Monitoring Interaction in Organisations

Mehdi Dastani[1], Leendert van der Torre[2], and Neil Yorke-Smith[3,4]

[1] Utrecht University, The Netherlands
M.M.Dastani@uu.nl
[2] University of Luxembourg, Luxembourg
leon.vandertorre@uni.lu
[3] Olayan School of Business, American University of Beirut, Lebanon
[4] University of Cambridge, United Kingdom
nysmith@aub.edu.lb

Abstract. In an organisational setting, such as an online marketplace, the organisation monitors agent interactions, and enforces norms by means of sanctions. This paper provides an operational semantics for agent interactions within such a setting, distinguishing constitutive norms for monitoring and sanction rules for enforcement of norms. Our contribution emphasizes a more detailed exploration of the processes of monitoring commitments created through agent interactions and imposition of sanctions when commitments are violated. We consider both agent–agent and agent–environment interactions, focusing on operationalizing enforcement of commitment-based norms. We provide a generic way to develop operational semantics from specific definitions of norm behaviour. For an example set of norm behaviours, we sketch some formal properties that follow from our semantics, such as continuity, (non-)interference, and (non-)redundancy.

1 Introduction

In an open multi-agent environment, agents can coordinate their interactions by means of communication. Following earlier works concerned with multi-agent organisations, we define the semantics of agent interactions, including communication actions, in terms of social commitments [20]. Such commitments constitute institutional facts of an organisation, and it is within an organisational setting such as 2OPL [12] that agents interact. For our purposes, the organisation consists of two main processes: the monitoring process checks for compliance to norms while the enforcement process ensures imposition of sanctions.[1] Agent interactions affect the state of the institutional facts (i.e., commitments). The norms, represented as counts-as rules, define both constitutive as well as regulative norms. Finally, sanction rules are responsible for updating the organisation state as a consequence of detected violations.

Our focus, then, is on what has been called *agreement technologies*, i.e., how coordination is achieved between autonomous computational entities [2]. Following [12] one can specify norms to govern the interaction between agents. However, we believe that interaction by means of communication should respect a set of generic norms which are

[1] Other organisation attributes and processes, such as environmental interaction, or roles, entities, and the relationships between them, are orthogonal to our purpose in this paper.

J. Simão Sichman and H. Aldewereld (Eds.): COIN 2012, LNAI 7756, pp. 17–34, 2013.

inherent in communication actions. We follow the line of works that posit that communication actions create and operate on social commitments [20,11,15]. In our setting, an organisation will manage the commitments according to the utterances (messages) and actions of the agents, and regiment the compliance of agents with the commitments. Since we represent norms as counts-as rules, we specify the relations between communication actions and commitments explicitly via counts-as rules.

Our contribution emphasizes a more detailed exploration of the processes of monitoring and enforcement through an organisation's tracking of institutional state defined in terms of commitments, which are established and modified through agent communication acts. By adopting a simple set of such actions, and not focusing on the protocol or semantic concerns of a full Agent Communication Language, we are able develop a more technical account for a full operational semantics, and explore its properties.

A common methodology to define an operational semantics in such a paradigm specifies the typical or desired behaviour of the system. For example, explicit acceptance of commitments (i.e., must the creditor accept a commitment, or not?), fulfilment of a commitment's consequent before its antecedent (i.e., can a non-detached commitment be satisfied?), or synchronization details (e.g., if a consequent is fulfilled at the same moment its deadline occurs, is the commitment detached or expired?). Based on such choices, one defines norm behaviour, typically based on a finite state automaton; based on the automaton, one then defines an operational semantics.

Our methodology differs in that, given an automaton defined using rules, we develop a generic way to go from the automaton to the operational semantics. Thus, the semantics for any given description of norm behaviours follow automatically by means of a generic mechanism. Hence our generic methodology can be applied to different sets of counts-as rules, and further can admit constitutive norms that change dynamically.

We use the term 'organisation' and 'organisational setting' in this paper. Some authors draw a distinction between institutional concepts and organisational concepts, placing for example in the former, concepts such as brute and institutions facts, counts-as relations, and the constitutive nature of commitments, and placing in the latter concepts such as roles, organisational objectives, and terminologies or conventions. In such a view, to be explored would be the relationship between the organisation and the institutional reality created by the conventions we study.

The remainder of the paper is organized as follows. Sections 2 and 3 present our formal setting and operational semantics of the organisation from the coordination point of view. Section 4 illustrates on an insurance scenario. Section 5 studies properties of our approach and investigates the management of agent interaction and commitments. Section 6 places our work in context of the existing literature, Section 7 presents topics for future research, and Section 8 concludes the paper with a summary.

2 Normative Organisation

Our model of an exogenous organisation allows individual agents to interact with each other and with their shared environment. The function of an organisation, as we will use the term, is to monitor and regulate the interaction between agents, i.e., it observes the (inter)actions of agents, evaluates their consequences and imposes sanctions if needed.

The agents' (inter)actions are assumed to be observed/received by the organisation as events. The evaluation of agents' (inter)actions as well as their sanctioning are realized based on the norms and sanctions that specify the organisation. In this paper, we focus on commitments as a specific type of norms and study agents' interactions that influence the generation and state of commitments.

For example, consider an organisation such as a marketplace where agents can buy and sell goods from each other. In such an organisation, agent i offers agent j to make a payment $p_{(b,20)}$ (i.e., paying 20 euro for book b) within 5 days if agent j first sends the book $s_{(b,i)}$ within 2 days. We assume that the marketplace has the credit card information of the participating agents and that the agents can give a payment order to the organisation when they have received their goods. We also assume that the marketplace is allowed to withdraw money from a credit card without the order from the card holder only in designated cases, e.g., when the buyer of a good, who has asked the seller to send him the good, returns the good back to the seller, the organisation can then subtract the shipping cost paid by the seller from the credit card of the buyer agent.

The organisation can monitor the agents' interactions, determine the commitments of agents, and ensure that the agents fulfill their commitments—or otherwise take necessary measures such as putting the violating agent on a blacklist. The organisation, as an exogenous process, cannot *intervene* in the decision making of individual agents by either disallowing them to perform actions or forcing them to perform specific actions, i.e., agents are autonomous and decide their own actions.

2.1 Agent Interactions

We begin by identifying possible actions that agents can perform to interact with each other or with their shared environment. These include pure communication actions, as well as non-communicative actions that change the actual state of their environment. It is important to notice that it is not our purpose to define an Agent Communication Language (ACL) or a communications protocol [14,1]. Instead, we select a representative set of actions that influence the generation and state of commitments. The following six actions will prove sufficient for justifying the adequacy of our organisational model for managing and enforcing commitments. We will use variables $x,\ y$, etc. to range over the agent names $i,\ j$, etc.; propositional variables $p,\ q$, etc. to range over propositions $a,\ b$, etc.; and finally $d,\ d'$, etc. to range over deadlines $t_m,\ t_n$, etc. where $m, n \in \mathbb{N}$.

- *offer*(x, y, p, q, d_1, d_2) — x tells y that x will make q true in the environment by deadline d_2 if p becomes true in the environment by deadline d_1
- *tell*(x, y, p) — x tells y that p is true in the environment
- *cancel*(x, y, q) — x tells y that x will not make q true in the environment
- *release*(y, x, q) — y tells x that x need not make q true in the environment
- *failure*(x, y, p) — x tells y that p cannot be made true in the environment
- *do*(x, p) — x performs an action to make proposition p true in the environment

In our running example, agent i can *offer* agent j to *do* the payment $p_{(b,20)}$ before time t_5 if agent j sends him the book $s_{(b,i)}$ before time t_2. Agent i can offer this deal to agent j by performing the following action: *offer*$(i, j, s_{(b,i)}, p_{(b,20)}, t_2, t_5)$.

