

Composing and Verifying Commitment-Based Multiagent Protocols

Matteo Baldoni
 Università di Torino
 Torino, Italy
 baldoni@di.unito.it

Cristina Baroglio
 Università di Torino
 Torino, Italy
 baroglio@di.unito.it

Amit K. Chopra
 Lancaster University
 Lancaster, UK
 akchopra.mail@gmail.com

Munindar P. Singh
 NC State University
 Raleigh, NC, USA
 singh@ncsu.edu

Abstract

We consider the design and enactment of multi-agent protocols that describe collaboration using “normative” or “social” abstractions, specifically, commitments. A (multiagent) protocol defines the relevant social states and how they progress; each participant maintains a local projection of these states and acts accordingly. Protocols expose two important challenges: (1) how to compose them in a way that respects commitments and (2) how to verify the compliance of the parties with the social states. Individually, these challenges are inadequately studied and together not at all.

We motivate the notion of a social context to capture how a protocol may be enacted. A protocol can be verifiably enacted when its participants can determine each other’s compliance. We first show the negative result that even when protocols can be verifiably enacted in respective social contexts, their composition cannot be verifiably enacted in the composition of those social contexts. We next show how to expand such a protocol so that it can be verifiably enacted. Our approach involves design rules to specify composite protocols so they would be verifiably enactable. Our approach demonstrates a use of dialectical commitments, which have previously been overlooked in the protocols literature.

1 Introduction

How can we address the problem of achieving secure collaboration? Secure collaboration relies upon an effective means for capturing the standards of correct interaction between autonomous parties and for providing an effective basis for deciding if those standards are met [Singh, 2013]. Accordingly, we take as our point of departure extensive previous work on multiagent protocols specifying interactions between autonomous *principals* (humans or organizations) facilitated by their computational *agents*. Singh [2013] proposes that the standards alluded to above can be effectively represented via directed normative relationships capturing who is accountable to whom for what and when.

Of the possible such normative relationships, the concept of (social) commitments has received the most attention in

connection with protocols. Specifically, we consider commitment protocols [Yolum and Singh, 2002] where the *social state* of an interaction is expressed using commitments and the actions of a protocol in terms of how they affect commitments. In particular, $C^p(x, y, r, u)$ means a commitment from a debtor x to a creditor y that if the antecedent r holds, then the consequent u will hold [Singh, 2008]. The social state is a notional state of an ongoing interaction. Each agent acts based on its local *projection* of that state using which it takes its decisions on how to proceed.

We focus on one important correctness criterion for interactions, namely, compliance with respect to the above-mentioned standards of interaction. Achieving and judging compliance is essential to realizing secure collaboration. An agent is *compliant* if and only if it discharges all commitments of which it is the debtor [Venkatraman and Singh, 1999]. We are not interested directly in whether an agent participating in a protocol is compliant, because that is dependent upon the agent’s internal decision making. Instead, we consider whether a protocol itself is such that, when enacted, it supports each participant *verifying* the compliance of the other participants.

For verification, a creditor of a commitment must be able to observe relevant events so it can determine the outcome of the commitment. The relevant events include operations on commitments (such as create or cancel) and events whose occurrence affects the truth or falsity of the commitment’s antecedent or consequent. When a commitment is minimally expressed, those events are precisely the atoms in the antecedent or consequent. Compliant behavior by a debtor presumes the debtor can observe the same events, otherwise they may be unaware of their engagements. For this reason, verifiability of compliance ought to incorporate considerations of alignment [Chopra and Singh, 2015b].

Loosely inspired by human interactions, we introduce the idea of a *social context*. A (social) context corresponds to a possible transaction whose constituent events the participants in the context can *observe*. We formally define the verifiability of a protocol with respect to a context. Roughly, an *enactment* of a multiagent system is any sequence of events in which the members of the system participate. An agent can *determine* whether a commitment enters a particular state if it can compute this state exclusively from the events it can observe. Then, a protocol is *verifiable* in a context if and only if

for any enactment that takes place in the context, if an agent determines that a commitment mentioned in the protocol is created of which it is the creditor, then that agent can determine if the commitment was discharged, expired, or violated.

Original Contributions We make the following contributions. (1) A characterization of protocol enactment in social contexts. (2) An approach to determine if a protocol is verifiable with respect to a social context. We show that, importantly, a composite protocol need not be verifiable in the union of contexts in which its constituent protocols are verifiable. (3) A way to bridge social contexts via dialectical commitments. (4) An extensible set of pragmatic rules with which to produce verifiable composite protocols.

Organization Section 2 discusses observability of events related to compliance and verifiability. Section 3 sets up our approach, introducing protocol enactments and explaining how to bridge enactments by way of dialectical commitments and social recognition. Section 4 concludes this paper with a discussion of the literature and some future directions.

2 Compliance and Verifiability

We explain the impact of event observability on commitments with the help of two simple scenarios.

Example 1 (Purchase at shop) *Consider a commitment made by a shopkeeper toward a client to give the client the goods she paid for. Payment at the cash register creates in the client the expectation to receive the goods. Moreover, cash and goods passing at the counter are events that can be observed by all of the parties involved. Consequently, the interaction between the client and the shopkeeper makes the commitment between the two progress for both of them.*

Since all relevant events can be observed by all principals, each principal knows which commitments are active (or expired, detached, or violated). Each principal knows about its own and the other principal’s *compliance*. We say that the two principals enact an interaction that is *verifiable*. However, it is not always the case that all concerned parties observe all of the events relevant to their commitments.

Example 2 (Delivery at home) *Suppose delivery is made by a courier. Thus, it is not directly observable by the shopkeeper. The shopkeeper’s commitment for delivery of goods is satisfied as soon as the courier completes the delivery. The client will see such a commitment as satisfied and will know the shopkeeper to be compliant. The shopkeeper cannot observe this event and will not be able to do the same.*

This case exhibits *two social contexts*: one that involves only the courier and the client, and one that involves the shopkeeper, client, and courier. Delivery belongs to the first context, not to the second, where its occurrence is therefore *not verifiable*. The problem arises because we would like an event from the former context to have an impact on the other. The principals’ views of the state of the interaction cannot therefore be “aligned” [Chopra and Singh, 2015b], hence yielding nonverifiability.

