Responsibilities in Organizations

Lambèr Royakkers
Eindhoven University of Technology, The Netherlands
L.M.M.Royakkers@tm.tue.nl

Davide Grossi
Utrecht University, The Netherlands
davide@cs.uu.nl

Frank Dignum
Utrecht University, The Netherlands
dignum@cs.uu.nl

Abstract
This paper addresses the issue of responsibilities within organizations of agents. The main aim of the research is to provide a formal analysis of the connections between collective obligations to individual responsibilities. Which individual agent in a group should be held responsible if an obligation directed to the whole group is not fulfilled? To this aim, concepts from planning literature (like plan and task allocation) and organization theory are used in order to conceptualize collective agency and the organizational structures. These concepts are formalized in a dynamic deontic logic framework, which allows us to study the connections between all these concepts. We will give a formal account of the notion of coordination, power and control intended as management of interdependencies among agents' activities and show how these organizational relations together with a specific task decomposition determine the responsibilities within an organization.

1. Introduction

Organizations "represent rationally ordered instruments for the achievement of stated goals" (Selznick 1948). These goals can be, for example, to fulfill an obligation directed to the whole group, i.e., a collective obligation. Since, such a collective obligation is beyond the capacity of an individual agent, the agents have to communicate, cooperate, coordinate and negotiate with each other, to achieve the collective task: the fulfillment of the obligation. So, the cooperation between agents is an important issue. For that purpose several theories have been developed about joint goals, plans and intentions. All these joint attitudes try to capture some of the team elements of agents that work together. In this paper we follow up on that work and will look at the individual responsibilities: What are the responsibilities of agents if the organization in which agents perform their tasks gets an obligation to fulfill a certain goal? Notions of responsibilities in an organization can enable agents to make decisions about what each member of the organization is expected to do, and to anticipate the tasks of others.

E.g. a program committee of ICAIL2005 may have the collective responsibility to review all submitted papers before a certain time and notify all the authors of acceptance. First, the program committee has to have a plan to organize its task by the allocation of the different tasks among the members of the program committee. In that allocation, it is preferable that at least each researcher is responsible for a certain task he is allocated at the appropriate time. However, what happens if the person that is supposed to distribute the submitted papers to the program committee members does not perform his task? Are all the other individual tasks (obligations, responsibilities) void because the papers were not distributed? Do the other researchers have to find an alternative solution? Who has to check whether the papers are distributed or not? We touch upon some fundamental aspects of responsibilities in an organizational structure.

We are interested in formalizing the concept of responsibility in (norm-governed) organizations. The concept of responsibility is a central concept to all legal systems and norm-governed organizations. Analyzing this concept is therefore fundamental if we aim at improving the behavior of these systems or organizations. Obtaining a formal representation of responsibility, however, is quite complex because of the very different meanings of this concept can take. Our concept of responsibility is restricted to the analysis of organizational performance. Therefore, we clarify and classify some meanings of responsibility and we relate them to the three relevant dimensions of an organizational structure we isolated in Grossi, Dignum, Royakkers and Dastani (2005) following intuitions developed in foundational studies in organizations and social theory (cf. Selznick 1948, Morgenstern 1951, Giddens 1984). These three relevant dimensions are power, coordination and control, with their matching actions 'to delegate', 'to inform' and 'to monitor'. The coordination actions are actually only one type of meta actions that should be considered. Besides the plan to achieve the content of the obligation the group should create that plan, allocate agents to parts of the plan, create a plan for what to do when the original plan fails, etc. These meta actions should also be coordinated again creating in the end an infinite regression of meta actions. In this paper we will not take all these layers into account, but will limit us to the coordination actions that are necessary to indicate the several notions of responsibility. In Grossi, Dignum, Royakkers and Dastani (2005) we proposed a logical framework for the organizational structure, which will be used to formalize these notions.

This paper is structured along the following lines. In Section 2, we will analyze the central notions of this paper in an intuitive way. In order to formalize these notions we use a modal logic which will be described in Section 3. Several notions of responsibility (given a plan) will be discussed formally in Section 4. How the individual responsibilities relate to the underlying structure of an organization will be discussed in Section 5. In Section 6, we show how the formal framework can be used to model an example and how certain individual responsibilities can be derived depending on characteristics of the organizational structure and the task allocation. In the last section, we will draw some conclusions and give directions for future research.

2. Preliminaries: organizations, structures and responsibilities

In this section we introduce and explain, at a pre-formal level, the notions we are aiming at analyzing.

