

Commitments with Regulations: Reasoning about Safety and Control in REGULA

Elisa Marengo
Università degli Studi di Torino
emarengo@di.unito.it

Amit K. Chopra
Università degli Studi di Trento
chopra@disi.unitn.it

Matteo Baldoni
Università degli Studi di Torino
baldoni@di.unito.it

Viviana Patti
Università degli Studi di Torino
patti@di.unito.it

Cristina Baroglio
Università degli Studi di Torino
baroglio@di.unito.it

Munindar P. Singh
North Carolina State Univ.
singh@ncsu.edu

ABSTRACT

Commitments provide a flexible means for specifying the business relationships among autonomous and heterogeneous agents, and lead to a natural way of enacting such relationships. However, current formalizations of commitments incorporate conditions expressed as propositions, but disregard (1) temporal regulations and (2) an agent's control over such regulations. Thus, they cannot handle realistic application scenarios where time and control are often central because of domain conventions or other requirements.

We propose a new formalization of commitments that builds on an existing representation of events in which we can naturally express temporal regulations as well as what an agent can control, including indirectly as based on the commitments and capabilities of other agents. Our formalization supports a notion of commitment safety. A benefit of our consolidated approach is that by incorporating these considerations into commitments we enable agents to reason about and flexibly enact the regulations.

The main contributions of this paper include (1) a formal semantics of commitments that accommodates temporal regulations; (2) a formal semantics of the notions of innate and social control; and (3) a formalization of when a temporal commitment is safe for its debtor. We evaluate our contributions using an extensive case study.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—Multiagent systems; H.1.0 [Information Systems]: Models and Principles—General

General Terms

Theory

Keywords

Business process modeling, business protocols

1. INTRODUCTION

Previously, commitments have been studied over propositional languages [6, 7, 11, 19]. But in a number of practical settings, com-

Cite as: Commitments with Regulations: Reasoning about Safety and Control in REGULA, E. Marengo, M. Baldoni, C. Baroglio, A. K. Chopra, V. Patti, M. P. Singh, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Tumer, Yolum, Sonenberg and Stone (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. XXX-XXX. Copyright © 2011, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

mitments involve rich temporal structure. Consider the following examples taken from a healthcare setting.

EXAMPLE 1. *An insurance company commits to reimbursing a covered patient for a health procedure provided the patient obtains approval from the company prior to the health procedure. Presumably, the patient would delay going in for the procedure until after having obtained an approval.*

EXAMPLE 2. *An insurance company commits to paying an in-network surgeon for a procedure only after a covered patient has undergone the procedure. Presumably, the surgeon would bill the insurance company after performing the procedure.*

As the following examples illustrate, temporal commitments can also involve more than two parties.

EXAMPLE 3. *A physician commits to a patient that if the patient has any sign of heart trouble after signing up with him, then the patient will be immediately referred to a laboratory for tests, the results of which will be evaluated by a specialist.*

EXAMPLE 4. *A pharmacy commits to provide medicine only if the patient obtains a prescription for that medicine.*

EXAMPLE 5. *For an out-of-network surgeon, an insurance company commits to paying the patient (instead of the surgeon) but only after the surgeon performs the procedure, the patient pays the surgeon, and the patient submits receipts to the insurance company.*

Temporal constraints such as those alluded to in Examples 1–5 are traditionally captured as procedural workflows. Instead, following recent approaches [8, 16], we think of such constraints more broadly as regulations and express them more flexibly in a logical notation. The commitments among autonomous parties capture their business relationships naturally. In contrast with existing approaches [3, 8, 10], we incorporate regulations as contents of commitments. By thus reifying regulations into business relationships, we bring normative force to the specification, thereby providing a clear basis for the participants to guide their actions locally and to judge the compliance of their counterparties. For example, if a regulation says that a physician's referral should precede a surgeon's procedure, it is not clear whether the physician is responsible for moving first or the surgeon is responsible for moving second. By placing the regulations in commitments, we make it explicit that it is the debtor of the commitment who needs to ensure its satisfaction. Further, in doing so, we can capture the business relationships (and concomitant regulations) in a flexible manner that avoids unnecessarily coupling or constraining the participants.

The expression $C(\text{debtor}, \text{creditor}, \text{antecedent}, \text{consequent})$ means that the debtor commits to the creditor that if the antecedent holds, the consequent will hold. The antecedent and the consequent, i.e., the contents of a commitment, are typically logical expressions over states of the world [19], although some have added an explicit temporal component for expressing deadlines [5, 13].

In the present development, the content is specified over events. We use ‘.’ (center dot) as *before*, our main temporal operator on events, where $a \cdot b$ means that event a occurs before event b (though both occur eventually). Then we may express the commitments in Example 1 and 2 as $C(\text{ins}, \text{pat}, \text{approve} \cdot \text{perform}, \text{reimburse})$ and $C(\text{ins}, \text{sur}, \text{perform} \cdot \text{bill}, \text{pay})$, respectively. The commitment in Example 3 would be $C(\text{phy}, \text{pat}, \text{signup} \cdot \text{heartTrouble}, \text{test} \cdot \text{evaluate})$.

1.1 Challenges: Progression, Control, Safety

Placing temporal regulations within commitments enables us to identify precisely the responsibilities of the agents individually, and offer flexibility in terms of how the agents enact their commitments. By contrast, a purely temporal approach, such as Singh’s [16], curtails the autonomy of the participants once the desired computations are specified. However, placing regulations inside commitments, as we did in the above examples, leads to new challenges. One, we must formalize the progression that is, the life cycle, of commitments, bearing in mind the events that have occurred. For example, we say that an active commitment $C(x, y, r, u)$ progresses to *discharged* when u occurs. Analogously, we would like to say that $C(x, y, r, e \cdot f)$ progresses to $C(x, y, r, f)$ when e occurs. The challenge is to formalize general progression rules for an expressive event language. Two, a regulation expresses a constraint over the occurrence of events in a distributed system. The capability for bringing about the events, that is, the *control* of the events, would also generally be distributed among the agents. In Example 3, the physician commits to the patient for both testing and evaluation, but has control of neither—he must rely on a laboratory and a specialist for these tasks. Further, the physician commits to the coordination constraint that the testing will occur before the evaluation. Clearly, the physician is committing to activities over whose performance he apparently has no control. What would make the physician’s commitment reasonable?

In general, an agent would want to commit only to temporal conditions over which it exercises adequate control. We distinguish between two kinds of control: *innate* and *social*. In the above example, the laboratory and the specialist have innate control over testing and evaluation, respectively. Social control ties in with commitments: for regulations specified over events that the debtor does not have control over, the debtor would need the appropriate commitments from those who have control. The physician would have control over testing and evaluation if he could get the appropriate commitments from the laboratory and the specialist. For example, $C(\text{lab}, \text{phy}, \top, \text{test})$ and $C(\text{specialist}, \text{phy}, \top, \text{evaluate})$ would give the physician social control over the two events; however, the physician really needs $C(\text{specialist}, \text{phy}, \top, \text{test} \cdot \text{evaluate})$ from the specialist in order to ensure the appropriate event order.

