

# Reasoning with Group Norms in Software Agent Organisations

Huib Aldewereld<sup>1</sup>, Virginia Dignum<sup>1</sup>, and Wamberto Vasconcelos<sup>2</sup>

<sup>1</sup> Delft University of Technology, The Netherlands,  
{H.M.Aldewereld, M.V.Dignum}@tudelft.nl

<sup>2</sup> University of Aberdeen, UK, w.w.vasconcelos@abdn.ac.uk

**Abstract.** Norms have been used to represent desirable behaviours that software agents should exhibit in sophisticated multi-agent solutions. An important open research issue refers to group norms, i.e. norms that govern groups of agents. Depending on the interpretation, group norms may be intended to affect the group as a whole, each member of a group, or some members of the group. Moreover, upholding group norms may require coordination among the members of the group. We have identified three sets of agents affected by group norms, namely, i) the addressees of the norm, ii) those that will act on it, and iii) those that are responsible to ensure norm compliance. We present a formalism to represent these, connecting it to a minimalist agent organisation model. We use our formalism to develop a reasoning mechanism which enables agents to identify their position with respect to a group norm, so as to further support agent autonomy and coordination when deciding on possible courses of action.

## 1 Introduction

Norms have been used to represent, in compact ways, desirable behaviour that autonomous components should have (alternatively, undesirable behaviour they should not have), so as to provide overall guarantees for distributed, open, and heterogeneous computing solutions. Research on norms has tackled important issues, ranging from logic-theoretic aspects (*e.g.*, [22]), to more pragmatic concerns (*e.g.*, [14]).

Up until recently, the study of norms has mostly been limited to abstractions via the use of *roles* from the individual to make norms stable over extended periods of time. However, while addressing multiple agents at once (namely each agent enacting the role), it is important to realise that these norms do not address these agents *together*. A main difference is that when addressing a group of agents, it is necessary to consider aspects as responsibility and fulfilment, that are not typically addressed by most norm representations (since the agent addressed is also responsible, and is also the one to act). To illustrate this difference, we consider an obligation for children under the age of 16 to attend school. While the norm addresses children under the age of 16, who are also the ones who must perform the task of going to school, the responsibility and

blame lays with their parents/guardians. Consequently, by saying that “group  $G$  should achieve outcome  $\varphi$ ”, it is not clear who in the group should actually perform the actions that lead to  $\varphi$ , and who is to blame if the outcome is not achieved. Another example is a removal company obliged by contract to move the contents of someone’s house, including a piano. If moving the piano requires specialised qualifications, even though the removal company is the addressee of the obligation, the company will not be able to act on the norm by itself and must outsource the task.

Such group norms, explicitly differentiating groups of agents targeted by the norm, those acting upon it, and those responsible for the outcome, raise coordination issues not typically seen in norms addressing individuals or roles. The agents responsible for the norm will avoid blame, and thus have to ensure that the agents supposed to act upon it are indeed doing what they are supposed to do (or conversely, avoiding forbidden behaviours). The acting agents might need to coordinate whether each of them has to do it, one of them has to do it, or even all of them have to do it together.

In addition to coordination issues, group norms also present challenges in norm reasoning. Reasoning about norms is essential to the regulation of behaviour in multiagent systems [4]. Work on models of norm-governed practical reasoning agents have so far studied the case of norms aimed at one agent (or role) [21]. That is, the cases in which the agent is both addressed and responsible for a norm. In this paper, we present work towards reasoning about group norms. These norms require that the agent is both able to reason about its relation to the norm (*i.e.*, is it addressed, responsible or actor?) and able to coordinate with other agents affected by the norm to determine how to handle the norm.

In the next section we present a formalisation of group norms. In section 3 we present a minimalistic model of agent organisations for the concepts of action, role, and power. In section 4 we show how agents can use this model to reason about how to act when they are addressed by, responsible for or actor of a norm. Section 5 discusses issues pertaining to the reasoning processes we sketch. We contrast our research with related work in Section 6, and in Section 7 we draw conclusions, discuss relevant issues, and give avenues for future investigation.

## 2 Group Norms

Norms are a natural way of constraining behaviours of groups and individuals. However, by simply stating “group  $G$  should achieve  $\varphi$ ”, it is not made explicit who is to act (that is, whether each group member individually, if only one group member, or if all of them together), and who is to blame when violations occur. We can take again the example from the introduction: the obligation for children under the age of 16 to attend school. While the norm *addresses* children under the age of 16, who are also the ones that must *act* upon it, the *responsibility* and blame lie with the parents.

In the rest of this section we introduce a language of set definitions to precisely establish the notion of group, present a representation for group norms and provide its semantics using temporal logic.

## 2.1 Set Definitions

We propose to represent groups as set definitions and operations. We assume the existence of a non-empty and finite universal set  $Agents = \{ag_1, \dots, ag_n\}$  consisting of the unique identifier of each agent in our society, and define a language  $\mathcal{L}_\Sigma$  of set definitions that captures some of the common operations of naïve set theory [17], namely, union, intersection, difference, and absolute complement (with respect to the universal set  $Agents$ ).

