

BIROn - Birkbeck Institutional Research Online



Artale, A. and Gnatenko, A. and Ryzhikov, Vladislav and Zakhariyashchev, Michael (2025) On deciding the data complexity of answering Linear Monadic Datalog Queries with LTL Operators. In: Roy, S. and Kara, A. (eds.) 28th International Conference on Database Theory (ICDT 2025). Leibniz International Proceedings in Informatics (LIPIcs) 328. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 1-19. ISBN 9783959773645.

Downloaded from: <https://eprints.bbk.ac.uk/id/eprint/55539/>



Usage Guidelines:

Please refer to usage guidelines at <https://eprints.bbk.ac.uk/policies.html> or alternatively contact lib-eprints@bbk.ac.uk.

On Deciding the Data Complexity of Answering Linear Monadic Datalog Queries with LTL Operators

Alessandro Artale  


Faculty of Engineering, Free University of Bozen-Bolzano, Italy

Anton Gnatenko  

Faculty of Engineering, Free University of Bozen-Bolzano, Italy

Vladislav Ryzhikov  

Birkbeck, University of London, UK

Michael Zakharyashev  

Birkbeck, University of London, UK

Abstract

Our concern is the data complexity of answering linear monadic datalog queries whose atoms in the rule bodies can be prefixed by operators of linear temporal logic *LTL*. We first observe that, for data complexity, answering any connected query with operators \bigcirc/\bigcirc^- (at the next/previous moment) is either in AC^0 , or in $ACC^0 \setminus AC^0$, or NC^1 -complete, or L-hard and in NL. Then we show that the problem of deciding L-hardness of answering such queries is PSPACE-complete, while checking membership in the classes AC^0 and ACC^0 as well as NC^1 -completeness can be done in EXPSpace. Finally, we prove that membership in AC^0 or in ACC^0 , NC^1 -completeness, and L-hardness are undecidable for queries with operators \Diamond/\Diamond^- (sometime in the future/past) provided that $NC^1 \neq NL$ and $L \neq NL$.

2012 ACM Subject Classification Theory of computation \rightarrow Database query processing and optimization (theory)

Keywords and phrases Linear monadic datalog, linear temporal logic, data complexity

Digital Object Identifier 10.4230/LIPIcs.ICDT.2025.31

Related Version *Extended version incl. full proofs:* <https://arxiv.org/abs/2501.13762> [4]

1 Introduction

We consider monadic datalog queries, in which atoms in the rule bodies can be prefixed by the temporal operators \bigcirc/\bigcirc^- (at the next/previous moment) and \Diamond/\Diamond^- (sometime in the future/past) of linear temporal logic *LTL* [17]. This query language, denoted $datalog_m^{\bigcirc\Diamond}$, is intended for querying temporal graph databases and knowledge graphs in scenarios such as virus transmission [25, 45], transport networks [28], social media [38], supply chains [46], and power grids [37]. In this setting, data instances are finite sets of ground atoms that are timestamped by the moments of time they happen at. The rules in $datalog_m^{\bigcirc\Diamond}$ queries are assumed to hold at all times, with time being implicit in the rules and only accessible via temporal operators. We choose *LTL* for our formalism rather than, say, more expressive metric temporal logic *MTL* [30, 1] because *LTL* has been a well established query language in temporal databases since the 1990s (see [39, 41, 14, 32] and the discussion therein on point versus interval-based query languages), also suitable in the context of temporal knowledge graphs as recently argued in [19].



© Alessandro Artale, Anton Gnatenko, Vladislav Ryzhikov, and Michael Zakharyashev;
licensed under Creative Commons License CC-BY 4.0

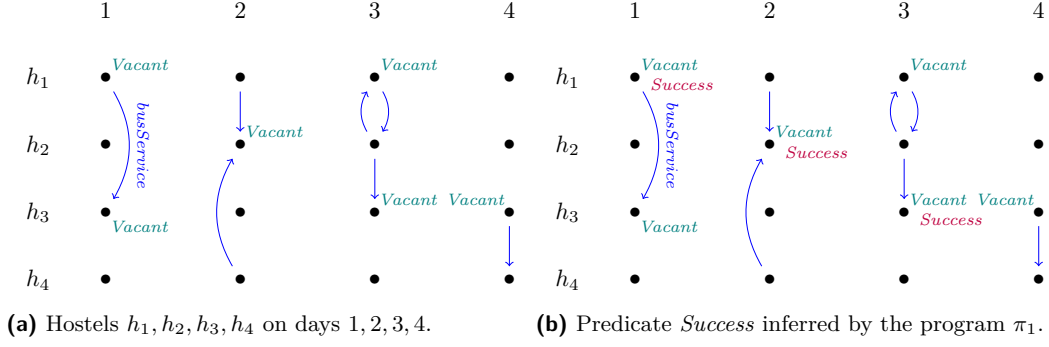
28th International Conference on Database Theory (ICDT 2025).

Editors: Sudeepa Roy and Ahmet Kara; Article No. 31; pp. 31:1–31:19

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** Illustrations for the query $(\pi_1, Success)$.

► **Example 1.** Imagine a PhD student working on a paper while hostel hopping. However, finishing the paper requires staying at the same hostel for at least two consecutive nights. Bus services between hostels, which vary from one day to the next, and hostel vacancies are given by a temporal data instance with atoms of the form $busService(a, b, n)$ and $Vacant(a, n)$, where a, b are hostels and $n \in \mathbb{N}$ a timestamp (see Fig. 1 for an illustration). The following $datalog_m^{\circ\Diamond}$ query $(\pi_1, Success)$ finds pairs (x, t) such that having started hopping at hostel x on day t , the student will eventually submit the paper:

$$\begin{aligned} \pi_1: \quad & Success(X) \leftarrow Vacant(X) \wedge \bigcirc busService(X, Y) \wedge \bigcirc Success(Y), \\ & Success(X) \leftarrow Vacant(X) \wedge \bigcirc Vacant(X). \end{aligned}$$

It is readily seen that answering this query is NL-complete for data complexity. If, however, we drop the next-time operator \bigcirc from π_1 , it will become equivalent to $Vacant(X)$. The next query $(\pi_2, Promising)$ simply looks for pairs (x, t) with x having vacancies for two consecutive nights some time later than t :

$$\begin{aligned} \pi_2: \quad & Promising(X) \leftarrow \Diamond Vacant42Nights(X), \\ & Vacant42Nights(X) \leftarrow Vacant(X) \wedge \bigcirc Vacant(X). \end{aligned}$$

This $datalog_m^{\circ\Diamond}$ query can be equivalently expressed as the two-sorted first-order formula

$$\exists t' ((t < t') \wedge Vacant(x, t') \wedge Vacant(x, t' + 1)),$$

where x ranges over objects (hostels) and t, t' over time points ordered by $<$. Formulas in such a two-sorted first-order logic, denoted $FO(<)$, can be evaluated over finite data instances in AC^0 for data complexity [26]. \lrcorner

Our main concern is the classical problem of deciding whether a given temporal monadic datalog query is equivalent to a first-order query (over any data instance). In the standard, atemporal database theory, this problem, known as predicate boundedness, has been investigated since the mid 1980s with the aim of optimising and parallelising datalog programs [27, 16]. Thus, predicate boundedness was shown to be undecidable for binary datalog queries [24] and 2EXPTIME-complete for monadic ones (even with a single recursive rule) [15, 8, 29].

Datalog boundedness is closely related to the more general rewritability problem in the ontology-based data access paradigm [11, 3], which brought to light wider classes of ontology-mediated queries (OMQs) and ultimately aimed to decide the data complexity of

answering any given OMQ and, thereby, the optimal database engine needed to execute that OMQ. Answering OMQs given in propositional linear temporal logic *LTL* is either in AC^0 , or in $ACC^0 \setminus AC^0$, or NC^1 -complete for data complexity [40], the classes well known from the circuit complexity theory for regular languages. For each of these three cases, deciding whether a given *LTL*-query falls into it is $EXPSPACE$ -complete, even if we restrict the language to temporal Horn formulas and atomic queries [31]. The data complexity of answering atemporal monadic datalog queries comes from four complexity classes $AC^0 \subsetneq L \subseteq NL \subseteq P$ [16, 15].

Our 2D query language $datalog_m^{\circ\Diamond}$ is a combination of datalog and *LTL*. It can be seen as the monadic fragment, without negation and aggregate functions, of $DEDALUS_0$, a language for reasoning about distributed systems that evolve in time [2]. It is also close in spirit to temporal deductive databases, TDDs [13, 12], extending their monadic fragment with the eventuality operator \Diamond . The main intriguing question we would like to answer in this paper is whether deciding membership of 2D queries in, say, AC^0 can be substantially harder than deciding membership in AC^0 of the corresponding 1D monadic datalog and *LTL* queries.

With this in mind, we focus on the sublanguage $datalog_{lm}^{\circ\Diamond}$ of $datalog_m^{\circ\Diamond}$ that consists of linear queries, that is, those that have at most one IDB (intensional, recursively definable) predicate in each rule body. While the full language inherits from TDDs a $PSPACE$ -complete query answering problem (for data complexity), we prove that linear queries can be answered in NL . By $datalog_{lm}^{\circ}$ and $datalog_{lm}^{\Diamond}$ we denote the fragments of $datalog_{lm}^{\circ\Diamond}$ that only admit the \circ/\circ^- and \Diamond/\Diamond^- operators, respectively. All of our queries are assumed to be connected in the sense that the graph induced by the body of each rule is connected. These fragments retain practical interest: as argued in [15], atemporal datalog programs used in practice tend to be linear and connected. For example, SQL:1999 explicitly supports linear recursion [20], which together with connectedness is a common constraint in the context of querying graph databases and knowledge graphs [35, 42], where the focus is on path queries [21, 23].

It is known that deciding whether a linear monadic datalog query can be answered in AC^0 is $PSPACE$ -complete [15, 43] (without the monadicity restriction, the problem is undecidable [24]). The same problem for the propositional *LTL* fragments of $datalog_{lm}^{\circ}$ and $datalog_{lm}^{\Diamond}$ is also $PSPACE$ -complete [31].

Our main results in this paper are as follows:

- Answering $datalog_{lm}^{\circ\Diamond}$ queries is NL -complete for data complexity.
- It is undecidable whether a given $datalog_{lm}^{\Diamond}$ -query can be answered in AC^0 , ACC^0 , or NC^1 (if $NC^1 \neq NL$); it is undecidable whether such a query is L -hard (if $L \neq NL$).
- Answering any connected $datalog_{lm}^{\circ}$ -query is either in AC^0 or in $ACC^0 \setminus AC^0$, or NC^1 -complete, or L -hard – and anyway it is in NL .
- It is $PSPACE$ -complete to decide whether a connected $datalog_{lm}^{\circ}$ -query is L -hard; checking whether it is in AC^0 , or in $ACC^0 \setminus AC^0$, or is NC^1 -complete can be done in $EXPSPACE$.

