

Verification of Semantically-Enhanced Artifact Systems*

Extended Version

Babak Bagheri Hariri, Diego Calvanese, Marco Montali,
Ario Santoso, and Dmitry Solomakhin

KRDB Research Centre for Knowledge and Data, Free University of Bozen-Bolzano
lastname@inf.unibz.it

Abstract. Artifact-Centric systems have emerged in the last years as a suitable framework to model business-relevant entities, by combining their static and dynamic aspects. In particular, the Guard-Stage-Milestone (GSM) approach has been recently proposed to model artifacts and their lifecycle in a declarative way. In this paper, we enhance GSM with a Semantic Layer, constituted by a full-fledged OWL 2 QL ontology linked to the artifact information models through mapping specifications. The ontology provides a conceptual view of the domain under study, and allows one to understand the evolution of the artifact system at a higher level of abstraction. In this setting, we present a technique to specify temporal properties expressed over the Semantic Layer, and verify them according to the evolution in the underlying GSM model. This technique has been implemented in a tool that exploits state-of-the-art ontology-based data access technologies to manipulate the temporal properties according to the ontology and the mappings, and that relies on the GSMC model checker for verification.

1 Introduction

In the last decade, the marriage between processes and data has been increasingly advocated as a key objective towards a comprehensive modeling and management of complex enterprises [10]. This requires to go beyond classical (business) process specification languages, which largely leave the connection between the process dimension and the data dimension underspecified, and to consider data and processes as “two sides of the same coin” [23]. In this respect, artifact-centric systems [21,18] have lately emerged as an effective framework to model business-relevant entities, by combining in a holistic way their static and dynamic aspects. Artifacts are characterized by an “information model”, which maintains the artifact data, and by a lifecycle that specifies the allowed ways to progress the information model. Among the different proposals for artifact-centric process modelling, the Guard-Stage-Milestone (GSM) approach has been recently proposed to model artifacts and their lifecycle in a declarative, flexible way [19]. GSM is equipped with a formal execution semantics [14], which unambiguously characterizes the artifact progression in response to external events. Notably, several key constructs

* This research has been partially supported by the EU under the ICT Collaborative Project ACSI (Artifact-Centric Service Interoperation), grant agreement n. FP7-257593.

of the emerging OMG standard on Case Management and Model Notation ¹ have been borrowed from GSM.

Despite the tight integration between the data and process component, the artifact information model typically relies on relatively simple structures, such as (nested) lists of key-value pairs. This causes an abstraction gap between the high-level, conceptual view that business stakeholders have of domain-relevant entities and relations, and the low-level representation adopted inside artifacts. To overcome this problem, in [9] it is proposed to enhance artifact systems with a Semantic Layer, constituted by a full-fledged ontology linked to the artifact information models through mapping specifications. On the one hand, the ontology provides a conceptual view of the domain under study, and allows one to understand the evolution of the artifact system at a higher level of abstraction. On the other hand, mapping specifications allow one to connect the elements present in the Semantic Layer with the concrete data maintained in the artifact information models, relying on an approach recently proposed in the context of Ontology-Based Data Access (OBDA) [6].

To specify the ontology, we adopt the OWL 2 QL profile [20] of the standard Web Ontology Language (OWL) [3], since it supports efficient access to data while fully taking into account the constraints expressed in the ontology. Specifically, OWL 2 QL enjoys so-called *first-order rewritability* of query answering [8], which guarantees that conjunctive queries posed over the ontology can be rewritten into first-order queries that incorporate the ontological constraints, and thus do not require further inference for query answering. A rewritten query can then be *unfolded* with respect to the mappings, so as to obtain a query formulated over the underlying data that can be directly evaluated so as to obtain the so-called *certain answers* to the original query over the ontology [22].

We follow here an approach that is similar to the one proposed in [9], and that is based on specifying in terms of the Semantic Layer dynamic/temporal laws that the system should obey, and that need to be verified according to the evolution in the underlying artifact layer. However, differently from [9], in which the Semantic Layer is mainly used to *govern* the progression of artifacts, by forbidding the execution of actions that would lead to violate the constraints in the ontology, here we are primarily interested in exploiting the Semantic Layer to ease the specification of the dynamic/temporal laws. In this light, we extend the technique provided in [9] by relying on a more expressive verification formalism, which supports first-order epistemic queries embedded into an expressive temporal language, the first-order μ -calculus [16], while allowing for quantification across states. The latter makes it possible to predicate over the temporal evolution of individuals, and thus represents an enhancement that is fundamental for capturing many practical scenarios (cf. Section 6).

We then concretize this formal framework by relying on GSM as the artifact model, and by exploiting state of the art technologies, on the one hand for dealing with the ontology and the mappings, and on the other hand for performing verification. Specifically, we have developed the OBGSM tool (see Section 5), which is able to reformulate temporal properties expressed over the ontology in terms of the underlying GSM information model, and which adopts the state of the art OBDA system Quest [24] to efficiently rewrite and unfold the epistemic queries embedded in the temporal property to verify.

¹ <http://www.omg.org/spec/CMMN/>

For the actual verification task, we rely on the recently developed GSMC model checker for GSM [4], which is the only model checker currently available that is able to verify temporal formula over artifact systems. GSMC adopts as temporal formalism a variant of the first-order branching time logic CTL [12] with a restricted form of quantification across states². The restrictions imposed by GSMC made it necessary to suitably accommodate the mapping language and the temporal formalism over the Semantic Layer so as to ensure that the temporal formulas resulting from rewriting and unfolding can be effectively processed by GSMC.

2 Preliminaries

OWL 2 QL is a profile³ of the Web Ontology Language OWL 2 standardized by the W3C. OWL 2 QL is specifically designed for building an ontology layer to wrap possibly very large data sources. Technically, OWL 2 QL is based on the description logic $DL-Lite_{\mathcal{R}}$, which is a member of the $DL-Lite$ family [8], designed specifically for effective ontology-based data access, and which we adopt in the following.

In description logics (DLs) [1], the domain of interest is modeled by means of *concepts*, representing classes of objects, and *roles*, representing binary relations between objects⁴. In $DL-Lite_{\mathcal{R}}$, concepts C and roles U obey to the following syntax:

$$B ::= N \mid \exists U \quad C ::= B \mid \exists U.B \quad U ::= P \mid P^-$$

P denotes a *role name*, and P^- an *inverse role*, which swaps the first and second components of P . N denotes a *concept name*, and B a *basic concept*, which is either simply a concept name, or the projection of a role P on its first component ($\exists P$) or its second component ($\exists P^-$). In the concept $\exists U.B$, the projection on the first (resp., second) component of U can be further qualified by requiring that the second (resp., first) component of U is an instance of the basic concept B .

