

Organizational structure and responsibility

An analysis in a dynamic logic of organized collective agency

Davide Grossi · Lambèr Royakkers · Frank Dignum

Received: 5 October 2005 / Accepted: 1 May 2007 / Published online: 2 June 2007
© Springer Science+Business Media B.V. 2007

Abstract Aim of the present paper is to provide a formal characterization of various different notions of responsibility within groups of agents (Who did that? Who gets the blame? Who is accountable for that? etc.). To pursue this aim, the papers proposes an organic analysis of organized collective agency by tackling the issues of organizational structure, role enactment, organizational activities, task-division and task-allocation. The result consists in a semantic framework based on dynamic logic in which all these concepts can be represented and in which various notions of responsibility find a formalization. The background motivation of the work consists in those responsibility-related issues which are of particular interest for the theory and development of multi-agent systems.

Keywords Deontic logic · Multi-agent systems · Organizational structure · Responsibility

1 Introduction

The concept of responsibility is central to a theory of collective agency and organizations. Responsibility issues arise any time a group of agents acts collectively in order to achieve certain objectives. Plans are made for the collective

D. Grossi · F. Dignum
Utrecht University, Utrecht, The Netherlands

D. Grossi
e-mail: davide@cs.uu.nl

F. Dignum
e-mail: dignum@cs.uu.nl

L. Royakkers (✉)
Technology Management, Eindhoven University of Technology, P.O. Box 513, 5600 MB
Eindhoven, The Netherlands
e-mail: l.m.m.royakkers@tue.nl

action of the group and specific agents are stated to be “responsible” for certain tasks. If something goes wrong certain agents might be found “responsible” for what happened, they might be held “accountable” and be “blamed”. The notion of responsibility displays different nuances all related with particular aspects of collective agency and, predominantly, obligation and knowledge.

The way obligations and knowledge flow within groups of agents is in turn related with the organizational structure those groups display. The possibility to delegate tasks to subordinated agents, or to successfully inform other agents about the actual state of the organization, or the possibility to put effective monitoring and recovery mechanisms in place are all aspects influencing the assessment of responsibilities within organizations. If an agent is appointed to perform a specific task, but it does not get the necessary knowledge for correctly performing it, can it be held responsible for a failure in the execution of the plan? And in what sense precisely is it responsible? Again, if an agent is appointed to a task but it delegates it to a subordinated agent, does the failure of the subordinated agent determines a form of responsibility for the first agent? And in what sense? The paper proposes a formal analysis for grounding rigorous answers to this type of questions. The main thesis of the work consists in claiming that responsibility issues within groups of agents are essentially related with the way groups are *organized* in order to pursue their objectives. In a nutshell, the less a group of agents is organized, the more blurred becomes the assessment of responsibilities within the group.

To provide a formal understanding of responsibility issues within groups of agents is of definite importance for the theory and development of multi-agent systems (MAS). In fact, many methodologies for MAS, like for instance GAIA (Wooldridge et al. 2000) are based on organizational concepts as their cornerstones. A formal theory of responsibility would provide these methodologies with conceptual tools for interpreting, in organizational terms, faulty performances of a given MAS, and at the same time suggest guidelines for the design of MAS behaving in specific ways with respect to the assessment of responsibilities among the agents.

In order to provide the desired analysis we consider various ingredients that make a collective agency an organized collective agency. In particular, we will consider the way the objectives of the group of agents are translated into concrete plans via task division and task allocation and the type of organizational structure in force within the group. With respect to these notions the present paper builds on previous work we presented in (Grossi et al. 2004, 2005).

The paper is structured according to the following outline. In Sect. 2 the notion of organizational structure is extensively discussed and the analytical perspective presented in Grossi et al. (2005) is summarized. A first informal description of the notions of responsibility we are interested in is also provided. In Sect. 3 we introduce the language we are going to use and present its formal semantics. This section expands on the work presented in Grossi et al. (2004). Then, in Sect. 4.1, a formal semantics of the basic organizational activities of delegating, informing and monitoring is presented. In the same section a thorough analysis of the notion of plan as intended in this work is exposed. Section 5 provides a formalization of some notions of responsibility and their logical relations with the notion of organizational structures is studied by means of some propositions. In Sect. 6 similarities and

differences with some related work are discussed and finally, in Sect. 7, some concluding remarks follow.

2 Informal preliminaries

2.1 Organizational structure and organizational activities

Organizational structures are sets of relations between the roles of an organization. A typical abstract example of such structures is the so-called “vertical differentiation” or “authority structure” of organizations, usually considered to be a “hierarchy” structure. These abstract types of structures are traditionally studied in the branch of sociology known as mathematical sociology (Fararo 1997; Sørensen 1978).

Work on organizations (especially in MAS¹) presents organizational structure as something essentially mono-dimensional, though it often, but only implicitly, considers a multiplicity of structured aspects: authority, communication, delegation, responsibility, control, power, etc. The thesis we hold here, which is inspired by foundational work on social and organization theory like Selznick (1948), Morgenstern (1951) and Giddens (1984), is that organizations do not exhibit one single structural dimension, but that they are instead multi-structured objects. In particular, we view organizational structure as hiding at least three relevant dimensions which we call: power, coordination and control.

These different structural dimensions are linked with as many specific activities that take place within any organized group of agents acting to pursue some goals. These activities, which we call *organizational activities*, consists in “managing the interdependencies between the activities” (Decker and Lesser 1995) of the group. In other words they guarantee the group ‘to act in an organized way’. We will analyze three of these activities: delegation, information and monitor, each of them related with one specific structural dimension. The *delegation activity*, concerning the flow of obligations within an organization is related with the structural dimension of **power**. The information activity, concerning instead the flow of knowledge within the group of agents, is related with the **coordination** dimension. Finally, the monitoring activity, concerning the recovery functionalities of the organization, is related with the **control** dimension.

As a result of this analysis, organizations will be represented as explicitly displaying a triple structure constrained on the basis of the interplay between the three notions of power, coordination, and control (see Definition 3.1). Although we do not pretend to give full definitions of these relations, we will characterize these relations in terms of the consequences they bear for the performance of organizational activities. In a nutshell, our thesis is that organizational structure makes organizational activities possible, that is, what organizational structure does is to make it possible for a group of agents to act as an organization (see Sect. 4.1).

Let us thus have a closer look to what we called organizational activities. A quote from Selznick (1948) provides interesting hints:

¹ See (Horling and Lesser (2004) for an exhaustive survey.

“Delegation is the primordial organizational act, a precarious venture which requires the continuous elaboration of formal mechanisms of *coordination* and *control*”.

In fact, delegation concerns the redistribution of tasks within an organization. To realize their objective organizations need to attribute tasks to agents according to specific plans for the realization of those objectives (see Sect. 4.4). Delegation introduces a dynamic within this attribution of tasks transferring tasks from agents to agents when the recipient of the transfer plays a somehow subordinated role within the organization. The transfer takes place in the form of a directed obligation (Dignum 1999) of the agent enacting the first role to the agent enacting the subordinated one. This second agent is thus obliged to perform a task which belonged to the first agent. The possibility of delegating goals constitutes one of the essential aspects of organizations. The power structure concerns exactly the channels through which this task flow can take place: “who can (successfully) delegate to whom?”

The quote from Selznick (1948) mentioned above emphasizes the importance of the coordination and control issues within organized groups. In particular, every organization has to handle the knowledge problem concerning the state of the organization itself (or of part of it²) at a given moment. Agents should know when to act, that is, they should be informed about the status of the activities of the organization on which their activities depend³, and what they are obliged to do. As we observed above, delegation introduces dynamics in the task distribution of an organization. The point is that once a task is delegated and a correspondent obligation arises for a specific agent, a certain amount of information might be required for that agent to include that task in its own goal base and to pursue it. Because of this, an information mechanism which can keep track of this dynamics is crucial for the performance of an organization. To quote Morgenstern (1951):

“The description of a delegation system [power structure] is incomplete unless the simultaneous signaling system [coordination structure] applied to it is also explicitly described”.

The coordination structure should then guarantee that each agent has a representation of the actual state of the organization which is sufficient for it to accomplish its tasks. The question is then how the access and sharing of information is structured within the organization: who can (successfully) inform whom?

Objectives and norms also determine to a great extent the *control* structure. Since agents’ tasks can be accomplished or not, the monitoring activity is an indispensable activity in which any organization has to engage. This is the case also for MAS

² Notice, in passing, that the amount of knowledge to be propagated through the organization also constitutes an important issue:

“If every competence [role] had full information about every other it might help but not necessarily; it would clearly be wasteful, if not physically impossible, for most organizations” (Morgenstern 1951).

³ This issue has been formally investigated in Grossi et al. (2004).

organizations where agents, even if “benevolent”, are anyway subjected to the possibility of failure. In its simplest form, control consists in a monitoring activity triggering appropriate reactions to determinate failures or violations. If an agent fails in performing one of the stated or delegated tasks, a kind of supervisor agent should engage in the performance of that task: organization calls for a form of supervision activity (Giddens 1984). Because of this, control can be seen as “an organization within an organization” (Morgenstern 1951). With respect to the control issue the relevant structural question is: who can (successfully) monitor whom?

