Towards formal semantics for reorganization

Virginia Dignum

Frank Dignum

institute of information and computing sciences, utrecht university

technical report UU-CS-2006-060

www.cs.uu.nl

Towards formal semantics for reorganization

Virginia Dignum, Frank Dignum
email: {virginia, dignum}@cs.uu.nl

December 21, 2006

Abstract

Organization concepts and models are increasingly being adopted for the design and specification of multi-agent systems. Just like agents, organizations should also be able to adapt themselves to changing environments. In order to develop a theory on how this reorganization should be performed we need a theoretical framework in which both organizational performance as well as the reorganization itself can be described. In this paper, we present a formal model for the specification of organizations and organizational change. The model is sufficiently generic to enable the comparison of different existing organizational approaches to MAS, while having enough descriptive power to describe realistic organizations.

1 Introduction

Researchers agree that organizations allow and support an individual (be it a person, a computer system, or an institution) to recognize its role, and the roles of others, in accomplishing collective goals. Organizational Theory research recognizes that organizations are instruments of purpose, that is they are seen as coordinated by intentions and goals [18]. In similar fashion, MAS researchers realize that the specification of an organization for a MAS helps coordinating the agents' autonomous behavior [14]. In this sense, an organization can be defined as a set of entities regulated by mechanisms of social order and created by more or less autonomous actors to achieve common goals [6].

Organizations and their environments are not static. Agents can migrate, organizational objectives can change, or operational behavior can evolve. That is, as circumstances change, organizational structures must also be able to change, disappear or grow. Models for MAS must therefore not only cater for adaptive agents [15] but also be able to describe organizations dynamically adapting to changes in the environment. In fact, organizations are active entities, capable not only of adapting to the environment but also of changing that environment, which leads to the question of how and why these decisions are made in organizations [11].

The motivations for developing an abstract organizational model are twofold. In the one hand, the need for a formal, provable system that in an abstract way enables to represent organizations, their environment, objectives and agents in a way that enables their partial contributions to the performance of the organization in a changing environment. On the other hand, such a model must be realistic enough to incorporate the more 'pragmatic' considerations faced by real organizations. Most existing formal models lack this realism, either by ignoring temporal issues, by taking a very restrictive view on the controllability of agents, or by assuming complete control and knowledge within the systems (cf. [24], [23], [22]). Another important realism requirement, is the notion that agent activity has a cost, that is, choosing one or the other course of action is not only dependent on agent capabilities but the costs of the action must compare positively to its benefits. To sum up, formal models for organizations that are able to deal with realistic situations, must meet the following requirements: (1)represent notions of ability and activity of agents, without requiring knowledge about the specific actions available to a specific agent (open environments); (2) represent ability and activity of a group of agents; (3) deal with temporal issues, especially the fact that activity takes time; (4) accept limitedness of agent capability; (5) represent the notion of responsibility for the achievement of a given state of affairs; (6) represent organizational (global) goals and its dependency on agents' activities (organizational structure); (7) relate activity and organizational structure; (8) represent organizational dynamics (evolution of organization over time, changes on agent population); (9) deal with resource limitations and the dependency of activity on resources (e.g. costs); (10) deal with normative issues (representation of boundaries for action and the violation thereof).

In this paper, we will introduce a formal model for organizations that will be increasingly extended to include these requirements. That is, we start with a simple model that only implements some of the requirements and extend it successively to incorporate more complex requirements. We conclude the paper with a case study demonstrating the applicability of the proposed model.

2 Ability and Activity

The notions of agent capability and action have been widely discussed in MAS. The intuition is that an agent possesses capabilities that make action possible. In the literature, there are many approaches to the formalization of these definitions[1]. Concerning the theory of action, two main perspectives can be distinguished. The first aims at the explicit representation of action by a specific agent, in terms of dynamic logic [13], or situation calculus [17]; whereas the second is concerned with representing the fact that a certain result has been achieved, such as in the *stit* theories [21] or in the notion of agency by Elgesem [9]. In both types of approaches, the notion of action is strongly linked to that of ability. However, there is also no consensus on the meaning of ability which is taken to mean competence (the capability of making a certain proposition true), possibility (conditions are right for that activity), opportunity (both competence and possibility), or even permission (there are no prohibitions or constraints on the activity.). As these distinctions are important for organizational theory, we will develop a theory in which all these concepts can be expressed properly.

We start by defining the basic environment as a set of partially ordered worlds, each one representing a particular state of affairs that are true in that state.

Definition 1 ***(Semantic Model)***
Given a set of atomic propositions, Φ, a semantic model on Φ is defined as $M = (W, T, \pi)$, where (W, T) is a poset, and π is a valuation function which associates each $w \in W$ with the set of atomic propositions from Φ that are true in that world.

Each world describes the propositions of Φ that are true in it. Conversely, each proposition of Φ corresponds to a set of worlds where it is true. A transition between worlds represents an update of the truth value of propositions in Φ. In this sense, for each transition t in T, $t = (w_i, w_{i+1})$ we can identify the set of propositions that change value between w_i and w_{i+1}, defined as:
$t^\pi = \{p \in \Phi : \pi(w_i, p) = 1 - \pi(w_{i+1}, p)\}$.

Well-formed formulae are built out of the set of atomic propositions through combinations using the classical proposition connectives $\vee$ ('or') and $\neg$ ('not'). The language also contains the temporal operators X ('next') and $\Diamond$ ('always in the future'), the constants $true$, $false$, and the operators C_a (ability or control of agent a), IC_a (agent a is in control of a situation) and E_a, (*stit*, agent a 'sees to it that'). Semantics are given by temporal logic CTL, extended with the semantics for C_a and E_a as defined in the following subsections.