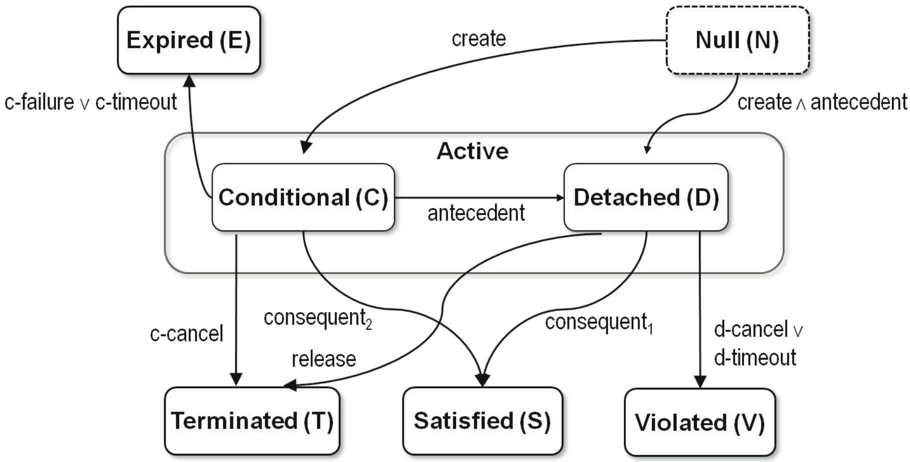

Fig. 1. State transitions of commitment lifecycle

2.2 Commitments

A *joint commitment* (hereafter *commitment*) is defined as a tuple $C = \mathsf{C}(x, y, p, q, d_1, d_2)$. Agent x (as debtor) tells agent y (as creditor) that if proposition p (the antecedent) is brought about by deadline d_1 then x will bring about q (the consequent) by deadline d_2. In the rest of the paper, we assume that deadline d_1 will be satisfied before deadline d_2. It is important to note that C is neither a second-order predicate nor a modal operator and that we do not aim at devising a logic to reason about the internal structure of commitments. $\mathsf{C}(x, y, p, q, d_1, d_2)$ is an atomic proposition denoting a specific commitment.

Fig. 1 shows the states of a lifecycle of a commitment, adapted from [21] (we omit suspension and delegation of commitments). Boxes indicate states and arrows indicate transitions. A commitment, once created, moves to state Conditional. Should the antecedent become true, the commitment moves to Detached. Should the consequent become true, it moves to Satisfied. However, should the antecedent not become true by d_1, then the commitment is Expired. Likewise if the consequent of a detached commitment does not become true by d_2, then the commitment is Violated. It is likewise Violated if x cancels a detached commitment. The commitment is Terminated if x cancels it before it is detached, or y releases x from C once it has been detached.

We will write commitment state with superscript, i.e., $\mathsf{C}^{\mathrm{state}}$. It will be useful to distinguish violation because of cancel (vc) and violated because of timeout (vt).

2.3 Organisation

An organisation is specified by facts, norms (including commitments), and sanctions. We distinguish *brute* and *institutional* facts. Brute facts denote the state of the shared environment (e.g., b_j denoting the fact that agent j has book b or $p_{(b,20)}$ denoting the fact that 20 euro is paid for book b), while institutional facts denote the normative state of an organisation (e.g., $\mathsf{C}^c(i, j, s_{(b,i)}, p_{(b,20)}, t_2, t_5)$ denoting the fact that agent i is committed to pay 20 euro before t_5 if agent j sends book b before t_2). In the following, we assume Π_b and Π_i to be finite disjoint sets of brute and institutional facts (constructed by two disjoint sets of brute and institutional atomic propositions), respectively. Moreover, we follow [12] and represent norms by means of *counts-as* rules that relate brute and institutional facts. The original version of the counts-as construct is of the form "ϕ counts as ψ in the context c". We represent a counts-as construct as a rule $\phi \wedge c \Longrightarrow_{cr} \psi$, where $\phi, c \in \Pi_b$ and $\psi \in \Pi_i$. For example, an offer by agent i to agent j to do a payment if j sends i a book *counts-as* the creation of a conditional commitment. In the following, we use counts-as rules to evaluate and determine the institutional consequences of a certain environment state. In the more general framework proposed in [12], the antecedent of counts-as rules can include institutional facts as well. This allows institutional facts to be created based on other institutional facts making it possible to specify more complex norms such as country-to-duty norms. Finally, we represent sanctions by rules of the form $\phi \Longrightarrow_{sr} \psi$, where $\phi \in \Pi_b \cup \Pi_i$ and $\psi \in \Pi_b$. Note the difference between counts-as and sanctions rules: the first relates brute to institutional facts, while the latter relates institutional facts to brute facts. In our running example, a delay in payment by agent i beyond day 5 will be considered as a violation and will be sanctioned by having agent i on the organisation's blacklist.

We would like to emphasize that the consequent of the sanction rules (in this case having an agent on the blacklist) are brute facts and not actions. In our proposed model, an organisation imposes sanctions by updating its brute state with the consequents of the applicable sanction rules. It is also important to note that sanction rules are not meant to create other commitments, but they are meant to realize the final punishments. The creation of new commitments based on other commitments can in principle be modelled through counts-as rules that allow institutional facts in their antecedents.

Definition 1. *An* organisation *is specified as* (F, cr, sr)*, where* $\mathsf{F} \subseteq \Pi_b$ *is a set of initial brute facts,* cr *is a set of counts-as rules, and* sr *is a set of sanction rules.* □

In our running example, the organisation is initially specified by some brute facts such as agent j wants to sell a book b, and the blacklist of the organisation is empty, i.e., $b_j \in \mathsf{F}$. The set of facts will change during the execution of multi-agent system, based on the interaction between agents, e.g., the fact $p_{(b,20)}$ will be added to F when an agent pays 20 euro for book b. Note that the institutional facts such as commitments will be generated only during the execution of multi-agent system as a consequence of agents' interactions. Although an organisation can be specified in terms of arbitrary norms [12], we suppose that specific agents' interaction creates and manipulate social commitments (institutional facts) and claim that manipulation of commitments should respect a specific set of norms. Therefore, we focus on commitment-based norms and represent them as counts-as rules defined in terms of specific actions.

Fig. 2 illustrates a set of commitment-based norms, represented as counts-as rules. We next explain these rules, which specify how interactions between agents operate on social commitments. The application of counts-as rules (and the resulting removal and addition of commitments) are explained later in this section when we present the operational semantics of our organisational model in transition rules **1**–**4**, which are introduced in Section 3, below.