(Note that dropping the past-time operators \circ^- and \Diamond^- from the languages has no impact on these complexity results.) Thus, the temporal operators \circ/\circ^- and \Diamond/\Diamond^- exhibit drastically different types of interaction between the object and temporal dimensions. To illustrate the reason for this phenomenon, consider first the following $datalog_{lm}^{\circ}$ -program:

$$G(X) \leftarrow A(X) \wedge \circ R(X, Y) \wedge \circ D(Y), \quad (1)$$

$$D(X) \leftarrow \circ D(X), \quad (2)$$

$$D(X) \leftarrow B(X). \quad (3)$$

Suppose a data instance consists of timestamped atoms $A(a, 0)$, $R(a, b, 1)$, $B(b, 5)$. We obtain $G(a, 0)$ by first applying rule (3) to infer $D(b, 5)$, then rule (2) to infer $D(b, 4)$, $D(b, 3)$, $D(b, 2)$, and $D(b, 1)$, and finally rule (1) to obtain $G(a, 0)$. Rules (2) and (3) are applied along the

timeline of a single object, b , while the final application passes from one object, b , to another, a . To do so, we check whether a certain condition holds for the joint timeline of a and b , namely, that they are connected by R at time 1. However, if we are limited to the operators \bigcirc/\bigcirc^- , the number of steps that we can investigate along such a joint timeline is bounded by the maximum number of nested \bigcirc/\bigcirc^- in the program. Therefore, there is little interaction between the two phases of inference that explore the object and the temporal domains. In contrast, rules with \Diamond/\Diamond^- can inspect both dimensions simultaneously as, for example, the rule

$$G(X) \leftarrow \Diamond R(X, Y) \wedge D(Y). \quad (4)$$

In this case, inferring $G(a, 0)$ requires checking the existence of an object b satisfying $D(b, 0)$ and $R(a, b, \ell)$ at some *arbitrarily distant* moment ℓ in the future. Given that our programs are monadic, the predicate $\Diamond R(X, Y)$ cannot be expressed using operators \bigcirc/\bigcirc^- only.

Our positive results are proved by generalising the automata-theoretic approach of [15]. As a by-product, we obtain a method to decompose every connected $\text{datalog}_{lm}^\bigcirc$ -query (π, G) into a plain datalog part (π_d, G) and a plain *LTL* part (π_t, G) , which are, however, substantially larger than (π, G) , so that the data complexity of answering (π, G) equals the maximum of the respective data complexities of (π_d, G) and (π_t, G) . This reinforces the “weakness of interaction” between the relational and the temporal parts of the query when the latter is limited to operators \bigcirc/\bigcirc^- . We also provide some evidence that, in contrast to the atemporal case, the automata-theoretic approach cannot be generalised to the case of disconnected queries. The undecidability of the decision problem for $\text{datalog}_{lm}^\Diamond$ -queries is proved by a reduction of the halting problem for Minsky machines with two counters [33].

The paper is organised as follows. In Section 2, we give formal definitions of data instances and queries, and prove that every temporal monadic datalog query can be answered in P, and in NL if it is linear. In Section 3, we show that checking whether a given query with operator \Diamond or \Diamond^- has data complexity lower than NL is undecidable. Section 4 considers $\text{datalog}_{lm}^\bigcirc$ -queries by presenting a generalisation of the automata-theoretic approach of [15], which is then used in Section 5 to provide the decidability results. We conclude with a discussion of future work and our final remarks in Section 6. Detailed proofs can be found in the full version of this paper [4].

2 Preliminaries

A *relational schema* Σ is a finite set of *relation symbols* R with associated arities $m \geq 0$. A *database* D over a schema Σ is a set of ground atoms $R(d_1, \dots, d_m)$, $R \in \Sigma$, m is the arity of R . We call d_i , $1 \leq i \leq m$, the *domain objects* or simply objects. We denote by Δ_D the set of objects occurring in D . We denote by $|D|$ the number of atoms in D . We denote by $[a, b]$ the set of integers $\{m \mid a \leq m \leq b\}$, where $a, b \in \mathbb{Z}$. A *temporal database* \mathcal{D} over a schema Σ is a finite sequence $\langle D_l, D_{l+1}, \dots, D_{r-1}, D_r \rangle$ of databases over this schema for some $l < r$, $l, r \in \mathbb{Z}$. Each database D_i , $l \leq i \leq r$, is called the i 'th *slice* of \mathcal{D} and i is called a *timestamp*. We denote $[l, r]$ by $\text{tem}(\mathcal{D})$. The *size* of \mathcal{D} , denoted by $|\mathcal{D}|$, is the maximum between $|\text{tem}(\mathcal{D})|$ and $\max\{|D_l|, \dots, |D_r|\}$. The domain of the temporal database \mathcal{D} is $\bigcup_{l \leq i \leq r} \Delta_{D_i}$ and is denoted by $\Delta_{\mathcal{D}}$. A homomorphism from \mathcal{D} as above to $\mathcal{D}' = \langle D_{l'}, D_{l'+1}, \dots, D_{r'-1}, D_{r'} \rangle$ is a function h that maps $\Delta_{\mathcal{D}} \cup [l, r]$ to $\Delta_{\mathcal{D}'} \cup [l', r']$ so that $R(d_1, \dots, d_m) \in D_\ell$ if and only if $R(h(d_1), \dots, h(d_m)) \in D'_{h(\ell)}$.

We will deal with *temporal conjunctive queries* (temporal CQs) that are formulas of the form $Q(\overline{X}) = \exists \overline{U} \varphi(\overline{X}, \overline{U})$, where $\overline{X}, \overline{U}$ are tuples of variables and φ , the *body* of the query, is defined by the following BNF:

$$\varphi ::= R(Z_1, \dots, Z_m) \mid (\varphi \wedge \varphi) \mid \mathcal{O}\varphi \quad (5)$$

where $R \in \Sigma$ and Z_1, \dots, Z_m are variables of $\overline{X} \cup \overline{U}$, and \mathcal{O} is any of \circ, \circ^-, \diamond and \diamond^- (the operators \circ/\circ^- mean “at the next/previous moment” and \diamond/\diamond^- “sometime in the future/past”). For brevity, we will use the notation \mathcal{O}^n , $\mathcal{O} \in \{\circ, \diamond\}$, for a sequence of n symbols \mathcal{O} if $n > 0$, of $|n|$ symbols $\mathcal{O}^- \in \{\circ^-, \diamond^-\}$ if $n < 0$, and an empty sequence if $n = 0$. We call \overline{X} the *answer variables* of Q . To provide the semantics, we define $\mathcal{D}, \ell \models \varphi(d_1, \dots, d_m)$ for $\ell \in \mathbb{Z}$ and $d_1, \dots, d_m \in \Delta_{\mathcal{D}}$ as follows:

$$\mathcal{D}, \ell \models R(d_1, \dots, d_m) \iff \ell \in [l, r] \text{ and } R(d_1, \dots, d_m) \in D_\ell \quad (6)$$

$$\mathcal{D}, \ell \models \varphi_1 \wedge \varphi_2 \iff \mathcal{D}, \ell \models \varphi_1 \text{ and } \mathcal{D}, \ell \models \varphi_2 \quad (7)$$

$$\mathcal{D}, \ell \models \circ\varphi \iff \mathcal{D}, \ell + 1 \models \varphi \quad (8)$$

$$\mathcal{D}, \ell \models \diamond\varphi \iff \mathcal{D}, \ell' \models \varphi \text{ for some } \ell' > \ell \quad (9)$$

and symmetrically for \circ^- and \diamond^- . Given a temporal database \mathcal{D} , a timestamp $\ell \in \mathbb{Z}$, and a query $Q(\overline{X}) = \exists \overline{U} \varphi(\overline{X}, \overline{U})$, we say that $\mathcal{D}, \ell \models Q(d_1, \dots, d_k)$ if there exist $\delta_1, \dots, \delta_s \in \Delta_{\mathcal{D}}$ such that $\mathcal{D}, \ell \models \varphi(d_1, \dots, d_k, \delta_1, \dots, \delta_s)$, where $k = |\overline{X}|$, $s = |\overline{U}|$.

The problem of *answering* a temporal CQ Q is to check, given \mathcal{D} , $\ell \in \text{tem}(\mathcal{D})$, and $\bar{d} = \langle d_1, \dots, d_k \rangle$, whether $\mathcal{D}, \ell \models Q(\bar{d})$. Answering temporal CQs is not harder than that for non-temporal CQs. Indeed, we show that any Q is FO($<$)-rewritable in the sense that there exists an FO($<$)-formula $\psi(\overline{X}, t)$ such that for all ℓ and \bar{d} as above $\mathcal{D}, \ell \models Q(\bar{d})$ whenever $\psi(\bar{d}, \ell)$ is true in the two-sorted first-order structure $\mathfrak{S}_{\mathcal{D}}$, whose domain is $\Delta_{\mathcal{D}} \cup \text{tem}(\mathcal{D})$, and where $R(\bar{d}, \ell)$ is true whenever $\mathcal{D}, \ell \models R(\bar{d})$ and $(\ell < \ell')$ is true whenever $\ell < \ell'$ (see [5] for details). It follows by [26] that the problem of answering a temporal CQ Q is in AC^0 .

We outline the construction of the rewriting. Let $Q(\overline{X}) = \exists \overline{U} \varphi(\overline{X}, \overline{U})$. The main issue with the construction is that temporal subformulas of φ , say $\diamond\varphi$, may be true for \bar{d} at $\ell \in \text{tem}(\mathcal{D})$, because $\mathcal{D}, \ell' \models \varphi(\bar{d})$ for $\ell' \notin \text{tem}(\mathcal{D})$. Had that not been the case, we could construct the rewriting for Q straightforwardly by induction of φ . To overcome this, let N be the number of temporal operators in φ . We use a property that for all tuples \bar{d} of objects from $\Delta_{\mathcal{D}}$ and subformulas φ of φ , we have

$$\begin{aligned} \mathcal{D}, r + N + 1 \models \varphi(\bar{d}) &\iff \mathcal{D}, \ell \models \varphi(\bar{d}) \text{ for all } \ell > r + N \\ \mathcal{D}, l - N - 1 \models \varphi(\bar{d}) &\iff \mathcal{D}, \ell \models \varphi(\bar{d}) \text{ for all } \ell < l - N \end{aligned} \quad (10)$$

Thus, for any subformula φ of φ , we construct, by induction, the formulas $\psi_{\varphi}(\bar{Z}, t)$ and $\psi_{\varphi}^i(\bar{Z})$ for $i \in [-N - 1, \dots, -1] \cup [1, \dots, N + 1]$, so that for any \mathcal{D} , $\text{tem}(\mathcal{D}) = [l, r]$, and any objects $\bar{d} \in \Delta_{\mathcal{D}}^{|\bar{Z}|}$,

$$\begin{aligned} \mathcal{D}, \ell \models \varphi(\bar{d}) &\iff \mathfrak{S}_{\mathcal{D}} \models \psi_{\varphi}(\bar{d}, \ell) && \text{for } \ell \in [l, r], \\ \mathcal{D}, (r + i) \models \varphi(\bar{d}) &\iff \mathfrak{S}_{\mathcal{D}} \models \psi_{\varphi}^i(\bar{d}) && \text{for } 1 \leq i \leq N + 1, \\ \mathcal{D}, (l + i) \models \varphi(\bar{d}) &\iff \mathfrak{S}_{\mathcal{D}} \models \psi_{\varphi}^i(\bar{d}) && \text{for } -N - 1 \leq i \leq -1. \end{aligned}$$

For the base case, we set $\psi_R(\bar{Z}, t) = R(\bar{Z}, t)$ and $\psi_R^i(\bar{Z}) = \perp$. For an induction step, e.g., $\psi_{\circ\varphi}(\bar{Z}) = \psi_{\varphi}(\bar{Z}, \min)$, $\psi_{\circ\varphi}^{N+1}(\bar{Z}) = \psi_{\varphi}^{N+1}(\bar{Z})$, $\psi_{\diamond\varphi}(\bar{Z}) = \psi_{\varphi}^{i+1}(\bar{Z})$ for all other i , and finally $\psi_{\diamond\varphi}(\bar{Z}, t) = \exists t' ((t' = t + 1) \wedge \psi_{\varphi}(\bar{Z}, t')) \vee ((t = \max) \wedge \psi_{\varphi}^1(\bar{Z}))$. Here, \min and \max are defined in FO($<$) as, respectively, $<$ -minimal and $<$ -maximal elements in $\text{tem}(\mathcal{D})$ ($(t' = t + 1)$ is FO($<$)-definable as well). The required rewriting of Q is then the formula $\exists \overline{U} \psi_{\varphi}(\overline{X}, \overline{U}, t)$.