Control in turn motivates the notion of the *safety* of a commitment. A commitment is safe if its debtor has established sufficient control to guarantee being able to discharge it (assuming others discharge commitments of which they are the debtors). For example, without the above commitments from the laboratory and the specialist, the physician’s commitment to the patient would be unsafe. As our examples illustrate, the notions of control and safety are especially relevant for understanding engagements among more than two parties. How can we determine whether the debtor of a com-

mitment is able to apply the requisite control so as to ensure that its commitments have the support of the other agents so that together they satisfy a given regulation, thereby accomplishing the coordination envisaged by the regulation?

1.2 Contributions

Our contributions may be summarized as follows. First, we formalize commitments with regulations in a simple but expressive event-based model. We formalize rules that capture the *progression* of a commitment over runs (sequences of events), using a previous sound and complete residuation reasoner. Second, we formalize *control* and *safety* in the same event-based model. In particular, we formalize a notion of social control via commitments, which naturally matches multiagent settings. Third, we *declaratively formalize* the *life cycle* of commitments, captured by Theorem 1. Moreover, we connect the notion of commitment progression with the notions of control and of safety (Theorems 2 and 3). Specifically, as long as the agents cause events that are expected by the application of the definition of safety itself, safety is preserved. We evaluate the proposed notions by formalizing Robert’s Rules of Order [14] (RONR), one of the best known set of laws for managing the proceedings of democratic parliamentary assemblies.

Organization

The paper is organized as follows: Section 2 presents the theoretical background necessary for our formalization; Section 3 contains the main theoretical results, concerning the notions of progression, control and safety; Section 4 reports our case study; Section 5 concludes with a discussion and a review of the relevant literature.

2. TECHNICAL FRAMEWORK

Previous works on events and on commitments are relevant here. However, the approaches of interest are not mutually compatible at a technical level. On the one hand, to reason incrementally about control and progression, we need a powerful notion of events and residuation whose semantics is given with respect to an event run [16]. On the other hand, to represent conditional, active commitments, we need an approach based on a state-based semantics given with respect to a state and an index on it [17]. Thus one of the challenges our framework addresses is reconciling the above. As a result, although our approach borrows ideas from Singh [17], our formal model and its details are novel to this paper.

2.1 Precedence logic

Precedence logic is an event-based logic [16]. It has three primary operators for specifying requirements about the occurrence of events: ‘ \vee ’ (choice), ‘ \wedge ’ (concurrency), and ‘ \cdot ’ (before). The *before* operator enables one to express specifications such as *approve*·*perform*: both *approve* and *perform* must occur and in the specified order. The specifications are interpreted over runs. Each run is a sequence of events. Figure 1 shows a schematic of our model. The transitions correspond to event occurrences (the \bullet symbols merely identify place holders between consecutive events: on each run, each \bullet corresponds to an index). The model shows several runs, of which it identifies τ_0 , τ_1 , and τ_2 , which all begin with ab . Additional runs include all the suffixes of τ_0 , τ_1 , and τ_2 —for example, *bef* and *bcd_x* (an event subscript indicates which agent has the capability to perform the event; thus, d_x means that x has the capability to perform d). The same point may be identified with different indices on different runs. For example, the point after b has index 2 on $\tau_0 = abcd_x \dots$ and index 0 on $gh \dots$ (the top branch).

Let e be an event. Then \bar{e} , the complement of e , is also an event. Initially, neither e nor \bar{e} hold. On any run, either e or \bar{e} may oc-




Figure 1: A schematic of runs with common prefixes.

cur, not both. We assume that events are nonrepeating. In practice, transaction IDs or timestamps differentiate multiple instances of the same event. This yields the following advantages. First, since we want to talk about precedence, it helps avoid the confusion where a preceding b would be consistent with b preceding a . Second, it supports negative events as occurring, and distinct from (and stronger than) a positive event not having occurred yet. Third, it facilitates a simpler language and logic that is nevertheless adequate for capturing several regulations of practical interest.

2.2 Language

We distinguish between physical and social (or institutional) events. Our specifications are limited to the physical events (those that are publicly performed by an agent). For example, the events in Figure 1 are physical events: think of a as a waiter pushing a menu card to a customer over a counter in a diner and b as the customer pushing \$5 back to the waiter. However, we supplement physical events with a set of means axioms, which capture the notion of *counts as* [2, 15]. We take the social events corresponding to a physical event to occur concurrently with the physical event. In this manner, we respect Goldman’s notion of (conventional) *generation* [12]. For example, a in Figure 1 may mean the creation of an offer to sell a coffee for \$5, which thus happens simultaneously with a . Likewise, b may mean accepting the offer created by a .

We propose a language, REGULA, in which the antecedents and consequents of commitments are event expressions. In intuitive terms, a commitment itself remains a state expression and we do not express it directly in our language of events. Instead, we think of the operations of commitments as first-class entities, i.e., as events, and let the resulting commitments stay in the background. In other words, we can think of an operation such as $\text{Create}(x, y, r, u)$ as a social event that brings about the corresponding commitment.

Below, \mathcal{X} yields agent names, \mathbf{E} yields event types, and param yields domain values using which an event instance is specified from an event type. In conceptual terms, an event type may be either (1) of sort *physical*, in which case we optionally specify the agent who has the capability to perform it (an unspecified agent indicates we do not care about the agent) or (2) of sort *social*, in which case it is an operation on commitments. For brevity, “event” means “event instance” throughout. The syntax of REGULA is:

```

REGULA  $\rightarrow$  axiom { , axiom }
axiom  $\rightarrow$   $\langle\langle$ physical means social $\rangle\rangle$ 
physical  $\rightarrow$   $\mathbf{E}(\{\mathcal{X}, \text{param}^*\}$ 
social  $\rightarrow$  op( $\mathcal{X}, \mathcal{X}, \text{regulations}, \text{regulations}$ )
op  $\rightarrow$  Create | Cancel | Release | Assign | Delegate
regulations  $\rightarrow$  regulation {  $\wedge$  regulation }
regulation  $\rightarrow$  sequence {  $\vee$  sequence }
sequence  $\rightarrow$  0 |  $\top$  | physical | physical  $\cdot$  physical

```

That is, a REGULA specification is a set of axioms asserting which physical events count as which social events. In well-formed axioms, we require that the performing agent of the physical and corresponding social event be the same.