Besides the activities of delegation, information and control, and their correlated organizational structures, a range of yet more activities can be isolated which play a crucial role within organized agency. One of them concerns the distribution of the necessary capabilities within the organization for agents to perform the required tasks. Such issue is somehow analogous to the information issue since it concerns requirements each tasks presupposes in order to be accomplished. A basic type of capabilities lies in the amount of resources that agents should have at their disposal. A second kind of capabilities plays a central role in organizations, namely those concerning the so-called *institutional power* (Jones and Sergot 1996; Castelfranchi 2003). Again the problem is related with the dynamics introduced by the delegation activity: delegating a task may require a parallel enabling or empowering activity such as making the relevant resources accessible, e.g., electronic money, and providing the required form of institutional empowerment, e.g., a suitable document. In this case the relevant structural question is: who can (successfully) enable or empower whom? We choose however not to complicate matters further and we assume in this work that all agents have the capabilities needed to perform the tasks they are appointed to.

2.2 Notions of responsibility: a sketch

The analysis of organizational structures concerns organizations at their role level. Responsibilities, instead, concern agents and arise in relation with task-allocation and structure once there are agents enacting the roles of a given organization.

Given a task-allocation allocating a specific subtask to a role, and given that an agent is enacting that role, the agent is then said to be responsible for that task or *task-based responsible*. In other words, the allocation of subtasks to roles determines a distribution of what we call task-based responsibilities over the set of agents enacting the roles of the organization. Task-based responsibilities can also arise via delegation under a power link or of a monitoring action under a control link. Thanks to a power structure an agent can appoint a new task to a subordinated agent and because of a control structure a monitoring agent can happen to be appointed to a task whose accomplishment failed.

Being autonomous, agents can independently decide whether to perform the subtasks to which they are appointed or not, and whether to perform them in the expected way. In this case the fulfillment of the organizational objectives is put in jeopardy by the conduct of some agent that is said then to be *causally responsible* for the failure occurred.

In organizations an agent can happen to be causally responsible of some failure without actually being *accountable* for that in the eyes of the organization. This can happen if an agent which is task-based responsible for performing a task, delegates the performance to a subordinate agent which fails to execute the delegated task. The second agent is causally responsible for the failure but it is the delegating agent that will be held accountable for the failure since it was the one appointed to the task in the first place.

The notion of accountability eminently reveals an interplay between the notions of responsibility isolated above, and dimensions of social structure such as the possibility to delegate allocated tasks, i.e., what we called power relation in the previous section. The presence of a power structure within an organization causes a difference between the two notions of task-based and causal responsibility: ‘I may have not performed the task you delegated to me, but you were the one appointed to it’. However, this is not the only case in which a given organizational bears consequences for the assessment of responsibilities within a group. As we have seen in the previous section organizational structure influences also the flow of knowledge within an organization and knowledge is an essential ingredient of yet another way of ‘being responsible’ for something, namely, *blameworthiness*. An agent can be causally responsible for a failure without actually being blameworthy for that. This is the case, for instance, if it just was not informed about the task it was supposed to perform. The acknowledgment of such a gap calls for the distinction of yet another meaning of the notion of responsibility which we call *blameworthiness*.

All the informal notions of responsibility just sketched are formalized in [Sect. 5](#).

3 A semantic framework for organized action

The framework we are going to present develops the logic for collective agency presented in Grossi et al. (2004) and which in turn built on Royakkers (1998). The key idea consists in enabling the possibility to formalize not only notions concerning the activity of groups of agents, but also notions concerning the activity of *organized* groups of agents, that is to say of groups of agents displaying some form of organizational structure.

Technically, the framework will handle event expressions, that is expressions about the performance of some action by some agents, and their composition, epistemic expressions, deontic expressions concerning event expressions (and thus ought-to-do types of deontic notions), and predicates on events. The framework expands the proposal contained in Grossi et al. (2004) in one essential direction, namely, adding special propositions for describing organizational structures. These will denote the existence of power coordination and control links between roles, intuitively, that a role is under the power of another role, that a role can coordinate with another role, finally, that a role can control another role.

3.1 Language \mathcal{L}^{ORG}

The alphabet of \mathcal{L}^{ORG} consists first of all of a set of agent identifiers Ag (groups of agents identifiers are denoted by X, Y, \dots), and a set AR of roles identifiers. The set

P of propositional symbols (p) of \mathcal{L}^{ORG} contains at least the propositional constant V (violation) and all organizational structure propositions, that is, for any $r, s \in AR$ and $i \in Ag$: $Power(r, s)$ (there exists a power-link between r and s), $Coord(r, s)$ (there exists a coordination-link between r and s), $Control(r, s)$ (there exists a control-link between r and s) and $rea(i, r)$ (the agent denoted by i enacts the role denoted by r). Besides, \mathcal{L}^{ORG} contains a set \mathcal{A} of atomic action symbols typically denoted by \underline{a} ,⁴ the family of epistemic operators $\{K_i\}_{i \in Ag}$ (“agent i knows that ...”), the dynamic operator $[]$ (“after each execution of ... it holds that”), the LTL (linear time temporal logic) operators $@_{start}$ (“it holds at the initial state of the run”), $@_{-1}$ (“it holds at the previous state of the run”) and $@_{+1}$ (“it holds at the next state of the run”) and the operator DO denoting what event is going to happen next. The set \mathcal{A} of atomic actions contains at least the “do nothing” action expression $skip$ and all organizational action expressions, that is: $delegate(i, \alpha)$ (delegating action α to agent i), $inform(i, \phi)$ (informing agent i that ϕ is the case), $monitor(i, \alpha)$ (monitoring the execution of action α by agent i) with $i \in Ag$, $\phi \in \mathcal{L}^{ORG}$ and $\alpha \in \mathcal{A}$.

The language \mathcal{L} is based on three types of syntactic constructs that we are now going to define.

The set Act of action expressions (α) is defined through the following BNF:

$$\alpha ::= \underline{a} \mid skip \mid \overline{\alpha} \mid \alpha_1 + \alpha_2 \mid \alpha_1 \& \alpha_2 \mid \alpha_1 ; \alpha_2,$$

where $skip$ represents a “doing nothing” action, the overline stands for the event negation operator, $+$ stands for the indeterministic choice operator, $\&$ for the parallel performance operator and $;$ for the sequencing operator.

The set Evt of event expressions (ξ) is defined through the following BNF:

$$\xi ::= X : \alpha \mid \overline{X : \alpha} \mid \xi_1 + \xi_2 \mid \xi_1 \& \xi_2 \mid \xi_1 ; \xi_2.$$

Notice that the same notation for actions and event operators (negation, $+$, $\&$, $;$) is used. It is nevertheless obvious that they belong to different categories of operators. Notice finally that event expressions consist of an action expression (α) plus an index denoting a group of agents (X). Events performed by individuals should be indexed via singletons ($\{a\}$), however, in order to keep the notation light, we will often omit the singleton notation.

A subset of Evt is of particular interest for our purposes, that is the set of events of the form $X : \alpha$ such that X is a singleton and α is an organizational action in \mathcal{A} . We call this set of events $OrgEvt$, i.e., organizational events.

The set Ass of assertions (ϕ) is defined through the following BNF:

$$\begin{aligned} \phi ::= & p \mid V \mid DO(\xi) \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid [\xi]\phi \mid K_i\phi \mid @_{start}\phi \mid @_{-1}\phi \\ & @_{+1}\phi \mid rea(r, i) \mid Power(r, s) \mid Coord(r, s) \mid Control(r, s), \end{aligned}$$

where $i \in Ag$ and $r, s \in AR$. The other boolean connectives can be defined as usual.

⁴ In order to keep the notation as light as possible we will often omit the underlining, when this does not cause any ambiguity.

3.2 Models

With \mathcal{L}^{ORG} we want to be able to talk about a number of different notions and in particular: organizational structures, hypothetical performances of actions by agents (i.e., events and their combinations in plans) and actual performances of actions by agents resulting in a run of the agent system. Our models should thus be rich enough to give a precise semantics to all these ingredients.

To model organizational structures we make use of the theory of directed graphs. However, we introduce only some basic elements of it which are strictly of use for the development of the article.⁵

Definition 3.1 (*Organizational structures*) An organizational structure OS is a tuple:

$$\langle Roles \cup Agents, R_{Power}, R_{Coord}, R_{Control}, Rea \rangle$$

where $Roles \cup Agents$ is the finite set of roles and agents of the organization, and R_{Power} , R_{Coord} , $R_{Control}$ are three irreflexive binary relations on $Roles$ characterizing the Power, Coordination and Control structures. Rea is a subset of $Agents \times Roles$ and indicates which agents play which roles. It is called *role enactment configuration* (or role adoption configuration).

The notion of organizational structure is an essential ingredient of the models for \mathcal{L}^{ORG} . They are defined as follows.

Definition 3.2 (*Models*) A model M is a structure:

$$M = \langle \mathcal{P}^+(Agents), \mathbb{A} \cup \text{skip}, \mathbb{W}, [\]_R, \{\mathcal{K}_i\}_{i \in Agents}, run, \pi, OS, J \rangle$$

where:

- $\mathcal{P}^+(Agents)$ is the non-empty powerset of the finite set of actors $Agents$, that means the possible groups of actors.
- $\mathbb{A} \cup \text{skip}$ is the set of actions.
- \mathbb{W} is the set of possible states.
- $[\]_R$ is a function f s.t. $f : Evt \times \mathbb{W} \longrightarrow \mathcal{P}(\mathbb{W})$, to each event expression-world couple it associates the set of states to which the performance of that event in that world leads. It consists of a composition of the two functions $[\]$ and R which will be introduced in Sect. 3.3.
- $\{\mathcal{K}_i\}_{i \in Agents}$ is a family of reflexive symmetric and transitive accessibility relations which are indexed by actors indicating the accessible worlds representing the epistemic alternatives of agent a_i .
- run is a structure $run = \langle \mathbb{W}_0, \prec \rangle$ modeling an actual run of the agent system and such that:
 - \mathbb{W}_0 denotes the set of states of \mathbb{W} reached by the run ($\mathbb{W}_0 \subseteq \mathbb{W}$);

⁵ For comprehensive expositions we refer the reader to Harary et al. (1965), Harary (1969) and Ross and Wright (1992).