2.1. Task-allocation, roles and agents

As stated in section 1, 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. In other words, in order for organizations to fulfill their objectives, subtasks are isolated via a form of organizational planning and distributed in a way which defines the roles agents can play in contributing to the performance of the organization. We call this designing process of the activity of an organization *task-allocation*. Roles can then be seen as sort of placeholders in a rationally designed activity of an organizations: an agent taking part to the organization will occupy one of these places, that is, will play a role.[1] In this work, agents playing a role in an organization are called, following Dignum (2003), *role enacting agents* or *rea*’s.

2.2. Organizational structure

Following Selznick (1948), Morgenstern (1951), and Giddens (1984), we view organizational structure as hiding at least three relevant dimensions which we call: power, information and control. These three structures will be analyzed in relation with the basic organizational activities with which they are related: *power* in relation with the delegation activity, *coordination* in relation with the knowledge and information issues, and *control* in relation with the monitoring and recovery issues. That these activities and dimensions are of essential importance to the overall performance of an organization is common thesis in the theory of organizations:

“delegation is the primordial organizational act, a precarious venture which requires the continuous elaboration of formal mechanisms of *coordination* and *control*” (Selznick 1948).

As a result, organizations are viewed as explicitly displaying a triple structure constrained on the basis of the interplay between the three notions of power, coordination, and control: task-allocation defines the roles that are played in the organization; the power structure regulates the delegation activity within the organization; the coordination structure explicits how information flows within the organization; finally, the control structure specifies how the performance of the organization is monitored and kept stable.

2.3. Responsibilities in organizational performance

Thus far we have dealt with organizations at their role level, where the task-allocation and the organizational structure range. Responsibilities concern instead 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. Being autonomous, agents can independently decide whether to perform the subtasks to which they are allocated 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 blamed by the organization. This can happen if an agent i which is task-based responsible for performing a task, *delegates* the performance to a subordinate agent j which fails or jeopardizes the execution of the delegated task. This observation 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. Social structure in relation with responsibility will be discussed in detail in Section 5. Here it suffices to notice that 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’. Therefore, if an organizational task is not performed, the one being *socially responsible* in front of the organization, the one who gets the blame for the failure, is not necessarily the one causally responsible for it, but it is the one to which that task was appointed. The acknowledgment of such a gap calls for the distinction of yet another meaning of the notion of responsibility which we call *failure-based responsibility*: who should control the performance of an agent to check whether a failure occurs and take countermeasures if that is the case?

Although certainly not complete, the above intuitive picture shows all the types of responsibility notions we are interested in for this paper. In Section 4 all the notions we introduced above will be formally characterized and their interplay analyzed. Before getting to that, something about the logical framework we are going to use is said in the next section.

3. Sketch of a logic for organizational structure and responsibilities

The logical framework we need in order to provide a formal characterization of the concepts introduced in Section 2 requires key expressivity capabilities for dealing with collective agency issues and with organizational structure issues. With respect to collective agency, expressions should be available to talk about: the performance of complex actions by agents and their effects (dynamic component); the obligations directed to agents (deontic component); the knowledge agents are endowed with (epistemic component). With respect to organizational structure issues, the relevant expressions concern essentially the possibility to represent relational structures depicting respectively the power, the coordination and the control structures (relational component). In this work we avoid to mention any technical aspect of the framework, limiting ourselves to a description of the language and its expressivity. For an exposition of the semantics, which is quite rich, we refer the reader to the work we exposed in Grossi, Dignum, Royakkers and Dastani (2005), and Grossi, Dignum, Royakkers and Meyer (2004).

In what follows we provide a list of the family of expressions available which will be used in Section 4, providing an intuitive reading of them which suggest their underlying semantics. As we will see, each expression is parameterized by or concerns either an agent or a role. The sets of agent identifiers and of role identifiers are denoted, respectively, by *Ag* (agents) and *RA* (role agents). More in particular, in the system it is possible to express and reason about:

- A set of agents *Ag* and a set of roles *RA*
- Dynamic formulas such as $[\alpha_i]\varphi$, meaning that after each execution of α_i formula φ holds, where α_i is a parameterized construct of the type $i{:}\alpha$ denoting the performance of action α by agent *i*, or a composed construct such as: $i{:}\alpha_1 \;;\; j{:}\alpha_2$ (subsequent performance), $i{:}\alpha_1 \;\&\; j{:}\alpha_2$ (parallel performance), $\underline{i{:}\alpha}$ (*i* refrains from performing α).[2]
- Assertions such as "α has just been performed by agent *i*" (DONE($i{:}\alpha$)), "α is going to be the next action performed by agent *i*" (DO($i{:}\alpha$)), and "α lies in the capabilities of agent *i*" ($CAN_i(\alpha)$).
- Epistemic assertions to the effect that agent *i* knows that φ ($K_i(\varphi)$), and deontic assertions to the effect that agent *i* ought to perform action α (O($i{:}\alpha$)).
- Special propositions: *Power*(*r*,*s*) which indicate that 'the agent enacting role *r* has the agent enacting role *s* in its power' (i.e., the agent playing role *r* can delegate tasks to the agent playing role *s*), *Coordination*(*r*,*s*) which indicate that 'the agent enacting role *s* has access to the information accessed by the agent enacting role *r*', and *Control*(*r*,*s*) which indicate that 'the agent enacting role *r* controls the agent enacting role *s*' (i.e., the agent playing role *r* is responsible for the agent playing role *s*). Note that relations are defined on roles, i.e., on members of *RA*. We denote the fact that agent *i* enacts role *r*, by the special proposition *rea*(*i*,*r*).

