A Temporal Description Logic for Reasoning over Conceptual Schemas and Queries

Alessandro Artale[1], Enrico Franconi[2], Frank Wolter[3], Michael Zakharyaschev[4]

[1] Dept. of Computation, UMIST, Manchester, UK; artale@co.umist.ac.uk
[2] Dept. of Computer Science, Univ. of Manchester, UK; franconi@cs.man.ac.uk
[3] Inst. für Informatik, Univ. of Leipzig, D; wolter@informatik.uni-leipzig.de
[4] Dept. of Computer Science, Kings College, London, UK; mz@dcs.kcl.ac.uk

Abstract. This paper introduces a new logical formalism, intended for temporal conceptual modelling, as a natural combination of the well-known description logic $\mathcal{DLR}$ and point-based linear temporal logic with *Since* and *Until*. The expressive power of the resulting $\mathcal{DLR}_{\mathcal{US}}$ logic is illustrated by providing a characterisation of the most important temporal conceptual modelling constructs appeared in the literature. We define a query language (where queries are non-recursive Datalog programs and atoms are complex $\mathcal{DLR}_{\mathcal{US}}$ expressions) and investigate the problem of checking query containment under the constraints defined by $\mathcal{DLR}_{\mathcal{US}}$ conceptual schemas—i.e., $\mathcal{DLR}_{\mathcal{US}}$ knowledge bases—as well as the problems of schema satisfiability and logical implication. Although it is shown that reasoning in full $\mathcal{DLR}_{\mathcal{US}}$ is undecidable, we identify the decidable (in a sense, maximal) fragment $\mathcal{DLR}^{-}_{\mathcal{US}}$ by allowing applications of temporal operators to formulas and class expressions only (but not to relation expressions). We obtain the following hierarchy of complexity results: (a) reasoning in $\mathcal{DLR}^{-}_{\mathcal{US}}$ with atomic formulas is EXPTIME-complete, (b) satisfiability and logical implication of arbitrary $\mathcal{DLR}^{-}_{\mathcal{US}}$ formulas is EXPSPACE-complete, and (c) the problem of checking query containment of non-recursive Datalog queries under $\mathcal{DLR}^{-}_{\mathcal{US}}$ constraints is decidable in 2EXPTIME.

1 Introduction

Temporal information systems are information systems that store historical information, i.e., past, present, and potential future data [33]. Many formalisations have been proposed for temporal information systems which are based on first-order temporal logic [17]. Although these formalisations can be very useful for characterising semantical problems arising in temporalised ontologies and in temporal databases, like conceptual modelling or querying, usually they are computationally unfeasible for performing deduction tasks (for example, logical implication in the first-order temporal logic of the flow of time $\langle \mathbb{Z}, < \rangle$ or $\langle \mathbb{N}, < \rangle$ is not even recursively enumerable). Note that we are interested in deduction rather than model checking. An obvious solution to this problem would be to look for well-behaved fragments of first-order temporal logic (see e.g. [17] and references therein); however this way has not been successful—the only promising approach we know of is the recent paper [28]. Another idea is to deviate from the first-order paradigm and start from computationally more friendly languages such as description logics which have been used in the area of non-temporal information management to characterise in a uniform framework both conceptual modelling and queries [34, 12, 9, 10].

The *temporal description logic* $\mathcal{DLR}_{\mathcal{US}}$ we devise in this paper is based on the expressive and decidable description logic $\mathcal{DLR}$ which allows the logical reconstruction and the extension of representational tools such as object-oriented data models (e.g., class diagrams in UML and ODMG), semantic data models (e.g., extended entity-relationship, EER), frame-based and web ontology languages (e.g., OIL and DAML+OIL) [13, 14, 19, 11, 31]. In this setting, an interesting feature of $\mathcal{DLR}$ is the ability to completely *define* classes and relations as $\mathcal{DLR}$ views over other classes and relations of the conceptual schema. Moreover, $\mathcal{DLR}$ formulas can express a large class of integrity constraints that are typical in databases, for instance, existence dependencies, exclusion dependencies, typed inclusion dependencies without projection of relations, unary inclusion dependencies, full key dependencies [12]. Logical implication in $\mathcal{DLR}$ is EXPTIME-complete [12]; practical correct and complete algorithms exist [30], used in in conceptual modelling applications [32, 21].

$\mathcal{DLR}$ is not only a very powerful language for conceptual modelling. The problem of view-based query processing under $\mathcal{DLR}$ constraints has also been studied [12]. View-based query answering requires to answer a query over a virtual database (constrained by a $\mathcal{DLR}$ theory playing the role of the conceptual schema and of the integrity constraints) for which the only information comes from a set of materialised views over the same database; this problem with non-recursive Datalog queries and views is a co-NP-complete problem (in data complexity) under the closed world assumption. Checking query containment of non-recursive Datalog queries under $\mathcal{DLR}$ constraints is decidable in 2EXPTIME [12, 29].

Given all these nice features of $\mathcal{DLR}$, it is natural to try to extend it with a temporal dimension, to understand the expressive power of the resulting hybrid with respect to the needs of temporal conceptual modelling and view based query processing, and to investigate its computational properties. This paper reports the results of such an attempt. We construct $\mathcal{DLR}_{\mathcal{US}}$ as an organic combination of $\mathcal{DLR}$ and the propositional linear temporal logic with *Since* and *Until* [40, 23] (which usually serves as the temporal component in the first-order approach) by allowing applications of temporal operators to all syntactical terms of $\mathcal{DLR}$: classes, relations, and formulas. We then investigate computational properties of reasoning with $\mathcal{DLR}_{\mathcal{US}}$ by analysing schema, class, and relation satisfiability, logical implication, and query containment for non-recursive Datalog queries under $\mathcal{DLR}_{\mathcal{US}}$ constraints.

The full $\mathcal{DLR}_{\mathcal{US}}$ turns out to be undecidable. The main reason for this is the possibility to postulate that a binary relation does not vary in time—a very small fragment of $\mathcal{DLR}_{\mathcal{US}}$ (say, $\mathcal{DLR}$ augmented with a single time invariant binary relation) can encode the undecidable tiling problem (cf. [47, 28]). The fragment $\mathcal{DLR}^{-}_{\mathcal{US}}$ of $\mathcal{DLR}_{\mathcal{US}}$ deprived of the ability to talk about temporal persistence of n-ary relations, for $n \geq 2$, is still very expressive, as is illustrated by examples in this paper, but its computational behaviour is much better. We obtain the following non-trivial novel complexity results: (1) reasoning in $\mathcal{DLR}^{-}_{\mathcal{US}}$ with atomic formulas is EXPTIME-complete, (2) satisfiability and logical implication of arbitrary $\mathcal{DLR}^{-}_{\mathcal{US}}$ formulas is EXPSPACE-complete, and (3) the problem of checking query containment of non-recursive Datalog queries under $\mathcal{DLR}^{-}_{\mathcal{US}}$ constraints is decidable in 2EXPTIME with an EXPSPACE lower bound.

The results obtained in this paper are novel for several reasons. Previous approaches to temporal description logics considered much weaker languages having only binary relations (i.e., roles), without the cardinality constructs, without the inverse construct (which $\mathcal{DLR}_{\mathcal{US}}$ implicitly is able to express by considering only binary relations), and without ever considering the ability to express queries [39, 46, 49, 42]. In this paper for the first time an upper bound for the complexity of reasoning in a temporal description logic with both future and past operators is proved, leading to a tight EXPSPACE-completeness result which automatically holds for the weaker basic temporal description logic $\mathcal{ALC}^{-}_{\mathcal{US}}$ as well. In this paper for the first time a decidability and a complexity result are found for the problem of temporal query containment under complex constraints. For a survey of the previous various approaches to temporal description logics see [4].

To illustrate the expressive power of $\mathcal{DLR}_{\mathcal{US}}$, we provide a brief semantic characterisation—by means of $\mathcal{DLR}_{\mathcal{US}}$ theories—of the most important modelling constructs (for the valid time representation) for temporal conceptual modelling. These can be applied to various temporally extended conceptual data models, such as, for example, entity-relationship and UML class diagrams. In this paper we extend the results presented in [3], where only the constructs for snapshot and temporary entities and relationships were proposed.

The paper is organised as follows. Section 2 introduces the syntax and the semantics of $\mathcal{DLR}_{\mathcal{US}}$ and of the query language, and provides a worked out example. In Section 3 decidability and complexity results are proved for satisfiability and logical implication in $\mathcal{DLR}_{\mathcal{US}}$ and $\mathcal{DLR}^{-}_{\mathcal{US}}$, as well as for the query containment problem. Section 4 discusses the expressivity of $\mathcal{DLR}_{\mathcal{US}}$ as a temporal conceptual model.

2 The temporal conceptual model

In this paper, we adopt the classical *snapshot* representation of abstract temporal databases (see e.g. [17]). The flow of time $\mathcal{T} = \langle \mathcal{T}_p, < \rangle$, where $\mathcal{T}_p$ is a set of time points (or chronons) and $<$ a binary precedence relation on $\mathcal{T}_p$, is assumed to be isomorphic to $\langle \mathbb{Z}, < \rangle$. Thus, a temporal database can be regarded as a mapping from time points in $\mathcal{T}$ to standard relational databases, with the same interpretation of constants and the same domains along time.

As a language for expressing temporal conceptual schemas we use the combination of the propositional temporal logic with *Since* and *Until* [23] and the (non-temporal) description logic $\mathcal{DLR}$ [12]. The resulting $\mathcal{DLR}_{\mathcal{US}}$ *temporal description logic* can be regarded as a rather expressive fragment of the first-order temporal logic $L^{\{\textbf{since, until}\}}$; cf. [17, 28] and Section 3 below; note that $\mathcal{DLR}$ itself is neither in the guarded fragment of FOL nor in the two variable variables fragment of FOL with counting quantifiers. The expressive capabilities of $\mathcal{DLR}_{\mathcal{US}}$ for temporal conceptual modelling are illustrated in Section 4. Section 2.2 extends $\mathcal{DLR}_{\mathcal{US}}$ with a query language.

