

Norm Conflicts and Inconsistencies in Virtual Organisations

Martin J. Kollingbaum, Timothy J. Norman, Alun Preece, and Derek Sleeman

University of Aberdeen, UK,
email: {mkolling, tnorman, apreece, sleeman}@csd.abdn.ac.uk

Abstract. Organisation-oriented approaches to the formation of multi-agent systems use roles and norms to describe an agent's social position within an artificial society or Virtual Organisation. Norms are descriptive information for a role – they determine the obligations and social constraints for an agent's actions. A legal instrument for establishing such norms are contracts signed by agents when they adopt one or more roles. A common problem in open Virtual Organisations is the occurrence of conflicts between norms – agents may sign different contracts with conflicting norms or organisational changes may revoke permissions or enact dormant obligations. Agents that populate such Virtual Organisations can remain operational only if they are able to resolve such conflicts. In this paper, we discuss, how agents can identify these conflicts and resolve them.

1 Introduction

Organisation-oriented approaches to the formation of multi-agent systems assume that a community of agents form a Virtual Organisation. Its purpose is to facilitate resource sharing and problem solving among software and/or human agents [1, 2]. Virtual Organisations are defined by a set of roles, inter-role relationships and norms describing the obligations and social constraints for agents adopting such roles. Agents are regarded as *signing a contract* with the rest of the community when they are recruited into a specific role – they commit to act according to the normative specification of a role. By adopting a set of norms, the agent finds itself in a specific *normative* position – it takes on a social burden in terms of specific norms. This implies that agents must be *norm-governed* – they must be able to reason about the obligations, permissions and prohibitions that characterise their role (or set of roles) within a specific organisational context.

Virtual organisations are situated in a changing world and may, therefore, need to adapt to changes. This dynamic nature of organisations has to be taken into account in the design of agents that are recruited into organisational structures. Due to the dynamic nature of coalitions, the agent's normative position can change – the agent may have to adopt additional norms or revise existing ones. Such a change can lead to *conflicts*: an agent wants to perform an action that is simultaneously allowed and forbidden. Or it can lead to *inconsistencies*: the agent may suddenly be forbidden to perform an action that may be essential for fulfilling one of its obligations.

The NoA model of norm-governed agency [3, 4] is specifically designed to deal with such problems. NoA takes inspirations from classical BDI models [5], but has certain unique characteristics: (a) norms are first class entities that influence the practical

reasoning of an agent and (b) a specific form of deliberation, called *informed deliberation* [6], is used that enables agents to efficiently identify and resolve norm conflicts and inconsistencies. An agent based on the NoA model will analyse whether it can fulfill its obligations in a *norm-consistent* way. The agent has to investigate whether (a) all options of actions for such an obligation are allowed, (b) least one of them or (c) whether the agent will be forced to violate any other norms if it wants to fulfill an obligation. NoA agents do not filter out options for action that are norm-inconsistent. Instead, the deliberation process of the agent is *informed* about conflicts between permissions and prohibitions and the consistency situation of obligations. With such a norm-informed deliberation, a NoA agent becomes *norm-autonomous* [7] – an agent can decide whether to honour its norms or act against them.

The concepts of norm conflict and consistency of obligations are related. The agent can perform actions and fulfill its obligations in a *norm-consistent* manner only if there are no *conflicts* within the set of norms – the agent must first resolve conflicts between permissions and prohibitions regarding actions in order to be able to create a *complete partitioning* of the options for actions for fulfilling obligations. Allowing conflicts in the first place has practical benefits in the engineering of multi-agent systems – exceptional situations do not have to be anticipated in advance, but the agents themselves are endowed to deal with them. In fact, we argue [6] that it is not possible to ensure that an agent will be conflict-free in even simple scenarios. For that, NoA introduces mechanisms for detecting and classifying conflicts and proposes conflict resolution strategies the agent can employ to disambiguate its normative position so that it can then decide and select actions for fulfilling its obligations.

This paper addresses the critical issue of the occurrence of norm conflicts and how agents can remain operative in the face of such conflicts. If there is a conflict, it has to be resolved by the agents involved. A set of conflict resolution strategies has been proposed in [3, 4]. In this paper, we are interested how agents can *refine* their set of currently held norms (for example, via re-negotiating clauses in their contracts) in order to answer questions such as which obligations and prohibitions should be refined or removed or what additional permissions would ease a conflict situation and help an agent to remain operational.