We extend the language of set definitions  $\mathcal{L}_\Sigma$  to represent more sophisticated scenarios. It is common for certain norms to address groups with size restrictions, as in “gatherings of more than 5 people are prohibited”. We can formalise such requirements as  $|\Sigma| \circ n$ , where  $\circ$  is a comparison operator  $>, <, \geq, \leq, =$ , or  $\neq$  and  $n \in \mathbb{N}$  (a natural number). These set definitions can be seen as *constrained sets* and they place restrictions on which sets can be built. For instance, if  $Agents = \{a, b, c, d\}$  the definition  $|\{\alpha : \top\}| = 3$  (where  $\top$  stands for “true”, that is a property which is vacuously true for everyone) stands for all subsets of  $Agents$  with 3 elements, that is, all groups of 3 agents. A *set definition* gives rise to different *actual values* of groups, depending on the universal set of agents, denoted as  $value(\Sigma, Agents) \subseteq Agents$ . The definition of set operations over value sets are trivial (e.g.  $value(\Sigma' \cup \Sigma'', Agents) = value(\Sigma', Agents) \cup value(\Sigma'', Agents)$  and so on) and left out due to space limitations.

## 2.2 Group Norms and their Semantics

We formally capture three different groups as set expressions  $\Sigma$ , as introduced in the previous sub-section, as well as the usual components of norms, namely, the deontic modality and the target of the norm [14, 22]. We formalise group norms via Def. 1 of Fig. 1 and provide their semantics via an Anderson’s reduction [3] of the deontic modality to the reserved  $viol(G, A, R, \varphi)$  construct indicating that a violation has happened of  $G$ ’s norm on  $\varphi$  by (in)action of  $A$  under the responsibility of  $R$ .

The meaning of group obligations is stated in Def. 2 of Fig. 1. Intuitively, this definition expresses that the deadline  $\delta$  will occur at some point in time and for all paths *either*  $\varphi$  is achieved by the actors ( $stit(A, \varphi)$ ), in which case no violation of the obligation will ever occur ( $\bigcirc(\mathbf{A}\square \neg viol(G, A, R, \varphi))$ ), *or* the state is not achieved, the deadline occurs, and a violation happens ( $\delta \wedge viol(G, A, R, \varphi)$ ). Similarly, we define the meaning of group prohibitions (Def. 3 of Fig. 1).

Group prohibitions are similar to group obligations, except that the deadline  $\delta$  is better seen as a deactivation of the prohibition (and may therefore not actually occur in the future states, meaning that the prohibition is not deactivated). So, no violation happens until *either* a violation is triggered by seeing to it that the prohibited state is achieved before the deactivation ( $\neg \delta \wedge stit(A, \varphi) \wedge$

**Definition 1 (Group Norms).** Given an addressee group  $G$ , an acting group  $A$ , and a responsible group  $R$  of agents, a group norm is written as  ${}^A\mathbf{D}_G^R \varphi$  where  $\mathbf{D}$  is a deontic modality in  $\{\mathbf{O}, \mathbf{F}\}$  such that  ${}^A\mathbf{O}_G^R \varphi$  represents the obligation towards  $G$  to have  $A$  achieve  $\varphi$  under the responsibility of  $R$ ; and  ${}^A\mathbf{F}_G^R \varphi$  represents the prohibition towards  $G$  to not have  $A$  achieve  $\varphi$  under the responsibility of  $R$ .

**Definition 2 (Semantics of Obligation).**

$${}^A\mathbf{O}_G^R \varphi < \delta \stackrel{\text{def}}{=} A \left[ \diamond \delta \wedge \left( \begin{array}{c} \neg \delta \wedge \\ \neg \text{stit}(A, \varphi) \wedge \\ \neg \text{viol}(G, A, R, \varphi) \end{array} \right) \mathcal{U} \left( \left( \begin{array}{c} \neg \delta \wedge \\ \text{stit}(A, \varphi) \wedge \\ \text{O}(\text{A}\Box \neg \text{viol}(G, A, R, \varphi)) \end{array} \right) \vee \left( \begin{array}{c} \delta \wedge \text{viol}(G, A, R, \varphi) \end{array} \right) \right) \right]$$

**Definition 3 (Semantics of Prohibition).**

$${}^A\mathbf{F}_G^R \varphi < \delta \stackrel{\text{def}}{=} A \left[ \left( \begin{array}{c} \neg \delta \wedge \\ \neg \text{stit}(A, \varphi) \wedge \\ \neg \text{viol}(G, A, R, \varphi) \end{array} \right) \mathcal{U} \left( \left( \begin{array}{c} \neg \delta \wedge \\ \text{stit}(A, \varphi) \wedge \\ \text{viol}(G, A, R, \varphi) \end{array} \right) \vee (\delta \wedge \text{A}\Box \neg \text{viol}(G, A, R, \varphi)) \right) \right]$$

**Definition 4 (Action).** An action  $\mathbf{ac}$  is the triple  $\langle S, \mathbf{ac}, S' \rangle$  where  $S, S' \subseteq \mathcal{P}$  and  $\mathbf{ac}$  is an action label.

**Definition 5 (Role).** A role  $\mathbf{rl}$  is the pair  $\langle \mathbf{rl}, \text{Ac}' \rangle$  where  $\mathbf{rl}$  is the role label and  $\text{Ac}' \subseteq \text{Ac}$  is a set of action labels (cf. Def. 4).

**Definition 6 (Capabilities).** We define the set of  $\alpha$ 's capabilities (when enacting roles  $Rl'$ ), denoted as  $\text{cap}(\alpha, Rl')$ , as:

$$\text{cap}(\alpha, Rl') = \bigcup_{\langle S, \mathbf{ac}, S' \rangle \in \text{AllAc}} S' \text{ where } \text{AllAc} = \bigcup_{\langle \mathbf{rl}, \text{Ac} \rangle \in Rl'} \text{Ac}$$

That is, the capability of an agent  $\alpha$  undertaking roles  $Rl$  is the union of the post-conditions  $S'$  of all actions  $\text{AllAc}$  of all of  $\alpha$ 's roles.