Situations such as the above occur normally in the real world. Therefore, a key challenge is how to make information about events from one context available to another where that information is socially relevant. In the real world, similar situations are often tackled by introducing actions (like delivery tracking), whose meaning is a claim about a state of affairs (a position). Such claims imply the taking of *responsibility* of the truth of what is declared; for instance, the courier declares the state of delivery through the tracking service. They concern events that are relevant to social aims, and that are not directly observable by all of the involved parties—as in the case of the shopkeeper. In other words, these claims act as “bridges” between different contexts. The courier’s statement about the delivery makes information about an event that belongs to the context involving the courier and client available to the context involving the shopkeeper.

The commitments correspond to expectations by some principals of others, and each principal is accountable for the expectations it creates in others. For this reason, we represent the above-mentioned claims as *dialectical commitments*: A *dialectical* commitment $C^d(x, y, r, p)$ means that x (creditor) commits to y (debtor) to the truth of consequent p , provided the antecedent r holds. Practical and dialectical commitments have different normative natures. Practical commitments are most useful in settings where we expect the debtor to possess the requisite capabilities and powers (including of persuading others to exercise the requisite capabilities). Dialectical commitments are most useful where we expect the debtor to possess the relevant knowledge or be able to obtain it. By realizing claims as commitments, debtors become accountable for their declarations, and are held liable when the commitment is violated. We use $C(\dots)$ when it is not necessary to distinguish between the two kinds of commitment.

In our proposal, the shopkeeper and courier rely on a *pragmatic rule* specific to their micro-society because of which they can use dialectical commitments as equal to the conditions they concern. This is just one example of a pragmatic rule. Pragmatic rules are characteristic of the specific social contexts. Thus, the key challenge to interoperation that we tackle is how to determine how the practical commitments progress when some relevant events are not observable in the context. To this end:

1. we use *dialectical commitments* as *claims* concerning the occurrence of events in other contexts;
2. we define *pragmatic rules* that the agents of a context may share: such rules exploit the normative value of dialectical commitments and enable commitments to progress even based on events that are not directly observable.

Our approach addresses the composition of *verifiable interaction protocols*: whether the participants in an interaction can verify their own compliance and that of their interlocutors with respect to a social context in which the protocol is executed. Our problem differs from and complements existing works on formal approaches for protocols. Existing approaches do not consider the context of enactment. Research on monitoring ongoing interactions, e.g., [Chesani et al., 2013], concerns only a given execution trace and often

requires a centralized approach or a shared monitor. Others, e.g., [Dastani *et al.*, 2004], consider the agents that enact protocol roles, whereas we consider the enactment of a protocol (as a whole) inside a context. Chopra and Singh’s [2015b] approach specifies the messages the debtor and creditor must send to each other to ensure their alignment, but do not analyze protocols as such. Others, e.g., [El-Menshaway *et al.*, 2010; Astefanoaei, 2011; Lomuscio *et al.*, 2012; El-Menshaway *et al.*, 2013; Gerard and Singh, 2013] specify formal models of protocols (or protocol or service compositions) and verify using model checkers such as MCMAS. They consider protocol properties such as conformance, termination, and refinement but not verifiability, which is what we study here.

3 Verifiability of Protocol Enactments

A commitment is directed from a debtor to a creditor and involves an antecedent and a consequent [Singh, 2008]. The directionality of commitments is crucial for establishing accountability. Additionally, this directionality provides a basis for tackling the challenges of compliance and verifiability. We enhance Marengo *et al.*’s [2011] formalization to incorporate dialectical commitments formed over temporal expressions in *precedence logic* [Singh, 2003]. This logic has three primary operators: ‘ \vee ’ (choice), ‘ \wedge ’ (concurrency), and ‘ \cdot ’ (before). The *before* operator enables one to express conditions such as *pay · deliver*: both *pay* and *deliver* must occur and in the specified order. The specifications are interpreted over runs. Each run τ is a sequence of events. For simplicity, the events can be thought of as propositional but can be generated schematically as in *pay\$1, pay\$2, ...*

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 occur, not both. Moreover, we extend complementation to expressions in general. Given temporal expressions p and q , (i) $\overline{p \wedge q} = \bar{p} \vee \bar{q}$. (ii) $\overline{p \vee q} = \bar{p} \wedge \bar{q}$. (iii) $\overline{p \cdot q} = \bar{p} \vee \bar{q} \vee (q \cdot \bar{p})$. We assume that events are nonrepeating. In practice, transaction IDs or timestamps differentiate multiple instances of the same event type. Commitments are manipulated by the operations *create* (an agent creates a commitment toward someone), *cancel* (a debtor withdraws its commitment), *release* (an agent withdraws a commitment of which it is the creditor), *assign* (a new creditor is specified by the previous one), *delegate* (a new debtor is specified by the previous one), *discharge* (the commitment is resolved) [Yolum and Singh, 2002]. Figure 1 shows the commitment life cycle [Telang *et al.*, 2011].

A commitment is *Violated* when its antecedent is true but its consequent will forever be false, or it is canceled when *Detached* (in informal terms, the debtor is liable for the violation); *Satisfied*, meaning that the engagement is accomplished; *Expired*, meaning that it is no longer in effect and therefore the debtor would not fail to comply even if does not accomplish the consequent; Typically, a commitment should be *Active* when it is initially created: *Conditional* if its antecedent is not true and *Detached* if its antecedent is true. *Active* has two substates: *Conditional* (as long as the antecedent does not occur) and *Detached* (the antecedent has occurred).




Figure 1: Commitment life cycle.