- \prec is a finite path of length n on \mathbb{W}_0 (a sequence $\langle w_1, \dots, w_{n+1} \rangle$ of distinct elements of W_0 s.t. $\forall w_i 1 \leq i \leq n, w_i \prec w_{i+1}$) denoting the order in which worlds are reached in the run through actual performances of events. Path \prec has therefore always a first and a last state. We denote the first state of the run as w_{start} .

The ordering \prec is constrained as follows: (a) if $w_1 \prec w_2$ then $\langle w_1, w_2 \rangle \in \llbracket \xi \rrbracket_R$ for some $\xi \in Evt$; (b) if $w_1 \prec w_2$ and $\exists w_3$ s.t. $w_3 \mathcal{K}_i w_1$ or $w_3 \mathcal{K}_i w_2$ then $w_3 \mathcal{K}_i w_2$ and $w_3 \mathcal{K}_i w_1$. From an intuitive point of view, the first condition states that the run happens always through transitions that are labeled by some event, and the second condition guarantees the whole path of actual performances through \mathbb{W} to be epistemically accessible.

- π is a usual truth function $\pi : Ass \times W \longrightarrow \{1, 0\}$
- OS is an organizational structure (Definition 3.1).
- $J = \langle J_a, J_r \rangle$ where $J_a : Ag \longrightarrow Agents$, i.e., J_a is a function that maps agent names into agents, and $J_r : AR \longrightarrow Roles$, i.e., J_r is a function that maps role names into the corresponding roles.

Like in Meyer (1988) and Dignum et al. (1996) our semantics consists of two parts: first event expressions are interpreted as set theoretic constructs on \mathbb{A} where events get a so-called *open* interpretation; successively event expressions are interpreted as state-transition functions determining the accessibility relation $\llbracket \cdot \rrbracket_R$ on \mathbb{W} .

3.3 Synchronicity sets, steps, synchronicity traces, and worlds

The interpretation of events is based on the basic notion of *synchronicity set* (s-set).

Definition 3.3 (*s-set*) The set \mathcal{S} of s-sets is defined as follows $\mathcal{S} = \mathcal{P}^+(Agents) \times \{\text{skip}\} \cup \mathcal{P}^+(Agents) \times \mathcal{P}^+(\mathbb{A})$.

Synchronicity sets, that is elements of \mathcal{S} , are denoted by S_1, S_2, \dots . Informally, a s-set is nothing but a set of parallel executions of events by a group of agents, and formalizes the aforementioned open interpretation view on events. We will often refer to s-sets by making the group of agents explicit in an index (e.g. S_X). Based on the notion of s-set we define the notion of *step*.⁶

Definition 3.4 (*Step*) The set *Step* of steps is defined as follows:

$$\begin{aligned}
 Step = & \{ \{S_X\}_{X \in \mathcal{P}^+(Agents)} \mid \forall X \in \mathcal{P}^+(Agents) : S_X \in \mathcal{S} \\
 & \& \forall X, Y \in \mathcal{P}^+(Agents) : \\
 & Y \subseteq X \Rightarrow act(S_Y) \subseteq act(S_X) \\
 & \& \forall X, Y \in \mathcal{P}^+(Agents) : \\
 & act(S_Y) = \text{skip} \Rightarrow act(S_{X \cup Y}) = act(S_X) \}
 \end{aligned}$$

⁶ Notice that in Dignum et al. (1996) s-sets are called steps, and no notion of step as it will be defined in this work occurs there.

where act is a function that extracts the action component from a given s-set ($act(X: \{a_1, a_2\}) = \{a_1, a_2\}$).

Steps represent a sort of snapshot of the activity of each subgroup of *Agents* at a certain moment, depicting how all agents move one ‘‘step’’ ahead. Steps are therefore sets of s-sets of cardinality $2^n - 1$ where n is the number of agents in *Agents*. They are constrained in such a way that whatever action is performed by a subgroup is also performed by a supergroup, and subgroups remaining inactive are treated as performing a skip action. Steps, that is elements of *Step*, are denoted by st_1, st_2, \dots .

In order to provide a semantics for sequential expressions the concept of *synchronicity trace* (s-trace) is needed. Notice that this concept uses steps instead of s-sets like it was originally defined in Meyer (1988).

Definition 3.5 (*s-trace*) The set \mathcal{T} of s-traces is defined as follows:

$$\mathcal{T} = \{ \langle st_1, \dots, st_n, \dots \rangle \mid st_1, \dots, st_n, \dots \in Step \}.$$

The length of an s-trace t is denoted by $dur(t)$. We assume $dur(t)$ to be finite.

An event will be interpreted as a set of s-traces. The range for our interpretation of events is a set \mathcal{E} such that $\mathcal{E} = \mathcal{P}(\mathcal{T})$. Elements of \mathcal{E} (sets of s-traces) are denoted as T_1, T_2, \dots . The length $dur(T)$ of a set T is defined as $max\{dur(t) \mid t \in T\}$.

We can now introduce the operations that constitute the semantic counterpart of our syntactic operators.

Definition 3.6 (*Operations on events*) Let $T_1, T_2 \in \mathcal{T}$:

$$\begin{aligned} T_1 \circ T_2 &= \{ t_1 \circ t_2 \mid t_1 \in T_1, t_2 \in T_2 \} \\ T_1 \cap T_2 &= \bigcup \{ t_1 \cap t_2 \mid t_1 \in T_1, t_2 \in T_2 \} \\ T_1 \sqcup T_2 &= T_1 \cup T_2 \setminus \bigcup \{ t_1 \cap t_2 \mid t_1 \in T_1, t_2 \in T_2 \text{ and } t_1 \neq t_2 \} \\ \tilde{T} &= \begin{cases} \text{if } T \neq \emptyset, & \tilde{T} = \cap \{ st \mid st \in T \} \\ \text{if } T = \emptyset, & \tilde{T} = Step \end{cases} \end{aligned}$$

where

- $t_1 \circ t_2$ is defined as follows: if $t_1 = \langle st_1, \dots, st_n \rangle$ and $t_2 = \langle st'_1, \dots, st'_m \rangle$ then, $t_1 \circ t_2 = \langle st_1, \dots, st_n, st'_1, \dots, st'_m \rangle$.
- $t_1 \cap t_2$ is defined as follows: $t_1 \cap t_2 = \begin{cases} t_1 & \text{if } t_2 \in start(t_1) \\ t_2 & \text{if } t_1 \in start(t_2) \\ \emptyset & \text{otherwise} \end{cases}$ where $start$ is a function which associates to a given s-trace all its starting possible s-traces: $start(t) = \{ t' \mid t' = t \text{ or } \exists t'' \neq \emptyset \text{ s.t. } t' \circ t'' = t \}$.
- If $t = \langle st_1, \dots, st_n \rangle$ then \tilde{t} is defined as follows:

$$\tilde{t} = \bigcup_{1 \leq n \leq dur(t)} \langle st_1, \dots, st_n \rangle$$

where $\tilde{st} = Step - \{st\}$.⁷

Intuitively, we want \cup to yield the property: $a \equiv a + a; b$ for event expressions. In order to establish this property we cannot just use a union of the sets of s-traces representing a and $a; b$ but have to do some ‘‘cleaning up’’ by subtracting superfluous parts.

The semantics of events are obtained by means of a function $\llbracket \cdot \rrbracket : Evt \rightarrow \mathcal{E}$ such that:

Definition 3.7 (Semantics of events)

$$\begin{aligned} \llbracket X : a \rrbracket &= \{st \in Step \mid st = S_X, a \in act(S_X)\} \\ \llbracket \xi_1; \xi_2 \rrbracket &= \llbracket \xi_1 \rrbracket \circ \llbracket \xi_2 \rrbracket \\ \llbracket \xi_1 + \xi_2 \rrbracket &= \llbracket \xi_1 \rrbracket \cup \llbracket \xi_2 \rrbracket \\ \llbracket \xi_1 \& \xi_2 \rrbracket &= \llbracket \xi_1 \rrbracket \cap \llbracket \xi_2 \rrbracket \\ \llbracket \bar{\xi} \rrbracket &= \llbracket \tilde{\xi} \rrbracket \\ \llbracket skip \rrbracket &= \{skip\}. \end{aligned}$$

The basic clause stipulates that the meaning of an atomic event consists of the set of steps where that action at least is performed by that specific group of agents.

On the basis of this evaluation for events, an evaluation of groups performing complex actions is obtained:

Definition 3.8 (Semantics of collective actions)

$$\begin{aligned} \llbracket X : \alpha_1; \alpha_2 \rrbracket &= \llbracket X : \alpha_1 \rrbracket \circ \llbracket X : \alpha_2 \rrbracket \\ \llbracket X : \alpha_1 + \alpha_2 \rrbracket &= \llbracket X : \alpha_1 \rrbracket \cup \llbracket X : \alpha_2 \rrbracket \\ \llbracket X : \alpha_1 \& \alpha_2 \rrbracket &= \llbracket X : \alpha_1 \rrbracket \cap \llbracket X : \alpha_2 \rrbracket \\ \llbracket X : \bar{\alpha} \rrbracket &= \overline{\llbracket X : \alpha \rrbracket}. \end{aligned}$$

To connect this interpretation of events to a possible world semantics a function $R : \mathcal{E} \times \mathbb{W} \rightarrow \mathbb{W}$ is defined, which couples events with state-transitions.