2 Usage Scenario

A specific scenario is used to illustrate the importance of a normative approach to the use of Grid services. In this scenario, a research facility commits to achieve specific research goals for a company. Such a commitment has to be specified formally in the form of a contract to define the rights of the contracting partners. In our scenario, we assume that such a contract is established between the research facility and the company and includes an obligation for the researchers to deliver results of a specific analysis of a set of data. We also assume that this agreement describes a prohibition for the researchers to disclose any of these data (they have an obligation to observe confidentiality). In order to fulfill their obligation, they use services on the Grid to execute their scientific work. We assume that there are two different service providers operating on the Grid:

- a non-profit organisation provides the required service for free, but requires the user to make its data available for public use
- a commercial organisation provides the required service without such an obligation to disclose data, but the service itself is expensive

We assume that the fee for the commercial service is not covered by the budget of the research organisation – the contract with the industrial partner does not allow to spend money on such extra costs. The research organisation is, therefore, compelled to use the free service. This introduces a conflict, as the free service requires the data to be disclosed.

3 Norm-Governed Practical Reasoning

According to the model of norm-governed practical reasoning, as described in [3, 4], NoA agents are motivated by obligations to achieve a state of affairs or to perform an action. NoA agents operate with a reactive planning mechanism, where capabilities of an agent are expressed as a set of pre-specified plans. These plans are adapted to the needs of a norm-governed reasoning – they include explicit effect specifications to allow an agent to reason about the normative consequences of possible actions. In the development of this model, specific attention was given to the fact that agents may be confronted with conflicting norms in open environments. A conflict would normally render an agent unable to act. Therefore, NoA includes a model of *informed* deliberation that provides the agent with information about classes of norm conflicts and proposes conflict resolution strategies. This guarantees that NoA agents remain operational in the face of such conflicts.

3.1 Conflicts and Inconsistencies

We describe an interference between obligations and prohibitions as *norm inconsistency* and the interference between permissions and prohibitions as *norm conflict*. In order to show how norms interfere, we have to investigate how norms are specified in NoA. The NoA norm specification language provides constructs to specify obligations, permissions and prohibitions. As NoA allows universally quantified variables within norm specifications, such specifications may address whole sets of states or actions:

```
obligation (r,perform shift("a", "r", Z))
prohibition(r,perform shift("a", Y, Z))
```

Obviously, these two norms address sets of actions that possibly overlap – each of the norms is regarded as having a so-called *scope of influence*. By creating a graph over all partial and full instantiations of action (plan) `shift`, we can gain insights into these scopes of influence in more detail.

Figure 1 shows a part of a graph that outlines all partial and full instantiations of action `shift(X, Y, Z)`. It also shows the *scope* of influence of the prohibition for action `shift("a", Y, Z)`. This prohibition is regarded to be explicitly specified for `shift("a", Y, Z)` and *propagated* to each node contained in its scope – each of these

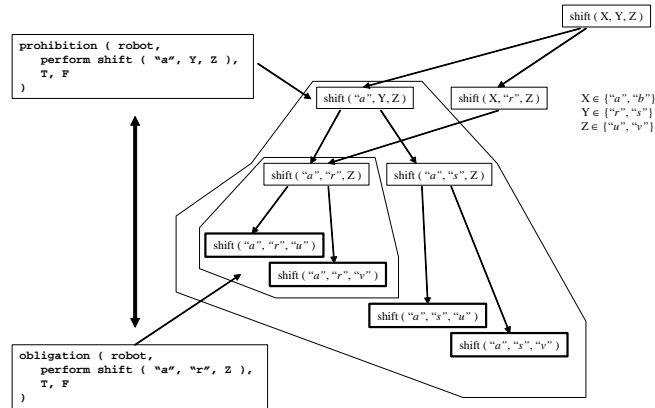


Fig. 1. Containment Relationship between Obligation and Prohibition.

nodes represents a specific partial instantiation of $\text{shift}(X, Y, Z)$ and each of these partial instantiations is regarded as being explicitly forbidden. The *instantiation set* in this depiction is the set of full instantiations that correspond to $\text{shift}("a", Y, Z)$. They are regarded as *inheriting* their normative status from their antecedents and represent those actions that are *explicitly forbidden* because of the adoption of a prohibition that contains an activity specification that addresses a whole set of actions. The instantiation set represents the set of actions (or states) that are actually allowed or forbidden. With this representation, we can regard norms as being *explicitly* introduced for a specific partial instantiation of an action (or state), represented as a node in this graph, and being *propagated* to all nodes in the scope of the norm. Nodes are interconnected according to their (partial) instantiation, with leaf nodes in this graph representing full instantiations. We see that the scope of influence of the prohibition covers the scope of influence of the obligation – the obligation demands actions that are forbidden.