2.1 Agent control and activity

Intuitively, in order to talk about agent activity, or, that agent a possesses the ability to make φ hold in some future state in a path from the current moment, we need to establish the control of the agent over the truth value of φ. Inspired by the work of Boutelier [2], and Cholvy and Garion [3], for each agent a we partition the set of atomic propositions Φ in any world w of M in two classes:

---

[1] A concise overview can be found in [4]

the set of atomic propositions that agent a can control, C_a, and the set of atomic propositions that a cannot control, $\bar{C}_a$. In the following, we will overload the operator C_a, respectively $\bar{C}_a$, and use $C_a\varphi$, respectively $\bar{C}_a\varphi$ to represent the fact that φ is controllable by a, respectively uncontrollable by a. Intuitively, given a world w, agent control identifies a subset of the worlds that can be reached from that world by possible activity of the agent. That is, we want to select the subset of transitions that is covered by C_a. The semantic model introduced in definition 1 is extended as follows:

Definition 2 ***(Selection Function Model)***
Given a set of atomic propositions, Φ, and a set $A = \{a_1, ..., a_n\}$ of agents, a selection function model $\mathcal{E}$ is defined as $\mathcal{E} = (W, T, \pi, F)$, where (W, T, π) is a semantic model, and $F = \{f_1, ..., f_n\}$ is a set of selection functions for each agent, each with signature $W \mapsto \mathcal{P}(T)$. Given a world $w \in W$ the result of application of selection function f_i on w is a set of transitions $f_i(w) = \{t \in T : (t^\pi \cap C_{a_i}(w) \neq \varnothing) \vee (t^\pi = \varnothing)\}$, and $C_{a_i}(w) \subseteq \bigcup_{t \in f_i(w)} t^\pi$

From this definition, it follows that $C_a(w)$ cannot contain tautologies. The selection functions identify the transitions departing from the current world that are possible by some activity of that agent, that is, those worlds which occurrence is controlled at least in part by that agent. For example, consider the language with 2 words $\Phi = \{p, q\}$ and agent a such that $C_a = \{p\}$. In a semantic model for Φ each world w has 4 possible transitions $r_1, ..., r_4$ identified by the following updates (depicted in figure 1): $r_1^\pi = \{\}, r_2^\pi = \{p\}, r_3^\pi = \{p, q\}$ and $r_4^\pi = \{q\}$. The selection function f_a selects those transitions that are covered by the set of propositions controlled by a, C_a, that is, $f_a(w) = \{r_1, r_2, r_3\}$. Evolution of the current world over these transitions is possible by some action of a.

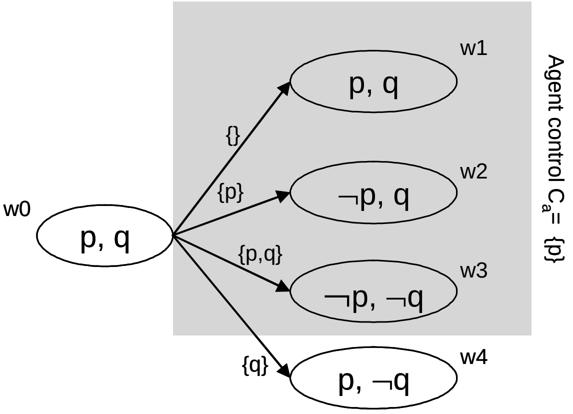

Figure 1: Scope of agent control (example).

A semantic selection function model $\mathcal{E}$ is a selection model that satisfies the following condition[2]:

(S1) $w \models_\mathcal{E} C_a\varphi$ iff $\exists w' : (w, w') \in f_a(w) : w' \models_\mathcal{E} \varphi$

(S1' new!) $w \models_\mathcal{E} C_a\varphi$ iff

1. $\exists w' : (w, w') \in f_a(w) : w' \models_\mathcal{E} \varphi$ and
2. $\exists w'' : (w, w'') \in f_a(w) : w'' \models_\mathcal{E} \neg\varphi$ and
3. if $w \models_\mathcal{E} \varphi$ then $\forall w_i : (w, w_i) \notin f_a(w), w_i \models_\mathcal{E} \varphi$ and
4. if $w \models_\mathcal{E} \neg\varphi$ then $\forall w_i : (w, w_i) \notin f_a(w), w_i \models_\mathcal{E} \neg\varphi$

---

[2] Plus the usual semantics of temporal logic.

As an example of control over formulas, consider again the example in figure 1 where $p \in C_a$ and $q \in \bar{C}_a$. In this case, proposition $p \wedge q$ is not controllable by a, because q is not a-controllable, and if q is false, agent a can never make $p \wedge q$ to become true.

In the special case in which all next possible states from a given state are influenced by an agent a, we say that a is ***in-control*** in w, represented by IC_a and defined formally as:

(S2) $w \models_{\mathcal{E}} IC_a$ iff $\forall t : t = (w, w')t \in f_a(w)$

The notion of controllability expresses the fact that an agent a has the capability to bring φ about. It does not mean that the agent will indeed ever do it. Obviously, agents must be able to *act* on the world and as such bring states of affairs to happen. Because we abstract from the internal motivations and capabilities of individual agents, we need ways to describe the result of agent action that are independent of particular actions available to the agent. The *stit* operator, $E_a\varphi$ ('agent a sees to it that φ), introduced by Pörn [21] allows to refer to the externally 'observable' consequences of an action instead of the action itself. *Stit* can be seen as an abstract representation of the family of actions that result in φ. In our approach we refine the definition of *stit* to include a temporal component that indicates the notion that action takes time. In the following we give our semantics of stit, in terms of a semantic selection function model $\mathcal{E}$ as introduced in definition 2.[3]

Definition 3 ***(Semantics of agent action)***
We extend the conditions of a model a semantic selection model $\mathcal{E}$ to describe the semantics of $E_a\varphi$ as follows:

(S3) $w \models_{\mathcal{E}} E_a\varphi$ *iff* $\forall w', t = (w, w') : t \in f_a(w) \wedge w' \models_{\mathcal{E}} \varphi$

$E_a\varphi$ represents the actual action of bringing φ to be. From the semantics above: if φ is controllable by a then $E_a\varphi$ causes φ to be true in all next controllable states from the current state (that is, states selected by the selection function for a). Furthermore, it is important to notice that we provide a temporal definition of E_a which differs from other authors, in particular [22].