**Definition 7 (Power).** Power  $\preceq \subseteq 2^{Rl}$  is a reflexive and transitive relation over the set  $Rl$  of roles. If  $\mathbf{rl}_1 \preceq \mathbf{rl}_2$  we say that  $\mathbf{rl}_2$  has power over  $\mathbf{rl}_1$  or alternatively that  $\mathbf{rl}_1$  is under the power of  $\mathbf{rl}_2$ .

**Definition 8 (Group Power).** Given sets  $\text{Agents}_1, \text{Agents}_2 \subseteq \text{Agents}$ , and a power relation  $\preceq \subseteq 2^{\text{Agents}}$  we say that  $\text{Agents}_2$  has power over  $\text{Agents}_1$ , denoted as  $\text{Agents}_1 \preceq \text{Agents}_2$ , if, and only if  $\forall \alpha' \in \text{Agents}_1, \exists \alpha'' \in \text{Agents}_2 : \alpha' \preceq \alpha''$ , that is, every member of  $\text{Agents}_1$  is under the power of at least one member of  $\text{Agents}_2$ .

**Fig. 1.** Definitions of Group Norms (Defs. 1–3) and Organisations (Defs. 4–8)

$\text{viol}(G, A, R, \varphi)$  or the prohibition is deactivated (after which no violation can occur  $(\delta \wedge \text{A}\Box \neg \text{viol}(G, A, R, \varphi))$ ).

To relate the groups and individuals of a norm, we formalise in Section 3 a notion of power – we address *social power* (*viz.*, a relation among individuals of a society, establishing who has authority or control over others [13]), as opposed to *institutional power* (*viz.*, whereby members of an institution are empowered to perform certain deeds [9, 20]). We represent power as a relation  $x \preceq y$  establishing that agent  $x$  is under the power of agent  $y$  (or conversely, that  $y$  has power over  $x$ ). This relation also applies to groups of agents, as presented later on in the paper.

### 3 A Minimalist Organisation Model

We introduce a simple organisation model to capture only those aspects necessary to explore the phenomena and mechanisms related to group norms and joint behaviour/coordination. We make use of the agent’s identity (*i.e.*, the set *Agents*), and we formalise the following aspects:

- *Roles* – representing organisation positions, abstracting from (groups of) individuals.
- *Capabilities* – roles are associated with sets of capabilities, represented as actions. These can be understood in two ways: *i*) individuals taking up a role should be able to perform what that role entails; *ii*) they specify what individuals are expected to do in the normal running of the organisation.
- *Power* – within organisations it is necessary to relate roles to one another, *i.e.* roles exert power (or influence) over other roles.

We make use of our propositional language  $\mathcal{L}_{\mathcal{P}}$ ; a set of propositional formulae  $\{\varphi_1, \varphi_2, \dots\}$  represents the *conjunction*  $\varphi_1 \wedge \varphi_2 \wedge \dots$ . We use two special propositions  $\top$  and  $\perp$  to represent, respectively, “true” and “false”. We assume that the meaning of propositions are captured with formulae establishing logical relations in a knowledge base (or, to use a more modern terminology, a reference ontology) shared by all stakeholders and components<sup>3</sup> (*e.g.*, engineers, designers, tools, software agents, and so on). We relate our formulae via logical entailment (formally, “ $\models$ ”) and deduction (formally, “ $\vdash$ ”): for any formulae  $\varphi, \psi$ , if  $\varphi \models \psi$  then  $\varphi \vdash \psi$  (completeness) and if  $\varphi \vdash \psi$  then  $\varphi \models \psi$  (correctness).

In particular, in our work we make use of logical implications represented as  $(p_1 \wedge \dots \wedge p_n) \rightarrow q$  to forge relationships among propositions, thus providing a background theory (or axioms). We denote as  $\Omega$ , a set of formulae from  $\mathcal{L}_{\mathcal{P}}$ , our background theory and we define the meaning of logical implication in terms of entailment as *if*  $((p_1 \wedge \dots \wedge p_n) \rightarrow q) \in \Omega$  *and*  $\Omega \models p_i, 1 \leq i \leq n$ , *then*  $\Omega \models q$ . A similar relation is defined for the “ $\vdash$ ” operator, if we assume its completeness.

We represent a repertoire of actions available during the enactment of an organisation. We propose an idealised representation for actions, and consider these as being *i*) instantaneous (*i.e.* they take one unit of time to be performed, that is, they do not have a duration or a period for their execution to be completed), *ii*) they are either executed or not (*i.e.*, we do not capture situations whereby actions are partially performed nor do we address scenarios in which actions are performed with degrees of success/quality). Actions are formalised in Def. 4 of Fig. 1 – the action labels uniquely identify actions. Our actions model pre-conditions  $S$  (a conjunction of propositions) which should be satisfied for action  $a$  to be performed, and the result of performing action  $ac$  ( $ac$ ’s post-conditions) is  $S'$  (a conjunction of propositions). We assume a universal, non-empty and finite set of actions  $Ac = \{\mathbf{ac}_1, \dots, \mathbf{ac}_n\}$ , such that no two actions have the same label.

<sup>3</sup> More realistically, the stakeholders and components have means to relate their knowledge bases (or, to re-phrase this in more modern terms, “align their ontologies”), thus being able to map their knowledge representation on to that of other parties.