2.1 Temporal conceptual schemas

The basic syntactical types of $\mathcal{DLR}_{\mathcal{US}}$ are *entities* (i.e., unary predicates, also known as *concepts* or *classes*) and n-ary *relations* of arity ≥ 2. Starting from a set of *atomic entities* (denoted by EN), a set RN of *atomic relations* (denoted by RN), and a set of *role symbols* (denoted by U) we define inductively (complex) entity and relation expressions as is shown in the upper part of Fig. 1, where the binary constructs ($\sqcap, \sqcup, \mathcal{U}, \mathcal{S}$) are applied to relations of the same arity, i, j, k, n are natural numbers, $i \leq n$, and j does not exceed the arity of R.

The non-temporal fragment of $\mathcal{DLR}_{\mathcal{US}}$ coincides with $\mathcal{DLR}$. For both entity and relation expressions all the Boolean constructs are available. The selection expression $U_i/n : E$ denotes an n-ary relation whose argument named U_i ($i \leq n$) is of type E; if it is clear from the context, we omit n and write $(U_i : E)$. The projection expression $\exists^{\lessgtr k}[U_j]R$ is a generalisation with cardinalities of the projection operator over the argument named U_j of the relation R; the plain classical projection is $\exists^{\geq 1}[U_j]R$. It is also possible to use the pure argument position version of the model by replacing role symbols U_i with the corresponding position numbers i.

The language of $\mathcal{DLR}_{\mathcal{US}}$ is interpreted in *temporal models* over $\mathcal{T}$, which are triples of the form $\mathcal{I} \doteq \langle \mathcal{T}, \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}(t)} \rangle$, where $\Delta^{\mathcal{I}}$ is non-empty set of objects (the *domain* of $\mathcal{I}$) and $\cdot^{\mathcal{I}(t)}$ an *interpretation function* such that, for every $t \in \mathcal{T}$, every entity E, and every n-ary relation R, we have $E^{\mathcal{I}(t)} \subseteq \Delta^{\mathcal{I}}$ and $R^{\mathcal{I}(t)} \subseteq (\Delta^{\mathcal{I}})^n$. The semantics of entity and relation expressions is defined in the lower part of Fig. 1, where $(u, v) = \{w \in \mathcal{T} \mid u < w < v\}$ and the operators $\Box^+$ (always in the future) and $\Box^-$ (always in the past) are the duals of $\Diamond^+$ (some time in the future) and $\Diamond^-$ (some time in the past), respectively, i.e., $\Box^+ E \equiv \neg\Diamond^+\neg E$ and $\Box^- E \equiv \neg\Diamond^-\neg E$, for both entities and relations. For entities,

$$
\begin{aligned}
R \to\ & \top_n \mid RN \mid \neg R \mid R_1 \sqcap R_2 \mid R_1 \sqcup R_2 \mid U_i/n : E \mid \\
& \Diamond^+ R \mid \Diamond^- R \mid \Box^+ R \mid \Box^- R \mid \oplus R \mid \ominus R \mid R_1\, \mathcal{U}\, R_2 \mid R_1\, \mathcal{S}\, R_2 \\
E \to\ & \top \mid EN \mid \neg E \mid E_1 \sqcap E_2 \mid E_1 \sqcup E_2 \mid \exists^{\leqslant k}[U_j]R \mid \\
& \Diamond^+ E \mid \Diamond^- E \mid \Box^+ E \mid \Box^- E \mid \oplus E \mid \ominus E \mid E_1\, \mathcal{U}\, E_2 \mid E_1\, \mathcal{S}\, E_2
\end{aligned}
$$

$$
\begin{aligned}
(\top_n)^{\mathcal{I}(t)} &\subseteq (\Delta^{\mathcal{I}})^n \\
RN^{\mathcal{I}(t)} &\subseteq (\top_n)^{\mathcal{I}(t)} \\
(\neg R)^{\mathcal{I}(t)} &= (\top_n)^{\mathcal{I}(t)} \setminus R^{\mathcal{I}(t)} \\
(R_1 \sqcap R_2)^{\mathcal{I}(t)} &= R_1^{\mathcal{I}(t)} \cap R_2^{\mathcal{I}(t)} \\
(U_i/n : E)^{\mathcal{I}(t)} &= \{ \langle d_1, \ldots, d_n \rangle \in (\top_n)^{\mathcal{I}(t)} \mid d_i \in E^{\mathcal{I}(t)} \} \\
(R_1\, \mathcal{U}\, R_2)^{\mathcal{I}(t)} &= \{ \langle d_1, \ldots, d_n \rangle \in (\top_n)^{\mathcal{I}(t)} \mid \exists v > t. (\langle d_1, \ldots, d_n \rangle \in R_2^{\mathcal{I}(v)} \wedge \forall w \in (t, v). \langle d_1, \ldots, d_n \rangle \in R_1^{\mathcal{I}(w)}) \} \\
(R_1\, \mathcal{S}\, R_2)^{\mathcal{I}(t)} &= \{ \langle d_1, \ldots, d_n \rangle \in (\top_n)^{\mathcal{I}(t)} \mid \exists v < t. (\langle d_1, \ldots, d_n \rangle \in R_2^{\mathcal{I}(v)} \wedge \forall w \in (v, t). \langle d_1, \ldots, d_n \rangle \in R_1^{\mathcal{I}(w)}) \} \\
(\Diamond^+ R)^{\mathcal{I}(t)} &= \{ \langle d_1, \ldots, d_n \rangle \in (\top_n)^{\mathcal{I}(t)} \mid \exists v > t. \langle d_1, \ldots, d_n \rangle \in R^{\mathcal{I}(v)} \} \\
(\oplus R)^{\mathcal{I}(t)} &= \{ \langle d_1, \ldots, d_n \rangle \in (\top_n)^{\mathcal{I}(t)} \mid \langle d_1, \ldots, d_n \rangle \in R^{\mathcal{I}(t+1)} \} \\
(\Diamond^- R)^{\mathcal{I}(t)} &= \{ \langle d_1, \ldots, d_n \rangle \in (\top_n)^{\mathcal{I}(t)} \mid \exists v < t. \langle d_1, \ldots, d_n \rangle \in R^{\mathcal{I}(v)} \} \\
(\ominus R)^{\mathcal{I}(t)} &= \{ \langle d_1, \ldots, d_n \rangle \in (\top_n)^{\mathcal{I}(t)} \mid \langle d_1, \ldots, d_n \rangle \in R^{\mathcal{I}(t-1)} \} \\
\top^{\mathcal{I}(t)} &= \Delta^{\mathcal{I}} \\
EN^{\mathcal{I}(t)} &\subseteq \top^{\mathcal{I}(t)} \\
(\neg E)^{\mathcal{I}(t)} &= \top^{\mathcal{I}(t)} \setminus E^{\mathcal{I}(t)} \\
(E_1 \sqcap E_2)^{\mathcal{I}(t)} &= E_1^{\mathcal{I}(t)} \cap E_2^{\mathcal{I}(t)} \\
(\exists^{\leqslant k}[U_j]R)^{\mathcal{I}(t)} &= \{ d \in \top^{\mathcal{I}(t)} \mid \sharp\{ \langle d_1, \ldots, d_n \rangle \in R^{\mathcal{I}(t)} \mid d_j = d \} \leqslant k \} \\
(E_1\, \mathcal{U}\, E_2)^{\mathcal{I}(t)} &= \{ d \in \top^{\mathcal{I}(t)} \mid \exists v > t. (d \in E_2^{\mathcal{I}(v)} \wedge \forall w \in (t, v). d \in E_1^{\mathcal{I}(w)}) \} \\
(E_1\, \mathcal{S}\, E_2)^{\mathcal{I}(t)} &= \{ d \in \top^{\mathcal{I}(t)} \mid \exists v < t. (d \in E_2^{\mathcal{I}(v)} \wedge \forall w \in (v, t). d \in E_1^{\mathcal{I}(w)}) \}
\end{aligned}
$$

Fig. 1. Syntax and semantics of $\mathcal{DLR}_{\mathcal{US}}$.

the temporal operators $\Diamond^+$, $\oplus$ (at the next moment), and their past counterparts can be defined via $\mathcal{U}$ and $\mathcal{S}$: $\Diamond^+ E \equiv \top\, \mathcal{U}\, E$, $\oplus E \equiv \bot\, \mathcal{U}\, E$, etc. However, this is not possible for relations of arity > 1, since $\top_n$—the top n-ary relation—can be interpreted by different subsets of the n-ary cross product $\top \times \cdots \times \top$ at different time points; the reason for this is the ability of $\mathcal{DLR}$ to talk only about difference between relations rather than of the complement of a relation. Then, we may have $\langle d_1, d_2 \rangle \in (\Diamond^+ R)^{\mathcal{I}(t)}$ because $\langle d_1, d_2 \rangle \in R^{\mathcal{I}(t+2)}$, but $\langle d_1, d_2 \rangle \notin (\top_2)^{\mathcal{I}(t+1)}$. The operators $\Diamond^*$ (at some moment) and its dual $\Box^*$ (at all moments) can be defined for both entities and relations as $\Diamond^* E \equiv E \sqcup \Diamond^+ E \sqcup \Diamond^- E$ and $\Box^* E \equiv E \sqcap \Box^+ E \sqcap \Box^- E$, respectively.

A *temporal conceptual database schema* (or a *knowledge base*) is a finite set Σ of $\mathcal{DLR}_{\mathcal{US}}$-formulas. Atomic formulas are formulas of the form $E_1 \sqsubseteq E_2$ and $R_1 \sqsubseteq R_2$, with R_1 and R_2 being relations of the same arity. If φ and ψ are $\mathcal{DLR}_{\mathcal{US}}$-formulas, then so are $\neg\varphi, \varphi \wedge \psi, \varphi\, \mathcal{U}\, \psi, \varphi\, \mathcal{S}\, \psi$. $E_1 \doteq E_2$ is used as an abbreviation for $(E_1 \sqsubseteq E_2) \wedge (E_2 \sqsubseteq E_1)$, for both entities and relations. The global atomic formula $E_1 \sqsubseteq^* E_2$ is used as an abbreviation for $\Box^*(E_1 \sqsubseteq E_2)$, for both entities and relations. Temporal conceptual database schemas specify the constraints for temporal databases.