The semantics of these expressions is obtained by mixing the semantics of the collective action logic we developed in Grossi, Dignum, Royakkers and Meyer (2004). with a graph-theoretical semantics modeling the organizational structures that can be represented via the special propositions *Power*(*r*,*s*), *Coordination*(*r*,*s*), *Control*(*r*,*s*) and *rea*(*i*,*r*)[3] which we presented in Grossi, Dignum, Royakkers and Dastani (2005).

4. A formal analysis

4.1. Task allocation

The distribution of the sub-tasks in an organization in order to achieve a certain goal or collective task τ depends on a plan of the organization, i.e., a concrete manner to achieve the goal (collective task). We can define a plan to achieve a certain goal τ as a decomposition of the complex action achieve(τ) by a sequence of (possible simultaneous) individual actions:

$$\text{Plan}(\text{achieve}(\tau)) = \langle \text{achieve}(\tau_1) * \text{achieve}(\tau_2) * \ldots * \text{achieve}(\tau_n) \rangle$$

such that

$$[\text{achieve}(\tau_1) * \text{achieve}(\tau_2) * \ldots * \text{achieve}(\tau_n)]\tau,$$

where * stands for either the simultaneous operator '&' or the sequential operator ';'. The action achieve(τ_1) & achieve(τ_2) stands for the simultaneous performance of achieve(τ_1) and achieve(τ_2), and action achieve(τ_1) ; achieve(τ_2) stands for the sequential composition of achieve(τ_1) and achieve(τ_2).

We need the simultaneous operator, since some actions have to be performed at the same time. The sequential operator is needed because some actions might depend on other ones: a certain action can only be performed if an other action is done. So, the plan must at least determine the *order* of sub-actions. For example, the notification of acceptance of a certain paper can only be done if it is reviewed by the delegated members of the program committee. The task-based responsibility of the performance of an action α by an agent depends not only on the individual who is committed to perform the action α, but also on agents who have to perform actions which are necessary to perform action α.[4]

Besides task division, task allocation is needed, which indicates which role of the organization has to achieve which sub-task of the complex task. We use the following definition for task allocation:

Definition 1 (Task allocation)
A task allocation for a task τ within the set of roles *RA* is defined as follows:

$$\langle r_1{:}\ \text{achieve}(\tau_1) * r_2{:}\text{achieve}(\tau_2) * \ldots * r_n{:}\text{achieve}(\tau_n) \rangle$$

such that

$$[r_1{:}\ \text{achieve}(\tau_1) * r_2{:}\text{achieve}(\tau_2) * \ldots * r_n{:}\text{achieve}(\tau_n)]\tau.$$

We refer to the task allocation of τ within *RA* as Plan(*RA*,τ). To indicate that task achieve(τ_1) has been allocated to role r_j in Plan(*RA*,τ) (for j = 1,2,...,n), we use the following notation: $\langle r_j : \tau_j \rangle \in \text{Plan}(RA,\tau)$.[5]

We will use the concept of task allocation as a starting point for framing the various notions we are interested in. In particular, as we will see in the coming section, it plays an essential role for the definition of the notion of task-based responsibility. Besides, we will analyze the notion of "failure" in the accomplishment of a task understanding it as an organizational variant of the notion of social harm described in Conte and Paolucci (2004). In our context, we define the *untoward event* $D\tau_r$ as the impossibility, or the reduction of the possibility to achieve the goal τ allocated to role r. The performance of an action α by an agent i enacting role r determining social harm can then be represented as $[i{:}\alpha]D\tau_r$, that means, after each execution of action α by agent i the social harm represented by $D\tau_r$ is the case.[6]

4.2. Responsibilities in form

We can now provide an action logic representation of the notions of responsibility isolated in Section 2

Causal responsibility
An agent is said to be *causally responsible* when it does something (or fails to do something) that causes the untoward event $D\tau$. We formalize causal responsibility as follows:

Definition 2 (Causal responsibility)
For all $i \in Ag$:

$$R^c_i(D\tau) := [i{:}\alpha]D\tau \wedge DO(i{:}\alpha) \wedge \neg D\tau$$

meaning that agent i is causally responsible for the untoward event if and only if agent i performs an action which necessarily determines the occurrence of the untoward event and, finally, the untoward event is not the case before the agent performs the action.

