

Counting Queries over \mathcal{ELHI}_\perp Ontologies ^{*} ^{**}

Meghyn Bienvenu¹[0000-0001-6229-8103], Quentin Manière¹[0000-0001-9618-8359],
and Michaël Thomazo²[0000-0002-1437-6389]

¹ CNRS, University of Bordeaux, Bordeaux INP, LaBRI, Talence, France
² Inria, DI ENS, ENS, CNRS, University PSL, Paris, France

Abstract. While ontology-mediated query answering most often adopts (unions of) conjunctive queries as the query language, some recent works have explored the use of counting queries coupled with DL-Lite ontologies. In the present paper, we initiate the study of counting queries for Horn description logics outside the DL-Lite family. Via a non-trivial adaptation of existing techniques, we devise a decision procedure for answering counting conjunctive queries over \mathcal{ELHI}_\perp ontologies, which yields the same upper bounds as are known for DL-Lite \mathcal{R} , namely, coNP in data complexity and coN2EXP w.r.t. combined complexity. We further show that the best known lower bounds for DL-Lite \mathcal{R} (coNP -hard for data, coNEXP -hard for combined) hold also for \mathcal{EL} .

Keywords: Ontology-mediated query answering · Counting queries

1 Introduction

In the context of ontology-mediated query answering, the most commonly considered queries are conjunctive queries (CQs), but several works have explored ways of equipping CQs with some form of counting [8,10,9]. A recent approach, proposed in [3] as a generalization of [10], considers *counting conjunctive queries (CCQs)* that are syntactically defined like standard CQs except that some variables may be designated as *counting variables*. In each model of the knowledge base, we can count the number of possible assignments to the counting variables that make the query hold. Certain answers are then defined as intervals that provide upper and lower bounds on the count values across all models.

The problem of answering CCQs over DL-Lite \mathcal{R} ontologies is intractable in general [10], and recent works have shown that intractability arises even in quite restricted settings [6,4]. However, some interesting tractable cases have also been identified, notably, rooted CCQs [3,6,11] and cardinality queries (= Boolean atomic CCQs) [4] coupled with DL-Lite $_{\text{core}}$ ontologies; query rewriting techniques have also begun to be explored [7]. To the best of our knowledge, however, CCQ answering has not yet been studied for ontologies outside the DL-Lite family.

^{*} Partially supported by ANR project CQFD (ANR-18-CE23-0003)

^{**} © 2021 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

This motivates us to extend the study of CCQ answering to other well-known Horn description logics, such as \mathcal{EL} and the more expressive \mathcal{ELHI}_\perp . While \mathcal{EL} enjoys good computational properties for CQ answering [1], the techniques developed for CCQs in the DL-Lite context do not easily transfer. Indeed, a key property of DL-Lite is that merging elements in a model preserves modelhood so long as disjointness axioms are not violated. This property does not hold in \mathcal{EL} , due to conjunctions of concepts on the left-hand side (LHS) of inclusions.

Through a non-trivial adaptation of the DL-Lite approach, we obtain similar upper bounds for CCQ answering over \mathcal{ELHI}_\perp ontologies as were known for DL-Lite $_{\mathcal{R}}$, namely, coNP membership w.r.t. data complexity and coN2EXP membership w.r.t. combined complexity. We further prove that even if we restrict to \mathcal{EL} , we have the same lower bounds as in DL-Lite $_{\mathcal{R}}$: coNP-hardness w.r.t. data complexity and coNEXP-hardness w.r.t. combined complexity.

2 Preliminaries

Knowledge Bases. We assume mutually disjoint sets N_C , N_R , and N_I of *concept*, *role*, and *individual names*. A *knowledge base (KB)* $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ consists of an ABox \mathcal{A} and a TBox \mathcal{T} . An *ABox* is a finite set of *concept assertions* $A(b)$ (with $A \in N_C$, $b \in N_I$) and *role assertions* $P(a, b)$ (with $P \in N_R$, $a, b \in N_I$). Let $\text{Ind}(\mathcal{A})$ denote the set of individuals occurring in an ABox \mathcal{A} .