In DLs, domain knowledge is divided into an intensional part, called *TBox*, and an extensional part, called *ABox*. Specifically, a $DL-Lite_{\mathcal{R}}$ ontology is a pair (T, A) , where the TBox T is a finite set of (concept and role) *inclusion assertions* of the forms $B \sqsubseteq C$ and $U_1 \sqsubseteq U_2$, and of *disjointness assertions* of the form $\text{disjoint}(B_1, B_2)$ and $\text{disjoint}(U_1, U_2)$. Instead, the ABox A is a finite set of *facts* (also called membership assertions) of the forms $N(c_1)$ and $P(c_1, c_2)$, where N and P occur in T , and c_1 and c_2 are constants.

The semantics of a $DL-Lite_{\mathcal{R}}$ ontology is given in terms of first-order *interpretations* $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, where $\Delta^{\mathcal{I}}$ is the interpretation domain and $\cdot^{\mathcal{I}}$ is an interpretation function that assigns to each concept C a subset $C^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ and to each role U a binary relation $U^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$, capturing the intuitive meaning of the various constructs (see [8] for details). An interpretation that satisfies all assertions in T and A is called a *model* of the ontology (T, A) , and the ontology is said to be *satisfiable* if it admits at least one model.

² Note that CTL can be expressed in the alternation-free fragment of the μ -calculus [13,16].

³ In W3C terminology, a *profile* is a sublanguage defined by suitable syntactic restrictions.

⁴ For simplicity of presentation, we do not distinguish here between roles, corresponding to OWL 2 object properties, and attributes, corresponding to OWL 2 data properties (see [6]). But all our results would carry over unchanged.

Queries. As usual (cf. OWL 2 QL), answers to queries are formed by terms denoting individuals explicitly mentioned in the ABox. The *domain of an ABox* A , denoted by $\text{ADOM}(A)$, is the (finite) set of terms appearing in A . A *union of conjunctive queries* (UCQ) q over a KB (T, A) is a FOL formula of the form $\bigvee_{1 \leq i \leq n} \exists \vec{y}_i. \text{conj}_i(\vec{x}, \vec{y}_i)$ with free variables \vec{x} and existentially quantified variables $\vec{y}_1, \dots, \vec{y}_n$. Each $\text{conj}_i(\vec{x}, \vec{y}_i)$ in q is a conjunction of atoms of the form $N(z), P(z, z')$, where N and P respectively denote a concept and a role name occurring in T , and z, z' are constants in $\text{ADOM}(A)$ or variables in \vec{x} or \vec{y}_i , for some $i \in \{1, \dots, n\}$. The (*certain*) *answers* to q over (T, A) is the set $\text{ans}(q, T, A)$ of substitutions σ of the free variables of q with constants in $\text{ADOM}(A)$ such that $q\sigma$ evaluates to true in every model of (T, A) . If q has no free variables, then it is called *boolean* and its certain answers are either true or false.

We compose UCQs using ECQs, i.e., queries of the query language *EQL-Lite*(UCQ) [7], which is the FOL query language whose atoms are UCQs evaluated according to the certain answer semantics above. An *ECQ* over T and A is a possibly open formula of the form

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

where q is a UCQ. The *answer to Q over (T, A)* , is the set $\text{ANS}(Q, T, A)$ of tuples of constants in $\text{ADOM}(A)$ defined by composing the certain answers $\text{ans}(q, T, A)$ of UCQs q through first-order constructs, and interpreting existential variables as ranging over $\text{ADOM}(A)$.

Finally, we recall that *DL-Lite_R* enjoys the *FO rewritability* property, which states that for every UCQ q , $\text{ans}(q, T, A) = \text{ans}(\text{REW}(q), \emptyset, A)$, where $\text{REW}(q)$ is a UCQ computed by the reformulation algorithm in [6]. Notice that this algorithm can be extended to ECQs [7], and that its effect is to “compile away” the TBox. Similarly, ontology satisfiability is FO rewritable for *DL-Lite_R* TBoxes [6], which states that for every TBox T , there exists a boolean first-order query $q_{\text{unsat}}(T)$ such that for every non-empty ABox A , we have that (T, A) is satisfiable iff $\text{ans}(q_{\text{unsat}}(T), T, A) = \text{false}$.

Ontology-Based Data Access (OBDA). In an OBDA system, a relational database is connected to an ontology representing the domain of interest by a mapping, which relates database values with values and (abstract) objects in the ontology (cf. [6]). In particular, we make use of a countably infinite set \mathcal{V} of values and a set Λ of function symbols, each with an associated arity. We also define the set \mathcal{C} of constants as the union of \mathcal{V} and the set $\{f(d_1, \dots, d_n) \mid f \in \Lambda \text{ and } d_1, \dots, d_n \in \mathcal{V}\}$ of *object terms*.

Formally, an OBDA system is a structure $\mathcal{O} = \langle R, T, \mathcal{M} \rangle$, where: (i) $R = \{R_1, \dots, R_n\}$ is a database schema, constituted by a finite set of relation schemas; (ii) T is a *DL-Lite_R* TBox; (iii) \mathcal{M} is a set of mapping assertions, each of the form $\Phi(\vec{x}) \rightsquigarrow \Psi(\vec{y}, \vec{t})$, where: (a) \vec{x} is a non-empty set of variables, (b) $\vec{y} \subseteq \vec{x}$, (c) \vec{t} is a set of object terms of the form $f(\vec{z})$, with $f \in \Lambda$ and $\vec{z} \subseteq \vec{x}$, (d) $\Phi(\vec{x})$, which also called as *source query*, is an arbitrary SQL query over R , with \vec{x} as output variables, and (e) $\Psi(\vec{y}, \vec{t})$, which also called as *target query*, is a CQ over T of arity $n > 0$ without non-distinguished variables, whose atoms are over the variables \vec{y} and the object terms \vec{t} .