► **Proposition 2.** *For a temporal CQ $Q(X_1, \dots, X_k)$, checking $\mathcal{D}, \ell \models Q(d_1, \dots, d_k)$ is in AC^0 for data complexity.*

In the non-temporal setting, a body of a CQ (a conjunction of atoms), can be seen as a database (a set of atoms). We have a similar correspondence in the temporal setting, for queries without operators \diamond and \diamond^- . Indeed, observe that $\circ^a(\varphi_1 \wedge \varphi_2) \equiv \circ^a\varphi_1 \wedge \circ^a\varphi_2$ and $\circ\circ^-\varphi \equiv \circ^-\circ\varphi \equiv \varphi$. Hence we can assume that any temporal CQ body is a conjunction of temporalised atoms of the form $\circ^k R(Z_1, \dots, Z_m)$. Given a temporal CQ Q of this form, let l be the least and r the greatest number such that $\circ^l R(\bar{Z})$ and $\circ^r R'(\bar{Z}')$ appear in Q , for some R, R', \bar{Z} and \bar{Z}' . Let \mathcal{D}_Q be a temporal database whose objects are the variables of Q , and $\text{tem}(\mathcal{D})$ equals $[l, r]$ if $0 \in [l, r]$, $[0, r]$ if $0 < l$, and $[l, 0]$ if $r < 0$. Furthermore, let $R(\bar{Z}) \in D_\ell$ whenever $\circ^\ell R(\bar{Z})$ is in Q , $\ell \in \text{tem}(\mathcal{D})$. Then we can, just as in the non-temporal case, characterise the relation \models in terms of homomorphisms:

► **Lemma 3.** *For any temporal CQ $Q(X_1, \dots, X_k)$ without \diamond and \diamond^- , $\mathcal{D}, \ell \models Q(d_1, \dots, d_k)$ if and only if there is a homomorphism h from \mathcal{D}_Q to \mathcal{D} such that $h(X_i) = d_i, 1 \leq i \leq k$, and $h(0) = \ell$.*

2.1 Temporal Datalog

We define a temporalised version of datalog to be able to use recursion in querying temporal databases. We call this language *temporal datalog*, or *datalog* $^{\circ\Diamond}$. A rule of *datalog* $^{\circ\Diamond}$ over a schema Σ has the form:

$$C(\bar{X}) \leftarrow \mathcal{O}^* R_1(\bar{U}_1) \wedge \dots \wedge \mathcal{O}^* R_s(\bar{U}_s) \quad (11)$$

where R_i and C are relation symbols over Σ and \mathcal{O}^* is an arbitrary sequence of temporal operators \circ, \circ^-, \diamond and \diamond^- . The part of the rule to the left of the arrow is called its *head* and the right-hand side – its *body*. All variables from the head must appear in the body.

A *program* is a finite set of rules. The relations that appear in rule heads constitute its IDB schema, $IDB(\pi)$, while the rest form the EDB schema, $EDB(\pi)$. A rule is linear if its body contains at most one IDB atom and *monadic* if the arity of its head is 1. A program π is linear (monadic) if so are all its rules. We say that the program is in plain datalog if it does not use the temporal operators and in plain LTL if all its relations have arity 1 and every rule uses just one variable. *Recursive rules* are those that contain IDB atoms in their bodies, other rules are called *initialisation rules*. The *arity* of a program is the maximal arity of its IDB atoms.

Our results are all about *connected* programs. Namely, define the *Gaifman graph* of a temporal CQ to be a graph whose nodes are the variables and where two variables are connected by an edge if they appear in the same atom. A rule body is connected if so is its Gaifman graph, and a program is connected when all rules are connected. The *size* of a program π , denoted $|\pi|$, is the number of symbols needed to write it down, where every relation symbol $R \in EDB(\pi) \cup IDB(\pi)$ is counted as one symbol, and a sequence of operators the form \mathcal{O}^k is counted as $|k|$ symbols.

When a program π is fixed, we assume that all temporal databases that we work with are defined over $EDB(\pi)$. So let π be a program and \mathcal{D} a temporal database. An *enrichment* of \mathcal{D} is an (infinite) temporal database $\mathcal{E} = \langle E_\ell \rangle_{\ell \in \mathbb{Z}}$ over the schema $EDB(\pi) \cup IDB(\pi)$ such that $\Delta_{\mathcal{E}} = \Delta_{\mathcal{D}}$ and for any $R \in EDB(\pi)$ and any $\ell \in \mathbb{Z}$, $R(d_1, \dots, d_m) \in E_\ell$ if and only if $R(d_1, \dots, d_m) \in D_\ell$. Thus, the only EDB atoms in \mathcal{E} are those of \mathcal{D} , but \mathcal{E} “enriches” \mathcal{D} with various IDB atoms at different points of time. We say that \mathcal{E} is a *model* of π and \mathcal{D} if (i) \mathcal{E} is an enrichment of \mathcal{D} ; and (ii) for any rule $C(\bar{X}) \leftarrow \psi(\bar{X}, \bar{U})$ of π , $\mathcal{E} \models C(\bar{X}) \leftarrow \psi(\bar{X}, \bar{U})$, i.e., for all $\ell \in \mathbb{Z}$ and any tuples $\bar{d} \in \Delta_{\mathcal{E}}^{|\bar{X}|}, \bar{\delta} \in \Delta_{\mathcal{E}}^{|\bar{U}|}$, $\mathcal{E}, \ell \models \psi(\bar{d}, \bar{\delta})$ implies $\mathcal{E}, \ell \models C(\bar{d})$. We write $\mathcal{D}, \pi, \ell \models C(\bar{d})$ if for every model \mathcal{E} of π and \mathcal{D} it follows that $\mathcal{E}, \ell \models C(\bar{d})$.

A *datalog*[◇] query is a pair (π, G) , where π is a *datalog*[◇] program and G an IDB atom, called the *goal predicate*. The *arity* of a query is the arity of G . Given a temporal database \mathcal{D} , a timestamp $\ell \in \text{tem}(\mathcal{D})$, a tuple $(d_1, \dots, d_k) \in \Delta_{\mathcal{D}}^k$ and a *datalog*[◇] query (π, G) of arity k , the pair $\langle (d_1, \dots, d_k), \ell \rangle$ is a *certain answer* to (π, G) over \mathcal{D} if $\mathcal{D}, \pi, \ell \models G(d_1, \dots, d_k)$. The *answering problem* for a *datalog*[◇] query (π, G) over a temporal database \mathcal{D} is that of checking, given a tuple $(d_1, \dots, d_k) \in \Delta_{\mathcal{D}}^k$, and $\ell \in \text{tem}(\mathcal{D})$, if $\langle (d_1, \dots, d_k), \ell \rangle$ is a certain answer to (π, G) over \mathcal{D} . We use the term “complexity of the query (π, G) ” to refer to the data complexity of the associated answering problem, and say, e.g., that (π, G) is complete for polynomial time (for P) or for nondeterministic logarithmic space (NL) if the answering problem for (π, G) is such.

Our main concern is how the data complexity of the query answering problem is affected by the features of π . The following theorem relies on similar results obtained for temporal deductive databases [13, 12] and temporal description logics [6, 7, 22].

► **Theorem 4.** *Answering (monadic) datalog[◇] queries is PSPACE-complete for data complexity; answering linear (monadic) datalog[◇] queries is NL-complete.*

Proof (Sketch). For full *datalog*[◇], PSPACE-completeness can be shown by reusing the techniques for temporal deductive databases [13, 12]. However, we prove that a *linear* query can be answered in NL. Indeed, fix a linear query (π, G) . Without loss of generality, we assume that temporalised IDB atoms in rule bodies of π have the form $\circ^k C(\bar{Y})$, where $|k| \leq 1$ (the cases of \diamond/\diamond^- and consecutive \circ/\circ^- can be expressed via recursion). Given a temporal database \mathcal{D} , tuples of objects \mathbf{c} and \mathbf{d} from $\Delta_{\mathcal{D}}$, and $\ell \in \mathbb{Z}$, we write $C(\mathbf{c}) \leftarrow_{\ell, k} D(\mathbf{d})$ if π has a rule $C(\bar{X}) \leftarrow \varphi(\bar{X}, \bar{Y}, \bar{U}) \wedge \circ^k D(\bar{Y})$ such that $\mathcal{D}, \ell \models \exists \bar{U} \varphi(\mathbf{c}, \mathbf{d}, \bar{U})$. Analogously, we write $C(\mathbf{c}) \leftarrow_{\ell}$ if there is an initialisation rule $C(\bar{X}) \leftarrow \varphi(\bar{X}, \bar{U})$ and $\mathcal{D}, \ell \models \exists \bar{U} \varphi(\mathbf{c}, \bar{U})$. Then, given a linear *datalog*[◇] query (π, G) , we have $\mathcal{D}, \pi, \ell \models G(\mathbf{d})$ if and only if $G = C_0$, $\mathbf{d} = \mathbf{c}_0$, and there exists a sequence

$$C_0(\mathbf{c}_0) \leftarrow_{k_0} C_1(\mathbf{c}_1) \leftarrow_{k_1} \dots \leftarrow_{k_{n-1}} C_n(\mathbf{c}_n) \quad (12)$$

such that $C_i \in \text{IDB}(\pi)$, tuples \mathbf{c}_i are from $\Delta_{\mathcal{D}}$, $\ell_0 = \ell$, $\ell_{i+1} = \ell_i + k_i$ for $0 \leq i < n$, $C_i(\mathbf{c}_i) \leftarrow_{\ell_i, k_i} C_{i+1}(\mathbf{c}_{i+1})$, and $C_n(\mathbf{c}_n) \leftarrow_{\ell_n}$. Let $\text{tem}(\mathcal{D}) = [l, r]$. Using property (10), we observe that a rule $C(\mathbf{c}) \leftarrow_{\ell, k} D(\mathbf{d})$ either holds or does not hold simultaneously for all $\ell > r + N$, where N is the number of temporal operators in π , and, similarly, for all $\ell < l - N$. Now, consider a sequence (12) and ℓ , where all $\ell_i > r + N$ (the case for $\ell < l - N$ is analogous). Any loop of the form $C_i(\mathbf{c}_i) \leftarrow_{k_i} \dots \leftarrow_{k_{j-1}} C_j(\mathbf{c}_j)$ in it with $C_i(\mathbf{c}_i) = C_j(\mathbf{c}_j)$ can be removed as long as in the resulting sequence

$$C_0(\mathbf{c}_0) \leftarrow_{k_0} C_1(\mathbf{c}_1) \leftarrow_{k_1} \dots \leftarrow_{k_{i-1}} C_i(\mathbf{c}_i) \leftarrow_{k_j} C_{j+1}(\mathbf{c}_{j+1}) \dots \leftarrow_{k_{n-1}} C_n(\mathbf{c}_n),$$

the sum of all k_t remains ≥ 0 . This allows us to convert any sequence (12) to a sequence with the same C_0 in the beginning, the same C_n in the end, and where all ℓ_i do not exceed $\ell + O(|\text{IDB}(\pi)| \cdot |\Delta_{\mathcal{D}}|^a)$, for a equal to the maximal arity of a relation in $\text{IDB}(\pi)$. This means that, for $\ell \in \text{tem}(\mathcal{D})$, we can check $\mathcal{D}, \pi, \ell \models G(\mathbf{d})$ using timestamps in the range $[l - O(|\pi| \cdot |\Delta_{\mathcal{D}}|^a), r + O(|\pi| \cdot |\Delta_{\mathcal{D}}|^a)]$. Clearly, the existence of such a derivation can be checked in NL. ◀