A *TBox* is a finite set of axioms. In our considered DL \mathcal{ELHI}_\perp , TBoxes consist of *concept inclusions* $B_1 \sqsubseteq B_2$, *positive role inclusions* $R_1 \sqsubseteq R_2$, and *negative role inclusions*³ $R_1 \sqcap R_2 \sqsubseteq \perp$, where the R_i are *roles* drawn from $N_R^\pm = \{P, P^- \mid P \in N_R\}$ and the B_i are (*complex*) *concepts* constructed as follows:

$$B := \perp \mid \top \mid A \mid B_1 \sqcap B_2 \mid \exists R.B \quad \text{where } A \in N_C, R \in N_R^\pm$$

Let $\text{sig}(\mathcal{T})$ denote the set of concept and role names appearing in TBox \mathcal{T} .

Semantics of KBs. An interpretation takes the form $\mathcal{I} = (\Delta^\mathcal{I}, \cdot^\mathcal{I})$, where $\Delta^\mathcal{I}$ is a non-empty set (called the domain) and $\cdot^\mathcal{I}$ is the interpretation function that maps each $A \in N_C$ to $A^\mathcal{I} \subseteq \Delta^\mathcal{I}$, each $P \in N_R$ to $P^\mathcal{I} \subseteq \Delta^\mathcal{I} \times \Delta^\mathcal{I}$, and each $a \in N_I$ to $a^\mathcal{I}$. In this paper, we will make the *Standard Names Assumption* by setting $a^\mathcal{I} = a$. Note however that our results only rely upon the weaker *Unique Names Assumption* (UNA), which stipulates that $a^\mathcal{I} \neq b^\mathcal{I}$ whenever $a \neq b$.

The function $\cdot^\mathcal{I}$ naturally extends to roles and complex concepts: $(P^-)^\mathcal{I} = \{(y, x) \mid (x, y) \in P^\mathcal{I}\}$, $\perp^\mathcal{I} = \emptyset$, $\top^\mathcal{I} = \Delta^\mathcal{I}$, $(B_1 \sqcap B_2)^\mathcal{I} = B_1^\mathcal{I} \cap B_2^\mathcal{I}$ and $(\exists P.B)^\mathcal{I} = \{d \mid (d, e) \in P^\mathcal{I}, e \in B^\mathcal{I}\}$. An inclusion $G \sqsubseteq H$ is satisfied in \mathcal{I} if $G^\mathcal{I} \subseteq H^\mathcal{I}$; an assertion $A(b)$ (resp. $P(a, b)$) is satisfied in \mathcal{I} if $b \in A^\mathcal{I}$ (resp. $(a, b) \in P^\mathcal{I}$). An interpretation is a *model* of a TBox \mathcal{T} (resp. KB \mathcal{K}) if it satisfies all axioms in \mathcal{T} (resp. axioms and assertions in \mathcal{K}). A KB is *satisfiable* if it has at least one model. An inclusion (resp. assertion) Φ is *entailed* from \mathcal{T} (resp. \mathcal{K}), written $\mathcal{T} \models \Phi$ (resp. $\mathcal{K} \models \Phi$), if Φ is satisfied in every model of \mathcal{T} (resp. \mathcal{K}).

³ We follow e.g. [2] by considering a version of \mathcal{ELHI}_\perp with negative role inclusions.

Example 1. Consider the ABox $\mathcal{A}_e = \{A(a), B(b)\}$ and the \mathcal{ELHI}_\perp TBox \mathcal{T}_e :

$$\begin{array}{llllll} A \sqsubseteq \exists P.A' & B \sqsubseteq \exists Q.B' & A' \sqcap B' \sqsubseteq A_0 & A' \sqsubseteq D & B' \sqsubseteq D \\ A_0 \sqsubseteq \exists R_1.A_1 & A_1 \sqsubseteq \exists R_2.A_2 & A_2 \sqsubseteq \exists R_3.A_3 & A_3 \sqsubseteq \exists S.B_0 & B_0 \sqsubseteq \exists V.B'_0 \\ B_0 \sqsubseteq \exists U.C_0 & U \sqsubseteq V & C_0 \sqsubseteq \exists V_1.C_1 & C_1 \sqsubseteq \exists V_2.C_2 & C_2 \sqsubseteq \exists V_3.D \end{array}$$

Our example KB is $\mathcal{K}_e := (\mathcal{T}_e, \mathcal{A}_e)$. Figures 1a, 1c, 1d and 1e depict models of \mathcal{K}_e ; its canonical model $\mathcal{C}_{\mathcal{K}_e}$ (formally defined later) is displayed in Figure 1a.

