Norm Refinement – Informing the Re-Negotiation of Contracts

Martin J. Kollingbaum and **Timothy J. Norman** and **Alun Preece** and **Derek Sleeman**[1]

Abstract.

Conflicts between norms are a common problem in open Virtual Organisations and has to be dealt with. Norm-governed agents that populate such VO's can remain operational only if they are able to resolve such conflicts. In this paper, we discuss how norm-governed agents based on the NoA architecture identify such conflicts and how the NoA model of norm conflicts can inform a re-negotiation of contracts and norms.

1 INTRODUCTION

In the formation of short-term coalitions or Virtual Organisations, agents have to negotiate and agree on their roles within such a coalition and what are their obligations and social constraints in such a role. 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. These agents must be able to process norms and contracts that define and determine organisational structures and the actions and interactions of agents taking on roles. These agents must be norm-governed – they must be able to take the obligations, prohibitions and permissions into account that arise from their role. By adopting a set of norms, the agent finds itself in a specific *social position*. Due to the dynamic nature of coalitions, this position can change – the agent may have to adopt additional norms or revise existing ones. Such a change can lead to conflicts – the agent may suddenly be forbidden to perform an action that may be essential to fulfil one of its obligations. For example, in e-Science, researchers utilise agents to find appropriate services for the performance of scientific tasks. Service providers may employ agents that answer such requests. The execution of a research task, for example the analysis of data with a tool implemented as a Grid service, will take place in the context of such a coalition between agents. This coalition has to be based on an agreement or contract that details this transaction – what the obligations of the participants are and what prohibitions constrain the actions of the agents. For example, a specific Grid service may demand that in case of its use the client has to make the data to be analysed publicly available in a repository. A common complication in such a scenario may be the fact that our research agent possibly uses this Grid service in order to fulfil a contract with an industrial partner who may not wish data to be disclosed. The agent, thus, has a conflict that must be resolved.

The NoA model of norm-governed agency [14, 17] is specifically designed to deal with such problems. NoA takes inspirations from classical BDI models [27], 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* [18], is used that enables agents to efficiently identify and resolve norm conflicts. An agent based on the NoA model will analyse whether it can fulfil its obligations in a *norm-consistent* way – are all options of actions for such an obligation allowed or at least one of them or would the agent be forced to violate any other norms if it wants to fulfil this obligation? NoA agents do not filter out options for action that are norm-inconsistent. Instead, the deliberation process of the agent is *informed* about the consistency situation of an obligation. With such a norm-informed deliberation, a NoA agent becomes *norm-autonomous* [5] – an agent can decide whether to honour its norms or act against them.

The agent can perform actions in a *norm-consistent* manner only if there are no *conflicts* within the set of norms – the agent must be able to create a *complete partitioning* of the options for actions in terms of consistency. An important aspect of the NoA model is to make agents robust against the potential threat of norm conflicts. Allowing conflicts in the first place has practical benefits in engineering such 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 [18] 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 decide how to fulfil 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 [14, 17]. In this paper, we give the re-negotiation of (parts of) a contract further consideration. Specifically, we are interested how to guide the agents in this negotiation – how 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

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

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 fulfil 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.

This scenario demonstrates that the use of Grid services requires agreements between service providers and clients and that these agreements or contracts introduce obligations and rights for the contracting partners.

3 NORM CONFLICTS

NoA [14, 17] is a model and architecture for norm-governed practical reasoning agents. 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.

The NoA model of conflict detection, classification and resolution is strongly influenced by the characteristics of the NoA architecture itself:

- NoA agents operate with a reactive planning mechanism – like classic implementations of the BDI architecture, the capabilities of a NoA agent are provided as a set of prespecified plans.
- NoA plans are adapted to the specific needs of norm-governed practical reasoning – agents must be able to reason about the normative consequences of possible actions. Therefore, NoA plans contain explicit effect specifications that describe the state of affairs that will occur due to the execution of this plan
- NoA agents distinguish between the achievement of goals and the performance of actions – an obligation may require the agent to either achieve a goal for which the agent has to select a plan that achieves this goal as one of its effects or the obligation directly demands a specific action for which a specific plan has to exist (otherwise the agent will not have the capability to fulfil this obligation)

Obligations are a motivator for a NoA agent to act, either to achieve a specific goal to to perform an action. If there is a conflict between norms the agent currently holds – for example, the agent is at the same time allowed and forbidden to perform a specific action – then it will not be clear to the agent whether it can act or not.