Causal responsibility can also be attributed to nonhuman events, for example, that a house is severely damaged in a storm. In this paper, we restrict ourselves to agents in an organizational context. Notice that an agent which is causally responsible, may not be considered *blameworthy*. For example, if the chairman of the program committee has forgotten to inform a program committee member i to review the three papers in one week, and agent i did not review the papers in one week, then the achievement of the goal of the program committee to notify of the results of the review within the deadline will be reduced. The member i would be considered responsible in the sense of having *caused the situation*, but he would not be responsible in the sense of *blameworthy*. An agent does something blameworthy, if he knows (or could have known) that the action he performs leads to the impossibility or the reduction of the possibility to achieve a goal τ:

Definition 3 (Causal blameworthiness)
For all $i \in Ag$:

$$Bl^c_i(D\tau) := [i{:}\alpha]D\tau \wedge DO(i{:}\alpha) \wedge \neg D\tau \wedge K_i([i{:}\alpha]D\tau)$$

The importance of the knowledge component in the dynamics of responsibilities within organizations is analyzed in detail in Section 5.

Task-based responsibility
The notion of *task-based responsibility* is somehow interchangeable with duty and refers to what individuals are expected to do in virtue of their social roles. We assume that task-based responsibility is a consequence of role adoption: 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 (cf. Conte and Paolucci 2004). In this paper, this notion of responsibility completely depends on the position an agent occupies in the performance of the organization.

Definition 4 (Task-based responsibility)
For all $i \in Ag$ and a task allocation $Plan(AR,\tau)$:

$$R^{tb}_i(\tau_j) := rea(i,r_j) \wedge \langle r_j{:}\ \tau_j \rangle \in Plan(AR,\tau) \wedge O(i{:}achieve(\tau_j)) \wedge \underline{[i{:}achieve(\tau_j)]}D\tau_j$$

The obligation $O(i{:}achieve(\tau_j))$ expresses that the organization entrusts agent i with his task τ_j ($rea(i,r_j) \wedge \langle r_j : \tau_j \rangle \in Plan(AR,\tau)$), and $[i{:}achieve(\tau_j)]D\tau_j$ expresses the empowerment of i to prevent the reduction of the possibility or the impossibility to achieve goal τ_j. So, an agent i fails to fulfill his task-based responsibility $R^{tb}_i(\tau_j)$ if he violates the norm $O(i{:}achieve(\tau_j))$ which leads to the untoward event $D\tau_j$. However, the agent is considered blameworthy when he actually knows (or could have known) that he has this obligation and that he can perform the action to achieve his task. For example, he has not received the information needed for the performance of his task, or the achievement of his task depends on an earlier task in the task allocation which is not performed. This notion of blameworthy can formally be described as follows:

Definition 5 (task-based blameworthiness)
For all $i \in Ag$ and a task allocation $Plan(AR,\tau)$:

$$\begin{gathered} Bl^{tb}_i(\tau_j) := rea(i,r_j) \wedge \langle r_j{:}\ \tau_j \rangle \in Plan(AR,\tau) \wedge \\ O(i{:}achieve(\tau_j)) \wedge \underline{[i{:}achieve(\tau_j)]}D\tau_j \wedge \\ K_i(O(i{:}achieve(\tau_j)) \wedge \underline{[i{:}achieve(\tau_j)]}D\tau_j) \wedge CAN_i(achieve(\tau_j)) \end{gathered}$$

Social responsibility
The notion of social responsibility builds, as anticipated in Section 2, on the notion of task-based responsibility, and it is somehow analogous to a notion of violation in standard deontic logic.

Definition 6
For all $i \in Ag$ and a task allocation $Plan(AR,\tau)$:

$$R^s_i(\tau_j) := R^{tb}_i(\tau_j) \wedge D\tau_j$$

that is to say, agent i has the responsibility to achieve τ_j and the achievement of τ_j is impossible or jeopardized. Notice that this notion of responsibility is very simple and is independent from the notion of causal responsibility.

