

Specification and Verification of Commitment-Regulated Data-Aware Multiagent Systems [★]

Marco Montali¹, Diego Calvanese¹, and Giuseppe De Giacomo²

¹ Free University of Bozen-Bolzano, Piazza Domenicani 3, 39100 Bolzano, Italy
lastname@inf.unibz.it

² Sapienza Università di Roma, Via Ariosto, 25, 00185 Rome, Italy
lastname@dis.uniroma1.it

Abstract. In this paper we investigate multiagent systems whose agent interaction is based on social commitments that evolve over time, in presence of (possibly incomplete) data. In particular, we are interested in modeling and verifying how data maintained by the agents impact on the dynamics of such systems, and on the evolution of their commitments. This requires to lift the commitment-related conditions studied in the literature, which are typically based on propositional logics, to a first-order setting. To this purpose, we propose a rich framework for modeling data-aware commitment-based multiagent systems. In this framework, we study verification of rich temporal properties, establishing its decidability under the condition of “state-boundedness”, i.e., data items come from an infinite domain but, at every time point, each agent can store only a bounded number of them.

1 Introduction

In this paper we investigate multiagent systems (MASs) whose agent interaction is based on social commitments that evolve over time, in presence of possibly incomplete data. MASs based on social commitments have been extensively studied in the literature [8]. Intuitively, a social commitment $CC(d, c, q_p, q_d)$ models a relationship between a debtor agent d and a creditor agent c , in which d commits towards c that, whenever condition q_p holds in the system, it will bring about condition q_d in the following course of interaction. Commitments provide a semantics for the agent interaction that abstracts away from the internal agent implementation, and can be thus employed to specify business protocols and contracts. The establishment of commitments is regulated by contracts, which depend on domain-specific events and conditions. Established commitments, in turn, have a lifecycle that is regulated by a so-called *commitment machine* [12] on the basis of such contracts. While in the literature, virtually all the work is based on propositional contents for such commitments [8], here we explicitly manage data described through first-order formalisms, in line with [7]. In other words, we study how data maintained by the agents impact on the dynamics of such systems, and on the evolution of their commitments. Technically, this requires to lift to first-order the notions related to contracts, commitments, and commitment machines.

As a result, we obtain a powerful framework of *data-aware commitment-based MASs* (DACMASs), which incorporates the typical notions of commitment-based MASs

[★] This paper is a short version of [10], to be presented at AAMAS 2014.

but in a rich, data-aware context. In our framework, the commitment machine itself becomes a special agent, called *institutional*, which is a agent, in charge of supporting the evolution of the system according to the commitments. In addition, this agent manages core information about the MAS itself, such as the list of participating agents, which changes over time as the system unfolds. In this light, it maintains and manipulates a common knowledge base, of interest for all the interacting agents. The data manipulated by the agents are described in terms of a domain ontology, expressed in a lightweight description logic (DL), tailored towards ontology-based data access. This ontology provides a common ground for the agent interaction and commitments, establishing the vocabulary that is shared by all of them. In particular, we rely on *DLR-Lite* [5], which is the n -ary version of the DL at the base of the OWL 2 QL profile of the OWL 2 semantic web standard. Each agent has its own data about the domain and the contracts it is involved in, expressed in terms of such ontology. Such data are manipulated through actions, in response to events and according to the commitments in place. At each point in time, only a finite number of data is present in the system. However, such data change over time: old data are removed by the agents, and new data (coming from a countably infinite domain Δ) are inserted.

Our main result is that, when a DACMAS is *state-bounded*, i.e., the number of data that are simultaneously present at each moment in time is bounded, verification of rich temporal properties becomes decidable. More specifically, we are able to check DACMASs against properties expressed in a sophisticated first-order variant of μ -calculus with a controlled form of quantification across states. We do this by exploiting recent results in [2], and reducing verification of state-bounded DACMASs to finite-state model checking through a faithful form of abstraction, essentially obtained by replacing real data items with a finite number of symbolic values, while correctly preserving the relationships among the real data items themselves.

2 Preliminaries

Description Logics (DLs) [1] are logics that represent the domain of interest in terms of *objects*, *concepts*, denoting sets of objects, and *relations* between objects. We consider here the DL *DLR-Lite* [5], which is a DL that belongs to the *DL-Lite* family of lightweight DLs and that is equipped with relations of arbitrary arity. In *DLR-Lite*, concepts C and relations R are built from atomic concepts N and atomic relations P (of arity ≥ 2):

$$C \longrightarrow N \mid P[i] \mid C \sqcap C \qquad R \longrightarrow P \mid P[i_1, \dots, i_h]$$

where $h \geq 2$ and for i_1, \dots, i_h , which denote pairwise distinct components of relation P , we have that $\{i_1, \dots, i_h\} \subseteq \{1, \dots, n\}$, where n is the arity of P . Similarly, $i \in \{1, \dots, n\}$. Intuitively, \sqcap denotes concept conjunction, while $P[i_1, \dots, i_m]$ denotes the projection of relation P on its components i_1, \dots, i_m . This results in a concept if $m = 1$ and in a relation otherwise.

Formally, the semantics of DLs is given in terms of first-order *interpretations* $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, where $\Delta^{\mathcal{I}}$ is a nonempty *interpretation domain*, and $\cdot^{\mathcal{I}}$ is an *interpretation function*, assigning to each concept C a subset $C^{\mathcal{I}}$ of $\Delta^{\mathcal{I}}$, and to each n -ary relation R an n -ary relation $R^{\mathcal{I}}$ over $\Delta^{\mathcal{I}}$ such that

$$\begin{aligned} (C_1 \sqcap C_2)^{\mathcal{I}} &= C_1^{\mathcal{I}} \cap C_2^{\mathcal{I}} \\ (P[i_1, \dots, i_m])^{\mathcal{I}} &= \{(o'_1, \dots, o'_m) \mid \text{there is } \mathbf{o} \in P^{\mathcal{I}} \text{ s.t. } \mathbf{o}[i_j] = o'_j, \text{ for } j \in \{1, \dots, m\}\} \end{aligned}$$

(Here, $\mathbf{o}[i]$ denotes the i -th component of tuple \mathbf{o} .) Also, $\cdot^{\mathcal{I}}$ assigns to each constant a an object $a^{\mathcal{I}}$ of $\Delta^{\mathcal{I}}$. We adopt the unique name assumption, i.e., $a_1 \neq a_2$ implies $a_1^{\mathcal{I}} \neq a_2^{\mathcal{I}}$.

In DLs, knowledge about the domain of interest is encoded in an ontology $\mathcal{O} = \langle \mathcal{T}, \mathcal{A} \rangle$, which is formed by a *TBox* \mathcal{T} , encoding intensional knowledge, and an *ABox* \mathcal{A} , encoding extensional knowledge about individuals objects.

A *DLR-Lite* TBox is a finite set of assertions of the form:

$$\begin{aligned} E_1 &\sqsubseteq E_2 && \text{(concept/relation inclusion assertion),} \\ E_1 &\sqsubseteq \neg E_2 && \text{(concept/relation disjointness assertion),} \\ (\text{key } i_1, \dots, i_\ell : R) &&& \text{(key assertion),} \end{aligned}$$

where R has arity n , and $1 \leq i_1 < i_2 < \dots < i_\ell \leq n$. To ensure decidability of inference, and good computational properties, we require that no relation P can appear both in a key assertion and in the right hand side of a relation inclusion assertion [11,5].