2.2 Group Control and Activity

Agents are limited on their capabilities, that is, on what parameters of the world they can control. This implies that in MAS certain states of affairs can only be reached if two or more agents cooperate to bring that state to be. One of the main ideas behind organizations is the notion that the combined action of two or more agents can result in an effect that none of the involved agents could bring about by themselves. In MAS there is very little research done on the notion of ability in a multi-agent context [4]. In the following, we define the concept of combined control, or *group control*. As for single agents, we need to start by defining the notion of group control over atomic propositions. The idea behind this definition, is that atomic propositions can be made 'small' enough to be controlled by a single agent. A group combines results from different agents.

Definition 4 ***(Group control over atomic propositions)*** *Given a set of agents $G = \{a_1, ..., a_n\} \subseteq A$ and the sets C_{a_i} of atomic propositions controllable by each agent $a_i \in G$, we define C_G as the union of the controllable propositions by all agents in G:*

$$C_G = \bigcup_{i=1}^{n} C_{a_i}$$

In the same way, we define $\bar{C}_G$ as the set of atomic propositions that group G cannot control. Given a semantic selection function model $\mathcal{E}$, a groups' selection function f_G is defined as:

---

[3] This notion of stit provides a *necessary* interpretation of action, and as such is related to the dynamic operator $[a]p$, meaning that after performing action a it is necessarily the case that p. The definition of a *possible* interpretation for stit will be part of future work.

$f_G(w) = \{t \in T : (t^\pi \cap C_G(w) \neq \varnothing) \vee (t^\pi = \varnothing)\}$ and $C_G(w) \subseteq \bigcup_{t \in f_G(w)} t^\pi$

Now we are able to define a semantic model $\mathcal{E}$ that incorporates the notion of group control as:

Definition 5 ***(Group Control)***
An expression φ is controlled by a group of agents $G = \{a_1, ..., a_n\}$, represented by $C_G\varphi$ or $C_{\{a_1,...,a_n\}}\varphi$,

(S4) $w \models_{\mathcal{E}} C_G\varphi$ *iff* $\exists w', t = (w, w'), t \in f_G(w) : w' \models_{\mathcal{E}} \varphi$

Furthermore, an expression is G-uncontrollable iff it is not G-controllable.

Consider again the example in figure 2, in which $(p \wedge q)$ was not a-controllable for an agent a such that C_ap. Now, suppose that there is an agent b such that C_bq. In the same way, $\neg C_b(p \wedge q)$. However, if we consider the group $G = \{a, b\}$ we get $C_G(p \wedge q)$, that is, together a and b have control over $(p \wedge q)$. This derives from the fact that function $f_{\{a,b\}}(w)$ completely covers all transitions from a world where $(p \wedge q)$ holds.

Note that a-controllable is a special case of group control when $G = \{a\}$. That is, all expressions controlled by one agent are also controlled by all groups in which that agent participates. This is different from the assumption made by e.g. [24],[23] who use group control to refer to those formulas that are only controlled by the whole group, (and not controlled by any of its subgroups). In our model, this is a special case of the definition 5 above, as follows:

Proposition 1 ***(Joint control)***
An expression φ is said to be joint controlled *by a group $G \subseteq A$ of agents iff $C_G\varphi$ and $\forall G' \subset G, \neg C_{G'}\varphi$*

Finally, we extend the definition of *stit* for groups: E_G (i.e. 'group G sees to it that'). The result of the combined activity of a group of agents G is defined as follows:

Definition 6 ***(Semantics of Group Activity)***
Given a semantic model $\mathcal{E}$ and a group G of agents, $G \subseteq A$, $E_G\varphi$ is defined as:

(S5) $w \models_{\mathcal{E}} E_G\varphi$ *iff* $\forall w', t = (w, w'), t \in f_G(w) : w' \models_{\mathcal{E}} \varphi$

A logic of ability and action for groups of agents has semantics given by a model $\mathcal{E}$ that satisfies **S1-S5** above, and the usual semantics of temporal logic.

2.3 Organization Control and Activity

The objectives of an organization are achieved through agent action. In order to make this possible, an organization must employ the relevant agents, so that it can 'enforce' the possibility of making its desires happen. Given a $\mathcal{E}$ model, an *organization O* is defined by a set of agents, a set of objectives (missions or desires), and a set of assets. Agents are the active entities that realize organizational activity. Organizational objectives are the issues that the organization 'wishes' to be true in the world. Organizational assets are the issues that are true (and relevant to the organization) at a given moment. Note that, for the purposes of this paper, we see agents purely as actors in a organization, with no goals of themselves. We assume that, by acting according to their capabilities, agents work towards organizational objectives[4]. Formally, an organization is defined as:

Definition 7 ***(Organization)***
Given a semantic model $\mathcal{E} = (W, T, \pi, F)$, an organization is defined by $O^w = \{A^w_O, D^w_O, S^w_O\}$ where $A^w_O = \{a_1, ..., a_n\}$, $D^w_{Org} \subseteq \Phi$, $S^w_O \subseteq \Phi$ and $w \in W$.

---

[4] That is, we abstract here from the motivation an individual agent may have to take up those organizational positions. More on this issue can be found in [5]).

The current state of the organization, S_O^w, corresponds to the set of formulas that are true in world w and relevant to O, and the objectives or desires of the organization, D_O^w, characterize the worlds (sets of formulas) that, at this moment, the organization wishes to reach. Note that, organizational change means that the organization's composition (agents) and objectives may differ from world to world. More on this in section 4. (From now on, whenever clear from the context, we'll drop the subscripts and superscripts).

Based on the definitions given in the previous sections, we are now in state to define organization control. In fact, an organization is only as good as its agents. That is, the scope of organizational control is defined by the union of the scopes of its agents' control together with the control of groups of its agents. Formally,

Definition 8 ***(Organization Control)***
Given a organization O such that A_O is the set of agents in O, organizational control C_O is defined as: $C_O = C_{A_O}$. Thus, $C_O\varphi$ iff $\exists G \subseteq A_O : C_G\varphi$.

