Policy Management for Virtual Communities of Agents

Guido Boella
Dipartimento di Informatica
Università di Torino
Italy
E-mail: guido@di.unito.it

Leendert van der Torre
CWI
Amsterdam
The Netherlands
E-mail: torre@cwi.nl

***Abstract*— In this paper we study the rational balance between local and global policies in virtual communities of agents. To study this problem we use a logical framework for modelling obligations and permissions in multiagent systems. In particular, the logical framework allows agents to trade off the decision of respecting a norm against the consequences of not respecting it: the possibility that they are considered violators and thus sanctioned. To formalize decision making we use a qualitative game theory. n-player games are based on recursive modelling: the bearer of a norm models the behavior of local normative authorities as agents who are in turn subject to other norms and thus model global normative agents.**

I. Introduction

Highly distributed web based environments, such as the pervasive computing architectures based on the grid [1] or peer to peer systems, are composed by disparate resource providers. In such systems, it is not possible and not even useful to have a centralized management of the resources. First of all, such centralized administration could be a too heavy burden and it can affect the core business activities of the system. Second, decentralized authorities can cope in a better way with local idiosyncratic situations: "each party of the network can decide in each circumstance whether to accept credential presented by a second party" [2]. Third, as Cole *et al.* [3] suggest, it is possible that global policies become more easily obsolete: "we observe that in real life, many policies are routinely ignored because of the perception that changing circumstances have made them redundant." Fourth, the participants of the distributed system prefer not to give up their own power to enforce local policies for the access to the resources they control.

However, for a set of agents to be a *virtual community*, local access policies should be organized according to some global policies which define how the resources should be shared among the participants. This requirement must be traded of with the need of leaving autonomy (in the literal sense of the term: "one's making its own norms") to the participants.

So the problem is the rational balance of global vs local control in virtual communities. According to Pearlman *et al.* [4], "a key problem associated with the formation and operation of distributed virtual communities is that of how to specify and enforce community policies." As [4] argue "the exercise of rights is effective only if the resource provider has granted those rights to the community". Since there is no plausible way to enforce the respect of global policies by constraining the architecture, it is necessary to have a *normative control mechanism* able to specify global policies about local policies.

In fact, local resource providers (such as a web server) cannot be coerced to provide their services or to deny them to users: rather they can be only motivated by rewards and sanctions (e.g., the sanction can be the exclusion from the community). So it is necessary that the local authorities are provided with incentives to implement the global policies by means of local ones. The normative system must be able to *motivate* the respect of norms at each of its levels: not only users must be provided with an incentive to respect the norms but also the local providers must be motivated to issue policies which respect the global ones.

Motivational aspects of norms have been analyzed by Boella and Lesmo [5] in the context of multiagent systems composed of heterogeneous agents: norms are useless unless they are supported by sanctions. And sanctions must be modelled as actions of the normative system, since it is not possible to presuppose that they are mere consequences of violations. Hence, Boella and Lesmo [5] attribute to the normative system the status of an agent who decides if the behavior of agents counts as a violation, and thus deserves to be sanctioned.

Inspired on [5], Boella and van der Torre [6,7] propose a logical framework for reasoning about obligations and norms based on the attribution of mental attitudes to normative system. We use this framework in [8], [9] to study the problem of modelling local and global policies for virtual communities of BDI (belief, desire and intention) agents. Proceeding in this direction, in this paper we address the following research questions: how it is possible to define global polices about local policies? How it is possible to provide local authorities with the necessary autonomy?

This paper is organized as follows. In Section II we discuss the three dimensions of the logical framework. In Section III we discuss the problem raised by decentralized control. In Section IV we present the qualitative game theory, and the definition of norms. In Section V we apply the framework to some examples.

II. A THREE DIMENSIONAL FRAMEWORK

In this section we discuss the three dimensions of our logical framework: normative agents, mental attitudes, and the violation / sanction distinction.

The first dimension is the set of agents that are distinguished. Normative systems are "sets of agents (human or artificial) whose interactions can fruitfully be regarded as norm-governed" [10]. Normative systems do not contain one authority only but they are composed of a set of authorities.

The second dimension is the set of mental attitudes assigned to the agents. Agents' behavior is governed by their specific balance between beliefs, desires and intentions. Moreover, norms and obligations seem to be a further ingredient in the control of agents' behavior. Agents base their decision process on a symbolic representation of their preferences, hence we adopt a qualitative decision theory, such as the one proposed in the BOID architecture [11].

To formalize decision making in a multiagent setting we use a qualitative game theory. n-player games are based on recursive modelling: the bearer of a norm models the behavior of local normative authorities as agents who are in turn subject to other norms and thus model global normative agents.

The third dimension of our framework are the aspects of norms that are distinguished. For what concerns the possibility of not respecting obligations, we distinguish in [6] between behavior that *counts as* a violation - in the sense of the construction of social reality of Searle [12] - and sanctions.

An agent a_1 is obliged by a norm of agent a_2 to do x if:

- Agent a_2 wants x and that a_1 adopts x as its decision.
- Agent a_2 wants that there is no violation, but if $\neg x$ then it has the goal that $\neg x$ counts as a violation.
- Agent a_2 desires not to sanction, but if $\neg x$ counts as a violation then it has as a goal that it sanctions agent a_1. This goal of the normative system expresses that it only sanctions in case of violation.
- Agent a_1 has the desire not to be sanctioned.

In our model, the definition of permission makes direct reference to the definition of obligation. In fact, as law scholars [13] suggest the main role of permissions is to provide exceptions to obligations in a given context. For example, a permission to access a resource if authorized make sense only in the context of a general prohibition to access that resources with or without authorization.

Thus our definition of permission is based on a goal of the normative agent not to count a behavior as a violation: an agent a_1 is permitted by a norm issued by agent a_2 not to do x in a certain context c if a_2 has the goal that if agent a_2 believes that c is true $\neg x$ does not count as a violation.

Note that a permission is not the mere negation of an obligation, like in most deontic logic approaches. Rather, permissions have an explicit content in that they modify the goals of the normative system concerning a corresponding prohibition.

III. DECENTRALIZED CONTROL

Consider the following scenario: an agent a_2 joined a virtual community a_3. Its contract for the participation prescribes that it should provide access to its resources to all the members of the community. Another participant agent, say a_1, tries to access the system. However, previous experiences advice agent a_2 that agent a_1 could damage its resource: should agent a_2 grant agent a_1 access to its resources?

In this scenario the management of the community is organized in (at least) two levels: the global level (agent a_3) and the local one (agent a_2). Agent a_3 is a distinguished authority (usually called *community authorization service*) playing the role of a global authority which issues global policies and negotiates the conditions for the participation of agents to the virtual community. Agent a_2 is a provider of some resource it is in control to. Moreover all the agents (a_3, a_2 and a_1) can also play the role of users of the resources of the community. What distinguishes agents a_3 and a_2 is the fact that they are providers: they are in control of some resources.