Given a database instance \mathcal{I} made up of values in \mathcal{V} and conforming to schema R , and given a mapping \mathcal{M} , the *virtual ABox* generated from \mathcal{I} by a mapping assertion $m = \Phi(\vec{x}) \rightsquigarrow \Psi(\vec{y}, \vec{t})$ in \mathcal{M} is $m(\mathcal{I}) = \bigcup_{v \in \text{eval}(\Phi, \mathcal{I})} \Psi[x/v]$, where $\text{eval}(\Phi, \mathcal{I})$ denotes

the evaluation of the SQL query $\bar{\Phi}$ over \mathcal{I} , and where we consider $\Psi[x/v]$ to be a set of atoms (as opposed to a conjunction). The ABox generated from \mathcal{I} by the mapping \mathcal{M} is $\mathcal{M}(\mathcal{I}) = \bigcup_{m \in \mathcal{M}} m(\mathcal{I})$. Notice that $\text{ADOM}(\mathcal{M}(\mathcal{I})) \subseteq \mathcal{C}$. As for ABoxes, the active domain $\text{ADOM}(\mathcal{I})$ of a database instance \mathcal{I} is the set of values occurring in \mathcal{I} . Notice that $\text{ADOM}(\mathcal{I}) \subseteq \mathcal{V}$. Given an OBDA system $\mathcal{O} = \langle R, T, \mathcal{M} \rangle$ and a database instance \mathcal{I} for R , a *model* for \mathcal{O} wrt \mathcal{I} is a model of the ontology $(T, \mathcal{M}(\mathcal{I}))$. We say that \mathcal{O} wrt \mathcal{I} is satisfiable if it admits a model wrt \mathcal{I} .

A UCQ q over an OBDA system $\mathcal{O} = \langle R, T, \mathcal{M} \rangle$ and a relational instance \mathcal{I} for R is simply an UCQ over $(T, \mathcal{M}(\mathcal{I}))$. To compute the certain answers of q over \mathcal{O} wrt \mathcal{I} , we follow the standard three-step approach [6]: (i) q is *rewritten* to compile away T , obtaining $q_r = \text{rew}(q, T)$; (ii) the mapping \mathcal{M} is used to *unfold* q_r into a query over R , denoted by $\text{UNFOLD}(q_r, \mathcal{M})$, which turns out to be an SQL query [22]; (iii) such a query is executed over \mathcal{I} , obtaining the certain answers. For an ECQ, we can proceed in a similar way, applying the rewriting and unfolding steps to the embedded UCQs. It follows that computing certain answers to UCQs/ECQs in an OBDA system is FO rewritable. Applying the unfolding step to $q_{\text{unsat}}(T)$, we obtain also that satisfiability in \mathcal{O} is FO rewritable.

3 Semantically-enhanced Artifact Systems

In this section we introduce Semantically-enhanced Artifact Systems (SASs), which generalize the framework of Semantically-Governed Data-Aware Processes (SGDAPs) presented in [9].

Intuitively, SAS models systems in which artifacts progress according to their lifecycles, and in which the evolution of the entire system is understood through the conceptual lens of an OWL 2 QL ontology. In accordance with the literature [15], it is assumed that artifacts are equipped with a relational information model. More specifically, a SAS is constituted by: (i) A *Relational Layer*, which account for the (relational) information models of the artifacts, and which employs a global transition relation to abstractly capture the step-by-step evolution of the system as a whole. (ii) A *Semantic Layer*, which contains an OWL 2 QL ontology that conceptually accounts for the domain under study. (iii) A set of *mapping assertions* describing how to virtually project data concretely maintained at the Relational Layer into concepts and relations modeled in the Semantic Layer, thus providing a link between the artifact information models and the ontology.

In the following, we assume a countably infinite set of values \mathcal{V} . Formally, a SAS \mathcal{S} is a tuple $\mathcal{S} = \langle R, \mathcal{I}_0, \mathcal{F}, T, \mathcal{M} \rangle$, where: (i) R is a database schema that incorporates the schemas of all artifact information models present in the Relational Layer; (ii) \mathcal{I}_0 is a database instance made up of values in \mathcal{V} and conforming to R , which represents the initial state of the Relational Layer; (iii) $\mathcal{F} \subseteq \Gamma \times \Gamma$ is the transition relation that describes the overall progression mechanism of the Relational Layer, where Γ is the set of all instances made up of values in \mathcal{V} and conforming to R ; (iv) T is a *DL-Lite_R* TBox; (v) \mathcal{M} is a set of mapping assertions that connect R to T , following the approach described in Section 2. By observing that the triple $\langle R, T, \mathcal{M} \rangle$ constitutes, in fact, an OBDA system, we notice that \mathcal{S} can be considered as an OBDA system equipped with a

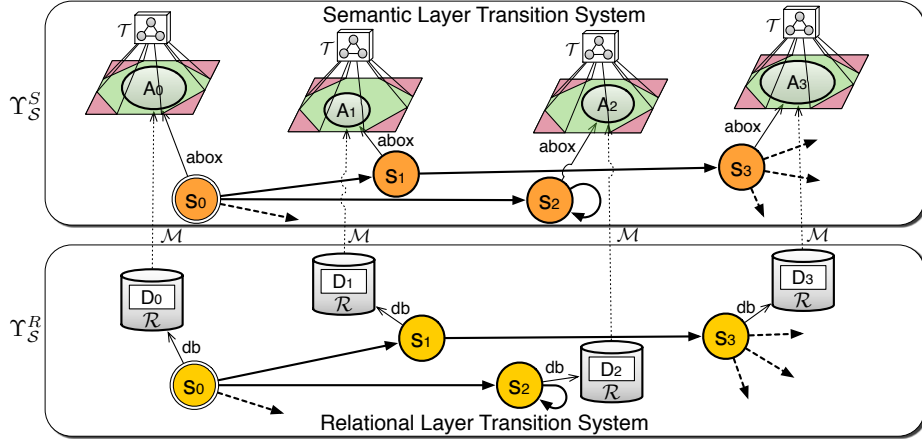


Fig. 1. SAS illustration

transition relation that accounts for the dynamics of the system at the level of R , starting from the initial state \mathcal{I}_0 .

3.1 Execution Semantics

The execution semantics of a SAS \mathcal{S} is provided by means of transition systems. While the temporal structure of such transition systems is fully determined by the transition relation of \mathcal{S} , the content of each state in the transition system depends on whether the dynamics of SASs are understood directly at the Relational Layer, or through the conceptual lens of the Semantic Layer ontology. In the former case, each state is associated to a database instance that represents the current snapshot of the artifact information models, whereas in the latter case each state is associated to an ABox that represents the current state of the system, as understood by the Semantic Layer.

Following this approach, the execution semantics of \mathcal{S} is captured in terms of two transition systems, one describing the allowed evolutions at the Relational Layer (*Relational Transition System - RTS*), and one abstracting them at the Semantic Layer (*Semantic Transition System - STS*). Figure 1 provides a graphical intuition about the RTS and STS, and their interrelations.