We use the following conventions: x , etc. are agents, e , etc. are physical events, r , s , u , etc. are regulations, τ , etc. are runs, and i , etc. are indices into runs. We drop agent names when they are understood. The above grammar limits sequences to two events each to simplify our formalization. However, in practice we write longer sequences because $e_1 \cdots e_n \equiv (e_1 \cdot e_2) \wedge \dots \wedge (e_{n-1} \cdot e_n)$.

2.3 Model and Semantics

We now describe the semantics of REGULA in terms of a model, $M = \langle \mathbb{E}, \mathbb{T}, \mathbb{C}, \mathbb{D}, \mathbb{X}, \mathbb{V} \rangle$. Here \mathbb{E} and \mathbb{T} describe the physical layer, \mathbb{C} describes the social (commitment) layer.

- \mathbb{E} is a denumerable set of possible event instances closed under complementation. That is, $e \in \mathbb{E}$ if and only if $\bar{e} \in \mathbb{E}$. For simplicity in our notation, we identify e and \bar{e} ; thus wherever we write e in the semantics, it applies both to e and \bar{e} . The set \mathbb{E} can itself be generated from event types and their parameters, as described above. Further, we introduce a special symbol ϵ for a null event.
- $\mathbb{T} = \{\tau \mid \tau : \mathbb{N} \mapsto \mathbb{E} \cup \{\epsilon\}, \tau \text{ is (1) injective, and (2) } (\forall i, j : \tau_i \neq \tau_j) \}$ is the set of possible event runs (we write the i^{th} event in τ as τ_i). \mathbb{N} is the set of natural numbers. Thus each member of \mathbb{T} is a sequence of events. The above constraints restrict \mathbb{T} to *legal* runs [16] wherein (1) no event repeats and (2) no event and its complement both occur.

We use the null event ϵ to indicate the termination of a run. If ϵ ever occurs on a run, all subsequent events are ϵ . That is, $(\forall i, j : j \geq i \text{ and } \tau_i = \epsilon \Rightarrow \tau_j = \epsilon)$. Below $|\tau|$ is the length of τ , and equals the smallest index i for which $\tau_i = \epsilon$ if i exists and is ω otherwise (indicating an infinite run).

Notice that \mathbb{T} is suffix-closed, meaning that if a run belongs to \mathbb{T} , then so does each suffix of it. Formally, using s as the successor function for \mathbb{N} and \odot as functional composition, we have $\{\tau \odot s \mid \tau \in \mathbb{T}\} \subseteq \mathbb{T}$. Below, $[i, j]$ refers to the subrun between the i^{th} and the j^{th} events, both inclusive. Likewise, \mathbb{T} is prefix-closed. That is, (using the fact that all natural numbers are less than ω), we have $\{\tau_{[0, j]} \mid \tau \in \mathbb{T} \text{ and } j \leq |\tau|\} \subseteq \mathbb{T}$.

- $\mathbb{C} : \mathbb{T} \times \mathbb{N} \times \mathcal{X} \times \mathcal{X} \times \wp(\mathbb{T}) \mapsto \wp(\wp(\mathbb{T}))$ is the standard for (active) commitments. That is, at each index on each run, for each debtor-creditor (ordered) pair of agents, \mathbb{C} assigns to a set of runs a set of set of runs. The intuition is that \mathbb{C} determines which conditional commitment is active from a debtor to a creditor at an index in a run. Given a potential antecedent, each of the consequents is placed in the set that is the output of this function. If the output set is empty at a run and an index that means no commitments are active there. We lack the space to include additional semantic (closure) constraints on \mathbb{C} along the lines of Singh [17].
- (For readability, we place the agents as subscripts.) If two runs are equal until i then \mathbb{C} yields the same result for each of them at index i . Formally, $(\forall \tau, \tau', i, R \subseteq \mathbb{T} : \tau_{[0, i]} = \tau'_{[0, i]} \Rightarrow \mathbb{C}_{x, y}(\tau, i, R) = \mathbb{C}_{x, y}(\tau', i, R))$.
- $\mathbb{D}, \mathbb{X}, \mathbb{V}$ with the same signature as \mathbb{C} are respectively the standards for discharged, expired, and violated commitments.

The following constraints on our model capture some of the essential intuitions about commitments. Let $\text{cone}(\tau, i) = \{\tau' \mid \tau_{[0, i]} = \tau'_{[0, i]}\}$. In intuitive terms, the cone of a run at an index includes all possible future branches given the history (the part of the run up to the index). Because \mathbb{T} contains all possible legal runs, the intuition

is that when a regulation is true (respectively, false) on all runs on a cone, then it is definitely true (respectively, definitely false). For example, e is false at index i of a run τ if it has not occurred yet; it is definitely false if e would occur on no runs in τ 's cone at i , meaning that \bar{e} must already have occurred.

- $U \in \mathbb{D}_{x,y}(\tau, i, R)$ only if $\tau_{[0,i]} \in U$
The consequent of a discharged commitment must be true.
- $U \in \mathbb{X}_{x,y}(\tau, i, R)$ only if $\mathbf{cone}(\tau, i) \subseteq (\mathbb{T} \setminus R)$
An expired commitment must have its antecedent definitely false. In other words, we do not just want that the antecedent is not yet true, we want it to never become true given the run so far.
- $U \in \mathbb{V}_{x,y}(\tau, i, R)$ only if $\mathbf{cone}(\tau, i) \subseteq (\mathbb{T} \setminus U)$ and $\tau_{[0,i]} \in R$
A commitment is violated if its antecedent holds but its consequent is false.
- $U \in \mathbb{C}_{x,y}(\tau, i, R)$ only if $U \notin (\mathbb{D}_{x,y}(\tau, i, R) \cup \mathbb{X}_{x,y}(\tau, i, R) \cup \mathbb{V}_{x,y}(\tau, i, R))$
An active commitment is one that is not discharged, expired, or violated.

Semantic postulates M₁–M₅, loosely based on Singh [16], address the temporal aspects of our language.

- M₁. $\tau \models_i \top$
- M₂. $\tau \models_i e$ iff $(\exists j \geq i : \tau_j = e)$, where e is a physical event
- M₃. $\tau \models_i r \vee u$ iff $\tau \models_i r$ or $\tau \models_i u$
- M₄. $\tau \models_i r \wedge u$ iff $\tau \models_i r$ and $\tau \models_i u$
- M₅. $\tau \models_i r \cdot u$ iff $(\exists j \geq i : \tau \models_{[i,j]} r \text{ and } \tau \models_{[j+1,|\tau|]} u)$

It is helpful to expand the notion of complementation to apply to regulations, not just to physical events. To this end, we define complementation via a set of inference rules as follows: (1) $\bar{r} \wedge \bar{u} = \bar{r} \vee \bar{u}$; (2) $\bar{r} \vee \bar{u} = \bar{r} \wedge \bar{u}$; (3) $\bar{r} \cdot \bar{u} = \bar{r} \vee \bar{u} \vee u \cdot r$; and (4) $\bar{\bar{e}} = e$. The interesting rule is the one for $\bar{r} \cdot \bar{u}$, which captures that $r \cdot u$ may fail to occur exactly when one of its components does not occur or they both occur but in the reverse order.