1 $\mathit{offer}(x, y, p, q, d_1, d_2) \Longrightarrow_{cr} \mathsf{C}^c(x, y, p, q, d_1, d_2)$
2 $\mathit{tell}(y, x, p) \wedge \mathsf{C}^c(x, y, p, q, d_1, d_2) \wedge \neg d_1 \wedge p \Longrightarrow_{cr} \mathsf{C}^d(x, y, p, q, d_1, d_2)$
3 $\mathit{do}(y, p) \wedge \mathsf{C}^c(x, y, p, q, d_1, d_2) \wedge \neg d_1 \wedge p \Longrightarrow_{cr} \mathsf{C}^d(x, y, p, q, d_1, d_2)$
4 $\mathit{tell}(x, y, q) \wedge \mathsf{C}^c(x, y, p, q, d_1, d_2) \wedge \neg d_1 \wedge q \Longrightarrow_{cr} \mathsf{C}^s(x, y, p, q, d_1, d_2)$
5 $\mathit{do}(x, q) \wedge \mathsf{C}^c(x, y, p, q, d_1, d_2) \wedge \neg d_1 \wedge q \Longrightarrow_{cr} \mathsf{C}^s(x, y, p, q, d_1, d_2)$
6 $\mathit{tell}(x, y, q) \wedge \mathsf{C}^d(x, y, p, q, d_1, d_2) \wedge \neg d_2 \wedge q \Longrightarrow_{cr} \mathsf{C}^s(x, y, p, q, d_1, d_2)$
7 $\mathit{do}(x, q) \wedge \mathsf{C}^d(x, y, p, q, d_1, d_2) \wedge \neg d_2 \wedge q \Longrightarrow_{cr} \mathsf{C}^s(x, y, p, q, d_1, d_2)$
8 $\mathit{cancel}(x, y, q) \wedge \mathsf{C}^d(x, y, p, q, d_1, d_2) \wedge \neg d_2 \wedge \neg q \Longrightarrow_{cr} \mathsf{C}^{vc}(x, y, p, q, d_1, d_2)$
9 $\mathsf{C}^d(x, y, p, q, d_1, d_2) \wedge d_2 \wedge \neg q \Longrightarrow_{cr} \mathsf{C}^{vt}(x, y, p, q, d_1, d_2)$
10 $\mathit{failure}(y, x, p) \wedge \mathsf{C}^c(x, y, p, q, d_1, d_2) \wedge \neg d_1 \wedge \neg p \Longrightarrow_{cr} \mathsf{C}^e(x, y, p, q, d_1, d_2)$
11 $\mathsf{C}^c(x, y, p, q, d_1, d_2) \wedge d_1 \wedge \neg p \Longrightarrow_{cr} \mathsf{C}^e(x, y, p, q, d_1, d_2)$
12 $\mathit{cancel}(x, y, q) \wedge \mathsf{C}^c(x, y, p, q, d_1, d_2) \wedge \neg d_1 \wedge \neg p \Longrightarrow_{cr} \mathsf{C}^t(x, y, p, q, d_1, d_2)$
13 $\mathit{release}(y, x, p) \wedge \mathsf{C}^d(x, y, p, q, d_1, d_2) \wedge \neg d_2 \wedge p \Longrightarrow_{cr} \mathsf{C}^t(x, y, p, q, d_1, d_2)$

Fig. 2. Counts-as rules specify the lifecycle of commitments based on actions and deadlines

- Performing action "x offers to y that x realizes q before d_2 if y realizes p before d_1" counts as creation of a conditional commitment (superscript c denotes conditional state of commitment; similar convention is used for other commitment states). The application of rule #1 by the organisation adds institutional fact $\mathsf{C}^c(x, y, p, q, d_1, d_2)$ to the institutional facts.
- Performing action "y tells x that p is realized" or "y does act and realizes p" when d_1 is still not passed counts as detaching the conditional commitment. The application of these rules #2 or #3 removes the conditional commitment from institutional facts and adds a corresponding detached commitment to it.[2]
- Performing action "x tells y that q is realized" or "x does act and realizes q" when d_1 is still not passed counts as satisfying the conditional commitment. The application of these rules #4 or #5 removes the conditional commitment from institutional facts and adds a corresponding satisfied commitment to it.
- Performing action "x tells y that q is realized" or "x does act and realizes q" when d_2 is still not passed counts as satisfying the detached commitment. The application of these rules #6 or #7 removes the detached commitment from institutional facts and adds a corresponding satisfied commitment to it.
- Performing action "x cancels to realizes q" when d_2 is still not passed counts as the violation of the detached commitment. The application of this rule #8 removes the detached commitment and adds a corresponding violated commitment.
- Elapsing deadline d_2 counts as the violation of a detached commitment. The application of this rule #9 removes the detached commitment from institutional facts and adds a corresponding violated commitment to it.
- Performing action "y fails to realize p" when d_1 is still not passed counts as expiration of the conditional commitment. The application of this rule #10 removes the conditional commitment and adds a corresponding expired commitment.

[2] Detaching a commitment based on telling assumes that y is a trusted agent, i.e., its utterances are according to its beliefs. Note that an organisation may develop a list of trusted agents.

- Elapsing deadline d_1 counts as expiration of a conditional commitment. The application of this rule #11 removes the conditional commitment from institutional facts and adds a corresponding expired commitment to it.
- Performing action "x cancels to realize q" when d_1 is still not passed counts as termination of the conditional commitment. The application of this rule #12 removes the conditional commitment and adds a corresponding terminated commitment to it.
- Performing action "y releases a detached commitment after p has been satisfied" when d_2 is still not passed counts as termination of the conditional commitment. The application of this rule #13 removes the conditional commitment from institutional facts and adds a corresponding terminated commitment to it.

Finally, organisations can be specified in terms of arbitrary sanctions represented by *sanction rules* [7]. Sanctions are defined in terms of specific violations and determine how a violated system state can be turned back to a 'normal' state by means of a system update. In contrast to commitments, sanctions are not generic and depend on application in hand. For example, in our marketplace organisation, a timed-out commitment caused by agent i who has failed to make a payment will be sanctioned by blacklisting the agent. This sanction can be represented as:

$$\mathsf{C}^{vt}(i, j, s_{(b,i)}, p_{(b,20)}, t_2, t_5) \Longrightarrow_{sr} \mathit{blacklist}_i$$

Another possible sanction can be designed to cope with a detached commitment that is cancelled by its debtor. For example, suppose agent j sends the book before deadline t_2 after which agent i cancels the commitment. The *cancel* action by i violates the detached commitment $\mathsf{C}^d(i, j, s_{(b,i)}, p_{(b,20)}, t_2, t_5)$, turning it into the canceled commitment $\mathsf{C}^{vc}(i, j, s_{(b,i)}, p_{(b,20)}, t_2, t_5)$. One may want to sanction such a violation by charging i the shipping cost paid by j (and possibly some additional administration costs). Such a sanction can be represented by the following rule, which indicates that agent i should pay the shipping cost ($\mathit{chargedShippingCost}_{(i,5)}$):

$$\mathsf{C}^{vc}(i, j, s_{(b,i)}, p_{(b,20)}, t_2, t_5) \Longrightarrow_{sr} \mathit{chargedShippingCost}_{(i,5)}$$

3 Operational Semantics

We now give the operational semantics for the normative organisation. The specification of an organisation determines the its initial state. The execution of an organisation is determined by a set of transition rules that specify possible transition steps. In the following, we first define the (initial) states of an organisation, followed by the set of four transition rules.