Conflicts and inconsistencies occur if norms are adopted with scopes of influence that overlap. In terms of the instantiation graph, norms are regarded as being introduced for different nodes within this graph at the same time, where (a) a norm addresses a specific partial instantiation of a state or action that is *contained* within the scope of another norm, (b) the scopes of two norms *intersect* or (c) a norm is adopted for a specific action that conflicts with norms adopted for states of affairs that are effects of this action. Three main categories of conflicts emerge [4]:

- *Containment.* The scope of a norm is contained within the scope of another norm. The norms themselves can be regarded as having a specialisation relationship – one norm contains an activity specification that addresses a subset of actions or states addressed by the second norm.
- *Intersection.* The scope of a norm intersects the scope of another norm. There is no specialisation relationship between the norms. The actions or states in the intersection of both scopes *inherit* both norms at the same time.
- *Indirect Conflict/Inconsistency.* As NoA distinguishes between the achievement of a state of affairs and the performance of an action explicitly, there can also be norms

formulated that address either an action or a state. For example, an obligation is adopted that demands the performance of a set of actions (its scope of influence) and some of these actions may have effects (produce states of affairs) that are forbidden. This is regarded as an *indirect inconsistency*. In an analogous fashion, an *indirect conflict* may occur.

With respect to these characteristics of NoA, a definition of *norm-consistent action* can be given. If T_F describes the set of currently forbidden actions, S_F the currently forbidden states and S_O the set of states that the agent is obliged to achieve, then the execution of an action (plan) α , where α is not a currently forbidden action (T_F), is consistent with the current set of norms of an agent, if none of the effects of α , expressed as $effects(\alpha)$, is currently forbidden and none of the effects of α counteracts any obligation currently held by the agent (expressed as $neg_effects(\alpha)$):

$$\begin{aligned} consistent(\alpha, T_F, S_F, S_O) \text{ iff } & \alpha \notin T_F \\ & \text{and } S_F \cap effects(\alpha) = \emptyset \\ & \text{and } S_O \cap neg_effects(\alpha) = \emptyset \end{aligned}$$

With the definition of norm-consistent action and the concept of the scope of influence of a norm regarding these actions, the consistency of obligations can be determined. In NoA, we distinguish three so-called *levels of consistency* for obligations. If we describe with $options(o)$ the set of options for action that would satisfy the obligation o and which represents the *scope of influence* for this obligation, then we can investigate the consistency of each element $\alpha \in options(o)$. There are three possible configurations for this set: (a) all elements in $options(o)$ are consistent, (b) at least one element in $options(o)$ is consistent or (c) all elements are inconsistent. According to these three possibilities, we introduce three so-called *consistency levels* for a specific obligation:

- *Strong Consistency*. An obligation is strongly consistent if all $\alpha \in options(o)$ are consistent.
- *Weak Consistency*. An obligation is weakly consistent if at least one candidate in the set $options(o)$ is consistent.
- *Inconsistency*. An obligation is inconsistent if no candidate in the set $options(o)$ is consistent.

With this consistency information, the agent can decide which actions to perform to remain in a situation of at least *weak consistency* regarding its obligations.

In accordance with our e-Science scenario, let us assume that the agent (representing the research institution) has signed a contract $C1$ (the research agent has to deliver a data analysis) (see figure 2) and, with that, committed to fulfil obligation O_{C1} and adhere to a prohibition F_{C1} (this can be, for example, a prohibition for the research agent to disclose data or to spend over budget). To fulfil its obligation, the agent has two Grid services available as options for action. To use one of these services, it has to accept a second contract $C2$ with one of the service providers. As outlined before, both service providers offer their services under conditions that counteract the original agreement between research agent and industrial partner. Let us assume that the agent

does not have the capability to fulfil its obligation and, therefore, has to subcontract with one of the service providers. As pointed out in figure 2, contract C_2 introduces a new obligation O_{C_2} , which is regarded as conflicting with the prohibition F_{C_1} of the original contract C_1 .




Fig. 2. Agent Signs an Additional Contract.

In this situation, the agent should re-negotiate one of these contracts. The contracting partners will try to change the norms specified within the contract. As our e-Science example shows, the research institution cannot act because of a conflict between obligations specified within different contracts. To resolve this conflict, certain obligations and prohibitions have to be changed. In our scenario, there are two options:

- the client lifts the non-disclosure agreement – with that, the contractee could use the free service
- the client makes additional allowances in the agreed budget, which makes the use of a commercial service possible (the data does not have to be disclosed)

Both partners need information about the best course of action in such a negotiation. For the contracting partners, it is important to be informed about the normative situation – what are the conflicting norms and how obligations and prohibitions can be “relaxed” in order to allow additional options for action.

4 Norm Refinement

The goal of the re-negotiation of contracts is to create or extend a set of options for actions for a contracting agent that are *consistent* with respect to its obligations and prohibitions. In order to make such a decision, additional information is needed.

