

Querying Temporal Databases via OWL 2 QL

Szymon Klarman and Thomas Meyer

Centre for Artificial Intelligence Research,
CSIR Meraka and University of KwaZulu-Natal, South Africa
{sklarman,tmeyer}@csir.co.za

Abstract. SQL:2011, the most recently adopted version of the SQL query language, has unprecedentedly standardized the representation of temporal data in relational databases. Following the successful paradigm of ontology-based data access, we develop a practical approach to querying the SQL:2011-based temporal data model via the semantic layer of OWL 2 QL. The interval-based *temporal query language* (TQL), which we propose for this task, is based on naturally characterizable combinations of temporal logic with conjunctive queries. As the central contribution, we present rules for sound and complete rewriting of TQL queries into two-sorted first-order logic, and consequently, into corresponding SQL queries, which can be evaluated in any existing relational database management system compliant with the SQL:2011 temporal data model. Importantly, the proposed rewriting is based on the direct reuse of the standard rewriting techniques for conjunctive queries under OWL 2 QL. This renders our approach modular and easily implementable. As a notable corollary, we show that the data complexity of TQL query answering remains in AC^0 , i.e., as in the usual, non-temporal case.

1 Introduction

The ability to manage the temporal aspects of information is critical for a variety of applications. One natural and prevailing scenario is that of representing and querying the *validity time* of data, i.e., the time during which data is deemed true about the application domain. The significance of this task is particularly visible in the area of semantic technologies, where the systematically growing number of proposed solutions, building on different levels of the Semantic Web architecture and differing in the flavour and depth of temporal reasoning they support, aim at addressing essentially the same problem [15,14,5,21,7,2]. A very similar proliferation of proposals was witnessed in the 1990s in the field of temporal databases. Intensive attempts to extend the traditional relational data model and SQL with temporal features inspired then a large body of candidate specifications, including such extensions as TSQL2, SQL3 or SQL/Temporal [23], which eventually failed to be adopted by the database community due to the persistent lack of consensus as to the preferred approach. Only very recently, that discussion has been picked up again and a compromise temporal extension has eventually found its way into SQL:2011 [19] — the newest standardization

of the SQL query language. This unprecedented circumstance offers an interesting opportunity to address the problem of reasoning with temporal semantic data from yet another angle, namely, by relating it via known links between relational databases and semantic technologies to its analogue in the database world, thus using the SQL:2011 standard as a leverage for the solution. In this paper, we contribute precisely to this research agenda by proposing a novel, temporal extension to the framework of ontology-based data access.

Ontology-based data access (OBDA) is a popular paradigm of managing information, which combines the data storage and querying capabilities offered by relational database management systems (RDBMSs) with the semantically enhanced view on the data provided by ontologies. The ontology language OWL 2 QL, based on the DL-*Lite* family of Description Logics, is a profile OWL 2 designed specifically to support optimally balanced OBDA. In a nutshell, conjunctive queries, posed over data under an OWL 2 QL ontology, can be rewritten into first-order logic using the ontology’s axioms, then translated to SQL and answered within an RDBMS, in such a way as if the ontology was mediating in the process [8]. As large portions of data available through the Web are in fact still hosted in relational datastores, OBDA provides a crucial channel for accessing this data from the level of the Semantic Web applications.

In this work, we establish an analogical OBDA interface between the semantic layer of OWL 2 QL and temporal data model endorsed by SQL:2011. The interval-based *temporal query language* (TQL), which we propose for this task, is based on naturally characterizable combinations of temporal logic with conjunctive queries, identified in [16]. TQL is tailored specifically to offer maximum expressivity while preserving the possibility of reuse central to OBDA first-order rewriting techniques and tools, developed specifically for the use with OWL 2 QL. While this technical compliance warrants the minimal implementation overhead for our approach, its well-defined logic foundations allow us to identify basic formal properties of TQL. In particular, we are able to demonstrate that under the finite time domain assumption the data complexity of query entailment remains in AC^0 , i.e., as in the case of standard (non-temporal) conjunctive queries, even though the combined complexity increases to PSPACE-complete. As the main contribution, we develop a rewriting of TQL queries in the presence of OWL 2 QL ontologies into two-sorted first-order logic, and consequently to SQL, which opens the way to efficient query answering by means of existing, commercially supported RDBMSs, such as IBM DB2 10.1, Oracle Database 11g Workspace Manager, or Teradata — all of which have by now adopted certain variants of the SQL:2011 standard¹.

The paper is organized as follows. In the next section we lay down the preliminaries. In Section 3, we introduce TQL and define the query entailment problem. In Section 4, we present the TQL query rewriting rules and, in Section 5, we study the formal properties of query answering queries via this rewriting. The related work is discussed in Section 6 and the paper is concluded in Section 7. Some proofs are included in the appendix.

¹ See <http://www.cs.arizona.edu/~rts/sql3.html> for an overview.

2 Basic Notions

We start by recapping the logic foundations of OBDA. Then we motivate and formally introduce the temporal extension of OBDA studied in this paper.

2.1 Ontology-Based Data Access

OWL 2 QL is a profile of OWL 2 based on the DL-*Lite* family of Description Logics (DLs) [8]. A DL vocabulary $\Sigma = (\mathbf{N}_I, \mathbf{N}_C, \mathbf{N}_R)$ consists of countably infinite sets of individual names (\mathbf{N}_I), concept names (\mathbf{N}_C) and role names (\mathbf{N}_R). An ABox \mathcal{A} is a finite set of assertions of type $A(a)$ and $r(a, b)$, for $a, b \in \mathbf{N}_I$, $A \in \mathbf{N}_C$ and $r \in \mathbf{N}_R$, which we also generically denote with $\alpha(\mathbf{a})$. A TBox \mathcal{T} is a finite set of concept inclusions $B \sqsubseteq C$ and role inclusions $r \sqsubseteq s$, where B, C and r, s are possibly complex concepts and roles, respectively, built using logical constructors allowed in OWL 2 QL, such as $\exists r. \top, A \sqcap B, r^-$, whose particulars are not of importance for this work.² The semantics is given in terms of DL interpretations $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, defined as usual [4]. An interpretation \mathcal{I} is a model of \mathcal{T} and \mathcal{A} , denoted as $\mathcal{I} \models \mathcal{T}, \mathcal{A}$, iff it satisfies every axiom in \mathcal{T} and \mathcal{A} .

In OBDA, the instance data, represented as an ABox, is accessed via an ontology, given as a TBox, using a designated query language such as, most commonly considered in that context, the language of conjunctive queries [13]. Let \mathbf{N}_V be a countably infinite set of variables. A *conjunctive query* (CQ) over a DL vocabulary Σ is a first-order formula:

$$q(\mathbf{y}) = \exists \mathbf{x}. \left(\bigwedge_{1 \leq j \leq n} \alpha_j(\mathbf{d}_j) \right)$$

where \mathbf{y} and \mathbf{x} are sequences of, respectively, free and existentially bounded variables occurring in $q(\mathbf{y})$ and every atom $\alpha_j(\mathbf{d}_j)$ is of the form $A(d)$ or $r(d_1, d_2)$, where $A \in \mathbf{N}_C$, $r \in \mathbf{N}_R$, and $d, d_1, d_2 \in \mathbf{N}_I \cup \mathbf{N}_V$. Whenever it is not confusing, we sometimes also abbreviate $q(\mathbf{y})$ to q . By $\text{term}(q)$ we denote the set of all terms occurring in q and by $\text{obj}(q)$ the set of all free variables. We call q grounded whenever $\text{obj}(q) = \emptyset$. A grounded CQ q is satisfied in \mathcal{I} iff there exists a mapping $\mu : \text{term}(q) \mapsto \Delta^{\mathcal{I}}$, with $\mu(d) = d^{\mathcal{I}}$ for every $d \in \mathbf{N}_I$, such that $\mu(d) \in A^{\mathcal{I}}$ and $(\mu(d_1), \mu(d_2)) \in r^{\mathcal{I}}$ for every $A(d)$ and $r(d_1, d_2)$ in q . Further, we say that q is *entailed* by \mathcal{T}, \mathcal{A} , denoted as $\mathcal{T}, \mathcal{A} \models q$ iff q is satisfied in every model of \mathcal{T}, \mathcal{A} . Whenever $\emptyset, \mathcal{A} \models q$ we also write $\mathcal{A} \models q$. An *answer* to q is a mapping $\sigma : \text{obj}(q) \mapsto \mathbf{N}_I$. By $\sigma(q)$ we denote the result of uniformly substituting every occurrence of x in q with $\sigma(x)$, for every $x \in \text{obj}(q)$. An answer σ is called *certain* over \mathcal{T}, \mathcal{A} iff $\sigma(q)$ is entailed by \mathcal{T}, \mathcal{A} .