Definition 2. *The* state of an organisation *is specified as the tuple $\langle \sigma_b, \sigma_i, cr, sr\rangle$, where $\sigma_b \subseteq \Pi_b$, $\sigma_i \subseteq \Pi_i$, cr is the set of counts-as rules, and sr is the set of sanction rules. Since the counts-as and sanction rules do not change during the execution of organisations, we omit cr and sr from the organisation states, representing state as $\langle \sigma_b, \sigma_i\rangle$. Let (F, cr, sr) be the initial specification of an organisation. The* initial state *of the organisation is $\langle \mathsf{F}, \emptyset\rangle$, i.e., the initial set of institutional facts is the empty set.* $\square$

In order to compute the set of commitments (and sanctions) that should be generated or modified by the counts-as rules cr (and the sanction rules sr) in a given state of an organisation, we follow [7] and define the *closure* of a set of propositions under a set of rules. Let $\mathsf{cond}(()r)$ and $\mathsf{cons}(()r)$ denote the condition and consequent of rule r, respectively. First, we determine the rules R that are applicable to a set of propositions X as:

$$\mathsf{App}(X, R) = \{r \in R \mid X \models \bigwedge \mathsf{cond}(r)\}$$

Using function $\mathsf{cl}_X^R(Y) = \{\mathsf{cons}(r) \mid r \in \mathsf{App}(X \cup Y, R)\}$ we define $\mathsf{cl}_X^R\!\uparrow^\omega$ as the smallest fixed-point of $\mathsf{cl}_X^R(\cdot)$, which exists due to Knaster/Tarski's fixed point theorem [7]. This fixed point provides the set of heads of all applicable rules.

Let cr be the set of counts-as rules presented in Fig. 2 and $\mathsf{cl}_X^{cr}\!\uparrow^\omega$ be the closure of set X under cr. The following transition rule specifies the interaction between two agents through communication action α. Note that the organisation updates the institutional facts (e.g., commitments) based on the performed communication action and by applying the above counts-as rules.

$$\frac{com(\alpha)\ \&\ \sigma_i' = \oplus(\sigma_i \cup \sigma_b \cup \{\alpha\})}{\langle \sigma_b, \sigma_i \rangle \rightarrow_c \langle \sigma_b, \sigma_i' \rangle} \tag{1}$$

1. $com(\alpha)$ indicates that α is the communicated message,
2. $\oplus(X) \ = \sigma_i \setminus \{c^x(\boldsymbol{V}) \mid c^y(\boldsymbol{V}) \in (\mathsf{cl}_X^{cr}\!\uparrow^\omega \setminus X)\} \cup \ (\mathsf{cl}_X^{cr}\!\uparrow^\omega \setminus X)$,
   for $x, y \in \{c, d, s, vc, vt, e, t\}$.

According to this transition rule, the state of a multi-agent organisation can make a transition when message α is communicated. In this transition the institutional state σ_i of the organisation is updated by applying the counts-as rules to determine new institutional facts. The new institutional facts are computed by taking the closure of institutional facts, brute facts, and the performed action under the counts-as rules, i.e., $\mathsf{cl}_X^{cr}\!\uparrow^\omega \setminus X$. Note that the original set of institutional facts, brute facts, and the performed action denoted by X are *removed* from the closure to obtain the new institutional facts. The second item in the rule specifies the update of the institutional facts by adding new institutional facts to σ_i while removing the corresponding institutional facts from σ_i. This operation guarantees that the state of commitments ($x, y \in \{c, d, s, vc, vt, e, t\}$) changes according to the transitions in the commitment lifecycle depicted in Fig. 1. Lastly, note that $\mathsf{cl}_X^{cr}\!\uparrow^\omega \setminus X = \emptyset$ if none of the counts-as rules are applicable. In our running example, the performance of action *offer*$(i, j, s_{(b,i)}, p_{(b,20)}, t_2, t_5)$ in an organisation state $\langle \sigma_b, \sigma_i \rangle$ causes a transition to state $\langle \sigma_b, \sigma_i \cup \{\mathsf{C}^c(i, j, s_{(b,i)}, p_{(b,20)}, t_2, t_5)\}\rangle$ where a commitment is created. This is accomplished by the application of rule 1.

The next transition rule specifies the performance of *non-communicative* actions, such as an agent sending a book. Note that counts-as rules 3 and 5 cover the cases where performing a non-communicative action by some agent counts as changing the state of a commitment.

$$\frac{act(\alpha)\ \&\ \sigma_b' = \otimes(\sigma_b, \alpha)\ \&\ \sigma_i' = \oplus(\sigma_i \cup \sigma_b' \cup \{\alpha\})}{\langle \sigma_b, \sigma_i \rangle \rightarrow_a \langle \sigma_b', \sigma_i' \rangle} \tag{2}$$

1. $act(\alpha)$ indicates that α is the action,
2. $\otimes$ is an update operation that changes the state of environment σ_b after performing α. We assume the existence of such an update operator.
3. $\oplus$ is the operator for updating institutional facts as defined above.

The new institutional facts σ_i' are determined based on σ_b', the result of realizing the effect of α on σ_b. In our running example, sending the book b by agent j (i.e., we have that $\alpha = s_{(b,i)}$) causes a transition of the organisation state from $\langle \sigma_b, \sigma_i \rangle$, where we have $\sigma_i \models \mathsf{C}^c(i, j, s_{(b,i)}, p_{(b,20)}, t_2, t_5)$, to $\langle \sigma_b', \sigma_i' \rangle$ with $\sigma_i' \not\models \mathsf{C}^c(i, j, s_{(b,i)}, p_{(b,20)}, t_2, t_5)$ and $\sigma_i' \models \mathsf{C}^d(i, j, s_{(b,i)}, p_{(b,20)}, t_2, t_5)$. Of course, we assume this is only possible if $\sigma_b \models s_{(b,i)} \wedge \neg t_2$ indicating that the sending action is indeed performed within the corresponding deadline t_2. This is accomplished by the application of counts-as rule 3.

Finally, we allow the environment state to change by the internal mechanism of the environment, e.g., the state of a clock changes automatically. This is essential for the application of counts-as rules 7 and 9, which are applicable when the deadline elapses.

$$\frac{\sigma_b \to \sigma_b' \;\&\; \sigma_i' = \oplus(\sigma_i \cup \sigma_b')}{\langle \sigma_b, \sigma_i \rangle \to_e \langle \sigma_b', \sigma_i' \rangle} \tag{3}$$

This transition rule ensures that elapse of the deadlines are possible and changes the state of the institutional facts accordingly. Let the organisation of our running example be in state $\langle \sigma_b, \sigma_i \rangle$ where $\sigma_i \models \mathsf{C}^d(i, j, s_{(b,i)}, p_{(b,20)}, t_2, t_5)$, $\sigma_b \not\models p_{(p,20)}$, and the environment makes a transition such that $\sigma_b' \models d_5$. In this case, we have $\sigma_i' \not\models \mathsf{C}^d(i, j, s_{(b,i)}, p_{(b,20)}, t_2, t_5)$ and $\sigma_i' \models \mathsf{C}^{vt}(i, j, s_{(b,i)}, p_{(b,20)}, t_2, t_5)$. This is accomplished by the application of the counts-as rule 7.