2.4 Residuation

In simple terms, residuation is a way for us to track progress in the real world. The residual of a regulation with respect to an event is the “remainder” regulation that would be left over from the original after the event, and whose satisfaction would guarantee the satisfaction of the original regulation.

For example, let $r = \bar{a} \vee b \cdot a$ be a regulation under consideration; r means that either a cannot occur, because \bar{a} occurred, or b and a both occur with b preceding a . If we residuate r by an event g , which does not occur in r , the result is the same as r , indicating that the desired regulation is unaffected by irrelevant events. Residuating r by b yields a , meaning that going forward a remains the only possibility. A subsequent occurrence of a would residuate this to \top , meaning that the regulation is satisfied on this execution. Residuating r directly by a yields 0, meaning that the occurrence of a has caused a violation of the regulation.

Following Singh [16], we can define the residual of a regulation r with respect to a physical event e as the maximal (most flexible) regulation such that an occurrence of e followed by an occurrence

of the residual guarantees the original regulation. A benefit of using the event-based semantics is that it supports a set of simple equations or rewrite rules through which we can symbolically calculate the residual of a regulation given an event. The following equations are due to Singh [16]. Here, r is a sequence expression, and e is a physical event or \top . Below, Γ_u is simply the set of literals and their complements mentioned in u . Thus $\Gamma_e = \{e, \bar{e}\} = \Gamma_{\bar{e}}$ and $\Gamma_{e \cdot f} = \{e, \bar{e}, f, \bar{f}\}$.

$$\begin{aligned} 0/e &\doteq 0 \\ \top/e &\doteq \top \\ (r \wedge u)/e &\doteq ((r/e) \wedge (u/e)) \\ (r \vee u)/e &\doteq ((r/e) \vee (u/e)) \\ (e \cdot r)/e &\doteq r, \text{ if } e \notin \Gamma_r \\ r/e &\doteq r, \text{ if } e \notin \Gamma_r \\ (e' \cdot r)/e &\doteq 0, \text{ if } e \in \Gamma_r \\ (\bar{e} \cdot r)/e &\doteq 0 \end{aligned}$$

The above equations characterize the progression of regulations under physical events. They have some important properties, including that (1) regulations not mentioning an event are independent of that event; (2) conjoined or disjointed regulations can be treated modularly; and (3) regulations can be incrementally progressed: hence a residuated regulation embodies the relevant history and no additional history need be represented.

We define the *intension* of r as the set of runs where it is true on index 0: $\llbracket r \rrbracket = \{\tau \mid \tau \models_0 r\}$. As an auxiliary definition, for a set of runs R , let $R \downarrow e = \{\nu \in \mathbb{T} \mid (\forall v : v \in \llbracket e \rrbracket \Rightarrow v\nu \in R)\}$. Then, following Singh [16], we can capture residuation semantically as $\llbracket r/e \rrbracket = \llbracket r \rrbracket \downarrow e$.

2.5 Commitments

When the antecedent is \top (*true*), we refer to the commitment as being *unconditional*. An unconditional commitment usually arises because a conditional commitment was detached. For example, $C(\text{merchant}, \text{customer}, \text{paid}, \text{goods})$ gives rise to $C(\text{merchant}, \text{customer}, \top, \text{goods})$ when *paid* holds. We briefly mention some important stages in the life cycle of a commitment, that is, its progression. A commitment holds either because of an explicit Create operation by the debtor or because an existing commitment was detached. It is considered *expired* if the antecedent has expired (it cannot be satisfied anymore), meaning that the creditor did not take up the offer entailed by the commitment. A commitment is considered *violated* if it is unconditional and its consequent has expired, meaning that the debtor did not fulfill the offer. Alternatively, if the consequent holds, it is considered *discharged*.

Although we adopt many of the intuitions of the previous works, our technical development is significantly different in that we model commitments in an event-based framework. Doing so is nontrivial but yields rewards in an improved characterization of the progression of commitments than previously available.

Now we enhance the above development to accommodate commitments. The commitment operator C is not included in our language but provides a useful basis for the social operations.

$$M_6. \tau \models_i C(x, y, r, u) \text{ iff } \llbracket u \rrbracket \in \mathbb{C}_{x,y}(\tau, i, \llbracket r \rrbracket)$$

The meaning of the physical events in terms of social events is defined as follows. This simply states that whenever a physical event occurs the corresponding social event occurs as well. We leave open the possibility that a social event could occur implicitly without the matching physical event.

$$M_7. \tau \models_i e_x \text{ means } Op(x, y, r, u) \text{ iff } (\forall j \geq i : \tau \models_j e_x \Rightarrow \tau \models_j Op(x, y, r, u))$$

Now we can state the meanings of the operations in terms of how they change the social state by manipulating commitments. We describe Create for concreteness and omit the rest for brevity.

$$M_8. \tau \models_i \text{Create}(x, y, r, u) \text{ iff } \tau \not\models_i C(x, y, r, u) \text{ and } \tau \models_{i+1} C(x, y, r, u)$$

We impose a restriction on our model capturing that commitments persist until they are discharged, expired, or violated. Formally, $\forall \tau, \tau', i, R, U \subseteq \mathbb{T}$ if $U \in \mathbb{C}_{x,y}(\tau, i, R)$ and $\tau_i = e$, then:

1. $U \downarrow e \in \mathbb{X}_{x,y}(\tau, i, R \downarrow e)$ if $R \downarrow e = \emptyset$;
2. $U \downarrow e \in \mathbb{V}_{x,y}(\tau, i, R \downarrow e)$ if $R \downarrow e = \mathbb{T}$ and $U \downarrow e = \emptyset$;
3. $U \downarrow e \in \mathbb{D}_{x,y}(\tau, i, R \downarrow e)$ if $U \downarrow e = \mathbb{T}$;
4. $U \downarrow e \in \mathbb{C}_{x,y}(\tau, i, R \downarrow e)$ otherwise.

3. THEORETICAL RESULTS

We now present the main theoretical results on REGULA.