Our semantics introduces the elements necessary for the interpretation of dialectical commitments. Briefly, the semantics is given in terms of a model, $M = \langle \mathbb{E}, \mathbb{T}, \mathbb{C}^p, \mathbb{D}^p, \mathbb{X}^p, \mathbb{V}^p, \mathbb{C}^d, \mathbb{D}^d, \mathbb{X}^d, \mathbb{V}^d \rangle$. \mathbb{E} is a denumerable set of possible events; \mathbb{T} is the set of possible event runs. \mathbb{C}^p and \mathbb{C}^d are the model standards for (active) commitments. That is, at each index on each run, for each debtor-creditor (ordered) pair of agents, \mathbb{C}^p (\mathbb{C}^d) assigns to a set of runs a set of sets of runs. The intuition is that the model determines which conditional commitment is active from a debtor to a creditor at an index in a run. $\mathbb{D}^p, \mathbb{X}^p, \mathbb{V}^p$ are respectively the standards for discharged, expired, and violated practical commitments. And, similarly $\mathbb{D}^d, \mathbb{X}^d, \mathbb{V}^d$ for dialectical commitments. We assume semantic postulates 1–8 from [Marengo *et al.*, 2011]. $\llbracket q \rrbracket$ is the intension of q , that is, the set of runs where it is true on index 0: $\llbracket q \rrbracket = \{ \tau \mid \tau \models_0 q \}$. And, $\tau_{i,j}$ refers to projection of τ from index i to j , both inclusive. We add the following semantic postulate to accommodate dialectical commitments:

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

Intuitively, a dialectical commitment is active at index i of run τ iff the intension of its consequent condition q belongs to the set of sets returned by $\mathbb{C}_{x,y}^d(\tau, i, \llbracket r \rrbracket)$. For operations, we define the following semantic postulates:

$$\begin{aligned}
& \tau \models_i \text{Create}(C(x, y, r, u)) \text{ iff} \\
& \quad \tau \not\models_i C(x, y, r, u) \text{ and } \tau \models_{i+1} C(x, y, r, u) \\
& \tau \models_i \text{Discharge}(C(x, y, r, u)) \text{ iff} \\
& \quad \tau \models_i C(x, y, r, u) \text{ and } \tau \models_{[0, i+1]} \bar{u} \\
& \tau \models_i \text{Expire}(C(x, y, r, u)) \text{ iff} \\
& \quad \tau \models_i C(x, y, r, u) \text{ and } \tau_{[0, i+1]} \models \bar{r} \\
& \tau \models_i \text{Violate}(C(x, y, r, u)) \text{ iff} \\
& \quad \tau \models_i C(x, y, r, u) \text{ and } \tau \models_{[0, i+1]} \bar{u} \\
& \tau \models_i \text{Detach}(C(x, y, r, u)) \text{ iff} \\
& \quad \tau \models_i C(x, y, r, u) \text{ and } \tau \models_{[0, i+1]} r
\end{aligned}$$

Commitments *persist* until they are discharged, expired, violated. When this happens, the corresponding commitment no longer holds (and never holds again). We adopt the notion of *residuation* [Marengo *et al.*, 2011; Singh, 2003] to track progress in the real world. The residual q/e of a condition q with respect to an event e is the remainder condition that would be left over after the event occurs, and whose satisfaction would guarantee the satisfaction of the original condition. Residuation helps compute commitment progress given an event occurrence. Below, when $c = C(x, y, r, u)$, c/e denotes $C(x, y, r/e, u/e)$. A commitment can progress toward its residual with respect to an event (its antecedent and consequent are residuated), or it can expire, be violated, or be discharged. This is captured by the following theorem [Marengo

et al., 2011], which shows how a commitment (either practical or dialectical) progresses.

Below, following [Singh, 2003], \top means the temporal expression satisfied by every run and 0 an expression that is satisfied by no run.

Theorem 1 *If $\tau \models_i C(x, y, r, u)$ and $\tau_{i+1} = e$, then*

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

3.1 Protocol enactments

Below, a principal refers to an active social entity such as a person. Let \mathcal{P} be a nonempty, finite set of principals. Let \mathcal{B} be a nonempty set of events. Let \mathcal{E} be the set of event temporal expressions generated from \mathcal{B} . Let $\mathcal{C} = \{C(x, y, r, u) : x, y \in \mathcal{P} \text{ and } r, u \in \mathcal{E}\}$, be the set of possible commitments. Let \mathcal{S} be the set of possible operations on commitments. Let $\mathcal{A} = \mathcal{B} \cup \mathcal{S}$ be a set of events. A (social) context relates principals with the events they observe in the context. Intuitively, \mathcal{C} is the set of commitment specifications for the context.

Definition 1 *A context is a tuple $X = \langle \alpha, \mathcal{C}, \mathcal{P} \rangle$ where α is a set of events in \mathcal{A} and $\mathcal{C} \subseteq \mathcal{C}$ is a set of commitments (over \mathcal{P} and \mathcal{E}), and \mathcal{P} is a set of principals. A context is nonempty when $\alpha \neq \emptyset$, and $\mathcal{P} \neq \emptyset$. Given an event f , the progression of X under f is a context $X/f = \langle \alpha, \{c/f : c \in \mathcal{C}\} \rangle$. Given two contexts, $X = \langle \alpha, \mathcal{C}, \mathcal{P} \rangle$ and $X' = \langle \alpha', \mathcal{C}', \mathcal{P}' \rangle$, their union is the context $X \cup X' = \langle \alpha \cup \alpha', \mathcal{C} \cup \mathcal{C}', \mathcal{P} \cup \mathcal{P}' \rangle$.*

Definition 2 *Let $X = \langle \alpha, \mathcal{C}, \mathcal{P} \rangle$ be a context. For an event $a \in \mathcal{A}$, the observers of a in X , $\text{obs}(a) \subseteq \mathcal{P}$, is a set of principals who can observe the occurrence a .*

Definition 3 *An event $e \in \mathcal{E}$ is relevant to a commitment $C(x, y, r, u) \in \mathcal{C}$ iff one of the following holds: $r/e \neq r$, $r/\bar{e} \neq r$, $u/e \neq u$, or, $u/\bar{e} \neq u$.*

Intuitively, Definition 3 states that e is relevant to $C(x, y, r, u)$ when e is involved either in r or u . That is, e is significant in the expiration, violation, discharge, detachment, or progression of the commitment. Clearly, event observability determines the observability of commitment operations, e.g., the observability of a commitment creation is determined by the observability of the message that brings about its creation, and the observability of a detach is determined by the one of the events that are relevant to its antecedent.