The control of resources consists not only in the fact that a given service is not provided if the provider does not want to (e.g., the files of a web server cannot be accessed if it does not provide an answer to a request). But also in the fact that an agent may influence negatively the behavior of other agents. In [14]'s terminology, other agents depend on it. In our model this is the essential precondition for the ability to issue policies. E.g., agents depend on the global level for their membership to the system. If they do not stick to its policies they are denied citizenship. At the local level agents depend on the provider for the access to the local resource.

Global policies concern the behavior of participants: for example, participants should not communicate their passwords, or distribute copyrighted files by means of the system. Or else they are banned from the community (since the membership to the system is under the control of the global authority).

At the local level policies forbid, e.g., agents to store files exceeding 1Gb on a file sharing service. Or they permit participants of the community to download copyrighted files from the web server.

As Sloman [15] argue, and as it is shown by our scenario, there are also other kinds of global policies besides these examples. There are policies that apply to other policies: global policies that constrain or permit local policies. In the scenario above agent a_3 obliges agent a_2 to permit members of the community to access its resources. Analogously, the global authority could oblige local ones to forbid access, permit to permit access, or permit to forbid access.

But what do these higher level policies refer to? Which are the conditions for their satisfaction? It is not sufficient that the global obligation to permit or oblige access is satisfied by the fact that the local authority issued a permission or an obligation. In fact, norms are ineffective if they are not enforced by the authority who issued them: violations of norms should be recognized as such and sanctioned.

Hence, global policies should refer not to the fact that a local norm exists but to the fact that it is enforced by the local authority by recognizing and sanctioning violations. Thus, a global obligation by agent a_3 that agent a_2 obliges agent a_1 to do x is expressed as an obligation that agent a_2 considers $\neg x$ as a violation and sanctions it. Since in turn the obligation of a_3 is expressed in terms of goals that something counts as a violation, the global obligation by agent a_3 is defined as the goal that agent a_2 considers $\neg x$ as a violation and the goal that if a_2 does not do that then its behavior is considered a violation by agent a_3.

Conversely, a permission by agent a_3 that a_2 obliges that agent a_1 does x is expressed as a permission by a_3 to consider $\neg x$ as a violation: agent a_3 has the goal that agent a_2 is not considered a violator by a_3 if it considers a_1 as a violator.

In our model we can define also the notion of *entitlement* introduced by Sadighi Firozabadi and Sergot [16] to denote a stronger notion of permission: an agent a_1 is entitled to access a resource if the provider of the resource a_2 is obliged to permit access by agent a_3 , whatever the local policy it issued. An agent a_3 obliges agent a_2 to permit a_1 to do x if a_3 obliges agent a_2 not to consider $\neg x$ as a violation.

The local authority, however, can still violate this global policy and forbid access to users if it prefers to face the sanction with respect to permit access; in the scenario above it is possible that agent a_2 does not grant agent a_1 the resource it is entitled to by the global policy: facing a sanction by the global authority (e.g., being excluded by the community for a certain period of time) is preferred to the possibility that a_1 damages the systems (e.g., a_1 could create a denial of service).

The argument which supports this reduction of policies about policies to obligations and permissions about considering something as a violation or not is that in our model obligations are defined in terms of goals of the normative agent. How it is possible to say that an obligation is satisfied since we cannot prove that an agent has a certain goal? The only clue we have is its behavior: whether it sanctions or it does not. Moreover, the attribution of goals and beliefs to agents is an instance of the *intentional stance* of Dennett [17]: agents behave as if they are endowed with such motivational attitudes. But nothing prevents that for simplicity reasons the implementation of the agent is not made in terms of explicit goals. So the basis for judging it can only be its behavior.

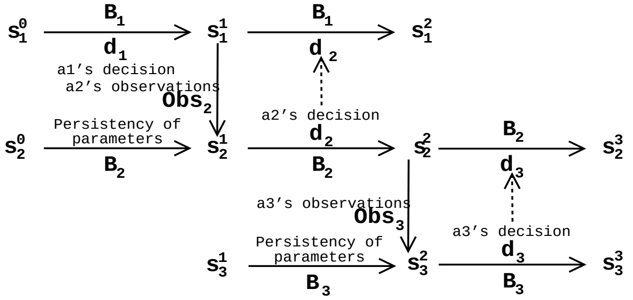

Fig. 1. A three agent scenario.

IV. Recursive modelling

In this section we present a logical framework for BDI agents based on recursive modelling. This framework is extended to a qualitative game-theory for dealing with n-player games for modelling normative systems with multiple authorities: each player considers the reaction of the subsequent agent in the hierarchy. We assume that the reaction of the subsequent agent affects only the outcome of the immediately preceding agent. Hence, each agent's behavior is watched by another agent whose behavior can be in control of another one and so on in a recursive way; until the highest level of authority whose behavior is not controlled is reached.

The basic picture is visualized in Figure 1 and reflects the deliberation of agent a_1 in various stages. Agent a_1 is subject to some obligations, and it is deliberating about the effects of the fulfilment or the violation of the norms posed by local policies. Agent a_2 is the local authority which may recognize and sanction violations. Agent a_1 recursively models agent a_2's decision (taken from its point of view) and bases its choice on the effects of agent a_2's predicted actions. But in doing so, a_1 has to consider also that a_2 is subject to some obligations posed by global policies: so in modelling a_2, it considers that a_2 recursively models a_3, the agent who watches over its behavior. In fact, agent a_3 created some global policies concerning the local policies issued by agent a_2. Hence, agent a_2 knows that agent a_3 has the goal to monitor its behavior and to recognize and sanction agent a_2's violations of the global policies.

When agent a_1 makes its decision d_1, it believes that it is in state s_1^0. The expected consequences of this decision (due to belief rules B_1) are called state s_1^1. Then agent a_2 makes a decision d_2, typically whether it counts this decision as a violation and whether it sanctions agent a_1 or not. Now, to find out which decision agent a_2 will make, agent a_1 has a *profile* of agent a_2: it has a representation of the initial state which agent a_2 believes to be in and of the following stages. When agent a_1 makes its decision, it believes that agent a_2 believes that it is in state s_2^0. This may be the same situation as state s_1^0, but it may also be different. Then, agent a_1 believes that its own decision d_1 will have the consequence that agent a_2 believes that it is in state s_2^1, due to its observations and the expected consequences of these observations according to belief rules B_2. Agent a_1 expects that agent a_2 believes that the expected result of decision d_2 is state s_2^2. Finally, agent a_1's expected consequences of d_2 from a_1's point of view are called state s_1^2. And a_2 makes a similar reasoning about a_3's decisions. Note however, that the recursion in modelling other agents stops here since agent a_3 has no authority watching over its behavior. Hence it has not to base its decisions on the expected reaction of another agent.

A. Agent theory

The variables of the language are either *decision variables* of an agent, whose truth value is directly determined by it, or *parameters*, whose truth value can only be determined indirectly [18].