Definition 3.9 (Function R) $R(T, w_1) = \{w_2 \mid \exists t \in T \text{ s.t. } w_2 = R(t, w_1)\}$ where R on transitions is inductively defined as follows:

$$\begin{aligned} R(st_1, w_1) &= reach(st_1, w_1) \\ R(t_1 \circ t_2, w_1) &= R(t_2, R(t_1, w_1)) \end{aligned}$$

⁷ Negation of sequences constitutes a delicate matter. For a deeper discussion of this issue we refer to Dignum and Meyer (1990).

function $reach : Step \times \mathbb{W} \rightarrow \mathbb{W}$ being a function that, given a state, returns the following states reachable through a given step, and such that $reach(\{X : skip\}_{X \in \mathcal{P}^+(Agents)}, w) = w$.

3.4 Evaluating formulae

The meaning of formulae ϕ in a world w , given the structure M , is defined as usual. We report here only the clauses for the organizational structure propositions, the epistemic, dynamic and temporal operators, and the *DO* operator.

Definition 3.10 (*Satisfaction relation*) In the following let $dur(\llbracket \xi_1 \rrbracket) = 1$,

- $M, w \models Power(r, s)$ iff $R_{Power}(J(r), J(s))$
- $M, w \models Coord(r, s)$ iff $R_{Coord}(J(r), J(s))$
- $M, w \models Control(r, s)$ iff $R_{Control}(J(r), J(s))$
- $M, w \models rea(a, r)$ iff $Rea(J(a), J(r))$
- $M, w_1 \models [\xi]\phi$ iff $\forall w_2 \in \llbracket \xi \rrbracket_R(w_1) : M, w_2 \models \phi$
- $M, w_1 \models K_a\phi$ iff $\forall w_2, w_1 K_a w_2 : M, w_2 \models \phi$
- $M, w \models @_{start}\phi$ iff $M, w_{start} \models \phi$
- $M, w_1 \models @_{-1}\phi$ iff $\exists w_2 \in \mathbb{W}, w_2 \prec w_1 : M, w_2 \models \phi$
- $M, w_1 \models @_{+1}\phi$ iff $\exists w_2 \in \mathbb{W}, w_1 \prec w_2 : M, w_2 \models \phi$
- $M, w_1 \models DO(\xi_1)$ iff $\forall w_2 \in \mathbb{W}, w_1 \prec w_2 \Rightarrow w_1 \in \llbracket \xi_1 \rrbracket_R w_2$
- $M, w_1 \models DO(\xi_1; \xi)$ iff $\forall w_2 \in \mathbb{W}, w_1 \prec w_2 \Rightarrow (M, w_2 \models DO(\xi) \text{ and } M, w_1 \models$

Informally *rea* assertions are an isomorphic representation within the language of the *Rea* relation in *OS*. The same holds for *Power*, *Coord* and *Control* assertions. As to the dynamic operator, a sentence $[\xi]\phi$ is true in w iff ϕ is true in every world accessible through a performance of ξ . The semantics of $@_{start}$, $@_{-1}$ and $@_{+1}$ is quite simple and it is just based on the \prec path contained in M . As to the semantics of $DO(\xi)$, the two clauses should be read as a basis and an induction step: intuitively, a sentence $DO(\xi)$ is evaluated as true in a world w_1 iff state w_2 in the run can be reached via the sequence ξ of events. Notice that a backward looking operator *DONE*, denoting what event took place in order to reach the present state, can be defined as follows: $DONE(\xi) := @_{-1}DO(\xi)$. In what follows we will thus sometimes use *DONE*-expressions as abbreviations for $@_{-1}DO(\xi)$.

3.5 Deontic notions

Essential for our purposes is the possibility to express a notion of obligation, in particular of obligation for a group of agents to execute a given plan. We express deontic notions making use of a reduction strategy in the classical fashion of Meyer (1988).

Definition 3.11 (*Deontics*) The deontic operator for obligation O is defined as follows:

$$O(\xi) \equiv [\bar{\xi}]V$$

For an extensive account the type of deontic logic generated by this reduction we refer to Royakkers (1998). As shown in Grossi et al. (2004), the reduction makes the formulae in Table 1 valid.

4 Organizational actions and organizational structure

4.1 Semantics of organizational actions

In Sect. 2 we have informally touched upon a number of activities that typically take place within an organized group of agents: delegation, information, monitor. In what follows we give a formal semantics of these activities, aiming at capturing some of their essential features. We do not consider our analysis, however, to exhaust all the aspects involved in the notions of delegation, control and monitor within organizations. Our aim is rather to capture those aspects that look more relevant in relation with the notions of responsibility which have been introduced in Sect. 2 and which are formally investigated in Sect. 5.

Essentially, the semantics we propose formalizes the connection between the organizational structures of power, coordination and control and the organizational activities of *delegation*, *information* and *monitor*. The existence of structural links between roles guarantees the successful performance of those organizational actions. In other words, in order for a group of agents to act in an organized way, that is, to be able to manage their collective endeavors, specific structures between the roles are necessary in order to guarantee the effectiveness of organizational activities.

If a power relation holds between roles r and s , all delegation acts performed by an agent a enacting role r on agents enacting role s succeed in creating an obligation for these agents. Analogously, if a coordination relation holds between roles r and s , all information acts performed by agents enacting role r to agents enacting role s are successful in the sense that they create knowledge in these agents. Finally, if a control relation holds between roles r and s , all monitoring acts performed by agents enacting role r on agents enacting role s do not only create knowledge in the controller about the

Table 1 Some validities concerning deontic notions

$O(X \cup Y : \bar{\gamma}) \rightarrow O(X : \bar{\gamma})$	(1)
$O(X : \gamma) \rightarrow O(X \cup Y : \gamma)$	(2)
$O(X : \gamma_1 + Y : \gamma_1) \rightarrow O(X \cup Y : \gamma)$	(3)
$O(X : \gamma_1 \& Y : \gamma_2) \rightarrow O(X \cup Y : \gamma_1 \& \gamma_2)$	(4)
$O(X : \gamma_1 ; Y : \gamma_2) \rightarrow O(X \cup Y : \gamma_1 ; \gamma_2)$	(5)
$O(X : \alpha_1 \& Y : \alpha_2) \leftrightarrow O(X : \alpha_1) \wedge O(Y : \alpha_2)$	(6)
$O(X : \alpha_1 ; Y : \alpha_2) \leftrightarrow O(X : \alpha_1) \wedge [X : \alpha_1]O(Y : \alpha_2)$	(7)
$O(X : \alpha_1) \vee O(Y : \alpha_2) \rightarrow O(X : \alpha_1 + Y : \alpha_2)$	(8)

relevant state of affairs, but they also determine an obligation for the controller in case the controlled agent did not perform the action that is monitored.

Definition 4.1 (*Semantic constraints for $a : delegate(b, \alpha)$*) For any step st and s-set S such that $a : delegate(b, \alpha) \in act(S)$ and $S \in st$ (with $a, b \in Ag$ and $\alpha \in Act$) and any $w \in W$:

$$reach(st, w) = \begin{cases} \{w' \mid w' \in reach(st', w) \text{ and } M, w' \models O(b : \alpha)\} \\ \text{if } M, w \models K_a O(a : \alpha) \text{ and } M, w \models Power(r, s) \\ \quad \wedge rea(a, r) \wedge rea(b, s) \\ reach(st', w), \text{ otherwise} \end{cases}$$

where st' is the step obtained from st removing all occurrences of $a : delegate(b, \alpha)$ from each of its s-sets.

Intuitively, if a power relation exists between roles that are enacted by two agents and the delegating agent about knows the to-be-delegated obligation, then a delegate action has as effect an obligation for the recipient. A delegate action implements therefore, given an appropriate power link, a form of ‘‘your wish is my command’’ principle.

More technically, the definition states that, provided that the preconditions in the first clause hold, the set of state-transitions generated by a step st where event $a : delegate(b, \alpha)$ is performed, is the subset of the state-transitions generated by the step st' where no $a : delegate(b, \alpha)$ takes place in which all transitions end up satisfying $O(b : \alpha)$. In other words what *delegate* actions do, with respect to all the other actions being performed in the step, is *just* creating obligations, given that the necessary preconditions hold. If the preconditions do not hold, then it is as if the action was never performed (this is the intuitive meaning of the ‘‘otherwise’’ clause).

Definition 4.2 (*Semantic constraints for $a : inform(b, \phi)$*) For any step st and s-set S such that $a : inform(b, \phi) \in S$ and $S \in st$ (with $a, b \in Ag$ and $\alpha \in Act$) and any $w \in W$:

$$reach(st, w) = \begin{cases} \{w' \mid w' \in reach(st', w) \text{ and } M, w' \models K_b \phi\} \\ \text{if } M, w \models K_a @_{+1} \phi \text{ and } M, w \models Coord(r, s) \\ \quad \wedge rea(a, r) \wedge rea(b, s) \\ reach(st', w), \text{ otherwise} \end{cases}$$

where st' is the step obtained from st removing all occurrences of $a : inform(b, \phi)$ from each of its s-sets.

Intuitively, if there exists a coordination link between the role enacted by the informing agent and the role enacted by the recipient, and provided that the informing agents knows that the to-be-communicated content is going to be the case in the next state reached by the system ($K_a @_{+1} \phi$), then an *inform* action always results in the creation of the corresponding epistemic state in the recipient. A coordination relation enables thus agents with reliable and trustworthy information

channels. Typical to-be-communicated contents are obligations to perform some action in the next state, or information about actions just undertaken.