5. Responsibilities and organization structure

Although the notion of organization structure does not play a direct role for the definition of the different types of responsibilities, it does play an important role in the dynamics of responsibilities within an organization. First of all, the causal and task-based responsibilities depend on a given task allocation, which splits the joint task (arising from the collective obligation) over the individual agents of the organization. The mechanism of this task allocation depends on the organizational structures. E.g., in an organization without any hierarchy it might be that all agents bid on some sub-task(s) of the tasks and allocation is done through choosing the optimal allocation based on these bids. In a purely hierarchical organization the allocation might be done through a delegation of sub-tasks through the hierarchy. So, these mechanisms depend on the existing power relations between the agents in an organization.

While formally task-based responsibilities can be allocated through the task allocation and the role enactment mechanisms, we already saw that agents should also be made aware of their responsibilities. The mechanism can be implicit, e.g., through the playing of a role in an organization an agent knows it has certain responsibilities that come with the role, but can also be explicit. For the latter an information structure should be present in the organization that facilitates the right dissemination of the information, such that agents also know they have a responsibility and can be held responsible when things go wrong. The latter point ties in with the last dynamic aspect of responsibilities. When an agent is task-based responsible for a task it should also perform the task. The organization should, in some way, monitor the progress of the performance such that both the blame for non-performance can be attributed correctly to an agent and also appropriate measures can be taken to repair the situation. This monitoring and repair (in case of failures) is structured along the control structure of the organization.

We cannot hope to provide a full account of all interactions between responsibilities and organizational structures. However, in the rest of this section we aim to capture some essential traits of those interconnections. We understand those relations essentially as guaranteeing some effects to the basic actions of *delegate*, *inform* and *monitor*, which play an essential role with respect to responsibilities and their development in organizations.

The following definitions characterize the influence of the organization relations on the actions above. Through these basic properties we can also formally analyze some consequences of them on the notions of responsibilities studied in the previous section.

Definition 7 (Power)
For all $i,j \in Ag$ s.t. $i \neq j$ and $r,s \in AR$:

$$(Power(\mathrm{r,s}) \wedge rea(i,r) \wedge rea(j,s)) \rightarrow$$
$$[i{:}delegate(j, \mathrm{achieve}(\varphi))]\mathrm{O}(j : \mathrm{achieve}(\varphi))$$

If a power relation exists between roles that are enacted by two agents then a *delegate* action will have as effect an obligation for the recipient, that is, a form of “your wish is my command” principle. Intuitively, if a power relation holds between roles *r* and *s*, all delegation acts performed by an agent *i* enacting role *r* on agents enacting role *s* succeed in creating an obligation for these agents.

Task-based responsibility cannot be delegated. If Agent *i* has, according the task allocation, to achieve task φ and has a power relation with agent *j*, he can delegate his task to *j*, but he remains task-based responsible for the achievement of φ. Since φ is not the original task of agent *j* according to the given task allocation (see definition 4). Agent *j*, however, can be causally responsible if he fails to fulfill his delegated obligation.

A difference between an individual task and a collective task is that in an individual task all information is readily available and can be reasoned about. However, when a collective task is divided over the individuals of that collective, they might not know the whole plan, typically do not have information about actions that are performed, etc. Therefore, we need a coordination structure.

Definition 8 (Coordination)
For all $i,j \in Ag$ s.t. $i \neq j$, $r,s \in AR$:

$$(Coordination(r,s) \wedge rea(i,r) \wedge rea(j,s)) \wedge$$
$$\mathrm{DONE}(i{:}monitor(\varphi)) \wedge \mathrm{K}_i\varphi \rightarrow$$
$$\mathrm{O}(i{:}\ inform(\mathrm{j},\varphi)) \wedge [i{:}inform(\mathrm{j},\varphi)]\mathrm{K}_j\varphi$$

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, in these last agents, the knowledge they acquired via monitoring the occurrence of a certain fact: the *inform* action will automatically lead to the corresponding epistemic state in the recipient. Further, there is a normative aspect: agent *i* *should* inform another agent *j* about φ if they are connected through a coordination link and if agent *i* has monitored (checked) φ.

On this basis, a coordination-related type of responsibility can be defined.

Definition 9 (Coordinational responsibility)
For all $i,j \in Ag$ and a task allocation Plan(AR,τ):

$$\mathrm{R}^{coor}_i\ (inform(j,\varphi)){:=}$$
$$\mathrm{rea}(j, r_l) \wedge < r_l{:}\ \tau_l > \in \mathrm{Plan}(AR,\tau) \wedge$$
$$\mathrm{O}(i{:}achieve(\mathrm{K}_j\varphi)) \wedge [\underline{i{:}achieve(\mathrm{K}_j\varphi)}]\mathrm{D}\tau_l$$

This notion of responsibility has some resemblance to the task-based responsibility (see definition 4), with as task for agent *i* to inform agent *j* about φ. On the basis of the coordination structure, there is a specific allocation of the information actions, which is needed for the achievement of the individual tasks in the task allocation. Given this definition, we can say agent *i* is responsible to inform agent *j*, when the knowledge of φ is a necessary means to the achievement of τ_l and that agent j does not have that knowledge.