Given a formula φ, an interpretation $\mathcal{I}$, and a time point $t \in \mathcal{T}$, the truth-relation $\mathcal{I}, t \models \varphi$ (φ holds in $\mathcal{I}$ at moment t) is defined inductively as follows:

$\mathcal{I}, t \models E_1 \sqsubseteq E_2$ iff $E_1^{\mathcal{I}(t)} \subseteq E_2^{\mathcal{I}(t)}$ $\quad$ $\mathcal{I}, t \models \varphi \,\mathcal{U}\, \psi$ iff $\exists v > t.(\mathcal{I}, v \models \psi \land \forall w \in (t, v).\mathcal{I}, w \models \varphi)$
$\mathcal{I}, t \models R_1 \sqsubseteq R_2$ iff $R_1^{\mathcal{I}(t)} \subseteq R_2^{\mathcal{I}(t)}$ $\quad$ $\mathcal{I}, t \models \varphi \,\mathcal{S}\, \psi$ iff $\exists v < t.(\mathcal{I}, v \models \psi \land \forall w \in (v, t).\mathcal{I}, w \models \varphi)$
$\mathcal{I}, t \models \neg\varphi$ iff $\mathcal{I}, t \not\models \varphi$
$\mathcal{I}, t \models \varphi \land \psi$ iff $\mathcal{I}, t \models \varphi$ and $\mathcal{I}, t \models \psi$

A formula φ is called *satisfiable* if there is a temporal model $\mathcal{I}$ such that $\mathcal{I}, t \models \varphi$, for some time point t. A conceptual schema Σ is *satisfiable* if the conjunction $\bigwedge \Sigma$ of all formulas in Σ is satisfiable (we write $\mathcal{I}, t \models \Sigma$ instead of $\mathcal{I}, t \models \bigwedge \Sigma$); in this case $\mathcal{I}$ is called a *model* of Σ. We say that Σ is *globally satisfiable* if there is $\mathcal{I}$ such that $\mathcal{I}, t \models \Sigma$ for every t ($\mathcal{I} \models \Sigma$, in symbols). An entity E (or relation R) is *satisfiable* if there is $\mathcal{I}$ such that $E^{\mathcal{I}(t)} \neq \emptyset$ (respectively, $R^{\mathcal{I}(t)} \neq \emptyset$), for some time point t. Finally, we say that Σ (*globally*) *implies* φ and write $\Sigma \models \varphi$ if we have $\mathcal{I} \models \varphi$ whenever $\mathcal{I} \models \Sigma$.

Note that an entity E is satisfiable iff $\neg(E \sqsubseteq \bot)$ is satisfiable. An n-ary relation R is satisfiable iff $\neg(\exists^{\geq 1}[i]R \sqsubseteq \bot)$ is satisfiable for some $i \leq n$. A conceptual schema Σ is globally satisfiable iff $\Box^*(\bigwedge \Sigma)$ is satisfiable. And $\Sigma \models \varphi$ iff $\Box^*(\bigwedge \Sigma) \land \neg\varphi$ is not satisfiable. Thus, all reasoning tasks connected with the notions introduced above reduce to satisfiability of formulas.

2.2 Temporal queries

In this Section we extend $\mathcal{DLR}_{\mathcal{US}}$ with a temporal query language, and we define the problem of evaluating a temporal query under $\mathcal{DLR}_{\mathcal{US}}$ constraints and the problem of temporal query containment under constraints (see, e.g., [17, 15, 1] for a survey and a discussion about temporal queries). A *non-recursive Datalog query* (i.e., a disjunction of conjunctive queries or SPJ-queries) over a $\mathcal{DLR}_{\mathcal{US}}$ schema Σ is an expression of the form

$$\mathcal{Q}(\overrightarrow{x}) :- \bigvee_j \mathcal{Q}_j(\overrightarrow{x}, \overrightarrow{y_j}, \overrightarrow{c_j}), \qquad \text{where} \qquad \mathcal{Q}_j(\overrightarrow{x}, \overrightarrow{y_j}, \overrightarrow{c_j}) \equiv \bigwedge_i P_j^i(\overrightarrow{x}^i_j, \overrightarrow{y}^i_j, \overrightarrow{c}^i_j),$$

P_j^i are $\mathcal{DLR}_{\mathcal{US}}$ entity or relation expressions possibly occurring in Σ, $\overrightarrow{x}^i_j$, $\overrightarrow{y}^i_j$, and $\overrightarrow{c}^i_j$ are sequences of distinguished variables, existential variables, and constants, respectively, the number of which is in agreement with the arity of P_j^i. The variables $\overrightarrow{x}$ in the head are the union of all the distinguished variables in each $\mathcal{Q}_j$; the existential variables are used to make coreferences in the query, and constants are fixed values. $\mathcal{Q}$ is a n-ary relation (not appearing in Σ) whose arity is the number of variables in $\overrightarrow{x}$.

It is to be noted that we allow entities and relations in the query to occur in the conceptual schema Σ. This approach is similar to that of [12], where atoms in a query can be constrained by means of schema formulas. Furthermore, query expressions do not directly manipulate explicit temporal attributes, but time is implicit in each query expression. Indeed, the temporal dimension is handled by means of the temporal modal operators in $\mathcal{DLR}_{\mathcal{US}}$. In this perspective, the query language is in strict relation with the First-Order Temporal Logic with *since* and *until*, $L^{\{\textbf{since, until}\}}$, used in [17] for querying temporal databases.

The semantics of queries is based on the snapshot representation of a temporal database, and defined as follows. Given a temporal schema Σ, let $\mathcal{I}$ be a temporal model, and t be a time point in $\mathcal{T}$ such that $\mathcal{I}$ satisfies Σ at t, i.e., $\mathcal{I}, t \models \Sigma$. The snapshot interpretation

$$\mathcal{I}(t) = \left\langle \Delta^{\mathcal{I}}, \{E^{\mathcal{I}(t)} \mid E \in \mathit{EN}\}, \{R^{\mathcal{I}(t)} \mid R \in \mathit{RN}\} \right\rangle$$

can be regarded as a usual first-order structure (i.e., a snapshot, non-temporal, database at time t conforming in a sense to the conceptual schema), and so the whole $\mathcal{I}$ as a first-order temporal model (with constant domain $\Delta^{\mathcal{I}}$ in which some values of the query constants are specified). The *evaluation* of a query $\mathcal{Q}$ of arity n, under the constraints Σ, in the model $\mathcal{I}$ that satisfies Σ at moment t, and the *answer* to the query $\mathcal{Q}$, are respectively the sets:

$$\mathsf{eval}(\mathcal{Q},\mathcal{I}(t)) = \{\overrightarrow{d} \in (\Delta^{\mathcal{I}})^n \mid \mathcal{I},t \models \bigvee_j \exists\overrightarrow{y_j}.\mathcal{Q}_j(\overrightarrow{d},\overrightarrow{y_j},\overrightarrow{c_j})\}$$
$$\mathsf{ans}(\mathcal{Q},\mathcal{I}) = \{\langle t,\overrightarrow{d}\rangle \in \mathcal{T} \times (\Delta^{\mathcal{I}})^n \mid \overrightarrow{d} \in \mathsf{eval}(\mathcal{Q},\mathcal{I}(t))\}$$

We obtain a so called *sequenced semantics* [8] for queries, which is based on the view of a database as a time-indexed collection of snapshots. The query language is also *snapshot-reducible* [8] in the sense that non-temporal queries—i.e., queries without any temporal connective—are still valid queries, and are interpreted using the sequenced semantics. Our language allows also for *upward compatible* queries [8]. Intuitively, a non-temporal query is upward compatible if the answer set on a temporal database is the same as the answer set on an associated non-temporal database. This is a temporal *slice* of the temporal database at the current time: $\mathsf{eval}(\mathcal{Q},\mathcal{I}(\mathsf{now}))$.

Given two queries (of the same arity) $\mathcal{Q}_1$ and $\mathcal{Q}_2$ over Σ, we say that $\mathcal{Q}_1$ is *contained* in $\mathcal{Q}_2$ under the constraints Σ, and write $\Sigma \models \mathcal{Q}_1 \subseteq \mathcal{Q}_2$, if, for every temporal model $\mathcal{I}$ of Σ we have $\mathsf{ans}(\mathcal{Q}_1,\mathcal{I}) \subseteq \mathsf{ans}(\mathcal{Q}_2,\mathcal{I})$. The *query satisfiability problem*—given a query $\mathcal{Q}$ over a schema Σ, to determine whether there are $\mathcal{I}$ and t such that $\mathcal{I},t \models \Sigma$, and $\mathsf{eval}(\mathcal{Q},\mathcal{I}(t)) \neq \emptyset$—is reducible to query containment: $\mathcal{Q}$ is satisfiable iff $\Sigma \not\models \mathcal{Q}(\overrightarrow{x}) \subseteq P(\overrightarrow{x}) \wedge \neg P(\overrightarrow{x})$, where P is a $\mathcal{DLR}_{\mathcal{US}}$-relation of the same arity as $\mathcal{Q}$.