An atomic protocol involves two roles (sender and receiver), a single message between them, and a set of commitments. The message meaning is its effect on some commitment. For example, a payment is an atomic protocol as is the physical delivery of an item by a courier to a recipient. The social events corresponding to a physical event occur concurrently with the physical event [Goldman, 1970].

A protocol involves two or more roles and describes what protocols those roles participate in, and a set of commitments. For example, the purchase protocol describes how a customer, merchant, and courier interact to carry out their

business transaction and can be built as the composition of three pair-wise protocols.

Definition 4 *A protocol is a triple $\langle R, M, I \rangle$, where R is a set of roles, M is a set of messages, and I is a set of commitments (that hold at the beginning of the interaction). An atomic protocol is a protocol where M is a singleton. Below, the suffix k can either be p or d , to specify the kind of commitment on which operations are applied.*

PROTOCOL \rightarrow axiom { , axiom }
 axiom \rightarrow $\langle\langle$ message { means social { , social } } $\rangle\rangle$
 social \rightarrow op(R, R , condition, condition)
 op \rightarrow Create^k | Cancel^k | Release^k | Assign^k | Delegate^k
 condition \rightarrow tempcond { \wedge tempcond }
 tempcond \rightarrow sequence { \vee sequence }
 sequence \rightarrow 0 | \top | elem | elem \cdot elem
 elem \rightarrow message | C^k(R, R , condition, condition)

R yields role names (a role is a placeholder for a principal who will enact the protocol), *message* yields physical events and *social* is an operation on commitments. We assume that every physical event has a corresponding social event, that for simplicity we assume having the same name. Below, we denote by M_P and C_P respectively the set of social events and of commitments that can possibly be generated during a protocol enactment.

Definition 5 *Let $P = \langle R, M, I \rangle$ be a protocol and let $X = \langle \alpha, \mathcal{C}, \mathcal{P} \rangle$ be a context. X enacts P iff for all $r \in R$ there is a principal $p \in \mathcal{P}$ that plays the role r , $M_P \subseteq \alpha$, $C_P \subseteq \mathcal{C}$, and $I \subseteq \mathcal{C}$. We denote the enactment of P by X as $P \triangleright X$. An event $a \in \mathcal{A}$ arises in an enactment iff there is a $p \in \mathcal{P}$ that plays a role $r \in R$ that can make action a occur. We denote by $\mathcal{P} \dagger R$ the set of principals that play protocol P 's roles in the enactment.*

Definition 6 *An enactment $E = P \triangleright X$ is social iff for every $a \in \mathcal{A}$ that arises in E , $\text{obs}(a) \supseteq \mathcal{P} \dagger R$.*

Definition 7 states that in a closed enactment, for any relevant event, there is at least one principal from the context who can make it occur.

Definition 7 *An enactment $P \triangleright X$ is closed iff for every commitment $c \in C_P$, every event $a \in \mathcal{A}$ that is relevant to c arises in $P \triangleright X$. A context that is not closed is open.*

Proposition 1 (Stability of Closure under Progress) *Let E be a closed social enactment. Let f be an event. Then E/f is a closed social enactment.*

Proof. From induction over the structure of formulas. ■

Definition 8 *Let $P \triangleright X$ be an enactment. Then P is strictly verifiable in X iff the detach and discharge events of any commitment in C_P can be observed in X .*

Clearly, a closed and social enactment is strictly verifiable but in all other cases strict verifiability does not necessarily hold because some relevant events may not be observable (or at least not be observable by the debtor and the creditor). Strict verifiability is too strong a condition: It generally works only for protocols being enacted in a monolithic context. For example, in a purchase protocol where the merchant commits to the customer but it is a courier who makes the delivery, strict

verifiability would hold only for contexts where the merchant is present to see the delivery done. In general, ensuring the merchant’s presence is impractical—the courier is often hired because the merchant can’t make the delivery directly.

Indeed, it is reasonable to think of real-world enactments as composable: that is, an enactment is generally the composition of simpler enactments where simpler protocols are executed in possibly different contexts. The aim of enactment composition is to create closed enactments, where all the relevant events arise.

Definition 9 Given two enactments $E = P \triangleright X$ and $E' = P' \triangleright X'$, where $P = \langle R, M, I \rangle$, $P' = \langle R', M', I' \rangle$, $X = \langle \alpha, C, \mathcal{P} \rangle$, and $X' = \langle \alpha', C', \mathcal{P}' \rangle$, the composition $E \diamond E' = Q \triangleright Z$, where $Q = \langle R \cup R', M \cup M', I \cup I' \rangle$ and Z is the context $\langle \alpha \cup \alpha', C \cup C', \mathcal{P} \cup \mathcal{P}' \rangle$.

Interestingly, while the composition of a set of closed enactments is a closed enactment, the composition of social enactments is not necessarily social. The property depends on the observability of events which does not necessarily extend to the whole set of principals playing roles in the composed protocol. Conversely, the composition of a nonsocial enactment with other enactments yields a nonsocial enactment because composition does not modify the set of observers of an event. In such a situation, however, principals may not remain aligned in their views of commitments and of the progression of these commitments due to the partial observability of events.