A *DLR-Lite* ABox is a finite set of assertions of the form:

$$\begin{aligned} N(a_1) &&& \text{(concept membership assertion),} \\ P(a_1, \dots, a_n) &&& \text{(relation membership assertion),} \end{aligned}$$

where P has arity n , and a_1, \dots, a_n denote constants.

The semantics of an ontology is given by stating when an interpretation \mathcal{I} satisfies an assertion, where \mathcal{I} satisfies: $E_1 \sqsubseteq E_2$, if $E_1^{\mathcal{I}} \subseteq E_2^{\mathcal{I}}$; $E_1 \sqsubseteq \neg E_2$, if $E_1^{\mathcal{I}} \cap E_2^{\mathcal{I}} = \emptyset$; (key $i_1, \dots, i_\ell : R$), if there are no two distinct tuples in $R^{\mathcal{I}}$ that agree on all their components i_1, \dots, i_ℓ ; $N(a_1)$, if $a_1^{\mathcal{I}} \in N^{\mathcal{I}}$; and $P(a_1, \dots, a_n)$, if $(a_1^{\mathcal{I}}, \dots, a_n^{\mathcal{I}}) \in P^{\mathcal{I}}$. A *model* of an ontology $\mathcal{O} = \langle \mathcal{T}, \mathcal{A} \rangle$ is an interpretation that satisfies all assertions in \mathcal{T} and \mathcal{A} . An ontology \mathcal{O} is *satisfiable* if it has at least one model, and it *logically implies* an assertion α , written $\mathcal{O} \models \alpha$, if all models of \mathcal{O} satisfy α .

Next we introduce queries, whose answers, as usual in ontologies, are formed by constants denoting individuals explicitly mentioned in the ABox. A *union of conjunctive queries* (UCQ) q over an ontology $\langle \mathcal{T}, \mathcal{A} \rangle$ is a FOL formula of the form $\bigvee_{i=1}^n \exists \mathbf{y}_i \text{conj}_i(\mathbf{x}, \mathbf{y}_i)$ with free variables \mathbf{x} and existentially quantified variables $\mathbf{y}_1, \dots, \mathbf{y}_n$. Each $\text{conj}_i(\mathbf{x}, \mathbf{y}_i)$ in q is a conjunction of atoms of the form $N(z)$, $P(z)$, where N and P respectively denote a concept and a role name occurring in \mathcal{T} , and z, z are constants of \mathcal{A} or variables in $\mathbf{x}, \mathbf{y}_1, \dots, \mathbf{y}_n$. The *certain answers* to q over $\langle \mathcal{T}, \mathcal{A} \rangle$ is the set $\text{ANS}(q, \mathcal{T}, \mathcal{A})$ of substitutions θ of the free variables of q with constants in \mathcal{A} such that $q\theta$ evaluates to true in every model of $\langle \mathcal{T}, \mathcal{A} \rangle$, denoted $\langle \mathcal{T}, \mathcal{A} \rangle \models q\theta$. *DLR-Lite* enjoys nice computational properties, in particular w.r.t. query evaluation: computing the certain answers to a UCQ can be done in polynomial time in the size of $\langle \mathcal{T}, \mathcal{A} \rangle$, and in AC^0 in the size of \mathcal{A} alone (i.e., in data complexity) [5]. Such result is based on the *FOL rewritability* property of *DLR-Lite* [5], which states that for every UCQ q and TBox \mathcal{T} , we can *rewrite* q into a new UCQ $\text{rew}_{\mathcal{T}}(q)$ such that $\text{ANS}(q, \mathcal{T}, \mathcal{A}) = \text{ANS}(\text{rew}_{\mathcal{T}}(q), \emptyset, \mathcal{A})$, for every ABox \mathcal{A} . In other words, the TBox can be “compiled away”.

We also consider ECQs, which are FOL queries whose atoms are UCQs evaluated according to the certain answer semantics above [4]. An *ECQ* over \mathcal{T} and \mathcal{A} is a possibly open formula of the form (where q is a UCQ):

$$Q \longrightarrow [q] \mid \neg Q \mid Q_1 \wedge Q_2 \mid \exists x.Q$$

The *certain answers* to Q over $\langle \mathcal{T}, \mathcal{A} \rangle$, is the set of substitutions θ of the free variables of Q with constants in \mathcal{A} defined by composing the certain answers of the UCQs q in Q

through first-order constructs, and interpreting existential variables as ranging over the constants in \mathcal{A} . Hence, the first-order constructs in ECQs are interpreted under a (weaker) epistemic semantics. ECQs over *DLR-Lite* ontologies enjoy the same computational properties as UCQs, in particular FOL rewritability of query answering [4].

3 Framework

We introduce now our framework for modeling DACMASs. Formally, a DACMAS is a tuple $\langle \mathcal{T}, \mathcal{E}, \mathcal{X}, \mathcal{I}, \mathcal{C}, \mathcal{B} \rangle$, where: (i) \mathcal{T} is a global *DLR-Lite* TBox; (ii) \mathcal{E} is a set of predicates denoting *events* (where the predicate name is the event type, and the arity determines the content/payload of the event); (iii) \mathcal{X} is a finite set of *agent specifications*; (iv) \mathcal{I} is a (partial) *specification for the institutional agent*; (v) \mathcal{C} is a *contractual specification*; (vi) and \mathcal{B} is a *Commitment Box* (CBox).

3.1 The Global TBox

The global TBox represents the key concepts, relations and constraints characterizing the domain in which the agents operate, so as to provide a *common ground* for the agent interaction. Part of this TBox is fixed for every agent system, and is used to model core notions related to the system itself. The extension of such core notions is maintained by a single, special *institutional* agent, which is also responsible for the manipulation of commitments (cf. Section 3.6). The data maintained by such an agent are publicly available and can be queried by the other agents, but only modified by the institutional agent. Specifically, the institutional agent maintains data about the following relations:

- Agent denotes the set of (names of) agents that currently participate to the system. Since the institutional agent is always part of the system, we fix its name as *inst*, and enforce that *inst* always belongs to the extension of Agent.
- Spec, whose extension is immutable, denotes the set of agent specification names mentioned in \mathcal{X} (cf. Section 3.2).
- *hasSpec* connects agents with their current specification(s): $\text{hasSpec}[1] \sqsubseteq \text{Agent}$, $\text{hasSpec}[2] \sqsubseteq \text{Spec}$.

Each agent, including the institutional agent, maintains a proprietary *DLR-Lite* ABox, in which it stores its own data. Such data can be queried only by the agent itself and by the institutional agent, which exploits the results of such queries to keep track of the evolution of commitments. Furthermore, each agent progresses its own ABox during the execution in such a way that it is always consistent with the global TBox \mathcal{T} . Notice that the overall collection of ABoxes is not assumed to be consistent with the TBox, i.e., the TBox assertions are only required to be satisfied by each agent individually.

Since, in general, queries may involve the ABoxes of several agents, to disambiguate to which ABox a query atom refers, we augment the vocabulary of the TBox with a *location argument* that points to an agent. We use $R@_a(x)$ to denote an atomic query returning the extension of R in the ABox of a . If a does not point to an agent currently in the system, then $R@_a(x)$ evaluates to empty. Beside the special constant *inst*, we also use *self* to implicitly refer to the agent that is posing the query (similarly to “this” in object-oriented programming). When clear from the context, we omit $@_{\text{self}}$ and just use relations without the location argument. We denote with UCQ_ℓ (resp., ECQ_ℓ) the language obtained from UCQ (resp., ECQ) by extending atoms with a location argument.