Sanctions. A common scenario is that the organisation imposes sanctions when commitments are violated. For example, a non-paying agent is added to a blacklist. We represent sanctions by rules connecting commitment violations to a specific brute state of the organisation. In order to allow context-dependent sanctions, the antecedent of rules can be composed of both brute and institutional facts. Given a set of sanction rules sr and the state of an organisation $\langle \sigma_b, \sigma_i \rangle$, we define the enforcement of sanctions as the closure of state under the sanction rules, denoted by $\mathsf{cl}^{sr}_{(\sigma_b \cup \sigma_i)}{\uparrow^\omega} \setminus \sigma_i$. The following transition rule ensures that the organisation imposes sanctions when commitments are violated.

$$\frac{\sigma_b' = \mathsf{cl}^{sr}_{(\sigma_b \cup \sigma_i)}{\uparrow^\omega} \setminus \sigma_i \;\;\&\;\; \sigma_i' = \sigma_i \setminus \mathit{sanctioned}(\sigma_i, \sigma_b', sr)}{\langle \sigma_b, \sigma_i \rangle \to_s \langle \sigma_b', \sigma_i' \rangle} \tag{4}$$

The function $\mathit{sanctioned}(\sigma_i, \sigma_b', sr)$ removes all violated commitments from σ_i for which sanctions are imposed. Consider again the sanction for the payment violation in our running example, which is represented by the rule

$$\mathsf{C}^{vt}(i, j, s_{(b,i)}, p_{(b,20)}, t_2, t_5) \Longrightarrow_{sr} \mathit{blacklist}_i$$

and let the organisation be in the state $\langle \sigma_b, \sigma_i \rangle$ where $\sigma_i \models \mathsf{C}^{vt}(i, j, s_{(b,i)}, p_{(b,20)}, t_2, t_5)$. The application of this transition rule causes a transition to the state $\langle \sigma_b', \sigma_i' \rangle$ where $\sigma_b' \models \mathit{blacklist}_i$ (i.e., i is blacklisted) and $\sigma_i' \not\models \mathsf{C}^{vt}(i, j, s_{(b,i)}, p_{(b,20)}, t_2, t_5)$.

Multi-agent System Execution. Recall that an organisation has two main processes to perform: (1) monitoring the interaction between agents, and between agents and the environment, to detect norm violations; and (2) enforcing norms by means of sanctions when violations are detected. In our framework, transitions **1–3** are responsible for the monitoring task while transition **4** is responsible for imposing sanctions. The executions of a multi-agent organisation are determined by the transition system which is specified by the transition rules **1–4**; the system consists of all possible computational runs.

Definition 3. *Given transition rules **1–4** and an initial state of an organisation* $c_0 = \langle \sigma_b, \sigma_i \rangle$*, a* computational run $CR(c_0)$ *is an infinite sequence* $c_0, c_1, \ldots$ *where* c_i *is a multi-agent organisation state and* $\forall_{i>0} : c_i \rightarrow c_{i+1}$ *is a transition derived by applying transition rules **1–4**. The* execution *of a multi-agent organisation with initial state* c_0 *is a set of all possible computational runs* $CR(c_0)$*.* □

Possible computational runs of an execution are due to different orders of the applications of transition rules **1–4**. One may constrain the set of possible computational runs by specifying one specific order for applying transition rules. For example, one order may apply transition rules **1–4** consecutively while a second order applies transition rules **1–3** interspersed by the application of transition rule **4**. According to the first order, the organisation allows (1) agents to communicate, (2) agents to interact with the environment, (3) the environment to make a transition, and (4) the organisation to impose sanctions. By contrast, the second order imposes sanctions immediately after each of the three activities. Note that the first order allows different violations to occur before sanctions are imposed while the second order applies the sanctions immediately. The consequences of these different orders will be studied in the next section.

4 Example

Chopra et al. [9] describe an insurance claim process involving a vehicle repair, with the actors being the driver claimant (assumed to be not at fault, and whose policy has no deductible), the driver's insurance company, a car repair garage, and a damage assessor. In our example scenario, the agents are a Customer (who owns the car), an Insurer (who pays for repairs), a Repairer (who conducts the repair), and an Assessor (who decides how much the repair should cost). Commitments between agents model the business protocols of the process:[3]

- C_1: insurer to repairer: if insurance has been validated and the repair has been reported, then the insurer will have paid and approved the assessment within 7 days
- C_2: insurer to assessor: if the assessment has been done, the assessment will have been paid within 5 days
- C_3: assessor to repairer: if damages have been reported and the insurance has been validated, a damage assessment will have been performed within 2 days
- C_4: repairer to customer: if the insurance has been validated and the car was damaged, then the car will have been repaired within 8 days

[3] Note there is an 'implied' commitment from insurer to customer: if the insurance has been validated and the car was damaged, then the car will have been repaired.

– C_5: insurer to customer: if the premium has been paid, then the insurance will have been validated within 8 days

Hence, formally, we have a set of social commitments as follows. Note that we specify relative deadlines for the consequent of commitments using the notation $+x$ to indicate x days after the antecedent becomes true. $\emptyset$ indicates an absence of a deadline.

– $C_1 = \mathsf{C}(\mathsf{Insurer}, \mathsf{Repairer}, \text{insurance-validated} \wedge \text{repair-reported}, \text{assessment-approved} \wedge \text{payment-done}, \emptyset, +7)$
– $C_2 = \mathsf{C}(\mathsf{Insurer}, \mathsf{Assessor}, \text{assessment-done}, \text{assessment-paid}, \emptyset, +5)$
– $C_3 = \mathsf{C}(\mathsf{Assessor}, \mathsf{Repairer}, \text{damages-reported} \wedge \text{insurance-validated}, \text{assessment-done}, \emptyset, +2)$
– $C_4 = \mathsf{C}(\mathsf{Repairer}, \mathsf{Customer}, \text{insurance-validated} \wedge \text{car-damaged}, \text{car-repaired}, \emptyset, +8)$
– $C_5 = \mathsf{C}(\mathsf{Insurer}, \mathsf{Customer}, \text{premium-paid}, \text{insurance-validated}, \emptyset, +8)$

In the scenario, the car owner reports an accident to her insurance company, and takes the car to a garage. On certification that the insurance is valid (because the customer has paid the insurance premium), the repair garage accepts the damaged car and contacts the assessor. Since the insurance is valid, the repair garage commences the repair. The assessor reports to the insurance company, which approves the assessment of damage and pays the assessor for its work. The repair garage reports to the insurance company when it has completed the repair, and with the approval of the insurer, the garage then tells the customer that the car is ready. The insurer pays the repair garage.

Consider now a set of agent actions corresponding to this scenario. These actions are manifest in communication actions, and the impact of the actions upon the commitment store. The trace in Table 1 begins with the *offer*s that create the five commitments. Column Rule gives the Counts-as Rule. We abbreviate the logical propositions by their initials, e.g., pp for premium-paid. In line 6, the customer proves payment of the insurance. In line 9, she reports car damage. In line 14, the garage has completed the repair. In line 17, the garage tells the customer the car is ready.

5 Properties

The propositions in this section show various kind of properties following from our approach and operational semantics. We describe them but leave the proofs and further formalization for future work.

5.1 Temporal Properties

We study first *temporal properties*, i.e., the dynamics of the commitment states. These properties illustrate that the commitment states follow the commitment lifecycle in Fig. 1. In this subsection, we assume that each *offer*(x, y, p, q, d_1, d_2) is done with unique p and q, such that there is no interference. We discuss the lifting of this assumption in the following subsection.

We begin by stating continuity properties about conditional and detached commitments: they hold until they are fulfilled or the deadline is reached. We use Linear Time Logic LTL with the *release* operator, not to be confused with the release communication action of Section 2.1: ϕ releases ψ if ψ is true until the first position in which ϕ is true (or forever if such a position does not exist). We assume execution traces to be fair.