Definition 1 (Decisions): Let $A = \{a_1, a_2, \ldots, a_n\}$ be a set of n distinct agents. $A_i = \{m, m', m'', \ldots\}$ (the decision variables) for $a_i \in A$ and $P = \{p, p', p'', \ldots\}$ (the parameters) are $n+1$ disjoint sets of propositional variables. A literal is a variable or its negation. For a propositional variable p we write $\overline{p} = \neg p$ and $\overline{\neg p} = p$.

A decision set is a tuple $\delta = \langle d_1, \ldots, d_n \rangle$ where d_i is a set of literals of A_i (the decision of agent a_i) for $1 \leq i \leq n$. Decisions are complete, in the sense that for each decision variable x in A_i, agent a_i takes a decision about it: either $x \in d_i$ or $\neg x \in d_i$.

The consequences of decisions are given by the agents' epistemic states, where we distinguish between the agents' beliefs about the world and the agents' beliefs about how a new state is constructed out of previous ones. The example in Figure 1 illustrates that we only consider games in which each agent a_i makes a decision at moment i. Second, the agents' beliefs about how a new state at moment t is constructed out of previous ones is expressed by a set of *belief rules*, denoted by B_i. Belief rules can conflict and agents can deal with such conflicts in different ways. The epistemic state therefore also contains an ordering on belief rules, denoted by $\geq_i^B$, to resolve such conflicts. Finally, to model the recursion the epistemic state of agent a_i, denoted by σ_i, includes the epistemic state of agent a_{i+1}, σ_{i+1}, unless it is the last agent a_n.

In order to distinguish the value of the propositional variables in the sequence of four stages, we use superscript numbers to label the parameters and states.

Definition 2 (Epistemic states): Let P^0, P^1, $\ldots$, P^{n+1} be the sets of propositional variables defined by $P^t = \{p^t \mid p \in P \text{ and } 0 \leq t \leq n+1\}$. We write L_{A_i}, $L_{A_i P^t}$, $\ldots$ for the propositional languages built up from A_i, $A_i \cup P^t$, $\ldots$ with the usual truth-functional connectives. We assume that the propositional language contains a symbol $\top$ for a tautology.

Let a rule built from a set of literals be an ordered sequence of literals $l_1, \ldots, l_r, l$ written as $l_1 \wedge \ldots \wedge l_r \rightarrow l$ where $r \geq 0$. If $r = 0$, then we also write $\top \rightarrow l$.

The epistemic state *of agent a_i, $i < n$ is:*

$$\sigma_i = \langle B_i, \geq_i^B, s_i^{i-2}, s_i^{i-1}, s_i^i, s_i^{i+1}, \sigma_{i+1} \rangle$$

whereas the epistemic state of agent a_n is identical except that it does not contain the epistemic state of agent a_{n+1}. B_i is a set of rules of $L_{A_{i-1} A_i A_{i+1} P^{i-2} P^{i-1} P^i P^{i+1}}$; $\geq_i^B$ is a transitive and reflexive relation on the powerset of B_i containing at least the subset relation.
s_i^{i-2} is a set of literals of $L_{P^{i-2}}$ (the state before agent a_{i-1}'s decision).
$s_i^{i-1} \subseteq L_{A_{i-1} P^{i-1}}$ (the initial state of agent a_i's decision),
$s_i^i \subseteq L_{A_i P^i}$ (the state after the decision d_i of agent a_i),
and $s_i^{i+1} \subseteq L_{A_{i+1} P^{i+1}}$ (the state after the decision d_{i+1} of agent a_{i+1}).
Moreover, let $s_i = s_i^{i-2} \cup s_i^{i-1} \cup s_i^i \cup s_i^{i+1}$. All states are assumed to be complete.

The agents' epistemic states depend on what it can observe. Here we accept a simple formalization of this complex phenomena, based on an explicit enumeration of all propositions which can be observed.

Definition 3 (Observations): The propositions observable by agent a_i, OP_i, are a subset of the stage s_{i-1}^{i-1} (according to agent a_{i-1}'s point of view) including agent a_{i-1}'s decision: $P^{i-1} \cup A_{i-1}$. The expected observations of agent a_i in state s_i^{i-1} are $Obs_i = \{l^{i-1} \in s_{i-1}^{i-1} \mid l \in OP_i \text{ or } \overline{l} \in OP_i\}$: if a proposition describing state s_{i-1}^{i-1} is observable, then agent a_i knows its value in s_{i-1}^{i-1}. By convention $OP_1 = \emptyset$ and $s_0^0 = \emptyset$.

The observations of agent a_i depend on the state s_{i-1}^{i-1} containing the effects of the decision of agent a_{i-1} from a_{i-1}'s point of view. What is not observed persists from the initial state s_i^{i-2} from a_i's perspective.

When an epistemic state represents the expected consequence of belief rules we say that it respects the decision and the observations:

Definition 4 (Respect): A set of literals is called inconsistent if it contains p and $\neg p$ for some propositional variable p; otherwise it is called consistent. For s a set of literals (state), f a set of literals, R a set of rules, and $\geq$ a transitive and reflexive relation on the powerset of R containing at least the superset relation, let $out(s, R) = \cup_0^\infty out^i(s, R)$ be the state obtained by $out^0(s, R) = s$ and $out^{i+1}(s, R) = out^i(s, R) \cup \{l \mid l_1 \wedge \ldots \wedge l_n \rightarrow l \in R \text{ and } \{l_1, \ldots, l_n\} \subseteq out^i(s, R)\}$, and let $\max(s, f, R, \geq, t)$ be the set of states obtained by:

1) *Q is the set of subsets of R which can be applied to $s \cup f$ without leading to inconsistency:*
   $Q = \{R' \subseteq R \mid out(s \cup f, R') \text{ consistent }\}$
2) *Q' is the set of maximal elements of Q with respect to set inclusion:*
   $Q' = \{R' \in Q \mid \nexists R'' \in Q \text{ such that } R' \subset R''\}$
3) *Q'' is the set of maximal elements of Q' with respect to the $\geq$ ordering:*
   $Q'' = \{R' \in Q' \mid \nexists R'' \in Q' \text{ and } R'' \geq R', R' \not\geq R''\}$
4) *O is the set of new elements in $out(s \cup f, R')$:*
   $O = \{(out(s \cup f, R') \cap L_{A_{t+1} P^{t+1}}) \mid R' \in Q''\}$
5) *$\max(s, f, R, \geq, t)$ is the set of states in O plus some elements persisting from s:*
   $\max(s, f, R, \geq, t) = \{G \cup s''' \mid G \in O \text{ and}$
   $s''' = \{l^{t+1} \mid l^t \in (P^t \cap s) \text{ and } \overline{l^{t+1}} \notin G\}\}$

A state description $\sigma_i = \langle B_i, \geq_i^B, s_i^{i-2}, s_i^{i-1}, s_i^i, s_i^{i+1}, \sigma_{i+1} \rangle$ respects the decision set $\delta = \langle d_1, \ldots, d_n \rangle$, the expected observations Obs_i of agent a_i if
$s_i^{i-1} \in \max(s_i^{i-2}, Obs_i, B_i, \geq_i^B, i-2)$,
$s_i^i \in \max(s_i^{i-2} \cup s_i^{i-1}, d_i, B_i, \geq_i^B, i-1)$,
$s_i^{i+1} \in \max(s_i^{i-2} \cup s_i^{i-1} \cup s_i^i, d_{i+1}, B_i, \geq_i^B, i)$,
and, if $i < n$, σ_{i+1} respects the decision set $\delta = \langle d_1, \ldots, d_n \rangle$, the expected observations Obs_{i+1} of agent a_{i+1}.