3.2 Agent Specifications

In a DACMAS, agents interact by exchanging messages. A message is sent by a sender agent to a receiver agent, and is about the occurrence of an *event* with a *payload*, containing data to be communicated. All agents but the institutional one are only aware of the events they send and receive. As for data, the institutional agent has instead full visibility of all exchanged messages, so as to properly handle the evolution of commitments. Agents determine the events they may send, and also how they react to events, through proactive and reactive rules. Such rules are grouped into behavioural profiles called *agent specifications*, and model: (i) the possible, proactive emission of an event, directed to another agent (*communicative rule*); (ii) conditional internal (re)actions, which lead to update the agent ABox when sending/receiving an event to/from another agent (*update rule*). The update could result in the insertion of new data items (from the countably infinite domain Δ), not already present in the system.

The exchange of a message represents a synchronization point among the sender, receiver and institutional agent. Hence, the reaction of the three agents is interpreted as a sort of transaction, such that each of them effectively enforces the update on its own ABox only if each of the three resulting ABoxes is consistent with \mathcal{T} . An inconsistency could potentially arise when reacting to an event either because the same data item is asserted to be member of two disjoint classes, or because a key assertion is violated.

Formally, an agent specification is a tuple $\langle sn, II \rangle$, where sn is the specification name, and II is a set of *communicative* and *update rules*. Such rules are defined over the vocabulary of \mathcal{T} and \mathcal{B} , and are applied over the ABoxes of the agent and of *inst*. This allows the agent to query the status of commitments and obtain the names of the other participants. A *communicative rule* has the form

$$Q(r, \mathbf{x}) \text{ enables } EV(\mathbf{x}) \text{ to } r$$

where Q is an ECQ_ℓ , and $EV(\mathbf{x})$ is an event supported by the system, i.e., predicate $EV/|\mathbf{x}|$ belongs to \mathcal{E} . The semantics of a communicative rule is as follows. Whenever $Q(r, \mathbf{x})$ evaluates positively, the agent autonomously selects one of the answers θ returned by the query, using it to determine the event receiver and its payload. This states that the ground event $EV(\mathbf{x})\theta$ can be sent by the agent to $r\theta$, provided that $r\theta$ points to an actual agent name in the system (including the two special names *inst* and *self*).

Example 1. Consider a DACMAS where customers and sellers interact to exchange goods. We model the behavioural rules for customers and sellers using two agent specifications. To buy from a seller, a customer must register to that seller. A registration request is modeled in the customer specification as:

$$\text{Spec@inst}(sel, seller) \text{ enables } \text{REQ_REG} \text{ to } sel$$

Assuming that each seller maintains its customers and items respectively in relations *MyCust* and *Item*, the proposal of an item to a customer is modeled in the seller specification as: $\text{MyCust}(m) \wedge \text{Item}(i) \text{ enables } \text{PROPOSE}(i) \text{ to } m$. ■

Update rules are ECA-like rules of the form:

- **on** $EV(\mathbf{x})$ **to** r **if** $Q(r, \mathbf{x})$ **then** $\alpha(r, \mathbf{x})$ (on-send)
- **on** $EV(\mathbf{x})$ **from** s **if** $Q(s, \mathbf{x})$ **then** $\alpha(s, \mathbf{x})$ (on-receive)

where $EV/|\mathbf{x}|$ is an event type from \mathcal{E} , Q is an ECQ_ℓ , and α is an update action with parameters (described below). Each such rule triggers when an event is sent to/received from another agent, and Q holds. This results in the application of α using the actual event payload and receiver/sender. Action α queries the ABox of the agent and of inst, using the answers to add and remove facts to the ABox.

Formally, an *update action* is an expression $\alpha(\mathbf{p}) : \{e_1, \dots, e_n\}$, where $\alpha(\mathbf{p})$ is the action signature (constituted by the name α and by a list \mathbf{p} of parameters), and $\{e_1, \dots, e_n\}$ are update effects, each of which has the form

$$[q^+(\mathbf{p}, \mathbf{x})] \wedge Q^-(\mathbf{p}, \mathbf{x}) \rightsquigarrow \mathbf{add} A, \mathbf{del} D$$

– q^+ is an UCQ_ℓ , and Q^- is an ECQ_ℓ whose free variables occur all among those of q^+ ; intuitively, q^+ selects a set of tuples from the agent ABox and that of inst, while Q^- filters away some of them.³ During the execution, the effect is applied with a ground substitution \mathbf{d} for the action parameters, and for every answer θ to the query $[q^+(\mathbf{d}, \mathbf{x})] \wedge Q^-(\mathbf{d}, \mathbf{x})$.

– A is a set of facts (over the alphabet of \mathcal{T} and \mathcal{B}) which include as terms: free variables \mathbf{x} of q^+ , action parameters \mathbf{p} and/or Skolem terms $f(\mathbf{x}, \mathbf{p})$. We use $\text{SKOLEM}(A)$ to denote all Skolem terms mentioned in A . At runtime, whenever a ground Skolem term is produced by applying θ to A , the agent autonomously substitutes it with a possibly new data item taken from Δ . This mechanism is exploited by the agent to inject new data into the system. The ground set of facts so obtained is *added* by the agent to its ABox.

– D is a set of facts which include as terms free variables \mathbf{x} of q^+ and action parameters \mathbf{p} . At runtime, whenever a ground fact in D is obtained by applying θ , it is *removed* from the agent ABox.

As in STRIPS, we assume that additions have priority over deletions (i.e., if the same fact is asserted to be added and deleted during the same execution step, then the fact is added). The “**add** A ” (resp. “**del** D ”) part can be omitted if $A = \emptyset$ (resp., if $D = \emptyset$).

Example 2. Consider three possible reaction rules for the seller. The fact that the seller makes every agent that sends a request become one of its customers is modeled as:

on ASK_REG **from** c **if** *true* **then** *makeCust*(c)

where $\textit{makeCust}(x) : \{\{true\} \rightsquigarrow \mathbf{add}\{\text{MyCust}(x)\}\}$. Assume now that the seller maintains the item cart for a customer, using relation $\text{InCart}(i, c)$ to model that item i is in the cart of c . The seller reaction to an “empty cart” request is modeled as:

on EMPTY_CART_REQ **from** c **if** $\text{MyCust}(c)$ **then** *doEmpty*(c)

where $\textit{doEmpty}(c) : \{\{\text{InCart}(i, c)\} \rightsquigarrow \mathbf{del}\{\text{InCart}(i, c)\}\}$. Note that the effect is applied to each i in the cart of c . Consider now the case where the seller receives a new item i to be sold. It reacts by adding i and deciding its price (denoted with Skolem $p(i)$):

on NEW_ITEM(i) **from** a **if** *true* **then** *addItem*(i)

where $\textit{addItem}(i) : \{\{true\} \rightsquigarrow \mathbf{add}\{\text{Item}(i), \text{Price}(i, p(i))\}\}$. ■

³ The distinction between q^+ and Q^- is needed for technical reasons, borrowed from [2].