Example 3 Consider purchase at shop, involving a shopkeeper and client, as explained in Example 1. We capture this as a context X , where $\mathcal{P} = \{\text{shopkeeper}, \text{client}\}$; $\alpha \supseteq \{\text{pay}, \text{give}, \text{offer}\}$, where *pay*, *give* and *offer* are observable by both principals, $C = \{C(\text{shopkeeper}, \text{client}, \text{pay}, \text{give})\}$. Let *PAY* be the atomic protocol $\langle \{\text{buyer}, \text{seller}\}, \{\text{pay}\}, \emptyset \rangle$ and suppose that shopkeeper and client respectively play the roles *seller* and *buyer*. Let *GIVE* be the atomic protocol $\langle \{\text{buyer}, \text{seller}\}, \{\text{give}\}, \{C^p(\text{seller}, \text{buyer}, \text{pay}, \text{give})\} \rangle$.

$\text{PAY} \triangleright X$ is a closed, social enactment. $\text{PAY} \circ \text{GIVE} \triangleright X$ is a composite enactment that realizes a purchase. It is a closed, social (thus, strictly verifiable) enactment.

Example 4 Let *OFFER* be the protocol $\langle \{\text{buyer}, \text{seller}\}, \{\text{offer means create}(C^p(\text{seller}, \text{buyer}, \text{pay}, \text{deliver}))\}, \emptyset \rangle$. When *deliver* does not arise in the context (e.g. *delivery at home is not performed directly by the seller*), $\text{OFFER} \triangleright X$ is an open but social enactment.

Let *DELIVER* be the atomic protocol $\langle \{\text{buyer}, \text{courier}\}, \{\text{deliver}\}, \emptyset \rangle$, and Y be a context, where $\mathcal{P} = \{\text{client}, \text{courier}\}$; $\alpha \supseteq \{\text{deliver}\}$, *deliver* is observable by client and courier, and $C = \emptyset$. $\text{DELIVER} \triangleright Y$ is a closed, social enactment.

$\text{PAY} \cup \text{OFFER} \cup \text{DELIVER} \triangleright X \cup Y$ is a closed, nonsocial enactment where the discharge of the commitment created by *offer* cannot be observed by shopkeeper because he does not observe *deliver*. Strict verifiability does not hold. However, if courier is trusted, shopkeeper could accept a claim by courier that *delivery occurred*.

A proposition is *acceptable* to an agent in an enactment if and only if it can support the proposition by direct observation or by a dialectical commitment from a trusted party.

Definition 10 Let $P \triangleright X$ be an enactment. Then P is weakly verifiable in X iff for each commitment in C_P , whose detach and discharge belong to the set of events α from X , its creditor accepts discharge and its debtor accepts detach.

All strictly verifiable enactments are weakly verifiable.

Now we come to the main negative result of this paper. In general, strict verifiability is not preserved by enactment composition: even in the case in which one composes two closed, social (and, thus, strictly verifiable) enactments, the resulting enactment is not necessarily social (as observed above). Thus, the observability of detach and discharge is not granted. Weak verifiability substitutes direct observability with the creation (and progress) of dialectical commitments. Similarly to the previous case, weak verifiability is not preserved by enactment composition. Again, event observability is not granted in the composed enactment.

Proposition 2 (Verifiability under Composition) There exist protocols P, P' and contexts X, X' such that P is weakly verifiable in X and P' is weakly verifiable in X' , yet $P \cup P'$ is not weakly verifiable in $X \cup X'$.

Consider the composed enactment $E \diamond E'$ where E and E' are weakly verifiable, and their sets of events α and α' are disjoint. Suppose events a, b, c arise in E : a is observed only by x ; b creates $C^p(x, z, \top, a)$; c creates $C^d(x, z, \top, a)$. Suppose E' is an open enactment because there is an event that creates $C1 = C^p(x', z', a, d)$, but a does not belong to E' . Then, in $E \diamond E'$ the principals playing z and z' are different and, by construction, neither of the principals playing x' and z' can observe a or c . Then, z' cannot observe $C1$ detach but this yields that $E \diamond E'$ is not weakly verifiable. ■

Proposition 2 shows that it is possible to create composite protocols that are not verifiable even though their constituent protocols are. We could require a monolithic context, but that would be counterproductive. Instead, we provide a way to develop new protocols that relies upon pragmatic rules and upon the addition of new messages that create dialectical commitments. Each pragmatic rule (Section 3.2) is a way to compose protocols. Potentially, each pragmatic rule imposes distinct conditions on contexts to ensure weak verifiability.

3.2 Bridging Enactments by Pragmatic Rules

We now show how an enactment that is not weakly verifiable can be turned into one that is *weakly verifiable*: here, dialectical commitments are used as claims that support the progress of practical commitments with respect to events that are not observable. Based on the above, we can relax the notion of satisfaction of a commitment to the following: that its creditor accepts that it is discharged (independent of whether it is in fact discharged). A good protocol would make sure that, for any commitment, the creditor can determine whether the debtor discharges it. In the present setting, we relax this requirement to the creditor *accepting* that the commitment is discharged.

To realize this view we introduce *pragmatic rules*: patterns of pragmatic reasoning that principals may or may not adopt

in an enactment. Adopting a pragmatic rule means that all principals in the context share the rule. Specifically, we use pragmatic rules to capture the interplay between practical and dialectical commitments. We present the pragmatic rules extending the operational semantics in [Telang *et al.*, 2011] with dialectical commitments. The configuration of an agent x is $S_x = \langle \mathcal{B}, \mathcal{G}, \mathcal{C} \rangle$ where \mathcal{B} is its set of beliefs about the current snapshot of the world, and \mathcal{C} its set of commitments. \mathcal{G} is x 's set of goals: it is not used in the following. The operational semantics is given via guarded rules in which the S_i are configurations and $S_1 \longrightarrow S_2$ is a transition:

$$\frac{\text{guard}}{S_1 \longrightarrow S_2}.$$

In most settings, it is possible to specify a family of transitions as an action. For example, for a commitment C , $\text{suspend}(C)$ refers to the set of transitions $S_i \longrightarrow S_j$ where $C \in S_i.\mathcal{C}$ (the set of commitments of configuration S_i) and $\text{suspend}(C) \in S_j.\mathcal{C}$. For actions a and b , $a \wedge b$ indicates that both must be performed. The same guard may enable multiple transitions $S_i \longrightarrow S_j$ with the same S_i , indicating choice of the agent involved.