This reasoning of an agent can be supported by information derived from the instantiation graph. Inconsistency of an obligation means that the scope of influence of an obligation is completely contained within the scope of influence of a prohibition (see figure 1). To achieve a shift from inconsistency to, at least, weak consistency for an obligation, the scopes of influence have to be changed so that such a containment does not occur. There are three options:

- *Extending the Scope of Influence.* Change an obligation so that it becomes a motivator for additional actions that do not have any prohibitions.

- *Reducing the Scope of Influence.* Change a prohibition so that additional actions motivated by obligations become free of conflict.
- *Overriding prohibitions.* Introduce new permissions that override prohibitions to “allow” additional actions for the fulfilment of obligations.

To achieve a shift from this level of inconsistency to, at least, *weak consistency*, the *scope of influence* of either the obligation, the prohibition or both has to be changed. Figure 1 shows, that the two norm specifications can change their scope of influence by becoming either more specialised or more general. For example, if the prohibition forbids the action `shift("a", "s", Z)` instead of the more general `shift("a", Y, Z)`, no interference with the obligation would occur – the obligation would become *strongly consistent*. Similarly, if the obligations would be re-negotiated from `shift("a", "r", Z)` to `shift(X, Y, Z)`, then its set *options(o)* is extended and it becomes *weakly consistent*.




Fig. 3. Possible Changes to Norms to achieve a state of Weak Consistency.

Figure 3 shows the transition from the initial situation of *inconsistency* to a situation of *weak consistency* by either re-negotiating F_{C1} to transform it into F'_{C1} (reducing its scope of influence) or re-negotiating O_{C2} to transform it into O'_{C2} (extending its scope of influence). Figure 3 also shows a third option. By introducing a new Permission P_{C1} with a scope of influence that intersects with the scope of F_{C1} , options for action can be made permitted to allow the fulfilment of obligation O_{C2} . The obligation O_{C2} is operating at a level of *weak consistency*. Translated into our e-Science example, the research agent will try to utilise the commercial service as an option for action, but has to re-negotiate additional budget allowances to cover the costs of its use. With that, it is able to fulfill its obligation of payment towards the commercial service.

To achieve *strong consistency*, those norms with intersecting scopes have to be separated completely. Figure 4 shows the transition from the left-most case of figure 3 into a situation of strong consistency. This can be achieved by further specialising prohibitions or generalising obligations or by introducing specific permissions for those

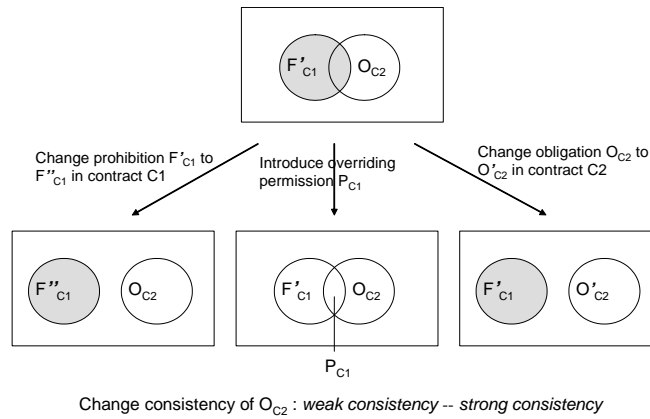


Fig. 4. Possible Changes to Norms to achieve a state of Strong Consistency.

options for action, where the scopes of obligations and prohibitions intersect. In figure 3, we introduced a new permission to free at least one option for action from the overpowering prohibition. This produced a level of *weak consistency* for the obligation O_{C2} . To achieve a transition to *strong consistency* from such a situation, a re-negotiation is necessary that transforms P_{C1} into a more general permission P'_{C1} with a scope of influence that is identical with the scope of obligation O_{C2} . This is shown in figure 5.




Fig. 5. Introducing a Permission to achieve a state of Strong Consistency.

As these examples show, the instantiation graph is a device that can display issues of conflict and inconsistencies. It shows, how prohibitions have to be specialised or obli-

gations generalised to achieve a partial or complete separation of their scopes or how the normative position of an agent can be eased by introducing a specific permission.

In order to operationalise such a refinement of norms, which may take place through a process of re-negotiation, the deliberation of an agent has to be informed about the problems occurring and the options available for solving them. An important device in NoA is the cross-referencing of actions and norms with a label that annotates actions with an accumulation of all its obligations, permissions and prohibitions.