3.3 Institutional Agent Specification

The institutional agent *inst* manages the core information of the DACMAS. Its behaviour is (partially) captured by the institutional agent specification \mathcal{I} , which differs from the other agent specifications in two respects. First, since *inst* is aware of all messages exchanged by the other agents and can query their ABoxes, its specification is not only constituted by communicative rules and on-send/on-receive reactive rules, but also by *on-exchange* rules of the form:

$$\text{on EV}(x) \text{ from } s \text{ to } r \text{ if } Q(s, r, x) \text{ then } \alpha(s, r, x)$$

where Q and α can query the internal ABox of the institutional agent, and the ABoxes of s and r . To conveniently specify reactions of *inst* that do not depend on a specific event, but trigger whenever an event is exchanged, we use:

$$\text{on any event from } s \text{ to } r \text{ if } Q(s, r) \text{ then } \alpha(s, r)$$

Second, \mathcal{I} is only a *partial* specification for *inst*. In fact, *inst* is also responsible for the manipulation of commitments, which results in a set of additional on-exchange rules that, starting from the contractual specification (cf. Section 3.4), encode the commitment machines for the commitments involved in the contract. These rules are automatically extracted from the contractual specification (cf. Section 3.6).

Example 3. Consider a portion of institutional agent specification, modeling the creation of a new agent whenever *inst* receives a request (whose payload denotes the specification to be initially followed by that agent). To handle this request, *inst* uses relation *NewA* to store a newly created agent together with its initial specification. The axiom $\text{NewA}[1] \sqsubseteq \neg\text{Agent}$ is part of \mathcal{T} , and enforces that a new agent has indeed a new name. The behaviour is defined in two steps. In the first step, *inst* reacts to a creation request by choosing an agent name (using Skolem term $n()$). The reaction is applied only if there is no pending new agent to be processed.

$$\begin{aligned} &\text{on AG_REQ}(s) \text{ from } a \text{ if } \neg(\exists x \exists y. \text{NewA}(x, y)) \text{ then } \text{create}(s) \\ &\text{create}(s) : \{ [true] \rightsquigarrow \mathbf{add}\{\text{NewA}(n(), s)\} \} \end{aligned}$$

Note that axiom $\text{NewA}[1] \sqsubseteq \neg\text{Agent}$ ensures that the update is blocked if the chosen name is already used in the system.

In the second step, *inst* informs itself that a new agent has to be processed; the corresponding reaction finalizes the insertion of the new agent, moving it to the set of participating agents:

$$\begin{aligned} &\text{NewA}(a, s) \text{ enables INSERT_AG}(a, s) \text{ to self} \\ &\text{on INSERT_AG}(a, s) \text{ from self if } true \text{ then } \text{do_ins}(a, s) \\ &\text{do_ins}(a, s) : \{ [true] \rightsquigarrow \mathbf{add}\{\text{Agent}(a), \text{Spec}(a, s)\}, \mathbf{del}\{\text{NewA}(a, s)\} \} \end{aligned}$$

3.4 Contractual Specification

The contractual specification \mathcal{C} consists of a set of *commitment rules*, which are reactive rules similar to on-exchange rules. The main difference is that, instead of actions, they

describe (first-order) conditional commitments and their creation:

$$\begin{array}{l} \text{on } EV(\mathbf{x}) \text{ from } s \text{ to } r \text{ if } Q_c(s, r, \mathbf{x}) \\ \text{then } CC_n(s, r, [q_p^+(s, r, \mathbf{x}, \mathbf{y})] \wedge Q_p^-(s, r, \mathbf{x}, \mathbf{y}), Q_d(s, r, \mathbf{x}, \mathbf{y})) \end{array} \quad (*)$$

where n is the commitment name, the ECQ_ℓ Q_c is the condition for the creation of the conditional commitment, $[q_p]^+ \wedge Q_p^-$ (where, as in update effects, q_p^+ is a UCQ_ℓ , and Q_p^- is an ECQ_ℓ whose free variables all occur among those of q_p^+) is the *precondition* determining the generation of a corresponding base-level commitment, and the ECQ_ℓ Q_d is the *discharge condition* for such base-level commitment. All the aforementioned queries can be posed over the ABoxes of s, r , and $inst$. We use $GET-CC(\mathcal{C})$ to extract the set of conditional commitments contained in \mathcal{C} .

According to the literature, commitments are manipulated either explicitly via specific events (such as a commitment cancellation or delegation), or implicitly when the commitment precondition or discharge condition becomes true. The allowed commitment manipulations, together with the resulting commitment states, are captured by means of a *commitment machine* [12]. In this work, we consider a simple commitment machine, inspired by [12,13], and show how to lift it to a first-order setting, taking into account that in our framework the precondition and the discharge condition are specified through queries over the data of the involved agents. More elaborated commitment machines, in terms of events and states, can be seamlessly incorporated.

Specifically, every commitment in $GET-CC(\mathcal{C})$ is associated to a specific first-order commitment machine, which is activated using the corresponding commitment rule in \mathcal{C} of the form above, instantiated possibly multiple times, depending on the agent data. The machine evolves as follows:

1. When an event of type EV is sent by agent a to agent b with payload \mathbf{d} , if $Q_c(a, b, \mathbf{d})$ is satisfied, an *instance of the conditional commitment* n is created. The debtor, creditor, and payload of this instance are respectively a, b , and \mathbf{d} .
2. Such instance is explicitly or implicitly manipulated by the involved agents. Explicit manipulation is done via specific message exchanges; we consider in particular the case of delegation from the debtor a to a new debtor, and the case of cancellation. Implicitly, instead, the instance can generate one or more corresponding base-level commitment instances: whenever $[q_p^+(a, b, \mathbf{d}, \mathbf{v})] \wedge Q_p^-(a, b, \mathbf{d}, \mathbf{v})$ is satisfied with actual values \mathbf{v} for variables \mathbf{y} , the conditional commitment instance *creates a base-level commitment instance* with payload \mathbf{d} and \mathbf{v} . Such base-level instance is put into the *active* state. The discharge condition for this instance is the instantiation of Q_d with the involved agents and specific payload, i.e., a is committed to bring about $Q_d(a, b, \mathbf{d}, \mathbf{v})$.
3. Also a base-level commitment instance is explicitly and implicitly manipulated by the involved agents. Explicit manipulation of an active base-level instance resembles that of conditional commitment instances, with the difference that, when canceled, a base-level commitment instance enters into the *violated* state. Implicit manipulation determines instead the *discharge* of the instance as soon as $Q_d(a, b, \mathbf{d}, \mathbf{v})$ holds, moving the instance from active to *satisfied*.