The responsibility of an agent i to inform some agent j about a certain aspect φ can follow from the coordination link between these agents if the knowledge of φ is necessary for the achievement of the task of agent j according the task allocation and i can monitor or check φ. This shows, in particular, how a given task allocation needs to be integrated with a suitable allocation of coordinational responsibilities in order to guarantee the information necessary for the correct functioning of the organization. This property can be formalized as follows:
For all $i,j \in Ag$ s.t. $i \neq j$, $r_k, r_l \in AR$ and task allocation Plan(AR,τ):

$$\begin{gathered} Coordination(r_k, r_l) \wedge rea(i, r_k) \wedge rea(j, r_l) \wedge \\ < r_l : \tau_l > \in \mathrm{Plan}(AR,\tau) \wedge \\ (\neg K_j\varphi \rightarrow \neg\, \mathrm{CAN}_j(achieve(\tau\_l)) \wedge \mathrm{CAN}_i(monitor(\varphi)) \\ \rightarrow R^{coor}{}_i\,(inform(j,\varphi)) \end{gathered}$$

So, agent i is responsible to inform agent j about φ if there is a coordination link between the roles r_l and r_k they respectively enact, and without the information about φ agent i_l cannot perform his task according to the task allocation. If agent i does not inform agent j, it follows that agent j cannot perform his task, which can lead to $D\tau_l$. So, agent i can be causally responsible if he does not inform agent j about φ (see definition 2). Note, that agent j is still task-based responsible with respect to τ_l, but not blameworthy, when he does not get the information necessary for the achievement of τ_l (see Section 4).

Finally, we get to a characterization of the dimension of control in organizational structure:

Definition 10 (Control)
For all $i,j \in Ag$ s.t. $i \neq j$ and r_k, $r_l \in AR$:

$$\begin{gathered} Control(r_k, r_l) \wedge rea(i, r_l) \wedge \mathrm{rea}(j, r_l) \rightarrow \\ [i{:}monitor(\mathrm{DONE}(j{:}achieve(\varphi))](D\tau_1 \rightarrow O(i{:}\ achieve(\tau_1)) \end{gathered}$$

If a control relation exists then the *monitor* action will have as further consequence the generation of an obligation for the controller in case the controlled actor did not achieve the relevant state causing the untoward event. On this basis, the notion of *failure-based responsibility* can be defined.

Definition 11 (Failure-based responsibility)
For all $i,j \in Ag$ s.t. $i \neq j$ and $r,s \in AR$:

$$\begin{gathered} R^{control}{}_i\,(monitor(j,\varphi)) := \\ Control(r,s) \wedge rea(i,r) \wedge rea(j,s) \end{gathered}$$

This type of responsibility depends completely on the control relation.

The control responsibility has another normative aspect: if an agent has control over another agent he is obliged to monitor the controlled agent whenever he knows the controlled agent has an obligation. Formally,

$$\begin{gathered} R^{control}{}_i\,(monitor(j,\varphi)) \wedge K_i(O(j{:}\mathrm{achieve}(\varphi)) \rightarrow \\ O(i{:}\ monitor(\mathrm{DONE}(j{:}\mathrm{achieve}(\varphi)))). \end{gathered}$$

We can imagine that an agent who has delegated his task to agent j, has the obligation to monitor whether the delegated agent has done the task, since he might be responsible to monitor agent j and he knows that the delegated agent j has the obligation.

6. An example

In this section we provide a few examples on how our formalism can be used in order to formalize and understand some organizational aspects of the simple scenario we have been often referring to through the paper.

Consider the program committee of the ICAIL2005 conference, with the roles of chairman (c), secretary (s), and reviewer (a_1, ...,a_n), such that $AR = \{c,s, a_1,...,a_n\}$. The collective task τ of the program committee is to notify the authors of the acceptance of their papers.