A prominent property of CQs is their *first-order* (FO) *rewritability* in OWL 2 QL, formally defined as follows.

Definition 1 (FO rewritability [8]). *For every CQ q and a TBox \mathcal{T} , there exists a FO formula $q^{\mathcal{T}}$, called the FO rewriting of q in \mathcal{T} , such that for every ABox \mathcal{A} and answer σ to q , it holds that:*

² See http://www.w3.org/TR/owl2-profiles/#OWL_2_QL for full details.

$$\mathcal{T}, \mathcal{A} \models \sigma(q) \text{ iff } db(\mathcal{A}) \models \sigma(q^{\mathcal{T}}),$$

where \models is the FO satisfaction relation and $db(\mathcal{A})$ denotes \mathcal{A} considered as a database/FO interpretation, i.e., a structure $(\mathbf{N}_I, \cdot^{\mathcal{D}})$, where \mathbf{N}_I is the data domain and $\cdot^{\mathcal{D}}$ is an interpretation function defined as $\alpha^{\mathcal{D}} = \{\mathbf{a} \mid \alpha(\mathbf{a}) \in \mathcal{A}\}$, for every $\alpha \in \mathbf{N}_C \cup \mathbf{N}_R$.

By the standard techniques the FO rewriting of q in \mathcal{T} is a union of conjunctive queries, i.e., a formula $q^{\mathcal{T}}(\mathbf{y}) = \bigvee_{1 \leq i \leq m} q_i(\mathbf{y})$, where every $q_i(\mathbf{y})$ is a CQ.

FO rewritability is particularly significant from the practical perspective. It implies that answering CQs in OWL 2 QL can be effectively performed in existing RDBMSs via a translation to SQL, as the ontological component in the task can be always compiled out in the query rewriting, without loss of soundness or completeness. As a theoretical corollary, it follows also that the data complexity of query answering in this setup is AC^0 , as in first-order logic (FOL).

In this work, we focus exclusively on OWL 2 QL TBoxes, even though some of the presented results should clearly transfer to other fragments of DLs warranting the FO rewritability property for CQs. Without always stating it explicitly, we assume that every TBox or ontology mentioned in the remainder of this paper is expressed in OWL 2 QL.

2.2 Ontology-Based Access to Temporal Data

In this paper, we study ontology-based access to temporal data, in the sense of an extension to the OBDA paradigm whose prototypical application could be illustrated with the following scenario.

Consider a temporal database (TDB) presented in Table 1, with columns **from** and **to** marking the limits of the validity periods of the respective records. Such databases can be naturally mapped to (virtual) temporal ABoxes, such as given in Table 2. Our goal is to define a dedicated language for querying temporal ABoxes, which would combine support for two essential functionalities: representation of temporal constraints over the validity periods of data and semantically enhanced access to that data via ontologies. For instance, given the ontology $\mathcal{T} = \{Emp \sqsubseteq Person, department \sqsubseteq worksAt, location \sqsubseteq basedIn\}$, the language should be able to support queries such as:

(Q) Find all persons X and times Y , such that X worked in a department based in Barcelona during Y and in a department based in Madrid some time earlier.

The expected set of answers should then include *e1* as X with the associated period [1999, 2000] as Y . The practical rationale behind ontology-based access to temporal data defined in this way is to eventually enable such queries to be translated to SQL and answered within existing RDBMSs.

Formally, the temporal data model under consideration is grounded in the point- and (derived) interval-based time domains.

Emp				
id	name	department	from	to
e1	john	d1	1998	2000
e1	john	d3	2000	2003
e2	mark	d2	1999	2002

Dep				
id	type	location	from	to
d1	financial	madrid	1998	1999
d1	financial	barcelona	1999	2003
d2	hr	barcelona	2000	2003
d3	hq	london	2000	2003

Table 1. SQL:2011 temporal database.

[1998, 2000] : $Emp(e1)$
[1998, 2000] : $name(e1, john)$
[1998, 2000] : $department(e1, d1)$
...

[1998, 1999] : $Dep(d1)$
[1998, 1999] : $type(d1, financial)$
[1998, 1999] : $location(d1, madrid)$
...

Table 2. Temporal ABoxes corresponding to the temporal relations in Table 1.

Definition 2 (Time domain, time intervals). A time domain is a tuple $\mathfrak{T} = (T, <)$, where T is a nonempty set of elements called time points and $<$ is an irreflexive, linear ordering on T . A time interval $\tau = [\tau^-, \tau^+]$ over \mathfrak{T} is a set of time points $\{t \in T \mid \tau^- \leq t \leq \tau^+\}$, where $\tau^-, \tau^+ \in T$, such that $\tau^- \leq \tau^+$. The points τ^- and τ^+ are called the beginning and the end of τ .³ The set of all time intervals over \mathfrak{T} is denoted by I .

In the SQL:2011 standard, every temporal relation extends a non-temporal one with two additional attributes storing the beginning and the end time of the validity period of a given tuple [19] — exactly as in the example from Table 1. Such a model supports a representation of so-called concrete TDBs, i.e., finite syntactic encodings of temporal data. The actual meaning of these encodings is captured by possibly infinite abstract TDBs [12]. In the OBDA setting these two notions translate naturally into the corresponding types of ABoxes.

Definition 3 (Concrete and abstract temporal ABoxes). A temporal assertion is an expression $\tau : \alpha(\mathbf{a})$, where $\alpha(\mathbf{a})$ is an ABox assertion and $\tau \in I$, stating that $\alpha(\mathbf{a})$ is valid in every time point in τ . A concrete temporal ABox (CTA) \mathfrak{A} is a finite set of temporal assertions. For a concrete temporal ABox \mathfrak{A} there exists a corresponding abstract temporal ABox (ATA), obtained by means of a mapping $\|\cdot\|$, such that $\|\mathfrak{A}\| = (\mathfrak{A}_t)_{t \in T}$, where $\mathfrak{A}_t = \{\alpha(\mathbf{a}) \mid \tau : \alpha(\mathbf{a}) \in \mathfrak{A} \text{ and } t \in \tau\}$.

The link between concrete temporal ABoxes and SQL:2011 TDBs is defined in a strict analogy to the non-temporal OBDA.

³ Note that SQL:2011 adopts a closed-open semantics for the validity periods, i.e., for any $\tau \in I$ it holds that $\tau^- \in \tau$ and $\tau^+ \notin \tau$. This is a technically insignificant difference which we omit here for clarity of presentation.

Definition 4 (TDB). Let \mathfrak{A} be a CTA over the time domain $\mathfrak{T} = (T, <)$. By $tdb(\mathfrak{A})$ we denote \mathfrak{A} considered as a temporal database (TDB) over the signature $\Gamma = \{R_\alpha \mid \alpha \in \mathbf{N}_C \cup \mathbf{N}_R\}$, i.e., the structure $(\mathbf{N}_I, T, <, \cdot^{\mathcal{D}})$, where \mathbf{N}_I is the data domain, $(T, <)$ is the time domain, and $\cdot^{\mathcal{D}}$ is the interpretation defined as $R_\alpha^{\mathcal{D}} = \{(\mathbf{a}, \tau^-, \tau^+) \mid \tau : \alpha(\mathbf{a}) \in \mathfrak{A}\}$, for every $\alpha \in \mathbf{N}_C \cup \mathbf{N}_R$.