Example 4. Consider a commitment rule establishing a conditional commitment that the seller takes whenever it accepts the registration of a customer c . The conditional commitment is about the delivery of items paid by c . Specifically, for each item sold by the seller, if c has paid that item, then the seller commits to ensure that c will hold *that*

item. Note that the two conditions are *correlated* by the same item, and that a base-level commitment is created for each paid item. This cannot be expressed in propositional logic. Assuming that the seller stores a fact $\text{Paid}(i, c)$ if c has paid for i , and that the customer stores a fact $\text{Owned}(i)$ whenever it owns i , the commitment rule can be specified as:

on ACCEPT_REG **from** s **to** c **if** $\text{MyCust}@_s(c)$
then $\text{CC}_{\text{Delivery}}(s, c, [\text{Item}@_s(i) \wedge \text{Paid}@_s(i, c)], \text{Owns}@_c(i))$

Note the use of location arguments, reflecting that payments are maintained by the seller, whereas the items owned by the customer are maintained by the customer itself. ■

3.5 Commitment Box

The commitment box \mathcal{B} is a set of relations used by *inst* to maintain the concrete instances of conditional commitments, and the instances of their corresponding base-level commitments (with their states). In fact, due to the presence of data, commitments do not only require to keep track of the involved agents, but also of the payload associated to each of their instances. Such relations are extracted from the contractual specification as follows. Each commitment $\text{CC}_n(s, r, [q_p^+(s, r, \mathbf{x}, \mathbf{y})] \wedge Q_p^-(s, r, \mathbf{x}, \mathbf{y}), Q_d(s, r, \mathbf{x}, \mathbf{y}))$ in $\text{GET-CC}(\mathcal{C})$ induces two relations in \mathcal{B} , on the basis of the commitment name n and the payloads \mathbf{x} and \mathbf{y} : (i) $n\text{CC}/ar$, where $ar = 2 + |\mathbf{x}|$ for debtor, creditor, and conditional commitment payload; (ii) $n\text{C}/ar$, where $ar = 3 + |\mathbf{x}| + |\mathbf{y}|$ for debtor, creditor, state, and base-level commitment payload.

Example 5. The commitment in Example 4 induces the following relations in \mathcal{B} : $\text{DeliveryCC}(\text{debtor}, \text{creditor})$ and $\text{DeliveryC}(\text{debtor}, \text{creditor}, \text{state}, \text{item})$. ■

3.6 Commitment Machine Formalization

As anticipated in Section 3.3, the specification of *inst* must be complemented with a set of additional on-exchange rules, used to properly manipulate the evolution of commitments as the interaction unfolds. Commitment instances are stored by *inst* using the vocabulary of the CBox \mathcal{B} , and evolved through the application of these rules. Specifically, these rules ground the (first-order) commitment machine described in Section 3.4 to each specific commitment of $\text{GET-CC}(\mathcal{C})$, according to the “templates” described in the remainder of this section. We denote with $\text{CC-RULES}(\mathcal{C})$ all the commitment manipulation rules produced from \mathcal{C} .

When discussing the templates, we refer to a commitment rule $\rho \in \mathcal{C}$ of the form (*) in Section 3.4. Notice that, when n , \mathbf{x} and \mathbf{y} are mentioned in the rule templates, they are meant to be replaced with the actual commitment name and payload variables.

CC creation. For each $\rho \in \mathcal{C}$, a corresponding creation rule is obtained, depending on n and \mathbf{x} . When the rule triggers, a new instance of the conditional commitment $n\text{CC}$ is created, with the actual agents and payload:

on $\text{EV}(\mathbf{x})$ **from** s **to** r **if** $Q_c(s, r, \mathbf{x})$ **then** $\text{create_nCC}(s, r, \mathbf{x})$
 $\text{create_nCC}(s, r, \mathbf{x}) : \{ [\text{true}] \rightsquigarrow \text{add}\{n\text{CC}(s, r, \mathbf{x})\} \}$

CC delegation. The delegation of a conditional commitment instance for commitment n is triggered when the old debtor d_o sends to the new debtor d_n a DELEGATE_nCC

event, specifying in the event payload the creditor and the payload of the instance to be delegated. If such an instance exists, the debtor is updated by inst:

on DELEGATE_nCC(c, \mathbf{x}) **from** d_o **to** d_n **if** nCC(d_o, c, \mathbf{x}) **then** *changedeb_nCC*(d_o, d_n, c, \mathbf{x})
changedeb_nCC(d_o, d_n, c, \mathbf{x}) : { [true] \rightsquigarrow **add**{nCC(d_n, c, \mathbf{x})}, **del**{nCC(d_o, c, \mathbf{x})} }

CC cancelation. The cancelation of a conditional commitment instance for commitment n is triggered when the debtor sends to the creditor a CANCEL_nCC event, providing the instance payload. If the instance exists, it is removed:

on CANCEL_CC(\mathbf{x}) **from** d **to** c **if** nCC(d, c, \mathbf{x}) **then** *delete_nCC*(d, c, \mathbf{x})
delete_nCC(d, c, \mathbf{x}) : { [true] \rightsquigarrow **del**{nCC(d, c, \mathbf{x})} }

C creation. Every conditional commitment instance for relation nCC creates a base-level commitment instance whenever the precondition (whose variables \mathbf{x} are grounded with the instance payload) holds with an answer substitution θ for variables \mathbf{y} . This results in the creation of a new tuple for relation nC with the actual, full payload. This does not depend on the specific exchanged event, but only on the actual configuration of the data. Hence, a single “any-event” rule can be used to manage the creation of all base-level instances at once:

on any event from d **to** c **if** true **then** *createC*(d, c)

where, for each commitment $CC_n(s, r, [q_p^+(s, r, \mathbf{x}, \mathbf{y})] \wedge Q_p^-(s, r, \mathbf{x}, \mathbf{y}), Q_d(s, r, \mathbf{x}, \mathbf{y}))$ in GET-CC(\mathcal{C}), action *createC*(d, c) contains the following detachment effect:

$[nCC(d, c, \mathbf{x}) \wedge q_p^+(d, c, \mathbf{x}, \mathbf{y})] \wedge Q_p^-(d, c, \mathbf{x}, \mathbf{y}) \rightsquigarrow$ **add** {nC($d, c, \text{active}, \mathbf{x}, \mathbf{y}$)}

Differently from the propositional formalization of a commitment machine, in which the conditional commitment detaches to a base-level one, in our setting the conditional commitment instance is maintained, and keeps waiting for other situations matching the precondition with different data.

C delegation. It resembles the CC delegation:

on DELEGATE_nC($c, \mathbf{x}, \mathbf{y}$) **from** d_o **to** d_n
if nC($d_o, c, \text{active}, \mathbf{x}, \mathbf{y}$) **then** *changedeb_nC*($d_o, d_n, c, \mathbf{x}, \mathbf{y}$)
changedeb_nC($d_o, d_n, c, \mathbf{x}, \mathbf{y}$) :
{ [true] \rightsquigarrow **add**{nC($d_n, c, \text{active}, \mathbf{x}, \mathbf{y}$)}, **del**{nC($d_o, c, \text{active}, \mathbf{x}, \mathbf{y}$)} }

C cancelation. It determines a transition for the base-level commitment instance from the active to the violated state:

on CANCEL_C(\mathbf{x}, \mathbf{y}) **from** d **to** c **if** nC($d, c, \mathbf{x}, \mathbf{y}$) **then** *viol_nC*($d, c, \mathbf{x}, \mathbf{y}$)
viol_nC($d, c, \mathbf{x}, \mathbf{y}$) : { [true] \rightsquigarrow **add**{nC($d, c, \text{viol}, \mathbf{x}, \mathbf{y}$)}, **del**{nC($d, c, \text{active}, \mathbf{x}, \mathbf{y}$)} }