However, individual queries may be easier to answer than in the general case. Since *datalog*[◇] combines features of plain datalog and of linear temporal logic, its queries can correspond to a variety of complexity classes. Recall, for example, queries (π_1, Good) and $(\pi_2, \text{Satisfactory})$ from Section 1, the first of which is hard for logarithmic space (L-hard) and

the second lies in AC^0 , the class of problems decidable by unbounded fan-in, polynomial size and constant depth boolean circuits. Furthermore, by using unary relations and operator \circ , one can simulate any regular language, giving rise to queries that lie in ACC^0 , the class obtained from AC^0 by allowing “MOD m ” gates, or are complete for NC^1 , the class defined similarly for bounded fan-in polynomial circuits of logarithmic depth. Intuitively, the problems in AC^0 , ACC^0 , and NC^1 are solvable in short (constant or logarithmic) time on a parallel architecture; see [40] for more details.

In the remainder of the paper we focus on deciding the data complexity for the linear monadic fragment of $datalog^{\circ\Diamond}$, denoted $datalog_{lm}^{\circ\Diamond}$. It is well-known that a plain datalog query can be characterised via an infinite set of conjunctive queries called its expansions [34]. We define expansions for $datalog_{lm}^{\circ\Diamond}$ and use them as the main tool in our (un)decidability proofs.

2.1.1 Expansions for Linear Monadic Queries

Let π be a $datalog_{lm}^{\circ\Diamond}$ program and $Q(X)$ be a unary temporal conjunctive query with a single answer variable and containing a unique IDB atom, say $D(Y)$, from π . Let $P(X)$ be another temporal conjunctive query with a single answer variable, and let $P'(Y)$ be obtained from $P(X)$ by substituting X by Y and all other variables with fresh ones. A *composition* of Q and P , denoted $Q \circ P$, has the form of Q with $D(Y)$ substituted with $P'(Y)$. We note that the variables of Q remain present in $Q \circ P$ and X is an answer variable of $Q \circ P$. If P contains an IDB atom and $K(X)$ is another temporal conjunctive query, the composition can be extended in the same fashion to $(Q \circ P) \circ K$, and so on. Note that, up to renaming of variables, $(Q \circ P) \circ K$ and $Q \circ (P \circ K)$ are the same queries, so we will omit the brackets and write $Q \circ P \circ K$.

Expansions are compositions of rule bodies of a program π . Let B_1, B_2, \dots, B_{n-1} be such that B_i is the body of the recursive rule:

$$C_i(X) \leftarrow A_i(X, Y, U_1, \dots, U_{m_i}) \wedge \circ^{k_i} C_{i+1}(Y), \quad (13)$$

and B_n is the body of an initialization rule

$$C_n(X) \leftarrow B(X, V_1, \dots, V_{m_n}). \quad (14)$$

The composition $B_1 \circ \dots \circ B_n$ is called an *expansion* of (π, C_1) , and n is its *length*. The set of all expansions of (π, C_1) is denoted $expand(\pi, C_1)$. Moreover, let Γ_π^r be the set of all recursive rule bodies of π and Γ_π^i be the set of all initialization rule bodies. Then each expansion can be regarded (by omitting the symbol \circ) as a word in $(\Gamma_\pi^r)^* \Gamma_\pi^i$, and $expand(\pi, C_1)$ as a sublanguage of $(\Gamma_\pi^r)^* \Gamma_\pi^i$. Adopting a language-theoretic notation, we will use small Latin letters w, u, v , etc. to denote expansions. To highlight the fact that each expansion is a temporal conjunctive query with the answer variable X , we sometimes write $w(X) \in expand(\pi, C)$.

It is a direct generalization from the case of plain datalog [34] that $\mathcal{D}, \pi, \ell \models C(d)$ if and only if there exists $w(X) \in expand(\pi, C)$ such that $D_\ell \models w(d)$.

3 Undecidability for Queries with \Diamond

Our first result about deciding the complexity of a given query is negative.

► **Theorem 5.** *It is undecidable whether a given $datalog_{lm}^{\circ\Diamond}$ -query can be answered in AC^0 , ACC^0 , or NC^1 (if $NC^1 \neq NL$). It is undecidable whether the query is L-hard (if $L \neq NL$).*

The proof is by a reduction from the halting problem of 2-counter machines [33]. Namely, given a 2-counter machine M we construct a query (π_M, G) that is in AC^0 if M halts and NL-complete otherwise.

Recall, that a *2-counter machine* is defined by a finite set of states $S = \{s_0, \dots, s_n\}$, with a distinguished *initial state* s_0 , two counters able to store non-negative integers, and a transition function Θ . On each step the machine performs a *transition* by changing its state and incrementing or decrementing each counter by 1 with a restriction that their values remain non-negative. The next transition is chosen according to the current state and the values of the counters. However, the machine is only allowed to perform zero-tests on counters and does not distinguish between two different positive values. Formally, transitions are given by a partial function Θ :

$$S \times \{0, +\} \times \{0, +\} \rightarrow S \times \{-1, 0, 1\} \times \{-1, 0, 1\}. \quad (15)$$

Let $\text{sgn}(0) = 0$ and $\text{sgn}(k) = +$ for all $k > 0$. A *computation* of M is a sequence of *configurations*:

$$(s_0, a_0, b_0), (s_1, a_1, b_1), (s_2, a_2, b_2) \dots (s_m, a_m, b_m), \quad (16)$$

such that for each $i, 0 \leq i < m$, holds $\Theta(s_i, \text{sgn}(a_i), \text{sgn}(b_i)) = (s_{i+1}, \varepsilon_1, \varepsilon_2)$ and $a_{i+1} = a_i + \varepsilon_1$, $b_{i+1} = b_i + \varepsilon_2$. We assume that $a_0, b_0 = 0$ and Θ is such that $a_i, b_i \geq 0$ for all i . We call m the *length* of the computation. We say that M *halts* if $\Theta(s_m, \text{sgn}(a_m), \text{sgn}(b_m))$ is not defined in a computation. Thus, M either halts after m steps or goes through infinitely many configurations.

Let M be a 2-counter machine. We construct a connected linear query (π_M, G) using the operator \diamond only, such that its evaluation is in AC^0 if M halts, but becomes NL-complete otherwise. The construction with the operator \diamond^- instead of \diamond is symmetric.

We set the EDB schema $\Sigma = \{T, U_1, U_2\}$ of three relations which stand for “transition”, “first counter update”, and “second counter update”, respectively. Intuitively, domain elements of a temporal database represent configurations of M , and role triples of T , U_1 and U_2 , arranged according to certain rules described below, will play the role of transitions. A sequence of nodes connected by such triples will thus represent a computation of M . Our program π_M will generate an expansion along such a sequence, trying to assign to each configuration an IDB that represents a state of M , which will be possible while the placement of the connecting roles on the temporal line follows the rules of Θ . If the machine halts, there is a maximum number of steps we can make, so the check can be done in AC^0 . If M does not halt, however, it has arbitrarily long computations, and the query evaluation becomes NL-complete.

Here are the details. A configuration (s, a, b) is represented by an object d and three timestamps ℓ_0, ℓ_1, ℓ_2 such that $\ell_1 = \ell_0 + a$ and $\ell_2 = \ell_0 + b$. The values of the counters a and b are indicated, respectively, by existence of connections $U_1^{\ell_1}(d, d')$ and $U_2^{\ell_2}(d, d')$ to some object d' that is supposed to represent the next configuration in the computation. For the transition to happen, we also require $T^{\ell_0}(d, d')$. Given a computation of the form (16), the corresponding *computation path* is a pair $(\langle d_0, \dots, d_n \rangle, \ell_0)$, where for each $i, 0 \leq i < n$, there are $T^{\ell_0}(d_i, d_{i+1})$, $U_1^{\ell_0+a_i}(d_i, d_{i+1})$, and $U_2^{\ell_0+b_i}(d_i, d_{i+1})$ in the database, with an additional requirement that $a_0 = b_0 = 0$. Two types of problems may be encountered on such a path. A *representation violation* occurs when an object has more than one outgoing edge of type T , U_1 , or U_2 . A *transition violation* of type (s_i, α, β) is detected when there are two consecutive nodes d_i, d_{i+1} , $\alpha = \text{sgn}(a_i)$, $\beta = \text{sgn}(b_i)$, and $\Theta(s_i, \alpha, \beta) \neq (s_{i+1}, a_{i+1} - a_i, b_{i+1} - b_i)$. The program π_M will look for such violations. It will have an IDB relation symbol S_i per each state

31:10 Deciding the Data Complexity of LTL Monadic Datalog Queries

s_i of M . The initialisation rules allow to infer any state IDB from a representation violation, and the IDB S_i from a transition violation of type (s_i, α, β) . The recursive rules then push the state IDB up a computation path following the rules of Θ and tracing “backwards” a computation of M . Once π_M infers the initial state IDB at a position with zero counter values, we are done. It remains to give the rules explicitly.