Importantly, we note that it is possible to model norms addressing *properties of states*, rather than actions. For any action  $\langle S, ac, S' \rangle$ , where  $S' = \{p'_1, \dots, p'_n\}$ , we have  ${}^A\mathbf{D}_G^R ac \leftrightarrow {}^A\mathbf{D}_G^R (p'_1 \wedge \dots \wedge p'_n)$ , that is, a norm on an action is equivalent to a norm on its post-conditions. In the case when  $S' = \emptyset$ , we have  ${}^A\mathbf{D}_G^R ac \leftrightarrow {}^A\mathbf{D}_G^R \top$ , that is, a norm on an action without any effect is equivalent to a norm on the vacuously true proposition “ $\top$ ”, as the empty set is a sub-set of any set,  $\emptyset \subseteq S'$ .

We represent roles (Def. 5 in Fig.1) as labels available to individual agents when they join the organisation during the enactment. We associate with each role a possibly empty set of action labels, depicting what the role requires doing. When an agent joins an organisation it takes up one or more roles; by taking up a role the agent agrees to perform any of the actions associated with that role, whenever it is required (or whenever the agent is asked to). We assume a universal, non-empty and finite set of roles  $Rl = \{\mathbf{rl}_1, \dots, \mathbf{rl}_m\}$ , such that no two roles have the same label.

When individual agents join organisations they take up roles which they will incorporate during the life-time of the organisation enactment. We thus consider agents associated with a set of roles,  $\langle \alpha, Rl' \rangle, \alpha \in Agents, Rl' \subseteq Rl$ . We define an agent’s *capabilities* – the properties of the states that the agent can bring about based on the roles the agent has adopted and the actions associated with these roles, captured by Def. 6 in Fig. 1. We formally relate roles via *power*, following the approach by [11], as in Def. 7 in Fig. 1. When an agent enacts a role  $\mathbf{rl}_2$  which “has power over” another role  $\mathbf{rl}_1$  then that agent may request the help of any agent enacting  $\mathbf{rl}_1$  to achieve a particular state of affairs. This request for help is, within a formal organisation, equivalent to *delegation*, since power relations should be followed without question. Power and delegation is best understood via the “see to it” (*stit*) operator [19],  $stit(\mathbf{rl}, \varphi)$  standing for “role  $\mathbf{rl}$  sees to it that  $\varphi$ ”.

If  $stit(\mathbf{rl}, \varphi)$  and  $\mathbf{rl} = \langle rl, Ac' \rangle$  hold, then one of the following properties must also hold:

1. Role  $\mathbf{rl}$  has associated actions with combined post-conditions logically entailing  $\varphi$ :  $S^* = \bigcup_{\langle S, ac, S' \rangle \in Ac'} S'$  and  $S^* \models \varphi$
2. Role  $\mathbf{rl}$  has power over roles  $\mathbf{rl}'_i$  each of which can see to it that  $\varphi'_i$ , and these combined  $\varphi'_i$  logically entail  $\varphi$ :  $\Phi' = \{\varphi'_i : \mathbf{rl}'_i \preceq \mathbf{rl} \wedge stit(\mathbf{rl}'_i, \varphi'_i)\}$  and  $\Phi' \models \varphi$

If the power relation is symmetric, then property 2 above will also address scenarios in which agents delegate some responsibility over the achievement of  $\varphi'$  but they also retain some of it through their actions.

The power relation can be extended to relate individual agents: let there be two agents  $ag_1, ag_2$  with associated sets of roles  $\langle ag_1, Rl_1 \rangle, \langle ag_2, Rl_2 \rangle$ ; if there is a role  $\mathbf{rl}'' \in Rl_2$  for which there is a role  $\mathbf{rl}' \in Rl_1$  such that  $\mathbf{rl}' \preceq \mathbf{rl}''$ , then we say  $ag_1 \preceq ag_2$ . That is,  $ag_2$  has power over  $ag_1$  if at least one of  $ag_2$ ’s roles has power over one of  $ag_1$ ’s roles. We notice that this is a “weak” definition of power which could, in some situations, lead to loops in delegation – this is an undesirable feature of an organisation specification to which designers should be

alerted. A stronger definition would require that, in addition to the requirements above, we also had  $\mathbf{rl}'' \not\preceq \mathbf{rl}'$ , for all roles  $\mathbf{rl}', \mathbf{rl}''$ .

We further extend the power relation to account for groups (sets) of agents, via Def. 8 in Fig. 1. In order to model realistic scenarios, group norms  ${}^A\mathbf{D}_G^R \varphi$  (where  $\mathbf{D}$  is either  $\mathbf{O}$  or  $\mathbf{F}$ ) should fulfil the following properties:

1.  $A \preceq R$  – the group of actors  $A$  must be under the power of the responsible group  $R$ . This property ensures that those responsible for the norm should be able to delegate to actors.
2.  $A \preceq G$  – the group of actors  $A$  must be under the power of the group  $G$  addressed by the norm. This property ensures that addressees are also able to delegate to actors.

Both properties above can be checked at run-time, when the groups are instantiated with specific members. Interestingly, we do not insist on  $R \preceq G$  since in some scenarios we studied (*e.g.*, a norm aimed at a group of junior engineers  $G$  has a senior manager responsible  $R$  for it, and operators as actors  $A$ )  $R$  and  $G$  do not directly relate power-wise.