C discharge. Similarly to the case of C creation, the discharge of base-level commitment instances is handled by a single “any-event” rule, which checks the discharge condition

for each active commitment instance with the actual payload, evolving the instance to the satisfied state if it holds:

on any event from d to c if $true$ then $dischargeC(d, c)$

where, for each $CC_n(s, r, [q_p^+(s, r, \mathbf{x}, \mathbf{y})] \wedge Q_p^-(s, r, \mathbf{x}, \mathbf{y}), Q_d(s, r, \mathbf{x}, \mathbf{y}))$ in $CC(\mathcal{C})$, action $dischargeC(d, c)$ contains:

$$[nC(d, c, \text{active}, \mathbf{x}, \mathbf{y})] \wedge Q_d(d, c, \mathbf{x}, \mathbf{y}) \\ \rightsquigarrow \mathbf{add}\{nC(d, c, \text{sat}, \mathbf{x}, \mathbf{y})\}, \mathbf{del}\{nC(d, c, \text{active}, \mathbf{x}, \mathbf{y})\}$$

C removal. A last “any-event” reactive rule is used by *inst* to remove those instances of base-level commitments that already achieved a final state (sat or viol):

on any event from a to b if $true$ then $removeFinal()$

where, for each base-level commitment relation nC in \mathcal{B} , action $removeFinal()$ contains:

$$[nC(d, c, s, \mathbf{x}, \mathbf{y})] \wedge (s = \text{sat} \vee s = \text{viol}) \rightsquigarrow \mathbf{del}\{nC(d, c, s, \mathbf{x}, \mathbf{y})\}$$

Example 6. Assume that the only rule in \mathcal{C} is that of Example 4. The following CC creation rule is produced

on ACCEPT_REG from s to c if $MyCust@s(c)$ then $create_DeliveryCC(s, c)$
 $create_DeliveryCC(s, c) : \{[true] \rightsquigarrow \mathbf{add}\{DeliveryCC(s, c)\}\}$

Furthermore, the following C creation and C discharge update actions are produced:

$$createC(d, c) : \{[DeliveryCC(d, c) \wedge Item@d(i) \wedge Paid@d(i, c)] \\ \rightsquigarrow \mathbf{add}\{DeliveryC(d, c, \text{active}, i)\}\}$$

$$dischargeC(d, c) : \{[DeliveryC(d, c, \text{active}, i)] \wedge Owns@c(i) \\ \rightsquigarrow \mathbf{add}\{DeliveryC(d, c, \text{sat}, i)\}, \mathbf{del}\{DeliveryC(d, c, \text{active}, i)\}\}$$

4 Execution Semantics

The execution semantics of a DACMAS is defined in terms of a *transition system* that, starting from a given initial state, accounts for all the possible system dynamics, considering all the (possibly infinite) sequences of message exchanges, and all the possible substitutions that the agents choose during the application of update actions to provide concrete values for the Skolem terms. Given a DACMAS $\mathcal{S} = \langle \mathcal{T}, \mathcal{E}, \mathcal{X}, \mathcal{I}, \mathcal{C}, \mathcal{B} \rangle$ and an initial state σ_0 , the execution semantics of \mathcal{S} over σ_0 is defined by a transition system $\mathcal{T}_{\mathcal{S}}^{\sigma_0} = \langle \Delta, \mathcal{T} \cup \mathcal{B}, \Sigma, \sigma_0, \Rightarrow \rangle$, where:

- Σ is a (possibly infinite) set of states. Each state $\sigma \in \Sigma$ is equipped with a function *abox* that, given the name *a* of an agent, returns the ABox $\sigma.abox(a)$ of *a* in σ , if and only if *a* participates to the system in state σ , i.e., *a* belongs to the extension Agent in $\sigma.abox(inst)$. Hence, $\sigma.abox(inst)$ is always defined.
- $\sigma_0 \in \Sigma$ is the initial state. We assume that every ABox \mathcal{A} in σ_0 is such that $(\mathcal{T}, \mathcal{A})$ is satisfiable, and that $Spec(sn) \in \sigma_0.abox(inst)$ if and only if $\langle sn, - \rangle \in \mathcal{X}$.
- $\Rightarrow \subseteq \Sigma \times \Sigma$ is a transition relation.

Instrumental to the definition of the transition system is the extension of answering ECQ_ℓ queries so as to take into account location arguments. Formally, given a TBox \mathcal{T} , we define that $R@b(\mathbf{x})$ holds in state σ from the perspective of agent a under substitution θ for \mathbf{x} , written $\mathcal{T}, \sigma, a, \theta \models R@b(\mathbf{x})$, if:⁴

$$\begin{cases} \sigma.abox(a) \text{ is defined and } (\mathcal{T}, \sigma.abox(a)) \models R(\mathbf{x})\theta, & \text{if } b\theta = \text{self} \\ \sigma.abox(b\theta) \text{ is defined and } (\mathcal{T}, \sigma.abox(b\theta)) \models R(\mathbf{x})\theta, & \text{if } b\theta \neq \text{self} \end{cases}$$

Note that the semantics supports a sort of *dynamic binding* of location arguments, using θ to substitute a variable location argument with an agent name. This relation extends in the natural way to UCQ_ℓ and ECQ_ℓ , considering that quantification ranges over the active domain $\text{ADOM}(\sigma)$ of σ , which is defined as the union of the active domains of the ABoxes maintained by the agents present in σ . This, in turn, allows us to define the *certain answers to Q obtained by agent a in state σ* , denoted $\text{ANS}_\ell(Q, \mathcal{T}, \sigma, a)$, as the set of substitutions θ for the free variables in Q such that Q holds in state σ from the perspective of a , i.e., $\text{ANS}_\ell(Q, \mathcal{T}, \sigma, a) = \{\theta \mid \mathcal{T}, \sigma, a, \theta \models Q\}$.

The construction of the transition system $\mathcal{T}_S^{\sigma_0}$ is given in Figure 1.

5 Verification of DACMAS

To specify dynamic properties over DACMASs, we use a first-order variant of μ -calculus [14,6]. μ -calculus is virtually the most powerful temporal logic used for model checking of finite-state transition systems, and is able to express both linear time logics such as LTL and PSL, and branching time logics such as CTL and CTL* [9]. In our variant of μ -calculus, local properties are expressed as ECQ_ℓ queries over the current state of the DACMAS. We allow for a controlled form of first-order quantification across states, inspired by [2], where the quantification ranges over data items across time only as long as such items persist in the active domain. Formally, we define the logic $\mu\mathcal{L}_p^{ECQ_\ell}$ as:

$$\Phi ::= Q_\ell \mid \neg\Phi \mid \Phi_1 \wedge \Phi_2 \mid \exists x. \text{LIVE}(x) \wedge \Phi \mid \text{LIVE}(x) \wedge \langle \neg \rangle \Phi \mid \text{LIVE}(x) \wedge [\neg] \Phi \mid Z \mid \mu Z. \Phi$$

where Q is a (possibly open) ECQ_ℓ query, in which the only constants that may appear are those in the initial state of the system, Z is a second order predicate variable (of arity 0), and $\text{LIVE}(x_1, \dots, x_n)$ abbreviates $\bigwedge_{i \in \{1, \dots, n\}} \text{LIVE}(x_i)$. For $\mu\mathcal{L}_p^{ECQ_\ell}$, the following assumption holds: in $\text{LIVE}(x) \wedge \langle \neg \rangle \Phi$ and $\text{LIVE}(x) \wedge [\neg] \Phi$, the variables \mathbf{x} are exactly the free variables of Φ , once we substitute to each bounded predicate variable Z in Φ its bounding formula $\mu Z. \Phi'$. We adopt the usual abbreviations, including $\nu Z. \Phi$ for greatest fixpoints. Intuitively, the use of $\text{LIVE}(\cdot)$ in $\mu\mathcal{L}_p^{ECQ_\ell}$ ensures that data items are only considered if they persist along the system evolution, while the evaluation of a formula with data that are not present in the current state trivially leads to false or true. This is in line with DACMASs, where the evolution of a commitment instance persists until the commitment is discharged or canceled, and where an agent name is meaningful only while it persists in the system: when an agent leaves the system and its name is canceled by *inst*, *inst* could reuse a in the future to identify another agent.