We use a shortcut \diamond^* to mean a “reflexive” version of \diamond , i.e. $\diamond^*\varphi \equiv \diamond\varphi \vee \varphi$. Clearly, every rule \diamond^* can be rewritten to an equivalent set of rules without it. We need the rules:

$$S_i(X) \leftarrow \diamond^*RV(X), \quad 0 \leq i \leq n \quad RV(X) \leftarrow T(X, Y) \wedge \diamond T(X, Z) \quad (17)$$

$$RV(X) \leftarrow U_1(X, Y) \wedge \diamond U_1(X, Z), \quad RV(X) \leftarrow U_2(X, Y) \wedge \diamond U_2(X, Z) \quad (18)$$

to detect representation violations. For transition violations, we first define IDBs NE_c^ϵ , where $c \in 1, 2$ stand for the respective counter and $\epsilon \in \{-1, 0, 1\}$ stands for the change of that counter value in a transition. Each NE_c^ϵ detects situations when a correct transition was not executed, e.g. having $\epsilon = -1$, the timestamps ℓ and ℓ' that are marked by an outgoing U_c in consecutive configurations satisfy $\ell - 1 > \ell'$, $\ell = \ell'$, or $\ell < \ell'$, which they should not. The rules are the following:

$$NE_c^{-1}(Y) \leftarrow U_c(Y, Z) \wedge U_c(X, Y) \quad NE_c^{-1}(Y) \leftarrow U_c(Y, Z) \wedge \diamond \diamond U_c(X, Y) \quad (19)$$

$$NE_c^{-1}(Y) \leftarrow U_c(X, Y) \wedge \diamond U_c(Y, Z) \quad NE_c^1(Y) \leftarrow U_c(Y, Z) \wedge U_c(X, Y) \quad (20)$$

$$NE_c^1(Y) \leftarrow U_c(Y, Z) \wedge \diamond U_c(X, Y) \quad NE_c^1(Y) \leftarrow U_c(X, Y) \wedge \diamond \diamond U_c(Y, Z) \quad (21)$$

$$NE_c^0(Y) \leftarrow U_c(Y, Z) \wedge \diamond U_c(X, Y) \quad NE_c^0(Y) \leftarrow U_c(X, Y) \wedge \diamond U_c(Y, Z) \quad (22)$$

for $c \in \{1, 2\}$. Then, again, any state can be inferred from a violation:

$$S_i(X) \leftarrow \diamond^*NE_1^{\epsilon_1}(Y) \wedge \diamond^\alpha U_1(X, Y) \quad S_i(X) \leftarrow \diamond^*NE_2^{\epsilon_2}(Y) \wedge \diamond^\beta U_2(X, Y) \quad (23)$$

for each transition $\Theta(s_i, \alpha, \beta) = (s_j, \epsilon_1, \epsilon_2)$, and for all ϵ_1, ϵ_2 when $\Theta(s_i, \alpha, \beta)$ is not defined. Finally, we require

$$S_i(X) \leftarrow T(X, Y) \wedge \diamond^\alpha U_1(X, Y) \wedge \diamond^\beta U_2(X, Y) \wedge S_j(Y), \quad (24)$$

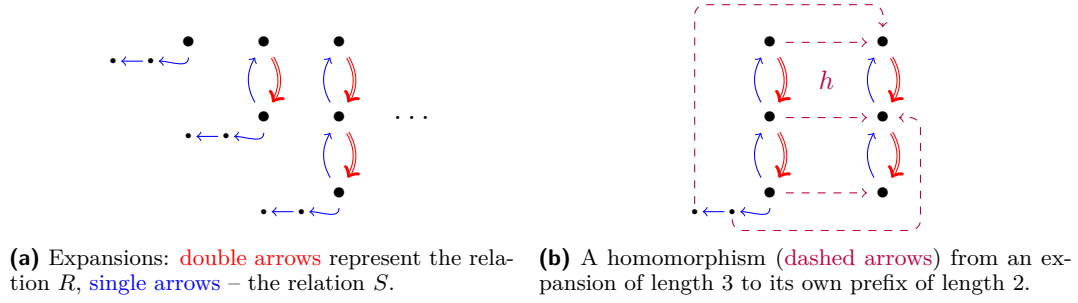
$$G(X) \leftarrow S_0(X) \wedge U_1(X, Y) \wedge U_2(X, Y). \quad (25)$$

This finalises the construction of (π_M, G) . Any expansion of (π_M, G) starts with the body of the rule (25) and then continues by a sequence of bodies of the rules of the form (24) and ends with a detection either of a representation violation, defined by (17) – (18), or of a transition violation, defined by (23). If M halts in m steps, it is enough to consider expansions containing no more than m bodies of (24). Thus, in this case (π_M, G) is in AC^0 . If, on the contrary, M does not halt, we can use expansions representing arbitrarily long prefixes of its computation for a reduction from directed reachability problem, rendering the query NL-hard.

► **Lemma 6.** *If M halts, (π_M, G) is in AC^0 . Otherwise, it is NL-hard.*

4 Automata-Theoretic Tools for Queries with \bigcirc

From now on, we focus on queries that use operators \bigcirc and \bigcirc^- only. In this section, we develop a generalisation of the automata-theoretic approach to analysing query expansions proposed in [15]. In Section 5 we use this approach to study the data complexity of this kind of queries.



■ **Figure 2** Illustrations for the query $\{(26), (28)\}, G$ of Example 7.

Recall from Section 2.1.1 the definitions of composition and expansion of rule bodies, alphabets Γ_π^r and Γ_π^i , and the language $\text{expand}(\pi, G) \subseteq (\Gamma_\pi^r)^*(\Gamma_\pi^i)$. We observe that for any sequence of rule bodies $B_1, \dots, B_{n-1} \in \Gamma_\pi^r$ and $B_n \in \Gamma_\pi^i$ the compositions $B_1 \circ \dots \circ B_{n-1}$ and $B_1 \circ \dots \circ B_n$ are well-defined. They are words in the languages $(\Gamma_\pi^r)^*$ and $(\Gamma_\pi^r)^*(\Gamma_\pi^i)$, respectively. Note that $B_1 \circ \dots \circ B_n$ is a temporal CQ over schema $EDB(\pi)$, while $B_1 \circ \dots \circ B_{n-1}$ is that over $EDB(\pi) \cup IDB(\pi)$, since it contains the IDB atom of B_{n-1} . For $w \in (\Gamma_\pi^r)^*(\Gamma_\pi^i)$, \mathcal{D}_w is defined as the (temporal) database corresponding to (temporal) CQ w , while for $w \in (\Gamma_\pi^r)^*$ we define \mathcal{D}_w as the database corresponding to the CQ obtained from w by omitting the IDB atom. For either w , \mathcal{D}_w is over the schema $EDB(\pi)$. Having that, we define the language $\text{accept}(\pi, G) \subseteq (\Gamma_\pi^r)^* \cup (\Gamma_\pi^r)^*(\Gamma_\pi^i)$ of all words $w \in (\Gamma_\pi^r)^* \cup (\Gamma_\pi^r)^*(\Gamma_\pi^i)$ such that $\mathcal{D}_w, \pi, 0 \models G(X)$.

Plain datalog queries are either in AC^0 (called *bounded*), or L-hard (*unbounded*), and a criterion of unboundedness can be formulated in language-theoretic terms [15]: a connected linear monadic plain datalog query (π, G) is *unbounded* if and only if for every k there is $w \in \text{expand}(\pi, G)$, $|w| > k$, such that its prefix of length k is *not* in $\text{accept}(\pi, G)$.

► **Example 7.** The query (π, G) , where π is given by the following plain datalog rules, is unbounded.

$$G(X) \leftarrow R(X, Y) \wedge S(Y, X) \wedge G(Y) \quad (26)$$

$$G(X) \leftarrow A(X) \quad (27)$$

However, if we substitute the rule (27) with another initialisation rule

$$G(X) \leftarrow S(X, Y) \wedge S(Y, Z) \quad (28)$$

it becomes bounded because for every $w \in \text{expand}(\pi, G)$, $|w| > 2$ its prefix of length 2 is in $\text{accept}(\pi, G)$. To see this, consider the expansions of the query, i.e., $\text{expand}(\{(26), (28)\}, G)$ given in Figure 2.

The authors of [15] construct finite state automata for languages $\text{expand}(\pi, G)$ and $\text{notaccept}(\pi, G)$, the complement of $\text{accept}(\pi, G)$, and use them to check their criterion in polynomial space. Our goal is to generalise this technique to the temporalised case. However, we can not use finite automata to work directly with sequences of rule bodies, since the language $\text{accept}(\pi, G)$, in the presence of time, may be non-regular, as demonstrated by the following example.

► **Example 8.** Consider a program

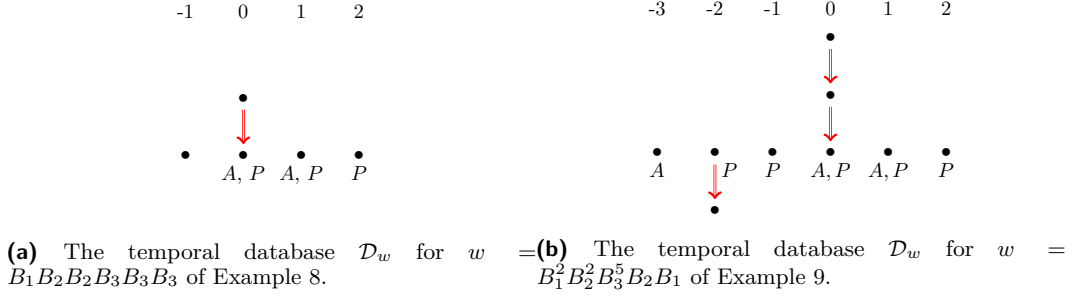
$$G(X) \leftarrow R(X, Y) \wedge G(Y)$$

$$G(X) \leftarrow P(X) \wedge \bigcirc^- G(X)$$

$$G(X) \leftarrow A(X) \wedge \bigcirc G(X)$$

$$G(X) \leftarrow P(X) \wedge R(Y, X).$$

31:12 Deciding the Data Complexity of LTL Monadic Datalog Queries



■ **Figure 3** Expansions of datalog_m° queries.

Denote its rule bodies $B_1(X, Y) = R(X, Y) \wedge G(Y)$, $B_2(X) = A(X) \wedge \bigcirc G(X)$, and $B_3(X) = P(X) \wedge \bigcirc^- G(X)$, and the composition $w = B_1 B_2 B_2 B_3 B_3 B_3$. The composition w is in $\text{accept}(\pi, G)$, and the corresponding database \mathcal{D}_w is given in Figure 3a. In general, a composition of the form $B_1 B_2^n B_3^k$ is in $\text{accept}(\pi, G)$ if and only if $n < k$, which, by a simple application of the pumping lemma, is a non-regular language.