Algorithm 1 Enactment Closure under Claim

Require: $E = P \triangleright X$ is an enactment; $X = \langle \alpha, C, \mathcal{P} \rangle$; $P = \langle R, M, I \rangle$

Require: A' is the set of events a' , each of which is relevant to some practical commitment $c \in C_P$, and for each of which $\exists p' \in \mathcal{P} \uparrow R$ and $p' \notin \text{obs}(a')$

Require: for any $a' \in A'$, there is at least one $p \in \mathcal{P}$ such that $p \in \text{obs}(a')$

```

1: function CLOSURE( $X = \langle \alpha, C, \mathcal{P} \rangle$ ,  $P = \langle R, M, I \rangle$ ,  $A'$ )
2:    $\alpha' \leftarrow \alpha$ ;  $C' \leftarrow C$ ;  $M' \leftarrow M$ ;
3:   for all  $a' \in A'$  do
4:      $x \leftarrow$  one of  $p \in \mathcal{P}$ , such that  $p \in \text{obs}(a')$ 
5:      $Y \leftarrow \mathcal{P} \uparrow R - \text{obs}(a')$ 
6:      $M' \leftarrow M' \cup$ 
7:        $\{ \text{claim}_{a'} \text{ means } \text{Create}(C^d(x, Y, \top, a')) \}$ 
8:      $\alpha' \leftarrow \alpha' \cup \{ \text{Create}(C^d(x, Y, \top, a')), \dots \}$ 
9:      $C' \leftarrow C' \cup \{ C^d(x, Y, \top, a') \}$ 
10:     $\text{obs}(\text{claim}_{a'}) \leftarrow \mathcal{P} \uparrow R$ 
11:   end for
12:    $X' \leftarrow \langle \alpha', C', \mathcal{P} \rangle$ ;  $P' \leftarrow \langle R, M', I \rangle$ 
13:   return  $P' \triangleright X'$ 
14: end function
```

Structural rules Telang *et al.* [2011] specify the progression of a commitment. Their guards are events that occur in the social state and are observed by the agent. We extend such rules to manage residuation because in our setting, the occurrence of events makes commitments progress via residuation:

$$\frac{\mathcal{B}_x \models t}{C(x, y, r, q) \longrightarrow C(x, y, r/t, q/t)}$$

We now introduce the pragmatic rule *Social Recognition* that enables considering dialectical commitments to events as sufficient claims that can be used to progress commitments.

$$\frac{C^d(w, z, \top, t)}{C^p(x, y, r, q) \longrightarrow C^p(x, y, r/t, q/t)} \text{ SR}$$

The intent is detaching or discharging practical commitments based upon the creation of appropriate dialectical commitments; alternatively, causing the violation or expiration of practical commitments in presence of appropriate dialectical commitments. For instance, the shopkeeper cannot physically observe the delivery of goods by the courier but since the shopkeeper trusts the courier, he accepts a dialectical commitment from the latter that the goods were delivered. To exploit the pragmatic rule SR, we require that:

$$\tau \models_i C^d(w, z, \top, t) \wedge \tau \models_i C^p(x, y, r, u) \Rightarrow \tau \models_{i+1} C^p(x, y, r/t, u/t) \quad (\text{R1})$$

To make an enactment weakly verifiable it is necessary to (1) close the enactment by composing the enactment with further enactments where relevant events, that do not arise in the first enactment, can occur; (2) make the enactment social via dialectical commitments, so that SR can be applied to support commitment progression. Given a closed, nonsocial enactment $E = P \triangleright X$, Algorithm 1 creates an enactment that extends it, and that is weakly verifiable.

Proposition 3 (Weak verifiability under Claim)

Consider a closed, nonsocial enactment $E = P \triangleright X$, where $X = \langle \alpha, C, \mathcal{P} \rangle$ and $P = \langle R, M, I \rangle$, such that: (1) for any relevant event that does not arise in E , some principal in $P \uparrow R$ observes it; (2) each $p \in P \uparrow R$ adopt Social Recognition. The enactment $E' = P' \triangleright X'$ that is produced by applying Algorithm 1 to E is weakly verifiable.

Proof. Suppose that there is a commitment $c \in C_{P'}$ and a pair of principals $p', p'' \in P'$, such that c is discharged for p but not for p' . Then, there must be an event a_k in the run that is observed by p but not by p' . However, either a_k arises in $P \triangleright X$ and is observable by all principals (by our hypothesis since the enactment is social), or it is a “claim” event, added by Algorithm 1, and which means the creation of a dialectical commitment. The event is observed by all principals by construction. SR makes the commitments progress in the same way for both p and p' , so the hypothesis is absurd. And, analogously for the other commitment state changes. ■

When the initial enactment is not closed, Algorithm 1 completes it by introducing the means for the principals to claim the occurrence of only part of the relevant events that cannot be observed. In order to tackle those events that do not arise in the enactment it is necessary to compose the enactment with others where such events arise. It is possible to make an enactment closed and social by iteratively alternating composition and Algorithm 1. A dialectical commitment *bridges two enactments*: one in which the event occurs and another in which the event is socially relevant but where it does not arise.