4.1 Labeling Actions

Actions are regarded as being motivated by obligations and may also be, at the same time, prohibited as well as permitted. The normative state of an action may, therefore, be determined by a set of obligations, permissions and prohibitions. In order to cross-reference actions with norms and to indicate potential interference between these norms, these three sets are collected into a label for actions. We can describe a label for an action α as a tuple

$$l = \langle \alpha, \text{MOTIVATORS}, \text{PROHIBITORS} \rangle,$$

where

- α is the labelled candidate action for a set of motivating obligations
- *MOTIVATORS* is the set of obligations that motivate the consideration of this action as a candidate for execution, because (a) one of its effects achieves the state of affairs demanded by this obligation or (b) it is the action demanded by these obligations
- *PROHIBITORS* is the set of prohibitions or obligations that conflict with all the obligations in the set *MOTIVATORS*.

As an obligation may address a whole set of actions (see figure 1), it will be a motivator for these actions and, therefore, be an element of the set *MOTIVATORS* for each of these actions. A label for an action tries to accumulate information about conflicting norms in relation to an action. Therefore, the sets of motivators and prohibitors contain norms that are in conflict. In the set of motivators there may be obligations that are in conflict with only a subset of the prohibitors. To account for this situation, multiple labels have to be established for an action for each subset of obligations and prohibitors that are in conflict. If the set of prohibitors is empty, then a label expresses that an obligation is, at least, *weakly consistent*, as there is at least one option for action to fulfil this obligation without creating conflicts. If the set of prohibitors is empty in all labels, where an obligation occurs in the set of motivators, then this obligation is *strongly consistent*.

4.2 Detecting Conflicts

In investigating our previous example of an obligation,

```
obligation(r,perform shift("a","r",Z))
```

conflicting with a prohibition $\text{prohibition}(r, \text{perform shift}("a", Y, Z))$, we can conclude that there is a conflict between these norms if the terms representing the shift operation in both norm specifications can be unified. Consequently, a conflict can be resolved if the agent finds a substitution so that such a unification fails. For guiding the re-negotiation, the agent has to find the set $\Sigma^{failed} = \{\sigma_1, \dots, \sigma_n\}$ where the $\sigma_i, 1 \leq i \leq n$, are substitutions that are not unifiers for terms occurring in our norm specifications.

For example, if we assume $X \in \{"a", "b", "c"\}$, $Y \in \{"r", "s", "t"\}$ and $Z \in \{"u", "v", "w"\}$ for action $\text{shift}(X, Y, Z)$, then a substitution $\sigma = \{X/t_1, Y/t_2\}$ with $t_1 \in \{"b", "c"\}$ and/or $t_2 \in \{"s", "t"\}$ would be an appropriate argument for the agent to be proposed in its effort to refine either the prohibition or obligation.

The introduction of new permissions may be used to override prohibitors. A permission partially or completely overrides a prohibitor (covers parts of or the complete scope of influence), if the agent can find a substitution so that unification is successful. The agent has to find the set $\Sigma^{success} = \{\sigma_1, \dots, \sigma_n\}$ where the $\sigma_i, 1 \leq i \leq n$, are substitutions that are unifiers for terms occurring in our norm specifications. For example, if we assume $X \in \{"a", "b", "c"\}$, $Y \in \{"r", "s", "t"\}$ and $Z \in \{"u", "v", "w"\}$ for action $\text{shift}(X, Y, Z)$, then a substitution $\sigma = \{X/t_1, Y/t_2, Z/t_3\}$ with $t_1 \in \{"a"\}$ and $t_2 \in \{"r", "s", "t"\}$ and $t_3 \in \{"u", "v", "w"\}$ would be an appropriate proposal for a new permission.

In the set of prohibitors, only those prohibitions are contained that do not have a conflict with a permission. Therefore, if there is a conflict (for example, by introducing a permission that overrides a prohibition), then such a prohibition is removed from the set of prohibitors.

4.3 Options for Re-negotiation

The label expresses consistency for an action. As expressed before, an action is annotated with multiple labels, each displaying a conflict between obligations and prohibitions that expresses a situation of inconsistency for this action. An action is consistent if the set $PROHIBITORS$ is empty in all of its labels.

As outlined in figure 2, we assume a scenario where a new contract introduces a conflict. To simplify the explanations and avoid an overload with indexing, we assume that each contract introduces a single obligation, prohibition and / or permission and that each norm is indexed with a contract identifier to express the relationship between a norm and a contract.

$$\begin{aligned} l_{\alpha_1} &= \langle \alpha_1, \{O_{C1}\}, \{\} \rangle \\ l_{\alpha_2} &= \langle \alpha_2, \{O_{C3}\}, \{\} \rangle \\ l_{\alpha_3} &= \langle \alpha_3, \{O_{C1}, O_{C2}\}, \{F_{C1}\} \rangle \\ l_{\alpha_4} &= \langle \alpha_4, \{O_{C1}, O_{C3}\}, \{F_{C2}\} \rangle \\ l_{\alpha_5} &= \langle \alpha_5, \{O_{C2}, O_{C3}\}, \{F_{C1}\} \rangle \end{aligned}$$

In this scenario, a set of norms $N = \{O_{C1}, O_{C2}, O_{C3}, F_{C1}, F_{C2}, F_{C3}\}$ motivate and, partially, prohibit the performance of actions from the set of actions $A = \{\alpha_1, \dots, \alpha_5\}$. For example, obligation O_{C1} can be fulfilled by candidate actions α_1, α_3 and α_4 – obligation O_{C1} is a motivator for these actions.