3 Towards more realistic Organizations

In the previous section, we have presented a fairly simple model for organizations, which only considers the current organizational state, the agents participating in the organization, and its objectives. In particular, that model of organization does not indicate how organizational objectives are to be achieved. One of the main reasons for creating organizations is efficiency, that is, to provide the means for coordination that enable the achievement of global goals in an efficient manner. In the model above, even if the agents in the organization have group control over all organizational objectives, they have no means to coordinate their activities in order to efficiently achieve those objectives. In most cases, the objectives of the organization are only known to a few of the agents in the organization, who may have no control over those objectives. It is therefore necessary to organize agents in the organization in a way that enables objectives to be passed to those agents that can effectively realize them. Organization structure is intended to describe and support distribution of work necessary to attain organizational objectives. Tasks or objectives are allocated to different organizational positions (or roles), that are enacted by individuals through the assignment of responsibilities [6]. The above considerations suggest that we need to extend our notion of organization to include some representation of interaction structure, and that we need to be able to refer to the responsibilities of agents in the organization. We therefore extend our initial definition of organization by introducing a relation between the agents in set A_O.

Definition 9 *(Structured Organization)*
Given a semantic model $\mathcal{E} = (W, T, \pi, F)$, a structured organization is defined by $O^w = \{A_O^w, \leq_O^w, D_O^w, S_O^w\}$ where $A_O^w = \{a_1, ..., a_n\}$, $(A_O^w, \leq_O^w)$ is a poset with the partial order relation reflecting the structure of the organization, $D_O^w \subseteq \Phi$, $S_O^w \subseteq \Phi$ and $w \in W$.

Intuitively, the ordering relation in the set of agents stands for the interaction possibilities between agents. Organizational structures influence the way that agents in the organization can interact. The relation $a \leq_O b$ indicates that a is able to interact with b in order to request or demand some result from b. In this paper, we will not further detail the types of interactions between agents (delegation, request, bid, ...) but assume that the relationship will achieve some result, through a more or less complex interaction process. More on this issue can be found in [7]

In order to describe the notion of responsibility and delegation as described above, we introduce two new operators, H_a and R_a. Informally the meaning of $H_a\varphi$ is that a attempts (possibly without success) to achieve φ, and $R_a\varphi$ means that a is responsible for, or in charge of, achieving φ. In the following, we will draw from the work of Santos et al.[22] and Governatori et al. [12]. In short, their work assumes that in organizations not all capabilities are always conductive of successful action - one can delegate or request something from someone else but this does not say anything about successful result. Using our formalism we are able to semantically define attempted activity as follows: Using our formalism we are furthermore able to semantically define attempted activity as follows:

Definition 10 ***(Attempt)***
Given a selection function $\mathcal{E}$ as before, the semantics of $H_a\varphi$ are defined as:

(S6) $w \models_{\mathcal{E}} H_a\varphi$ *iff* $\forall t \in f_a(w)$ *if* $t = (w, w')$ *then* $w' \models_{\mathcal{E}} \varphi$

Note that we can now, redefine the E_a operator in terms of the H_a operator and IC_a as follows: $E_a\varphi \equiv H_a\varphi \wedge IC_a$.

That is, in worlds where agent a is in-control, an attempt to achieve φ is always successful. We are now ready to formally define responsibility of agent a for making it the case that φ, $R_a\varphi$. Informally, by responsibility we mean that an agent or group has to make sure that a certain state of affairs is achieved, either by realizing it itself or by delegating that result to someone else. That is, a is given the responsibility for φ but one has no guarantee that a will be successful.

Definition 11 ***(Responsibility)***
Given an organization $O = (A_O, \leq_O, D_O, S_O^w)$ in a model $\mathcal{E}$, and an agent $a \in A_O$, responsibility $R_a\varphi$ is such that:
$R_a\varphi \equiv \Diamond H_a\varphi \vee H_a R_b\varphi$*, for some* $b \in A_O$

Furthermore, responsibility can be divided between several agents, That is, if $\varphi = \varphi_1 \wedge ... \wedge \varphi_n$ and $G = \{b_1, ..., b_n\}$ is a group of agents: $H_a R_G\varphi \equiv H_a R_{b_1}\varphi_1 \wedge ... \wedge H_a R_{b_n}\varphi_n$.

In a structured organization, order relations between agents indicate how agent can interact in order to specify responsibility to make it the case that a result is achieved. Formally,

Definition 12 ***(Structured delegation)***
Given a semantic model $\mathcal{E}$ and a structured organization $O = (A_O, \leq_O, D_O, S_O^w)$ in $\mathcal{E}$, delegation of φ between two agents $a, b \in A$ is defined as: $(a \leq_O b) \rightarrow C_a R_b\varphi$.

The responsibility operator is also defined for a group G of agents, R_G in a similar way. In the following, we provide an example of structured organization. Consider the organization $O = ((A, \leq_O), D, S^0)$, where:
$A = \{a, b, c, d\}$,
$\leq_O$ is such that $(a \leq_O b, a \leq_O c \leq_O d))$,
$D = \{\rho\}$, where $\rho = (p \wedge q) \vee r$,
$S^0 = \{R_a\rho, C_b p, C_d q, C_{\{a,c\}} r\}$

Note that the initial organizational state S^0 indicates the capabilities of the agents in A and that agent a is responsible for the achievement of the organizational goal. This example, also shows that organizations are dependent on the capabilities of their agents to achieve their objectives. In this case, without agent b, the organization can never achieve its goals. In the next section, we will discuss the issue of reorganization that enables organizations to modify their own characteristics. In the above organization, there are several ways for agent a to realize the organizational goal ρ, for example:

s_1: $E_a R_b p \wedge E_a R_c q$
... .
s_i: $R_b p \wedge E_c R_d q$
s_{i+i}: $R_b p \wedge R_d q$
... .
s_N: $p \wedge q$

Different properties can be defined for the responsibility operator, which identify different types of organizations. For example, a *well-defined organization* is such that $\forall\varphi \in D_O, \exists a \in A_O : R_a\varphi$. Or, in a *lazy bureaucratic organization* where everybody only acts if 'forced', the property $E_a\varphi \rightarrow R_a\varphi$ holds.