In order to manage such conflict situations, a NoA agent must be able to detect and classify them in order to resolve them.

3.1 Norm Specification in NoA

The NoA norm specification language provides constructs to specify obligations, permissions and prohibitions. Following two examples (using a simple blocks world scenario) shows an obligation in its two forms, either to achieve a specific state of affairs or to perform an action:

```
obligation (r,achieve on ("a","c"))
prohibition(r,perform shift("a","b","c"))
```

This example illustrates that a norm specification contains a so-called *activity specification* that determines either the achievement of a goal or the performance of an action. In this example, the prohibition forbids that a specific state of affairs is achieved in the blocks world, whereas the obligation determines that a specific $shift$ operation must be performed. In this case, the obligation demands an action with an outcome that is forbidden by the prohibition – both norms are in conflict. This assumes that the agent holds a plan $shift$ that provides a capability for fulfilling the obligation.

To specify more general norms that allow to address a whole range of activities, norm specification contain variables. Accordingly, a prohibition

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

specifies that an agent is prohibited from shifting block a from any block to any other block. As a more general specification, it covers the above prohibition as well. This second prohibition prohibits a whole set of actions the agent could deploy. Obviously, the *scope of influence* of the second prohibitions on the actions of the agent is much broader than that of the first prohibition. In general, norm specifications in NoA contain universally quantified variables and, therefore, address sets of actions. The activity statement expressed in a norm specification is regarded as being partially instantiated. This has to be taken into account in the investigation of norm conflicts.

From this, we can derive that norms have a *scope of influence* on the actions of the agent. These scopes may intersect, which either indicates a specialisation relationship between norms of the same type (as in case of the two prohibitions above) or a conflict between norms of different type.

To illustrate the *scope of influence* of norms and to provide more insight into the conflict situation in terms of partially instantiated norm specifications, we introduce the so-called *instantiation graph* as a device to map out all possible (partial) instantiations of actions (as well as states) and to explain and visualise possible conflict scenarios.

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 nodes represents a specific partial instantiation of `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 `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

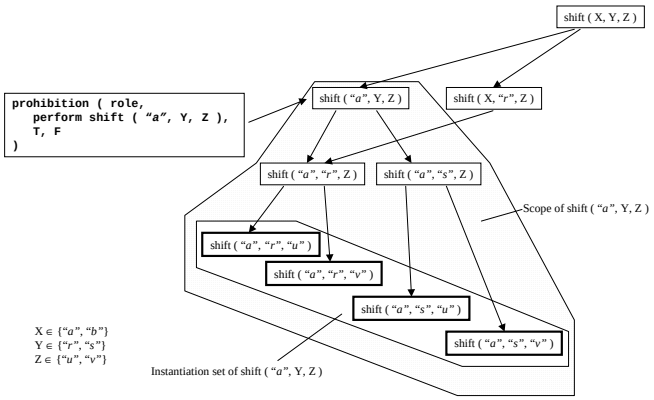

Figure 1. Instantiation Graph and Scope of Influence of Norms.

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.

Conflicts occur if norms are adopted with scopes of norm 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:

- *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*. A norm is adopted for a specific action that conflicts with norms adopted for states of affairs that are effects of this action.

In [14, 17], a set of conflict resolution startegies are outlined. One of these conflict resolution strategies is the re-negotiation of a contract.

4 INFORMING THE RE-NEGOTIATION

In the process of re-negotiating their contract, 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 allowences 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.

The goal of such a re-negotiation is to create or extend a set of options for actions for a contracting agent that are *consistent* with its currently held norms. For NoA agents, norm consistent action is defined in the following way:

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) p, where p 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 p is currently forbidden and none of the effects of p counteracts any obligation currently held by the agent:

$$\begin{array}{lll} consistent(p, T_F, S_F, S_O) & \textit{iff} & p \notin T_F \\ & \textit{and} & S_F \cap effects(p) = \emptyset \\ & \textit{and} & S_O \cap neg\_effects(p) = \emptyset \end{array}$$

Via the definition for the *consistency* of options for action (in NoA called the *candidate plans*), the consistency of obligations can be characterised. 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 $p \in options(o)$. 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 $p \in options(o)$ 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.