3.1 Residuation for Commitment Progression

Although commitment expressions are not event expressions in REGULA, we can use residuation to compute the progression of a commitment. For example, consider a commitment $c_1 = C(x, y, b, d \cdot c \cdot a)$ and assume that events d, c, a occur in this order. In intuitive terms, after d , the antecedent of c_1 would be unaffected whereas its consequent would progress to $c \cdot a$. If c were to occur then, the antecedent would still be unaffected, but the consequent would reduce to a ; then when a occurs, the consequent would reduce to \top , indicating that c_1 is discharged. Alternatively, assume that events d, a, b, c occur in this order. Now after d and a , the antecedent would still be unaffected, but the consequent would reduce to \emptyset , indicating that commitment c_1 is violated.

In essence, the idea is that a commitment progresses as its antecedent and consequent are residuated by the events as they occur. The foregoing intuition can be thought of as distributing residuation into the antecedent and consequent of a commitment. This leads us to Theorem 1 on commitment progression. This theorem is technically trivial but important because it shows how a commitment progresses. Here, we assume that operators exp , vio , and dis are defined analogously to C though based on \mathbb{X} , \mathbb{V} , \mathbb{D} , respectively.

THEOREM 1. *If $\tau \models_i C(x, y, r, u)$ and $\tau_i = e$, then*

$$\begin{array}{ll} \tau \models_{i+1} \text{exp}(x, y, r/e, u/e) & \text{if } r/e \doteq 0 \\ \text{vio}(x, y, r/e, u/e) & \text{if } r/e \doteq \top, u/e \doteq 0 \\ \text{dis}(x, y, r/e, u/e) & \text{if } u/e \doteq \top \\ C(x, y, r/e, u/e) & \text{otherwise} \end{array}$$

Proof sketch: Trivial by construction of \mathbb{C} , \mathbb{X} , \mathbb{V} , and \mathbb{D} .

3.2 Control

Consider $C(x, y, \top, a_x \cdot b_y)/a_x$, yielding $C(x, y, \top, b_y)$. Now x is committed to b_y , for which it depends on y . The key challenge here is one of *control*, whether an agent can bring about an event or complex action so as to detach or discharge a given commitment. Our intuition is that control is a combination of capability and opportunity. An agent may control an event innately, i.e., based on which events it can perform, or socially, i.e., based on the commitments of others and what they control. We define the intuitive notion of control $\xi(\cdot, \cdot)$ (our primary definition) in two mutually recursive parts. First, we capture the base cases of control through $\zeta(\cdot, \cdot)$, which captures the notion of innate control.

$$M_9. \tau \models_i \zeta(x, \top)$$

$$M_{10}. \tau \models_i \zeta(x, e_x) \text{ iff } (\exists \tau' \in \mathbf{cone}(\tau, i) : \tau' \models_i e_x)$$

Notice we refer to $\mathbf{cone}(\tau, i)$ above because it serves as a surrogate for notion of state, which is otherwise not present in our framework. Thus, merely finding a τ' on which event e_x occurs at index i would not be enough.

$$M_{11}. \tau \models_i \zeta(x, e_y), \text{ where } x \neq y, \text{ iff } (\exists r : \tau \models_i \xi(x, r) \text{ and } \tau \models_i C(y, x, r, e_y) \text{ and } \tau \models_i \xi(y, e_y))$$

Second, we formulate control recursively using the above.

$$M_{12}. \tau \models_i \xi(x, r \vee u) \text{ iff } \tau \models_i \xi(x, r) \text{ or } \tau \models_i \xi(x, u)$$

$$M_{13}. \tau \models_i \xi(x, r \wedge u) \text{ iff } \tau \models_i \xi(x, r) \text{ and } \tau \models_i \xi(x, u)$$

$$M_{14}. \tau \models_i \xi(x, r \cdot u) \text{ iff } \tau \models_i \xi(x, r) \text{ and } (\forall \tau' \in \mathbf{cone}(\tau, i) : \tau' \models_i r \Rightarrow \tau' \models_i (r \cdot \xi(x, u)))$$

$$M_{15}. \tau \models_i \xi(x, r) \text{ iff } \tau \models_i \zeta(x, r), \text{ when } r \text{ is not of the form } r \vee u, r \wedge u, \text{ or } r \cdot u$$

Given that an agent controls a regulation, the question is whether it is possible that such control be propagated along the execution, as the regulation is residuated based on the occurred events.

THEOREM 2. *If $\tau \models_i \xi(x, r)$ then $(\exists \tau', e : \tau' \in \mathbf{cone}(\tau, i) \text{ and } \tau'_i = e, \text{ and } \tau' \models_{i+1} \xi(x, r/e))$.*

Proof sketch: Follows directly from the definition of control.

Notice that the preservation of control requires some cooperation of the involved agents. For instance, $\xi(x, e_y)$ requires that y continues to support x . In other words, either y causes e_y at some point or y does not cancel its commitment to x to execute e_y when some condition becomes true.

Informally, by a *Create*, the debtor provides social control to the creditor; by performing a *Cancel* the debtor takes back social control; by performing a *Release* the creditor relinquishes social control. Likewise, *Assign* and *Delegate* transfer control suitably.

3.3 Safety

Safety is a property of commitments. Since the regulations embedded in commitments involve many actors, it is important for the potential debtor to understand when it is “reasonable” for it to commit. Intuitively, this is reasonable when the agent controls the events that are part of the regulation. In other words, a commitment is safe for its debtor when the coordination necessary to fulfill the regulation is supported by commitments by the other agents involved. We can thus define the *safety of a commitment* $C(x, y, r, u)$ for the debtor agent x as $\sigma(x, C(x, y, r, u))$ as follows:

$$M_{16}. \tau \models_i \sigma(x, C(x, y, r, u)) \text{ iff } (\forall \tau' \in \mathbf{cone}(\tau, i) \text{ and } (\tau' \models_i \xi(x, \bar{r}) \text{ or } (\mu j \geq i : \tau' \models_{[i,j]} r \Rightarrow \tau' \models_j \xi(x, u/\tau'_{[i,j]}))))$$

Residuation by subrun $\tau'_{[i,j]}$ is a shorthand notation standing for the residuations, in sequence, by all the events in the subrun. By μj we refer to a generalized quantifier that selects the least such j index. So, a commitment is safe for its debtor if either the debtor controls the negation of the antecedent or whenever the antecedent holds, the debtor controls the residuation of the consequent. In the former case, the debtor can act so as to avoid letting the commitment become active. In the latter case, instead, when the commitment becomes active, there is a way to satisfy it. Residuation is necessary because at j some event that is in the consequent might have occurred already.

Notice that the definition of safety does not depend directly on the given run but consider all runs of the same history. The intuition here is to capture the fact that safety is essentially a state property, even though we express it in an event-based model. The definition is symmetric between all the runs that have the same history as τ .

Moreover, safety does not mean that no matter how bad a decision an agent takes, success is guaranteed; just that the agent is not subject to the whims of another agent. In other words, the agent can prevent a bad situation, not that a bad situation is impossible.