4 Organizational change

In order to keep effective, organizations must maintain a good fit with the environment. Changes in the environment lead to alterations on the effectiveness of the organization and therefore in a need to reorganize, or in the least, the need to consider the consequences of the change to the organization's effectiveness and possibly efficiency. On the other hand, organizations are active entities, capable not only of adapting to the environment but also of changing that environment. This means that organizations are in state of, to a certain degree, altering environment conditions to meet their aims and requirements, which leads to the question of how and why reorganization decisions should be reached. *Adaptation* in Organizational Theory literature means different things ranging from strategic choice to environmental determinism. Our concept of organizational change is more related to the first meaning, in the sense that we treat adaptation as a *design* issue that requires an (explicit) action resulting in the modification of some organizational characteristics, whereas the latter meaning refers to the *emergence* of organizational patterns. Moreover, decisions for adaptation are of two kinds: proactive, preparing the organization in advance for an expected future, and reactive, making adjustments after the environment has changed [8].

In terms of the formal model of organizations introduced in the previous sections, changes in the environment are represented as (temporal) transitions between two different worlds. Given a world $w \in W$, many different events may happen that change some proposition in that world resulting in a different world (the relation T between two worlds represents this). Because not all parameters in w are O-controllable, the current state of the organization is not necessarily, and in fact in most cases not, completely controlled by the organization itself. That is, changes are not always an effect of (planned) organizational activity. We therefore distinguish between exogenous and endogenous change. In general, at any given moment, an organization cannot fully control its state nor the way it will evolve. Formally this is represented by: $S_O \neq C_O$, or $\exists\varphi \in S_O : \neg C_O\varphi$. Some parameters may be controlled by external parties, but others may be completely uncontrollable (disasters, natural forces, etc.). Those external agents which can control parameters that are part of the organization state are said to have an *influence* on the organization. Figure 2 shows the relation between organization desires and states, and controllable parameters both the current situation, and the ideal situation.

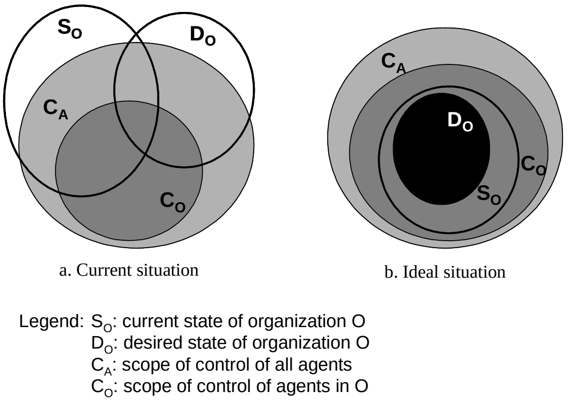

a. Current situation

b. Ideal situation

Figure 2: States and control.

Organizations are constantly trying to identify the optimal design for their environment, and will choose a change strategy that they believe will improve their current situation. In an ideal successful organization the set of desires D_O will be a subset of the organizational state S_O, which is under the control of the agent in the organization. Reorganization activities aim therefore at aligning these sets, either by attempting to change the current state, or by altering the set of desires. In summary, reorganization consists of two activities. Firstly, the formal representation and evaluation of current organizational state and its 'distance' to desired state, and, secondly, the formalization of reorganization strategies, that is, the purposeful change of organizational constituents (structure, agent population, objectives) in order to make a path to desired state possible and efficient.

4.1 Monitoring for change

Organizations and environments are closely related, and different organizations will fare better or worse in different environments. There is no one best way to organize or structure the organization, but not all structures are equally effective, that is, organizational structure is one determinant of organizational performance. Performance of the organization can then be seen as the measure to which its objectives are achieved at a certain moment. Because environments evolve, performance will vary. Many organizational studies are therefore concerned with the evaluation of performance, identifying triggers for change and determine the influence of environment change in the organizational performance and indicating directions to improve performance.

Our aim is to be able to specify performance, based on the model introduced in the previous sections. Intuitively, performance is a value function on the environment (current world), agents and organizational capability, and on the task (desired state of affairs). Formally, we define a function *perform* with signature $W \times \mathcal{P}(A_{\leq}) \times \Phi \mapsto \mathbb{R}$, such that $perform(w, G_{\leq}, \varphi)$ returns the value of the performance in world w of structured group $G_{\leq}$ for φ, indicating how well G can realize φ. We assume that for each agent and each world, the performance for each atomic proposition p is fixed. That is, $\forall w, a, p, \exists c \in \mathbb{R} : perform(w, a, p) = c$. The function *perform* has the following properties:

- $perform(w, G_{\leq}, p) + perform(w, G_{\leq}, q) \leq$
  $perform(w, G_{\leq}, p \wedge q)$, for atomic propositions p, q
- $\neg C_{G_{\leq}} \varphi \rightarrow (perform(w, G_{\leq}, \varphi) = \infty)$
- $perform(w, G_{\leq}, \varphi) \leq perform(w, G_{\leq}, \varphi \wedge \psi)$
- $(\varphi \rightarrow \psi) \rightarrow (perform(w, G_{\leq}, \psi) \leq perform(w, G_{\leq}, \varphi))$

The *perform* function can be seen as the cost associated with a transition in the world W. Informally, the first property above, represents the fact that an agent can get ‘tired’. That is, the cost of performing a sequence of activities can be higher than the sum of the costs associated with each activity, for that agent [4]. Some authors have used finite state machines to describe organization transitions [19]. An important difference between our work and such approaches is the temporal dimension of our model, that is, an organization can never move to a previous state even if conceptually equivalent, all states are different in time.