Let us then suppose that the program committee has selected the following plan for the notification of acceptance: the secretary collects the submitted papers and divides the papers among the reviewers; the reviewers a_1 up to and including a_n review the papers they have received from the secretary and send their results to the secretary. Then the secretary makes a summary of the results and send it to the chairman. The chairman makes the final decision which papers are selected for the workshop and informs the authors about the decision. This corresponds with the following task allocation:

Plan(AR,τ) = < s : collect papers ; s : divide papers ;
a_1 : review$_1$ & ... & a_n : review$_n$;
s : collect reviews ; s : send reviews to c ;
c : decide ; c : inform authors >

In the task allocation only the ‘performance functions’ (or the ‘direct work’) needed for the collective task, i.e., notification of acceptance, are described. All the agents are task-based responsible for their task defined by the task allocation. In the task allocation, the special actions based on the three relevant dimensions of an organizational structure, i.e., power, coordination and control, are not defined. The difference between these kinds of special actions is very important for the issue of responsibility.

Power
Suppose $rea(j,a_1)$, so agent j enacts role a_1 (the role of reviewer). Agent j is a professor with an overfull agenda, and he is not able to review the papers before the deadline. Since he has a power relation with his PhD-students, he can delegate his task to a PhD-student i. From definition 7 it follows that if professor j delegates his task to the PhD-student i, the PhD-student is obliged to review the papers: $O(i{:}\mathrm{review}_1)$. He is, however, not task-based responsible, since it was not a task of the PhD-student according the task allocation (note that the PhD-student does not enact role a_1).

If the PhD-student does not review the papers, he will be causally responsible for the (social) harm, in this example, for the delay of the process for reviewing the papers (Dreview_1), since now it holds that

$$[\underline{i\text{:achieve}(\text{review}_1)}]\text{Dreview}_1 \wedge \text{DO}(\underline{i\text{:achieve}(\text{review}_1)}) \wedge \neg\text{Dreview}_1$$

which equivalent is with $R^c_i(D\tau)$ (see definition 2). The professor will then be as well as task-based responsible as socially responsible (see definitions 4 and 6), but not causally responsible.
From his task-based responsibility it follows that

$$\text{O}(j\text{:achieve}(\text{review}_1)) \wedge [\underline{j\text{:achieve}(\text{review}_1)}]\text{Dreview}_1$$

which can determine a social responsibility of j when Dreview_1 actually occurs, when, e.g., the PhD-student has not reviewed the papers (see definition 6).

Coordination
In the coordination relation the informative actions are central. In the example, we can imagine a coordination relation between the secretary and the reviewers, that the secretary has to inform the reviewers that the deadline has been changed; instead of two weeks the papers has to be reviewed within three days. If a reviewer, say our professor *j*, has not received this information, the secretary, say *k* (so *rea*(*k*,*s*) holds) is causally responsible that the papers are not reviewed within the deadline, and so for the delay of the process for reviewing the papers (Dreview_1), which follows from definition 2:

$$[\underline{k\text{:}\mathit{inform}(j,\text{deadline})}]\text{Dreview}_1 \wedge \text{DO}(\underline{k\text{:}\mathit{inform}(j,\text{deadline})}) \wedge \neg\text{Dreview}_1$$

where *k*:*inform*(*j*,deadline) denotes that agent *k* informs *j* about the advancing of the deadline.

However, the secretary *k* is coordinationally responsible:

$$R^{coor}_k(\mathit{inform}(j,\text{deadline}))$$

since given the coordination link between the secretary and the reviewers, and the given task allocation, the following formula is true:

$$\neg K_j\text{deadline} \rightarrow \neg\text{CAN}_j(\text{achieve}(\text{review}_1)) \wedge \text{CAN}_k(\mathit{monitor}(\text{deadline}))$$

where $\text{CAN}_k(\mathit{monitor}(\text{deadline}))$ expresses that the secretary k is capable to check whether the deadline has been changed.

Control
In the control relation, the monitoring activity has a normative effect. In our scenario, a control relation can be assumed to hold between the secretary and the group of the reviewers. This means that if the professor *j* does not review the papers sent to him, secretary *k* should somehow take the task over as shown in definition 10 because of his failure-based responsibility:

$$[k\text{:}\mathit{monitor}(\text{DONE}(j\text{:}\mathit{achieve}(\text{revie}\underline{\text{w}_1})))](\text{Dreview}_1 \rightarrow \text{O}(k\text{:}\mathit{achieve}(\text{review}_1)))$$

Professor *j* would be causally responsible, task-based responsible, as well as socially responsible for Dreview_1. Notice that, instead, if also secretary *k* does not accomplish the task review_1, he would violate the obligation, and would only be causally responsible for Dreview_1.

7. Conclusions

We have provided some elementary notions of responsibility in its interconnection with the structure of an organization. The notions of responsibilities are useful in the process of designing an organizational structure, and conversely in understanding how a given organization is structured (Conte and Paolucci 2004). We have shown that responsibilities have an impact both on what agents should do within an organization, but also on whom to turn to when things go wrong.