To overcome this, we introduce a larger alphabet and define more general versions of the languages $\text{expand}(\pi, G)$ and $\text{accept}(\pi, G)$ to regain their regularity. Recall that every composition w of rule bodies gives rise to a temporal database \mathcal{D}_w . Instead of working with w , we use an exponentially larger alphabet $\Omega = 2^{\Gamma_\pi^r \cup \Gamma_\pi^i \cup \{\perp, \top\}}$ to describe \mathcal{D}_w as a word and define analogues of $\text{expand}(\pi, G)$ and $\text{accept}(\pi, G)$ over that alphabet. Consider a recursive rule

$$D(X) \leftarrow \bigcirc^{k_1} R_1(\bar{U}_1) \wedge \cdots \wedge \bigcirc^{k_s} R_s(\bar{U}_s) \wedge \bigcirc^k E(Y), \quad (29)$$

where $E(Y)$ is the unique IDB atom in the rule body and $k_i, k \in \mathbb{Z}$. We call such a rule *horizontal* if $X = Y$ and *vertical* otherwise. In a composition $w \in (\Gamma_\pi^r)^* \cup (\Gamma_\pi^r)^* \Gamma_\pi^i$, a *vertical (horizontal) segment* is a maximal subword that consists of vertical (respectively, horizontal) rule bodies. For every composition w we define the *description* $[w] \in \Omega^*$ of the respective database \mathcal{D}_w as follows. Let $w = x_1 y_1 x_2 y_2 \dots x_n y_n$, where x_i are vertical segments and y_i are horizontal segments, with x_1 and y_n possibly empty. For each $x_i = B_1 \dots B_n$ we set $[x_i]$ to be the sequence $\{B_1\} \dots \{B_n\}$ of singleton sets, each of which contains a vertical rule body. For each $y_i = B_1 \dots B_n$, we construct $[y_i]$ in n steps, as follows. Recall that B_j , $1 \leq j < n$, has the form $A_j(X, U) \wedge \bigcirc^{m_j} D(X)$, where $D(X)$ is the unique IDB atom in B_j . Let $\ell_1 = 0$ and $\ell_j = \sum_{i=1}^{j-1} m_i$ for $j \leq n$. Intuitively, ℓ_j is the moment of time where the body B_j lands in the composition $B_1 \circ \dots \circ B_n$. Let ℓ'_1, \dots, ℓ'_s be the ordering of the numbers in the set $\{\ell_j\}_{j=1}^n$ in the increasing order. We set, $\alpha_{\ell'_k}$ to be $\{B_j \mid \ell_j = \ell'_k\}$ for $\ell'_k \in \{\ell_1, \ell_n\}$; $\{B_j \mid \ell_j = \ell_1\} \cup \{\perp\}$ for $\ell'_k = \ell_1 \neq \ell_n$; $\{B_j \mid \ell_j = \ell_n\} \cup \{\top\}$ for $\ell'_k = \ell_n \neq \ell_1$; and $\{B_j \mid \ell_j = \ell_1\} \cup \{\top, \perp\}$ for $\ell'_k = \ell_1 = \ell_n$. Additionally, we set $\alpha_k = \emptyset$ for $k \in [\ell'_1, \ell'_s] \setminus \{\ell'_1, \dots, \ell'_s\}$. Now we take $[y_i] = \alpha$.

Finally, $[w] = [x_1][y_1] \dots [x_n][y_n]$. We use the symbol Λ to refer to letters of $[w]$. We call letters of the form $\{B\}$, where B is a vertical rule body, *vertical letters*, and the rest – *horizontal letters*. Consequently, we can speak of vertical and horizontal segments of $[w]$, meaning maximal segments composed of vertical (respectively, horizontal) letters only.

Intuitively, $[w]$ describes \mathcal{D}_w , which can be seen as composed of $\mathcal{D}_{x_1}, \mathcal{D}_{y_1}, \dots, \mathcal{D}_{x_n}, \mathcal{D}_{y_n}$, described by $[x_1], [y_1], \dots, [x_n], [y_n]$, respectively. The symbol \perp , representing a vertical line meeting a horizontal one, marks the point in time where \mathcal{D}_{x_i} is connected to \mathcal{D}_{y_i} , while the symbol \top , analogously, shows where \mathcal{D}_{y_i} is connected to $\mathcal{D}_{x_{i+1}}$.

► **Example 9.** Recall the rule bodies of Example 8 and consider the composition $w = B_1^2 B_2^2 B_3^5 B_2 B_1$. The corresponding \mathcal{D}_w is depicted in Figure 3b. Then $x_1 = B_1^2$, $y_1 = B_2^2 B_3^5 B_2$, $x_2 = B_1$ and y_2 is empty. Thus, $[w]$ is equal to

$$\{B_1\}\{B_1\}\{B_2\}\{\top, B_3\}\{B_3\}\{\perp, B_2, B_3\}\{B_2, B_3\}\{B_3\}\{B_1\} \quad (30)$$

Not every word over Ω correctly describes a temporal database. This motivates the following definition: $\alpha \in \Omega^*$ is *correct* if (i) every symbol Λ is either a singleton (standing for a vertical rule body) or a set of horizontal rule bodies, with a possible addition of \perp , \top , and (ii) every horizontal segment of α preceded by a vertical segment contains exactly one \perp , and every horizontal segment followed by a vertical one – exactly one \top . A correct word $\alpha \in \Omega^*$ describes a temporal database \mathcal{D}_α similarly to how $[w]$ describes \mathcal{D}_w . Formally, break α into vertical/horizontal segments $\chi_1 v_1 \dots \chi_n v_n$. For a vertical segment $\chi_i = \{B_1\} \dots \{B_n\}$, set $Q_{\chi_i} = B_1 \circ \dots \circ B_n$. For a horizontal segment $v_i = \Lambda_1 \dots \Lambda_s$, let $\Lambda_{j\perp}$ be the one containing \perp and $\Lambda_{j\top}$ – containing \top . Then Q_{v_i} is the conjunction of rule bodies from v_i , where each $B \in \Lambda_j$ is prefixed by $\circ^{j-\perp}$, plus, if $i \neq n$, an IDB atom $\circ^{j\top-j\perp} D(X)$. Finally, set $\mathcal{D}_\alpha = \mathcal{D}_{Q_\alpha}$ for $Q_\alpha = Q_{\chi_1} \circ Q_{v_1} \circ \dots \circ Q_{\chi_n} \circ Q_{v_n}$. This Q_α will be also useful further.

We are now ready to define languages over Ω that will be useful to study the data complexity of our queries. Let $Accept(\pi, G)$ be the language of correct words α such that $\mathcal{D}_\alpha, \pi, 0 \models G(X)$, with $X \in \Delta_{\mathcal{D}_\alpha}$, and $NotAccept(\pi, G)$ be its complement. We need to define the language of expansions over the alphabet Ω that we will use together with the language $NotAccept(\pi, G)$ to formulate a criterion for L-hardness similar to the one of [15]. It would be natural to take the language of expansions as $\{[w] \mid w \in expand(\pi, G)\}$. However, it is harder to define an automaton recognising such a language than it is for the language of $[w]$ s as above where each horizontal letter $[w]_i$ may be extended (as a set) with arbitrary “redundant” rule bodies, and each horizontal segment may be extended, from the left and right, by “redundant” horizontal letters. For example, we will include into the language of expansions the word $\{B_1\}\{B_1\}\{B_3\}\{B_2, B_3\}\{\top, B_3\}\{B_2, B_3\}\{\perp, B_2, B_3\}\{B_2, B_3\}\{B_3\}\{B_1\}$ alongside (30). It turns out that the latter language works for our required criteria as good as the former one. Formally, if $\alpha, \beta \in \Omega^*$, we write $\alpha \preceq \beta$ if $\alpha = x_1 y_1 \dots x_n y_n$ and $\beta = x_1 y'_1 \dots x_n y'_n$, where x_i are vertical segments and y_i, y'_i are horizontal segments, and if $y_i = a_1 \dots a_m$, $y'_i = b_1 \dots b_s$, then there is $k \geq 0$ such that $m+k \leq s$ and $a_j \subseteq b_{j+k}$, $1 \leq j \leq m$. We define $Expand_{\preceq}(\pi, G)$ to be the set of correct words $\alpha \in \Omega^*$ such that $[w] \preceq \alpha$ for some $w \in expand(\pi, G)$.

It is important for our purposes that these languages are regular. For $Expand_{\preceq}(\pi, G)$, the rules of the program π may be naturally seen as transition rules of a two-way automaton whose states are the IDBs of π (plus a final state). The initial state is the one associated with G , and the final state is reached by an application of an initialisation rule.

► **Lemma 10.** *For any datalog $_{lm}^\circ$ -query (π, G) , the language $Expand_{\preceq}(\pi, G)$ is regular.*

The case of $NotAccept(\pi, G)$ is more involved. Generalising from [15], for an automaton to recognise if $\alpha \in NotAccept(\pi, G)$ it suffices to guess (by nondeterminism) an enrichment \mathcal{E} of \mathcal{D}_α , and check that \mathcal{E} is a model of π and that \mathcal{E} does not contain the atom $G(X)$. Moreover, since π is connected, for \mathcal{E} to be a model it is enough that every piece of \mathcal{E} of radius $|\pi|$, in terms of Gaifman graph, satisfies all the rules. The idea is to precompute the answer for each such piece and encode them in the state-space of the automaton. The problem, specific to the temporalised case, is that enrichments are infinite in the temporal dimension. To resolve this, we observe that there are still finitely many EDB atoms in \mathcal{E} . Since, once again, π is connected, to check that rules with EDBs are satisfied it is enough to consider only those pieces of \mathcal{E} that contain an EDB atom. For the rest of the rules, it suffices to check if such a

finite piece *can* be extended into an infinitely in time to give a model of π . The rules without EDBs can be seen as plain *LTL* rules, so we can employ satisfiability checking for *LTL* to perform this check.

► **Lemma 11.** *For any $\text{datalog}_{lm}^\circ$ -query (π, G) , the language $\text{NotAccept}(\pi, G)$ is regular.*

► **Corollary 12.** *For any $\text{datalog}_{lm}^\circ$ -query (π, G) , the language $\text{Accept}(\pi, G)$ is regular.*

5 Decidability for Connected Linear Queries with \circ

We use the automata introduced in the previous section to prove a positive result.

► **Theorem 13.** *(i) Every connected $\text{datalog}_{lm}^\circ$ query is either in AC^0 , or in $\text{ACC}^0 \setminus \text{AC}^0$, or NC^1 -complete, or L-hard. (ii) It is PSPACE-complete to check whether such a query is L-hard; whether it belongs to AC^0 , ACC^0 , or is NC^1 -complete can be decided in EXPSpace.*

We first deal with (i). Intuitively, L-hardness is a consequence of the growth of query expansions in the relational domain. If this growth is limited, the query essentially defines a certain temporal property, which can be checked in NC^1 . Formally, given a word $\alpha \in \Omega^*$, we define the $\text{height}(\alpha)$ as the number of vertical letters in α . Then, we call a query *vertically unbounded* if for every k there is a word $\alpha \in \text{Accept}(\pi, G)$, $\text{height}(\alpha) > k$, such that every prefix of α of height k is in $\text{NotAccept}(\pi, G)$. Otherwise, the query is called *vertically bounded*.

Vertically unbounded queries can be shown to be L-hard by a direct reduction from the undirected reachability problem. Namely, take the deterministic automata for $\text{NotAccept}(\pi, G)$ and $\text{Accept}(\pi, G)$, supplied by Lemma 11 and Corollary 12, and apply the pumping lemma to obtain words ξ, v, ζ, γ such that $\text{height}(v) > 0$, $\xi v^i \zeta \in \text{NotAccept}(\pi, G)$ and $\xi v^i \zeta \gamma \in \text{Accept}(\pi, G)$, for all $i \geq 0$. Then, given a graph \mathcal{G} and two nodes s, t , use copies of \mathcal{D}_v to simulate the edges of \mathcal{G} , and attach \mathcal{D}_ξ to s and $\mathcal{D}_{\zeta\gamma}$ to t . Thus, you will obtain a temporal database $\mathcal{D}_\mathcal{G}$, where $\mathcal{D}_\mathcal{G}, \pi, 0 \models G(s)$ if and only if there is a path from s to t .