The formula $\mu Z. \Phi$ denotes the least fixpoint of the formula Φ . As usual in μ -calculus, formulae of the form $\mu Z. \Phi$ must obey to the *syntactic monotonicity* of Φ w.r.t. Z , which

⁴ We assume that θ is the identity on data items (including the special constants *self* and *inst*).

```

1: procedure BUILD-TS
2: input: DACMAS  $\mathcal{S} = \langle \mathcal{T}, \mathcal{E}, \mathcal{X}, \mathcal{I}, \mathcal{C}, \mathcal{B} \rangle$  and initial state  $\sigma_0$ 
3: output: Transition system  $\langle \Delta, \mathcal{T} \cup \mathcal{B}, \Sigma, \sigma_0, \Rightarrow \rangle$ 
4:    $\Sigma := \{\sigma_0\}, \Rightarrow := \emptyset$ 
5:   while true do
6:     pick  $\sigma \in \Sigma$  and  $a \in \{\text{ag} \mid \text{Agent}(\text{ag}) \in \sigma.\text{abox}(\text{inst})\}$ 
7:     Fetch all current behavioural rules for  $a$ ; Calculate enabled events with receivers for  $a$ 
8:     if There exists at least an enabled event then
9:       pick an enabled event  $\text{EV}(\mathbf{e})$  for  $a$  with receiver  $b$ 
10:       $\mathcal{A}_i := \text{APPLY}(\mathcal{S}, \sigma, \text{inst}, a, b, \text{EV}(\mathbf{e}))$  ▷ New inst ABox
11:       $\Sigma := \Sigma \cup \{\sigma'\}$  ▷ Tentatively add a new state  $\sigma'$ 
12:      for all  $x \in \{\text{ag} \mid \text{Agent}(\text{ag}) \in \mathcal{A}_i\}$  do
13:         $\sigma'.\text{abox}(x) := \text{APPLY}(\mathcal{S}, \sigma, x, a, b, \text{EV}(\mathbf{e}))$ 
14:        if for every  $x \in \{\text{ag} \mid \text{Agent}(\text{ag}) \in \mathcal{A}_i\}, \langle \mathcal{T}, \sigma'.\text{abox}(x) \rangle$  is satisfiable then
15:           $\Rightarrow := \Rightarrow \cup \langle \sigma, \sigma' \rangle$ 
16:        else  $\Sigma := \Sigma \setminus \{\sigma'\}$  ▷ Inconsistent execution step
17:   function APPLY( $\mathcal{S}, \sigma, x, a, b, \text{EV}(\mathbf{e})$ )
18:   output: new ABox for  $x$  after reacting to  $\text{EV}(\mathbf{e})$  from  $a$  to  $b$ 
19:   if  $x \notin \{\text{inst}, a, b\}$  then return  $\sigma.\text{abox}(x)$ 
20:   Fetch all current behavioural rules for  $x$ 
21:   if  $x = a$  then ▷  $x$  is the sender agent
22:     Fetch on-send and “self” on-receive rules and compute actions with actual params
23:   if  $x = b$  then ▷  $x$  is the receiver agent
24:     Fetch on-receive and “self” on-send rules and compute actions with actual param
25:   if  $x = \text{inst}$  then ▷  $x$  is the institutional agent
26:     Fetch matching/“any-event” on-exchange rules and compute actions with actual param
27:    $\text{TOADD} := \emptyset, \text{TODEL} := \emptyset$ 
28:   for all  $\alpha(\mathbf{v}) \in \text{ACT}$  do ▷  $\text{ACT}$  = set of fetched actions
29:      $\text{TOADDSK} := \emptyset$ 
30:     for all effect “ $[q^+(\mathbf{p}, \mathbf{x})] \wedge Q^-(\mathbf{p}, \mathbf{x}) \rightsquigarrow \text{add } A, \text{del } D$ ” in the definition of  $\alpha$  do
31:       for all  $\theta \in \text{ANS}_\ell([q^+(\mathbf{v}, \mathbf{x})] \wedge Q^-(\mathbf{v}, \mathbf{x}), \mathcal{T}, \sigma, x)$  do
32:          $\text{TOADDSK} := \text{TOADDSK} \cup A\theta[\mathbf{p}/\mathbf{v}]$ 
33:          $\text{TODEL} := \text{TODEL} \cup D\theta[\mathbf{p}/\mathbf{v}]$ 
34:       pick a substitution  $\theta_{sk}$  of the Skolem terms with data
35:        $\text{TOADD} := \text{TOADD} \cup \text{TOADDSK}\theta_{sk}$ 
36:       if  $x = \text{inst}$  then  $\text{TOADD} := \text{TOADD} \cup \{\text{Agent}(\text{inst})\}$ 
37:   return  $(\sigma.\text{abox}(x) \setminus \text{TODEL}) \cup \text{TOADD}$ 

```

Fig. 1. Transition system construction

states that every occurrence of the variable Z in Φ must be within the scope of an even number of negation symbols. This ensures that the least fixpoint $\mu Z.\Phi$ always exists.

The semantics of $\mu\mathcal{L}_p^{ECQ_\ell}$ formulae is defined over a possibly infinite transition system $\mathcal{T} = \langle \Delta, \mathcal{T} \cup \mathcal{B}, \Sigma, \sigma_0, \Rightarrow \rangle$ (cf. Section 4), assuming that ECQ_ℓ queries are posed from the point of view of inst . This does not prevent the possibility to query the ABoxes of the other agents, thanks to the dynamic binding for location arguments. Since $\mu\mathcal{L}_p^{ECQ_\ell}$ contains formulae with both individual and predicate free variables, we introduce an *individual variable valuation* v , i.e., a mapping from individual variables x to Δ , and a *predicate variable valuation* V , i.e., a mapping from the predicate variables Z to subsets of Σ . With these three notions in place, we assign meaning to formulae by