RTS. Given a SAS $\mathcal{S} = \langle R, \mathcal{I}_0, \mathcal{F}, T, \mathcal{M} \rangle$, its RTS Υ_S^R is defined as a tuple $\langle R, \Sigma, s_0, db, \Rightarrow \rangle$, where: (i) Σ is a set of states, (ii) $s_0 \in \Sigma$, (iii) db is a function that, given a state in Σ , returns a corresponding database instance (conforming to R), (iv) $\Rightarrow \subseteq \Sigma \times \Sigma$ is the transition relation. The components Σ, \Rightarrow and db of Υ_S^R are defined by simultaneous induction as the smallest sets satisfying the following conditions:

- $db(s_0) = \mathcal{I}_0$;
- for every databases instance \mathcal{I}' such that $\langle db(s), \mathcal{I}' \rangle \in \mathcal{F}$:
 - if there exists $s' \in \Sigma$ such that $db(s') = \mathcal{I}'$, then $s \Rightarrow s'$;
 - otherwise, if $\mathcal{O} = \langle R, T, \mathcal{M} \rangle$ is *satisfiable* wrt \mathcal{I}' , then $s' \in \Sigma$, $s \Rightarrow s'$ and $db(s') = \mathcal{I}'$, where s' is a fresh state.

The satisfiability check done in the last step of the RTS construction accounts for the semantic *governance* (cf. Section 1): a transition is preserved in the RTS only if the target state does not violate any ontological constraints of the Semantic Layer, otherwise it is rejected [9].

STS. Given a SAS $\mathcal{S} = \langle R, \mathcal{I}_0, \mathcal{F}, T, \mathcal{M} \rangle$, its STS $\Upsilon_{\mathcal{S}}^{\mathcal{S}}$ is defined as a tuple $\langle R, \Sigma, s_0, db, \Rightarrow \rangle$, which is similar to an RTS, except from the fact that states are attached to ABoxes, not database instances. In particular, $\Upsilon_{\mathcal{S}}^{\mathcal{S}}$ is defined as a “virtualization” of the RTS $\Upsilon_{\mathcal{S}}^{\mathcal{R}} = \langle R, \Sigma, s_0, db, \Rightarrow \rangle$ at the Semantic Layer: it maintains the structure of $\Upsilon_{\mathcal{S}}^{\mathcal{R}}$ unaltered, reflecting that the progression of the system is determined at the Relational Layer, but it associates each state to a virtual ABox obtained from the application of the mapping specification \mathcal{M} to the database instance associated by $\Upsilon_{\mathcal{S}}^{\mathcal{R}}$ to the same state. Formally, the transition relation \mathcal{M} is equivalent to the one of the \mathcal{S} , and the *abox* function of $\Upsilon_{\mathcal{S}}^{\mathcal{S}}$ is defined as follows: for each $s \in \Sigma$, $abox(s) = \mathcal{M}(db(s))$.

4 Verification of Semantically-Enhanced Artifact Systems

Given a SAS \mathcal{S} , we are interested in studying verification of semantic dynamic/temporal properties specified over the Semantic Layer, i.e., to be checked against the STS $\Upsilon_{\mathcal{S}}^{\mathcal{S}}$. As verification formalism, we consider a variant of first-order μ -calculus [16,25], called $\mu\mathcal{L}_A^{\text{EQL}}$ [2,11]. We observe that μ -calculus is one of the most powerful temporal logics: it subsumes LTL, PSL, and CTL* [12]. The logic $\mu\mathcal{L}_A^{\text{EQL}}$ extends propositional μ -calculus by allowing to query the states of the STS using the first-order epistemic queries introduced in Section 2. In $\mu\mathcal{L}_A^{\text{EQL}}$, first-order quantification is restricted to objects that are present in the current ABox, and can be used to relate objects across states. The syntax of $\mu\mathcal{L}_A^{\text{EQL}}$ is as follows:

$$\Phi ::= Q \mid \neg\Phi \mid \Phi_1 \wedge \Phi_2 \mid \exists x.\text{LIVE}(x) \wedge \Phi \mid \langle - \rangle\Phi \mid Z \mid \mu Z.\Phi$$

Where Q is an ECQ over T , Z is a second-order variable denoting a 0-ary predicate, μ is the least fixpoint operator, and the special predicate $\text{LIVE}(x)$ is used to indicate that x belongs to the current active domain, i.e., it is mentioned in some concept or role of the current ABox. For a detailed semantics of $\mu\mathcal{L}_A^{\text{EQL}}$, refer to [2,11].

Given a SAS \mathcal{S} , we show that verification of $\mu\mathcal{L}_A^{\text{EQL}}$ properties over the STS $\Upsilon_{\mathcal{S}}^{\mathcal{S}}$ can be reduced to verification of $\mu\mathcal{L}_A$ [10] properties over the RTS $\Upsilon_{\mathcal{S}}^{\mathcal{R}}$, where $\mu\mathcal{L}_A$ is a logic similar to $\mu\mathcal{L}_A^{\text{EQL}}$, except for the local formula Q , which is an (open) first-order query over the database schema in the Relational Layer.

Theorem 1. *For every SAS \mathcal{S} and $\mu\mathcal{L}_A^{\text{EQL}}$ property Φ , there exists a $\mu\mathcal{L}_A$ property Φ' such that $\Upsilon_{\mathcal{S}}^{\mathcal{R}}$ satisfies Φ if and only if $\Upsilon_{\mathcal{S}}^{\mathcal{S}}$ satisfies Φ' .*

Proof (sketch). The proof extends the one in [9]. We recap such a proof here, and show how to extend it to the logic $\mu\mathcal{L}_A^{\text{EQL}}$. The construction of Φ' is done by *rewriting* and *unfolding* the Φ , by suitably extending the notions of rewriting and unfolding in OBDA to the case of temporal properties [22].

As for the rewriting step, the TBox T does not focus on the dynamics of the system, but only on the information maintained by \mathcal{S} . Therefore, the rewriting of Φ wrt the TBox

T is done by separating the treatment of the dynamic part, which is simply maintained unaltered, to the one of the embedded ECQs, which is handled as in standard OBDA.