The function *perform* provides a way to establish the cost of achievement of a certain state of affairs, given the current state and group of agents. Based on this function, different strategies can be adopted that determine the course of action to follow in order to choose a organization with better performance with respect to a given goal. In the next subsection, we describe the actions that can be taken in order to effectively change the organization.

4.2 Acting for change

In this section, we deal with planned reorganization. That is, we want to specify the activities that are consciously taken by an organization in order to attempt to modify its characteristics. Note that, under planned reorganization, we consider both *endogenous* reorganization, that is, the reorganization is result of activity by the agents themselves (and thus realized in run-time), as *exogenous* reorganization, in which reorganization is achieved by activity outside of the system, for example by the designer (and thus off-line).

In human organizations, reorganization strategies take different forms, such as hiring new personnel, downsizing, or reassign tasks or personnel.Organizations can also choose to modify their mission or objectives. Studies suggest that such second-order change, that is, a shift from one strategic orientation to another, is atypical even in times of highly dynamic environmental behavior [16]. Because organizations aim at making certain states of affairs to be the case, and only agents can bring affairs to be, it is important for the organization to make sure it ‘hires’ and organizes an adequate set of agents $(A_O, \leq_O)$ such that the combined action of those agents has the potentiality to bring about the desired state of affairs D_O. The dependency relation $\leq_O$ between the agents must allow for the desired states to be achieved, that is, dependencies must be sufficient for responsibilities to be passed to the appropriate agents, that is, the agents that have the necessary capabilities. If that is not the case, the organization should take the steps needed to decide and implement reorganization, such that the resulting organization O' is indeed able to realize its objectives $D_{O'}$. In practice, reorganization activities can be classified in three groups:

- **Staffing**: Changes on the set of agents: adding new agents, or deleting agents from the set. Corresponding to personnel activities in human organizations (hiring, firing and training).
- **Structuring** Changes on the ordering structure of the organization. Corresponding to infrastructural changes in human organizations: e.g. changes in composition of departments or positions.
- **Strategy** Changes on the objectives of the organization: adding or deleting desired states. Corresponding to strategic (or second-order) changes in human organizations: on the mission, vision, or charter of the organization.

In our model, we formally define the reorganization activities above as follows:

Definition 13 ***(Reorganization Operations)***
Given an organization $O = (A, \leq, D, S)$*, an agent* a*, such that* $C_a\varphi$*, and a semantic model* $\mathcal{E} = (W, T, \pi, F)$ *for* O*, such that* $m, m' \in M$ *and* $(w, w') \in T$*, the reorganization operations over* O *in* M *are:*

- *If* $a \notin A : w \models_{\mathcal{E}}$ *staff*$^+(O, a)$ *iff* $\forall w' : (w, w') \in T$,
  $w' \models_{\mathcal{E}} O'$*, where* $O' = (A \cup \{a\}, \leq, D, S \cup C_a\varphi)$,
  (Note that $(O', w') \models_{\mathcal{E}} C_a\varphi$ *holds)*
- *If* $a \in A : w \models_{\mathcal{E}}$ *staff*$^-(O, a)$ *iff* $\forall w' : (w, w') \in T$,
  $w' \models_{\mathcal{E}} O'$*, where* $O' = (A/\{a\}, \leq, D, S/C_a)$,
- *If* $a, b \in A : w \models_{\mathcal{E}}$ *struct*$^+(O, (a \leq b))$ *iff* $\forall w' : (w, w') \in T$,
  $w' \models_{\mathcal{E}} O'$*, where* $O' = (A, \leq \cup\{(a \leq b)\}, D, S)$,
- *If* $a, b \in A$ *and* $(a \leq b) \in \leq$: $w \models_{\mathcal{E}}$ *struct*$^-(O, (a \leq b))$ *iff* $\forall w' : (w, w') \in T, w' \models_{\mathcal{E}} O'$*, where* $O' = (A, \leq /\{(a \leq b)\}, D, S)$,
- *If* $\neg(d \wedge D) \rightarrow \perp$: $w \models_{\mathcal{E}}$ *strateg*$^+(O, d)$ *iff* $\forall w'$:
  $(w, w') \in T, w' \models_{\mathcal{E}} O'$*, where* $O' = (A, \leq, D \cup \{d\}, S)$
- *If* $d \in D$: $w \models_{\mathcal{E}}$ *strateg*$^-(O, d)$ *iff* $\forall w'$:
  $(w, w') \in T, w' \models_{\mathcal{E}} O'$*, where* $O' = (A, \leq, D/d, S)$

The above definition gives a fairly simple (naive) description of model updates. This specially the case for the strategic reorganization operations. The solution chosen above is to specify that strategic decisions cannot be realized if they yield in a contradiction. In reality, a more elaborate definition, must consider belief revision problems that result from the addition and/or deletion of model components. This is an issue for future research and will not be further discussed in this paper.

The operations described above enable the reorganization of an agent organization. The issue of deciding about reorganization still remains. That is, how do organizations reach the decision to reorganize? What should then be reorganized? When should one reorganize? An informal foundation for reasoning about dynamic reorganization is given in [8]. In section 4.3 we will further discuss this issue.

Reorganization operations are just operations, that is, effect the value of some variables. That is, reorganization can also be either endogenous or exogenous. In the *exogenous* case, reorganization lays outside the control of any agent in the world, and is often realized by action of the system designer. In the *endogenous* case, agents are able to achieve the states resulting from reorganization operations (that is, can control reorganization results). In this case, we can specify an agent a such that $C_a\sigma$ where σ is one of the reorganization results specified above. For example, the fact that agent a is able to hire other agents is represented by $C_a \mathit{staff}^+(O, b)$. We can thus cover both the concept of engineered reorganization (exogenous) and dynamic reorganization (endogenous). Furthermore, as discussed in [8, 1], in the endogenous case, reorganization can either be *collaborative*, that is, there is a group of agents $G \subseteq A$, such that $C_G\rho$, or role-based, that is, a single agent controls the reorganization result, $C_a\rho$.