In accordance with our e-Science scenario, let us assume that the agent (representing the research institution) has signed a contract $C1$ (see figure 2).

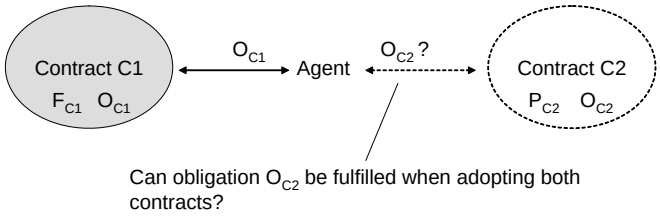

Figure 2. Agent Signs an Additional Contract.

Contract $C1$ specifies an obligation O_{C1} (the research agent has to deliver a data analysis) and 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. In figure 2, contract $C2$ introduces a new obligation O_{C2} (in case of the free service, this would be the obligation to disclose the data, in case this contract is negotiated with the commercial service, this would be the obligation to pay for the service). In figure 2, the obligation O_{C2} is regarded as conflicting with the prohibition F_{C1} of the original contract $C1$.

According to the NoA model of norm-consistent action and consistency levels of obligations, an obligation can be fulfilled in a norm-abiding manner, if it is at least *weakly consistent*. This means that there is at least one option for action that is consistent with the agent's currently held norms. The contracting partners have to decide how to change the set of norms so that the obligations can be fulfilled.

We will use the instantiation graph, shown in figure 3, as as device to guide these negotiations. Figure 3 illustrates a possible relationship between an obligation and a prohibition regulating the activities in a blocks world. Figure 3 shows that the *scope of influence* of the obligation is completely *contained* within the scope of the prohibition. This means that the set of options for action $options(o)$ for this obligation contains only inconsistent options. The obligation is, therefore, *inconsistent*.

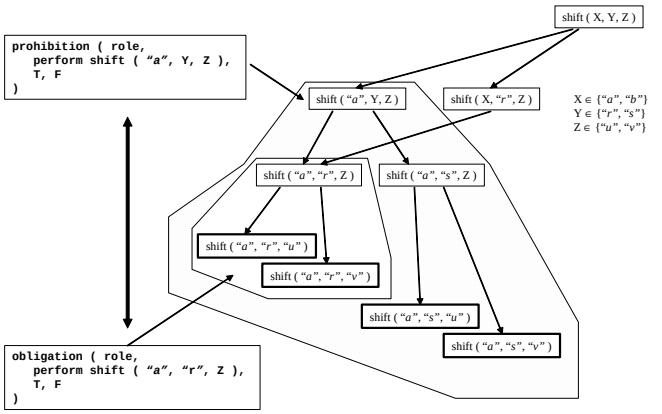

Figure 3. Containment Relationship between Obligation and Prohibition.

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 3 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*.

Figure 4 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

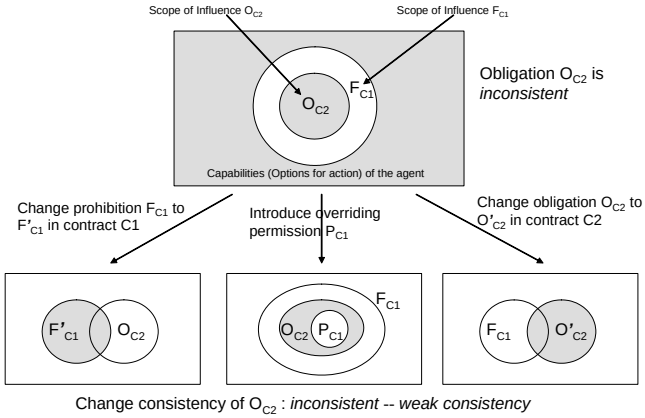

Figure 4. Possible Changes to Norms to achieve a state of Weak Consistency.

influence). Figure 4 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.