The unfolding step is managed in a similar way. Both Υ_S^S and Υ_S^R have the same structure, reflecting that the dynamics of the SAS are determined at the Relational Layer. Hence, the unfolding of $rew(\Phi, T)$ is handled by maintaining the dynamic component unaltered, and unfolding all the embedded ECQs wrt the mapping assertion \mathcal{M} . The correctness of this approach has been proven in [9], which however requires here to extend the proof to the following two cases, related to first-order quantification: (i) the unfolding of $LIVE(x)$; (ii) the unfolding of quantification across states, i.e., $\exists x.LIVE(x) \wedge \Phi$. As for $LIVE(x)$, it can be defined in terms of the following UCQ, which can then be unfolded in the usual way:

$$LIVE(x) \equiv \bigvee_{N \text{ in } T} N(x) \vee \bigvee_{P \text{ in } T} (P(x, _) \vee P(_, x))$$

Quantification across states is dealt as in standard OBDA, i.e.:

$$\begin{aligned} UNFOLD(\exists x.LIVE(x) \wedge \Phi, \mathcal{M}) = & \bigvee_{(f/n) \in FS(\mathcal{M})} \exists x_1, \dots, x_n. UNFOLD(LIVE(f(x_1, \dots, x_n)), \mathcal{M}) \\ & \wedge UNFOLD(\Phi[x/f(x_1, \dots, x_n)], \mathcal{M}) \end{aligned}$$

where $FS(\mathcal{M})$ is the set of function symbols used to construct object terms in \mathcal{M} . To conclude the proof, we need to observe that the unfolded formula is in $\mu\mathcal{L}_A$, i.e., it preserves the property that quantification ranges over the current active domain. This can be easily confirmed by noticing that a live object term in a state of Υ_S^S can only be produced starting from live values in the corresponding state of Υ_S^R .

5 SAS Instantiation: The OBGSM Tool

In this section, we show how the formal framework of SASs can be concretely instantiated, in the case where the transition relation at the Relational Layer is obtained from artifacts specified using the Guard-Stage-Milestone (GSM) approach. In particular, we show how existing techniques and tools have been suitably combined into a tool, called *OBGSM*, which enables the verification of GSM-based artifacts with a Semantic Layer.

GSM Overview. For the sake of space, here we provide a general overview of the GSM methodology and we refer to [19,14] for more detailed and formal definitions. GSM is a declarative modeling framework that has been designed to simplify design and maintenance of business models. The GSM information model uses (possibly nested) attribute/value pairs to capture the domain of interest. The key elements of a lifecycle model are *stages*, *milestones* and *guards* (see Figure 6). Stages (represented as rounded rectangles) are possibly hierarchical clusters of activities, intended to update and extend the data of the information model. They are associated to milestones (circles), business operational objectives which can be achieved while the stage is under execution. Each stage has one or more guards (diamonds), which control the activation of stages and, like milestones, are described in terms of data-aware expressions, involving conditions over the artifact information model.

The Ontology-Based GSM (OBGSM) Tool. OBGSM exploits two already existing tools to provide its functionalities: (i) -ONTOP⁵, a JAVA-based framework for OBDA.

⁵ <http://ontop.inf.unibz.it/>

(ii) the *GSMC model checker*, developed within the EU FP7 Project ACSII⁶ to verify GSM-based artifact-centric systems against temporal/dynamic properties [4]. An important observation is related to semantic governance in this setting: since the construction of the RTS for GSM is handled internally by GSMC, it is not possible (at least for the time being) to prune it so as to remove inconsistent states. Therefore, in the following we assume that all the states in the RTS are consistent with the constraints of the Semantic Layer. This can be trivially achieved by, e.g., avoiding to use negative inclusion assertions in the TBox, which are the only source of inconsistency for OWL 2 QL. A more elaborated discussion on this topic is included in Section 7.

The main purpose of OBGSM is: given a temporal property specified over the Semantic Layer of the system, together with mapping assertions whose language is suitably shaped to work with GSMC, automatically rewrite and unfold the property by producing a corresponding translation that can be directly processed by the GSMC model checker. This cannot be done by solely relying on the functionalities provided by -ONTOP-, for two reasons: (i) OBGSM deals with temporal properties specified in a fragment of $\mu\mathcal{L}_A^{\text{EQL}}$, and not just (local) ECQs; (ii) the mapping assertions are shaped so as to reflect the specific query language supported by GSMC, guaranteeing that the rewriting and unfolding process produces a temporal property expressed in the input language of GSMC. In particular, we note that GSMC is not able to process the entire $\mu\mathcal{L}_A$ logic, but only its CTL fragment, denoted by CTL_A . This requires also to restrict the $\mu\mathcal{L}_A^{\text{EQL}}$ verification formalism accordingly, in particular focusing on its CTL fragment, denoted by $\text{CTL}_A^{\text{EQL}}$.

As depicted in the Figure 2, the workflow of OBGSM is as follows: (i) The tool reads and parses the input conceptual temporal property Φ , the input ontology, and the input mapping declaration \mathcal{M} . (ii) It then *rewrites* the input conceptual temporal property Φ based on the input ontology (TBox) in order to compile away the TBox. This step produces a *rewritten temporal property* $\text{REW}(\Phi, T)$. (iii) The rewritten property $\text{REW}(\Phi, T)$ is *unfolded* by exploiting \mathcal{M} . The final temporal property $\Phi_{\text{GSM}} = \text{UNFOLD}(\text{REW}(\Phi, T), \mathcal{M})$ obeys to the syntax expected by GSMC, and is such that verifying Φ over the transition system of the GSM model under study after projecting its states into the Semantic Layer through \mathcal{M} , is equivalent to verifying Φ_{GSM} directly over the GSM model (without considering the Semantic Layer). (iv) GSMC is invoked by passing Φ_{GSM} and the specification file of the GSM model under study.

Notice that the correctness of the translation is guaranteed by the fact that OBGSM manipulates the local components of the query Φ according to the standard rewriting and unfolding algorithms, while maintaining untouched the temporal structure of the property. This has been proven to be the correct way of manipulating the property (cf. Theorem 1). Furthermore, since the temporal component of the formula is left untouched, it is guaranteed that $\text{CTL}_A^{\text{EQL}}$ properties are rewritten and unfolded into corresponding CTL_A properties.