$$\begin{aligned}
(Q\ell)_{v,V}^{\mathcal{Y}} &= \{\sigma \in \Sigma \mid \mathcal{T}, \sigma, \text{inst}, v \models Q\ell\} \\
(\neg\Phi)_{v,V}^{\mathcal{Y}} &= \Sigma \setminus (\Phi)_{v,V}^{\mathcal{Y}} \\
(\Phi_1 \wedge \Phi_2)_{v,V}^{\mathcal{Y}} &= (\Phi_1)_{v,V}^{\mathcal{Y}} \cap (\Phi_2)_{v,V}^{\mathcal{Y}} \\
(\exists x.\text{LIVE}(x) \wedge \Phi)_{v,V}^{\mathcal{Y}} &= \{\sigma \in \Sigma \mid \exists d \in \text{ADOM}(\sigma). \sigma \in (\Phi)_{v[x/d],V}^{\mathcal{Y}}\} \\
(\text{LIVE}(\mathbf{x}) \wedge \langle \neg \rangle \Phi)_{v,V}^{\mathcal{Y}} &= \{\sigma \in \Sigma \mid \mathbf{x}/\mathbf{d} \in v \text{ implies } \mathbf{d} \subseteq \text{ADOM}(\sigma) \\
&\quad \text{and } \exists \sigma'. \sigma \Rightarrow \sigma' \text{ and } \sigma' \in (\Phi)_{v,V}^{\mathcal{Y}}\} \\
(\text{LIVE}(\mathbf{x}) \wedge [\neg] \Phi)_{v,V}^{\mathcal{Y}} &= \{\sigma \in \Sigma \mid \mathbf{x}/\mathbf{d} \in v \text{ implies } \mathbf{d} \subseteq \text{ADOM}(\sigma) \\
&\quad \text{and } \forall \sigma'. \sigma \Rightarrow \sigma' \text{ implies } \sigma' \in (\Phi)_{v,V}^{\mathcal{Y}}\} \\
(Z)_{v,V}^{\mathcal{Y}} &= V(Z) \\
(\mu Z.\Phi)_{v,V}^{\mathcal{Y}} &= \bigcap \{\mathcal{E} \subseteq \Sigma \mid (\Phi)_{v,V[Z/\mathcal{E}]}^{\mathcal{Y}} \subseteq \mathcal{E}\}
\end{aligned}$$

Fig. 2. Semantics of $\mu\mathcal{L}_p^{ECQ\ell}$

associating to \mathcal{Y} , v , and V an *extension function* $(\cdot)_{v,V}^{\mathcal{Y}}$, which maps formulae to subsets of Σ . Formally, the extension function $(\cdot)_{v,V}^{\mathcal{Y}}$ is defined inductively as shown in Figure 2. When Φ is a closed formula, $(\Phi)_{v,V}^{\mathcal{Y}}$ does not depend on v or V , and we denote the extension of Φ simply by $(\Phi)^{\mathcal{Y}}$. A closed formula Φ holds in a state $s \in \Sigma$ if $s \in (\Phi)^{\mathcal{Y}}$. In this case, we write $\mathcal{Y}, s \models \Phi$. Given DACMAS \mathcal{S} , an initial state σ_0 and a $\mu\mathcal{L}_p^{ECQ\ell}$ formula Φ , we are interested in the following *verification* problem: $\mathcal{Y}_{\mathcal{S}^{\sigma_0}}, \sigma_0 \models \Phi$.

Example 7. Consider the contract of Example 4. Assume that \mathcal{T} contains that gold customers are seller customers: $\text{MyGoldCust} \sqsubseteq \text{MyCust}$. The $\mu\mathcal{L}_p^{ECQ\ell}$ property

$$\begin{aligned}
&\nu Z. (\forall s, c, i. \text{DeliveryC}(s, c, \text{active}, i) \wedge \text{MyGoldCust}@s(c) \\
&\quad \rightarrow \mu Y. (\text{DeliveryC}(s, c, \text{sat}, i)) \vee (\text{LIVE}(s, c, i) \wedge \langle \neg \rangle Y)) \wedge [\neg] Z
\end{aligned}$$

models that, for every delivery commitment instance a seller has towards a gold customer, there must exist a run where the instance persists in the system until it is satisfied. ■

The number of states of $\mathcal{Y}_{\mathcal{S}^{\sigma_0}}$ is in general infinite, and verification of (even propositional) temporal properties of simple forms (e.g., reachability) turns out to be undecidable [2,6]. This calls for identifying interesting classes of DACMASs for which verification is decidable. Recently, the notion of *state-bounded* system has been proposed in the context of both data-aware business processes [2] and MASs [3], as an interesting condition that ensures decidability of verification for rich first-order temporal properties, while reflecting naturally occurring working assumptions in real-world systems. Intuitively, state-boundedness allows for encountering infinitely many different data during the evolution of the system, provided that such data do not accumulate in a single state.

We take this general notion and adapt it to DACMASs. In particular, a DACMAS is *state-bounded* if, for every agent active in the system, there exists a bound on the number of data items simultaneously stored in its ABox. Since the ABox of inst stores the names of the active agents, this implicitly bounds also the number of simultaneously active agents. Observe, however, that the overall number of data items (and hence also agents) encountered across and along the runs of the system can still be infinite. With this notion in place, we obtain the following fundamental result:

Theorem 1 ([10]). *Verifying state-bounded DACMASs against $\mu\mathcal{L}_p^{ECQ\ell}$ properties is decidable and reducible to finite-state model checking.*

This means that state-bounded DACMASs can be verified, in principle, using standard model checkers for propositional μ -calculus.

6 Conclusion

DACMASs are readily implementable in standard technologies such as JADE (which supports dynamic agent creation) and lightweight ontologies. Observe that a system execution requires polynomial time at each step (actually logspace w.r.t. the data, as any system based on relational databases). Only offline verification of the system is (as usual) exponential in the representation. Our framework complements that of [7], which employs data-aware commitments to monitor a system execution and track the state of commitment instances, but cannot be exploited for static analysis.

We consider extending our framework with the possibility of checking epistemic properties, in the line of [3]. Notice that, if instead of relying on the μ -calculus, we rely on CTL, we can relax the persistence requirement in the logic, as in [3]. We also intend to study how to derive skeletons for the local agent specifications from a global, choreographic commitment-based protocol.

References

1. Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P.F. (eds.): The Description Logic Handbook.: Cambridge University Press (2003)
2. Bagheri Hariri, B., Calvanese, D., De Giacomo, G., Deutsch, A., Montali, M.: Verification of relational data-centric dynamic systems with external services. In: Proc. of PODS (2013)
3. Belardinelli, F., Lomuscio, A., Patrizi, F.: An abstraction technique for the verification of artifact-centric systems. In: Proc. of KR (2012)
4. Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M., Rosati, R.: EQL-Lite: Effective first-order query processing in description logics. In: Proc. of IJCAI (2007)
5. Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M., Rosati, R.: Data complexity of query answering in description logics. *Artificial Intelligence* 195, 335–360 (2013)
6. Calvanese, D., De Giacomo, G., Montali, M., Patrizi, F.: Verification and synthesis in description logic based dynamic systems. In: Proc. of RR (2013)
7. Chesani, F., Mello, P., Montali, M., Torroni, P.: Representing and monitoring social commitments using the event calculus. *AutonAgent and Multi-Agent Syst.* 27(1), 85–130 (2013)
8. Chopra, A.K., Singh, M.P.: *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, chap. Agent Communication. MIT Press (2013)
9. Clarke, E.M., Grumberg, O., Peled, D.A.: *Model checking*. MIT Press (1999)
10. Montali, M., Calvanese, D., De Giacomo, G.: Verification of data-aware commitment-based multiagent systems. In: Proc. of AAMAS (2014), to appear
11. Poggi, A., Lembo, D., Calvanese, D., De Giacomo, G., Lenzerini, M., Rosati, R.: Linking data to ontologies. *J. on Data Semantics X*, 133–173 (2008)
12. Singh, M.P.: Formalizing communication protocols for multiagent systems. In: Proc. of IJCAI (2007)
13. Singh, M.P., Chopra, A.K., Desai, N.: Commitment-based service-oriented architecture. *IEEE Computer* 42(11), 72–79 (2009)
14. Stirling, C.: *Modal and Temporal Properties of Processes*. Springer (2001)