THEOREM 3. *If $\tau \models_i \sigma(x, C(x, y, r, u))$ then $(\exists \tau' \in \mathbf{cone}(\tau, i), e$ and $\tau' \models_i \bar{r}$ or $(\tau' \models_i r$ and $\tau'_i = e$ and $\tau' \models_{i+1} \sigma(x, C(x, y, r/e, u/e)))$.*

In words, Theorem 3 states that if at some point on a run a commitment is safe for an agent, then there is a possible continuation such that the residuation of the commitment remains safe.

Proof: Follows from Theorem 2 and the definition of safety. Suppose that $\tau \models_i \sigma(x, C(x, y, r, u))$ holds. Therefore, $(\forall \tau' \in \mathbf{cone}(\tau, i)$ and $\tau' \models_i \xi(x, \bar{r})$ or $(\mu j \geq i : \tau' \models_{[i,j]} r \Rightarrow \tau' \models_j \xi(x, u/\tau'_{[i,j]}))$ by M16. We have two cases. *First case.* Let us suppose that $\tau' \models_i \xi(x, \bar{r})$. By Theorem 2, $(\exists \tau'' \in \mathbf{cone}(\tau', i), e$ and $\tau''_i = e$, and $\tau'' \models_{i+1} \xi(x, \bar{r}/e)$) and this proves the case. *Second case.* Let us suppose that $(\mu j \geq i : \tau' \models_{[i,j]} r \Rightarrow \tau' \models_j \xi(x, u/\tau'_{[i,j]}))$. Thus, whenever $\tau' \models_{[i,j]} r$ we have $\tau' \models_j \xi(x, u/\tau'_{[i,j]})$. By Theorem 2, $(\exists \tau'' \in \mathbf{cone}(\tau', j+1), e$ and $\tau''_i = e$, and $\tau'' \models_{j+1} \xi(x, (u/\tau'_{[i,j]})/e)$). Let us consider the previous τ'' and e and assume that $\tau'' \models_{[i,j+1]} r/e$ holds. Then, we have that $\tau'' \models_i \xi(x, (u/\tau'_{[i,j]})/e)$. This proves the theorem.

4. CASE STUDY

We apply our approach on *Robert's New Rules of Order* (RONR) [14], a system of parliamentary laws. RONR posits two roles: *chair* and *participants*. The activity of the assembly consists of discussing a motion at a time, and then voting. The rules are aimed at guaranteeing that the assembly works in a democratic way. Among other rules, in particular, it specifies that voting will not take place until all the participants who raised their hand for expressing their opinion have spoken. Different members are (not allowed to speak at the same time and, in particular, in order to speak one must have the floor. As long as everybody *behaves* according to the rules, the assembly works in a democratic way. In other terms, RONR not only specifies the actions but also governs the behavior of the participants and the chair (specifying the contexts in which the execution of actions makes sense) so as to guarantee the success of the assembly if all the agents behave according to RONR. Each participant autonomously decides whether to conform to the rules, but doing so confers some rights on the participant.

The first column of Table 1 lists physical events that can occur during an enactment of RONR (the subscript indicates which agent directly controls the event: c stands for chair and p_i generically stands for participant). Recall that given two agents p_1 and p_2 , the event instance e_{p_1} is different from e_{p_2} . So, for instance, $askFloor_{p_1}$ is different from $askFloor_{p_2}$. A possible specification of the semantics of events is given in terms of their effects on the social state. The second column of Table 1 reports event effects in terms of commitments that are created by their occurrence. Antecedents and consequents are written in REGULA. The social effects are operations on commitments. Besides Create, in the example, we also use Assign and Delegate: a participant can delegate its vote or assign its time slot for speaking to another participant.

Notice that if the meaning were given using propositional commitments, one could not express temporal regulations. For instance,

Physical event	Means these social events
openAssembly _c	$\forall p_i \in P, \text{Create}(C(c, p_i, \top, \text{exposeMotion}_c(m) \cdot \text{openDebate}_c(m))) \wedge \forall p_i, p_j \neq p_i \in P, \text{Create}(C(c, p_i, \text{discuss}_{p_j} \wedge \text{giveFloor}_c(p_j) \vee \text{discuss}_{p_j} \cdot \text{giveFloor}_c(p_j), \text{punish}_c(p_j)))$
openDebate _c (m)	$\forall p_i \in P, \text{Create}(C(c, p_i, \text{askFloor}_{p_i}, \text{askFloor}_{p_i} \cdot \text{giveFloor}_c(p_i))) \wedge \forall p_i \in P, \text{Create}(C(c, p_i, \text{askFloor}_{p_i} \cdot \text{giveFloor}_c(p_i) \cdot \text{discuss}_{p_i} \vee \text{askFloor}_{p_i}, \text{askFloor}_{p_i} \cdot \text{giveFloor}_c(p_i) \cdot \text{discuss}_{p_i} \cdot \text{cfv}_c \vee \text{askFloor}_{p_i} \cdot \text{cfv}_c))$
cfv _c	none
enterAssembly _p	Create($C(p, c, \text{cfv}_c, \text{cfv}_c \cdot \text{vote}_p)$)
askFloor _p	Create($C(p, c, \top, \text{discuss}_p)$)
exposeMotion _c (m)	none
discuss _p	none
giveFloor _c (p)	none
passFloor _{p_i} (p _j)	Assign($p_j, C(c, p_i, \top, \text{giveFloor}_c(p_i))$)
vote _p	none
delegateVote _{p_i} (p _j)	Delegate($p_j, C(p_i, c, \top, \text{vote}_{p_i})$)
close_cfv _c	none
closeAssembly _c	none
punish _c (p)	none

Table 1: RONR physical events mapped to their social effects. Here p_i are participants, c the chair, and m a motion.

to express that the floor is given after it is asked for, the commitment $C(c, p_i, \text{askFloor}_{p_i}, \text{giveFloor}_c(p_i))$ would be inadequate since it does not ensure that the two events occur in the expected order. Potentially, the chair could give the floor to p_i before p_i asked for it and the commitment would be discharged. Such apparent flexibility may be desirable in some settings but not where it violates a regulation.