3 Temporal Query Language

The *Temporal Query Language* (TQL), presented in this section, is defined using a generic construction method for temporal query languages in DLs, explored also in [16,3,7,18]. TQL is a combination of a temporal language with CQs, obtained by substituting CQs for the atoms in temporal formulas. This design allows for very flexible interleaving of data queries with temporal constraints, while benefiting from the expressive power of both components. As the temporal language we use first-order monadic logic of orders, which is at least as expressive as most common linear temporal logics [22]. Two specific characteristics of TQL, which distinguish it from other similar proposals, are:

- the use of an interval-based variant of the temporal language, rather than a point-based, which enables direct querying of concrete TDBs, without requiring intermediate translations, such as studied in [24];
- the use of the epistemic semantics for embedding CQs in the temporal language, as suggested in [16], which renders the language more expressive and computationally well-behaved, as explained further.

By \mathbf{l}_V we denote a countably infinite set of variables ranging over I .

Definition 5 (TQL). Temporal query language (TQL) is the smallest set of formulas induced by the grammar:

$$\psi ::= [q](u) \mid u^* < v^* \mid \neg\psi \mid \psi_1 \wedge \psi_2 \mid \exists y.\psi$$

where q is a CQ, $u, v \in I \cup \mathbf{l}_V$, $y \in \mathbf{l}_V$, and $* \in \{-, +\}$. An i -substitution is a mapping $\pi : I \cup \mathbf{l}_V \mapsto I$ such that $\pi(\tau) = \tau$, for every $\tau \in I$. By $\text{obj}(\psi)$ we denote the set of free individual variables and by $\text{int}(\psi)$ the set of free interval variables in ψ . A TQL formula ψ is grounded iff $\text{obj}(\psi) = \text{int}(\psi) = \emptyset$. The satisfaction relation for TQL formulas, w.r.t. a TBox \mathcal{T} , a CTA \mathfrak{A} , and an i -substitution π , is defined inductively as follows:

$$\begin{aligned} (\dagger) \quad \mathcal{T}, \mathfrak{A}, \pi \models [q](u) & \quad \text{iff} \quad \mathcal{T}, \mathfrak{A}_t \models q, \text{ for every } t \in \pi(u), \\ \mathcal{T}, \mathfrak{A}, \pi \models u^* < v^* & \quad \text{iff} \quad \pi(u)^* < \pi(v)^*, \\ \mathcal{T}, \mathfrak{A}, \pi \models \neg\psi & \quad \text{iff} \quad \mathcal{T}, \mathfrak{A}, \pi \not\models \psi, \\ \mathcal{T}, \mathfrak{A}, \pi \models \psi_1 \wedge \psi_2 & \quad \text{iff} \quad \mathcal{T}, \mathfrak{A}, \pi \models \psi_1 \text{ and } \mathcal{T}, \mathfrak{A}, \pi \models \psi_2, \\ \mathcal{T}, \mathfrak{A}, \pi \models \exists y.\psi & \quad \text{iff} \quad \text{there exists } \tau \in I, \text{ such that} \\ & \quad \mathcal{T}, \mathfrak{A}, \pi[y \mapsto \tau] \models \psi, \end{aligned}$$

where $\pi[y \mapsto \tau]$ denotes the i -substitution exactly as π except for that we fix $\pi(y) = \tau$. We say that $\mathcal{T}, \mathfrak{A}$ entail a grounded TQL formula ψ , denoted as $\mathcal{T}, \mathfrak{A} \models \psi$, iff there exists an i -substitution π , such that $\mathcal{T}, \mathfrak{A}, \pi \models \psi$.

TQL formulas with free variables can naturally serve as queries over concrete temporal ABoxes. We refer to such formulas as *concrete TQL queries* (CTQs). As an example of a CTQ, consider a rephrasing of the query (**Q**) from Section 2.2:

$$\psi(x, y) := [\exists z.(Person(x) \wedge worksAt(x, z) \wedge basedIn(z, barcelona))](y) \wedge \exists v.(v^+ < y^- \wedge [\exists z.(worksAt(x, z) \wedge basedIn(z, madrid))](v))$$

As one of its answers, $\psi(x, y)$ should return $\{x \mapsto e1, y \mapsto [1999, 2000]\}$. The certain answer semantics for such queries is defined as expected.

Definition 6 (CTQ answering). *Let \mathcal{T} be a TBox, \mathfrak{A} a CTA and ψ a CTQ with free variables $\text{obj}(\psi)$ and $\text{int}(\psi)$. An answer to ψ is a mapping σ such that $\sigma : \text{obj}(\psi) \mapsto \mathbf{N}_I$ and $\sigma : \text{int}(\psi) \mapsto I$. By $\sigma(\psi)$ we denote the result of uniformly substituting every occurrence of x in ψ with $\sigma(x)$, for every $x \in \text{obj}(\psi) \cup \text{int}(\psi)$. An answer σ is called certain over $\mathcal{T}, \mathfrak{A}$ iff $\mathcal{T}, \mathfrak{A} \models \sigma(\psi)$.*

The key to the design of TQL is condition (\dagger), in Definition 5, which ensures the epistemic interpretation of the CQs embedded in TQL queries. The formula $[q](\tau)$, for a grounded CQ q and $\tau \in I$, reads as “ q is entailed in all time points in τ ”. Analogically, $\neg[q](\tau)$ is interpreted as negation-as-failure: “it is not true that q is entailed in all time points in τ ”. This approach of combining FO-based query languages has been originally proposed by Calvanese et al. [9], giving rise to a family of lightweight query languages, which permit well-behaved, modular answering algorithms and support the use of negation (interpreted as negation-as-failure), which otherwise easily leads to undecidability in the context of querying DL ontologies.

As a consequence, condition (\dagger) can be effectively replaced with its equivalent, which involves the standard FO rewritability techniques in the sense of Definition 1:

$$\mathcal{T}, \mathfrak{A}, \pi \models [q](u) \quad \text{iff} \quad db(\mathfrak{A}_t) \models q^{\mathcal{T}}, \text{ for every } t \in \pi(u),$$

where $q^{\mathcal{T}}$ is an FO rewriting of q in \mathcal{T} . What it eventually means, is that all occurrences of CQs in a CTQ can be replaced with their FO rewritings, so that the ontology \mathcal{T} can be dropped while the formula can be evaluated exclusively over the temporal ABox seen as a sequence of FO interpretations. However, such point-based rewriting strategy, also explored in [7], is highly impractical when applied over concrete TDBs, as it necessitates either unfolding a concrete TDB into an abstract one, or a further translation from point-based to an interval-based language. Instead, here we pursue a direct interval-based approach, which requires a more sophisticated rewriting technique, capable of handling the well-known problems of computing temporal joins and coalescing, as explained in detail in the following section.

4 Query Rewriting

We start by defining a two-sorted first-order language, tailored specifically for talking about temporal relations of TDBs.