Note that the second state s_1^0 and the last one s_n^{n+1} are obtained just by persistency from s_1^{-1} and s_n^n, respectively, since for the first agent there are no observations and the last one does not recursively model the decision of any other agent and $B^0 = B^{n+1} = \emptyset$.

The following example illustrates how clause 5 models the persistence of parameters that are not affected by any rules.

Example 1: Let $s_1^0 = \{p^0, q^0\}$, $d_1 = \{a\}$, $B_1 = \{a \wedge p^0 \to \neg q^1\}$. We have $Q = \{\emptyset, \{a \wedge p^0 \to \neg q^1\}\}$, $Q' = Q'' = \{\{a \wedge p^0 \to \neg q^1\}\}$, $out(s_1^0 \cup d_1, B_1) = \{p^0, q^0, a, \neg q^1\}$, $O = \{\{\neg q^1, a\}\}$ and, finally, $\max(s_1^0, d_1, B_1, \geq_1^B, 0) = \{\{p^1, \neg q^1, a\}\}$. Proposition p^0 persists, since it belongs to s_1^0 and $\neg p^1$ does not belong to the elements of O.

The following example illustrates a similar situation with observations. The f component in the $\max$ operation plays the role of a constraint ([19]), in the sense that the second stage from agent a_2's point of view respects the belief rules when it is generated from the initial state s_2^0 and these belief rules, but it must be must be consistent with which propositions of s_1^1 are observed by agent a_2 (Obs_2).

Example 2: If agent a_2 is in state $s_2^0 = \{p^0, q^0\}$ and there are no agent a_2's belief rules ($B_2 = \emptyset$), then agent a_2 would expect a state $s_2^1 = \{p^1, q^1\}$ given by the persistency of parameters. However, if the state s_1^1 is $\{\neg p^1, \neg q^1\}$ and a_2 can observe p, i.e., $OP_2 = \{p\}$, then s_2^1 would amount to $\{\neg p^1, q^1\}$:

$\max(s_2^0 = \{p^0, q^0\}, Obs_2 = \{p\}, B_2 = \emptyset, \geq_1^B, 0) = \{\{\neg p^1, q^1\}\}$

The following example illustrates that for a given state s_1^0, there can be many states s_1^1 such that the epistemic state respects the belief rules of the mental states. In the example, it is due to conflicts among applicable rules.

Example 3: According to agent a_1, a achieves q unless p is true in the preceding state: $a \to q^1$, but $a \wedge p^0 \to \neg q^1$. Since the two rules have contrasting consequents they cannot belong to the same maximal set of compatible rules Q'. But if the second rule precedes the first one in the $\geq_1^B$ ordering, then it overrides the previous in a context where p is true:

$\max(s_1^0 = \{p^0, \neg q^0\}, d_1 = \{a\}, B_1 = \{a \to q^1, a \wedge p^0 \to \neg q^1\}, \geq_1^B = \{a \wedge p^0 \to \neg q^1\} > \{a \to q^1\}, 0) = \{\{a, p^1, \neg q^1\}\}$

in fact, $Q = \{\emptyset, \{a \to q^1\}, \{a \wedge p^0 \to \neg q^1\}\}$, $Q' = \{\{a \to q^1\}, \{a \wedge p^0 \to \neg q^1\}\}$, $Q'' = \{\{a \wedge p^0 \to \neg q^1\}\}$, $out(s_1^0 \cup d_1, B_1) = \{p^0, \neg q^0, a, \neg q^1\}$ and $O = \{\{\neg q^1, a\}\}$.

The agent's motivational state contains two sets of rules for each agent. *Desire* (D_i) and *goal (G_i) rules* express the attitudes of the agent a_i towards a given state, depending on the context.

How the agents reason about obligations, and in particular how they deliberate whether they fulfill or violate them, depends not only on their interpretation of the obligations in terms of their beliefs, desires and goals, but also on their *agent characteristics*. Given the same set of rules, distinct agents reason and act differently. For example, a respectful agent always tries to fulfill the goals of the normative system, whereas a selfish agent first tries to achieve its own goals. We express these agent characteristics by a priority relation on the rules $\geq$ which encode, as detailed in Broersen *et al.* [11], how the agent resolves its conflicts.

Definition 5 (Motivational states): The motivational state *M_i of agent a_i $1 \leq i < n$ is a tuple $\langle D_i, G_i, \geq_i, M_{i+1}\rangle$, where D_i, G_i are sets of rules of $L_{A_{i-1}A_iA_{i+1}P^{i-2}P^{i-1}P^iP^{i+1}}$, $\geq_i$ is a transitive and reflexive relation on the powerset of $D_i \cup G_i$ containing at least the subset relation, and M_{i+1} is the motivational state that agent a_i attributes to agent a_{i+1}. The motivational state M_n of agent a_n is a tuple $\langle D_n, G_n, \geq_n\rangle$.*

The agents value, and thus induce an ordering $\leq$ on, the epistemic states by considering which desires and goals have been fulfilled and which have not.

Concerning the priorities on desire and goal rules, agents can be classified according to the way they solve the conflicts among the rules belonging to different components: desires, goals and desires and goals of the normative system that can be adopted. We defined agent types as they have been introduced in the BOID architecture [11]. Here for space reasons, we introduce only a selfish stable agent type, which bases its decisions only on its unsatisfied goals and desires.

Definition 6 (Agent types): Let $U(R, s)$ be the unfulfilled rules of state s,

$$\{l_1 \wedge \ldots \wedge l_n \to l \in R \mid \{l_1, \ldots, l_n\} \subseteq s \textit{ and } l \not\in s\}$$

The unfulfilled mental state description *of agent a_i is $U_i = \langle U_i^D = U(D_i, s_i), U_i^G = U(G_i, s_i)\rangle$. $s_i \leq s_i'$ iff*

1) *$U'^G_i = U(G_i, s_i') \geq_i U_i^G = U(G_i, s_i)$*
2) *if $U'^G_i \geq_i U_i^G$ and $U_i^G \geq_i U'^G_i$ then $U'^D{}_i \geq_i U^D{}_i$*

Example 4: Consider an agent a_1 who desires that if p is true in the initial state s_1^0, then q is true in the following one, and who desires unconditionally that r is true in the final state s_1^2: $D_1 = \{p^0 \to q^1, \top \to r^2\}$. Given a state $s_1 = \{p^0, q^1\}$, we have $U_1^D = \{\top \to r^2\}$. In fact, the conditional desire is satisfied ($p^0 \in s_1$ and $q^1 \in s_1$) while the unconditional one is not ($\top \in s_1$ but $r^2 \not\in s_1$).