The following is an example commitment whose antecedent or consequent use all the allowed operators of REGULA. In this case, c commits to each p_i that c would punish any other p_j if p_j starts speaking when the chair refused to give it the floor or speaks before having the floor.

$$C_1 = \forall p_i, p_j \neq p_i \in P, C(c, p_i, \text{discuss}_{p_j} \wedge \overline{\text{giveFloor}_c(p_j)} \vee \text{discuss}_{p_j} \cdot \text{giveFloor}_c(p_j), \text{punish}_c(p_j))$$

4.1 Simulation of a Possible Enactment

In order to explain the notions of safety and of control, let us suppose that, instead of the commitments in Table 1, the physical event *openDebate* creates the following commitments:

$$\forall p_i \in P, C_2(p_i) = C(c, p_i, \top, \text{askFloor}_{p_i} \cdot \overline{\text{giveFloor}_c(p_i)} \cdot \text{discuss}_{p_i} \cdot \text{cfv}_c \vee \text{askFloor}_{p_i} \cdot \text{cfv}_c)$$

Given that P denotes the set of all participants to the assembly, the formula specifies the set of *unconditional* commitments of c to all the participants to the assembly to call for votes (cfv_c) after each participant has either (1) asked for the floor, obtained it, and discussed or (2) declined the possibility to speak.

Figure 2 shows a tree corresponding to some of the possible runs that can be obtained by RONR. Since the RONR events have no preconditions, all interleavings where each event instance occurs at most once are possible. The chair, by executing *openDebate* commits unconditionally to the regulation $u = \text{askFloor}_{p_i} \cdot \text{giveFloor}_c$




Figure 2: A schematic of some runs with common prefixes for RONR.

$\langle p_i \rangle \cdot discuss_{p_i} \cdot cfv_c \vee askFloor_{p_i} \cdot cfv_c$. Let us consider the bottom run, which involves the chair and two participants (p_1 and p_2): participant p_1 asks for the floor, receives it and speaks; instead, p_2 asks for the floor but starts speaking *before* the chair gives it the floor. This causes a violation of the commitment of the chair. This violation is due to p_2 , who, however, did not have any commitment to wait for the floor before speaking. Therefore, the chair has no right to expect this behavior by p_2 . Indeed, it was not safe for the chair to adopt the above commitment by opening the debate. More formally, we can show that $\tau' \not\models_5 \sigma(c, C_2(p_2))$. Safety holds iff $(\forall \tau'' \in \mathbf{conc}(\tau', 5) : (\exists j \geq 5 : \tau'' \models_{[5,j]} \top \Rightarrow \tau'' \models_j \xi(c, u/\tau''_{[5,j]})))$. With $j = 5$ this simplifies to $(\forall \tau'' \in \mathbf{conc}(\tau', 5) : \tau'' \models_5 \xi(c, u))$, which does not hold. The participant has not adopted any commitment toward the chair to execute any of the actions under its control. In particular, by repeatedly applying the definition of control, it is easy to see that, in order for the commitment to be safe, it is necessary that after a certain step, for all $\tau, \tau \models_i askFloor_{p_2} \cdot giveFloor_c(p_2) \cdot \xi(c, discuss_{p_2} \cdot cfv_c)$. In other words, after $askFloor_{p_2} \cdot giveFloor_c(p_2)$ occurs, c must have control over $discuss_{p_2} \cdot cfv_c$. If there were a commitment of kind $C(p_2, c, \top, askFloor_{p_2} \cdot giveFloor_c(p_2) \cdot discuss_{p_2})$ at Step 5, $C_2(p_2)$ would be safe because after $askFloor_{p_2} \cdot giveFloor_c(p_2)$ it would have residiuated to $C(p_2, c, \top, discuss_{p_2})$. The fact that no similar commitment is adopted by all agents can lead to the violation described above.

Let us suppose that the commitment adopted after *openDebate* were instead: $C(c, p_i, askFloor_{p_i} \cdot giveFloor_c(p_i) \cdot discuss_{p_i}, askFloor_{p_i} \cdot giveFloor_c(p_i) \cdot discuss_{p_i} \cdot cfv_c)$. Can the chair cause the event *openDebate* without worrying about the above commitment? Again, the answer depends on whether the commitment is safe for c . From the definition, it is easy to see that when the antecedent is satisfied and, thus, $askFloor_{p_i} \cdot giveFloor_c(p_i) \cdot discuss_{p_i}$ occurs, safety depends on the fact that $\xi(c, cfv_c)$ which is trivially true. The effect of *openDebate* in Table 1 is safe.

The notion of control enables other kinds of reasoning. Let us consider the antecedent of the second commitment that is created as an effect of the action *openAssembly*: $discuss_{p_i} \wedge giveFloor_c(p_i) \vee discuss_{p_i} \cdot giveFloor_c(p_i)$. The effect of the action will be a punishment applied to p_i . Does p_i have control over the set of runs respecting these dependencies, so as to avoid activating the commitment of the chair to punish it? The application of the M_{12} rule allows for deconstructing the expression into (a) $discuss_{p_i} \wedge giveFloor_c(p_i)$ and (b) $discuss_{p_i} \cdot giveFloor_c(p_i)$. Condition (a) is a conjunction where $discuss_{p_i}$ is a physical action which is controlled by p_i . The conjunction can be made false by avoiding contributing to a discussion when $giveFloor_c(p_i)$. Condition (b) is a sequence that is started by a physical action, which is controlled by p_i : this agent can avoid that the condition becomes true by avoiding to contribute to the discussion until c gives it the floor.

5. DISCUSSION

The RONR case study validates some important claims about REGULA. First, it shows that commitments with regulations better help a group of agents coordinate their interactions than traditional propositional commitments would.

In general, because of the autonomy of the agents, no agent x may legitimately expect that another agent y would satisfy a particular regulation. However, the existence of a commitment whose debtor is y and creditor x , and whose consequent is the given regulation precisely specifies and legitimizes such an expectation. The placement of regulations within the antecedents and consequents of commitments helps make the regulations explicit within the system of interacting agents and thereby facilitates their coordination. In particular, autonomous agents can potentially trade off the regulations that they respectively prefer with such trade offs expressed in the commitments they make to one another. Lastly, using events makes regulations, control, safety computationally precise while preserving flexibility of commitments: no event trace is explicitly dictated.

Safety is a means for deciding whether it is reasonable for an agent to adopt a commitment in a given state. We showed that when safety holds for the current state, then there is a possible evolution to another safe state. Further, it is possible to exploit the notion of control in other kinds of reasoning. For instance, as the above example shows, if an agent wants to prevent another agent's commitment from being activated, it can check whether it controls the antecedent of such a commitment.

5.1 Relevant Literature

We now review some previous efforts at enriching commitments with time. Following Searle [15], we can classify norms as either constitutive or regulative. Clearly *openAssembly*, *openDebate*, and so on are *constitutive*: their meaning is defined in terms of commitments using the *means* construct, which amounts to a *counts-as* relation. However, the commitments also have a naturally regulative flavor, which we enhance thanks to the coordination requirements arising from the temporal nature of their content. Thus, in effect, both kinds of norms are grounded in communication in our framework. By contrast, Boella and van der Torre [4] define both kinds of norms in terms of agents' mental states.