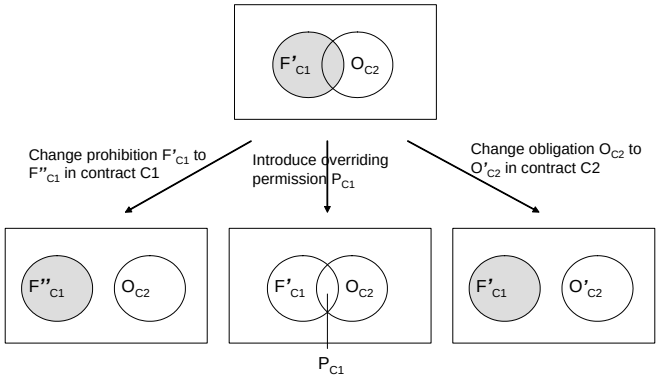

Figure 5. Possible Changes to Norms to achieve a state of Strong Consistency.

To achieve *strong consistency*, those norms with intersecting scopes have to be separated completely. Figure 5 shows the transition from the left-most case of figure 4 into a situation of strong consistency. This can be achieved by further specialising prohibitions or generalising obligations or by introducing specific permissions for those options for action, where the scopes of obligations and prohibitions intersect. In figure 4, 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 6.

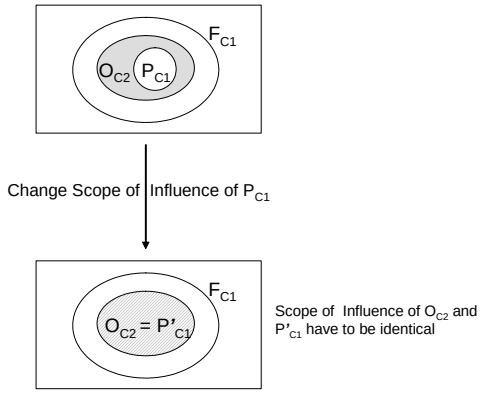

Figure 6. Introducing a Permission to achieve a state of Strong Consistency.

As these examples show, the instantiation graph is a device that can guide the re-negotiation of norms. It shows, how prohibitions have to be specialised or obligations 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.

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 [4, 2] 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. [4, 5], 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 [5], 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. Cavedon and Sonenberg [3] 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 then other roles and, therefore, translates into a stronger social commitment for the agent. Panzarasa et al. [24, 23] 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. [1], Dastani and van der Torre [9, 8], the model of a normative agent described by Lopez et al. [21] and, specifically, the NoA system as presented in this paper and elsewhere [17, 14] 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 [13], Lindahl [19] and Jones and Sergot [12, 25] 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 [19], 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 [12, 22].

Dignum et al. [10] 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 [6], 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 [22] 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 [15, 17, 16, 14], also find attention in the work of Esteva et al. [11], Lopez and Luck [20] and Skarmeas [26]. Dastani et al. [7] investigate conflicts that can occur during the adoption of a role by an agent. Esteva et al. [11] present a computational approach for determining the consistency of an electronic institution. As shown in [17], 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 renegotiation of their contracts.

REFERENCES

[1] J. Broersen, M. Dastani, J. Hulstijn, Z. Huang, and L. van der Torre, 'The BOID architecture: Conflicts between Beliefs, Obligations, Intentions and Desires', in *Proceedings of Autonomous Agents 2001*, pp. 9–16, (2001).

[2] C. Castelfranchi, 'Modelling Social Action for AI Agents', *Artificial Intelligence*, **103**, 157–182, (1998).

[3] L Cavendon and L. Sonenberg, 'On social Commitments, Roles and Preferred Goals', in *Proceedings ICMAS 98*, (1998).

[4] R. Conte and C. Castelfranchi, *Cognitive and Social Action*, UCL Press, 1995.

[5] R. Conte, R. Falcone, and G. Sartor, 'Agents and Norms: How to fill the Gap?', *Artificial Intelligence and Law*, **7**(1), (March 1999).

[6] M. Dastani, F. Dignum, and V. Dignum, 'Organizations and Normative Agents', in *Proceedings of the First Eurasian Conference on Advances in Information and Communication Technology Eurasia ICT*, (2002).