4.3 Deciding about change

A reorganization strategy should take into account the current performance and determine which characteristics of the organization should be modified in order to achieve a better performance. The general idea behind reorganization strategies, is that one should be able to evaluate the utility of the current state of affairs (that is, what happens if nothing changes), and the utility of future states of affairs that can be obtained by performing reorganization actions. The choice is then to choose the future with the highest utility. By applying the $perform$ function defined in section 4.1, one is able to describe the cost of reorganization. This can be used to decide about when to reorganize. A possible strategy to decide whether to realize reorganization operation σ is if $perform(w, O, \sigma) + perform(w', O', D_O) \leq perform(w, O, D_O)$, where O' is the organization in w' resulting from the realization of σ on organization O in w. Informally, this strategy says that if the cost of reorganization plus the cost of achieving D_O by the reorganized organization is less than the cost of achieving D_O without reorganizing, then reorganization should be chosen.

5 Cases

We have applied the model to formally describe several existing agent systems that simulate organizational adaptation. Cases were taken from different domains, using different frameworks and were developed independently by other research groups. Due to space limitations, we will only discuss here one of the studied cases.

5.1 Gas Supply Chain

Economic reforms from the last decade have impacted large reorganization of supply chain processes. The opening of traditional monopolist markets to competition is an example of such reorganizations. Agents have often be advocated for the modelling of supply chain processes [10]. In this section, we demonstrate the application of our formal reorganization model to the supply chain domain by modelling a simplified version of the Dutch gas pipeline transport market. Our model is roughly based on the work described by Pelletier et al. [20].

The classic model for the gas market was a simple linear value chain, controlled by a (local) *monopolist* that organized and direct the transmission of gas from producer to consumer by controlling and directing all parties involved, namely *trader* (responsible for exchanging gas from wellhead producer to consumer), *shipper* (responsible for transport of high-pressure gas from origin to destination) and *local transport manager* (responsible for the overall capacity and flow of gas). This situation is formally modelled by:

$O = ((\{m, t, s, l\}, \{m \leq_O t, m \leq_O s, m \leq_O l\}), D_O, S_O\}$ where:

- agent m represents the monopolist
- agent t represents the trader, C_t *buy-gas*
- agent s represents the shipping, C_s *transport-gas*
- agent l represents the local transport manager, C_l *local-flow*
- $\leq_O$ specifies the power of m to delegate tasks to the other partners in the supply chain, t, s, l
- $D_O = \mathit{buy\text{-}gas} \wedge \mathit{transport\text{-}gas} \wedge \mathit{local\text{-}flow}$ is the objective of getting gas to the end-user
- $S_O = \{R_m D_O, C_t \mathit{buy\text{-}gas}, C_s \mathit{transport\text{-}gas}, C_l \mathit{local\text{-}flow}\}$

It is straightforward to see that this organization is in state of realizing its objective of supplying gas to its costumer. However, as all monopolies, the process is fully determined and controlled by the monopolist and the other agents are not able to make any agreements between themselves outside the control of the monopolist. Following the political decision on the liberalization of the gas market in the Netherlands, the role of the monopolist disappears and the possibility for the other partners to directly contract each others and for multiple parties to enter the market is created. By removing the monopolist agent m in the formal model above, it can easily be proven that the organization is no longer in state of realizing its objectives: even though the remaining agents have sufficient capabilities to achieve the organization objectives, they lack the possibilities to coordinate their activity. There is therefore the need to extend the agents structuring capabilities. Formally, the liberalization process can be represented by the following exogenous reorganization operations:

$$\begin{array}{lll}
staff^-(O, M, m), & & \\
staff^-(O, M, t), & & \\
staff^-(O, M, s), & & \\
staff^-(O, M, l), & & \\
staff^+(O, M, t') & \text{such that:} & C_{t'} \mathit{buy\text{-}gas} \wedge \\
 & & C_{t'} \mathit{struct}(t' \leq s') \wedge \\
 & & C_{t'} \mathit{struct}(t' \leq l') \\
staff^+(O, M, s') & \text{such that:} & C_{s'} \mathit{transport\text{-}gas} \wedge \\
 & & C_{s'} \mathit{struct}(s' \leq t') \wedge \\
 & & C_{s'} \mathit{struct}(s' \leq l') \\
staff^+(O, M, l') & \text{such that:} & C_{l'} \mathit{local\text{-}flow} \wedge \\
 & & C_{l'} \mathit{struct}(l' \leq t') \wedge \\
 & & C_{l'} \mathit{struct}(l' \leq s')
\end{array}$$

That is, the monopolist is removed and the renewed trader, shipping and local manager agents get the capability of creating dependencies to other agents (in the list above, *struct* represents both struct^+ and struct^-), representing the fact that they can negotiate service contracts. The resulting, reorganized organization, representing the liberalization of the market, is formally given by:

$$O' = ((\{t', s', l'\}, \{\}), D_O, S'_O),$$

where $S'_O = \{(R_{t'} D_O \vee R_{s'} D_O \vee R_{l'} D_O)\} \cup C_{t'} \cup C_{s'} \cup C_{l'}$ and the organization's objective is the same as above. Endogenous reorganization activities (i.e. the possibility for the different agents to realize *struct* operations) enable runtime adaptation. In this organization, each of the agents can take the responsibility to realize the organization's objective and uses its structuring capabilities to establish the necessary relations with the other agents. For example, the local manager can now take the initiative to organize gas distribution in its area, by performing the following operations:

$E_{l'} \mathit{struct}^+(l' \leq t')$
$E_{l'} \mathit{struct}^+(l' \leq s')$
$E_{l'} R_{t'} \mathit{buy\text{-}gas}$
$E_{l'} R_{s'} \mathit{transport\text{-}gas}$
$E_{l'} \mathit{local\text{-}flow}$

6 Conclusions

Dynamic reorganization of agent systems is needed in order to enable systems to enforce or adapt to changes in the environment. This issue has been discussed by many researchers and several domain-oriented solutions have been proposed. However, such solutions often lack a formal basis. This prohibited the development of theories about reorganization and it prevented comparison or adaptation to other domains or situations. In this paper we presented a first attempt at a formal model for organizational concepts and the reorganization process, based on modal temporal logic. We have applied the proposed model to existing domain-specific systems proposed in recent literature. The current model is based on the notions of controllability, *stit*, attempt and responsibility. In the future, we will extend the model to include deontic concepts and their relation to the operational concepts presented in this paper. We are currently engaged in the full axiomatization of the logic presented here as well as an implementation for simulating the reorganization process.