Definition 7 (2FO). *The language of two-sorted first-order logic (2FO) over the signature $\Gamma = \{R_1, R_2, \dots\}$ is the smallest set of formulas induced by the grammar:*

$$\varphi ::= R(d_1, \dots, d_n, t_1, t_2) \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid t_1 < t_2 \mid \exists x.\varphi \mid \exists y.\varphi$$

where $R \in \Gamma$, $d_1, \dots, d_n \in \mathbf{N}_I \cup \mathbf{N}_V$, $t_1, t_2 \in T \cup \mathbf{T}_V$, $x \in \mathbf{N}_V$ and $y \in \mathbf{T}_V$. A d -substitution δ is a mapping $\delta : \mathbf{N}_I \cup \mathbf{N}_V \mapsto \mathbf{N}_I$, such that $\delta(a) = a$ for every $a \in \mathbf{N}_I$. A t -substitution ν is a mapping $\nu : T \cup \mathbf{T}_V \mapsto T$, such that $\nu(t) = t$ for every $t \in T$. A 2FO formula φ is called grounded whenever it does not have any free variables. The satisfaction relation for 2FO formulas, w.r.t. a TDB $\text{tdb}(\mathfrak{A}) = (\mathbf{N}_I, T, <, \cdot^{\mathcal{D}})$, a d -substitution δ and a t -substitution ν , is defined inductively as follows:

$$\begin{aligned} \text{tdb}(\mathfrak{A}), \delta, \nu \Vdash R(d_1, \dots, d_n, t_1, t_2) & \text{ iff } (\delta(d_1), \dots, \delta(d_n), \nu(t_1), \nu(t_2)) \in R^{\mathcal{D}}, \\ \text{tdb}(\mathfrak{A}), \delta, \nu \Vdash \neg\varphi & \text{ iff } \text{tdb}(\mathfrak{A}), \delta, \nu \not\Vdash \varphi, \\ \text{tdb}(\mathfrak{A}), \delta, \nu \Vdash \varphi_1 \wedge \varphi_2 & \text{ iff } \text{tdb}(\mathfrak{A}), \delta, \nu \Vdash \varphi_1 \text{ and } \text{tdb}(\mathfrak{A}), \delta, \nu \Vdash \varphi_2, \\ \text{tdb}(\mathfrak{A}), \delta, \nu \Vdash t_1 < t_2 & \text{ iff } \nu(t_1) < \nu(t_2), \\ \text{tdb}(\mathfrak{A}), \delta, \nu \Vdash \exists x.\varphi & \text{ iff for } x \in \mathbf{N}_V, \text{ there exists } a \in \mathbf{N}_I \text{ such} \\ & \text{ that } \text{tdb}(\mathfrak{A}), \delta[x \mapsto a], \nu \Vdash \varphi, \\ \text{tdb}(\mathfrak{A}), \delta, \nu \Vdash \exists y.\varphi & \text{ iff for } y \in \mathbf{T}_V, \text{ there exists } t \in T \text{ such that} \\ & \text{tdb}(\mathfrak{A}), \delta, \nu[y \mapsto t] \Vdash \varphi, \end{aligned}$$

where $\delta[x \mapsto a]$ ($\nu[y \mapsto t]$) denotes the substitution exactly as δ (ν) except for that we fix $\delta(x) = a$ ($\nu(y) = t$). We say that a 2FO formula φ is satisfied in $\text{tdb}(\mathfrak{A})$, denoted as $\text{tdb}(\mathfrak{A}) \Vdash \varphi$, iff there exist d -/ t -substitutions δ, ν , such that $\text{tdb}(\mathfrak{A}), \delta, \nu \Vdash \varphi$.

Next, we define the rules for rewriting CTQs into 2FO formulas.

Definition 8 (2FO rewriting of CTQs). *The 2FO rewriting of a CTQ ψ is a formula $[\psi]^{2FO}$ obtained from ψ by applying the transformation $[\cdot]^{2FO}$, defined inductively as follows:*

$$\begin{aligned} [[q(\mathbf{d})](u)]^{2FO} &= \exists t_1, t_2. (R_{q^T}^{\text{coal}}(\mathbf{d}, t_1, t_2) \wedge t_1 \leq u^- \wedge u^+ \leq t_2), \\ [u^* < v^*]^{2FO} &= u^* < v^*, \\ [\neg\psi]^{2FO} &= \neg[\psi]^{2FO}, \\ [\psi_1 \wedge \psi_2]^{2FO} &= [\psi_1]^{2FO} \wedge [\psi_2]^{2FO}, \\ [\exists y.\psi]^{2FO} &= \exists y^-, y^+. (y^- \leq y^+ \wedge [\psi]^{2FO}), \end{aligned}$$

where the involved syntactic abbreviations are as follows:

$$\begin{aligned} R_{q^T}^{\text{coal}}(\mathbf{d}, u, v) &\triangleq \exists t_1, t_2. R_{q^T}(\mathbf{d}, u, t_1) \wedge R_{q^T}(\mathbf{d}, t_2, v) \wedge \\ &\quad \neg\exists t_3, t_4. (R_{q^T}(\mathbf{d}, t_3, t_4) \wedge t_3 < u \wedge u \leq t_4) \wedge \\ &\quad \neg\exists t_3, t_4. (R_{q^T}(\mathbf{d}, t_3, t_4) \wedge t_3 \leq v \wedge v < t_4) \wedge \\ &\quad \neg\exists t_3, t_4. (R_{q^T}(\mathbf{d}, t_3, t_4) \wedge u < t_3 \wedge t_4 \leq v \wedge \\ &\quad \neg\exists t_5, t_6. (R_{q^T}(\mathbf{d}, t_5, t_6) \wedge t_5 < t_3 \wedge t_3 \leq t_6)) \end{aligned} \tag{1}$$

where for $q^T = \bigvee_{1 \leq i \leq m} q_i$:

$$R_{q^T}(\mathbf{d}, u, v) \triangleq \bigvee_{1 \leq i \leq m} R_{q_i}(\mathbf{d}, u, v) \tag{2}$$

and for every $q_i(\mathbf{d}) = \exists \mathbf{x}. (\bigwedge_{1 \leq j \leq n} \alpha_j(\mathbf{d}_j))$:

$$R_{q_i}(\mathbf{d}, u, v) \triangleq \exists t_1, \dots, t_{2n}. \exists \mathbf{x}. \bigwedge_{1 \leq j \leq n} (R_{\alpha_j}(\mathbf{d}_j, t_j, t_{n+j})) \wedge u = \max(t_1, \dots, t_n) \wedge v = \min(t_{n+1}, \dots, t_{2n}) \wedge u \leq v \quad (3)$$

where:

$$\begin{aligned} u = \max(t_1, \dots, t_n) &\triangleq \bigwedge_{1 \leq i \leq n} ((\bigwedge_{1 \leq j \leq n} t_j \leq t_i) \rightarrow u = t_i) \\ v = \min(t_1, \dots, t_n) &\triangleq \bigwedge_{1 \leq i \leq n} ((\bigwedge_{1 \leq j \leq n} t_i \leq t_j) \rightarrow v = t_i) \end{aligned} \quad (4)$$

The pivotal part of CTQ rewriting is the translation of the embedded CQs. In this respect, the proposed approach builds directly on the standard FO rewritings of CQs, which are obtainable via existing techniques [8]. Given an FO rewriting q^T of a CQ q , all atoms in q^T are temporalized (3) and further incorporated in special formula templates (1)-(4), in order to meet two key challenges inherent to querying concrete TDBs:

- computing *temporal joins*, i.e., identifying maximal time intervals over which conjunctions of atoms are satisfied,
- applying *coalescing*, i.e., merging overlapping and adjacent intervals for the (intermediate) query results.

To illustrate these issues and the roles played by the formula templates, consider the example presented in Figure 1, which addresses the following setup:

$$\begin{aligned} \text{CTA: } \mathfrak{A} &= \{[1, 7] : B(a), [1, 3] : C(a), [3, 10] : C(a), [6, 12] : D(a)\}, \\ \text{TBox: } \mathcal{T} &= \{D \sqsubseteq B\}, \\ \text{CTQ: } [q(a)](u) &, \text{ where } q(x) = B(x) \wedge C(x) \text{ and } u = [1, 10]. \end{aligned}$$

Under the assumed TBox, the FO rewriting of q can be formulated as $q^T(x) = q_1(x) \vee q_2(x)$, where $q_1(x) = B(x) \wedge C(x)$ and $q_2(x) = D(x) \wedge C(x)$. By Definition 8, the satisfaction of condition $tdb(\mathfrak{A}) \Vdash [q(a)](u)$ is equivalent to verifying whether $(a, t_1, t_2) \in (R_{q^T}^{coal})^{\mathcal{D}}$, for some $t_1, t_2 \in T$, such that $t_1 \leq u^- \leq u^+ \leq t_2$. Computing $(R_{q^T}^{coal})^{\mathcal{D}}$ can be conceptually divided into three consecutive phases.