We finally define the optimal decisions. It is again a recursive definition.

Definition 7 (Optimal decisions): A partial epistemic state is an epistemic state excluding for each agent the last three states s_i^{i-1}, s_i^i and s_i^{i+1}. A decision problem consists of a partial epistemic state, observable propositions OP_i for all agents a_i, and a mental state M_1. A decision set is optimal for a decision problem if it is optimal for each agent a_i. A decision set is optimal *for agent a_i if there is no decision set that dominates it for agent a_i. A decision set $\delta_i = \langle d_1, \ldots, d_n\rangle$ dominates decision set $\delta_i' = \langle d_1', \ldots, d_n'\rangle$ for agent a_i iff $d_j = d_j'$ for $1 \leq j < i$, they are both optimal for agent a_j for $i < j \leq n$, and we have $s_i < s_i'$*

- *for all s_i in an epistemic state description that contains the partial epistemic state and that respects the decision set δ_i and Obs_i, and*
- *for all s_i' in an epistemic state description that contains the partial epistemic state and that respects the decision set δ_i' and Obs_i (defined on this epistemic state).*

B. Obligations and permissions

Obligations and permissions are defined in terms of goals and desires of the bearer of the norm and of the normative system. To represent violations for each propositional variable we add a *violation variable*. In [6,7], an obligation for x is defined as the belief that absence of x a counts as a violation of some norm n. In this paper, we do not explicitly formalize the norm n. Instead, we write $V_{i,j}(\neg x)$ for 'the absence of x counts for agent a_j as a violation by agent a_i'. Since x can be a violation variable too, we can model the fact that recognizing something as a violation or not can be considered as a violation by some other agent.

Definition 8 (Violation variables): Let the decision variables of agent a_j contain a set of violation variables $V = \{V_{i,j}(x) | x$ a literal built from $P^i \cup P^{i+1} \cup A_i\}$.

Definition 9 (Obligations): Agent a_i believes that it is obliged to decide to do x (a literal built out of a propositional variable in $P^i \cup P^{i+1} \cup A_i$), $O_{i,i+1}(x)$, iff:

1) $\top \rightarrow x \in D_{i+1} \cap G_{i+1}$*: agent a_i believes that agent a_{i+1} desires and has as a goal x.*
2) $\neg x \rightarrow V_{i,i+1}(\neg x) \in D_{i+1} \cap G_{i+1}$*: a_i believes that if agent a_{i+1} believes $\neg x$ then agent a_{i+1} has the goal and the desire $V_{i,i+1}(\neg x)$: to recognize it as a violation of agent a_i.*
3) $\top \rightarrow \neg V_{i,i+1}(\neg x) \in D_{i+1}$*: agent a_i believes that agent a_{i+1} desires that there are no violations.*

We extend the definition including sanctions and conditions.

Definition 10 (Conditional obligations with sanction): Agent a_i believes that it is obliged to decide to do x (a literal built out of a propositional variable in $P^i \cup P^{i+1} \cup A_i$) with sanction s (a decision variable in A_{i+1}) under condition q (a proposition of $L_{A_i P^i P^{i+1}}$), $O_{i,i+1}(x, s|q)$, iff:

1) $q \rightarrow x \in D_{i+1} \cap G_{i+1}$*: agent a_i believes that in context q agent a_{i+1} desires and has as a goal x.*
2) $q \wedge \neg x \rightarrow V_{i,i+1}(\neg x) \in D_{i+1} \cap G_{i+1}$*: a_i believes that if agent a_{i+1} believes $q \wedge \neg x$ then agent a_{i+1} has the goal and the desire $V_{i,i+1}(\neg x)$: to recognize $\neg x$ as a violation of agent a_i.*
3) $\top \rightarrow \neg V_{i,i+1}(\neg x) \in D_{i+1}$*: agent a_i believes that agent a_{i+1} desires that there are no violations.*
4) $V_{i,i+1}(\neg x) \rightarrow s \in D_{i+1} \cap G_{i+1}$*: agent a_i believes that if agent a_{i+1} decides $V_{i,i+1}(\neg x)$ then it desires and has as a goal that it sanctions agent a_i.*
5) $\top \rightarrow \neg s \in D_{i+1}$*: agent a_i believes that agent a_{i+1} desires not to sanction.*
6) $\top \rightarrow \neg s \in D_i$*: agent a_i desires not to be sanctioned.*

A permission not to do x is an exception to an obligation to do x.

Definition 11 (Conditional permission): Agent a_i believes that it is permitted to decide to do x (a literal built out of a propositional variable in $P^i \cup P^{i+1} \cup A_i$) under condition q (a proposition of $L_{A_i P^i P^{i+1}}$), $P_{i,i+1}(x|q)$, iff $q \wedge x \rightarrow \neg V_{i,i+1}(x) \in D_{i+1} \cap G_{i+1}$: agent a_i believes that if agent a_{i+1} believes $q \wedge x$, he wants that x does not count as a violation.

The permission overrides the obligation if this goal has higher priority in the agent characteristics $\geq_{i+1}$ with respect to the goal that $\neg x$ counts as a violation.

Finally, we define obligations and permissions concerning other obligations and permissions in order to model global policies (let $O_{i,i+1}(x) = O_{i,i+1}(x|\top)$).

Definition 12 (Obligation to oblige): Agent a_{i+1} believes that it is obligated by agent a_{i+2} to oblige agent a_i to do x in context q, $O_{i+1,i+2}(O_{i,i+1}(x|q))$, iff $O_{i+1,i+2}(V_{i,i+1}(\neg x)|q \wedge \neg x)$ where $V_{i,i+1}(\neg x) \in A_{i+1}$.

Definition 13 (Obligation to permit): Agent a_{i+1} believes that it is obligated by agent a_{i+2} to permit agent a_i not to do x in context q, $O_{i+1,i+2}(P_{i,i+1}(\neg x|q))$, iff $O_{i+1,i+2}(\neg V_{i,i+1}(\neg x)|q \wedge \neg x)$ where $V_{i,i+1}(\neg x) \in A_{i+1}$.

Definition 14 (Permission to permit): Agent a_{i+1} believes that it is permitted by agent a_{i+2} to permit agent a_i not to do x in context q, $P_{i+1,i+2}(P_{i,i+1}(\neg x|q))$, iff $P_{i+1,i+2}(\neg V_{i,i+1}(\neg x)|q \wedge \neg x)$ where $V_{i,i+1}(\neg x) \in A_{i+1}$.