[7] M. Dastani, V. Dignum, and F. Dignum, 'Role-assignment in open Agent Societies', in *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems*, (2003).

[8] M. Dastani and L. van der Torre, 'A Classification of Cognitive Agents', in *Proceedings of the 24th Annual Meeting of the Cognitive Science Society CogSci 2002*, pp. 256–261, (August 2002).

[9] M. Dastani and L. van der Torre, 'What is a normative Goal?', in *Proceedings of the International Workshop on Regulated Agent-based Social Systems: Theories and Applications (AAMAS 2002 Workshop)*, (2002).

[10] V. Dignum, J-J. Meyer, H. Weigand, and F Dignum, 'An Organization-Oriented Model for Agent Societies', in *Proceedings of AAMAS 2002*, (2002).

[11] M Esteva, W. Vasconcelos, C. Sierra, and J.A. Rodriguez-Aguilar, 'Verifying Norm Consistency in Electronic Institutions', in *Proceedings of the AAAI 2004 Workshop onf Agent Organisations: Theory and Practice*, (2004).

[12] A.J.I. Jones and M. Sergot, 'A Formal Characterisation of Institutionalised Power', *Journal of the IGPL*, **4**(3), 429–445, (1996).

[13] S. Kanger and H. Kanger, 'Rights and Parliamentarism', *Theoria*, **32**, 85–115, (1966).

[14] M.J. Kollingbaum, *Norm-governed Practical Reasoning Agents*, Ph.D. dissertation, University of Aberdeen, 2005.

[15] M.J. Kollingbaum and T.J. Norman, 'Norm Consistency in Practical Reasoning Agents', in *PROMAS Workshop on Programming Multiagent Systems*, eds., M. Dastany and J. Dix, (2003).

[16] M.J. Kollingbaum and T.J. Norman, 'Norm Adoption and Consistency in the NoA Agent Architecture', in *Programming Multiagent Systems: Languages, Frameworks Techniques, and Tools*, eds., M. Dastani, J. Dix, and A. Seqhrouchni, volume 3067 of *Lecture Notes in Artificial Intelligence*, Springer-Verlag, (2004).

[17] M.J. Kollingbaum and T.J. Norman, 'Strategies for Resolving Norm Conflict in Practical Reasoning', in *ECAI Workshop CEAS 2004*, (2004).

[18] M.J. Kollingbaum and T.J. Norman, 'Anticipating and Resolving Conflicts during Organisational Change, Technical Report AUCS/TR0505', Technical report, University of Aberdeen, (2005).

[19] L. Lindahl, *Position and Change: A Study in Law and Logic*, D. Reidel Publishing Company, 1977.

[20] F. Lopez y Lopez and M. Luck, '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 RASTA02*, Bologna, (2002).

[21] F. Lopez y Lopez, M. Luck, and M. dÍnverno, 'Constraining autonomy through norms', in *In Proceedings of the First International Joint Conference on Autonomous Agents and Multi-agent Systems*, pp. 647–681, (2002).

[22] O. Pacheco and J. Carmo, 'A Role Based Model for the Normative Specification of Organized Collective Agency and Agents Interaction', *Autonomous Agents and Multi-Agent Systems*, **6**(2), 145–184, (2001).

[23] P. Panzarasa, N.R. Jennings, and T.J. Norman, 'Social Mental Shaping: Modelling the Impact of Sociality on the Mental States of Autonomous Agents', *Computational Intelligence*, **17**(4), 738–782, (2001).

[24] P. Panzarasa, T.J. Norman, and N.R. Jennings, '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.*, pp. 202–206. World Scientific Publishing, (1999).

[25] M. Sergot, 'A Computational Theory of Normative Positions', *ACM Transactions on Computational Logic*, **2**(4), 581–622, (2001).

[26] N. Skarmeas, 'Organisations through Roles and Agents', in *Proceedings of the International Workshop on the Design of Cooperative Systems COOP95*, (1995).

[27] M. Wooldridge, *An Introduction to MultiAgent Systems*, John Wiley & Sons, LTD, 2002.