Firstly, we compute temporal joins over the sets of ABox assertions entailing each CQ q_i , i.e., we identify all intervals τ , such that there exist a set of assertions $S = \{\tau_j : \alpha_j(\mathbf{d}_j) \in \mathfrak{A}\}$, where the atoms $\alpha_j(\mathbf{d}_j)$ in S provide the exact matches for the conjuncts of q_i , while $\bigcap_j \tau_j = \tau$. Relation R_{q_i} , defined via formula template (3), augmented with (4), which fixes an FO formalization of the functions max and min, selects exactly such intervals for each answer to each q_i . In our example, R_{q_1} contains tuples $(a, 1, 3)$ and $(a, 3, 7)$, while R_{q_2} tuple $(a, 6, 10)$. In the second phase, captured by the template (2), all such tuples become automatically instances of the common relation R_{q^T} . Finally, R_{q^T} is coalesced, resulting in the relation $R_{q^T}^{coal}$. This phase is executed by means of the template (1), which is based on a straightforward FO formalization of the coalescing mechanism, known in the context of TDBs [6]. Technically, for

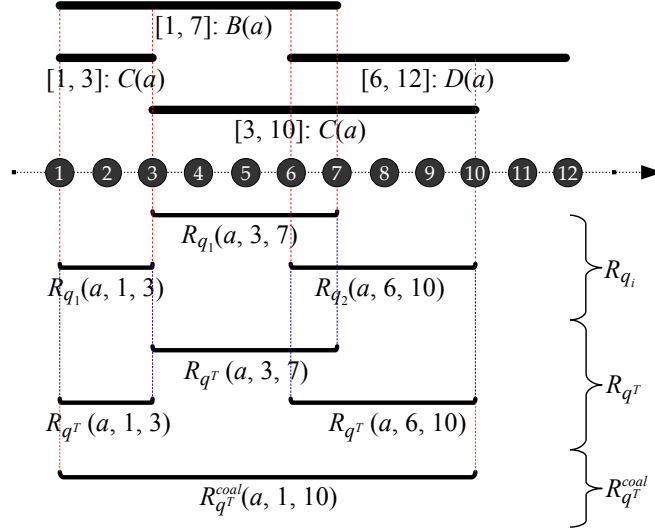


Fig. 1. Computing R_{q_i} , R_{q^T} and $R_{q^T}^{coal}$, for $q(x) = B(x) \wedge C(x)$ and $T = \{D \sqsubseteq B\}$.

each tuple \mathbf{a} comprising an answer to q , the relation $R_{q^T}^{coal}$ selects the minimal and maximal time points, such that for every point t between the two there exists a tuple $(\mathbf{a}, t_1, t_2) \in (R_{q^T})^D$ with $t_1 \leq t \leq t_2$. In the presented example, $R_{q^T}^{coal}$ contains exactly one tuple $(\mathbf{a}, 1, 10)$, which is the result of coalescing tuples $(\mathbf{a}, 1, 3)$, $(\mathbf{a}, 3, 7)$ and $(\mathbf{a}, 6, 10)$ in R_{q^T} . Consequently, the CTQ $[q(\mathbf{a})](u)$ evaluates to true over \mathcal{T} and \mathfrak{A} .

Since 2FO queries, in the shape defined above, can be naturally considered as formulas of the relational calculus, the final step of restating them into the SQL syntax can be carried out following broadly adopted translation strategies [17]. In practice, to ensure appropriate quantification over the time domain in RDBMSs, one would usually need to include it as an explicitly represented (and stored in the TDB) monadic relation R_T , such that $R_T^D = T$, which should be further introduced in the translation of the temporal quantifiers from the Definition 1, as follows:

$$\lceil \exists y. \psi \rceil^{2FO} = \exists y^-, y^+. (R_T(y^-) \wedge R_T(y^+) \wedge y^- \leq y^+ \wedge \lceil \psi \rceil^{2FO}).$$

Accordingly, every free temporal variable $x \in \mathbb{T}_V$ occurring in a 2FO query should be additionally guarded by the restriction $R_T(x)$ in the resulting translation, in order to be properly handled by the substitution mechanisms implemented in typical RDBMSs.

Agreeably, even though the presented rewriting is mathematically correct, as we demonstrate in the following part, it can be expected to be suboptimal in terms of the query answering efficiency. For instance, the naive coalescing mechanism, captured by the formula template (1), is known to be highly inefficient and can be substantially improved using more sophisticated approaches

[25]. The usefulness of this and other potential optimizations, on which we also shortly remark at the end of the next section, can be ultimately assessed only on the grounds of empirical evaluation, which is outside the scope of the current work.

5 Query Answering via 2FO Rewriting

The next theorem ensures correctness of CTQ answering via the 2FO rewriting.

Theorem 1 (Correctness of 2FO rewriting). *For every TBox \mathcal{T} , CTA \mathfrak{A} , CTQ ψ , and an answer σ to ψ , it holds that:*

$$\mathcal{T}, \mathfrak{A} \models \sigma(\psi) \text{ iff } \text{tdb}(\mathfrak{A}) \Vdash [\sigma(\psi)]^{2FO}.$$

The proof, presented in the appendix, follows by structural induction over the syntactic cases addressed in the rewriting rules. In the technically most demanding case of $[[q(\mathbf{d})](u)]^{2FO}$ we essentially formalize the discussion from the previous section. The remaining ones are largely straightforward.

Importantly, the definitions of CTQs and their 2FO rewritings, guarantee that query answering remains insensitive to the particular ways the CTAs might encode temporal data, as long as the data is semantically equivalent. This result is due to the fact that the semantics of CTQs is defined in terms of the underlying ATAs, hence whenever $\|\mathfrak{A}\| = \|\mathfrak{A}'\|$, for any two CTAs \mathfrak{A} and \mathfrak{A}' , the answers over them must always coincide.

Theorem 2 (CTQs are $\|\cdot\|$ -generic). *For every two CTAs \mathfrak{A} and \mathfrak{A}' over a time domain $\mathfrak{T} = (T, <)$ such that $\|\mathfrak{A}\| = \|\mathfrak{A}'\|$, a TBox \mathcal{T} , a CTQ ψ and an answer σ to ψ , it holds that:*

1. $\mathcal{T}, \mathfrak{A} \models \sigma(\psi)$ iff $\mathcal{T}, \mathfrak{A}' \models \sigma(\psi)$,
2. $\text{tdb}(\mathfrak{A}) \Vdash [\sigma(\psi)]^{2FO}$ iff $\text{tdb}(\mathfrak{A}') \Vdash [\sigma(\psi)]^{2FO}$.

Further, we address the complexity and algorithmic aspects of the CTQ entailment. We start by observing that, similarly as in the case of CQ rewriting, 2FO rewritings of CTQs can be in principle exponential in the size of the original queries. In fact, as the next proposition shows, this exponential blow-up is exclusively due to the embedded CQs, as the extra temporal layer itself adds only linearly to the overall size of the resulting 2FO formulas. This observation follows directly by scrutinizing the rewriting rules from Definition 8.