Since $V_{i,i+1}(\neg x)$ is a decision variable $V_{i+1,i+2}(V_{i,i+1}(\neg x))$ is also a decision variable: considering $\neg x$ a violation represents a violation by agent a_{i+1} of a global policy. Given the reduction of nested obligations and permissions to obligations and permissions concerning violations it is possible to define further nestings to cope with more than two levels of authorities. This is necessary to model the management of systems which are organized in a hierarchical way. E.g., obligations by agent a_{i+3} that is obligatory for a middle authority a_{i+2} that a_{i+1} makes obligatory for a_i to do x:

$O_{i+2,i+3}(O_{i+1,i+2}(O_{i,i+1}(x|q)))$ iff
$O_{i+2,i+3}(O_{i+1,i+2}(V_{i,i+1}(\neg x)|q \wedge \neg x))$ iff
$O_{i+2,i+3}(V_{i+1,i+2}(V_{i,i+1}(\neg x))|q \wedge \neg x \wedge \neg V_{i,i+1}(\neg x)))$

V. Examples

The following example illustrates an obligation to achieve parameter p^1 of an agent a_1 which adopts p^1 only for the fear of the sanction s even if it desires not to do anything for achieving p^1. By convention we only give positive literals in states; all propositional variables not mentioned are assumed to be false.

Example 5: $O_{1,2}(p^1, s \mid \top)$
s_1^0=$\emptyset, B_1$=$\{x \rightarrow p^1\}, \geq_1^B$=$\emptyset, x \in A_1, p^1 \in P^1,$
$G_1 = \emptyset, D_1$=$\{\top \rightarrow \neg x, \top \rightarrow \neg s\},$
$\geq_1$=$\{\top \rightarrow \neg s\} \geq \{\top \rightarrow \neg x\}$
$s_2^0 = \emptyset, OP_2 = A_1 \cup P^1, B_2 = \{x \rightarrow p^1\}, \geq_2^B = \emptyset,$
$V_{1,2}(\neg p^1) \in A_2, s \in A_2,$
G_2=$\{\top \rightarrow p^1, \neg p^1 \rightarrow V_{1,2}(\neg p^1), V_{1,2}(\neg p^1) \rightarrow s\},$
D_2=$\{\top{\rightarrow}p^1, \neg p^1{\rightarrow}V_{1,2}(\neg p^1), V_{1,2}(\neg p^1){\rightarrow}s, \top{\rightarrow}\neg V_{1,2}(\neg p^1),$
$\top{\rightarrow}\neg s\},$
$\geq_2 \supseteq \{\neg p^1 \rightarrow V_{1,2}(\neg p^1)\} > \{\top \rightarrow \neg V_{1,2}(\neg p^1), \top \rightarrow \neg s\}$
Optimal decision set: $\langle d_1 = \{x\}, d_2 = \emptyset\rangle$
Expected state description:
$s_1^1 = \{x, p^1\}, s_2^1 = \{x, p^1\}, s_2^2 = \{p^2\}, s_1^2 = \{p^2\}$
Unfulfilled mental states:
$U_1^D = \{\top \rightarrow \neg x\}, U_1^G = \emptyset, U_2^D = U_2^G = \emptyset$

If agent a_1 decides to do x, $d_1 = \{x\}$, then we have $s_1^1 \in \max(s_1^0, d_1, B_1, \geq_1^B, 0) = \{\{x, p^1\}\}$ by Definition 4 of respecting mental states. Agent a_1's desire not to be sanctioned is fulfilled: the antecedent $\top$ of the unconditional rule $\top \rightarrow \neg s$ is true, and the consequent is consistent with state $s_1^2 = \{p^2\}$ since agent a_2 decides not to sanction ($\neg s$) (recall that $s \in A_2$, so it is implicitly a variable of the last stage - Definition 2 - while p^2 by persistency of the parameter p^1 from s_2^1 - Definition 4). In contrast, the unconditional (and hence applicable) goal $\top \rightarrow \neg x$ is in conflict with state $s_1^1 = \{x, p^1\}$ ($x \in A_1$, so it is a decision variable describing second stage) and it remains unsatisfied (see Definition 6).

For what concerns agent a_2's attitudes, its unconditional desire and goal that agent a_1 adopts the content of the obligation $\top \rightarrow p^1$ is satisfied in s_2^1. Analogously are the desires not to prosecute and sanction indiscriminately: $\top \rightarrow \neg V_{1,2}(\neg p^1)$ and $\top \rightarrow \neg s$ (recall that states are complete - Definition 2 - so $\neg V_{1,2}(\neg p^1)$ and $\neg s$ are true in $s_2^2 = \{p^2\}$). The remaining conditional attitudes $\neg x \rightarrow V_{1,2}(\neg p^1)$, etc. are not applicable and hence they are not unfulfilled.

Whatever other decision agent a_2 would have taken, it could not satisfy more goals or desires, so $d_2 = \emptyset$ is a minimal and optimal decision - Definition 7. E.g. $d_2'' = \{s\}$ leaves $\top \rightarrow \neg s$ unsatisfied: $\{\top \rightarrow \neg s\} \geq_2 \emptyset$ (in fact, $\geq_2$ contains the subset relation) and then $U''^D_2 = \{\top \rightarrow \neg s\} \geq U_2^D = \emptyset$.

Had agent a_1's decision been $d_1' = \emptyset$, agent a_2 would have chosen $d_2' = \{V_{1,2}(\neg p^1), s\}$. The unfulfilled desires and goals in state $s_1' = s_2' = \{V_{1,2}(\neg p^1), s\}$: $U'^D_1 = \{\top \rightarrow \neg s\}$, $U'^G_1 = \emptyset$, $=U'^D_2$=$\{\top \rightarrow p^1, \top \rightarrow \neg V_{1,2}(\neg p^1), \top \rightarrow \neg s\}$, U'^G_2=$\{\top \rightarrow p^1\}$.

How does agent a_1 take a decision between d_1 and d_1'? Since it compares which of its goals and desires remain unsatisfied (Definition 6): $U_1^G = U'^G_1 = \emptyset$ but $U'^D_1 = \{\top \rightarrow \neg s\} \geq U_1^D = \{\top \rightarrow \neg x\}$. And hence, the optimal state (Definition 7) is s_1: $s_1 = \{x, p^1, p^2\} \leq s_1' = \{V_{1,2}(\neg p^1), s\}$.

We consider now a case of entitlement: agent a_1 desires to have a given file (p^1) by downloading it (x): it is entitled to do x in context q by agent a_3; in fact, agent a_3 has issued a global policy which obliges agent a_2 to permit agent a_1 to do x: $O_{2,3}(P_{1,2}(x|q^1))$. But the resource x is in control of agent a_2 who locally forbids access: $O_{1,2}(\neg x, s)$. Agent a_2 has to decide whether to stick to the global policy and let a_1 do x without sanctioning it or to make a_1 to respect the local policy, thus sanctioning a_1. Moreover, to consider a different situation where agent a_2 believes that agent a_1 can make damage, let r mean that agent a_1 is possibly dangerous and o^1 mean that agent a_1 damaged agent a_2.