Responsibilities are closely related to the specific task allocation within an organization. Although the task allocation can be determined dynamically through the process of delegation, some of it is predetermined through the role structure of the organization which assigns typical tasks to certain roles. The organizational structure plays an even greater role in the monitoring and control of execution of the tasks for which the agents are responsible. The logical framework we presented offers a semantics for the notions of responsibility that is necessary for determining at least some interconnections between organizational structure and responsibilities. It gives some insides into when an agent can really be held responsible for when tasks are not (or wrongly) performed. These observations might lead to guidelines for the design of an organizational structure given that one wants some responsibilities to be covered at all times. In this paper we just offered a glance of these observations through the example. However, we hope to extend this area in future work.

Another line of future work concerns the logical formalism that was used to describe the notions in this paper. Although it is sufficient to denote most of the basic properties and relations between them, we also touched upon some fundamental issues. In general responsibility is closely related to “causal” relations (cf. Lehmann 1999). I.e. who is actually causing some action or state. This is a famous problem to represent in logic and we can only hope to give a close approximation that is good enough for the present purpose.

A second related point is the use of temporal relations. In further work we will explore the dynamics of the responsibilities, their persistence and evolution over time. In order to do this we need to combine the current formalism with a temporal framework in the same vain as was done for temporal dynamic deontic logic in Dignum, Broersen, Dignum and Meyer (2004) and Broersen, Dignum and Dignum (2004).

References

J. Broersen, F. Dignum, and V. Dignum. Designing a deontic logic of deadlines. In A. Lomuscio and D. Nute, editors, Proceedings of DEON'04, pages 43–56, 2004.

R. Conte and M. Paolucci. Responsibility for societies of agents. Journal of Artificial Societies and Social Simulation, 7, 2004.

F. Dignum, J. Broersen, V. Dignum, and Meyer J-J. Ch. Meeting the deadline: Why, when and how. To be published in Proceedings of FAABS III Workshop, Washington, April 2004.

V. Dignum. A Model for Organizational Interaction. SIKS Dissertation Series, 2003.

A. Giddens. Social Theory and Modern Sociology. Polity Press, 1984.

D. Grossi, F. Dignum, L. Royakkers and M. Dastani. Foundations of organizational structure in multi-agent systems. Acceoted for AAMAS'05., 2005.

D. Grossi, F. Dignum, L. Royakkers, and J-J. Ch. Meyer. Collective obligations and agents: Who gets the blame. In A. Lomuscio and D. Nute, editors, Proceedings of DEON'04, pages 129–145, 2004.

F. Harary. Graph Theory. Addison-Wesley, London, 1969.

F. Harary, R.Z. Norman, and D. Cartwright. Structural Models: An Introduction to the Theory of Directed Graphs. John Wiley & Sons, New York, 1965.

J. Lehmann. Towards the Formalization of Legal Causal Reasoning. DEXA Workshop 1999, pages 780-784

J.-J. Ch. Meyer and R.J. Wieringa. Deontic Logic in Computer Science: Normative Systems Specification. John Wiley and sons, 1991.

O. Morgenstern. Prolegomena to a theory of organizations. Manuscript, 1951.

L. Nunes de Barros, A. Valente, V. R. Benjamins. Modeling Planning Tasks. AIPS 1996, pages 11-18.

K. A. Ross and C. R. B. Wright. Graph Theory. Prentice-Hall, New Jersey, 1992.

S. Russell and P. Norvig. Artificial Intelligence. A Modern Approach. Prentice Hall International, 2001.

P. Selznick. Foundations of the theory of organization. American Sociological Review, 13:25–35, 1948.

[1] We presuppose a distinction between two ways of intending the notion of role within an organizations: role as role-type, and role as role-token. Examples of role-types are the university roles of 'professor' or 'PhD Student'. Role-tokens are instead the specific 'professor' and 'PhD student' positions, like 'professor of x at department y' etc. The notion of roles as placeholders in the organizational activity correspond to the notion of role-token.

[2] It may be instructive to notice that these are just parametrized dynamic logic constructs which enable to represent collective agency.

[3] For comprehensive expositions of directed graph theory applied to organizations see Harary (1969), and Harary, Norman and Cartwright (1965), and Ross and Wright (1992).

[4] These ideas about the notion of plan are quite standard in the literature about planning in Artificial Intelligence (Russell and Norvig 2001). See also Nunes de Barros, Valente and Benjamins (1996).

[5] Note that the function of the numeric index j consists in denoting the position within the task allocation sequence.

[6] Note that this modelling of social harm is formally analogous to the dynamic logic reduction of deontic logic. See for example Meyer and Wieringa (1991).