2.3 Conceptual Schema and Query Examples

As an example, let us consider the following conceptual schema Σ:

$\mathtt{Works\text{-}for} \sqsubseteq^* \mathtt{emp}/2 : \mathtt{Employee} \sqcap \mathtt{act}/2 : \mathtt{Project}$
$\mathtt{Manages} \sqsubseteq^* \mathtt{man}/2 : \mathtt{TopManager} \sqcap \mathtt{prj}/2 : \mathtt{Project}$
$\mathtt{Employee} \sqsubseteq^* \exists^{=1}[\mathtt{from}]\mathtt{PaySlipNumber} \sqcap \exists^{=1}[\mathtt{from}](\mathtt{PaySlipNumber} \sqcap \mathtt{to}/2 : \mathtt{Integer}) \sqcap$
$\qquad \exists^{=1}[\mathtt{from}]\mathtt{Salary} \sqcap \exists^{=1}[\mathtt{from}](\mathtt{Salary} \sqcap \mathtt{to}/2 : \mathtt{Integer})$
$\top \sqsubseteq^* \exists^{\leq 1}[\mathtt{to}](\mathtt{PaySlipNumber} \sqcap \mathtt{from}/2 : \mathtt{Employee})$
$\mathtt{Manager} \sqsubseteq^* \mathtt{Employee} \sqcap (\mathtt{AreaManager} \sqcup \mathtt{TopManager})$
$\mathtt{AreaManager} \sqsubseteq^* \mathtt{Manager} \sqcap \neg\mathtt{TopManager}$
$\mathtt{TopManager} \sqsubseteq^* \mathtt{Manager} \sqcap \exists^{=1}[\mathtt{man}]\mathtt{Manages}$
$\mathtt{Project} \sqsubseteq^* \exists^{\geq 1}[\mathtt{act}]\mathtt{Works\text{-}for} \sqcap \exists^{=1}[\mathtt{prj}]\mathtt{Manages}$
$\mathtt{Employee} \sqcap \neg(\exists^{\geq 1}[\mathtt{emp}]\mathtt{Works\text{-}for}) \sqsubseteq^* \mathtt{Manager}$
$\mathtt{Manager} \sqsubseteq^* \neg(\exists^{\geq 1}[\mathtt{emp}]\mathtt{Works\text{-}for}) \sqcap (\mathtt{Qualified}\ \mathcal{S}\ (\mathtt{Employee} \sqcap \neg\mathtt{Manager}))$

The theory introduces `Works-for` as a binary relation between employees and projects, and `Manages` as a binary relation between managers and projects. Employees have exactly one pay slip number and one salary each, which are represented as binary relations (with `from` and `to` roles) with an integer domain; moreover, a pay slip number uniquely identifies an employee (it acts as a key). It is stated that managers are employees, and are partitioned into area managers and top managers. Top Managers participate exactly once in the relation `Manages`, i.e., every top manager manages exactly one project. Projects participate at least once to the relation `Works-for` and exactly once in the relation `Manages`. Finally, employees not working for a project are exactly the managers, and managers should be qualified, i.e., should have passed a period of being employees. The meaning of the above conceptual schema (with the exception of the last two formulas) is illustrated by the left part of the diagram in Fig. 2.

The conceptual schema Σ globally logically implies that, for every project, there is at least one employee who is not a manager, and that a top manager worked in a project before managing some (possibly different) project:

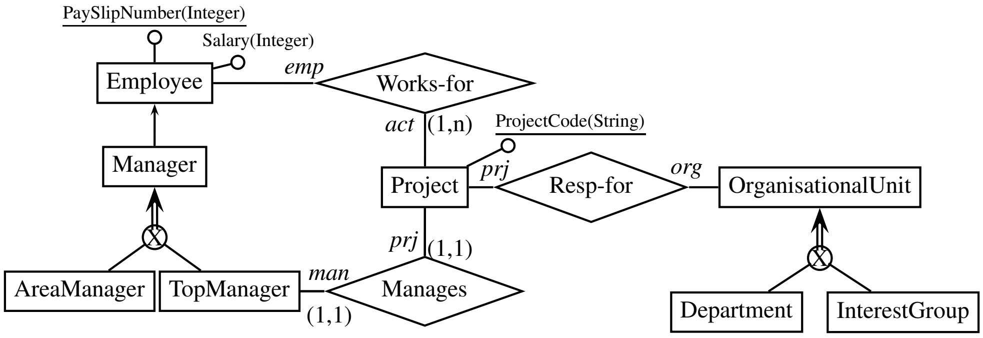

Fig. 2. The example EER diagram.

$\Sigma \models \texttt{Project} \sqsubseteq^* \exists^{\geq 1}[\texttt{act}](\texttt{Works-for} \sqcap \texttt{emp} : \neg\texttt{Manager})$
$\Sigma \models \texttt{TopManager} \sqsubseteq^* \Diamond^- \exists^{\geq 1}[\texttt{emp}](\texttt{Works-for} \sqcap \texttt{act} : \texttt{Project})$

Note also that if we add to Σ the formula

$\texttt{Employee} \sqsubseteq^* \exists^{\geq 1}[\texttt{emp}]\texttt{Works-for}$

saying that every employee should work for at least one project, then all the entities and the relations mentioned in the conceptual schema are interpreted as the empty set in every model of Σ, i.e., they are not satisfiable relative to Σ.

The expressivity of the query language can be understood with the following examples, taken from [17]:

"*Find all people who have worked for only one project*"
$\mathcal{Q}(\texttt{x}) :\!- (\exists^{=1}[\texttt{emp}](\Diamond^*\texttt{Works-for}))(\texttt{x})$
"*Find all managers whose terminal project has code* prj342"
$\mathcal{Q}(\texttt{x}) :\!- \texttt{Manager}(\texttt{x}) \wedge \texttt{Manages}(\texttt{x}, \text{prj342}) \wedge (\Box^+ \neg\texttt{Manages})(\texttt{x}, \texttt{y})$
"*Find all project-hoppers—people who never spent more than two consecutive years in a project*"
$\mathcal{Q}(\texttt{x}) :\!- (\Box^* \neg\exists^{\geq 1}[\texttt{emp}](\texttt{Works-for} \sqcap \oplus\, \texttt{Works-for} \sqcap \oplus \oplus\, \texttt{Works-for}))(\texttt{x})$
"*Find all people who did not work between two projects*"
$\mathcal{Q}(\texttt{x}) :\!- (\Diamond^- \exists^{\geq 1}[\texttt{emp}]\texttt{Works-for})(\texttt{x}) \wedge (\neg\exists^{\geq 1}[\texttt{emp}]\texttt{Works-for})(\texttt{x}) \wedge (\Diamond^+ \exists^{\geq 1}[\texttt{emp}]\texttt{Works-for})(\texttt{x})$

We now consider an example of query containment under constraints, where the constraints are expressed by the above schema Σ. Consider the following queries:

$\mathcal{Q}_1(\texttt{x}, \texttt{y}) :\!- \neg\texttt{AreaManager}(\texttt{x}) \wedge \texttt{Manages}(\texttt{x}, \texttt{z}) \wedge \texttt{Project}(\texttt{z}) \wedge \texttt{Resp-for}(\texttt{y}, \texttt{z}) \wedge \texttt{Department}(\texttt{y})$
$\mathcal{Q}_2(\texttt{x}, \texttt{y}) :\!- (\Diamond^- \exists^{\geq 1}[1]\texttt{Works-for})(\texttt{x}) \wedge \texttt{Manages}(\texttt{x}, \texttt{z}) \wedge \texttt{Resp-for}(\texttt{y}, \texttt{z}) \wedge \neg\texttt{InterestGroup}(\texttt{y})$

It is not hard to see that $\mathcal{Q}_1$ is contained in $\mathcal{Q}_2$ under the constraints in Σ, i.e., $\Sigma \models \mathcal{Q}_1 \subseteq \mathcal{Q}_2$.

3 Decidability and complexity

In this section, we present and briefly discuss our main results on the computational behaviour of $\mathcal{DLR}_{\mathcal{US}}$ and its fragments over the flow of time $\langle \mathbb{Z}, < \rangle$.

Unfortunately, full $\mathcal{DLR}_{\mathcal{US}}$, even restricted to *atomic* formulas, turns out to be undecidable.

Theorem 1. *The global satisfiability problem for $\mathcal{DLR}_{\mathcal{US}}$ conceptual schemas containing only atomic formulas is undecidable.*

Proof. The proof is by reduction of the well-known undecidable tiling problem [37]: given a finite set of square tiles of fixed orientation and with coloured edges, decide whether it can tile the grid $\mathbb{Z} \times \mathbb{N}$. Suppose $\mathfrak{T} = \{T_1, \ldots, T_k\}$ is a set of tiles with colours $left(T_i)$, $right(T_i)$, $up(T_i)$, and $down(T_i)$. Consider the following schema Σ, where $D_1, \ldots, D_k$ are concepts and R is a binary relation:

- $R \doteq \Box^+ R,\ R \doteq \Diamond^+ R,\ \top \doteq \exists R.\top$,
- $D_i \sqsubseteq \neg D_j,\ \top \doteq D_1 \sqcup \cdots \sqcup D_k$, for $i \neq j$,
- $D_i \sqsubseteq \bigsqcup_{right(T_i)=left(T_j)} \forall R.D_j$, for $i \leq k$,
- $D_i \sqsubseteq \bigsqcup_{up(T_i)=down(T_j)} \oplus D_j$, for $i \leq k$.

(Here $\exists R.C = \exists^{\geq 1}[1](R \sqcap 2/2 : C)$, $\forall R.C = \neg\exists R.\neg C$.) It is readily checked that Σ is globally satisfiable iff $\mathfrak{T}$ tiles $\mathbb{Z} \times \mathbb{N}$.