Acknowledgements. This work has been partially done during a sabbatical of the authors at the University of Melbourne. The research of the first author is funded by the Netherlands Organization for Scientific Research (NWO), through Veni-grant 639.021.509.

References

[1] K. S. Barber and C. E. Martin. Dynamic reorganization of decision-making groups. In *Proceedings of the 5th Autonomous Agents*, 2001.

[2] C. Boutelier. Toward a logic for qualitative decision theory. In J. Doyle, E. Sandewall, and P. Torasso, editors, *Principles of Knowledge Representation and Reasoning (KR'94)*, 2004.

[3] L. Cholvy and C. Garion. An attempt to adapt a logic of conditional preferences for reasoning with contrary to duties. *Fundamenta Informaticae*, 47, 2001.

[4] L. Cholvy, C. Garion, and C. Saurel. Ability in a multi-agent context: a model in the situation calculus. In *Proceedings of Sixth International Workshop on Computational Logic in Multi-Agent Systems (CLIMA VI)*, 2005.

[5] M. Dastani, V. Dignum, and F. Dignum. Role assignment in open agent societies. In *Proceedings of AAMAS03, Second International Joint Conference on Autonomous Agents and Multi-agent Systems*. ACM Press, July 2003.

[6] V. Dignum. *A Model for Organizational Interaction: based on Agents, founded in Logic.* SIKS Dissertation Series 2004-1. Utrecht University, 2004. PhD Thesis.

[7] V. Dignum and F. Dignum. Coordinating tasks in agent organizations. or: Can we ask you to read this paper? In *Workshop on Coordination, Organization, Institutions and Norms in Agent Systems (COIN@ECAI'06)*, 2006.

[8] V. Dignum, F. Dignum, and L. Sonenberg. Towards dynamic organization of agent societies. In G. Vouros, editor, *Workshop on Coordination in Emergent Agent Societies, ECAI 2004*, pages 70–78, 2004.

[9] D. Elgesem. The model logic of agency. *Nordic Journal of Philosophical Logic*, 2(2):1–46, 1997.

[10] Mark S. Fox, Mihai Barbuceanu, and Rune Teigen. Agent-oriented supply-chain management. *International Journal of Flexible Manufacturing Systems*, 12(2–3):165–188, 2000.

[11] Henk Gazendam and John L.Simons. An analysis of the concept of equilibrium in organization theory. In *Proceedings of the Computational and Mathematical Organization Theory Workshop*, 1998.

[12] Guido Governatori, Jonathan Gelati, Antonino Rotolo, and Giovanni Sartor. Actions, institutions, powers. preliminary notes. In G. Lindemann et al., editor, *Proceedings of RASTA'02*, volume 318 of *Mitteilung*, pages 131–147. Fachbereich Informatik, Universitt Hamburg, 2002.

[13] D. Harel. Dynamic logic. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic*, volume 2, pages 497–6049. Reidel.

[14] Jomi Fred Hübner, Jaime Simo Sichman, and Olivier Boissier. S-moise+: A middleware for developing organised multi-agent systems. In Olivier Boissier et al., editor, *Coordination, Organization, Institutions and Norms in Agent Systems I*, volume 3913 of *LNAI*, pages 64–78. Springer, 2006.

[15] N. R. Jennings, K. Sycara, and M. Wooldridge. A roadmap of agent research and development. *JAAMAS*, 1(1):7–38, 1998.

[16] J. Hunt K. Boal, S. Fox-Wolfgramm. Organizational adaptation to institutional change: A comparative study of first-order change in prospector and defender banks. *Administrative Science Quarterly*, 43, 1998.

[17] J. MacCarthy and P. Hayes. Some philosophical problems from the standpoint of artificial intelligence. *Machine Intelligence*, 4:463–502, 1969.

[18] J. March and R. Sutton. Organizational performance as a dependent variable. *Organization Science*, 8(6):698–706, Nov. Dec. 1997.

[19] Eric Matson and Scott A. DeLoach. Formal transition in agent organizations. In *IEEE International Conference on Knowledge Intensive Multiagent Systems (KIMAS'05), year = 2005, location =*.

[20] C.M.P. Pelletier, C. de Snoo, and L. Maruster. Modelling the gas transport supply chain with an agent-based approach. In *Proceedings of the Supply Chain Management and ICT Workshop*, 2005.

[21] I. Pörn. Some basic concepts of action. In S. Stenlund, editor, *Logical Theory and Semantical Analisys*. Reidel, 1974.

[22] Filipe Santos, Andrew Jones, and Jose Carmo. Action concepts for describing organised interaction. In Jr. R.A. Sprague, editor, *Proc. of the Thirtieth Annual Hawaii International Conference on System Sciences*, volume V, pages 373–382. IEEE Computer Society Press, 1997.

[23] Luigi Sauro, Jelle Gerbrandy, Wiebe van der Hoek, and Michael Wooldridge. Reasoning about action and cooperation. In Gerhard Weiss Peter Stone, editor, *Proceedings of AAMAS-06*. ACM, 2006.

[24] Michael Wooldridge Wiebe van der Hoek. On the logic of cooperation and propositional control. *Artificial Intelligence*, 24(1-2):81–119, 2005.