Example 5 Example 4 introduces the closed, non social enactment $\text{PAY} \cup \text{OFFER} \cup \text{DELIVER} \triangleright X \cup Y$, *deliver* cannot be observed by shopkeeper. Algorithm 1 makes it weakly verifiable by adding the protocol action $\text{claim}_{\text{deliver}}$ means $\text{Create}(C^d(\text{courier}, \text{shopkeeper}, \top, \text{deliver}))$ using which courier can state that delivery occurred, taking liability for the truth of its claim. The commitment $C^d(\text{courier}, \text{shopkeeper}, \top, \text{deliver})$ is added to the set of all possible commitments of the context.

This leads us to Theorem 2 on commitment progression. The proof is simple by construction. The theorem shows that practical commitments progress when adopting Social Recognition and considering dialectical commitments as evidence of event occurrence.

Theorem 2 *Given a weakly verifiable enactment $P \triangleright X$, and assuming requirement (R1), if $\tau \models_i C(x, y, r, u)$ and $\tau_i = e$, then:*

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

4 Conclusions

A proper treatment of the challenges of secure collaboration and governance requires that we model sociotechnical systems [Singh, 2013; Sommerville, 2010] in formal terms. Crucial to such systems is the incorporation of a societal perspective. It helps to understand sociotechnical systems as performing a *social computation*, taken as the sum of the independent contributions of autonomous and heterogeneous parties [Singh, 2014; Baldoni *et al.*, 2014].

Our approach addresses the challenges of passing to a societal perspective by studying the impact of event observability on commitments and by supplying a notion of progression of commitments that relies upon dialectical commitments, and on the interplay between these and practical commitments, captured through a pragmatic rule. To the best of our knowledge, the notion of observability is original to this paper.

Although we emphasize just one pragmatic rule in this paper, social recognition, other such rules could be defined. It would be interesting to study others as well as their joint use, in order, e.g., to turn practical commitments into dialectical ones, or to create dialectical commitments based on dialectical commitments by other agents.

Our approach provides a valuable tool for cross-organizational process modeling, and it would be interesting to combine it with design methodologies where the specification of commitment-based protocols and of their composition is fundamental, like Amoeba [Desai *et al.*, 2009] and the 2CL Methodology [Baldoni *et al.*, 2014]. In the former, for instance, a real-life insurance claim processing case is analyzed, which involves a number of principals (insurance company, call center, mechanic, claim handler). Once the individual protocols are identified, they are composed into an overall process, which may still be modified if some changes occur, e.g. as in the case of subcontracting or of outsourcing. Clearly, observability and weak verifiability play a fundamental role in this kind of application.

Currently, dialectical commitments are used in formal dialog systems, dialog games, and argumentation, e.g., [McBurney and Parsons, 2002; Norman *et al.*, 2004; Chaib-draa *et al.*, 2006; Wells and Reed, 2005; Singh, 2008]. In contrast

to such works, we adopt a conception wherein the connection to real world is key and introduce an indirection from objective reality to a commitment by an agent about the reality; further, we enable principals to commit both dialectically and practically, and exploit dialectical commitments to support practical ones.

Our approach opens the way to the study of protocols based on normative relationships that are modular and composable. In this respect, we claim that social contexts provide abstraction spaces: each social context potentially involves its own vocabulary of interaction that defines the facts and actions of relevance to the context. That is, the heart of a social context lies not in the specific principals who participate in it but in the interpretation of their interactions. We mean to study relationships between contexts, e.g., that a context observes another context or that an event in one context counts as another event in a different context, as well as how contexts can be composed into more complex contexts.

Our approach and its possible future advancements bear interesting implications in the development of electronic institutions [Arcos *et al.*, 2007; Esteva *et al.*, 2004; Fornara *et al.*, 2008]. We provide a unified framework for reasoning about violations, be they captured by practical or by dialectical commitments. One development that we mean to pursue is to introduce a first-order representation of commitment conditions and to allow commitment conditions to progress in the presence of incomplete information based on works such as [Montali *et al.*, 2014]. A related direction of future work involves modeling explicit deadlines [Chesani *et al.*, 2009; Chopra and Singh, 2015a]. Deadlines would be useful to identify classes of violations (and the liable parties) in real-world applications. We will also study the relations between dialectical commitments, reputation, and trust. Another future development is to study whether the notion of alignment due to information propagation in [Chopra and Singh, 2015b], and the notion of alignment due to nonuniform observability, as presented in this paper, can be unified into a single framework.

Acknowledgments

We thank the anonymous reviewers for helpful comments. Munindar Singh’s research was partially supported by the US Department of Defense through a Science of Security Label and partially by the National Science Foundation under grant IIS-0910868. Matteo Baldoni and Cristina Baroglio are partially supported by the EU FP7 project EUCASES, grant agreement n. 611760.

References

- [Arcos *et al.*, 2007] J. L. Arcos, P. Noriega, J. A. Rodríguez-Aguilar, and C. Sierra. E4MAS through electronic institutions. In *Proc. E4MAS, LNCS 4389*. Springer, 2007.
- [Astefanoaei, 2011] Lacramioara Astefanoaei. *An executable theory of multi-agent systems refinement*. PhD thesis, Leiden University, 2011.
- [Baldoni *et al.*, 2014] M. Baldoni, C. Baroglio, and F. Cappuzzimati. A commitment-based infrastructure for pro-