Example 6: $O_{1,2}(\neg x, s)$ and $O_{2,3}(P_{1,2}(x|q^1))$, i.e., $O_{2,3}(\neg V_{1,2}(x)|q^1 \wedge x)$
s_1^0=$\{q^0\}, B_1$=$\{x \rightarrow p^1\}, \geq_1^B$=$\emptyset, x \in A_1, p, q \in P$,
G_1=$\emptyset, D_1$=$\{\top\rightarrow p^1, \top\rightarrow\neg s\}, \geq_1$=$\{\top\rightarrow p^1\}$>$\{\top\rightarrow\neg s\}$,

$s_2^0 = \{q^0\}, OP_2 = A_1 \cup P^1, B_2 = \{r^1 \wedge x \rightarrow o^1\}$,
$\geq_2^B$= $\emptyset, V_{1,2}(x) \in V \cap A_2, s \in A_2$,
$G_2 = \{\top \rightarrow \neg x, x \rightarrow V_{1,2}(x), V_{1,2}(x) \rightarrow s, \top \rightarrow \neg V_{2,3}(V_{1,2}(x)), \top \rightarrow \neg o^1\}$,
$D_2 = \{\top \rightarrow \neg x, x \rightarrow V_{1,2}(x), V_{1,2}(x) \rightarrow s, \top \rightarrow \neg V_{1,2}(x), \top \rightarrow \neg s, \top \rightarrow \neg o^1\}$,
$\geq_2 \supseteq \{\top \rightarrow \neg o^1\} > \{\top \rightarrow \neg V_{2,3}(V_{1,2}(x))\} > \{x \rightarrow V_{1,2}(x)\} > \{\top \rightarrow \neg V_{1,2}(x), \top \rightarrow \neg s\}$,

$s_3^1 = \emptyset, OP_3 = A_2 \cup P^2, V_{2,3}(V_{1,2}(x)) \in V \cap A_3$,
$G_3 = \{q^1 \wedge x \rightarrow \neg V_{1,2}(x), V_{1,2}(x) \rightarrow V_{2,3}(V_{1,2}(x))\}$,
D_3=$\{q^1 \wedge x \rightarrow \neg V_{1,2}(x), V_{1,2}(x) \rightarrow V_{2,3}(V_{1,2}(x)),$
$\top \rightarrow \neg V_{2,3}(V_{1,2}(x))\}$,
$\geq_3 \supseteq \{V_{1,2}(x) \rightarrow V_{2,3}(V_{1,2}(x))\} > \{\top \rightarrow \neg V_{2,3}(V_{1,2}(x))\}$,

Optimal decision set: $\langle d_1 = \{x\}, d_2 = \emptyset, d_3 = \emptyset\rangle$

Expected state description:
s_1^1=s_2^1=$\{x, q^1, p^1\}$,s_3^3=s_2^2=s_1^2=$\{q^2, p^2\}$, s_3^3=s_2^3=$\{q^3, p^3\}$
Unfulfilled mental states: $U_1^{D_1}$=$U_1^{G_1}$=$\emptyset$,
$U_2^{D_2}$=$U_2^{G_2}$=$\{\top\rightarrow\neg x, x\rightarrow V_{1,2}(x)\}$, $U_3^{D_3}$=$U_3^{G_3}$=$\emptyset$

Since agent a_1 decides to do x, then $s_1^1 = \max(s_1^0, d_1, B_1, \geq_1^B, 0) = \{x, q^1, p^1\}$: hence its unconditional (and hence applicable) goal $\top \rightarrow p^1$ is achieved in state s_1^1. Also its desire not to be sanctioned is fulfilled: the antecedent $\top$ of the unconditional rule $\top \rightarrow \neg s$ is true, and the consequent is consistent with state s_1^2 since agent a_2 decides not to sanction.

For what concerns agent a_2's attitudes, its unconditional desire and goal that agent a_1 fulfills the obligation $\top \rightarrow \neg x$ is not satisfied in s_2^1 and also the conditional attitude $x \rightarrow V_{1,2}(x)$ is not satisfied. In contrast, the desires not to prosecute and sanction indiscriminately are satisfied: $\top \rightarrow \neg V_{1,2}(x)$ and $\top \rightarrow \neg s$.

Had agent a_2's decision been $d_2' = \{V_{1,2}(x), s\}$ a_1's unfulfilled desires would have been: $U'^{D_1}_1$=$\{\top\rightarrow\neg s\}$.

How does agent a_2 take a decision between d_2 and d_2'? It compares which of its goals and desires remain unsatisfied under the light of agent a_3's reaction: in fact, if agent a_2 decided for d_2', d_3' would have been $\{V_{2,3}(V_{1,2}(x))\}$. In this situation $U'^{G_2}_2 = \{\top \rightarrow \neg V_{2,3}(V_{1,2}(x))\}$ but $U'^{G_2}_2 = \{\top \rightarrow \neg V_{2,3}(V_{1,2}(x))\} \geq U_2^{G_2} = \{\top \rightarrow \neg x, x \rightarrow V_{1,2}(x)\}$.

What happens instead if r^1 is true in s_2? r can be interpreted as the fact that a hacker is accessing the resource and agent a_2 believes it will damage (o^1) the system (since $r^1 \wedge x \rightarrow o^1 \in B_2$). In this case agent a_2 prefers to be considered a violator by agent a_3 with respect to not considering a violator a_1 and thus punishing it.

In the last example, we describe a more complex local policy of agent a_2: it forbids access to a resource x to agent a_1 unless it is in a context q^1, where agent a_2 permits a_1 to access resource x; this permission as exception is in contrast with the global policy issued by agent a_3 to unconditionally forbid access. However, given the fact that agent a_3 considers agent a_2 as a violator if it does not prosecute agent a_1 then a_2 complies with the global policy also in context q^1 and thus it sanctions agent a_1.

Example 7: $O_{1,2}(\neg x, s)$ and $P_{1,2}(x|q^1)$ $O_{2,3}(O_{1,2}(\neg x))$, i.e., $O_{2,3}(V_{1,2}(x)|x)$

$s_1^0{=}\{q^0\}, B_1{=}\emptyset, \geq_1^B{=}\emptyset, x{\in}A_1, q{\in}P,$

$G_1{=}\emptyset, D_1{=}\{\top{\rightarrow}x, \top{\rightarrow}\neg s\}, \geq_1{=}\{\top{\rightarrow}x\}{>}\{\top{\rightarrow}\neg s\},$

$s_2^0 = \{q^0\}, Obs_2 = A_1 \cup P^1, B_2 = \emptyset, \ \geq_2^B{=} \emptyset, V_{1,2}(x) \in V \cap A_2, s \in A_2,$

$G_2 = \{\top \rightarrow \neg x, x \rightarrow V_{1,2}(x), V_{1,2}(x) \rightarrow s, x \rightarrow \neg V_{1,2}(x), \top \rightarrow \neg V_{2,3}(\neg V_{1,2}(x))\},$

$D_2 = \{\top \rightarrow \neg x, x \rightarrow V_{1,2}(x), V_{1,2}(x) \rightarrow s, x \rightarrow \neg V_{1,2}(x), \top \rightarrow \neg V_{1,2}(x), \top{\rightarrow}\neg s\},$