Proposition 1 (Size of rewriting). *Let \mathcal{T} be a TBox, and ψ a CTQ with q_1, \dots, q_n being all the CQs occurring in ψ . Then the size of $[\psi]^{2FO}$ is linear in the joint size of ψ and the FO rewritings q_1^T, \dots, q_n^T of the CQs.*

Clearly, the fragment of two-sorted first-order logic used for the CTQ rewriting is as expressive as FOL. At the same time, it is obviously not more expressive than that, as one can apply a commonly known, linear reduction, involving the

Algorithm 1 Incremental computation of relations $R_{q^T}^{coal}$.

Input: CTA \mathfrak{A} , CTQ ψ

Output: A database $\mathcal{R} = (\mathbb{N}_1, T, <, \cdot^{\mathcal{D}})$ over the signature $\{R_{q^T}^{coal} \mid q \in \text{cq}(\psi)\}$

```

1: for all  $q \in \text{cq}(\psi)$  do
2:   for all  $q_i(\mathbf{a}) \in \text{cq}(q^T)$  do
3:     for all  $t_1, t_2 \in T$  such that  $\text{tdb}(\mathfrak{A}) \Vdash R_{q_i}(\mathbf{a}, t_1, t_2)$  do
4:        $(R_{q^T}^{coal})^{\mathcal{D}} := (R_{q^T}^{coal})^{\mathcal{D}} \cup \{(\mathbf{a}, t_1, t_2)\}$ 
5:       while applicable do  $\{\text{coalescing}\}$ 
6:         for all  $(\mathbf{a}, t_1, t_2), (\mathbf{a}, t_3, t_4) \in (R_{q^T}^{coal})^{\mathcal{D}}$  do
7:           if  $[t_1, t_2] \cap [t_3, t_4] \neq \emptyset$  then
8:              $(R_{q^T}^{coal})^{\mathcal{D}} := (R_{q^T}^{coal})^{\mathcal{D}} \setminus \{(\mathbf{a}, t_1, t_2), (\mathbf{a}, t_3, t_4)\}$ 
9:              $(R_{q^T}^{coal})^{\mathcal{D}} := (R_{q^T}^{coal})^{\mathcal{D}} \cup \{(\mathbf{a}, \min(t_1, t_3), \max(t_2, t_4))\}$ 
10:          end if
11:        end for
12:      end while
13:    end for
14:  end for
15: end for

```

introduction of two designated predicates for representing the respective domains, which are used to guard the scopes of the two sorts of quantifiers (*cf.* the last paragraph of the previous section). Based on these observations, we obtain two results concerning the complexity of CTQ query entailment.

Corollary 1 (Data complexity). *The data complexity of CTQ entailment in CTAs in the presence of TBoxes, over finite time domains, is in AC^0 .*

The restriction to finite time domains in this and the following case is a natural strategy of ensuring that the first-order structures, over which queries are evaluated, are indeed finite and, consequently, that model-checking can be effectively performed over them. Note, that in terms of data complexity, deciding CTQ entailment via 2FO rewriting remains precisely as hard as deciding CQ entailment via FO rewriting. The combined complexity, however, increases from NP- to PSPACE-complete.

Theorem 3 (Combined complexity). *The combined complexity of CTQ entailment over CTAs in the presence of TBoxes, over finite time domains, is PSPACE-complete.*

The result transfers from the entailment problem for boolean FO queries over relational databases, which is known to be PSPACE-complete [10]. However, due to the exponential blow-up in the size of the 2FO rewritings, explained in Proposition 1, answering CTQs in PSPACE requires a somewhat more sophisticated strategy than just a straightforward evaluation of the rewritten queries over a TDB. In the approach described in the following proof, we decide the entailment of a grounded CTQ ψ , by first incrementally computing the interpretations of

relations $R_{q^T}^{coal}$, for each CQ q embedded in ψ , and then evaluate over them a restricted 2FO rewriting of ψ of at most polynomial size.

Proof. Let $\mathbf{cq}(\psi)$ denote all CQs occurring in ψ and grounded as in ψ , and let $\mathbf{cq}(q^T)$, for every $q \in \mathbf{cq}(\psi)$, be the set of all CQs comprising the FO rewriting q^T . With every such $q_i(\mathbf{a}) \in \mathbf{cq}(q^T)$, we associate the (query) formula $R_{q_i}(\mathbf{a}, y_1, y_2)$ defined via the templates (3) and (4). Note that every such formula is linear in the size of ψ . Algorithm 1 constructs a TDB $\mathcal{R} = (\mathbf{N}_I, T, <, \cdot^{\mathcal{D}})$ over the signature $\{R_{q^T}^{coal} \mid q \in \mathbf{cq}(\psi)\}$, by collecting and coalescing all matches to the formulas $R_{q_i}(\mathbf{a}, y_1, y_2)$ over $tdb(\mathfrak{A})$. Observe that by applying coalescing on $R_{q^T}^{coal}$ every time another tuple is added, we guarantee that the size of each $(R_{q^T}^{coal})^{\mathcal{D}}$, and therefore of \mathcal{R} , is linear in the size of T , and in fact cannot be larger than \mathfrak{A} . Therefore the algorithm runs in PSPACE. Once \mathcal{R} is computed, we decide $\mathcal{R} \models [\psi]_R^{2FO}$, where $[\cdot]_R^{2FO}$ is a rewriting specified exactly as $[\cdot]^{2FO}$ in Definition 8, but without employing any of the formula templates (1)-(4), and instead, retaining all $R_{q^T}^{coal}$ as actual predicates in the resulting 2FO formulas. Clearly, $[\psi]_R^{2FO}$ is linear in the size of ψ . Since deciding $\mathcal{R} \models [\psi]_R^{2FO}$ is a variant of the (PSPACE-hard [10]) entailment problem for boolean FO queries over relational databases (observe, that $[\cdot]_R^{2FO}$ permits n -ary predicates, for $n \geq 1$), it follows that deciding the CTQ entailment over CTAs in the presence of TBoxes, over finite time domains, is PSPACE-complete. Note that since the interpretations of $R_{q^T}^{coal}$ in both \mathcal{R} and $tdb(\mathfrak{A})$ must coincide, the procedure preserves soundness and completeness of query answering. \square

The query answering algorithm implied by the above decision procedure paves the way towards efficient practical implementations, based on the well-known database technique of *materialized view maintenance* [20], particularly in the context of large or often changing TDBs. By maintaining the intermediate views containing the answers to the embedded CQs, the approach should facilitate more localized and fine-grained updating of the answers in response to database updates and query refinements (be it on the CQ or temporal level), as well as forms of incremental, anytime query answering.

6 Related Work

This work is naturally related to the research on TDBs, conducted largely in the 1990s [11,12]. Although the advancements in that area provide some theoretical grounding for our proposal, they are obviously agnostic about the ontology-based approach on the problem, which we concentrate on here.

The research on temporal extensions to OBDA has been taken up only very recently by Borgwardt et al. [3,7] and Artale et al. [2]. In [3], the authors study the same prototypical scenario as addressed here, but focus on its more foundational aspects. They consider a more expressive DL \mathcal{ALC} as the background ontology language and adopt a less restrictive definition of temporal queries. This offers a richer setting, yet without apparent application prospects in the context of existing TDBs. The work in [7] is closer aligned with ours. It also

considers *DL-Lite* ontologies and studies a rewriting approach based on the use of standard CQ rewriting techniques. However, the authors define their query language without involving the epistemic interpretation of embedded CQs and consider only its positive fragment. It is not difficult to show that under OWL 2 QL, such design leads to a query language strictly less expressive from the one proposed here. Observe, that due to the condition (\dagger), in Definition 5, the class of negation-free CTQs preserves the semantics regardless of the use of the epistemic interpretation of CQs, and so it coincides with the fragment considered in [7]. Furthermore, the authors focus exclusively on abstract TDBs and do not address the problem of direct querying of concrete representations, which is the main goal of our work. In [2], the authors study an orthogonal approach to temporal OBDA. Instead of adding temporal features on the query level, they propose temporal extensions to OWL 2 QL, in the spirit of temporal DLs [1].