The verification language of OBGSM relies on CTL, in accordance to the input verification language of GSMC [4,17] with some adjustment. Since we are querying over an OWL ontology, we rely for the local query language on SPARQL, in accordance to the query language supported by -ONTOP-. More specifically, the syntax is as follows:

⁶ “Artifact-Centric Service Interoperation”, see <http://www.acsi-project.eu/>




Fig. 2. OBGSM System Architecture

$$\begin{aligned} \Phi ::= & [Q] \mid (\Phi) \mid \Phi_1 \text{ AND } \Phi_2 \mid \Phi_1 \text{ OR } \Phi_2 \mid \Phi_1 \rightarrow \Phi_2 \mid !\Phi \mid \text{AG } \Phi \mid \text{EG } \Phi \\ & \mid \text{AF } \Phi \mid \text{EF } \Phi \mid \text{AX } \Phi \mid \text{EX } \Phi \mid \text{A } (\Phi_1 \text{ UNTIL } \Phi_2) \mid \text{E } (\Phi_1 \text{ UNTIL } \Phi_2) \\ & \mid \text{FORALL } \mathbf{Var} . \text{fQuantification} \mid \text{EXISTS } \mathbf{Var} . \text{eQuantification} \end{aligned}$$

$$\text{fQuantification} ::= [Q] \rightarrow \Phi \mid \text{FORALL } \mathbf{Var} . \text{fQuantification} \mid [Q]$$

$$\text{eQuantification} ::= [Q] \text{ AND } \Phi \mid \text{EXISTS } \mathbf{Var} . \text{eQuantification} \mid [Q]$$

where “!” denotes the negation, and $[Q]$ is a SPARQL 1.1⁷ SELECT query, which only allows for: (i) prefixes declarations, (ii) variables selections, and (iii) a WHERE clause that only contains triple patterns and some filters on variable/constant comparison.

The semantics of the temporal operators is the one of CTL [12]. Additionally, we also impose the following restrictions on first-order quantification: (i) Only closed temporal formulae are supported for verification. (ii) Each first-order quantifier must be “guarded” in such a way that it ranges over individuals present in the current active domain. This active domain quantification is in line with GSMC, and also with the $\mu\mathcal{L}_A^{\text{EQL}}$ logic. As attested by the grammar above, this is syntactically guaranteed by requiring quantified variables to appear in $[Q]$ according to the following guidelines: (a) $\forall \vec{x}. Q(\vec{x}) \rightarrow \phi$ (b) $\exists \vec{x}. Q(\vec{x}) \wedge \phi$. (iii) Quantified variables must obey to specific restrictions as follows: for each variable y ranging over values, there must be at least one variable x that ranges over object terms and that appears in the first component of the corresponding attribute (i.e., $\text{Attr}(x, y)$ is present in the query, with Attr being an attribute of the ontology), such that x is quantified “before” y . For example, $\forall x. C(x) \implies \exists y. \text{Attr}(x, y)$ satisfies this condition, whereas $\exists y \exists x. \text{Attr}(x, y)$ does not. These restrictions have been introduced so as to guarantee that the conceptual temporal property can be translated into a corresponding GSMC temporal property. In fact, GSMC poses several restrictions on the way values (i.e., attributes of artifacts) can be accessed.

The structure of the OBGSM mapping language is mainly borrowed from -ONTOP-. Each mapping assertion is described by three components: (i) `mappingId`, which provides a unique identifier for the mapping assertion. (ii) `target`, which contains the *target query* (query over ontology) with the syntax adopted from -ONTOP-. (iii) `source`, which describes the *source query*. For the grammar of the source query we used a

⁷ <http://www.w3.org/TR/sparql11-query/>

modified grammar of the GSMC input language [17], with extensions that allow to “extract” artifact identifiers and their value attributes, so as to link them to the ontology. For more detail about OBGSM, please refer to [5].

6 Case Study: Energy Scenario

As a case study to demonstrate our approach we refer to the fragment of the Energy Use Case Scenario developed within the ACSI Project [26]. We show how the Semantic Layer can be exploited in order to facilitate the specification of temporal properties of interest, and discuss how these are automatically translated into properties that can be directly verified by GSMC over the Energy GSM model.

Energy Scenario at a Glance. We sketch the main aspects of the Energy Use Case, and refer to [26] for further details. This use case focuses on the electricity supply exchange process between electric companies inside a distribution network. The exchange occurs at *control points* (CP). Within a CP, a measurement of electricity supply exchange takes place in order to calculate the fair remuneration that the participating companies in the CP should receive. The measurement is done by a *meter reader company*, which corresponds to one of the companies pertaining to that particular CP. The measurement results from the CP are then submitted to the *system operator*, who is in charge of processing the results and publishing a *CP monthly report*. A participating company can raise an objection concerning the published measurement. Once all the risen objections are resolved, the report is closed. The collection of CP monthly reports is then represented as a *CP monthly report accumulator*.

In order to implement the sketched scenario using the GSM methodology, we model the *Control Point Monthly Report (CPMR)* as an artifact which contains the information about hourly measurements done in a particular CP within a certain month. The lifecycle of an instance of this artifact begins once the hourly measurements are provided, and runs until the liquidation for the CP measurements has started. The artifact lifecycle consists of three root stages: (i) *CPMRInitialization*, which is activated when a new instance of the CPMR artifact is created. (ii) *Claiming*, which handles the submission of measurements, and the subsequent reviewing procedure during which objections may be raised. (iii) *MeasurementUpdating*, which is activated once an event is raised, requesting for updating the measurement results.

Figure 3 shows a GSM model of the *Claiming* stage, on which we focus here. It contains the following five sub-stages: (i) *Drafting*, which models the process of submitting and processing the CP hourly measurements and processing a draft of the CPMR. (ii) *Evaluating*, which is responsible for official measurement publication. (iii) *Reviewing*, which models the process of either accepting the official measurement, or deciding whether the difference is acceptable or not. Such decision may result in an objection either by the Metering Office (MO) or the Energy Control Department (ECD). (iv) *CreateObjection*, which corresponds to creation of a new Objection instance. (v) *Closing*, which is responsible for closing the objection period. Once this period finishes and there are no requested objections, the monthly report is finalized and the *CPMRFinished* milestone is achieved. The collection of CP monthly reports for a specific month is then modeled by the *CPMR Monthly Accumulator (CPMRMA)* artifact.




Fig. 3. The Claiming Stage in the GSM Model for the CPMR Artifact

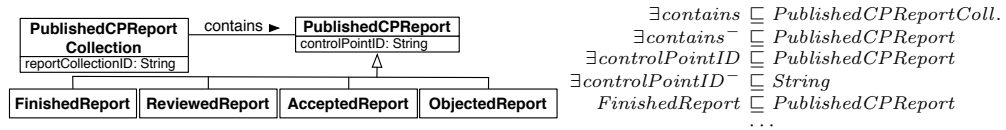


Fig. 4. Ontology for the CPMR reviewing process in the Energy scenario

The Semantic Layer for the Energy Use Case. Here we provide a Semantic Layer built on top of the GSM model for the Energy scenario, restricting our attention to the Control Point Measurement Report (CPMR).