Counting Queries. We consider counting queries as defined in [3] (which generalizes the queries considered in [10,6]). A *counting conjunctive query (CCQ)* takes the form $q(\mathbf{x}) = \exists \mathbf{y} \exists \mathbf{z} \psi(\mathbf{x}, \mathbf{y}, \mathbf{z})$, where $\mathbf{x}, \mathbf{y}, \mathbf{z}$ are tuples of *answer, existential, and counting variables*, respectively, and ψ is a conjunction of concept and role atoms with terms from $\mathbb{N}_1 \cup \mathbf{x} \cup \mathbf{y} \cup \mathbf{z}$. A CCQ q is *Boolean* if $\mathbf{x} = \emptyset$.

A *match* for a CCQ q in an interpretation \mathcal{I} is a homomorphism from q into \mathcal{I} . If a match π maps \mathbf{x} to \mathbf{a} , then the restriction of π to \mathbf{z} is called a *counting match (c-match)* of $q(\mathbf{a})$ in \mathcal{I} . The set of *answers* to q in \mathcal{I} , denoted $q^\mathcal{I}$, contains all pairs $(\mathbf{a}, [m, M])$, with $m, M \in \mathbb{N} \cup \{+\infty\}$, such that the number of distinct c-matches of $q(\mathbf{a})$ in \mathcal{I} belongs to the interval $[m, M]$. A *certain answer* to q w.r.t. \mathcal{K} is an answer in every model of \mathcal{K} , that is a pair from $\bigcap_{\mathcal{I} \models \mathcal{K}} q^\mathcal{I}$.

As usual, it is sufficient to consider the Boolean case: $(\mathbf{a}, [m, M])$ is a certain answer to a CCQ $q(\mathbf{x})$ iff $(\emptyset, [m, M])$ is a certain answer to the Boolean CCQ $q(\mathbf{a})$ obtained by replacing \mathbf{x} with \mathbf{a} . Thus, from now on, we *focus on Boolean CCQs*, and work with candidate answers $[m, M]$ in place of $(\emptyset, [m, M])$.

We further observe that since \mathcal{ELHI}_\perp cannot restrict the size of models, the value M in a certain answer $[m, M]$ is: 0 if the underlying CQ is unsatisfiable w.r.t. \mathcal{T} , any number greater than 1 if q has a match in every model but $\mathbf{z} = \emptyset$; and $+\infty$ otherwise. As the first two cases can be readily handled using existing techniques, we *focus on identifying certain answers of the form $[m, +\infty]$* .

Example 2. Let $q_e := \exists z D(z)$ be a Boolean CCQ. Intervals $[0, +\infty]$ and $[1, +\infty]$ are certain answers to q_e over \mathcal{K}_e . Interval $[3, +\infty]$ is not a certain answer as the models depicted on Figures 1a, 1c, 1d and 1e contain only 2 matches for q_e .

Complexity. Given a \mathcal{ELHI}_\perp knowledge base $\mathcal{K} = (\mathcal{T}, \mathcal{A})$, a Boolean CCQ q , and an integer $m \geq 0$ (in binary), we are interested in the complexity of deciding whether $[m, +\infty]$ is a certain answer to q w.r.t. \mathcal{K} . We will consider the two usual complexity measures: combined complexity which is in terms of the size of the whole input, and data complexity which is only in terms of the size of \mathcal{A} and m (\mathcal{T} and q are treated as fixed). If O is a TBox, ABox, KB, or CCQ, then the size of O , denoted $|O|$, is the number of occurrences of concept and role names in O .