Table 1. Execution trace in car insurance scenario

Step	Agent	Communication Act	Time	Rule	Institutional State	Brute Facts
0	–	–	0	–	$\{\}$	$\{\}$
1	I	*offer*(I, R, iv ∧ rr, aa ∧ pd)	1	1	$\{C_1^C\}$	$\{\}$
2	I	*offer*(I, A, ad, ap)	2	1	$\{C_1^C, C_2^C\}$	$\{\}$
3	A	*offer*(A, R, dr ∧ iv, ad)	3	1	$\{C_1^C, C_2^C, C_3^C\}$	$\{\}$
4	R	*offer*(R, C, iv, cr)	4	1	$\{C_1^C, C_2^C, C_3^C, C_4^C\}$	$\{\}$
5	I	*offer*(I, C, pp, iv)	5	1	$\{C_1^C, C_2^C, C_3^C, C_4^C, C_5^C\}$	$\{\}$
6	C	*do*(C, pp)	6	3	$\{C_1^C, C_2^C, C_3^C, C_4^C, C_5^D\}$	$\{pp\}$
7	C	*do*(I, iv)	7	5	$\{C_1^C, C_2^C, C_3^C, C_4^C, C_5^S\}$	$\{iv, pp\}$
8	–	–	8	–	$\{C_1^C, C_2^C, C_3^C, C_4^C, C_5^S\}$	$\{cd, iv, pp\}$
9	C	*tell*(C, R, iv ∧ cd)	9	2	$\{C_1^C, C_2^C, C_3^C, C_4^D, C_5^S\}$	$\{cd, iv, pp\}$
10	R	*tell*(R, A, dr ∧ iv)	10	2	$\{C_1^C, C_2^C, C_3^D, C_4^D, C_5^S\}$	$\{cd, dr, iv, pp\}$
11	A	*do*(A, ad)	12	5	$\{C_1^C, C_2^C, C_3^S, C_4^D, C_5^S\}$	$\{ad, cd, dr, iv, pp\}$
12	A	*tell*(A, I, ad)	12	2	$\{C_1^C, C_2^D, C_3^S, C_4^D, C_5^S\}$	$\{ad, cd, dr, iv, pp\}$
13	I	*tell*(I, A, ap)	16	4	$\{C_1^C, C_2^S, C_3^S, C_4^D, C_5^S\}$	$\{ad, ap, cd, dr, iv, pp\}$
14	R	*do*(R, rr)	16	3*	$\{C_1^C, C_2^S, C_3^S, C_4^D, C_5^S\}$	$\{ad, ap, cd, dr, iv, pp, rr\}$
15	R	*tell*(R, I, iv)	17	2*	$\{C_1^D, C_2^S, C_3^S, C_4^D, C_5^S\}$	$\{ad, ap, cd, dr, iv, pp, rr\}$
16	I	*tell*(I, R, aa)	17	4*	$\{C_1^D, C_2^S, C_3^S, C_4^D, C_5^S\}$	$\{aa, ad, ap, cd, dr, iv, pp, rr\}$
17	R	*tell*(R, C, cr)	18	2	$\{C_1^D, C_2^S, C_3^S, C_4^S, C_5^S\}$	$\{aa, ad, ap, cd, cr, dr, iv, pp, rr\}$
18	I	*do*(I, pd)	21	5*	$\{C_1^S, C_2^S, C_3^S, C_4^S, C_5^S\}$	$\{aa, ad, ap, cd, dr, iv, pd, pp, rr\}$

Proposition 1. *If the organisation updates the institutional facts based on the performed communication action as follows:*

- offer(x, y, p, q, d_1, d_2)*, then* $p \vee q \vee d_1$ *releases* $\mathsf{C}^c(x, y, p, q, d_1, d_2)$*.*
- *if* tell(y, x, p)*, and* $\mathsf{C}^c(x, y, p, q, d_1, d_2)$ *holds, then* $\neg d_1 \wedge p \rightarrow (q \vee d_2)$ *releases* $\mathsf{C}^d(x, y, p, q, d_1, d_2)$
- *if we have* do(y, p)*, and* $\mathsf{C}^c(x, y, p, q, d_1, d_2) \wedge \neg d_1 \wedge p$ *holds, then* $q \vee d_2$ *releases* $\mathsf{C}^d(x, y, p, q, d_1, d_2)$

Second, we can state continuity properties about reaching termination states: each conditional and detached commitment state will lead to precisely one termination state.

Proposition 2. *For any commitment, we have at most one of* $\mathsf{C}^c(x, y, p, q, d_1, d_2)$*,* $\mathsf{C}^d(x, y, p, q, d_1, d_2)$*,* $\mathsf{C}^s(x, y, p, q, d_1, d_2)$*,* $\mathsf{C}^{vc}(x, y, p, q, d_1, d_2)$*,* $\mathsf{C}^{vt}(x, y, p, q, d_1, d_2)$*,* $\mathsf{C}^e(x, y, p, q, d_1, d_2)$*, and* $\mathsf{C}^t(x, y, p, q, d_1, d_2)$ *true at any time.* □

Third, we can state continuity properties about some of the termination states: once they are true, they will stay true. This does not hold for the violation states, as the violation will be removed once the sanction is applied.

Proposition 3. *If* $\mathsf{C}^s(x, y, p, q, d_1, d_2)$*,* $\mathsf{C}^e(x, y, p, q, d_1, d_2)$*, or* $\mathsf{C}^t(x, y, p, q, d_1, d_2)$ *hold, then they will hold forever.* □

Further temporal properties can be defined to illustrate that the operational semantics behaves according to the commitment lifecycle in Fig. 1. Since the characterization of these properties is relatively straightforward, we turn here to two other classes of properties which can be defined for our operational semantics.

5.2 Interference

Figure 1 is a bit misleading in the sense that in practice there is not a single commitment cycle at the same time, but many of them operate in parallel. In this section we consider whether one commitment cycle can affect another one; we call this interference. We now consider the lifting of the assumption that each *offer*(x, y, p, q, d_1, d_2) is done with unique p and q. For example, consider the situation where an agent i makes two consecutive *offer* statements, *offer*$(i, j, s_{(b,i)}, p_{(b,20)}, t_2, t_5)$ and *offer*$(i, j, t_{(b,i)}, p_{(b,20)}, t_2, t_5)$.

Proposition 4. *If the organisation updates the institutional facts based on a performed communication action* tell(y, x, p) *or an action* do(y, p)*, then it can detach or satisfy multiple commitments at once.* □

Example 1. Assume the organisation updates the institutional facts based on the following actions: *offer*$(i, j, s_{(b,i)}, p_{(b,20)}, t_2, t_5)$, *offer*$(i, j, t_{(b,i)}, p_{(b,20)}, t_2, t_5)$, *do*$(j, s_{(b,i)})$, and *do*$(j, t_{(b,i)})$. This leads to two detached commitments, namely $\mathsf{C}^d(i, j, s_{(b,i)}, p_{(b,20)}, t_2, t_5)$ and $\mathsf{C}^d(i, j, t_{(b,i)}, p_{(b,20)}, t_2, t_5)$. Performing *do*$(i, p_{(b,20)})$ by agent i will now move both detached commitments to satisfied commitments, i.e., agent i de-commits itself by paying 20 euro instead of 40 euro.[4] □

[4] Another similar example is when an agent makes two offers that differ only in the deadlines. In that case, the commitment with the earlier deadline ‘subsumes’ the one with the later deadline.