## 4 Reasoning about Group Norms

We consider that organisations exist in open environments where heterogeneous agents, possibly developed by third parties, may join the organisation. We assume that agents have their own motivations to decide to join an organisation and which roles they will take on, but once a role enactment is fixed, the agent is able to act on the capabilities described for its role(s). That is, role enactment depends on the agent’s own “personality”, its interpretation of what is expected from it as enactor of the role, and how the agent decides about its role norms. For instance, an agent with a strong sense of responsibility will first consider the norms for which it belongs to the Responsible group, whereas an agent that has a strong sense of duty may start by considering the norms for which it is an Actor.

In the following, we describe, in pseudo-code, reasoning mechanisms for role enacting agents, where agents have access to the organisation specification, given by the minimalistic organisation model described in section 3. The organisation specification allows agents to figure out each other’s (as well as their own) roles, capabilities, and who has power over whom.

We initially present in Algorithm 1 the general reasoning mechanism, consisting of an assessment of the value of the norm groups and a check whether or not the agent belongs to these. Depending on which group the agent belongs to, separate sub-mechanisms are invoked, and these are explained in the remainder of this section. We assume that the mechanisms have access to a global set *Agents* comprising the organisation, as well as the specific actions, roles, capabilities and (group) power relations (cf. Defs. 4–8).

The mechanism above also caters for situations in which agents simultaneously belong to more than one of the groups  $A$ ,  $R$ , or  $G$ . As we show below, agents

**Algorithm:**  $groupNormReasoning(\alpha, {}^A\mathbf{D}_G^R \varphi)$   
**if**  $\alpha \in value(G, Agents)$  **then**  $addressment(\alpha, {}^A\mathbf{D}_G^R \varphi)$ ;  
**if**  $\alpha \in value(R, Agents)$  **then**  $responsibility(\alpha, {}^A\mathbf{D}_G^R \varphi)$ ;  
**if**  $\alpha \in value(A, Agents)$  **then**  $actorship(\alpha, {}^A\mathbf{D}_G^R \varphi)$ ;  
**Algorithm 1:** Group norm reasoning

in  $G$  that are addressed by the norm will “farm out” the norm among those responsible (in group  $R$ ) and those acting (in group  $A$ ); those agents responsible for the norm (in group  $R$ ) will require the help of acting agents  $A$ . When an agent is part of more than one group, then we will have the phenomenon of agents calling upon themselves to handle the norm under a different guise.

#### 4.1 Reasoning about Actorship

Group norms are ultimately “processed” by actors: these are agents belonging to the group  $A$  of norms  ${}^A\mathbf{D}_G^R \varphi$  and their behaviours should be affected by these norms. We recall that our group norms consider propositional formulae  $\varphi$  (cf. Def. 1), and we note that these can come about as a result of a coordinated joint action among various agents, each contributing some effort to achieve or avoid  $\varphi$ . We propose the reasoning mechanism depicted in Algorithm 2 to enable norm-aware decision-making and coordination among acting agents.

Line 1 computes all those sub-groups of actors whose capabilities (under their respective adopted roles in the organisation) when pooled together logically derive (or entail). Line 2 computes those coalitions to which  $\alpha$ , the agent executing the algorithm, belongs. Lines 3–13 describe the provisions for norm-compliant behaviours. Line 14 is a place holder for non-norm-compliant behaviours – these might include, for instance, having  $\alpha$  alerting other team-members (that is, all  $\alpha'$  such that  $\alpha' \preceq \alpha$  and  $\alpha \preceq \alpha'$ ) or informing its in-line manager (that is, an  $\alpha'$  such that  $\alpha \preceq \alpha'$  and  $\alpha' \not\preceq \alpha$ ) about its decision to not comply with a norm. We focus on norm-compliant behaviour: line 4 computes agent  $\alpha$ 's own capabilities  $MyCap$  within the organisation, based on its roles. Lines 5–13 loop through each minimal coalition to which  $\alpha$  belongs, checking whether the norm is an obligation (line 6) or a prohibition (line 9). In the case of an obligation,  $\alpha$  tries to coordinate with the coalition  $Agents'$  to *contribute* with its capabilities  $MyCap$  to achieve  $\varphi$  – it is enough for one such coordination attempts to succeed for the actorship algorithm to return  $\top$  (line 8). In the case of a prohibition (line 9),  $\alpha$  attempts to coordinate with  $Agents'$  to agree on who is to refrain from doing what in order to not achieve  $\varphi$  (and hence abide by the prohibition) – it is enough for one coordination attempt to fail (that is,  $\neg coord$  holds in line 9) for the actorship algorithm to also fail. In both cases, the loop is cut short and a result is returned. If, however, the loop in lines 5-10 explores all coalitions without returning anything, then the test in lines 11-13 confirms that agent  $\alpha$  was unsuccessful in coordinating to fulfill the obligation (line 12) or  $\alpha$  was successful in coordinating to abide by a prohibition (line 13), otherwise the commands in



line 8 (respectively, line 10) would have been performed and the flow of execution would never have reached line 12 (respectively, line 13). We use logical implications to capture domain axioms, as explained previously, and we assume completeness of the “+” operator (which appears in line 1 of Alg. 2).