Normal form. As is standard (see e.g. [2]), we work with \mathcal{ELHI}_\perp TBoxes in a convenient *normal form*, where every concept inclusion has one of the following restricted shapes:

$$A \sqsubseteq \perp \quad \top \sqsubseteq A \quad A_1 \sqcap A_2 \sqsubseteq A \quad A_1 \sqsubseteq \exists R.A_2 \quad \exists R.A_1 \sqsubseteq A_2$$

with $A, A_1, A_2 \in \mathbf{N}_C, R \in \mathbf{N}_R^\pm$. Through the introduction of fresh concept names, we can transform in polynomial time any TBox \mathcal{T} into a normal-form TBox \mathcal{T}' that is a model-conservative extension of \mathcal{T} (hence, indistinguishable from \mathcal{T} from the point of view of queries). We therefore *assume w.l.o.g. that all considered TBoxes are in normal form*.

Canonical model. It is well known that every satisfiable \mathcal{ELHI}_\perp KB admits a canonical (or universal) model that embeds homomorphically into each of its models. We recall how such a model $\mathcal{C}_\mathcal{K}$ can be constructed (see e.g. [5]). The domain $\Delta^{\mathcal{C}_\mathcal{K}}$ consists of all sequences $\mathbf{a}R_1.M_1 \dots R_n.M_n$ ($n \geq 0$) such that $\mathbf{a} \in \text{Ind}(\mathcal{A})$, each R_i belongs to \mathbf{N}_R^\pm , each M_i is a conjunction of concepts from $\mathbf{N}_C \cup \{\top\}$ (treated as a set when convenient), and the following conditions hold:

- If $n \geq 1$, then $\mathcal{T} \models M_0 \sqsubseteq \exists R_1.M_1$ where $M_0 = \{A \in \mathbf{N}_C \cup \{\top\} \mid \mathcal{K} \models A(\mathbf{a})\}$ and M_1 is maximal, as a set of concept names, for this property.
- For every $1 \leq i < n$, $\mathcal{T} \models M_i \sqsubseteq \exists R_{i+1}.M_{i+1}$ and M_{i+1} is maximal, as a set of concept names, for this property.

Individual names are interpreted as themselves ($\mathbf{a}^{\mathcal{C}_\mathcal{K}} = \mathbf{a}$), and concept and role names are interpreted as follows:

$$\begin{aligned} A^{\mathcal{C}_\mathcal{K}} &= \{\mathbf{a} \mid \mathcal{K} \models A(\mathbf{a})\} \cup \{e \cdot R.M \mid A \in M\} \\ P^{\mathcal{C}_\mathcal{K}} &= \{(\mathbf{a}, \mathbf{b}) \mid \mathcal{K} \models P(\mathbf{a}, \mathbf{b})\} \\ &\quad \cup \{(e, e \cdot P_0.M) \mid \mathcal{T} \models P_0 \sqsubseteq P\} \cup \{(e \cdot P_0.M, e) \mid \mathcal{T} \models P_0 \sqsubseteq P^-\} \end{aligned}$$

3 Upper Bounds via Countermodels

This section presents our main contribution: a decision procedure (and associated complexity upper bounds) for deciding whether a candidate interval $[m, +\infty]$ is a certain answer to a Boolean CCQ q and \mathcal{ELHI}_\perp KB $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ in normal form. It is more convenient in fact to focus on the complementary problem of checking whether $[m, +\infty]$ is *not* a certain answer, as the latter holds iff there exists a *countermodel*, i.e., a model of \mathcal{K} with fewer than m c-matches.

We start by observing that if m is large enough, a countermodel always exists.

Lemma 1. *There exists a countermodel for all $m \geq (|\text{Ind}(\mathcal{A})| + 3 |\mathcal{T}| 2^{|\mathcal{T}|})^{|q|}$.*

Proof (sketch). We exhibit a model of size at most $|\text{Ind}(\mathcal{A})| + 3 |\mathcal{T}| 2^{|\mathcal{T}|}$.

We may thus assume that the input m is such that $m \leq (|\text{Ind}(\mathcal{A})| + 3 |\mathcal{T}| 2^{|\mathcal{T}|})^{|q|}$.

The main ingredient underlying our decision procedure and upper bounds is the following result, which restricts the size of countermodels we need to consider.