In analogy with the definition of delegation it is stated that, provided that the preconditions in the first clause hold, the set of state-transitions generated by a step st where event $a : inform(b, \phi)$ is performed, is the subset of the state-transitions generated by the step st' where no $a : inform(b, \phi)$ takes place in which all transitions end up satisfying $K_b\phi$.

Definition 4.3 (*Semantic constraints for $a : monitor(b, \alpha)$*) For any step st and s-set S such that $a : monitor(b, \alpha) \in S$ and $S \in st$ (with $a, b \in Ag$ and $\alpha \in Act$) and any $w \in W$:

$$reach(st, w) = \begin{cases} \{w' \mid w' \in reach(st', w) \\ \text{and } M, w' \models K_a DONE(b : \alpha)\} \\ \text{if } M, w \models DO(b : \alpha) \wedge K_a O(b : \alpha) \\ \text{and } M, w \models Control(r, s) \wedge rea(a, r) \wedge rea(b, s) \\ \{w' \mid w' \in reach(st', w) \\ \text{and } M, w' \models K_a DONE(b : \bar{\alpha}) \wedge K_a O(a : \alpha)\} \\ \text{if } M, w \models DO(b : \bar{\alpha}) \wedge K_a O(b : \alpha) \\ \text{and } M, w \models Control(r, s) \wedge rea(a, r) \wedge rea(b, s) \\ reach(st', w), \text{ otherwise} \end{cases}$$

where st' is the step obtained from st removing all occurrences of $a : monitor(b, \alpha)$ from each of its s-sets.

Intuitively, if a control relation exists between the role enacted by the monitoring agent and the one enacted by the monitored agent, and the monitoring agent knows a certain action ought to be performed by the monitored agent, then the *monitor* action is:

- an *informative action* (Meyer and Van der Hoek 1995), i.e., after the performance of $a : monitor(b, \alpha)$ either $K_a DONE(b : \alpha) \vee K_a DONE(b : \bar{\alpha})$;
- an action generating a *recovery obligation* on the monitoring agent, in case the monitored agent did not performed the action whose performance is checked.

This semantics constraint models therefore the idea that a control link between two roles on the one hand enables the monitoring agents with the necessary tools and capabilities for being always able to ascertain whether the action to be checked was actually performed or not, and on the other it attributes to the monitoring agent tasks of a recovery kind. Notice that monitor actions are performed in parallel with the to-be-monitored actions.

Looking at the definition from a more technical point of view, we see again the same patterns used in Definitions 4.1 and 4.2. If the preconditions in the first clause hold true, then the set of state-transitions generated by a step s where event $a : monitor(b, \alpha)$ is performed, is the subset of the state-transitions generated by the step s' where no $a : monitor(b, \alpha)$ takes place in which all transitions end up satisfying $K_a DONE(b : \phi)$. If the preconditions in the second clause hold true, then the set of transitions leads to worlds all satisfying $\models K_a DONE(b : \bar{\alpha}) \wedge K_b O_b \alpha$. If

none of the preconditions hold, the monitor action does not influence the transition generated by the step.

It is easy to see that Definitions 4.1, 4.2 and 4.3 make Formulae 9, 10 and 11 in Table 2 valid in our models.

4.2 Organizational actions and knowledge

We consider organizational actions to be such that their necessary effects are always known to the agents performing them. In other words, the actions of delegating, informing, and monitoring, once executed by an agent, always determine the knowledge about their necessary effects which the agent expects. To use the terminology of Meyer and Van der Hoek (1995), these actions are always *accordant to plan*.

Definition 4.4 (*Knowing the effects of organizational actions*) If $a \in \text{OrgEvt}$, then for any step st s.t. $a \in st$:

$$\forall w_1, w_2 : w_2 \in R(st, w_1) \Rightarrow (\forall w_3 : (w_2 \mathcal{K}_i w_3 \Rightarrow (\exists w_4 : (w_1 \mathcal{K}_i w_4 \& w_3 \in R(st, w_4)))))$$

with i being the agent of event a .

Intuitively, the constraint guarantees that any world reachable via a concatenation of the transitions of step st and \mathcal{K}_i , is also reachable via a concatenation of \mathcal{K}_i and the transitions generated by s . It can be proven (see Meyer and Van der Hoek 1995) that such constraint validates Formulae 12, 13 and 14 in Table 2.

4.3 Organizational actions, knowledge, and deontics

Organizational actions are activities by means of which any collective agency can be managed. As we have seen they are dependent on the organizational structure of a group. When an agent, given a role-based plan and an enactment configuration, is

Table 2 Validities concerning organizational actions

$(\text{Power}(r, s) \wedge \text{rea}(a, r) \wedge \text{rea}(b, s) \wedge K_a O(a : \alpha))$	(9)
$\rightarrow [a : \text{delegate}(b, \alpha)] O(b : \alpha)$	(10)
$(\text{Coord}(r, s) \wedge \text{rea}(a, r) \wedge \text{rea}(b, s) \wedge K_a @_{+1} \phi) \rightarrow [a : \text{inform}(b, \phi)] K_b \phi$	(11)
$(\text{Control}(r, s) \wedge \text{rea}(a, r) \wedge \text{rea}(b, s) \wedge K_a O(b : \alpha))$	(12)
$\rightarrow [a : \text{monitor}(b, \alpha)] (K_a \text{DONE}(b : \alpha) \vee (K_a \text{DONE}(b : \bar{\alpha}) \wedge K_a O(a : \alpha)))$	(13)
$K_a ([a : \text{delegate}(b, \alpha)] \psi) \rightarrow [a : \text{delegate}(b, \alpha)] K_a \psi$	(14)
$K_a ([a : \text{inform}(b, \phi)] \psi) \rightarrow [a : \text{inform}(b, \phi)] K_a \psi$	(15)
$K_a ([a : \text{monitor}(b, \alpha)] \psi) \rightarrow [a : \text{monitor}(b, \alpha)] K_a \psi$	(16)
$O(a : \text{delegate}(b, \alpha)) \rightarrow K_a O(a : \text{delegate}(b, \alpha))$	(17)
$O(a : \text{inform}(b, \phi)) \rightarrow K_a O(a : \text{inform}(b, \phi))$	(18)
$O(a : \text{monitor}(b, \alpha)) \rightarrow K_a O(a : \text{monitor}(b, \alpha))$	(19)

appointed to perform a certain organizational action we consider reasonable to assume that it also knows about this appointment. In fact, we consider such knowledge to follow from the role enactment itself: if an agent enacts a role it acquires knowledge about its tasks. This motivates the following semantic constraint.

Definition 4.5 (*Knowing about organizational tasks*) If $a \in \text{OrgEvt}$, then for any step st s.t. $a \notin s$:

$$\forall w_1, w_2 : ((w_2 \in R(st, w_1) \Rightarrow M, w_2 \models V) \Rightarrow (\exists w_3 : w_1 \mathcal{K}_i w_3 \& w_2 \in R(st, w_3)))$$

with i being the agent of event a .

The constraint states that if all the worlds reachable via st satisfy the violation constant, then the same worlds are reachable via a concatenation of \mathcal{K}_i and transitions generated by s . It is easy to see that such a constraint makes Formulae 15, 16 and 17 in Table 2 valid.

4.4 Role-based and agent-based plans

Organizations “represent rationally ordered instruments for the achievement of stated goals” (Selznick 1948), that is, organizations arise in order to achieve specific objectives, and these objectives are pursued defining a number of subgoals contributing to the overall purpose of the organization. These subgoals identify the roles that are played in the organization. The relation between subgoals and overall objectives of the organization, i.e., the primitive decomposition of tasks within the organization, defines the essential form of organizational structure: “viewed in this light, formal organization is the structural expression of rational action” (Selznick 1948). Roles are the basic units over which this structure ranges determining the source of the “rational order” holding in the organization. The above quotes consider then the decomposition of tasks as the central source of structure within organizations: structure is necessary for each organization to pursue its objectives.

In order for the objectives of an organization to be realized the organization needs to “translate” them in concrete sub-goals to be systematically reached following specific plans, i.e., via complex collective actions which, once performed, guarantee the achievement of those objectives. Normally, this “translation” of objectives into plans is described via the two steps of *task division* and *task allocation*.

Through the *task division* process, goals are reduced to complex actions. Notice that task division consists of two steps. First a *raw plan* is found, which consists only of the atomic actions necessary for carrying out the organizational goal at issue:

$$\alpha_1 \bullet \dots \bullet \alpha_n$$

where \bullet stands for one of the event composition operators (so, $\bullet \in \{;, \&, +\}$), for all $1 \leq i \leq n$ $\alpha_i \in \mathcal{A}$ (we consider thus plans to be spelled out in terms of atomic actions). In terms of our running example, suppose that the program committee has

selected the following task division for the notification of acceptance: the chairman collects the submitted papers and divides the papers among the other PC members; the PC members review the papers they have received from the chairman and send their results to the chairman; the chairman makes the final decision which papers are selected for the workshop and informs the authors about the decision. Such a row plan does not include all the organizational actions necessary for the program committee to manage the performance of the plan itself. Coordination actions are required between the chairman and the PC members (PC members should know they have to review the papers appointed to them by the chairman) as well as monitoring actions (the chairman should control all reviewers do their job and possibly take appropriate counter measures in case of failure). We call row plans including all the necessary organizational actions *proper plans* or simply *plans*.