The goal of re-negotiation is to resolve conflicts in the set of norms that determine an agent's normative position. It must be possible for the agent to fulfil its obligations without violating other norms – all obligations have to be at least weakly consistent. To achieve this, the agent has to know which prohibitors to re-negotiate in order to resolve conflicts. In order to perform such an analysis, we will take a snapshot of the set of labels and investigate their sets of motivators and prohibitors.

As outlined before, a label with an empty set of prohibitors indicates that obligations in the set of motivators for this label are weakly consistent. As the scenario outlined above shows, actions α_1 and α_2 have a set of motivators only:

$$\begin{aligned} l_{\alpha_1} &= \langle \alpha_1, \{O_{C1}\}, \{\} \rangle \\ l_{\alpha_2} &= \langle \alpha_2, \{O_{C3}\}, \{\} \rangle \end{aligned}$$

O_{C1} and O_{C3} are the motivators for actions α_1 and α_2 . Their labels contain no prohibitors. Therefore, these two actions provide possibilities to fulfil these obligations without violating other norms – they make obligation O_{C1} and O_{C3} weakly consistent. With that, these two obligations do not have to be considered any more.

For a further analysis of the set of labels, all occurrences of these weakly consistent obligations are removed from the set of labels. If, after this cleanup, a label has an empty set of motivators, then this label will be removed from the set of labels. In our scenario, labels l_{α_1} , l_{α_2} and l_{α_4} are changed in this way and are removed:

$$\begin{aligned} l_{\alpha_3} &= \langle \alpha_3, \{O_{C2}\}, \{F_{C1}\} \rangle \\ l_{\alpha_5} &= \langle \alpha_5, \{O_{C2}\}, \{F_{C1}\} \rangle \end{aligned}$$

The resulting set of labels can now be used to derive the minimal set of prohibitors that the agent has to re-negotiate in order to achieve weak consistency for all obligations. A procedure is employed here that will select a prohibitor according to *occurrence* – the prohibitor with the highest occurrence is chosen, removed from the remaining labels and added to a set R of prohibitors to be re-negotiated. Such a prohibitor has a relationship to a set of obligations and, therefore, has to be added to the set R together with its related obligations. In our scenario, set R contains prohibitor F_{C1} together with obligation O_{C2} :

$$R = \{\{F_{C1}, O_{C2}\}\}$$

In general, the removal of such a prohibitor from all the labels where it occurs will, again, leave some labels with empty sets of prohibitors. The cleanup step described before must be repeated and such labels deleted. After that, again, a new prohibitor with maximal occurrence has to be selected, added to set R and removed from labels. Both the cleanup step and the selection of a prohibitor has to be repeated until all labels are removed. This creates a set R of prohibitors for re-negotiation, where a precedence relationship \prec exists between its members. The relationship of a prohibitor to its obligations has to be expressed accordingly:

$$\begin{aligned} R &= \{\{F_1, O_1, \dots, O_{m_1}\}, \dots, \{F_n, O_1, \dots, O_{m_n}\}\}, \\ F_1 &\prec F_2 \prec \dots \prec F_n \end{aligned}$$

Instead of selecting prohibitors according to *occurrence*, other criteria may be chosen for such a selection process. For example, the agent may hold a function

$$violate(F) : N \rightarrow \mathbb{R}$$

that calculates the cost of a violation of a specific prohibitor, which influences the selection in the elimination process described before.