Theorem 1. *If there exists a countermodel for input $[m, +\infty]$, CCQ q and KB \mathcal{K} , then there exists a countermodel with a polynomial-size (resp. double-exponential size) domain w.r.t. data complexity (resp. combined complexity).*

Theorem 1 generalizes analogous results for DL-Lite KBs in [10,3] and gives rise to a decision procedure that guesses an interpretation of bounded size and checks whether it is a countermodel, yielding the following upper bounds:

Theorem 2. *Deciding if input $[m, +\infty]$ is a certain answer for q over \mathcal{K} is in coNP (resp. in coN2EXP) w.r.t. data complexity (resp. combined complexity).*

The remainder of the section is devoted to proving Theorem 1. The high-level idea is to start from an arbitrary countermodel \mathcal{I} and merge its elements so as to reduce its size, while at the same time *not introducing any new query matches*. But how can we decide which elements of \mathcal{I} can be safely merged? Looking to existing DL-Lite approaches [10,3] for inspiration, we observe that they proceed in two steps. First, they define an intermediate model \mathcal{I}' (called *interleaving*) that, informally, retains the useful parts of \mathcal{I} (i.e., those involved in query matches or needed to satisfy the ABox) and replaces the rest with tree-shaped structures taken from the corresponding parts of the canonical model. With this more structured countermodel \mathcal{I}' , it is easier to identify, via a well-chosen equivalence relation, the elements that behave similarly and thus can be safely merged. In a second step, elements of \mathcal{I}' from the same equivalence class are merged to obtain the desired bounded-size countermodel.

A naive adaptation of the DL-Lite approach to \mathcal{ELHI}_\perp (or even \mathcal{EL}) fails already at the first step. Indeed, as the next example illustrates, due to conjunction in the LHS of concept inclusions, the interleaving need not be a model.

Example 3. *A countermodel \mathcal{I}_e of \mathcal{K}_e for q_e and candidate integer 3 is depicted in Figure 1c. Indeed \mathcal{I}_e yields only 2 possible values for z : α and β .*

The naive adaptation of the DL-Lite approach builds the interpretation depicted in Figure 1b and induced from $\mathcal{C}_{\mathcal{K}_e}$. Observe that α violates the axiom $A' \sqcap B' \sqsubseteq A_0$. Generally speaking, the issue is that the canonical model may not contain elements witnessing conjunctions of concepts that occur in the initial countermodel, so it is not enough to copy over parts of the canonical model.

We now show how to revamp the DL-Lite approach to make it work for \mathcal{ELHI}_\perp . To aid comprehension, we give a graphical overview in Figure 2.

3.1 Interlacing

We propose a new intermediate countermodel, called *interlacing*, which retains the desirable features of the interleaving but avoids the issues highlighted in Example 3. Essentially, the idea is to replace the canonical model by an alternative tree-shaped domain (called *existential extraction*) that is built from the countermodel \mathcal{I} by keeping track of the RHS existential concepts satisfied in \mathcal{I} .

The definition of existential extraction uses the alphabet Ω consisting of all R.A such that $\exists R.A$ is the RHS of an axiom in \mathcal{T} . Furthermore, it assumes that, for every R.A $\in \Omega$, we have chosen a function $\text{succ}_{R.A}^{\mathcal{I}}$ that maps every element $e \in (\exists R.A)^{\mathcal{I}}$ to an element $e' \in \Delta^{\mathcal{I}}$ such that $(e, e') \in R^{\mathcal{I}}$ and $e' \in A^{\mathcal{I}}$.