It follows, in particular, that (a) the problem of satisfiability of complex $\mathcal{DLR}_{\mathcal{US}}$ formulas is undecidable, and (b) the problem of global logical implication in $\mathcal{DLR}_{\mathcal{US}}$—even involving only atomic formulas—is undecidable as well. The main technical reason for undecidability is the possibility to 'temporalise' binary relations, cf. [47]. The fragment $\mathcal{DLR}^-_{\mathcal{US}}$, in which the temporal operators can be applied only to entities and formulas (but not to n-ary relations, with $n \geq 2$), exhibits a much better computational behaviour. The following theorem presents the complexity results we have obtained for schema and query reasoning in $\mathcal{DLR}^-_{\mathcal{US}}$:

Theorem 2. *Let the flow of time be $\langle \mathbb{Z}, < \rangle$. Then*

(1) *the problem of logical implication in $\mathcal{DLR}^-_{\mathcal{US}}$ involving only atomic formulas is EXPTIME-complete;*

(2) *the formula satisfiability problem (and so the problem of logical implication) in $\mathcal{DLR}^-_{\mathcal{US}}$ is EXPSPACE-complete;*

(3) *the query-containment problem for non-recursive Datalog queries under $\mathcal{DLR}^-_{\mathcal{US}}$-constraints is decidable in 2EXPTIME and is EXPSPACE-hard.*

In the remainder of the section we sketch a proof of this theorem. To make it more transparent, we confine ourselves to considering only the 'future fragment' $\mathcal{DLR}^-_{\mathcal{U}}$ of $\mathcal{DLR}^-_{\mathcal{US}}$. (From now on $\Box$ stands for $\Box^+$ and $\bigcirc$ for $\oplus$.)

Suppose that Γ consists of a finite set $f(\Gamma)$ of $\mathcal{DLR}^-_{\mathcal{U}}$-formulas and a finite set $c(\Gamma)$ of concepts, $f(\Gamma)$ is closed under sub-formulas, $c(\Gamma)$ under subconcepts, both are closed under (single) negation, and each concept occurring in $f(\Gamma)$ belongs $c(\Gamma)$. A *concept type* for Γ is a subset t of $c(\Gamma)$ such that

- $C \sqcap D \in t$ iff $C, D \in t$, for all $C \sqcap D \in c(\Gamma)$;
- $\neg C \in t$ iff $C \notin t$, for all $C \in c(\Gamma)$.

A *formula type* for Γ is a subset Φ of $f(\Gamma)$ such that

- $\psi \wedge \chi \in \Phi$ iff $\psi, \chi \in \Phi$, for all $\psi \wedge \chi \in f(\Gamma)$;
- $\neg\psi \in \Phi$ iff $\psi \notin \Phi$, for all $\psi \in f(\Gamma)$.

A pair $\langle T, \Phi \rangle$, where T is a set of concept types and Φ a formula type for Γ, is called a *quasistate candidate* for Γ. We say that the quasistate candidate $\mathfrak{C} = \langle T, \Phi \rangle$ is a *quasistate* for Γ if the following (non-temporal) $\mathcal{DLR}$-formula $\alpha_{\mathfrak{C}}$

$$\big(\bigsqcup_{t \in T} c(t) \doteq \top \big) \wedge \bigwedge_{t \in T} \neg (c(t) \doteq \bot) \wedge \bigwedge \Phi$$

is satisfiable. Here $c(t)$ denotes the conjunction of all concepts in t, concepts of the form $C \,\mathcal{U}\, D$ are regarded as atomic concepts $A_{C\mathcal{U}D}$, and formulas of the form $\varphi \,\mathcal{U}\, \psi$ in Φ are regarded as atomic formulas $A_{\varphi\mathcal{U}\psi} = \top$.

Consider now a sequence of quasistates

$$Q = \langle Q(n) : n \in \mathbb{Z} \rangle,$$

where $Q(n) = \langle T_n, \Phi_n \rangle$. A *run* in Q is a sequence $r = \langle r(n) : n \in \mathbb{Z} \rangle$ such that

1. $r(n) \in T_n$ for every $n \in \mathbb{Z}$;
2. for every $C \,\mathcal{U}\, D \in c(\Gamma)$ and every $n \in \mathbb{Z}$, we have $C \,\mathcal{U}\, D \in r(n)$ iff there is $l > n$ such that $D \in r(l)$ and $C \in r(k)$ for all $k \in (n, l)$.

Finally, Q is called a *quasimodel for Γ* if the following conditions hold:

3. for every $n \in \mathbb{Z}$ and every $t \in T_n$ there is a run r in Q such that $r(n) = t$;
4. for every $\psi \,\mathcal{U}\, \chi \in f(\Gamma)$ and every $n \in \mathbb{Z}$, we have $\psi \,\mathcal{U}\, \chi \in \Phi_n$ iff there is $l > n$ such that $\chi \in \Phi_l$ and $\psi \in \Phi_k$ for all $k \in (n, l)$.

Given a $\mathcal{DLR}^-_{\mathcal{U}}$-formula φ, we denote by $cl(\varphi)$ the closure under (single) negation of the set of subformulas and subconcepts of φ.

Theorem 3. *A $\mathcal{DLR}^-_{\mathcal{U}}$-formula φ is satisfiable iff there is a quasimodel for $cl(\varphi)$ such that $\varphi \in \Phi_0$.*

Proof. Suppose φ is satisfied in a model $\mathcal{I}$ with domain Δ. For every $n \in \mathbb{Z}$, define $Q(n) = \langle T_n, \Phi_n \rangle$ by taking

$$T_n = \{ t^n(x) : x \in \Delta \},\ \Phi_n = \{ \psi \in cl(\varphi) : \mathcal{I}, n \models \psi \},$$

where $t^n(x) = \{ C \in cl(\varphi) : x \in C^{\mathcal{I}(n)} \}$. It is not hard to see that $\langle Q(n) : n \in \mathbb{Z} \rangle$ is a quasimodel for φ. (Note that the sequence $\langle t^n(x) : n \in \mathbb{Z} \rangle$ is a run through $t^n(x)$, for every $n \in \mathbb{Z}$ and every $x \in \Delta$). To show the converse we require the following lemma.

Lemma 1. *For any cardinal $\kappa \geq \aleph_0$ and any quasistate $\mathfrak{C}$ for φ, the formula $\alpha_{\mathfrak{C}}$ is satisfied in a (non-temporal) $\mathcal{DLR}$-model $\mathcal{J}$ in which $\big| [x]^{\mathcal{J}} \big| = \kappa$ for all x in the domain Δ of $\mathcal{J}$, where*

$$[x]^{\mathcal{J}} = \{ y \in \Delta : \forall C \in cl(\varphi) (x \in C^{\mathcal{J}} \Leftrightarrow y \in C^{\mathcal{J}}) \}.$$

Proof. As $\mathcal{DLR}$ is a fragment of first-order logic, we have a countable $\mathcal{DLR}$-model $\mathcal{I}$ satisfying $\alpha_{\mathfrak{C}}$. Define $\mathcal{J}$ as the disjoint union of κ copies of $\mathcal{I}$; more precisely, let

$$\Delta^{\mathcal{J}} = \{ \langle x, \xi \rangle : x \in \Delta^{\mathcal{I}}, \xi < \kappa \},$$

$$P_i^{\mathcal{J}} = \{ \langle \langle x_0, \xi \rangle, \ldots, \langle x_n, \xi \rangle \rangle : \langle x_0, \ldots, x_n \rangle \in P_i^{\mathcal{I}}, \xi < \kappa \},$$

$$(\top_n)^{\mathcal{J}} = \{ \langle \langle x_0, \xi \rangle, \ldots, \langle x_n, \xi \rangle \rangle : \langle x_0, \ldots, x_n \rangle \in (\top_n)^{\mathcal{I}}, \xi < \kappa \}.$$

It is not hard to see that $\mathcal{J}$ is as required.

Suppose now that $\varphi \in \Phi_0$, for a quasimodel Q. Let κ be a cardinal exceeding the cardinality of the set Ω of all runs in Q and $\aleph_0$, and let

$$\Delta = \{\langle r, \xi \rangle : r \in \Omega, \xi < \kappa\}.$$

Note that $|\{\langle r, \xi \rangle \in \Delta : r(n) = t\}| = \kappa$, for every $n \in \mathbb{Z}$ and every $t \in T_n$. By Lemma 1, for every $n \in \mathbb{Z}$ there is a $\mathcal{DLR}$-model $\mathcal{J}(n)$ with domain Δ satisfying $\alpha_{Q(n)}$ and such that $\{C \in cl(\varphi) : \langle r, \xi \rangle \in C^{\mathcal{J}(n)}\} = r(n)$, for all $r \in \Omega$ and $\xi < \kappa$. It is not hard to see that the temporal $\mathcal{DLR}$-model $\mathcal{I} = \langle \mathbb{Z}, \Delta, \cdot^{\mathcal{I}(n)} \rangle$ defined by taking $\mathcal{I}(n) = \mathcal{J}(n)$, for every $n \in \mathbb{Z}$, satisfies φ at moment 0.

Thus, the satisfiability problem for $\mathcal{DLR}_{\mathcal{U}}^-$-formulas reduces to checking 'satisfiability' in quasimodels. Consider now a $\mathcal{DLR}_{\mathcal{U}}^-$-schema Σ and two queries

$$\mathcal{Q}_i(\boldsymbol{x}) :- \bigvee_j \mathcal{Q}_j^i(\boldsymbol{x}, \boldsymbol{y_{ij}}, \boldsymbol{w_{ij}}), \quad i = 1, 2.$$

Denote by $cl(\Sigma, \mathcal{Q}_1, \mathcal{Q}_2)$ the closure under (single) negation of the set of all formulas and concepts occurring in Σ, $\mathcal{Q}_1$ and $\mathcal{Q}_2$. Given a formula or a concept χ, denote by $\overline{\chi}$ the result of replacing all subformulas (subconcepts) in χ of the form $\chi_1 \,\mathcal{U}\, \chi_2$ with $A_{\chi_1 \mathcal{U} \chi_2} = \top$ (respectively, $A_{\chi_1 \mathcal{U} \chi_2}$). Thus, $\overline{\chi}$ is a $\mathcal{DLR}$ formula or concept, and the $\overline{\mathcal{Q}_i}$ are non-temporal $\mathcal{DLR}$-queries.

