of the norm \mathbf{a} and of the agent \mathbf{o} . The definition of obligation contains several clauses. The first one defines obligations of agents as goals of the normative agent, following the ‘Your wish is my command’ strategy, the remaining ones are instrumental to the respect of the obligation.

Agent $\mathbf{a} \in A$ is *obliged* by normative agent $\mathbf{o} \in OA \cup FA$ to decide to do $x \in L(X_{\mathbf{a}} \cup P)$ with sanction $s \in L(X_{\mathbf{o}} \cup P)$ if $Y \subseteq L(X_{\mathbf{a}} \cup P)$ in $NMAS$, written as $NMAS \models O_{\mathbf{a}\mathbf{o}}(x, s|Y)$, if and only if there is a $n \in N$ such that:

1. $Y \rightarrow x \in D_{\mathbf{o}} \cap G_{\mathbf{o}}$: if agent \mathbf{o} believes Y then it desires and has as a goal that x .
2. $Y \cup \{\sim x\} \rightarrow V_{\mathbf{o}}(n, \mathbf{a}) \in D_{\mathbf{o}} \cap G_{\mathbf{o}}$: if agent \mathbf{o} believes Y and $\sim x$, then it has the goal and the desire $V_{\mathbf{o}}(n, \mathbf{a})$: to recognize it as a violation by agent \mathbf{a} .
3. $Y \cup \{V_{\mathbf{o}}(n, \mathbf{a})\} \rightarrow s \in D_{\mathbf{o}} \cap G_{\mathbf{o}}$: if agent \mathbf{o} believes Y and decides $V_{\mathbf{o}}(n, \mathbf{a})$, then it desires and has as a goal that it sanctions agent \mathbf{a} .
4. $\top \rightarrow \sim s \in D_{\mathbf{a}}$: agent \mathbf{a} desires $\sim s$, which expresses that it does not like to be sanctioned.

Since obligations are defined in terms of mental states, they can be created by means of the creation actions C introducing new desires and goals, as shown by [6]. In this paper, we will use the shorthand $CR(\mathbf{o}, O_{\mathbf{a}\mathbf{o}}(x, s|Y))$ to represent the set of creation actions necessary to create an obligation $O_{\mathbf{a}\mathbf{o}}(x, s|Y)$.

4 Control and communication in organizations

Instead of having a single global collection of beliefs and motivations, modelling organizations as socially constructed agents allows allocating different beliefs B_a , desires D_a and goals G_a to separate agents $a \in A_o$ composing the organization $o \in OA$. Agents can be thought of as a locality for intentionality. In this way it is possible to distribute subgoals of G_o among the different functional areas and roles $a \in A_o$ to decompose problems in a hierarchical way and to avoid to overburden them with too much goals. In particular, the goals G_r attributed to role $r \in RO$ represent the responsibilities which agent $b \in A$ playing that roles ($PL(r) = b$) has to fulfill.

The beliefs attributed to the organization (B_o) and attributed by the organization to its components (B_m and $m \in A_o$) represent their know how and the procedures used to achieve the goals of the organization; these beliefs are represented for example by statutes and manuals of organizations. As in case of goals, different beliefs B_a can be distributed to functional areas and roles $a \in A_o$. In this way the organization can respect the incapsulation principle and preserve security and privacy of information, as requested by [10].

The beliefs, desires and goals of the components of an organization play also another role. They express the institutional relations among the different components: in particular, the control and communication relations among the functional areas and roles. Both issues will be addressed using the notion of *document*. Documents are the way information parcels are represented in organizations and represent also the records of decisions and information flow.

The institutional relations of control and communication among the components of an organization are defined in terms of the “counts as” relation. For Jones and Sergot [13], the “counts as” relation expresses the fact that a state of affairs or an action of an agent “is a sufficient condition to guarantee that the institution creates some (usually normative) state of affairs”. As [13] suggest this relation can be considered as “constraints of (operative in) [an] institution”. In Section 3 we propose to model “counts as” relations by means of belief rules of the socially constructed agents. They express how an organization, a functional area or a role provide an institutional classification of reality.

In an organization it is fundamental to specify how agents can control other agents by giving orders to them [10, 24]; the control is achieved by the command structure of an organization. In fact, organizations can be seen as bureaucracies according to [18]. Control has two dimensions: how the organization and its functional areas can pose obligations (commands) to roles, and who has the power to create these obligations (since, as organizations and their units are socially constructed agents, they do not act). For example, a production unit can decide to give a production order to its members and the decision of the production unit can be taken by a director of that unit. The basic block of control is the creation of obligations. As described in the conceptual model, an agent can change its own mental attitudes. In particular, an organization o can change its desires and goals so to create a new obligation $O_{ao}(x, s | Y)$ by means of the creation action $CR(o, O_{ao}(x, s | Y))$. It is possible to create sanction-based obligations addressed to agent $a \in A$ since the agents involved in organizations are

depended on them, for example, for the fact that organizations pay them salaries and decide benefits.