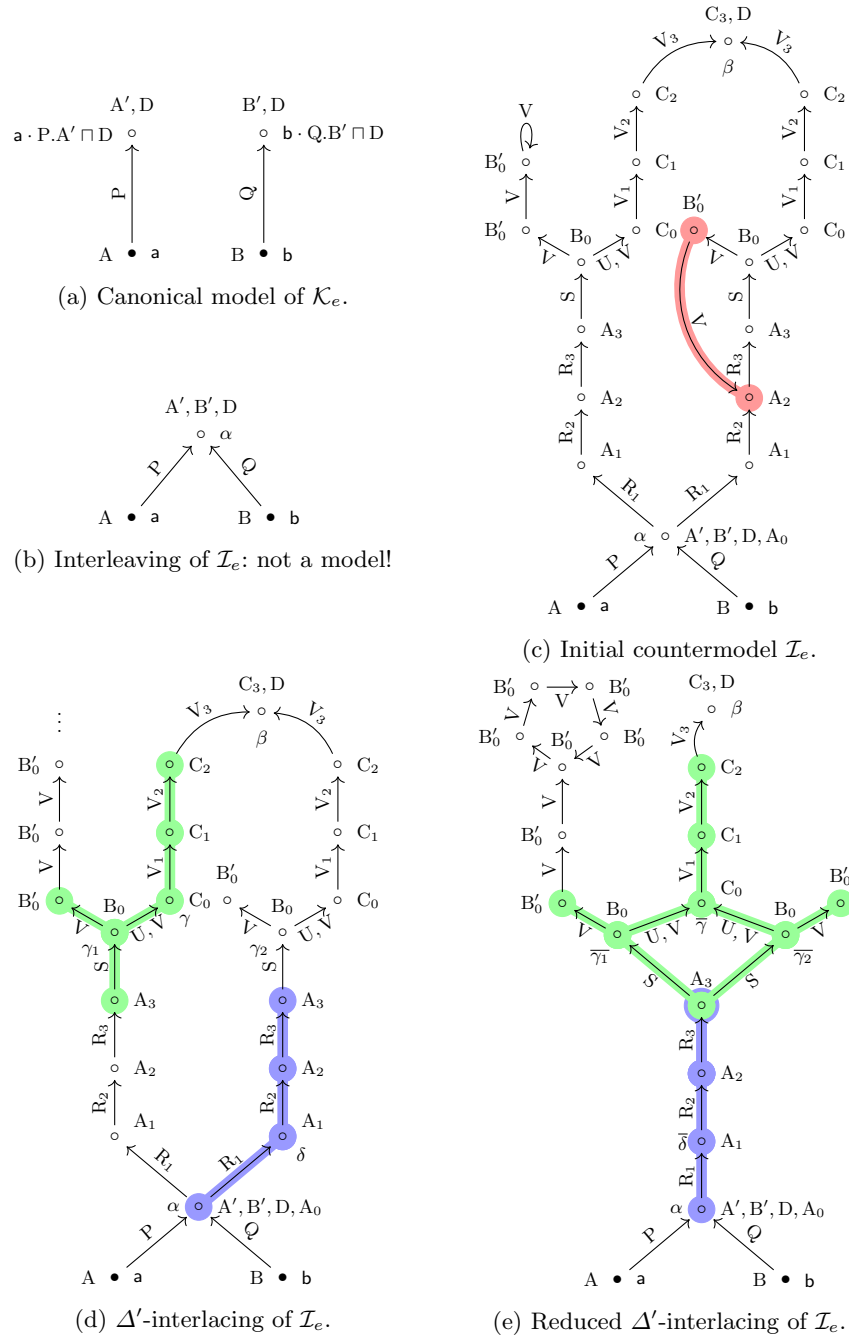


Fig. 1: Interpretations used along our examples.




Fig. 2: Models, domains, and mappings employed in our construction.

Definition 1. Over the set $\text{Ind}(\mathcal{A}) \cdot \Omega^*$, inductively build the following mapping:

$$f : \text{Ind}(\mathcal{A}) \cdot \Omega^* \rightarrow \Delta^{\mathcal{I}} \cup \{\uparrow\}$$

$$a \mapsto a$$

$$w \cdot \text{R.A} \mapsto \begin{cases} \uparrow & \text{if } f(w) = \uparrow \text{ or } f(w) \notin (\exists \text{R.A})^{\mathcal{I}} \\ \text{succ}_{\text{R.A}}^{\mathcal{I}}(f(w)) & \text{otherwise} \end{cases}$$

The existential extraction⁴ of \mathcal{I} is $\Delta^\circ := \{w \mid w \in \text{Ind}(\mathcal{A}) \cdot \Omega^*, f(w) \neq \uparrow\}$.

Remark 1. Δ° can be seen as the domain of a form of unravelling of \mathcal{I} starting from $\text{Ind}(\mathcal{A})$, in which we only follow the selected successors for the RHS existential concepts. The selective nature of the unravelling is visible in Figure 1d which is based on the existential extraction of \mathcal{I}_e : observe there is no edge that corresponds to following the red edge from \mathcal{I}_e (Figure 1c) since $\forall \text{A}_2 \notin \Omega$.