Theorem 4. *$\mathcal{Q}_1$ is not contained in $\mathcal{Q}_2$ relative to Σ iff there is a quasimodel Q for $cl(\Sigma, \mathcal{Q}_1, \mathcal{Q}_2)$ such that $\overline{\mathcal{Q}}_1$ is not contained in $\overline{\mathcal{Q}}_2$ relative to $\overline{\Sigma} \cup \{\overline{\alpha}_{Q(0)}\}$.*

Proof. ($\Rightarrow$) Without loss of generality we may assume that we have a model $\mathcal{I}$ such that $\mathcal{I}(0) \models \Sigma$ and $\mathsf{ans}(\mathcal{Q}_1, \mathcal{I}(0)) \not\subseteq \mathsf{eval}(\mathcal{Q}_2, \mathcal{I}(0))$. Construct a quasimodel Q for $cl(\Sigma, \mathcal{Q}_1, \mathcal{Q}_2)$ as in the proof of Theorem 3. To show that $\overline{\mathcal{Q}}_1$ is not contained in $\overline{\mathcal{Q}}_2$ relative to $\overline{\Sigma} \cup \{\overline{\alpha}_{Q(0)}\}$, it is enough to extend the (non-temporal) model $\mathcal{I}(0)$ to the new 'surrogate' atoms of the form $A_{C_1 \mathcal{U} C_2}$ and $A_{\chi_1 \mathcal{U} \chi_2}$ in accordance with their behaviour in $\mathcal{I}$ at time point 0: $A_{C_1 \mathcal{U} C_2}^{\mathcal{I}(0)} = (C_1 \mathcal{U} C_2)^{\mathcal{I}(0)}$ and

$$A_{\chi_1 \mathcal{U} \chi_2}^{\mathcal{I}(0)} = \begin{cases} \top \text{ if } \mathcal{I}(0) \models \chi_1 \,\mathcal{U}\, \chi_2 \\ \bot \text{ otherwise.} \end{cases}$$

($\Leftarrow$) is also proved similarly to Theorem 3. The only difference is that now we select $\mathcal{J}(0)$ so that $\mathcal{J}(0) \models \overline{\Sigma} \wedge \overline{\alpha}_{Q(0)}$ and $\mathsf{eval}(\overline{\mathcal{Q}}_1, \mathcal{J}(0)) \not\subseteq \mathsf{ans}(\overline{\mathcal{Q}}_2, \mathcal{J}(0))$.

So, the query-containment problem for $\mathcal{DLR}_{\mathcal{U}}^-$ reduces to satisfiability in quasimodels and the query-containment problem for (non-temporal) $\mathcal{DLR}$. The latter problem was shown to be decidable in 2EXPTIME time in [12]. But how to check satisfiability in quasimodels? First of all, we need a procedure deciding whether a quasistate candidate is a quasistate for a given set of formulas and concepts. The following proposition can be proved using the reduction in [12].

Proposition 1. (i) *Given a $\mathcal{DLR}_{\mathcal{U}}^-$-formula φ, it is decidable in EXPSPACE whether a quasistate candidate for $cl(\varphi)$ is a quasistate.*

(ii) *Given a $\mathcal{DLR}_{\mathcal{U}}^-$-schema Σ and queries $\mathcal{Q}_1$, $\mathcal{Q}_2$, it is decidable in 2EXPTIME whether $\overline{\mathcal{Q}}_1$ is contained in $\overline{\mathcal{Q}}_2$ relative to $\overline{\Sigma} \cup \{\overline{\alpha}_{\mathfrak{C}}\}$ for a quasistate candidate $\mathfrak{C}$ for $cl(\Sigma, \mathcal{Q}_1, \mathcal{Q}_2)$.*

Now, given a set Γ as defined above, we have at most $O(2^{2^{|\Gamma|}})$ distinct quasistates for Γ. The problem then is whether they can be properly arranged to form a quasimodel for Γ. As we have no past temporal operators, it is enough to consider the flow of time $\langle \mathbb{N}, < \rangle$ and quasimodels of the form $Q = \langle Q(n) : n \in \mathbb{N} \rangle$.

Let Q be a sequence of quasistates $Q(i) = \langle T_i, \Phi_i \rangle$, $i \in \mathbb{N}$, and r a sequence of elements from T_i such that $r(i) \in T_i$. Say that r *realises* $C\mathcal{U}D \in r(n)$ *in* m *steps* if there is $l \leq m$ such that $D \in r(n+l)$ and $C \in r(n+k)$ for all $k \in (0, l)$. A formula $\psi\mathcal{U}\chi \in \Phi_n$ is *realised in* m *steps* if there is $l \leq m$ such that $\chi \in \Phi_{n+l}$ and $\psi \in \Phi_{n+k}$ for all $k \in (0, l)$. We also say that a pair t, t' of concept types is *suitable* if for every $C\mathcal{U}D \in \Gamma$, $C\mathcal{U}D \in t$ iff either $D \in t'$ or $C \in t'$ and $C\mathcal{U}D \in t'$.

Suppose Q_1 and Q_2 are finite sequences of quasistates for Γ of length l_1 and l_2, respectively, and let $Q = Q_1 * Q_2^*$ (i.e., $Q = Q_1 * Q_2 * Q_2 * Q_2 * \dots$) with $Q(n) = \langle T_n, \Phi_n \rangle$. One can check that Q is a quasimodel for Γ if the following conditions hold:

(a) for every $i \leq l_1 + l_2$ and every $t' \in T_{i+1}$, there is $t \in T_i$ such that the pair t, t' is suitable;
(b) for every $i \leq l_1 + 1$ and every $t_i \in T_i$, all concepts of the form $C\mathcal{U}D \in t_i$ are realised in $l_1 + l_2 - i$ steps in some sequence $t_i, t_{i+1}, \dots, t_{l_1+l_2}$ in which $t_{i+j} \in T_{i+j}$ and every pair of adjacent elements is suitable;
(c) for every $i \leq l_1 + l_2$, and every $\psi\mathcal{U}\chi \in \Gamma$, $\psi\mathcal{U}\chi \in \Phi_i$ iff either $\chi \in \Phi_{i+1}$ or $\psi \in \Phi_{i+1}$ and $\psi\mathcal{U}\chi \in \Phi_{i+1}$;
(d) for every $i \leq l_1 + 1$, all formulas of the form $\psi\mathcal{U}\chi \in \Phi_i$ are realised in $l_1 + l_2 - i$ steps.

Moreover, given a quasimodel for Γ, one can always extract from it a subquasimodel $Q = Q_1 * Q_2^*$ which satisfies (a)–(d) above, all quasistates in Q_1 are distinct and $|Q_2| = O(2^{2^{|\Gamma|}})$.

Using this observation together with Proposition 1 one can construct an EXPSPACE formula-satisfiability checking algorithm and a 2EXPTIME query-containment checking algorithm. A proof of EXPSPACE-hardness of the formula-satisfiability problem (for a much weaker logic) can be found in [22]. It follows, in particular, that the query-containment problem is EXPSPACE-hard as well. It is an open problem, however, whether there exists an EXPSPACE algorithm deciding this problem.

Finally, we show EXPTIME-completeness of the logical implication for atomic formulas in $\mathcal{DLR}_{\mathcal{U}}^{-}$ by means of a polynomial reduction of $\mathcal{DLR}_{\mathcal{U}}^{-}$ to the logic $\mathcal{DLR}_{reg}$ of [12]. For our purposes, it is enough to know that $\mathcal{DLR}_{reg}$ allows one to form the transitive closure R^* of every binary relation R, and that the satisfiability problem in $\mathcal{DLR}_{reg}$ is in EXPTIME. To simplify presentation, we reduce here the fragment $\mathcal{DLR}_{\Box\bigcirc}^{-}$ of $\mathcal{DLR}_{\mathcal{U}}^{-}$ with the temporal operators $\Box$ and $\bigcirc$ only (the reader should not have problems to extend this reduction to the language with $\mathcal{U}$).

Fix a binary relation R and define a translation $\star$ from $\mathcal{DLR}_{\Box\bigcirc}^{-}$ to $\mathcal{DLR}_{reg}$ as follows: $P^\star = P$ for every atom P of $\mathcal{DLR}$,

$$(\bigcirc C)^\star = \forall R.C^\star$$
$$(\Box C)^\star = \forall R^*.C^\star;$$

$\star$ commutes with the remaining constructs, and $(P_1 \sqsubseteq P_2)^\star = P_1^\star \sqsubseteq P_2^\star$.

Lemma 2. *Suppose that $\Gamma \cup \{\varphi\}$ is a set of atomic $\mathcal{DLR}_{\bigcirc\Box}^{-}$-formulas and that R does not occur in $\Gamma \cup \{\varphi\}$. Then $\Gamma \models \varphi$ iff $\varphi^\star$ is a logical consequence of the following set Ξ of $\mathcal{DLR}_{reg}$-formulas*

$$\Gamma^\star,\ \exists^{=1}R.\top \doteq \top,\ \exists^{=1}R^-.\top \doteq \top,$$

where $\exists^{=1}R^-.C = \exists^{=1}[2](R \sqcap 1/2 : C)$.

Proof. Suppose $\Gamma \not\models \varphi$. Then there is a model $\mathcal{I}$ such that $\mathcal{I}, 0 \not\models \varphi$, but $\mathcal{I}, n \models \Gamma$ for all $n \in \mathbb{Z}$. Define a $\mathcal{DLR}$-model

$$\mathcal{J} = \left\langle \Delta', P_1^{\mathcal{J}}, \dots, R^{\mathcal{J}} \right\rangle$$

by taking

- $\Delta' = \Delta^{\mathcal{I}} \times \mathbb{Z}$,
- $\langle \langle x_1, n_1 \rangle, \dots, \langle x_l, n_l \rangle \rangle \in P_i^{\mathcal{J}}$ iff $n_i = n_j$, for $i, j \leq l$, and $\langle x_1, \dots, x_n \rangle \in P_i^{\mathcal{I}(n_1)}$,
- $\langle \langle x_1, n_1 \rangle, \langle x_2, n_2 \rangle \rangle \in R^{\mathcal{J}}$ iff $x_1 = x_2$ and $n_2 = n_1 + 1$.

It is readily checked $\mathcal{J} \models \Xi$ and $\mathcal{J} \not\models \varphi^\star$.