The Ontology TBox. We consider the following states of a control point measurement report: (i) a finished CPMR (when the milestone *CPMRFinished* is achieved), (ii) a reviewed CPMR (after finishing the review inside the *Reviewing* stage) (iii) an accepted CPMR (when the milestone *PublishedOK* is achieved) (iv) an objected CPMR. Following this intuition, a portion of the ontology TBox of the Semantic Layer is depicted in Figure 4, together with its representation as a UML Class Diagram.

The Mapping Assertions. We use mapping assertions as described in Section 5 in order to link the GSM information model to the Semantic Layer. Due to the space limitations, we only give an example of such a mapping in the OBGSM mapping specification language:

```

mappingId ReviewedReportMapping
target    <"&::cpmr/{x}/"> rdf:type :ReviewedReport .
source    get('x', 'CPMR') (GSM.isMilestoneAchieved('x', 'AcceptingPublishedOK') OR
                           GSM.isMilestoneAchieved('x', 'MOReviewingPublishedOK') OR
                           GSM.isMilestoneAchieved('x', 'ECDReviewingPublishedOK'))

```

Intuitively, this mapping populates the *ReviewedReport* concept with the CPMR artifact instances, given the achievement of either of the milestones: *AcceptingPublishedOK*, *MOReviewingPublishedOK*, or *ECDReviewingPublishedOK*. In the Semantic Layer, *ReviewedReport* intuitively represents a CPMR that has been reviewed. In the Artifact Layer, this corresponds to the situation in which the CPMR has been reviewed and accepted either by the electric company, or by the metering office (MO), or by the Electric Control Department (ECD). This example shows how such details can be abstracted away in the Semantic Layer, which does not show that a *ReviewedReport* is obtained by a (possibly complex) chaining of achieved milestones in the underlying GSM model.

Verification within the Energy Use Case. We demonstrate now how the presence of the Semantic Layer may help when specifying different temporal properties of interest. For example, consider a property saying that: “*All control point monthly reports will eventually be finished*”. This can be expressed as a conceptual temporal property over the Semantic Layer defined above:

$$\text{AG}(\text{FORALL } x. ([\text{PublishedCPRReport}(x)] \rightarrow \text{EF}[\text{FinishedReport}(x)]))$$

Then it can be encoded in the machine-readable SPARQL-based verification language, defined in Section 5, as follows:

```
AG (FORALL ?x. ([PREFIX : <http://acsi/example/ACSIEnergy/ACSIEnergy.owl#>
  SELECT ?x WHERE { ?x a :PublishedCPRReport} ] ->
  EF [PREFIX : <http://acsi/example/ACSIEnergy/ACSIEnergy.owl#>
  SELECT ?x WHERE { ?x a :FinishedReport} ]
));
```

We now show how this kind of high-level properties are processed by OBGSM into temporal properties formulated over the GSM model and which can be fed to the GSMC model checker. Notice that without the Semantic Layer and OBGSM, the user would have to construct the low-level properties manually.

We first perform a rewriting step, which expands concepts contained in the property of interest using the ontology TBox. In our example, this results in the conceptual property $\text{AG}(\text{PCR}(x) \rightarrow \text{EF}[\text{FinishedReport}(x)])$, in which:

$$\begin{aligned} \text{PCR}(x) = & \text{EXISTS } y. [\text{hasControlPointID}(x, y)] \text{ OR } [\text{ObjectedReport}(x)] \text{ OR} \\ & [\text{AcceptedReport}(x)] \text{ OR } [\text{PublishedCPRReport}(x)] \text{ OR } [\text{ReviewedReport}(x)] \\ & \text{OR } [\text{FinishedReport}(x)] \text{ OR } (\text{EXISTS } z. [\text{contains}(z, x)]) \end{aligned}$$

In particular, the rewriting step employs reasoning over the $DL\text{-Lite}_{\mathcal{R}}$ TBox by taking into account the following knowledge contained in the Semantic Layer: (i) Those objects that have a control point ID (i.e., are in the domain of the attribute *hasControlPointID*), are instances of *PublishedCPRReport*; (ii) *ObjectedReport* is a *PublishedCPRReport*; (iii) *AcceptedReport* is a *PublishedCPRReport*; (iv) *ReviewedReport* is a *PublishedCPRReport*; (v) Those objects that are in the range of the role *contains* are instances of *PublishedCPRReport*; (vi) *FinishedReport* is a *PublishedCPRReport*.

We then perform the unfolding step, which exploits the mapping assertions in order to reformulate the rewritten property to the temporal property over the GSM model. Such property, shown in Figure 5, can be now be fed to the GSMC model checker in order to validate the initial property claiming that “*all control point monthly reports will eventually be finished*”. Given the significant difference in complexity between the property formulation over the Semantic Layer and its corresponding translation over GSM, the advantage of the presence of the Semantic Layer becomes apparent even for this simple case study. Conceptual temporal property formulation allows to hide low-level details from the user and enables the modeler to rather focus on the domain under study using the vocabulary he/she is familiar with.

```

AG (forall('x', 'CPMR')(!(GSM.isMilestoneAchieved('x', 'Objected') OR
GSM.isMilestoneAchieved('x', 'ObjectionRequested') OR
GSM.isMilestoneAchieved('x', 'ObjectionCreated') OR
GSM.isMilestoneAchieved('x', 'MOReviewingObjectionRequested') OR
GSM.isMilestoneAchieved('x', 'ECDReviewingObjectionRequested') OR
GSM.isMilestoneAchieved('x', 'PublishedOK') OR
GSM.isMilestoneAchieved('x', 'AcceptingPublishedOK') OR
GSM.isMilestoneAchieved('x', 'MOReviewingPublishedOK') OR
GSM.isMilestoneAchieved('x', 'ECDReviewingPublishedOK') OR
exists('y', 'CPMRMA')(y./CPMRMA/CPMRDATA ->
exists(CPMRID == x./CPMR/ID)) OR
GSM.isMilestoneAchieved('x', 'CPMRFinished') OR
EF(GSM.isMilestoneAchieved('x', 'CPMRFinished'))))

```

Fig. 5. Resulting temporal property over the GSM model

7 Discussion and Conclusion

In this paper we have introduced a framework in which an artifact system based on the GSM model is connected through declarative mappings to a Semantic Layer expressed as an OWL 2 QL ontology, and temporal properties expressed over the Semantic Layer are reformulated in terms of the underlying GSM model and verified over its execution, using state of the art OBDA and model checking technology. We have assumed in our development that all states generated for the purpose of verification by the GSMC model checker (which we use as a black box), are consistent with the Semantic Layer. Indeed, inconsistent states could not be easily pruned away, as e.g., the approach proposed in [9] would require. Notice that in our case study presented in Section 6 such assumption holds since the TBox does not contain any disjointness assertions.