► **Lemma 14.** *If a query is vertically unbounded, then it is L-hard.*

Observe that $\mathcal{D}, \pi, \ell \models G(d)$ whenever $\mathcal{D}, \ell \models Q_\alpha(d)$ for some $\alpha \in \text{Accept}(\pi, G)$. If $\mathcal{D}, \ell \models Q_\alpha(d)$, then $x_1, y_1, \dots, y_n, x_n$, the vertical and horizontal segments of α , appear in \mathcal{D} , starting from d at time ℓ . Expand y_i to y'_i , the \preceq -maximal horizontal segment fitting in \mathcal{D} , and let $\beta = x_1 y'_1 \dots x_n y'_n$. Then $\beta \in \text{Accept}(\pi, G)$ and $\mathcal{D}, \ell \models Q_\beta(d)$. Thus, to check $\mathcal{D}, \pi, \ell \models G(d)$, it suffices to find vertical segments of some $\alpha \in \text{Accept}(\pi, G)$, taking \preceq -maximal horizontal segments, and check that $\beta \in \text{Accept}(\pi, G)$. If (π, G) is vertically bounded, there are finitely many vertical segments of interest. Finding vertical segments, as well as extracting \preceq -maximal horizontal ones, can be done by an AC^0 circuit.

► **Lemma 15.** *If (π, G) is vertically bounded, then the data complexity of (π, G) coincides with that of checking membership in $\text{Accept}(\pi, G)$, modulo reductions computable in AC^0 .*

Recall that $\text{Accept}(\pi, G)$ is a regular language, so checking membership in it is either in AC^0 , or $\text{ACC}^0 \setminus \text{AC}^0$, or NC^1 -complete [40]. This settles the part (i) of Theorem 13: a query is either vertically unbounded and thus L-hard, or vertically bounded and thus belongs to one of the three classes mentioned above.

It remains to address the part (ii) of Theorem 13. By a careful analysis of the proof of Lemma 11 one can show that $\text{NotAccept}(\pi, G)$ is recognised by a nondeterministic automaton of size $2^{\text{poly}(|\pi|)}$. Further, checking whether its language belongs to AC^0 , $\text{ACC}^0 \setminus \text{AC}^0$, or is NC^1 -complete, can be done via the polynomial-space procedure developed in [31]. Thus, given a vertically bounded query, its data complexity can be established in exponential space.

For the vertical boundedness itself, note that substituting $Accept(\pi, G)$ with $Expand_{\leq}(\pi, G)$ in the respective definition preserves all the results proven so far. For $Expand_{\leq}(\pi, G)$, we can get from Lemma 10 an exponential-size one-way automaton. Checking that the query is vertically unbounded can be done in nondeterministic space logarithmic in the size of the automata for $Expand_{\leq}(\pi, G)$ and $NotAccept(\pi, G)$, as it boils down to checking reachability in their Cartesian product.

► **Lemma 16.** *Checking if a connected $datalog_{lm}^{\circ}$ query is vertically bounded is in PSPACE.*

For the lower bound, we show that checking boundedness is already PSPACE-hard for connected linear monadic queries in plain datalog. In fact, PSPACE-hardness was proved in [15] for *program boundedness* of disconnected programs. A program π is called bounded if (π, G) is bounded for every IDB G in π . We were able to regain the connectedness of π by focusing on query boundedness (also called *predicate boundedness*) instead. The idea combines that of [15] with that of Section 3: define an IDB F that slides along a computation, this time of a space-bounded Turing machine, looking for an erroneous transition.

► **Lemma 17.** *Deciding boundedness of connected linear monadic queries in plain datalog is PSPACE-hard.*

Lemmas 14 and 15 bring us to an interesting consideration: to understand the data complexity of a given query one should analyse its behaviour in the relational domain separately from that in the temporal domain. This can be given a precise sense as follows.

► **Proposition 18.** *For every connected $datalog_{lm}^{\circ}$ query (π, G) there exist a plain datalog query (π_d, G) and a plain LTL query (π_t, G) , such that:*

1. (π, G) is vertically bounded if and only if (π_d, G) is bounded;
2. if (π, G) is vertically bounded then its data complexity coincides with that of (π_t, G) .

Technically, both π_d and π_t are obtained by simulating the deterministic automaton $\mathcal{A}_{(\pi, G)}$ for the language $Accept(\pi, G)$ provided by Corollary 12. In both programs, the IDBs correspond to the states of $\mathcal{A}_{(\pi, G)}$ and EDBs to the letters of Ω . For π_t these EDBs are unary and all the rules are horizontal, so that the expansions unwind fully in the temporal domain. In the case of π_d , the EDBs are binary and every vertical transition of $\mathcal{A}_{(\pi, G)}$ is a step by a binary relation, while the horizontal transitions are skipped (thus, vertical boundedness becomes just boundedness). In both programs, initialisation rules correspond to $\mathcal{A}_{(\pi, G)}$ reaching an accepting state.

We conclude this section with a consideration on disconnected queries. In [15], deciding boundedness is based on the fact that $accept(\pi, G)$ remains regular even when π has disconnected rules. This is not the case for $datalog_{lm}^{\circ}$, as can be seen from the following example.

► **Example 19.** Consider the program π of four rules:

$$\begin{array}{ll} G(X) \leftarrow A(X) \wedge \odot D(X) & D(X) \leftarrow \odot D(X) \\ D(X) \leftarrow R(X, Y) \wedge \odot^- B(Y) \wedge \odot^- D(Y) & D(X) \leftarrow A(Y) \wedge B(X) \end{array}$$

Let $B_1 = A(X) \wedge \odot D(X)$, $B_2 = \odot D(X)$, and $B_3 = R(X, Y) \wedge \odot^- B(Y) \wedge \odot^- D(Y)$. Then $B_1 B_2^n B_3^m \in accept(\pi, G)$, and, consequently, $\{\perp, B_1\}\{B_2\}^n\{\top\}\{B_3\}^m \in Accept(\pi, G)$, whenever $m > n \geq 1$. This property is not recognisable by any finite state automaton. For more general classes of automata suitable to recognise $accept(\pi, G)$ or $Accept(\pi, G)$ in the disconnected setting, the properties that are to be checked along the lines of [15], such as language emptiness or finiteness, become undecidable. Therefore, disconnected queries possibly require a different approach for analysing the data complexity.

6 Conclusions and Future Work

We have started investigating the complexity of determining the data complexity of answering monadic datalog queries with temporal operators. For linear connected queries with operators \bigcirc/\bigcirc^- , we have generalised the automata-theoretic technique of [15], developed originally for plain datalog, to establish an $AC^0/ACC^0/NC^1/L/NL$ classification of temporal query answering and proved that deciding L-hardness of a given query is PSPACE-complete, while checking its membership in AC^0 or ACC^0 can be done in EXPSpace. As a minor side product, we have established PSPACE-hardness of deciding boundedness of atemporal connected linear monadic datalog queries. Rather surprisingly and in sharp contrast to the \bigcirc/\bigcirc^- case, it turns out that checking (non-trivial) membership of queries with operators \diamond/\diamond^- in the above complexity classes is undecidable. The results of this paper lead to a plethora of natural and intriguing open questions. Some of them are briefly discussed below.

1. What happens if we disallow applications of \diamond/\diamond^- to binary EDB predicates in $datalog_{lm}^\diamond$ -queries? We conjecture that this restriction makes checking membership in the above complexity classes decidable. In fact, this conjecture follows from a positive answer to the next question.
2. Can our decidability results for $datalog_{lm}^\bigcirc$ be lifted to $datalog_m^\bigcirc$ -queries? Dropping the linearity restriction in the atemporal case results in the extra data complexity class, P, and the higher complexity, 2EXPTIME-completeness, of deciding boundedness. The upper bound was obtained using tree automata in [15], and we believe that this approach can be generalised to connected $datalog_m^\bigcirc$ -queries in a way similar to what we have done above.
3. On the other hand, dropping the connectedness restriction might turn out to be trickier, if at all possible, as shown by Example 19. Finding a new automata-theoretic characterisation for disconnected $datalog_{lm}^\bigcirc$ -queries remains a challenging open problem.
4. A decisive step in understanding the data complexity of answering queries mediated by a description logic ontology and monadic disjunctive datalog queries was made in [9, 18] by establishing a close connection with constraint satisfaction problems (CSPs). In our case, quantified CSPs (see, e.g., [47]) seem to be more appropriate. Connecting the two areas might be beneficial to both of them.
5. In the context of streaming data, it would be interesting to investigate the data complexity classes and the complexity of recognising them for datalog MTL -queries [10, 36, 44].

References

- 1 Rajeev Alur and Thomas A. Henzinger. Real-time logics: Complexity and expressiveness. *Inf. Comput.*, 104(1):35–77, 1993. doi:10.1006/inco.1993.1025.
- 2 Peter Alvaro, William R. Marczak, Neil Conway, Joseph M. Hellerstein, David Maier, and Russell Sears. Dedalus: Datalog in time and space. In Oege de Moor, Georg Gottlob, Tim Furche, and Andrew Jon Sellers, editors, *Datalog Reloaded - First International Workshop, Datalog 2010, Oxford, UK, March 16-19, 2010. Revised Selected Papers*, volume 6702 of *Lecture Notes in Computer Science*, pages 262–281. Springer, 2010. doi:10.1007/978-3-642-24206-9_16.
- 3 Alessandro Artale, Diego Calvanese, Roman Kontchakov, and Michael Zakharyashev. The dl-lite family and relations. *J. Artif. Intell. Res.*, 36:1–69, 2009. doi:10.1613/jair.2820.
- 4 Alessandro Artale, Anton Gnatenco, Vladislav Ryzhikov, and Michael Zakharyashev. On deciding the data complexity of answering linear monadic datalog queries with LTL operators (extended version), 2025. URL: <https://arxiv.org/abs/2501.13762>.