This example illustrates that the agents must be careful to distinguish propositions. In this case, the agents should syntactically distinguish the two propositions referring to pay 20 euro, and it is left to the agent to prevent this undesired interference.

5.3 Redundancy

Redundancy properties concern communication actions that, if omitted as part of a sequence of communication actions, do not change the final set of institutional facts [3]. We confine ourselves to an example of a redundancy property, as follows.

Consider a sanction system where agents are blacklisted regardless of the commitment they violated, i.e., containing the following two rules:

$$\mathsf{C}^{vt}(x, y, p, q, d_1, d_2) \Longrightarrow_{sr} add_{y,BL}$$
$$\mathsf{C}^{vc}(x, y, p, q, d_1, d_2) \Longrightarrow_{sr} add_{y,BL}$$

In the context of our running example, we say that an offer is *blacklist redundant* if it does not change the situations in which an agent will be blacklisted. Given a sequence of offers, which offers can be made in addition which are in this sense blacklist redundant?

Proposition 5. – *If* $\mathsf{C}^{c}(x, y, p, q, d_1, d_2)$*, then* offer$(x, y, p \wedge r, q, d_1, d_2)$ *is blacklist redundant. If I have offered you that if you do* p *I will do* q*, then it is redundant to offer you that if you do* $p \wedge r$ *then I will do* q*, under the same conditions.*
– *If* $\mathsf{C}^{c}(x, y, p, q \wedge r, d_1, d_2)$*, then* offer$(x, y, p, q, d_1, d_2)$ *is blacklist redundant. If I have offered you that if you do* p *I will do* $q \wedge r$*, then it is redundant to offer you that I will do* q*, under the same conditions.*
– *If* $\mathsf{C}^{c}(x, y, p, q, d_1, d_2) \wedge \mathsf{C}^{c}(x, y, r, q, d_1, d_2)$*, then* offer$(x, y, p \vee r, q, d_1, d_2)$ *is blacklist redundant. If I have offered you that if you do* p *I will do* q*, and if you do* r *I will do* q*, then it is redundant to offer you that if you do* $p \vee r$ *then I will do* q*, under the same conditions.*
– *If* $\mathsf{C}^{c}(x, y, p, q, d_1, d_2) \wedge \mathsf{C}^{c}(x, y, p, r, d_1, d_2)$*, then* offer$(x, y, p, q \wedge r, d_1, d_2)$ *is blacklist redundant. If I have offered you that if you do* p *I will do* q*, and if you do* p *I will do* r*, then it is redundant to offer you that I will do* $q \wedge r$*, under the same conditions.*

The redundancy properties give some deeper insight in the system, but they depend on the sanctioning system adopted by the organisation. The above properties still hold when the sanction system does not depend on the actual commitment that is violated, but only that *some* commitment has been violated. Another way of reasoning says that the sanction depends only on q—not on p, the agents x and y, or the deadlines. This yields a subset of the above properties. A further refinement is to assume relations among sanctions. For example, a sanction for violating the offer to do p should not be more than that for violating the offer to do $p \wedge q$. We leave this matter for future work.

6 Related Work and Discussion

The idea of non-communicative actions and the enforcement of sanctions by monitoring the state of commitments is something common in the agent literature. The schema of our model is standard when dealing with commitments in agents organizations: organizations consist of facts, norms, commitments and sanction, the agents can perform a set of actions which lead to the creation of commitments, and these commitments have rules in form of count-as rules and when the commitments are violated, sanctions apply [20,11,17,14,15,19,8,16,4,1,9,18,10].

Moreover, formalizations are often based on a lifecycle for commitments using an operational semantics. However, the aim is typically to give a semantics for a large variety of speech acts, for example for propose or request, and they therefore have to define their semantics in considerable detail. Further still, the trend is to make the languages ever more complex, for example by introducing higher order commitments, meta-commitments [22], embedded temporal regulations [1], goal of organizations, other norms than commitment-based ones, and so on.

Building on earlier work [11,14], Fornara et al. [15] propose an ACL based on communication actions. They define norms as "rules that manipulate commitments of the agents engaged in an interaction" and thus, unlike our work, define norms as event-driven rules and provide a more limited operational semantics. Our work is further distinguished in generality since we can define a range of operational semantics by changing the constitutive norms. Our purpose is not to define an ACL but to establish how norms are be operationalized in an organisation setting.

Compared to earlier works, our approach differs in the following regards. First, we use counts-as rules explicitly as technical constructs while Fornara et al. [14], for instance, treat counts-as relation primarily as linguistic conventions. Note that counts-as rules in our work can be seen as defeasible rules. Second, we provide an operational semantics for interactions (among which communication actions) within an organisational setting. Fornara et al. provide semi-formal specification of organisations and consider only communication actions. Further, we analyze the properties of interaction within an organisational setting. Third, we consider the effect of non-communicative actions as well as the elapse of deadlines. Fourth we have sanction rules while earlier works leave open the question of what should happens when commitments are violated. Fifth, in contrast to some works, we adopt a contemporary lifecycle of commitments, following Chopra and Singh [10] and precedents.

These last authors are interested in simplifying the specification of commitment-based protocols or requirements. Our aim is in line with the latter paper [10], and the two approaches are orthogonal. Chopra and Singh capture business requirements with commitment-based specifications, and then group these specifications into reusable methods. They consider protocol enactment but not the organisational setting, which is our focus. By contrast, we do not aim to directly model business requirements, but start from the lifecycle of commitments and define a minimal language to make all possible changes to the commitment base.

An alternative approach, that aims for the same level of genericity as we pursue, uses deontic-inspired specifications that are monitored and enforced, such as the recent work of Álvarez-Napagao and colleagues [13]. An interesting discussion for the future is to compare our commitment-based approach versus a deontic-inspired approach.

7 Future Work

Our on-going work is to characterise the properties of the operational semantics and to prove them. In this paper, we considered two principal communication actions: *offer* and *tell*. Alternative semantics can be defined for these actions using other constitutive rules, and compared to the ones we defined. Moreover, our model can be extended to a wider range of communication actions and provide their semantics in terms of social commitments. Extending the set of communication actions should be done carefully since commitments may create unwanted interferences.

Second, commitments, as we have adopted them, possess special slots for deadlines. Alternatively, as proposed by Marengo et al. [18], the deadline could be part of the propositional content (the antecedent and consequent) of the commitment. In that case, the deadline would not have a special status. For example, we could express the commitment that the insurer will have paid the assessment within 7 days, and (further) approved it within 14 days. Whether such composite commitments can be expressed by several commitments in our language, and in general the advantages and disadvantages of the two definitions from the point of view of the organisational setting, is left for further research.

Third, we considered here generic counts-as rules that were designed based on the semantics of specific communication actions (i.e., *offer* and *tell*). A possible extension would be to investigate the interaction between such generic counts-as rules and domain-specific counts-as rules representing domain-specific norms. For example, we can consider norms such as “buying a book while having no money on a credit card counts-as a violation”.

Fourth, a further issue not considered in this paper is the role of interaction protocols in agents’ interactions. The representation of an organisation could be enriched with—in addition to counts-as and sanction rules—interaction protocols that constrain the order of communication actions.

Finally, the dynamics of organisations is an important topic which we have not treated here. We think there is potential in exploring a mechanism to dynamically change the rules of the constitutive norms and sanctions [6,5].