A number of other frameworks have been developed for managing the validity time of Semantic Web data represented natively in RDF(S)/OWL languages [15,14,5,21]. These proposals can be seen as parallel to ours, in that they address variants of the same problem, but focus on the use of dedicated Semantic Web technologies, such as SPARQL and RDF triplestores, instead of involving a semantic view on the data managed within traditional RDBMSs. This rises interesting questions about the formal correspondences between the employed reasoning methods, regardless of their practical implementations. Unfortunately, most of these approaches are strongly technology-driven and often fall short of indicating their links to the logic foundations of temporal querying.

7 Conclusions

In this work, we have proposed a principled, yet practical approach to lifting the popular paradigm of OBDA to temporal case. The presented language TQL allows for querying temporal data stored in SQL:2011-compliant databases via a semantic layer of OWL 2 QL ontologies. We believe that this proposal strikes a good balance between the theoretical strength of its formal foundations and the feasibility of practical applications, warranted by the possibility of answering TQL queries in commercially supported RDBMSs via a translation to SQL.

Considering the growing interest in temporal extensions of OBDA, it is critical to continue the formal study of temporal features supported by the SQL standards, temporal extensions to logic-based query languages and ontology languages intended for use in practical OBDA scenarios, and finally, the relationships holding between all of them. In this respect, as part of future research, it is necessary to establish stronger links between the approach pursued here and those proposed recently in [3,7,2], with the prospect of gaining a clearer view on the landscape of technical possibilities regarding the temporal OBDA. In particular, the scenario of querying temporal data, as considered here and in [3,7], under temporalized ontologies, such as introduced by Artale et al. [2], deserves further attention and in-depth study.

References

1. Artale, A., Kontchakov, R., Lutz, C., Wolter, F., Zakharyashev, M.: Temporalising tractable description logics. In: Proceedings of the Fourteenth International Symposium on Temporal Representation and Reasoning (2007)
2. Artale, A., Kontchakov, R., Wolter, F., Zakharyashev, M.: Temporal description logic for ontology-based data access. In: Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-13) (2013)
3. Baader, F., Borgwardt, S., Lippmann, M.: Temporalizing ontology-based data access. In: Proceedings of the International Conference on Automated Deduction (CADE-13) (2013)
4. Baader, F., Calvanese, D., McGuinness, D.L., Nardi, D., Patel-Schneider, P.F.: The description logic handbook: theory, implementation, and applications. Cambridge University Press (2003)
5. Batsakis, S., Stravoskoufos, K., Petrakis, E.G.M.: Temporal Reasoning for Supporting Temporal Queries in OWL 2.0. In: Proc. of the Int. Conf. on Knowledge-based and Intelligent Information and Engineering Systems (KES-11) (2011)
6. Böhlen, M.H., Snodgrass, R.T., Soo, M.D.: Coalescing in temporal databases. *IEEE Computer* 19, 35–42 (1996)
7. Borgwardt, S., Lippmann, M., Thost, V.: Temporal query answering in the description logic DL-Lite. In: Proc. of the International Symposium on Frontiers of Combining Systems (FroCoS-13) (2013)
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., Giacomo, G.D., Lembo, D., Lenzerini, M., Rosati, R.: Eql-lite: Effective first-order query processing in description logics. In: Proc. of IJCAI-07 (2007)
10. Chandra, A.K., Merlin, P.M.: Optimal implementation of conjunctive queries in relational data bases. In: Proc. of the ACM Symposium on Theory of Computing (STOC-77) (1977)
11. Chomicki, J.: Temporal query languages: A survey. In: Proc. of the International Conference on Temporal Logic (ICTL-94) (1994)
12. Chomicki, J., Toman, D.: Temporal Databases. In: Handbook of Temporal Reasoning in Artificial Intelligence (Foundations of Artificial Intelligence). pp. 429–468. Elsevier Science Inc. (2005)
13. Glimm, B., Horrocks, I., Lutz, C., Sattler, U.: Conjunctive query answering for the description logic SHIQ. In: *Journal of Artificial Intelligence Research* (2008)
14. Grandi, F.: T-SPARQL: a TSQL2-like temporal query language for RDF. In: Proc. of the International Workshop on Querying Graph Structured Data (2010)
15. Gutierrez, C., Hurtado, C.A., Vaisman, A.A.: Introducing time into RDF. *IEEE Transactions on Knowledge and Data Engineering* 19(2), 207–218 (2007)
16. Gutiérrez-Basulto, V., Klarman, S.: Towards a unifying approach to representing and querying temporal data in description logics. In: Proc. of the Int. Conf. on Web Reasoning and Rule Systems (RR-12) (2012)
17. Kawash, J.: Complex quantification in Structured Query Language (SQL): A tutorial using relational calculus. *Journal of Computers in Mathematics and Science Teaching* 23(2), 169–190 (2004)
18. Klarman, S., Meyer, T.: Prediction and explanation over DL-Lite data streams. In: Proc. of the International Conference on Logic for Programming, Artificial Intelligence and Reasoning (LPAR-19) (2013)

19. Kulkarni, K., Michels, J.E.: Temporal features in SQL:2011. *SIGMOD Rec.* 41(3) (2012)
20. Mohania, M., Konomi, S., Kambayashi, Y.: Incremental maintenance of materialized views. In: *Database and Expert Systems Applications*, pp. 551–560. Springer Berlin Heidelberg (1997)
21. Motik, B.: Representing and querying validity time in RDF and OWL: A logic-based approach. *Journal of Web Semantics: Science, Services and Agents on the World Wide Web* 57(5), 1–62 (2012)
22. Reynolds, M.: The complexity of decision problems for linear temporal logics. *Journal of Studies in Logic* 3(1) (2010)
23. Snodgrass, R.T., Böhlen, M.H., Jensen, C.S., Steiner, A.: Transitioning temporal support in TSQL2 to SQL3. In: Etzion, O., Jajodia, S., Sripada, S. (eds.) *Temporal Databases: Research and Practice*, LNCS, vol. 1399, pp. 150–194 (1998)
24. Toman, D.: Point vs. interval-based query languages for temporal databases. In: *Proc. of the Symposium on Principles of Database Systems (PODS-96)* (1996)
25. Zhou, X., Wang, F., Zaniolo, C.: Efficient temporal coalescing query support in relational database systems. In: *Proc. of the International Conference on Database and Expert Systems Applications (DEXA-06)* (2006)

Appendix

Below we present the proofs omitted in the paper.

Proof of Theorem 1. Suppose $\mathfrak{A}, \mathcal{T}, \pi \models \sigma(\psi)$, for a CTA \mathfrak{A} , TBox \mathcal{T} , i -substitution π , CTQ ψ and an answer σ . We show that $tdb(\mathfrak{A}), \delta, \nu \Vdash [\sigma(\psi)]^{2FO}$, for some d -/ t -substitutions δ and ν . For simplicity, we assume ψ is already grounded with σ , and we consequently omit explicit references to σ in the proof.

To start with, w.l.o.g., we assume that no variable is used more than once in $\sigma(\psi)$, and similarly in $[\sigma(\psi)]$. Given that we can assume that $\pi[y \mapsto \tau] = \pi$, for every $y \in \mathbb{I}_V$, $\tau \in I$, and $\nu[y \mapsto t] = \nu$, for every $y \in \mathbb{T}_V$, $t \in T$. Then for every $y \in \mathbb{I}_V$ and $* \in \{-, +\}$, we fix $\nu(y^*) = \pi(y)^*$. Further, we introduce some auxiliary nomenclature. Let $q = \exists \mathbf{x}. (\bigwedge_{1 \leq j \leq n} \alpha_j(\mathbf{d}_j))$ be a grounded CQ. A

minimal set of support for q is an ABox \mathcal{A} such that there exists a mapping $\mu : \text{term}(q) \mapsto \mathbb{N}_I$, called the *support substitution*, where $\mu(a) = a$ for every $a \in \mathbb{N}_I$, such that $\mathcal{A} = \{\alpha_j(\mu(\mathbf{d}_j)) \mid 1 \leq j \leq n\}$, where by $\mu(\mathbf{d}_j)$ we denote the result of applying μ to every term in the sequence \mathbf{d}_j . Clearly, $\mathcal{A} \models q$. We next extend this notion towards an FO rewriting of q in a TBox \mathcal{T} . A *minimal set of support* for q^T is a minimal set of support for one of the disjuncts q_i in q^T .