Through the *task allocation* process, each atomic component of the complex action, which is intended to realize a specific objective of the organization, is allocated to one agent. Within groups displaying an explicit organizational structure, the task allocation process consists of two essential steps. First, given a plan, each action component of the plan is linked to a role of the organization. In this view, roles are therefore placeholders within a plan description. A plan in which the atomic action components are indexed with roles identifiers is called *role-based plan* and it looks like this:

$$r_1 : \alpha_1 \bullet \cdots \bullet r_n : \alpha_n$$

where \bullet stands for one of the event composition operators (so, $\bullet \in \{;, \&, +\}$), for all $1 \leq i \leq n$ $\alpha_i \in \mathcal{A}$ and $r_i \in AR$. Notice that, obviously, different actions can be indexed with a same role.

The second step in a task allocation consists in the so-called *role enactment* specifying which agent of the organization plays which role. Again, different roles can be enacted by a same agent.⁸ In this work, agents playing a role in an organization are called *role enacting agents* or *rea*'s. We have already introduced the notion of *role enactment configuration* in Definition 3.1 formalized by the relation *Rea* in OS structures and representable in \mathcal{L}^{ORG} via finite conjunctions of the form:

$$rea(a_1, r_1) \wedge \cdots \wedge rea(a_i, r_i) \wedge \cdots \wedge rea(a_n, r_n)$$

such that $\forall 1 \leq i \leq n, a_i \in Ag$ and $r_i \in AR$.

Given a role-based plan and a role-enactment configuration a corresponding *agent-based plan* can be obtained which specifies which agent of the organization has to play which role in the plan. In other words a role-based plan $Plan(AR, \tau)$ and a role enactment configuration *Rea* univocally determine an agent-based plan, i.e., a complex event description.

⁸ Schematically, a task allocation given a plan consists of a surjection of the set of roles AR of the organization onto the set of atomic action components of the plan, followed by a surjection of the set of agents Ag onto AR , provided that all roles are employed in the task allocation and that all agents enact at least a role.

Definition 4.6 (*Agent-based plan*) An agent-based plan $Plan(Ag, \tau)$ for a task τ within the set of agents Ag is a structure:

$$Plan(Ag, \tau) = \langle Plan(AR, \tau), Rea \rangle$$

As such, agent-based plans be represented in \mathcal{L}^{ORG} as an event expression of the form:

$$a_1 : \alpha_1 \bullet \cdots \bullet a_n : \alpha_n$$

where $\bullet \in \{;, \&, +\}$, for all $1 \leq i \leq n$ $\alpha_i \in \mathcal{A}$, $a_i \in Ag$ and such that:

- $Plan(Ag, \tau)$ is obtained from $Plan(AR, \tau)$ by substitution of the role indexes r_i with the agent indexes a_i according to Rea ,
- $M, w \models [a_1 : \alpha_1 \bullet \cdots \bullet a_n : \alpha_n]\tau$.

The definition makes explicit the translation step of the organizations' objective into concrete plans for groups of agents: from an organizational level (roles) to a collective agency level (agents). Complex event expressions can be seen as the result of an instantiation process of role-based plans via role enactment configurations.

4.5 Plans and structure

In the previous section we distinguished between raw plans, i.e., complex action descriptions not including any organizational action, and proper plans, i.e., complex action descriptions which include instead organizational actions. The step from raw plans to proper plans is the most typical feature of planning a collective activity with respect to planning an individual one. When a plan concerns only the performance of a single agent, organizational activities such as delegating, informing and monitoring loose their meaning since those activities just happens within the single mind of one individual agent. Groups have, instead, no single mind even though they can act as if they had one precisely by undertaking appropriate organizational activities. Given a raw plan, an organization always needs to elaborate a corresponding proper plan which can accordingly manage the knowledge flow and the control issue within the group.

As we have seen in Sect. 4.1, organizational actions require, in order to be successful, specific structural constraints among the roles of the organization and specific enactment configurations. So, if an agent-based plan requires a certain agent a to inform agent b about ϕ then a suitable coordination link between the roles enacted by a and b should be put in place, or otherwise the information action could fail in transferring the necessary knowledge to b . Analogously, if an agent-based plan requires a to monitor the performance of b with respect to action α , a suitable control link between the roles enacted by a and b should be effective, or otherwise the monitoring action could fail not creating the necessary knowledge in a . These observations have precise formal counterparts. In fact, it can easily be seen that if the suitable structural links and enactment configurations do not hold, the following formulae are satisfiable in the models for \mathcal{L}^{ORG} :

$$DO(a : inform(b, \phi)) \wedge \neg[a : inform(b, \phi)]\phi \quad (18)$$

$$DO(a : monitor(b, \alpha)) \wedge \neg[a : monitor(b, \alpha)](K_a DONE(b : \alpha) \vee (K_a \neg DONE(b : \alpha) \wedge K_a O(a : \alpha))). \quad (19)$$

Satisfiability of such formulae can be seen as a sign of faulty design of the organization, where the organizational structure is not tuned on the organizational activities needed for managing the collective agency.

There is however another face of the coin. Given a desired plan, a suitable organizational structure can be designed or, vice versa, given an organizational structure, appropriate plans can be designed to meet the objectives of the organization. Proper plans can be chosen on the basis of the available structural links and enactment configuration. The delegation activity can play an essential role in this sense, improving given plans via attributing tasks to more suitable agents. Again, this cannot successfully happen without appropriate structural links, and Formulae 18 and 19 have a delegation variant:

$$DO(a : delegate(b, \beta)) \wedge \neg[a : delegate(b, \beta)]O(b : \beta) \quad (20)$$

which is also satisfiable if no power link is put in place.

5 Responsibilities in form

5.1 Causal responsibility

An agent a_i is said to be *causally responsible* for a state of affairs ϕ by performing α iff it performs an action α whose necessary effect is the state of affairs ϕ and it is not the case that if it did not perform α then ϕ was anyway the case.

Definition 5.1 (*Causal responsibility*) The causal responsibility of $a_i \in Ag$ for ϕ by performing α is defined as follows:

$$\mathfrak{CausalR}(a_i, \phi, \alpha) := @_{-1}([a_i : \alpha]\phi \wedge \neg[\overline{a_i : \alpha}]\phi \wedge DO(a_i : \alpha))$$

Intuitively, a_i has just performed α in the last step of the run of the system and ϕ would have not been necessary the case if α was not performed by a_i .⁹ In other words, in the very previous state in the run the occurrence of ϕ marked the performance of α by a_i and an event of type $a_i : \alpha$ has led to the actual state (in which the causal responsibility is evaluated). A similar notion has been formalized, making use of a “bringing-it-about” modal logic, in Cholvy et al. (1997).

A particularly relevant case of causal responsibility is the *causal responsibility for a violation*. Such a notion is formalized by expressions of the form $\mathfrak{CausalR}(a_i, V, \alpha)$.

⁹ It might be worth recollecting that $[\overline{a_i : \alpha}] = [a_i : \overline{\alpha}]$.

5.2 Blameworthiness

An agent which is causally responsible, may not be considered *blameworthy*. For example, if the chairman of the Editorial Board has forgotten to inform a member a_i to review some papers in one week, and agent a_i did not review the papers in one week, then the achievement of the goal of the Editorial Board to notify of the results of the reviews within the deadline will not be met. The agent would be considered causally responsible, but it would not be considered *blameworthy*. An agent does something causally blameworthy, if it is causally responsible and if it knows that the action it performs leads to a violation which could be avoided by not performing the action.

Definition 5.2 (*Blameworthiness*) Blameworthiness of $a_i \in Ag$ for ϕ by performing α is defined as follows:

$$\mathbf{Blame}(a_i, V, \alpha) := \mathbf{CausalR}(a_i, V, \alpha) \wedge @_{-1}K_{a_i}([a_i : \alpha]V \wedge \neg[\overline{a_i : \alpha}]V \wedge DO(\overline{a_i : \alpha})).$$

That is to say, a_i is blameworthy iff it is causally responsible of V and it knew in the very previous state of the run that the performance of α would have resulted in the occurrence of V .

5.3 Task-based responsibility

The notion of *task-based responsibility* corresponds to the notion of duty and refers to what the individuals of the organization are expected to do in virtue of their roles. As proposed also in Conte and Paolucci (2004), we assume task-based responsibility to be a consequence of role enactment. As we have seen, a role-based plan together with an enactment configuration univocally determines an agent-based plan. An agent who accepts to play a given role in an organization takes a responsibility with regard to the accomplishment of that role, i.e., with the tasks associated to it Conte and Paolucci (2004).

Definition 5.3 (*Task-based responsibility*) An agent $a_i \in Ag$ is task-based responsible for action α_i iff there exists a to-be-executed agent-based plan $Plan_j(Ag, \tau) = a_1 : \alpha_1 \bullet \dots \bullet a_i : \alpha_i \bullet \dots \bullet a_n : \alpha_n$, that is, iff there subsists an obligation for a plan including event $a_i : \alpha$ as a component:

$$\mathbf{TaskR}(a_i, \alpha_i) := \bigvee_{j \in |PL(a_i : \alpha)|} O(Plan_j(Ag, \tau))$$

where $PL(a_i : \alpha)$ is the set of all finite plans having event $a_i : \alpha$ as one of their components.

In other words, an agent is task-based responsible for the performance of a given action iff the event $a_i : \alpha_i$ is part of the to-be-executed agent-based plan.