8 Summary

This papers focuses on the use of social commitments as institutional facts. The life-cycle of these commitments can be managed by an organisation based on agents’ interactions, and thus the commitments play a role in coordination of agents and their interaction, through the monitoring and enforcement processes of the organisation.

The state of an organisation consists of brute and institutional facts, a set of counts-as rules, and a set of sanction rules. In addition, we define the closure of a set of propositions under a set of rules. Then we introduce an operational semantics using four transition rules. The first transition rule defines the interaction between two agents through communication actions, the second specifies the performance of non-communicative actions, the third defines the change by the internal mechanism of the environment, such as the state of the clock, and the fourth ensures that the organisation imposes sanctions when commitments are violated. We define a multiagent system execution as a sequence of states of the organisation, where each transition is derived by applying one of the four transition rules.

We define a normative organisation using six actions of the agents: conditional offers to make something true (*offer*), telling that something is true (*tell*), performing an action to make something true (*do*), cancelling one's own offer to make something true (*cancel*), cancelling someone else's offer to make something true (*release*), and informing that something cannot be made true (*failure*). In addition, we define seven states for commitments: Null, Conditional, Detached, Expired, Terminated, Satisfied, and Violated. We introduce thirteen counts-as rules specifying the lifecycle of commitments based on communication actions, non-communicative actions, and deadline expirations. In addition, we give two examples of sanction rules.

We illustrate the operational semantics and the normative organisation using an insurance claim process. We define five commitments, from insurer to repairer, from insurer to assessor, from assessor to repairer, from repairer to customer, and from insurer to customer. An execution trace of seventeen actions and one deadline expiration illustrates how these commitments change over time.

We discuss *temporal properties*, *interference*, and *redundancy*. First, the temporal properties are the usual specification and verification properties defined for operational semantics. Second, the interference properties refer to the possible interaction between offers. For example, can a single communication act detach multiple commitments? Third, the redundancy properties, which are more involved than temporal and interference properties, give a form of implicit logical relations between the communication actions. We consider a sanction system where agents are blacklisted regardless of the commitment they violated, and we show, for example, that weaker commitments and some conjunctions of commitments are blacklist redundant.

Our approach presents a generic methodology that can be applied to different sets of counts-as rules: for instance, variants of Fig. 2 that reflect variants of the commitment lifecycle, e.g., if conditional commitments are created only after explicit acceptance by the creditor. It further can permit constitutive norms that change dynamically.

Acknowledgements. We thank the referees for their constructive comments, and the participants of the COIN'12 workshop at AAMAS'12 for the discussions.

References

1. Baldoni, M., Baroglio, C., Marengo, E.: Behavior-oriented commitment-based protocols. In: Proc. ECAI, pp. 137–142 (2010)
2. Billhardt, H., Centeno, R., Cuesta, C.E., Fernández, A., Hermoso, R., Ortiz, R., Ossowski, S., Pérez-Sotelo, J.S., Vasirani, M.: Organisational structures in next-generation distributed systems: Towards a technology of agreement. Multiagent and Grid Systems 7(2-3), 109–125 (2011)
3. Boella, G., Broersen, J., van der Torre, L.: Reasoning about Constitutive Norms, Counts-As Conditionals, Institutions, Deadlines and Violations. In: Bui, T.D., Ho, T.V., Ha, Q.T. (eds.) PRIMA 2008. LNCS (LNAI), vol. 5357, pp. 86–97. Springer, Heidelberg (2008)
4. Boella, G., Damiano, R., Hulstijn, J., van der Torre, L.: A common ontology of agent communication languages: Modelling mental attitudes and social commitments using roles. Applied Ontology 2(3-4), 217–265 (2007)

5. Boella, G., Pigozzi, G., van der Torre, L.: Normative framework for normative system change. In: Proc. AAMAS, pp. 169–176 (2009)
6. Boella, G., van der Torre, L.: Regulative and constitutive norms in normative multiagent systems. In: Proc. KR, pp. 255–266 (2004)
7. Bulling, N., Dastani, M.: Verification and implementation of normative behaviours in multi-agent systems. In: Proc. IJCAI 2011, pp. 103–108 (2011)
8. Carabelea, C., Boissier, O.: Coordinating agents in organizations using social commitments. Electronic Notes in Theor. Comp. Sci. 150(3), 73–91 (2006)
9. Chopra, A.K., Dalpiaz, F., Giorgini, P., Mylopoulos, J.: Modeling and Reasoning about Service-Oriented Applications via Goals and Commitments. In: Pernici, B. (ed.) CAiSE 2010. LNCS, vol. 6051, pp. 113–128. Springer, Heidelberg (2010)
10. Chopra, A.K., Singh, M.P.: Specifying and applying commitment-based business patterns. In: Proc. AAMAS, pp. 475–482 (2011)
11. Colombetti, M.: A commitment-based approach to agent speech acts and conversations. In: Proc. Workshop on Agent Languages and Communication Policies, pp. 21–29 (2000)
12. Dastani, M., Grossi, D., Meyer, J.-J.C., Tinnemeier, N.: Normative Multi-agent Programs and Their Logics. In: Meyer, J.-J.C., Broersen, J. (eds.) KRAMAS 2008. LNCS, vol. 5605, pp. 16–31. Springer, Heidelberg (2009)
13. Felipe, L.O., Álvarez-Napagao, S., Vázquez-Salceda, J.: Towards a framework for the analysis of regulative norm performance in complex networks. In: Proc. of 1st Intl. Conf. on Agreement Technogolies (AT 2012), pp. 103–104 (October 2012)
14. Fornara, N., Colombetti, M.: A commitment-based approach to agent communication. Applied Artificial Intelligence 18(9-10), 853–866 (2004)
15. Fornara, N., Viganò, F., Colombetti, M.: Agent communication and artificial institutions. Autonomous Agents and Multi-Agent Systems 14, 121–142 (2007)
16. Kibble, R.: Speech acts, commitment and multi-agent communication. Computational and Mathematical Organization Theory 12, 127–145 (2006)
17. Mallya, A.U., Yolum, P., Singh, M.P.: Resolving Commitments among Autonomous Agents. In: Dignum, F.P.M. (ed.) ACL 2003. LNCS (LNAI), vol. 2922, pp. 166–182. Springer, Heidelberg (2004)
18. Marengo, E., Baldoni, M., Baroglio, C., Chopra, A.K., Patti, V., Singh, M.P.: Commitments with regulations: Reasoning about safety and control in REGULA. In: Proc. AAMAS, pp. 467–474 (2011)
19. Pasquier, P., Chaib-draa, B.: Integrating Social Commitment-Based Communication in Cognitive Agent Modeling. In: Dignum, F.P.M., van Eijk, R.M., Flores, R. (eds.) AC 2005. LNCS (LNAI), vol. 3859, pp. 76–92. Springer, Heidelberg (2006)
20. Singh, M.P.: A Social Semantics for Agent Communication Languages. In: Dignum, F., Greaves, M. (eds.) Issues in Agent Communication. LNCS, vol. 1916, pp. 31–45. Springer, Heidelberg (2000)
21. Telang, P.R., Singh, M.P.: Specifying and verifying cross-organizational business models. IEEE Trans. Services Computing 4 (2011)
22. Venkatraman, M., Singh, M.P.: Verifying compliance with commitment protocols. Autonomous Agents and Multi-Agent Systems 2(3), 217–236 (1999)