- gramming socio-technical systems. *ACM Transactions on Internet Technology*, 14(4), 2014.
- [Baldoni *et al.*, 2014] M. Baldoni, C. Baroglio, E. Marengo, V. Patti, and F. Capuzzimati. Engineering commitment-based business protocols with the 2CL methodology. *Autonomous Agents and Multi-Agent Systems*, 28(4), 2014.
- [Chaib-draa *et al.*, 2006] B. Chaib-draa, M.-A. Labrie, M. Bergeron, and P. Pasquier. DIAGAL: An agent communication language based on dialogue games and sustained by social commitments. *Autonomous Agents and Multi-Agent Systems*, 13(1), 2006.
- [Chesani *et al.*, 2009] F. Chesani, P. Mello, M. Montali, and P. Torroni. Commitment tracking via the reactive event calculus. In *Proc. IJCAI*, pages 91–96, 2009.
- [Chesani *et al.*, 2013] Federico Chesani, Paola Mello, Marco Montali, and Paolo Torroni. Representing and monitoring social commitments using the event calculus. *Autonomous Agents and Multi-Agent Systems*, 27(1):85–130, 2013.
- [Chopra and Singh, 2015a] Amit K. Chopra and Munindar P. Singh. Cupid: Commitments in relational algebra. In *Proceedings of the AAAI*, pages 2052–2059, 2015.
- [Chopra and Singh, 2015b] Amit K. Chopra and Munindar P. Singh. Generalized commitment alignment. In *Proc. AAMAS*, 2015.
- [Dastani *et al.*, 2004] M. Dastani, B. van Riemsdijk, J. Hulstijn, F. Dignum, and J.-J. Ch. Meyer. Enacting and detecting roles in agent programming. In *Agent-Oriented Software Engineering V, Revised Selected Papers*, volume 3382 of *LNCS*, pages 189–204. Springer, 2004.
- [Desai *et al.*, 2009] N. Desai, A. K. Chopra, and M. P. Singh. Amoeba: A methodology for modeling and evolving cross-organizational business processes. *ACM Trans. on Software Engineering Methodologies*, 19(2), 2009.
- [El-Menshawey *et al.*, 2010] Mohamed El-Menshawey, Jamal Bentahar, and Rachida Dssouli. Verifiable semantic model for agent interactions using social commitments. In *LADS*, volume 6039 of *LNCS*. Springer, 2010.
- [El-Menshawey *et al.*, 2013] Mohamed El-Menshawey, Jamal Bentahar, Warda El Kholy, and Rachida Dssouli. Verifying conformance of multi-agent commitment-based protocols. *Expert Syst. Appl.*, 40(1):122–138, 2013.
- [Esteva *et al.*, 2004] M. Esteva, B. Rosell, J. A. Rodríguez-Aguilar, and J. L. Arcos. AMELI: An agent-based middleware for electronic institutions. In *Proc. AAMAS*, pages 236–243. IEEE Computer Society, 2004.
- [Fornara *et al.*, 2008] N. Fornara, F. Viganò, M. Verdicchio, and M. Colombetti. Artificial institutions: a model of institutional reality for open multiagent systems. *Artificial Intelligence and Law*, 16(1):89–105, 2008.
- [Gerard and Singh, 2013] Scott N. Gerard and Munindar P. Singh. Formalizing and verifying protocol refinements. *ACM TIST*, 4(2):21:1–21:27, 2013.
- [Goldman, 1970] Alvin I. Goldman. *A Theory of Human Action*. Prentice-Hall, Englewood Cliffs, NJ, 1970.
- [Lomuscio *et al.*, 2012] Alessio Lomuscio, Hongyang Qu, and Monika Solanki. Towards verifying contract regulated service composition. *Autonomous Agents and Multi-Agent Systems*, 24(3):345–373, 2012.
- [Marengo *et al.*, 2011] Elisa Marengo, Matteo Baldoni, Cristina Baroglio, Amit K. Chopra, Viviana Patti, and Munindar P. Singh. Commitments with regulations: Reasoning about safety and control in REGULA. In *Proc. AAMAS*, pages 467–474, 2011.
- [McBurney and Parsons, 2002] P. McBurney and S. Parsons. Dialogue games in multi-agent systems. *Informal Logic*, 22, 2002.
- [Montali *et al.*, 2014] Marco Montali, Diego Calvanese, and Giuseppe De Giacomo. Verification of data-aware commitment-based multiagent system. In *Proc. International Conference on Autonomous Agents and Multi-Agent Systems, (AAMAS)*, pages 157–164. IFAAMAS/ACM, 2014.
- [Norman *et al.*, 2004] Timothy J. Norman, Daniela V. Carbogim, Erik C. W. Krabbe, and Douglas N. Walton. Argument and multi-agent systems. In *Argumentation Machines*, volume 9 of *Argumentation Library*. Kluwer Academic Publishers, 2004.
- [Singh, 2003] Munindar P. Singh. Distributed enactment of multiagent workflows: temporal logic for web service composition. In *Proc. AAMAS*, pages 907–914. ACM, 2003.
- [Singh, 2008] Munindar P. Singh. Semantical considerations on dialectical and practical commitments. In *Proc. AAAI*. AAAI Press, 2008.
- [Singh, 2013] Munindar P. Singh. Norms as a basis for governing sociotechnical systems. *ACM TIST*, 5(1):21:1–21:23, December 2013.
- [Singh, 2014] Munindar P. Singh. Social computing: Principles, methods, and technologies, 2014. Invited talk at the First Int. Workshop on Multiagent Foundations of Social Computing.
- [Sommerville, 2010] Ian Sommerville. *Software Engineering*. Addison-Wesley, 9 edition, 2010.
- [Telang *et al.*, 2011] Pankaj R. Telang, Munindar P. Singh, and Neil Yorke-Smith. Relating goal and commitment semantics. In *ProMAS*, volume 7217 of *Lecture Notes in Computer Science*, pages 22–37. Springer, 2011.
- [Venkatraman and Singh, 1999] M. Venkatraman and M. P. Singh. Verifying compliance with commitment protocols: Enabling open Web-based multiagent systems. *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)*, 2(3):217–236, September 1999.
- [Wells and Reed, 2005] S. Wells and C. Reed. Formal dialectic specification. In *Argumentation in Multi-Agent Systems*, pages 31–43. Springer, 2005.
- [Yolum and Singh, 2002] Pinar Yolum and Munindar P. Singh. Commitment machines. In J.-J. Ch. Meyer and M. Tambe, editors, *Proc. ATAL 2001*, volume 2333 of *LNCS*, pages 235–247. Springer, 2002.