5 RELATED WORK

Norms have found increasing attention in the research community as a concept that drives the behaviour of agents within virtual societies. Conte and Castelfranchi [8, 9] investigate in detail how agents within a society reason about norms regarding their actions and what motivates them to honour their obligations and prohibitions and fulfill their commitments. Conte et al. [8, 7], argue that for a computational model of norm-governed agency, the internal representation of norms and normative attitudes, and models of reasoning about norms is a necessity. Norm-governed agents must be able to recognise norms as a social concept, represent them as mental objects and solve possible conflicts among them. Such agents should, in the words of [7], be truly *norm-autonomous* – they must be able to take a “flexible” approach towards norms: know existing norms, learn / adopt new ones, negotiate norms with peers, convey / impose norms on other agents, control and monitor other agents’ norm-governed behaviour, and be able to decide whether to obey or violate them. Cavendon and Sonenberg [10] use Castelfranchi’s concept of *social commitment* to investigate mechanisms of commitment, power and control within agent societies. Like in NoA, in their model obligations of agents are attached to the concept of a role. By adopting such a role the agent *commits* to pursue the attached obligations or “goals”. Such an adoption takes place, if agents engage in a specific relationship based on a social contract that assigns specific roles to the contracting partners. To solve conflicts between obligations due to the agent adopting multiple roles, a concept of *role influence* is used – one role is more influential to the agent’s acting than other roles and, therefore, translates into a stronger social commitment for the agent. Panzarasa et al. [11, 12] discuss the influence of a social context on the practical reasoning of an agent. They point out that the concept of *social commitment* as introduced by Castelfranchi and investigated by Cavendon and Sonenberg has to be extended to include issues of how social commitments and regulations inform and *shape* the internal mental attitudes of an agent to overcome the solipsistic nature of current BDI models. Work pursued by Broersen et al. [13], Dastani and van der Torre [14, 15], the model of a normative agent described by Lopez et al. [16] and, specifically, the NoA system as presented in this paper and elsewhere [4, 3] introduce concepts of norm influence into practical reasoning agent to make this transition from solipsistic to social agents. The NoA model of norm-governed agents takes strong inspirations from the work of Kanger [17], Lindahl [18] and Jones and Sergot [19, 20] for the representation of *rights* and the concept of a *normative position*. Members of a society adopt these norms and, ideally, operate according to them. Adopted norms determine the social or normative position of an individual [18], expressing duties, powers, freedom etc. under specific legal circumstances. This normative position can change any time with new

norms coming into existence or old ones removed. Relationships of power create organisational structures and hierarchies within a society, assigning specific roles to members of an organisation [19, 21].

Dignum et al. [22] describe the three basic aspects in the modelling of virtual societies of agents: (a) the overall purpose of such a community of agents, (b) organisational structure based on a set of roles and (c) norms for regulating the actions and interactions of the agents adopting such roles. In line of our previous argument that the solipsistic nature of agents has to be overcome for virtual organisations, they emphasise as well the importance of introducing a *collective* perspective on an agent's actions in a specific role within a society - the agent cannot not be solely driven by internal motivations, but it has to be socially aware in its practical reasoning. As also described in [23], Agents take on roles and responsibilities and are determined in their actions by *external* influences in the form of social regulations and norms. Pacheco and Carmo [21] describe the modelling of complex organisations and organisational behaviour based on roles and normative concepts. The creation of virtual societies is based on contracts between agents. Such a contract describes the set of norms that specify roles and agents adopting such roles commit to act according to these norms. Pacheco and Carmo emphasise the importance of these contracts as the central element to bind agents into societies.

Organisational change and the impact of these social dynamics on the normative position of the agent, as addressed in previous work [24, 4, 25, 3], also find attention in the work of Esteva et al. [26], Lopez and Luck [27] and Skarneas [28]. Dastani et al. [29] investigate conflicts that can occur during the adoption of a role by an agent. Esteva et al. [26] present a computational approach for determining the consistency of an electronic institution. As shown in [4], the NoA model includes a detailed classification of conflict situations that *informs* the deliberation of the agent about problems of norm conflicts and inconsistencies between the agents actions and its norms and can be used to guide the re-negotiation of contracts. With that, a NoA agent does not require a conflict-free set of norms to be operable, as it is provided with conflict resolution strategies to deal with conflicting norm sets.

6 CONCLUSION

In case of a norm conflict, agents may have to re-negotiate their contracts. The goal of such a re-negotiation must be a guarantee that obligations can be fulfilled by actions that do not violate any prohibitions. The NoA model and architecture for norm-governed practical reasoning agents takes specific care to inform the agent about the *norm consistency* of its options for actions for fulfilling its obligations and provides resolution strategies for conflicts between norms. In this paper, we illustrate how this model of norm-consistent action and norm conflicts can be used to inform the agents in the re-negotiation of their contracts.

References

1. Foster, I., Kesselman, C., Tuecke, S.: The Anatomy of the Grid: Enabling Scalable Virtual Organizations. *Int' J. Supercomputer Applications* **15** (2001) 209–235