- 5 Alessandro Artale, Roman Kontchakov, Alisa Kovtunova, Vladislav Ryzhikov, Frank Wolter, and Michael Zakharyashev. First-order rewritability of ontology-mediated queries in linear temporal logic. *Artif. Intell.*, 299:103536, 2021. doi:10.1016/j.artint.2021.103536.
- 6 Alessandro Artale, Roman Kontchakov, Vladislav Ryzhikov, and Michael Zakharyashev. The complexity of clausal fragments of LTL. In Kenneth L. McMillan, Aart Middeldorp, and Andrei Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning - 19th International Conference, LPAR-19, Stellenbosch, South Africa, December 14-19, 2013. Proceedings*, volume 8312 of *Lecture Notes in Computer Science*, pages 35–52. Springer, 2013. doi:10.1007/978-3-642-45221-5_3.
- 7 Alessandro Artale, Roman Kontchakov, Vladislav Ryzhikov, and Michael Zakharyashev. A cookbook for temporal conceptual data modelling with description logics. *ACM Trans. Comput. Log.*, 15(3):25:1–25:50, 2014. doi:10.1145/2629565.
- 8 Michael Benedikt, Balder ten Cate, Thomas Colcombet, and Michael Vanden Boom. The complexity of boundedness for guarded logics. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2015, Kyoto, Japan, July 6-10, 2015*, pages 293–304. IEEE Computer Society, 2015. doi:10.1109/LICS.2015.36.
- 9 Meghyn Bienvenu, Balder ten Cate, Carsten Lutz, and Frank Wolter. Ontology-based data access: A study through disjunctive datalog, csp, and MMSNP. *ACM Trans. Database Syst.*, 39(4):33:1–33:44, 2014. doi:10.1145/2661643.
- 10 Sebastian Brandt, Elem Güzel Kalayci, Vladislav Ryzhikov, Guohui Xiao, and Michael Zakharyashev. Querying log data with metric temporal logic. *J. Artif. Intell. Res.*, 62:829–877, 2018. doi:10.1613/jair.1.11229.
- 11 Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. Tractable reasoning and efficient query answering in description logics: The *DL-Lite* family. *J. Autom. Reason.*, 39(3):385–429, 2007. doi:10.1007/s10817-007-9078-x.
- 12 Jan Chomicki. Polynomial time query processing in temporal deductive databases. In Daniel J. Rosenkrantz and Yehoshua Sagiv, editors, *Proceedings of the Ninth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, April 2-4, 1990, Nashville, Tennessee, USA*, pages 379–391. ACM Press, 1990. doi:10.1145/298514.298589.
- 13 Jan Chomicki and Tomasz Imielinski. Temporal deductive databases and infinite objects. In Chris Edmondson-Yurkanan and Mihalis Yannakakis, editors, *Proceedings of the Seventh ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, March 21-23, 1988, Austin, Texas, USA*, pages 61–73. ACM, 1988. doi:10.1145/308386.308416.
- 14 Jan Chomicki, David Toman, and Michael H. Böhlen. Querying ATSQL databases with temporal logic. *ACM Trans. Database Syst.*, 26(2):145–178, 2001. doi:10.1145/383891.383892.
- 15 Stavros S. Cosmadakis, Haim Gaifman, Paris C. Kanellakis, and Moshe Y. Vardi. Decidable optimization problems for database logic programs (preliminary report). In Janos Simon, editor, *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, May 2-4, 1988, Chicago, Illinois, USA*, pages 477–490. ACM, 1988. doi:10.1145/62212.62259.
- 16 Evgeny Dantsin, Thomas Eiter, Georg Gottlob, and Andrei Voronkov. Complexity and expressive power of logic programming. *ACM Comput. Surv.*, 33(3):374–425, 2001. doi:10.1145/502807.502810.
- 17 Stéphane Demri, Valentin Goranko, and Martin Lange. *Temporal Logics in Computer Science: Finite-State Systems*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2016. doi:10.1017/CBO9781139236119.
- 18 Cristina Feier, Antti Kuusisto, and Carsten Lutz. Rewritability in monadic disjunctive datalog, mmsnp, and expressive description logics. *Log. Methods Comput. Sci.*, 15(2), 2019. doi:10.23638/LMCS-15(2:15)2019.
- 19 Valeria Fionda and Giuseppe Pirrò. Characterizing evolutionary trends in temporal knowledge graphs with linear temporal logic. In Jingrui He, Themis Palpanas, Xiaohua Hu, Alfredo Cuzzocrea, Dejing Dou, Dominik Slezak, Wei Wang, Aleksandra Gruca, Jerry Chun-Wei

- Lin, and Rakesh Agrawal, editors, *IEEE International Conference on Big Data, BigData 2023, Sorrento, Italy, December 15-18, 2023*, pages 2907–2909. IEEE, 2023. doi:10.1109/BigData59044.2023.10386573.
- 20 International Organization for Standardization. Information technology — database languages — sql — part 2: Foundation (sql/foundation), 1999. URL: <https://www.iso.org/standard/23532.html>.
 - 21 Nadime Francis, Alastair Green, Paolo Guagliardo, Leonid Libkin, Tobias Lindaaker, Victor Marsault, Stefan Plantikow, Mats Rydberg, Petra Selmer, and Andrés Taylor. Cypher: An evolving query language for property graphs. In Gautam Das, Christopher M. Jermaine, and Philip A. Bernstein, editors, *Proceedings of the 2018 International Conference on Management of Data, SIGMOD Conference 2018, Houston, TX, USA, June 10-15, 2018*, pages 1433–1445. ACM, 2018. doi:10.1145/3183713.3190657.
 - 22 Víctor Gutiérrez-Basulto, Jean Christoph Jung, and Roman Kontchakov. Temporalized EL ontologies for accessing temporal data: Complexity of atomic queries. In Subbarao Kambhampati, editor, *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pages 1102–1108. IJCAI/AAAI Press, 2016. URL: <http://www.ijcai.org/Abstract/16/160>.
 - 23 Steve Harris and Andy Seaborne. Sparql 1.1 query language. <https://www.w3.org/TR/sparql11-query/>, 2013. W3C Recommendation, 21 March 2013. URL: <https://www.w3.org/TR/sparql11-query/>.
 - 24 Gerd G. Hillebrand, Paris C. Kanellakis, Harry G. Mairson, and Moshe Y. Vardi. Undecidable boundedness problems for datalog programs. *J. Log. Program.*, 25(2):163–190, 1995. doi:10.1016/0743-1066(95)00051-K.
 - 25 Ismail Husein, Herman Mawengkang, Saib Suwilo, and Mardiningsih. Modeling the transmission of infectious disease in a dynamic network. *Journal of Physics: Conference Series*, 1255(1):012052, August 2019. URL: <https://dx.doi.org/10.1088/1742-6596/1255/1/012052>.
 - 26 Neil Immerman. *Descriptive complexity*. Graduate texts in computer science. Springer, 1999. doi:10.1007/978-1-4612-0539-5.
 - 27 Paris C. Kanellakis. Elements of relational database theory. In Jan van Leeuwen, editor, *Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics*, pages 1073–1156. Elsevier and MIT Press, 1990. doi:10.1016/b978-0-444-88074-1.50022-6.
 - 28 Kevin Cullinane. Modeling dynamic transportation networks: Bin ran and david boyce springer 1996 isbn 3540611398. *Journal of Transport Geography*, 6(1):76–78, 1998. doi:10.1016/S0966-6923(98)90041-2.
 - 29 Stanislav Kikot, Agi Kurucz, Vladimir V. Podolskii, and Michael Zakharyashev. Deciding boundedness of monadic sirups. In Leonid Libkin, Reinhard Pichler, and Paolo Guagliardo, editors, *PODS’21: Proceedings of the 40th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, Virtual Event, China, June 20-25, 2021*, pages 370–387. ACM, 2021. doi:10.1145/3452021.3458332.
 - 30 Ron Koymans. Specifying real-time properties with metric temporal logic. *Real Time Syst.*, 2(4):255–299, 1990. doi:10.1007/BF01995674.
 - 31 Agi Kurucz, Vladislav Ryzhikov, Yury Savateev, and Michael Zakharyashev. Deciding forewritability of regular languages and ontology-mediated queries in linear temporal logic. *J. Artif. Intell. Res.*, 76:645–703, 2023. doi:10.1613/jair.1.14061.
 - 32 Ling Liu and M. Tamer Özsu, editors. *Encyclopedia of Database Systems, Second Edition*. Springer, 2018. doi:10.1007/978-1-4614-8265-9.
 - 33 Marvin L. Minsky. *Computation: finite and infinite machines*. Prentice-Hall, Inc., USA, 1967. URL: <https://dl.acm.org/doi/book/10.5555/1095587>.
 - 34 Jeffrey F. Naughton. Data independent recursion in deductive databases. *J. Comput. Syst. Sci.*, 38(2):259–289, 1989. doi:10.1016/0022-0000(89)90003-2.

- 35 Juan L. Reutter, Adrián Soto, and Domagoj Vrgoc. Recursion in SPARQL. *Semantic Web*, 12(5):711–740, 2021. doi:10.3233/SW-200401.
- 36 Vladislav Ryzhikov, Przemyslaw Andrzej Walega, and Michael Zakharyashev. Data complexity and rewritability of ontology-mediated queries in metric temporal logic under the event-based semantics. In Sarit Kraus, editor, *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 1851–1857. ijcai.org, 2019. doi:10.24963/ijcai.2019/256.
- 37 Benjamin Schäfer, Dirk Witthaut, Marc Timme, and Vito Latora. Dynamically induced cascading failures in power grids. *Nature Communications*, 9(1):1975, May 2018. doi:10.1038/s41467-018-04287-5.
- 38 Brian Skyrms and Robin Pemantle. A dynamic model of social network formation. *Proceedings of the National Academy of Sciences*, 97(16):9340–9346, 2000. arXiv:https://www.pnas.org/doi/pdf/10.1073/pnas.97.16.9340.
- 39 Richard T. Snodgrass, Ilsoo Ahn, Gad Ariav, Don S. Batory, James Clifford, Curtis E. Dyreson, Ramez Elmasri, Fabio Grandi, Christian S. Jensen, Wolfgang Käfer, Nick Kline, Krishna G. Kulkarni, T. Y. Cliff Leung, Nikos A. Lorentzos, John F. Roddick, Arie Segev, Michael D. Soo, and Suryanarayana M. Sripada. TSQL2 language specification. *SIGMOD Rec.*, 23(1):65–86, 1994. doi:10.1145/181550.181562.
- 40 Howard Straubing. *Finite Automata, Formal Logic, and Circuit Complexity*. Birkhäuser, Boston, MA, 1994. URL: <http://link.springer.com/10.1007/978-1-4612-0289-9>.
- 41 David Toman. Point vs. interval-based query languages for temporal databases. In Richard Hull, editor, *Proceedings of the Fifteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 3-5, 1996, Montreal, Canada*, pages 58–67. ACM Press, 1996. doi:10.1145/237661.237676.
- 42 Valentina Urzua and Claudio Gutierrez. Linear recursion in G-CORE. In Aidan Hogan and Tova Milo, editors, *Proceedings of the 13th Alberto Mendelzon International Workshop on Foundations of Data Management, Asunción, Paraguay, June 3-7, 2019*, volume 2369 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2019. URL: <https://ceur-ws.org/Vol-2369/short07.pdf>.
- 43 Ron van der Meyden. Predicate boundedness of linear monadic datalog is in PSPACE. *Int. J. Found. Comput. Sci.*, 11(4):591–612, 2000. doi:10.1142/S0129054100000351.
- 44 Dingmin Wang, Pan Hu, Przemyslaw Andrzej Walega, and Bernardo Cuenca Grau. Meteor: Practical reasoning in datalog with metric temporal operators. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 5906–5913. AAAI Press, 2022. doi:10.1609/aaai.v36i5.20535.
- 45 Mincheng Wu, Chao Li, Zhangchong Shen, Shibo He, Lingling Tang, Jie Zheng, Yi Fang, Kehan Li, Yanggang Cheng, Zhiguo Shi, Guoping Sheng, Yu Liu, Jinxing Zhu, Xinjiang Ye, Jinlai Chen, Wenrong Chen, Lanjuan Li, Youxian Sun, and Jiming Chen. Use of temporal contact graphs to understand the evolution of covid-19 through contact tracing data. *Communications Physics*, 5(1):270, 2022. doi:10.1038/s42005-022-01045-4.
- 46 Mengkai Xu, Srinivasan Radhakrishnan, Sagar Kamarthi, and Xiaoning Jin. Resiliency of mutualistic supplier-manufacturer networks. *Scientific Reports*, 9, September 2019. doi:10.1038/s41598-019-49932-1.
- 47 Dmitriy Zhuk. \prod_2^P vs pspace dichotomy for the quantified constraint satisfaction problem. In *65th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2024, Chicago, IL, USA, October 27-30, 2024*, pages 560–572. IEEE, 2024. doi:10.1109/FOCS61266.2024.00043.