Definition 5.3 has interesting consequences concerning the relation between the notion of obligation and of task-based responsibility itself which can all be expressed in our language. First of all, task-based responsibility occurs in the very moment (situation, possible world) in which an obligation to execute a plan occurs which contains the relevant event (Formula 21). Besides, notice that task-based responsibility holds true at every execution stage of a plan (Formula 22) at least until the relevant event is performed for the last time (Formula 23).

Proposition 5.1 (*TasfR and O*) *Given a sequential plan $Plan(Ag, \tau) = a_1 : \alpha_1 ; \dots ; a_i : \alpha_i ; \dots ; a_n : \alpha_n$ in which all events are different from each other, the following are validities of our framework:*

$$O(Plan(Ag, \tau)) \rightarrow \text{TasfR}(a_i, \alpha); \quad (21)$$

$$O(Plan(Ag, \tau)) \rightarrow [a_1 : \alpha_1] \text{TasfR}(a_i, \alpha_i) \wedge (a_i, \alpha_i); \quad (22)$$

$$\dots \wedge [a_1 : \alpha_1; \dots; a_{i-1} : \alpha_{i-1}] \text{TasfR}$$

$$\neg(O(Plan(Ag, \tau)) \rightarrow [a_1 : \alpha_1; \dots; a_i : \alpha_i] \text{TasfR}(a_i, \alpha_i)). \quad (23)$$

Proof 5.1 Formula 21 follows directly from Definition 5.3. Formula 22 follows again from Definition 5.3, Formula 7 and Formula 21. Formula 23 is proven since there always exists a model satisfying $O(Plan(Ag, \tau)) \wedge \neg[a_1 : \alpha_1; \dots; a_i : \alpha_i] \text{TasfR}(a_i, \alpha_i)$, which intuitively corresponds to the model in which the task-based responsibility ceases to hold after the execution of the relevant event. \square

Formulae 21 and 23 deserve in particular some more words. It is worth noticing why Definition 5.3 guarantees the persistence of the responsibility through the execution of the plan. That depends on the fact that after each execution of a fragment of the plan a new obligation with respect to the rest of the plan holds (see Formula 7). And as far as these new obligations hold for plans which include the event concerning the task-based responsibility, that responsibility also holds. As soon as the event is performed, persistence is not guaranteed any more. As a matter of fact, it would be desirable to have a stronger version of Formula 23 such as $O(Plan(Ag, \tau)) \rightarrow [a_1 : \alpha_1; \dots; a_i : \alpha_i] \neg \text{TasfR}(a_i, \alpha_i)$. However, this last version is not a validity since we cannot rule out the possibility of violation constants holding also in those worlds reached via correct executions of the plan. Obligations just state that if certain actions are not performed, a violation necessarily occurs (Definition 3.11), but they do not exclude that violations can hold no matter what actions are performed.

5.4 Accountability

The notion of accountability concerns the interplay of causal and task-based responsibilities. We say that an agent is accountable for a violation if it caused the violation by performing an action α and if it was appointed to $\bar{\alpha}$ according to the plan that had to be executed at the beginning of the run. This notion is necessary because of the organizational actions of delegation and monitor which can lead to

the generation of new tasks throughout a run of the system which are different from the one attributed by the original plan at the beginning of the run. In other words, the original task allocation established at the beginning of the run of the system is the one used for assessing accountability. These considerations are captured in the following definition.

Definition 5.4 (*Accountability*) An agent $a_i \in Ag$ is accountable for a violation V by performing α iff it is blameworthy for V by performing α and it was, at the initial state of the run, task-based responsible for the execution α :

$$\mathfrak{AccountR}(a_i, V, \alpha) := \mathfrak{Blame}(a_i, V, \alpha) \wedge @_{start} \mathfrak{TaskR}(a_i, \alpha)$$

It might be interesting to spend a few words about the intuitive meanings of the two expressions $\mathfrak{TaskR}(a_i, \alpha)$ and $@_{start} \mathfrak{TaskR}(a_i, \alpha)$. The first formula can be read as “agent a_i will have sooner or later to perform α ”. The second formula captures the idea of an initial appointment of a task: “agent a_i has been appointed, according to the initial plan, to the performance of α ”. Notice that neither the first implies the second nor the second implies the first. They formalize logically unrelated notions. In fact, I can have a task even if I was not appointed to it by the initial plan, for instance because I have been addressee of a delegation action. On the other hand, if I was appointed to a task this does not imply that I have that task at the current state, because I have for instance already performed the required action.

5.5 Responsibilities and organizational structure

In this section we show how our analysis provides a way for understanding the influence of organizational structure on the various notions of responsibility formalized in the previous section.

First of all we show a quite obvious result, namely that the existence of a power structure can determine the occurrence of a causal responsibility if the addressee of a delegation act does not perform the required action. This is no surprise since we know (Definition 4.1) that the existence of a power structure determines the successful creation of directed obligations via delegation.

Proposition 5.2 (*The power structure is grounds for CausalR*) Let $a, b \in Ag$, $r, s \in AR$ and $\alpha, \beta \in A$. The following formula is a validity of our framework:

$$\begin{aligned} & (@_{-1} @_{-1} (Power(r, s) \wedge rea(a, r) \wedge rea(b, s) \wedge K_a O(a : \beta) \\ & \wedge DO(a : delegate(b, \beta)) \wedge @_{-1} DO(b : \bar{\beta}) \wedge \neg [b : \beta] V) \quad (24) \\ & \rightarrow \mathfrak{CausalR}(b, V, \bar{\beta}) \end{aligned}$$

Proof 5.2 We show that there is no countermodel of Formula 24, that is no model that makes the antecedent of Formula 24 true and the consequent false. A countermodel M should be such that $\exists w_1, w_2, w_3: w_3 \prec w_2 \prec w_1$ and $M, w_3 \models Power(r, s) \wedge rea(a, r) \wedge rea(b, s) \wedge K_a O(a : \beta)$; for Definition 4.1 and

the properties of $K_a, M, w_2 \models DONE(a : delegate(b, \beta) \wedge O(b : \beta) \wedge \neg [b : \beta]V \wedge DO(b : \bar{\beta}))$; and finally the consequent should be false, that is, $M, w_1 \models \neg @_{-1}([b : \bar{\beta}]V \wedge \neg [b : \beta]V \wedge DO(b : \bar{\beta}))$ which cannot be the case is given Definition 3.11 and Definition 5.1. \square

On the other hand a power structure is not enough for determining accountability. In fact, accountability depends on the initial task-allocation assumed by the organization and it is not influenced by the creation of new obligations which do not stem from that task-allocation.

Proposition 5.3 (The power structure is not grounds for $\mathfrak{A}ccount\mathfrak{R}$) *Let $a, b \in Ag, r, s \in AR$ and $\alpha, \beta \in \mathcal{A}$. The following formula is satisfiable in the logic:*

$$@_{-1}@_{-1}(Power(r, s) \wedge rea(a, r) \wedge rea(b, s) \wedge O(a : \beta) \wedge DO(a : delegate(b, \beta)) \wedge \mathfrak{B}lame(b, V, \bar{\beta}) \wedge \neg \mathfrak{A}ccount\mathfrak{R}(b, V, \bar{\beta})) \quad (25)$$

Proof 5.3 The desired model is a model M such that there exists a world $w: M, w \models \neg @_{start} \mathfrak{T}ask\mathfrak{R}(b : \beta)$. \square

The following proposition shows that the occurrence of blameworthiness requires the existence of a coordination structure, which is instead not relevant for the occurrence of causal responsibility.

Proposition 5.4 (Structural conditions for $\mathfrak{B}lame$ w.r.t. $\mathfrak{C}ausal\mathfrak{R}$) *Let $a, b \in Ag, r, s \in AR$ and $\alpha, \beta \in \mathcal{A}$. The following are validities of the framework:*

$$@_{-1}(K_b DO(b : \bar{\beta}) \wedge K_b \neg [b : \beta]V \wedge O(b : \beta) \rightarrow \mathfrak{C}ausal\mathfrak{R}(b, V, \beta); \quad (26)$$

$$(@_{-1}(K_b DO(b : \bar{\beta}) \wedge K_b \neg [b : \beta]V \wedge O(b : \beta)) \wedge @_{-1}@_{-1}(Coord(r, s) \wedge rea(a, r) \wedge rea(b, s) \wedge K_a @_{+1} O(b : \beta) \wedge DO(a : inform(b, O(b : \beta)))) \rightarrow \mathfrak{B}lame(b, V, \beta). \quad (27)$$

Proof 5.4 Proofs are given showing that countermodels are impossible. The negation of Formula 26 immediately results in a contradiction since the K_b operator obeys reflexivity ($K_b \phi \rightarrow \phi$). Formula 27 is considerably more complex but it is not difficult to show that its negation also results in an inconsistent formula. Since the antecedent of Formula 27 implies the antecedent of Formula 26, a countermodel M for Formula 27 should contain a world w_1 s.t. $M, w_1 \models \mathfrak{C}ausal\mathfrak{R}(b, V, \beta) \wedge \neg (@_{-1}K_b([b : \beta]V \wedge \neg [b : \bar{\beta}]V \wedge DO(b : \beta)))$ from which it follows, given what is stated to hold in the antecedent, that $M, w_1 \models \neg K_b O(b : \beta)$. However the antecedent states that $\exists w_2, w_3$ s.t. $w_3 \prec w_2 \prec w_1$ and $M, w_3 \models Coord(r, s) \wedge rea(a, r) \wedge rea(b, s) \wedge K_a @_{+1} O(b : \beta)$ and, thanks to Definition 4.2, $M, w_2 \models K_b O(b : \beta)$, which is impossible. \square

Formula 27 deserves in particular some more words. We see that in order to obtain blameworthiness the same condition determining causal responsibility is first of all required (first conjunct of the antecedent). The second conjunct of the antecedent states something about two steps backward in the run ($@_{-1}@_{-1}$). In that state an appropriate coordination configuration needs to be in place, and the informing agent should know that in the next state (i.e., one state backward in the run from the evaluation state) an obligation holds for the recipient of the information action. This complex statement is, thanks to the semantics of the information actions (Definition 4.2), guarantees that the recipient knows it is obliged and it therefore determines the necessary conditions for blameworthiness.