The proof follows by structural induction over the syntactic cases covered in the translation rules.

- Suppose $\mathfrak{A}, \mathcal{T}, \pi \models [q(\mathbf{a})](u)$. Then, by Definition 5, $\mathcal{T}, \mathfrak{A}_t \models q(\mathbf{a})$, for every $t \in \pi(u)$. Consequently, by Definition 1, $db(\mathfrak{A}_t) \Vdash q^T(\mathbf{a})$, for every $t \in \pi(u)$. Let us fix one such arbitrary $t \in \pi(u)$. By construction of $db(\mathfrak{A}_t)$, it follows that there must exist a minimal set of support $\mathcal{A} \subseteq \mathfrak{A}_t$, with an associated support substitution μ , for some disjunct $q_i(\mathbf{a})$ of $q^T(\mathbf{a})$. Then, by the construction of \mathfrak{A}_t , for every $\alpha(\mathbf{d}) \in \mathcal{A}$, there must exist $\tau : \alpha(\mathbf{d}) \in \mathfrak{A}$, and consequently $R_\alpha(\mathbf{d}, \tau^-, \tau^+) \in tdb(\mathfrak{A})$, such that $t \in \tau$. Then there must exist $t', t'' \in T$, and d -/ t -substitutions δ, ν , such that $tdb(\mathfrak{A}), \delta, \nu \Vdash \bigwedge_{1 \leq j \leq n} (R_{\alpha_j}(\mathbf{d}_j, t_j, t_{n+j})) \wedge t' = \max(t_1, \dots, t_n) \wedge t'' = \min(t_{n+1}, \dots, t_{2n}) \wedge t' \leq t''$. In particular, this must hold for δ , such that $\delta(d) = \mu(d)$, for every $d \in \text{term}(q_i)$. Then it must be also the case that $tdb(\mathfrak{A}), \delta, \nu \Vdash R_{q_i}(\mathbf{a}, t', t'')$, given it is defined via formula templates (3), (4) in Definition 8. Consequently, $R_{q^T}(\mathbf{a}, t', t'')$, as defined via formula template (2) in Definition 8, must be also satisfied over $tdb(\mathfrak{A})$. The same argument can be repeated for every $t \in \pi(u)$, and possibly involving other q_i in q^T . This means that the relation R_{q^T} can be eventually coalesced over the whole interval $\pi(u)$, as all tuples in R_{q^T} , referred to above, must be clearly adjacent or overlapping. Formally then, there must exist $v', v'' \in T$, such that $v' \leq \pi(u)^-$ and $\pi(u)^+ \leq v''$, and some d -substitution δ' , agreeing with the support substitutions for each involved $q_i(\mathbf{a})$, for which $tdb(\mathfrak{A}), \delta', \nu \Vdash R_{q^T}^{coal}(\mathbf{a}, v', v'')$, as defined by formula template (1) in Definition 8. From this it follows that $tdb(\mathfrak{A}), \delta', \nu \Vdash \exists v', v''. (R_{q^T}^{coal}(\mathbf{a}, v', v'') \wedge v' \leq u^- \wedge u^+ \leq v'')$, which concludes the argument.

- Suppose $\mathfrak{A}, \mathcal{T}, \pi \models u^* < v^*$. It follows immediately, that $tdb(\mathfrak{A}), \delta, \nu \Vdash u^* < v^*$, for an arbitrary δ and ν restricted as in the beginning of the proof.
- Suppose $\mathfrak{A}, \mathcal{T}, \pi \models \neg\psi$, and so that it is not the case that $\mathfrak{A}, \mathcal{T}, \pi \models \psi$. By inductive hypothesis we assume that it is also not the case that $tdb(\mathfrak{A}), \delta, \nu \Vdash [\psi]^{2FO}$, for some δ and ν . Then it follows immediately that $tdb(\mathfrak{A}), \delta, \nu \Vdash \neg[\psi]^{2FO}$.
- Suppose $\mathfrak{A}, \mathcal{T}, \pi \models \psi_1 \wedge \psi_2$, and so that $\mathfrak{A}, \mathcal{T}, \pi \models \psi_1$ and $\mathfrak{A}, \mathcal{T}, \pi \models \psi_2$. By inductive hypothesis we assume that $tdb(\mathfrak{A}), \delta, \nu \Vdash [\psi_1]^{2FO}$ and $tdb(\mathfrak{A}), \delta, \nu \Vdash [\psi_2]^{2FO}$, for some δ and ν . Then it follows immediately that $tdb(\mathfrak{A}), \delta, \nu \Vdash [\psi_1]^{2FO} \wedge [\psi_2]^{2FO}$.
- Suppose $\mathfrak{A}, \mathcal{T}, \pi \models \exists y.\psi$, and so that $\mathfrak{A}, \mathcal{T}, \pi \models \psi$, given the restriction on the use of variables in ψ explained in the beginning of the proof. By inductive hypothesis we assume that $tdb(\mathfrak{A}), \delta, \nu \Vdash [\psi]^{2FO}$. But then it must be the case that $tdb(\mathfrak{A}), \delta, \nu \Vdash \exists y^-, y^+. (y^- \leq y^+ \wedge [\psi]^{2FO})$, as ν is restricted as in the beginning of the proof. Observe that an interval variable y is translated into a pair of time-point variables y^-, y^+ , with an intended order imposed by the condition $y^- < y^+$.

The inverse direction follows by reverting the steps of the argument above. \square

Proof of Theorem 2. Consider two CTAs \mathfrak{A} and \mathfrak{A}' over a time domain $\mathfrak{T} = (T, <)$ such that $\|\mathfrak{A}\| = \|\mathfrak{A}'\|$, a TBox \mathcal{T} , a CTQ ψ and an answer σ to ψ . Suppose that $\mathcal{T}, \mathfrak{A}, \pi \models \sigma(\psi)$ for some i -substitution π . Observe that the semantics of TQL, given in Definition 5, is expressed purely in terms of the ATA $\|\mathfrak{A}\|$. Since $\|\mathfrak{A}\| = \|\mathfrak{A}'\|$ then for every subformula φ of $\sigma(\psi)$ it must be the case that $\mathcal{T}, \mathfrak{A}, \pi \models \varphi$ iff $\mathcal{T}, \mathfrak{A}', \pi \models \varphi$. Consequently, it follows that $\mathcal{T}, \mathfrak{A}', \pi \models \sigma(\psi)$. The argument clearly works in the same way in the opposite direction. Further, by Theorem 1, it holds that $\mathcal{T}, \mathfrak{A} \models \sigma(\psi)$ iff $tdb(\mathfrak{A}) \Vdash [\sigma(\psi)]^{2FO}$. By the previous claim, this implies that $tdb(\mathfrak{A}) \Vdash [\sigma(\psi)]^{2FO}$ iff $tdb(\mathfrak{A}') \Vdash [\sigma(\psi)]^{2FO}$. \square

Proof of Proposition 1. It is easy to see that both the basic rewriting rules in Definition 8, as well as the involved formula templates (1), (3) and (4), add at most linearly to the size of the original CTQ. Only template (2) can possibly introduce an exponential blow-up of the rewriting. This effect is exclusively due to the fact that the standard FO rewriting of a CQ q , handled in the template, might consist of exponentially many CQs. \square