Conversely, suppose that $\mathcal{J} = \left\langle \Delta, P_1^{\mathcal{J}}, \dots, R^{\mathcal{J}} \right\rangle$ is a model such that $\mathcal{J} \models \Xi$ but $\mathcal{J} \not\models \varphi^\star$. Let $\Sigma = \bigcup \{cl(\chi) : \chi \in \Gamma \cup \{\varphi\}\}$ and, for every $x \in \Delta$,

$$t(x) = \{C \in c(\Sigma) : x \in (C^\star)^{\mathcal{J}}\}.$$

Then the pair $\langle T, \Phi \rangle$, where $T = \{t(x) : x \in \Delta\}$ and $\Phi = \{\chi \in f(\Sigma) : \mathcal{J} \models \chi^\star\}$, is a quasistate for Σ. Define a map Q by taking $Q(n) = \langle T, \Phi \rangle$ for all $n \in \mathbb{Z}$. It is easy to see that Q is a quasimodel. Hence, by Theorem 3, we have a model $\mathcal{I}$ such that $\mathcal{I} \models \Gamma$ but $\mathcal{I}, 0 \not\models \varphi$.

4 Conceptual modelling

In this section we briefly show how the temporal description logic $\mathcal{DLR}_{\mathcal{US}}$ can provide a formal semantic characterisation of the most important temporal conceptual modelling constructs (for the valid time representation). We refer mostly to the temporal extended entity-relationship data model, for which a detailed literature exists [25, 24, 41, 44]. We consider the following constructs: snapshot and temporary entities, relationships, attributes (together with the temporal key attribute), temporal cardinalities, dynamic entities, liveliness and safety constraints, and the monotonic form of schema evolution.

The extended entity-relationship (EER) model—i.e., the standard entity-relationship data model, enriched with IsA links, disjoint and covering constraints, and full cardinality constraints [18]—may be viewed as a temporalised EER model which assigns to every construct a temporal interpretation but provides no explicit temporal constructs. Gregersen and Jensen [24] call this approach *implicit*, because the temporal dimension is hidden in the interpretation structure so that entities and relationships are always time-dependent. The non-temporal fragment of $\mathcal{DLR}_{\mathcal{US}}$, i.e., $\mathcal{DLR}$, is enough to capture the EER model with implicit time. For the non-temporal EER model, such an encoding, introduced by [13, 14], establishes a precise correspondence between legal database states of the EER diagram and models of the derived $\mathcal{DLR}$ theory. That this encoding is correct for the EER model with implicit time was shown in [2, 3], where $\mathcal{DLR}$ was interpreted by a temporal semantics. The example knowledge base in Section 2.3 (without the last two formulas) shows the exact encoding for the left-hand part of the temporally implicit EER diagram in Fig.2. This encoding could support the design of a temporal conceptual schema by exploiting the reasoning in $\mathcal{DLR}_{\mathcal{US}}$: it becomes possible to verify the conceptual specification, infer implicit facts and stricter constraints, and manifest any inconsistencies. Note that the same ideas presented here would apply to non-temporal UML class diagrams (encoded in $\mathcal{DLR}$ in [11]) and DAML+OIL ontologies [19].

We now introduce the basic temporal constructs that can be added on top of a temporally implicit model leading to a full fledged temporal conceptual model.

Temporal entities and relations. Both entity and relation instances in a temporal setting have an existence time associated to them. $\mathcal{DLR}_{\mathcal{US}}$-formulas can enforce either that entities (relations) cannot last forever—we call them *temporary entities* (*temporary relations*), or that their extension never changes in time—we call them *snapshot entities* (*snapshot relations*). Temporary entities and relations are captured by the following $\mathcal{DLR}_{\mathcal{US}}$-formulas:

$$E \sqsubseteq^* (\Diamond^+ \neg E) \sqcup (\Diamond^- \neg E) \qquad\qquad R \sqsubseteq^* (\Diamond^+ \neg R) \sqcup (\Diamond^- \neg R)$$

saying that there must be a past or a future time point where the entity (relation) does not hold. In other words, instances of temporary entities (relations) always have a limited lifetime. On the other hand, snapshot entities and relations are captured by the following $\mathcal{DLR}_{\mathcal{US}}$-formulas:

$$E \sqsubseteq^* (\Box^+ E) \sqcap (\Box^- E) \qquad\qquad R \sqsubseteq^* (\Box^+ R) \sqcap (\Box^- R)$$

saying that whenever the entity (relation) is true it is necessarily true in every past and every future time point, i.e, they never change along time. Snapshot entities and relations are used to capture the semantics of legacy non-temporal schemas when included in a temporal model, thus enforcing the upward compatibility. In our example (Fig. 2), `Employee,Department,Resp-For` could be constrained by snapshot $\mathcal{DLR}_{\mathcal{US}}$ formulas, while `Manager,Work-For` by temporary formulas.

Temporal attributes. At different points in time, an entity may have different values for the same attribute. Attributes can be forced either to remain unchanged in time (*snapshot* attribute), or to necessarily change (*temporary* attribute) with the following $\mathcal{DLR}_{\mathcal{US}}$-formulas, respectively:

$$E \sqsubseteq^* \exists^{\geq 1}[\mathtt{from}]\Box^* R_A \qquad\qquad E \sqsubseteq^* \exists^{=1}[\mathtt{From}](R_A \sqcap (\Diamond^+ \neg R_A \sqcup \Diamond^- \neg R_A))$$

When considering the interaction between the temporal behaviour of an attribute and that of the owner entity, it is consistent in $\mathcal{DLR}_{\mathcal{US}}$ to have both snapshot attributes of a temporary entity and temporary attributes of a snapshot entity. In the former case, the $\mathcal{DLR}_{\mathcal{US}}$ semantics says that during the lifespan of an entity the value of a snapshot attribute never changes. In the latter one, the meaning is that each instance always belongs to the snapshot entity but the value of the temporary attribute will change during its existence. In our running example, where `Employee` is a snapshot entity, `Salary` is modelled as a temporary attribute, while `Names,PaySlipNumber` are modelled as snapshot attributes. In particular, `PaySlipNumber` plays the role of a full fledged temporal key (see the formula for `PaySlipNumber` in Sect. 2.3).

Temporal cardinalities. Cardinality constraints limit the participation of entities in relationships. In a temporal setting, we can distinguish between *snapshot participation constraints* (true at each point in time) and *lifespan participation constraints* [43, 25, 41] (evaluated during the entire existence of the entity). While the standard $\mathcal{DLR}_{\mathcal{US}}$ cardinality construct captures snapshot participation constraints, the lifespan participation constraints are defined by the following $\mathcal{DLR}_{\mathcal{US}}$-formula:

$$E \sqsubseteq^* \exists^{\lessgtr n}[i]\Diamond^* R$$

i.e., over its lifespan, an instance of the entity E must participate as the ith argument in at least n (at most n, or precisely n) tuples of the relation R. Obviously, since for snapshot relations the set of instances does not change in time, there is no difference between snapshot and lifespan participation constraints with respect to snapshot relations. In our example, we could say for example that managers should manage at most 5 different projects in their entire existence as managers, while still being constrained in managing exactly one project at a time.

Dynamic entities. In the temporal conceptual modelling literature, two notions of dynamic transitions between a entities (also called object migrations) are considered [26, 27, 41]: *dynamic evolution*, when an object ceases to be an instance of a source entity, and *dynamic extension*, when an object continues to belong to the source. Let E_s be the source entity and E_t be the target one; the two cases are captured by the following formulas, respectively:

$$E_s \sqsubseteq^* \diamond^+(E_t \sqcap \neg E_s) \qquad\qquad E_s \sqsubseteq^* \diamond^+ E_t$$

An interesting consequence of dynamic evolution is that the source is necessarily a temporary entity.

$\mathcal{DLR}_{\mathcal{US}}$ is also able to capture *safety* and *liveness* constraints [36]. A *safety* constraint intuitively says that "*nothing bad ever happens*" and can be captured by the formula $\Box^*\neg(E \sqsubseteq \bot)$. On the other hand, *liveness* constraints saying that "*something good will happen*" can be expressed by existential temporal formulas: $\diamond^*\neg(E \sqsubseteq \bot)$.

Schema evolution. We consider here a simplified case of conceptual schema evolution [38, 20, 35], which can be called a *monotonic* approach. This allows only for changes in the schema such that the resulting conceptual schema is compatible with the previous *global* constraints.

Let Γ and ϕ_p be atomic formulas, or possibly a Boolean combination of atomic formulas, which introduce, respectively, a new schema portion and a condition to be checked. The formula $\Box^*(\neg\phi_p \vee \Box^+\Gamma)$ states that as soon as the property ϕ_p becomes true for the data, the conceptual schema will include the additional Γ constraint. A simple example is:

$\phi_{\mathtt{p}} \equiv (\mathtt{InterestGroup} \sqsubseteq \bot)$
$\Gamma \equiv (\exists^{\geq 1}[\mathtt{amount}](\mathtt{Salary} \sqcap \mathtt{payee/2} : \mathtt{TopManager}) \sqsubseteq \mathtt{LowAmount})$

meaning that as soon as the organisation does not include interest groups anymore, the salary of top managers should be in the "low amount" class.

Note that the $\mathcal{DLR}^-_{\mathcal{US}}$ with only atomic formulas (i.e., the EXPTIME-complete fragment of $\mathcal{DLR}_{\mathcal{US}}$) is enough to capture most of the modelling constructs discussed in this Section—in fact, in the case of (global) logical implication for atomic formulas, there is no difference between $\sqsubseteq$ and $\sqsubseteq^*$—with the exception of (a) schema evolution constraints, (b) safety and liveness conditions, (c) snapshot relations and attributes, and (d) temporal cardinalities. Full $\mathcal{DLR}^-_{\mathcal{US}}$ (i.e., the EXPSPACE-complete fragment of $\mathcal{DLR}_{\mathcal{US}}$) is able to express (a) and (b) as well. However, (c) and (d) require temporalised relations which, by Theorem 1, lead to undecidability.