Finally we show that the control structure can determine the occurrence of new task-based responsibilities. If an agent is found to be causally responsible for a violation by not performing β , then if an appropriate control structure is in place and a monitor action is performed to check whether β has been performed then the monitoring agent becomes task-based responsible for that action.

Proposition 5.5 *(The control structure is grounds for \mathfrak{TasfR})* Let $a, b \in Ag$, $r, s \in AR$ and $\alpha, \beta \in \mathcal{A}$. The following is a validity of the framework:

$$\begin{aligned} & \text{CausalR}(b, V, \bar{\beta}) \wedge @_{-1}(\text{Control}(r, s) \wedge \text{rea}(a, r) \wedge \text{rea}(b, s)) \\ & \wedge DO(a : \text{monitor}(b : \beta)) \rightarrow \mathfrak{TasfR}(a, \beta) \end{aligned} \quad (28)$$

Proof 5.5 The formula is easily proven considering Definition 4.3, Definition 5.1 and Definition 5.3. \square

Other results of this kind are obtainable in the framework. Our aim in this section was to provide a satisfactory sample of some intuitive relations among the concepts of responsibility and organizational structure which find a natural formalization in the framework.

6 Related work

The work presented in this paper moves from a number of precise ideas about the notion of responsibility within organization. Responsibility, in its various senses, is related with the notions of obligations and knowledge which, within organized groups of agents, are in turn related to three essential aspects of organized agency:

- the notions of objectives and plans (and therefore task division and task allocation) of the organization;
- the organizational actions of delegating, informing and monitoring, that is to say, with the issue of the *management of the collective activity*;
- the notions of role, organizational structure, role-enacting agent.

This perspective is essentially different from the work on formalization of responsibilities and other organization-related concepts presented in Cholvy et al. (1997) and Santos et al. (1997). In those works a more abstract view on

organizations is assumed as starting point. As explicitly stated in Santos et al. (1997), organizations are viewed there as “instances of normative systems”, i.e., as agents’ interaction patterns obeying the rules stated by a normative system. Organizations, and therefore responsibilities, are analyzed from the point of view of the rules to which the organized group of agents is subjected. We consider this perspective perfectly legitimate and sound. However, it would not suit our purposes. In fact, in our work we were not interested in the normative systems of which an organization is an instance, but on the structures of the organization itself which are one of the possible instances of a normative system. In a sense, we abstracted from the rules and we looked more concretely at their results, i.e., the organizational structures that the rules can impose on a group of agents. That is why organizational structures are a first-class citizen in our framework. By doing this we can better characterize what the effect of those structures, which are specified by rules, are on the activities of agents. This emerges clearly from the formal treatment of actions in our work and in Cholvy et al. (1997 and Santos et al. (1997)). While we were able to handle in quite fine-grained details the notion of plan having at disposals sets of atomic actions, in those works agents’ activities are captured via “bringing-it-about” modal constructs and organizational activities (in particular delegation) are described via modal operators modeling forms of indirect action or influence (“bringing-it-about indirectly that”).

We deem worth stressing, however, that the two approaches are perfectly bridgeable, the bridge being the way sets of rules defines sets of organizational structures. This is very interesting topic which is worth future researches.

7 Conclusions

The work has provided an analysis of some elementary notions of responsibility in connection with the structure of an organization. The study of the relations between these two notions showed what are the structural requirements grounding specific notions of responsibilities and, conversely, what kind of responsibilities can be assessed on the basis of a given structure. Such results can provide useful hints on possible guidelines for the design of agents’ organizations exhibiting desirable properties.

Future work will focus on a yet more detailed semantic characterization of the organizational actions and of the structures on which they are funded, and on the study of more of these actions (for instance “committing”) possibly making use of work based on similar formalisms such as, for instance, (Dunin-Keplicz and Verbrugge 2002) which provides an analysis of collective commitments in a dynamic logic settings. Another point worth of future developments is the introduction of a temporal dimension in relation with the notion of obligation, following the work done in Dignum et al. (2004) and Broersen et al. (2004). This would allow for a more flexible specification of plans and of the obligations for their executions. It could at the same time provide insights also in the phenomenon of the dynamics of responsibilities which has not been addressed in this work.

Acknowledgments We would like to thank Paolo Turrini for reading a first draft of this work and for the helpful suggestions he provided.

References

- Broersen J, Dignum F, Dignum V (2004) Designing a deontic logic of deadlines. In: Lomuscio A, Nute D (eds) *Deontic Logic in Computer Science*. Springer, Berlin, pp 43–56
- Castelfranchi C (2003) The micro-macro constitution of power. *ProtoSociology* 18:208–268
- Cholvy L, Cuppens F, Saure CI (1997) Towards a logical formalization of responsibility. In: International conference on artificial intelligence and law. ACM Press, New York, pp 233–242
- Conte R, Paoletti M (2004) Responsibility for societies of agents. *J Artif Soc Social Simul* 7
- Decker K, Lesser V (1995) Designing a family of coordination algorithms. In: Lesser V (ed) *Proceedings of ICMAS'95*, San Francisco, CA, USA. The MIT Press, Cambridge, MA, USA, pp 73–80
- Dignum F (1999) Autonomous agents with norms. *Artif Intell Law* 7:69–79
- Dignum F, Meyer J-JCh (1990) Negations of transactions and their use in the specification of dynamic and deontic integrity constraints. In: Kwiatkowska M, Shields MW, Thomas RM (eds) *Semantics for concurrency*. Springer Verlag, Berlin, pp 61–80
- Dignum F, Meyer J-J Ch, Wieringa R, Kuiper R (1996) A modal approach to intentions commitments and obligations: intention plus commitment yields obligation. In: Brown MA, Carmo J (eds) *Deontic logic, agency and normative systems*. Springer, London, pp 80–97
- Dignum F, Broersen J, Dignum V, Meyer J-J Ch (2004) Meeting the deadline: why, when and how. In: Hinchey MG, Rash JL, Truszkowski W, Rouff C (eds) *Formal approaches to agent-based systems*. Third international workshop, FAABS 2004, Greenbelt, MD, USA, April 26–27, 2004. Revised Selected Papers, volume 3228 of *Lecture notes in computer science*. Springer, Berlin, pp 30–40
- Dunin-Keplicz B, Verbrugge R (2002) Evolution of collective commitment during reconfiguration. In: *AAMAS '02: proceedings of the first international joint conference on autonomous agents and multiagent systems*. ACM Press, New York, NY, USA, pp 1141–1142
- Fararo TJ (1997) Reflections on mathematical sociology. *Sociol Forum* 12:73–101
- Giddens (1984) *A social theory and modern sociology*. Polity Press
- Grossi D, Dignum F, Royakkers L, Meyer J-JCh (2004) Collective obligations and agents: who gets the blame? In: Lomuscio A, Nute D (eds) *Deontic Logic in Computer Science*. Springer, Berlin, pp 129–145
- Grossi D, Dignum F, Dastani M, Royakkers L (2005) Foundations of organizational structure in multi-agent systems. In: Dignum F, Dignum V, Koenig S, Kraus S, Singh MP, Woolridge M (eds) *Proceedings of AAMAS'05*. ACM Press, New York, pp 690–697
- Harary F (1969) *Graph theory*. Addison-Wesley, London
- Harary F, Norman RZ, Cartwright D (1965) *Structural models: an introduction to the theory of directed graphs*. John Wiley & Sons, New York
- Horling B, Lesser V (2004) A survey of multi-agent organizational paradigms. *Computer science technical report 04-45*, University of Massachusetts, May 2004
- Jones AJL, Sergot M (1996) A formal characterization of institutionalised power. *J IGPL* 3:427–443
- Meyer JJCh (1988) A different approach to deontic logic: deontic logic viewed as a variant of dynamic logic. *Notre Dame J Formal Logic* 29(1):109–136
- Meyer J-JCh, van der Hoek W (1995) *Epistemic logic for AI and computer Science*, vol 41 of *Cambridge tracts in theoretical computer science*. Cambridge University Press
- Morgenstern O (1951) ASTIA Document Number ATI 210734 (unpublished manuscript)
- Ross KA, Wright CRB (1992) *Graph theory*. Prentice-Hall, New Jersey
- Royakkers L (1998) *Extending deontic logic for the formalization of legal rules*. Kluwer Academic Publishers, Dordrecht
- Santos F, Carmo J, Jones A (1997) Action concepts for describing organised interaction. In: Sprague RA (ed) *Thirtieth annual Hawaii international conference on system sciences*. IEEE Computer Society Press, Los Alamitos (California), pp 373–382
- Selznick P (1948) Foundations of the theory of organization. *Am Sociol Rev* 13:25–35
- Sørensen AB (1978) Mathematical models in sociology. *Annu Rev Sociol* 4:345–371
- Woolridge M, Jennings NR, Kinny D (2000) The gaia methodology for agent-oriented analysis and design. *Auton Agents Multi-Agent Syst* 3(3):285–312