2. Norman, T., Preece, A., Chalmers, S., Jennings, N., Luck, M., Dang, V., Nguyen, T., Deora, V., Shao, J., Gray, W., Fiddian, N.: Agent-based Formation of Virtual Organisations. *Knowledge Based Systems* **17** (2004) 103–111
3. Kollingbaum, M.: Norm-governed Practical Reasoning Agents. PhD thesis, University of Aberdeen (2005)
4. Kollingbaum, M., Norman, T.: Strategies for Resolving Norm Conflict in Practical Reasoning. In: ECAI Workshop CEAS 2004. (2004)
5. Wooldridge, M.: An Introduction to MultiAgent Systems. John Wiley & Sons, LTD (2002)
6. Kollingbaum, M., Norman, T.: Anticipating and Resolving Conflicts during Organisational Change, Technical Report AUCS/TR0505. Technical report, University of Aberdeen (2005)
7. Conte, R., Falcone, R., Sartor, G.: Agents and Norms: How to fill the Gap? *Artificial Intelligence and Law* **7** (1999)
8. Conte, R., Castelfranchi, C.: *Cognitive and Social Action*. UCL Press (1995)
9. Castelfranchi, C.: Modelling Social Action for AI Agents. *Artificial Intelligence* **103** (1998) 157–182
10. Cavendon, L., Sonenberg, L.: On social Commitments, Roles and Preferred Goals. In: Proceedings ICMAS 98. (1998)
11. Panzarasa, P., Norman, T., Jennings, N.: Modelling Sociality in the BDI Framework. In: *Intelligent Agent Technology: Systems, Methodologies, and Tools*. Proceedings of the First Asia-Pacific Conference on Intelligent Agent Technology., World Scientific Publishing (1999) 202–206
12. Panzarasa, P., Jennings, N., Norman, T.: Social Mental Shaping: Modelling the Impact of Sociality on the Mental States of Autonomous Agents. *Computational Intelligence* **17** (2001) 738–782
13. Broersen, J., Dastani, M., Hulstijn, J., Huang, Z., van der Torre, L.: The BOID architecture: Conflicts between Beliefs, Obligations, Intentions and Desires. In: *Proceedings of Autonomous Agents 2001*. (2001) 9–16
14. Dastani, M., van der Torre, L.: What is a normative Goal? In: *Proceedings of the International Workshop on Regulated Agent-based Social Systems: Theories and Applications (AAMAS 2002 Workshop)*. (2002)
15. Dastani, M., van der Torre, L.: A Classification of Cognitive Agents. In: *Proceedings of the 24th Annual Meeting of the Cognitive Science Society CogSci 2002*. (2002) 256–261
16. Lopez y Lopez, F., Luck, M., dInverno, M.: Constraining autonomy through norms. In: *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-agent Systems*. (2002) 647–681
17. Kanger, S., Kanger, H.: Rights and Parliamentarism. *Theoria* **32** (1966) 85–115
18. Lindahl, L.: *Position and Change: A Study in Law and Logic*. D. Reidel Publishing Company (1977)
19. Jones, A., Sergot, M.: A Formal Characterisation of Institutionalised Power. *Journal of the IGPL* **4** (1996) 429–445
20. Sergot, M.: A Computational Theory of Normative Positions. *ACM Transactions on Computational Logic* **2** (2001) 581–622
21. Pacheco, O., Carmo, J.: A Role Based Model for the Normative Specification of Organized Collective Agency and Agents Interaction. *Autonomous Agents and Multi-Agent Systems* **6** (2003) 145–184
22. Dignum, V., Meyer, J.J., Weigand, H., Dignum, F.: An Organization-Oriented Model for Agent Societies. In: *Proceedings of AAMAS 2002*. (2002)
23. Dastani, M., Dignum, F., Dignum, V.: Organizations and Normative Agents. In: *Proceedings of the First Eurasian Conference on Advances in Information and Communication Technology Eurasia ICT*. (2002)

24. Kollingbaum, M., Norman, T.: Norm Consistency in Practical Reasoning Agents. In Dastani, M., Dix, J., eds.: PROMAS Workshop on Programming Multiagent Systems. (2003)
25. Kollingbaum, M., Norman, T.: Norm Adoption and Consistency in the NoA Agent Architecture. In Dastani, M., Dix, J., Seqhrouchni, A., eds.: Programming Multiagent Systems: Languages, Frameworks Techniques, and Tools. Volume 3067 of Lecture Notes in Artificial Intelligence. Springer-Verlag (2004)
26. Esteva, M., Vasconcelos, W., Sierra, C., Rodriguez-Aguilar, J.: Verifying Norm Consistency in Electronic Institutions. In: Proceedings of the AAAI 2004 Workshop on Agent Organisations: Theory and Practice. (2004)
27. Lopez y Lopez, F., Luck, M.: Towards a Model of the Dynamics of Normative Multi-Agent Systems. In: Proceedings of the International Workshop on Regulated Agent-Based Social Systems: Theories and Applications RASTA'02, Bologna (2002)
28. Skarmees, N.: Organisations through Roles and Agents. In: Proceedings of the International Workshop on the Design of Cooperative Systems COOP'95. (1995)
29. Dastani, M., Dignum, V., Dignum, F.: Role-assignment in open Agent Societies. In: Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems. (2003)