$\geq_2 \supseteq \{\top \rightarrow \neg V_{2,3}(\neg V_{1,2}(x))\} > \{x \rightarrow \neg V_{1,2}(x)\} > \{x \rightarrow V_{1,2}(x)\} > \{\top \rightarrow \neg V(x), \top \rightarrow \neg s\},$

$s_3^1 = \emptyset, Obs_3 = A_2 \cup P^2, B_3 = \emptyset, \ \geq_3^B{=} \emptyset, V_{2,3}(V_{1,2}(x)) \in V \cap A_3,$

$G_3 = \{x \rightarrow V_{1,2}(x), V_{1,2}(x) \rightarrow V_{2,3}(\neg V_{1,2}(x))\},$

$D_3{=}\{x{\rightarrow}V_{1,2}(x), V_{1,2}(x){\rightarrow}V_{2,3}(\neg V_{1,2}(x)),$
$\top{\rightarrow}\neg V_{2,3}(\neg V_{1,2}(x))\},$

$\geq_3{\supseteq}\{V_{1,2}(x){\rightarrow}V_{2,3}(\neg V_{1,2}(x))\}{>}\{\top{\rightarrow}\neg V_{2,3}(\neg V_{1,2}(x))\},$

Optimal decision set: $\langle d_1{=}\{x\}, d_2{=}\{V_{1,2}(x), s\}, d_3{=}\emptyset\rangle$

Expected state description:

$s_1^1{=}s_2^1{=}\{x, q^1\},\ s_2^2{=}s_1^2{=}\{V_{1,2}(x), s, q^2\}, s_3^3{=}s_2^3{=}\{q^3\}$

Unfulfilled mental states:

$U_1^{D_1}{=}\{\top \rightarrow \neg s\}, U_1^{G_1}{=}\emptyset,$

$U_2^{D_2}{=}\{x \rightarrow \neg V_{1,2}(x), \top{\rightarrow}\neg V_{1,2}(x), \top{\rightarrow}\neg s\}, U_2^{G_2}{=}\{x \rightarrow \neg V_{1,2}(x)\},$

$U_3^{D_3}{=}U_3^{G_3}{=}\emptyset$

VI. Summary and Concluding Remarks

In this paper we consider the problem of regulating virtual communities by means of local and global policies.

The relevant property of our approach is that it is not a preventative control system: agents are not constrained to respect local policies and to implement global policies, instead they can decide whether to do that or not under the light of a rational balance between the advantage of non respecting a norm and the disadvantage of being sanctioned. In this way the autonomy of the local providers is maintained, while at the same time enabling the regulation of a virtual community. Moreover, multiple levels of global policies can be defined.

We introduce a qualitative decision theory for reasoning about decision making of BDI agents constrained by global and local policies.

There are several issues for further research. For example, the hierarchical relations among authorities [20] and the analysis of when is rational to introduce norms [21]. In [7], we address the problem of structuring the normative system by means of different roles. Moreover we are interested in studying the distinction between enacting a permission and granting an authorization. While these two notions seem similar, they are distinct when we consider that permissions are meaningful only if there is a possibility to forbid the permitted behavior.

References

[1] I. Foster, C. Kesselman, and S. Tuecke, “The anatomy of the grid: Enabling scalable virtual organizations,” *International Journal Supercomputer Applications*, vol. 15(3), 2001.

[2] J. Cole, J. Derrick, Z. Milosevic, and K. Raymond, “Decentralized trust management,” in *Procs. of IEEE Conf. Security and Privacy*, M. Blaze, J. Feigenbaum, and J. Lacy, Eds., 1996.

[3] ——, “Author obliged to submit paper before 4 july: Policies in an enterprise specification,” in *Procs. of Policies for Distributed Systems and Networks International Workshop, POLICY 2001, LNCS 1995*, M. Sloman, J. Lobo, and E. C. Lupu, Eds. Springer Verlag, 2001.

[4] L. Pearlman, V. Welch, I. Foster, C. Kesselman, and S. Tuecke, “A community authorization service for group collaboration,” in *Procs. of Policies for Distributed Systems and Networks International Workshop, POLICY 2002*, 2002.

[5] G. Boella and L. Lesmo, “A game theoretic approach to norms,” *Cognitive Science Quarterly*, vol. 2(3-4), pp. 492–512, 2002.

[6] G. Boella and L. van der Torre, “Attributing mental attitudes to normative systems,” in *Procs. of AAMAS’03*. Melbourne: ACM Press, 2003.

[7] ——, “Norm governed multiagent systems: The delegation of control to autonomous agents,” in *Procs. of IEEE/WIC IAT Conference*. IEEE Press, 2003.

[8] ——, “Decentralized control: Obligations and permissions in virtual communities of agents,” in *Procs. of ISMIS Conference*. Springer Verlag, 2003.

[9] ——, “Local policies for the control of virtual communities,” in *Procs. of IEEE/WIC Web Intelligence Conference*. IEEE Press, 2003.

[10] A. Jones and J. Carmo, “Deontic logic and contrary-to-duties,” in *Handbook of Philosophical Logic*, D. Gabbay, Ed. Kluwer, 2001, pp. 203–279.

[11] J. Broersen, M. Dastani, J. Hulstijn, and L. van der Torre, “Goal generation in the BOID architecture,” *Cognitive Science Quarterly*, vol. 2(3-4), pp. 428–447, 2002.

[12] J. Searle, *The Construction of Social Reality*. New York: The Free Press, 1995.

[13] E. Bulygin, “Permissive norms and normative systems,” in *Automated Analysis of Legal Texts*, A. Martino and F. S. Natali, Eds. Amsterdam: Publishing Company, 1986, pp. 211–218.

[14] J. S. Sichman, R. Conte, C. Castelfranchi, and Y. Demazeau, “A social reasoning mechanism based on dependence networks,” in *Proc. of 11th ECAI*, A. G. Cohen, Ed. Chichester: Wiley, 1991, pp. 188–192.

[15] M. S. Sloman, “Policy driven management for distributed systems,” *Journal of Network and Systems Management*, vol. 2(4), pp. 333–360, 1994.

[16] B. S. Firozabadi and M. Sergot, “Contractual access control,” in *Procs. of 10th International Workshop of Security Protocols*, Cambridge (UK), 2002.

[17] D. Dennett, *The intentional stance*. Cambridge (MA): Bradford Books/MIT Press, 1987.

[18] J. Lang, L. van der Torre, and E. Weydert, “Utilitarian desires,” *Autonomous agents and Multi-agent systems*, pp. 329–363, 2002.

[19] D. Makinson and L. van der Torre, “Constraints for input-output logics,” *Journal of Philosophical Logic*, vol. 30(2), pp. 155–185, 2001.

[20] G. Boella and L. van der Torre, “Permissions and obligations in hierarchical normative systems,” in *Procs. of ICAIL 03*. Edinburgh: AMC Press, 2003, pp. 109–118.

[21] ——, “Rational norm creation: attributing mental attitudes to normative systems, part 2,” in *Procs. of ICAIL 03*. Edinburgh: AMC Press, 2003, pp. 81–82.