Aldevereld et al. [2] use the counts-as relation to operationalize norms. By contrast, we use the counts-as relation to understand physical events in terms of the normative relations they create. Our framework includes significant operational elements. The notion of commitment progression is an operational one. Further, the notion of control may be also viewed as an operational tool for reasoning about norm compliance.

Alberti et al. [1] use events and expectations to model interaction protocols. Expectations help define a temporal relation between events. For example, one can state that if an event occurs at

a certain point in time, another event must occur afterward. Expectations, however, are not scoped by a debtor or creditor nor they are used inside commitments. By contrast, we include temporal regulations inside commitments. This enables precisely identifying who is responsible for each regulation and potentially liable for a violation. Moreover, we define control and safety properties both over regulations and commitments. Using these, an agent can determine whether it would be able to satisfy a (temporal) engagement and to reason about the opportunity of adopting specific commitments.

Baldoni et al. [3] define commitment-based protocols wherein the constitutive and the regulative specifications are decoupled. In particular, they assert regulations as temporal constraints but place them separate from commitments, not within the antecedents and consequents of commitments, as we have done here. By contrast, by including regulations inside commitment, we identify a debtor and a creditor with duties and rights, and propose a notion of control and of safety.

Commitment life cycles, that is progressions, have been variously formalized, especially by Fornara and Colombetti [10], Mallya et al. [13], and El-Menshawey et al. [9]. However, in general, these works neither provide a symbolic characterization of progression as we did above nor do they consider the interplay between control and commitment progression. And, the previous semantic approach work on commitments [17] considers only whether a commitment is active or not, and does not discuss the full life cycle.

Cranefield and Winikoff [7] formalize expectation progression in a linear temporal logic. However, unlike commitments, the expectation modality is not a relation between agents. Such a modality would not be able to support the notions of control and safety as we have formalized here.

Verification of protocols is an important theme. Giordano and Martelli [11] perform two kinds of verification: one, whether an agent's execution is compliant with the protocol, and two, whether the protocol specification itself satisfies some temporal property. Our notion of safety is a third category in that it helps an agent determine whether it has adequate control in order to be able to fulfill its commitments. Safety suggests that the protocol in question is well-designed and the agent's behavior complies with the protocol. We establish compliance at runtime through the notion of the progression of a commitment (Theorem 1).

van der Hoek and Wooldridge [18] reason about the abilities of a coalition of agents given each agent's control over certain variables. Moreover, control may be transferred via what they term a "delegate" operation. Our work embodies similar intuitions: commitments allow control to be passed among agents. Additionally, through the use of commitments, we can support cancel and release as ways to return control and delegate and assign as ways to propagate control.

5.2 Future Directions

The notions of control and of safety that we proposed concern single agents. Along the lines of van der Hoek and Wooldridge [18], a key future direction is to explore notions of *teamwork* and to extend the definitions of control and safety accordingly. It would be worth investigating a richer formal model and language in which we include both states and events as transitions between states. Another interesting question is: given a *specification* in terms of a set of temporal regulations, and knowledge of what events are performed by what agent, can we determine the safe commitments that the agents should adopt so that the resulting computation satisfies the original specification? Such set of commitments could be used to implement agents, interacting by means of commitment-based protocols [3, 19].

Acknowledgments

We thank the reviewers for their helpful comments. The Torino team was partially funded by Regione Piemonte, ICT4LAW project. Chopra was supported by a Marie Curie Trentino Fellowship.

6. REFERENCES

- [1] M. Alberti, F. Chesani, D. Daolio, M. Gavanelli, E. Lamma, P. Mello, and P. Torroni. Specification and Verification of Agent Interaction Protocols in a Logic-based System. *Scalable Computing: Pract. & Exp.*, 8(1):1–13, 2007.
- [2] H. Aldewereld, S. Álvarez-Napagao, F. Dignum, and J. Vázquez-Salceda. Making norms concrete. *AAMAS*, pp. 807–814, 2010.
- [3] M. Baldoni, C. Baroglio, and E. Marengo. Behavior-oriented commitment-based protocols. In *ECAI*, pp. 137–142, 2010.
- [4] G. Boella and L. W. N. van der Torre. Regulative and constitutive norms in normative multiagent systems. In *KR Conf.*, pp. 255–266, 2004.
- [5] F. Chesani, P. Mello, M. Montali, and P. Torroni. Commitment tracking via the reactive event calculus. In *IJCAI*, pp. 91–96, 2009.
- [6] A. K. Chopra and M. P. Singh. Contextualizing commitment protocol. In *AAMAS*, pp. 1345–1352, 2006.
- [7] S. Cranefield and M. Winikoff. Verifying social expectations by model checking truncated paths. In *COIN, LNCS 5428*, pp. 204–219. Springer, 2009.
- [8] N. Desai and M. P. Singh. On the enactability of business protocols. In *AAAI*, pages 1126–1131, July 2008.
- [9] M. El-Menshawey, J. Bentahar, and R. Dssouli. Verifiable semantic model for agent interactions using social commitments. In *Proc. Intl. WS Languages, Methodologies, and Development Tools for Multi-Agent Sys., LNCS 6039*, pages 128–152. Springer, 2010.
- [10] N. Fornara and M. Colombetti. Operational specification of a commitment-based agent communication language. In *AAMAS*, pp. 535–542, 2002.
- [11] L. Giordano, A. Martelli, and C. Schwind. Specifying and verifying interaction protocols in a temporal action logic. *Journal of Applied Logic*, 5(2):214–234, 2007.
- [12] A. I. Goldman. *A Theory of Human Action*. Prentice-Hall, Englewood Cliffs, NJ, 1970.
- [13] A. U. Mallya, P. Yolum, and M. P. Singh. Resolving commitments among autonomous agents. In *WS on Agent Communication, LNAI 2922*, pp. 166–182. Springer, 2003.
- [14] H. M. I. Robert, W. J. Evans, D. H. Honemann, and T. J. Balch. *Robert's Rules of Order*, 10th Ed. Da Capo Press, 2000.
- [15] J. R. Searle. *The Construction of Social Reality*. Free Press, New York, 1995.
- [16] M. P. Singh. Distributed enactment of multiagent workflows: Temporal logic for service composition. In *AAMAS*, 2003.
- [17] M. P. Singh. Semantical considerations on dialectical and practical commitments. In *AAAI*, pp. 176–181, 2008.
- [18] W. van der Hoek and M. Wooldridge. On the dynamics of delegation, cooperation, and control: a logical account. In *AAMAS*, pp. 701–708, 2005.
- [19] P. Yolum and M. P. Singh. Flexible protocol specification and execution: Applying event calculus planning using commitments. In *AAMAS*, pp. 527–534, 2002.