**Algorithm:**  $actorship(\alpha, {}^A\mathbf{D}_G^R \varphi)$

```

1 let  $\mathcal{A} = \{Agents'_0, \dots, Agents'_n\}$ , where each  $Agents'_i \subseteq value(A, Agents)$  is
  the smallest set s.t.  $\forall \alpha' \in Agents'_i, \langle \alpha', Rl' \rangle \in AssocRl \wedge (\bigcup cap(\alpha', Rl')) \vdash \varphi$ 
2  $\mathcal{A}^\alpha \leftarrow \{Agents'_i \in \mathcal{A} \mid \alpha \in Agents'_i\}$ 
3 if  $complyNorm({}^A\mathbf{D}_G^R \varphi, \mathcal{A}^\alpha)$  then
4    $MyCap \leftarrow cap(\alpha, Rl)$  where  $\langle \alpha, Rl \rangle \in AssocRl$ 
5   for  $Agents' \in \mathcal{A}^\alpha$  do
6     if  $\mathbf{D} = \mathbf{O}$  then
7       if  $coord(\alpha, (Agents' \setminus \{\alpha\}), contr, MyCap, \varphi)$  then
8         return  $\top$ 
9       else
10        if  $\neg coord(\alpha, (Agents' \setminus \{\alpha\}), refr, MyCap, \varphi)$  then
11          return  $\perp$ 
12      if  $\mathbf{D} = \mathbf{O}$  then
13        return  $\perp$ 
14      else
15        return  $\top$ 
  else
16   ... non-norm-compliant behaviour...

```

**Algorithm 2:** Reasoning about actorship

## 4.2 Reasoning about Responsibility

Agents belonging to group  $R$  of a norm  ${}^A\mathbf{D}_G^R \varphi$  are responsible for the norm, that is, they are to blame if the norm is violated. Those agents responsible enlist the help of acting agents belonging to the group  $A$  of our norms. However, agents responsible for the norm need to agree among themselves who will take the initiative to contact the actors. Moreover, the agents responsible for the norm should only contact actors over whom they have power. This process is represented in Algorithm 3.

Line 1 invokes a coordination mechanism whereby  $\alpha$  interacts with the other members of  $R$  regarding who should be ultimately responsible for the norm. This process should factor in the nature of  $\varphi$  – it might be the case that more than one agent should become involved in procuring actors to fulfill the norm. We note that the result of this coordination exercise could be the re-casting of the original norm into distinct norms whose overall effect, when they are complied with, is the same as the original norm<sup>4</sup> and we indicate this in the algorithm with a (possibly) different norm  ${}^{A'}\mathbf{X}_{G'}^{R'} \phi$  being agreed to (line 2) by the group  $R$  to

<sup>4</sup> We illustrate this with a norm (without the groups)  $\mathbf{OLiftTable}$  and axiom  $(liftEndA \wedge liftEndB) \leftrightarrow liftTable$ , which gives rise to  $\mathbf{OLiftEndA} \wedge \mathbf{OLiftEndB}$ .

have  $\alpha$  being in charge. Step 3 computes the set of actors  $\alpha'$  over which  $\alpha$  has power ( $\alpha' \preceq \alpha$ ). Line 4 establishes a loop over all actors, repeatedly invoking the actorship reasoning mechanism of algorithm 2, stopping (and returning “ $\top$ ”, that is, success) when the first of the acting agents handles the group norm. Otherwise, when we run out of choices for acting agents, the mechanism reports in line 7 a failure (“ $\perp$ ”).

**Algorithm:**  $responsibility(\alpha, {}^A\mathbf{D}_G^R \varphi)$

```

1  $\langle inCharge, {}^{A'}\mathbf{X}_{G'}^{R'} \phi \rangle \leftarrow coord'(\alpha, (value(R, Agents) \setminus \{\alpha\}), resp, {}^A\mathbf{D}_G^R \varphi)$ 
2 if  $inCharge = \alpha$  then
3    $ActorsSet \leftarrow \{\alpha' \mid \alpha' \in value(A', Agents) \wedge \alpha' \preceq \alpha\}$ 
4   for  $\alpha' \in ActorsSet$  do
5     if  $actorship(\alpha', {}^{A'}\mathbf{X}_{G'}^{R'} \phi)$  then
6       return  $\top$ 
7   return  $\perp$ 

```

**Algorithm 3:** Reasoning about Responsibility

### 4.3 Reasoning about Addressment

We finally consider the case when an agent is a member of the addressed group  $G$  of norm  ${}^A\mathbf{D}_G^R \varphi$ , depicted in Algorithm 4. In this case, the mechanism computes (line 1) the set of agents  $\alpha'$  responsible for the norm, and over which  $\alpha$  has power ( $\alpha' \preceq \alpha$ ). Line 2 starts a loop invoking, for each  $\alpha'$ , the *responsibility* mechanism depicted in algorithm 3, stopping when the first agent handles the norm. Lines 5–9 explores the exception to the responsibility mechanism, that is, a member  $\alpha$  of the addressed group  $A$ , directly takes responsibility over finding actors to deal with the norm – this part of the mechanism corresponds to lines 3–7 of the responsibility mechanism.

**Algorithm:**  $addressment(\alpha, {}^A\mathbf{D}_G^R \varphi)$

```

1  $ResponsibleSet \leftarrow \{\alpha' \in value(R, Agents) \wedge \alpha' \preceq \alpha\}$ 
2 for  $\alpha' \in ResponsibleSet$  do
3   if  $responsibility(\alpha', {}^A\mathbf{D}_G^R \varphi)$  then
4     return  $\top$ 
5  $ActorSet \leftarrow \{\alpha' \in value(A, Agents) \wedge \alpha' \preceq \alpha\}$ 
6 for  $\alpha' \in ActorSet$  do
7   if  $actorship(\alpha', {}^A\mathbf{D}_G^R \varphi)$  then
8     return  $\top$ 
9 return  $\perp$ 