The creation actions C of an organization \mathbf{o} are parameters, hence they are not directly controlled by it: the organization does not act directly, but only by means of the actions of the agents composing it. Creation actions achieve their effect to introduce new obligations if some other action “counts as” a creation action for the organization: this relation is expressed by a belief rule of the organization \mathbf{o} , e.g., $c \rightarrow CR(\mathbf{o}, O_{ao}(x, s | Y)) \in B_{\mathbf{o}}$. Since there is no other way for making true the creation action, only the organization itself can specify who create new obligations. In particular, $c \in X_r$ can be an action $CD(d)$ of a role $r \in RO$ of producing a document $d \in DC$: in this way the organization \mathbf{o} specifies that the role r has control over some other role $a \in RO$ such that $a \in A_{\mathbf{o}}$. The document d represents the record of the exercise of the power of agent r . Also functional areas are modelled as agents in an organization: hence, the same mechanism can be used to specify that an agent r has control over role $a \in RO$, where r and a can belong to the same functional area $m \in FA$ ($\{r, a\} \subseteq A_m \cap RO$).

Since the “counts as” relation can be iterated, it is possible to specify how a role $r \in RO$ belonging to a functional area $m \in FA$ ($r \in A_m$) of an organization $\mathbf{o} \in OA$ can create an obligation $O_{ao}(x, s | Y)$ directed to a functional area or role $a \in FA \cup RO$ directly belonging to the organization: $a \in A_{\mathbf{o}}$. This is possible since an action $c \in X_r$ of role r can count as an institutional fact $p \in I$ for the functional area m : $c \rightarrow p \in B_m$. In turn, the institutional fact p can count as the creation of an obligation $O_{ao}(x, s | Y)$ by the organization \mathbf{o} : $p \rightarrow CR(b, \mathbf{o}, O_{ao}(x, s | Y)) \in B_{\mathbf{o}}$; this obligation is directed towards agent a which belongs to the organization \mathbf{o} . These relations are only possible since the beliefs B_m of the functional area m are attributed to agent m by the organization \mathbf{o} itself, since $m \in A_{\mathbf{o}}$. For example, a decision of the CEO counts as an obligation of the entire organization since the direction functional area to which the CEO belongs considers the CEO’s decision as made by itself and the organization, in turn, considers the decision of the direction as having the obligation as a consequence. In this way, the organization, when it creates its components by attributing mental attitudes to them, at the same time, constructs its control structure.

The second issue is communication among roles. It is often claimed [10] that the organizational structure specifies the communication possibilities of agents. Agents can communicate almost by definition and standard communication languages have been defined for this aim [11]. What the organization can specify is their possibility to communicate to each other in an institutional way by means of documents; as Wooldridge *et al.* [22] claim, organizations specify “systematic institutionalized patterns of interactions”.

Communication among socially constructed agents is based on the same principle as control. It relies on the fact that the beliefs of a functional area or of a role are attributed to them by the higher level socially constructed agent which they are attributed mental attitudes by. In this way we can express the fact that a document created by a role $r \in RO$ communicates some belief p to an organization or functional area $m \in OA \cup FA$ it belongs to $r \in A_m$: $CD(d) \rightarrow p \in B_m$, where $CD(d) \in X_r$ is an action creating a document $d \in DC$. This is read as the fact the action of role r “counts as” the

official belief p of agent m . The document d represents the record of the communication between r and m .

Analogously, we can specify official communication among roles. A role $r \in RO$ communicates to a role $a \in RO$ that $p \in P$ if there is some action $CD(d) \in X_r$ creating a document $d \in DC$ such that $CD(d) \rightarrow p \in B_a$. Note that B_a are not the beliefs of the agent $b \in RA$ playing role a ($b = PL(a)$). Rather they are the beliefs attributed to the role by the functional area $m \in FA$: since the role a is created by the functional area m , those beliefs are attributed to a by the functional area m . When an agent $b \in RA$ which plays the role $a \in RO$ knows that document d has been created, it has to act as if it had the belief p , while it is not requested to be psychologically convinced that p is true. Otherwise agent b does not stick to its role anymore and it becomes liable to having violated its duties.

5 Summary

In this paper we propose a way to model the organizational structure of multiagent systems. Organizations are composed by functional areas and roles; functional areas, in turn, are composed by functional areas and roles. Roles are played by agents. Using the methodology of attributing mental attitudes to social entities, we show that organizations and their components can be described as agents: socially constructed agents. Since socially constructed agents are agents, they can construct, in turn, other agents which constitute their components. This strategy allows creating a decomposition structure as rich as the one in object orientation. Moreover, it allows progressively decomposing an organization in simpler agents described by beliefs and motivations to manage the complexity of a multiagent system. Finally, since agents can be subject to obligations and endowed with permissions and powers, all the social entities composing an organization can be the addressees of norms and powers; at the same time, socially constructed agents can be normative systems imposing obligations on their components, i.e., organizations can be modelled as bureaucracies [18].