5 Conclusion

This work introduces the temporal description logic $\mathcal{DLR}_{\mathcal{US}}$ and illustrates its expressive power. A temporal query language was defined and the problem of query containment under the constraints defined by a $\mathcal{DLR}_{\mathcal{US}}$ conceptual schema is investigated.

Tight complexity results were proved. In particular, reasoning in the full logic $\mathcal{DLR}_{\mathcal{US}}$ was shown to be undecidable, while decidability was obtained using a still expressive fragment, $\mathcal{DLR}^-_{\mathcal{US}}$. We have also shown that the problem of checking query containment of non-recursive Datalog queries under constraints with arbitrary $\mathcal{DLR}^-_{\mathcal{US}}$ formulas is decidable in 2EXPTIME with an EXPSPACE lower bound. This result is the first decidability result we are aware of on containment of temporal conjunctive queries under expressive constraints.

Acknowledgments

The work was partially funded by U.K. EPSRC grants GR/R42474/01, GR/R45369/01, GR/R04348/01, and GR/R09428/01.

References

1. S. Abiteboul, L. Herr, and J. Van den Bussche. Temporal versus first-order logic to query temporal databases. In *Proc. of the 15th ACM SIGACT SIGMOD SIGART Sym. on Principles of Database Systems (PODS'96)*, pages 49–57, 1996.
2. A. Artale and E. Franconi. Reasoning with enhanced temporal entity-relationship models. In *Proc. of the International Workshop on Spatio-Temporal Data Models and Languages*. IEEE Computer Society Press, August 1999.
3. A. Artale and E. Franconi. Temporal ER modeling with description logics. In *Proc. of the International Conference on Conceptual Modeling (ER'99)*. Springer-Verlag, November 1999.
4. A. Artale and E. Franconi. A survey of temporal extensions of description logics. *Annals of Mathematics and Artificial Intelligence*, 30(1-4), 2001.
5. A. Artale, E. Franconi, M. Mosurovic, F. Wolter, and M. Zakharyaschev. The DLRUS temporal description logic. In *Proc. of the 2001 International Workshop on Description Logics (DL'2001)*, 2001.
6. A. Artale, E. Franconi, M. Mosurovic, F. Wolter, and M. Zakharyaschev. Reasoning over conceptual schemas and queries in temporal databases. In *Proc. of the 9th Italian Database Conference (SEBD'2001)*, 2001.
7. F. Baader, D. McGuinness, D. Nardi, and P. F. Patel-Schneider, editors. *Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press, 2002.
8. M. Böhlen, C. Jensen, and R. Snodgrass. Temporal statement modifiers. *ACM Transactions On Database Systems*, 2000.
9. A. Borgida and R. J. Brachman. Conceptual modelling with description logics. In Baader et al. [7].
10. A. Borgida, M. Lenzerini, and R. Rosati. Description logics for databases. In Baader et al. [7].
11. A. Calì D. Calvanese, G. De Giacomo, and M. Lenzerini. A formal framework for reasoning on UML class diagrams. In *Proc. of the 13th Int. Sym. on Methodologies for Intelligent Systems (ISMIS 2002)*, 2002.
12. D. Calvanese, G. De Giacomo, and M. Lenzerini. On the decidability of query containment under constraints. In *Proc. of the 17th ACM SIGACT SIGMOD SIGART Sym. on Principles of Database Systems (PODS'98)*, pages 149–158, 1998.
13. D. Calvanese, M. Lenzerini, and D. Nardi. Description logics for conceptual data modeling. In Chomicki and Saake [16].
14. D. Calvanese, M. Lenzerini, and D. Nardi. Unifying class-based representation formalisms. *J. of Artificial Intelligence Research*, 11:199–240, 1999.
15. J. Chomicki. Temporal query languages: a survey. In *Proc. of the 1st International Conference on Temporal Logic (ICTL'94)*, pages 506–534, 1994.
16. J. Chomicki and G. Saake, editors. *Logics for Databases and Information Systems*. Kluwer, 1998.
17. J. Chomicki and D. Toman. Temporal logic in information systems. In Chomicki and Saake [16], chapter 1.
18. R. Elmasri and S. B. Navathe. *Fundamentals of Database Systems*. Benjamin/Cummings, 2nd edition, 1994.
19. D. Fensel, I. Horrocks, F. van Harmelen, S. Decker, M. Erdmann, and M. Klein. OIL in a nutshell. In *Proceedings of the European Knowledge Acquisition Conference (EKAW-2000)*, Lecture Notes In Artificial Intelligence. Springer-Verlag, 2000.
20. E. Franconi, F. Grandi, and F. Mandreoli. A semantic approach for schema evolution and versioning in object-oriented databases. In *Proc. of the 1st International Conf. on Computational Logic (CL'2000), DOOD stream*. Springer-Verlag, July 2000.
21. E. Franconi and G. Ng. The ICOM tool for intelligent conceptual modelling. In *Proc. of the 7^{th} International Workshop on Knowledge Representation meets Databases (KRDB'2000)*, 2000.
22. D. Gabbay, A.Kurucz, F. Wolter, and M. Zakharyaschev. *Many-Dimensional Modal Logics: Theory and Applications*. Studies in Logic. Elsevier, 2002. To appear.
23. D. M. Gabbay, I. Hodkinson, and M. Reynolds. *Temporal Logic: Mathematical Foundations and Computational Aspects*. Oxford University Press, 1994.
24. H Gregersen and J. S. Jensen. Temporal Entity-Relationship models - a survey. *IEEE Transactions on Knowledge and Data Engineering*, 11(3):464–497, 1999.
25. H. Gregersen and J.S. Jensen. Conceptual modeling of time-varying information. Technical Report TimeCenter TR-35, Aalborg University, Denmark, 1998.

26. R. Gupta and G. Hall. Modeling transition. In *Proc. of ICDE'91*, pages 540–549, 1991.
27. R. Gupta and G. Hall. An abstraction mechanism for modeling generation. In *Proc. of ICDE'92*, pages 650–658, 1992.
28. I. Hodkinson, F. Wolter, and M. Zakharyaschev. Decidable fragments of first-order temporal logics. *Annals of Pure and Applied Logic*, 106:85–134, 2000.
29. I. Horrocks, U. Sattler, S. Tessaris, and S. Tobies. How to decide query containment under constraints using a description logic. In *Proceedings of the 7th International Conference on Logic for Programming and Automated Reasoning (LPAR'00)*, 2000.
30. I. Horrocks, U. Sattler, and S. Tobies. Practical reasoning for expressive description logics. In *Proceedings of the 6th International Conference on Logic for Programming and Automated Reasoning (LPAR'99)*, pages 161–180, 1999.
31. C. Welty, I. Horrocks, D. McGuinness. Digital libraries and web based information systems. In Baader et al. [7].
32. M. Jarke, C. Quix, D. Calvanese, M. Lenzerini, E. Franconi, S. Ligoudistiano, P. Vassiliadis, and Y. Vassiliou. Concept based design of data warehouses: The DWQ demonstrators. In *2000 ACM SIGMOD Intl. Conference on Management of Data*, 2000.
33. C. S. Jensen and R. T. Snodgrass. Temporal data management. *IEEE Transactions on Knowledge and Data Engineering*, 111(1):36–44, 1999.
34. A. Y. Levy and M-C. Rousset. Combining horn rules and description logics in CARIN. *Artificial Intelligence*, 104(1-2):165–209, 1998.
35. R. J. Peters and M. T. Özsu. An Axiomatic Model of Dynamic Schema Evolution in Objectbase Systems. *ACM Transactions On Database Systems*, 22(1):75–114, 1997.
36. A. Pnueli. Application of temporal logic to the specification and verification of reactive systems: A survey of current trends. In *Current Trends in Concurrency*, volume 224 of *LNCS*, pages 510–584. Springer-Verlag, 1986.
37. R. Robinson. Undecidability and non-periodicity for tilings of the plan. *Invent. Math.*, 12:177–209, 1971.
38. J. F. Roddick. A Survey of Schema Versioning Issues for Database Systems. *Information and Software Technology*, 37(7):383–393, 1995.
39. K. Schild. Combining terminological logics with tense logic. In *Proceedings of the 6^{th} Portuguese Conference on Artificial Intelligence, EPIA'93*, October 1993.
40. A. S. Sistla and E. M. Clarke. The complexity of propositional linear temporal logics. *J. Assoc. Comput. Mach.*, 32(3):733–749, 1985.
41. S. Spaccapietra, C. Parent, and E. Zimanyi. Modeling time from a conceptual perspective. In *Int. Conf. on Information and Knowledge Management (CIKM98)*, 1998.
42. H. Sturm and F. Wolter. A tableau calculus for temporal description logic: the expanding domain case. *Journal of Logic and Computation*, 2001. To appear.
43. B. Tauzovich. Towards temporal extensions to the entity-relationship model. In *Proc. of the 10^{th} International Conference on the Entity-Relationship Approach (ER'91)*, pages 163–179, 1991.
44. C. Theodoulidis, P. Loucopoulos, and B. Wangler. A conceptual modelling formalism for temporal database applications. *Information Systems*, 16(3):401–416, 1991.
45. P. van Emde Boas. The convenience of tiling. Technical Report CT-96-01, ILLC—Univ. of Amsterdam, 1996.
46. F. Wolter and M. Zakharyaschev. Satisfiability problem in description logics with modal operators. In *Proc. of the 6 th International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*, pages 512–523, Trento, Italy, June 1998.
47. F. Wolter and M. Zakharyaschev. Modal description logics: Modalizing roles. *Fundamenta Informaticae*, 39(4):411–438, 1999.
48. F. Wolter and M. Zakharyaschev. Multi-dimensional description logics. In *Proc. of IJCAI'99*, pages 104–109, 1999.
49. F. Wolter and M. Zakharyaschev. Temporalizing description logics. In D. Gabbay and M. de Rijke, editors, *Frontiers of Combining Systems*, pages 379–401. Studies Press-Wiley, 1999.