```

**Algorithm 4:** Reasoning about Addressment

## 5 Discussion

Our representation of group norms caters for three distinct groups involved. Being able to differentiate among those *addressed* by the norm (*i.e.*, group  $G$ ),

those *responsible* for the norm (*i.e.*, group  $R$ ), and those *acting* on the norm (*i.e.*, group  $A$ ), allows us to formally capture interesting and realistic situations. For instance, a norm such as “anyone under the age of 16 is obliged to attend school”, can be represented as  ${}^A\mathbf{O}_G^R$  *attendSchool* where  $R$  is  $\{x : x \in \text{People} \wedge \text{parent}(x, y) \wedge \text{under16}(y)\}$ , that is, the group responsible for the norm consists of anyone who is a parent of an under-16 person;  $G = A$  and they are  $|\{y : y \in \text{People} \wedge \text{under16}(y)\}| = 1$ , that is, those responsible and the actors are individuals (*i.e.*, sets of size one) who are under-16. The norm “groups of more than 3 children are forbidden to be in a shop” is formalised as  ${}^A\mathbf{F}_G^R$  *inShop* where  $G$  is *Children*, that is, the norm is addressed at all kids;  $R$  is  $|\{x : x \in \text{Children}\}| = 1$ , that is, each child is individually responsible for the norm (hence the set has exactly one member);  $A$  is  $|\text{Children}'| > 3$ , that is, the actors are all groups of kids with 3 or more members

Our group norm representation has been put to use in mechanisms to support agents reasoning about actorship, responsibility and addressment. The reasoning invokes individual planning (during actorship reasoning), group coordination (during responsibility reasoning) and individual deliberation (during addressee reasoning). The reasoning is interleaved with message-passing (protocols) to enable coordination, as well as communication regarding who is taking up actorship, and to signal which norm has been violated.

Completeness is achieved as the net effect of our mechanism is that addressee agents exhaustively try to find someone responsible or someone to act (invoking responsibility and actorship analyses), the responsible agents exhaustively try to find actors, and finally the actors try to plan, factoring in the constraints of the norm (avoiding prohibited states, and aiming at obliged states). Termination of the process is guaranteed if there are no loops in the power relation, as all groups are finite, and so are the agents’ individual roles and actions, and the interaction (although not shown) converges with a successful action/plan or a message declining to help. The complexity of the three combined analyses, in the worst case, is the permutation of the elements of all three sets, that is,  $2^{|G \times R \times A|}$  – this is increased by the overall number of plans required to operate the organisation (instantiated with specific values when agents customise plans).

## 6 Related Work

Work on collective agency (*e.g.*, [7, 8, 26]) and collective obligations (*e.g.*, [15]) have addressed similar concerns as ours. These approaches represent norms over actions, establishing groups of agents to whom the norms apply. Some approaches regard group norms as a shorthand for a norm which applies to all/some members of the group (*e.g.*, [8]), whereas other approaches (*e.g.*, [15]) regard group norms (more specifically, collective obligations) as a shared complex action requiring individual contributions (*i.e.*, simpler actions) from those individuals of the group. However, these approaches only deal with the element of shared responsibility, neglecting the element of shared actorship. Research about the concept of shared actorship can be found in work on joint action and coalitions

(*e.g.*, [5, 1, 16]). This line of investigation is relevant as it looks into individual deliberation when coordination is required, whereas work on delegation (*e.g.*, [11, 25]) sheds light on how norms can be transferred among individuals and groups. When agents join organisations they will need to consider the implications of taking up roles, since these will determine to which groups agents will ultimately belong, and consequently which norms will be applicable, as well as how power and delegation will impact on the agents' choices. Research has addressed issues of expressiveness and reasoning complexity in various logics of coalition (*e.g.*, [6, 27]), establishing that even for simple propositional fragments, complexity is very high (*i.e.*, PSPACE in the size of the formula checked).

The notion of group association and imposing norms on groups of agents is closely related to the concept of *roles*. Roles have been explored in research on electronic institutions [12] and organisations [10, 18, 23, 26]. Roles describe collections of *stereotypical* individuals who, by adopting a role, become subject to any norms associated with that role. We note that norms addressing roles are a useful shorthand for specialised norms addressing individuals, that is, they stand for “any one who has adopted role  $r$  is subject to norm  $\nu$ ”. For instance, a norm such as “Soldiers are forbidden to enter area  $(x, y)$ ” and given agents  $a_1, \dots, a_n$  who have taken up the *soldier* role, stands for “Agent  $a_i$  is forbidden to enter area  $(x, y)$ ”, for each  $i, 1 \leq i \leq n$ . Importantly, in existing research role norms typically do not influence the joint behaviour of individuals and do not require coordination.

## 7 Conclusions and Future Work

In this paper we have proposed a representation for group norms, a topic largely ignored in the literature. Our proposal caters for three distinct types of stakeholders, namely, the addressees of the norm, those responsible for the norm, and those whose behaviours are impacted by the norm (the actors). Our representation has been influenced by a taxonomy of cases for group norms [2], with two dimensions – the individual and the collective – within a group. Certain norms, although addressed at groups, are fulfilled/violated by a single (or some) members; other norms are aimed at the group as a whole. Our reasoning mechanisms are a first attempt at defining how agents can factor in group/individual issues when deciding what to do within an organisation (hence there is a degree of predictability on the agents' part), presenting clear connections with generally agreed organisational concepts.

We are currently extending our mechanisms with the communication layer, using classic, off-the-shelf protocols such as the Contract-Net. We will connect our approach with existing planning techniques (*e.g.*, HTN [24]), to evaluate how our group norms can help agents agree on joint plans with fewer messages and in fewer rounds.