However, in general, the consistency of the states at the artifact level with the TBox in the Semantic Layer is not guaranteed. In such a situation the strategy proposed here, based on delegating the verification of CTL_A properties to the existing model checker, cannot be followed directly. One possible solution to this is to provide a mechanism for the model checker to exploit the ontology at the Semantic Layer to distinguish consistent from inconsistent states of the low-level system. The other possible solution is to consider less expressive temporal logics, i.e., fragments of CTL_A , and investigate whether the check for consistency can be embedded in the formula to verify. These latter scenarios provide us with interesting problems for future investigation.

References

1. Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P.F. (eds.): The Description Logic Handbook: Theory, Implementation and Applications. Cambridge University Press (2003)
2. Bagheri Hariri, B., Calvanese, D., Montali, M., De Giacomo, G., De Masellis, R., Felli, P.: Description logic Knowledge and Action Bases. J. of Artificial Intelligence Research (2013)
3. Bao, J., et al.: OWL 2 Web Ontology Language document overview (second edition). W3C Recommendation, World Wide Web Consortium (2012)

4. Belardinelli, F., Lomuscio, A., Patrizi, F.: Verification of GSM-based artifact-centric systems through finite abstraction. In: Proc. of ICSSOC 2012. LNCS, vol. 7636, pp. 17–31 (2012)
5. Calvanese, D., Bagheri Hariri, B., De Masellis, R., Lembo, D., Montali, M., Santoso, A., Solomakhin, D., Tessaris, S.: Techniques and tools for KAB, to manage action linkage with the Artifact Layer – Iteration 2. Deliverable ACSI-D2.4.2, ACSI Consortium (May 2013), available at <http://www.inf.unibz.it/~calvanese/papers-html/ACSI-D2.4.2-2013.html>
6. Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M., Poggi, A., Rodríguez-Muro, M., Rosati, R.: Ontologies and databases: The *DL-Lite* approach. In: Tessaris, S., Franconi, E. (eds.) RW 2009 Tutorial Lectures, LNCS, vol. 5689, pp. 255–356. Springer (2009)
7. 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. pp. 274–279 (2007)
8. Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M., Rosati, R.: Tractable reasoning and efficient query answering in description logics: The *DL-Lite* family. J. of Automated Reasoning 39(3), 385–429 (2007)
9. Calvanese, D., De Giacomo, G., Lembo, D., Montali, M., Santoso, A.: Ontology-based governance of data-aware processes. In: Proc. of RR 2012. pp. 25–41 (2012)
10. Calvanese, D., De Giacomo, G., Montali, M.: Foundations of data aware process analysis: A database theory perspective. In: Proc. of PODS 2013 (2013)
11. Calvanese, D., Kharlamov, E., Montali, M., Santoso, A., Zheleznyakov, D.: Verification of inconsistency-tolerant knowledge and action bases. In: Proc. of IJCAI 2013 (2013)
12. Clarke, E.M., Grumberg, O., Peled, D.A.: Model checking. The MIT Press (1999)
13. Dam, M.: CTL* and ECTL* as fragments of the modal μ -calculus. In: Proc. of the 17th Colloquium on Trees in Algebra and Programming (CAAP’92) (1992)
14. Damaggio, E., Hull, R., Vaculín, R.: On the equivalence of incremental and fixpoint semantics for business artifacts with Guard-Stage-Milestone lifecycles. Information Systems (2013)
15. Deutsch, A., Hull, R., Patrizi, F., Vianu, V.: Automatic verification of data-centric business processes. In: Proc. of ICDT 2009. pp. 252–267 (2009)
16. Emerson, E.A.: Model checking and the Mu-calculus. In: Immerman, N., Kolaitis, P. (eds.) Proceedings of the DIMACS Symposium on Descriptive Complexity and Finite Models. pp. 185–214. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, American Mathematical Society Press (1996), ISBN 0-8218-0517-7
17. Gonzales, P., Griesmayer, A., Lomuscio, A.: Model checking tool for artifact interoperations (MOCAI) – Iteration 3. Deliverable ACSI-D2.2.3, ACSI Consortium (May 2013)
18. Hull, R.: Artifact-centric business process models: Brief survey of research results and challenges. In: Proc. of ODBASE 2008. LNCS, vol. 5332, pp. 1152–1163 (2008)
19. Hull, R., Damaggio, E., De Masellis, R., Fournier, F., Gupta, M., Heath, III, F.T., Hobson, S., Linehan, M., Maradugu, S., Nigam, A., Sukaviriya, P.N., Vaculin, R.: Business artifacts with Guard-Stage-Milestone lifecycles: Managing artifact interactions with conditions and events. In: Proc. of the 5th ACM Int. Conf. on Distributed Event-Based Systems (DEBS 2011) (2011)
20. Motik, B., Cuenca Grau, B., Horrocks, I., Wu, Z., Fokoue, A., Lutz, C.: OWL 2 Web Ontology Language profiles (second edition). Tech. rep., W3C Recommendation (2012)
21. Nigam, A., Caswell, N.S.: Business artifacts: An approach to operational specification. IBM Systems J. 42(3), 428–445 (2003)
22. 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)
23. Reichert, M.: Process and data: Two sides of the same coin? In: Proc. of the On the Move Confederated Int. Conf. (OTM 2012). LNCS, vol. 7565, pp. 2–19 (2012)
24. Rodríguez-Muro, M., Calvanese, D.: Quest, an OWL 2 QL reasoner for ontology-based data access. In: Proc. of OWLED 2012. CEUR, ceur-ws.org, vol. 849 (2012)

25. Stirling, C.: Modal and Temporal Properties of Processes. Springer (2001)
26. Toribio Gómez, D., Murphy-O Connor, C., Leenheer, P.D., Malarne, P.: Deployment and evaluation of pilots using the ACSI Hub System – Results and evaluation. Deliverable ACSI-D5.5, ACSI Consortium (May 2013)