This paper is part of a wider project modelling normative multiagent systems. In [6] we model normative systems by means of the agent metaphor: we attribute them beliefs, desires and goals: beliefs represent the constitutive rules of the organization while regulative rules, like obligations, are modelled in terms of goals of the system. In [2] we extend the model to virtual communities and we use the agent metaphor to describe local and global policies. In [5], constitutive rules are used to define contracts and games among agents are extended to allow an agent to change the obligations enforced by the normative system. Roles have been introduced in [4]. This paper constitutes a step forward in this project in that the agent metaphor is used to explain how organizations can create other social entities like functional areas and roles and, at the same time, specify their behavior. In this way we account for their definitional dependency characteristic of social entities [16]. Our ontology of social reality is presented in [3].

Future work concerns defining the relation between roles described as agents and the agents playing those roles. Moreover, contracts, described in [5] can be introduced to regulate the possibility to create new obligations, new roles and new social entities inside an organization [10].

References

1. B. Bauer, J.P. Muller, and J. Odell. Agent UML: A formalism for specifying multiagent software systems. *Software Engineering and Knowledge Engineering*, 11(3):207–230, 2001.
2. G. Boella and L. van der Torre. Local policies for the control of virtual communities. In *Procs. of IEEE/WIC Web Intelligence Conference*, pages 161–167. IEEE Press, 2003.
3. G. Boella and L. van der Torre. An agent oriented ontology of social reality. In *Procs. of FOIS'04*, Torino, 2004.
4. G. Boella and L. van der Torre. Attributing mental attitudes to roles: The agent metaphor applied to organizational design. In *Procs. of ICEC'04*, 2004.
5. G. Boella and L. van der Torre. Contracts as legal institutions in organizations of autonomous agents. In *Procs. of AAMAS'04*, New York, 2004.
6. G. Boella and L. van der Torre. Regulative and constitutive norms in normative multiagent systems. In *Procs. of KR'04*, Whistler (CA), 2004.
7. G. Booch. *Object-Oriented Analysis and Design with Applications*. Addison-Wesley, Reading (MA), 1988.
8. D. Dennett. *The intentional stance*. Bradford Books/MIT Press, Cambridge (MA), 1987.
9. V. Dignum, J.-J. Meyer, and H. Weigand. Towards an organizational-oriented model for agent societies using contracts. In *Procs. of AAMAS'02*, pages 694–695, Bologna, 2002.
10. J. Ferber, O. Gutknecht, and F. Michel. From agents to organizations: an organizational view of multiagent systems. In *Procs. of AOSE'03*, pages 214–230. Springer Verlag, 2003.
11. T. W. Finin, Y. Labrou, and J. Mayfield. KQML as an agent communication language. In J. Bradshaw, editor, *Software Agents*. MIT Press, Cambridge, 1995.
12. N. R. Jennings. On agent-based software engineering. *Artificial Intelligence*, 117(2):277–296, 2000.
13. A. Jones and M. Sergot. A formal characterisation of institutionalised power. *Journal of IGPL*, 3:427–443, 1996.
14. J. Lang, L. van der Torre, and E. Weydert. Utilitarian desires. *Autonomous Agents and Multiagent Systems*, pages 329–363, 2002.
15. R.M. Lee. Documentary petri nets: A modeling representation for electronic trade procedures. In *Business Process Management*, pages 359–375, Berlin, 2000. Springer Verlag.
16. C. Masolo, L. Vieu, E. Bottazzi, C. Catenacci, R. Ferrario, A. Gangemi, and N. Guarino. Social roles and their descriptions. In *Procs. of KR'04*, Whistler (CA), 2004.
17. M. McCallum, T.J. Norman, and W.W. Vasconcelos. A formal model of organisations for engineering multi-agent systems. In *Procs. of CEAS Workshop at ECAI'04*, Valencia, 2004.
18. W.G. Ouchi. A conceptual framework for the design of organizational control mechanisms. *Management Science*, 25(9):833–848, 1979.
19. O. Pacheco and J. Carmo. A role based model of normative specification of organized collective agency and agents interaction. *Autonomous Agents and Multiagent Systems*, 6:145–184, 2003.
20. A. Ricci, A. Omicini, and E. Denti. Activity theory as a framework for mas coordination. In *Procs. of ESAW'02*, pages 96–110, 2002.
21. J.R. Searle. *The Construction of Social Reality*. The Free Press, New York, 1995.
22. M. Wooldridge, N.R. Jennings, and D. Kinny. The Gaia methodology for agent-oriented analysis and design. *Autonomous Agents and Multi-Agent Systems*, 3(3):285–312, 2000.
23. E. Yu. Agent orientation as a modelling paradigm. *Wirtschaftsinformatik*, 43(2):123–132, 2001.
24. F. Zambonelli, N.R. Jennings, and M. Wooldridge. Developing multiagent systems: The Gaia methodology. *IEEE Transactions of Software Engineering and Methodology*, 12(3):317–370, 2003.