## References

1. Thomas Ågotnes and Natasha Alechina. Reasoning about joint action and coalitional ability in kn with intersection. In *Procs. 12th Int'l Conf. on Computational Logic in Multi-Agent Systems (CLIMA'11)*, volume 6814 of *LNCS*. Springer-Verlag, 2011.
2. Huib Aldewereld, Virginia Dignum, and Wamberto W. Vasconcelos. We ought to; They do; Blame the management! A conceptualisation of group norms. In *Coordination, Organizations, Institutions, & Norms IX*, LNCS. Springer, 2014.
3. A.R. Anderson. A reduction of deontic logic to alethic modal logic. *Mind*, 67:100–103, 1958.
4. G. Boella and L. van der Torre. Normative multiagent systems. In *Proceedings of Trust in Agent Societies Workshop at AAMAS'04*, 2004.
5. Stefano Borgo. Coalitions in action logic. In *Procs. 20th Int'l Joint Conf. on Artificial Intelligence, IJCAI'07*, pages 1822–1827, San Francisco, CA, USA, 2007. Morgan Kaufmann Publishers Inc.
6. Jan Broersen, Andreas Herzig, and Nicolas Troquard. What groups do, can do, and know they can do: an analysis in normal modal logics. *Journal of Applied Non-Classical Logics*, pages 261–290, 2009.
7. José Carmo. Collective agency, direct action and dynamic operators. *Logic Journal of the IGPL*, 18(1):66–98, 2010.
8. José Carmo and Olga Pacheco. Deontic and action logics for organized collective agency, modeled through institutionalized agents and roles. *Fundam. Inform.*, 48(2-3):129–163, 2001.
9. Robert Demolombe and Vincent Louis. Norms, institutional power and roles: Towards a logical framework. In *Foundations of Intelligent Systems*, volume 4203 of *Lecture Notes in Computer Science*, pages 514–523. Springer, 2006.
10. Virginia Dignum. *A Model for Organizational Interaction: Based on Agents, Founded in Logic*. PhD thesis, Universiteit Utrecht, The Netherlands, 2004.
11. Virginia Dignum and Frank Dignum. A logic of agent organizations. *Logic Journal of IGPL*, 2011.
12. Marc Esteva, Juan A. Rodríguez-Aguilar, Carles Sierra, Pere Garcia, and Josep Lluís Arcos. On the formal specifications of electronic institutions. In *Agent Mediated Electronic Commerce*, volume 1991 of *Lecture Notes in Computer Science*, pages 126–147. Springer, 2001.
13. Noah E. Friedkin. A formal theory of social power. *Journal of Mathematical Sociology*, 12(2):103–126, 1986.
14. A. García-Camino, P. Noriega, and J.-A. Rodríguez-Aguilar. Implementing norms in electronic institutions. In *Procs. 4th Int'l Joint Conf. on Autonomous Agents & Multiagent Systems, AAMAS '05*, pages 667–673, New York, NY, USA, 2005. ACM.
15. Davide Grossi, Frank Dignum, Lambèr Royakkers, and Jean-Jules Meyer. Collective obligations and agents: Who gets the blame? In *Deontic Logic in Computer Science*, volume 3065 of *Lecture Notes in Computer Science*, pages 129–145. Springer, 2004.
16. Davide Grossi, Lambèr Royakkers, and Frank Dignum. Organizational structure and responsibility. *Artificial Intelligence and Law*, 15(3):223–249, 2007.
17. Paul Halmos. *Naïve Set Theory*. Van Nostrand, 1960. Reprinted by Springer-Verlag, Undergraduate Texts in Mathematics, 1974.

18. Mahdi Hannoun, Olivier Boissier, Jaime Simão Sichman, and Claudette Sayettat. MOISE: An organizational model for multi-agent systems. In *IBERAMIA-SBIA*, volume 1952 of *LNCS*. Springer, 2000.
19. John F. Horty. *Agency and Deontic Logic*. Oxford University Press, Oxford, 2001.
20. Andrew J. I. Jones and Marek J. Sergot. A formal characterisation of institutionalised power. *Logic Journal of the IGPL*, 4(3):427–443, 1996.
21. Martin Kollingbaum and Tim Norman. NoA – a normative agent architecture. In *Procs. 18th Int’l Joint Conf. on Artificial Intelligence, IJCAI’03*, pages 1465–1466. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2003.
22. Alessio Lomuscio and Marek Sergot. On multi-agent systems specification via deontic logic. In *Intelligent Agents VIII*, volume 2333 of *Lecture Notes in Computer Science*, pages 86–99. Springer, 2002.
23. Mairi McCallum, Wamberto Weber Vasconcelos, and Timothy J. Norman. Organizational change through influence. *Autonomous Agents and Multi-Agent Systems*, 17(2):157–189, 2008.
24. Dana Nau, Okhtay Ilghami, Ugur Kuter, J. William Murdock, Dan Wu, and Fusun Yaman. SHOP2: An HTN planning system. *Journal of Artificial Intelligence Research*, 20:379–404, 2003.
25. Timothy J. Norman and Chris Reed. A logic of delegation. *Artificial Intelligence*, 174:51–71, January 2010.
26. Olga Pacheco and José Carmo. A role based model for the normative specification of organized collective agency and agents interaction. *Autonomous Agents and Multi-Agent Systems*, 6:145–184, March 2003.
27. Nicolas Troquard. Reasoning about coalitional agency and ability in the logics of “bringing-it-about”. *Autonomous Agents and Multi-Agent Systems*, 28(3):381–407, 2014.