We proceed to define Δ -interlacings, parametrized by a set of interest Δ with $\text{Ind}(\mathcal{A}) \subseteq \Delta \subseteq \Delta^\circ$. Intuitively, we preserve the portion of \mathcal{I} corresponding to $f(\Delta)$ and complete it with (locally) tree-shaped structures issuing from Δ° .

Definition 2. The Δ -interlacing mapping of \mathcal{I} is:

$$f' : \Delta^\circ \rightarrow f(\Delta) \uplus (\Delta^\circ \setminus \Delta) \quad w \mapsto \begin{cases} f(w) & \text{if } w \in \Delta \\ w & \text{otherwise} \end{cases}$$

The Δ -interlacing \mathcal{I}' of \mathcal{I} is the interpretation given by $\Delta^{\mathcal{I}'} := f'(\Delta^\circ)$ and:

$$\begin{aligned} \text{A}^{\mathcal{I}'} &= \{f'(u) \mid u \in \Delta^\circ, f(u) \in \text{A}^{\mathcal{I}}\} (= f'(f^{-1}(\text{A}^{\mathcal{I}}))) \\ \text{P}^{\mathcal{I}'} &= \{(a, b) \mid a, b \in \text{Ind}(\mathcal{A}) \wedge \mathcal{K} \models \text{P}(a, b)\} && (\text{Role}'_{\mathcal{A}}) \\ &\cup \{(f'(u), f'(u \cdot \text{P}_0 \cdot \text{B})) \mid u, u \cdot \text{P}_0 \cdot \text{B} \in \Delta^\circ \wedge \mathcal{T} \models \text{P}_0 \sqsubseteq \text{P}\} && (\text{Role}'_{+}) \\ &\cup \{(f'(v \cdot \text{P}_0 \cdot \text{B}), f'(v)) \mid v, v \cdot \text{P}_0 \cdot \text{B} \in \Delta^\circ \wedge \mathcal{T} \models \text{P}_0^- \sqsubseteq \text{P}\} && (\text{Role}'_{-}) \end{aligned}$$

Like the interleaving, the Δ -interlacing is a model which embeds in \mathcal{I} .

Lemma 2. \mathcal{I}' is a model, and it embeds homomorphically in \mathcal{I} via the mapping:

$$\sigma : \Delta^{\mathcal{I}'} \rightarrow \Delta^{\mathcal{I}} \quad w \mapsto \begin{cases} w & \text{if } w \in f(\Delta) \\ f(w) & \text{otherwise, that is } w \in \Delta^\circ \setminus \Delta \end{cases}$$

⁴ While the definitions of f , Δ° , and later constructions depend on the choice of successor functions, all choices lead to the desired result.

Proof (sketch). Both statements rely upon the facts that (i) if $f'(u) \in \mathbf{A}^{\mathcal{I}'}$, then $f(u) \in \mathbf{A}^{\mathcal{I}}$, and (ii) if $(f'(u), f'(v)) \in \mathbf{R}^{\mathcal{I}'}$, then $(f(u), f(v)) \in \mathbf{R}^{\mathcal{I}}$.

Remark 2. *It is easily verified that $\sigma \circ f' = f$.*

To obtain a countermodel, we identify the crucial elements of \mathcal{I} , namely, $\Delta^* := \text{Ind}(\mathcal{A}) \cup \{e \in \Delta^{\mathcal{I}} \mid \exists \pi \in m(q, \mathcal{I}), \exists z \in \mathbf{z}, \pi(z) = e\}$, where $m(q, \mathcal{I})$ is the set of matches of q in \mathcal{I} . We then take $\Delta' := f^{-1}(\Delta^*)$ as our set of interest.

Lemma 3. *If \mathcal{I} is a countermodel, then so is its Δ' -interlacing \mathcal{I}' .*

Proof (sketch). We show that σ injectively maps matches in \mathcal{I}' to matches in \mathcal{I} .

Example 4. *Figure 1d depicts the Δ' -interlacing of \mathcal{I}_e . Like the initial model \mathcal{I}_e , it is a countermodel for q_e and integer 3. Notice \mathcal{I}'_e has an infinite domain.*

3.2 Reduced interlacing

It remains to merge elements of the Δ' -interlacing \mathcal{I}' to obtain a countermodel of the required size. To identify similar elements, we consider their neighbourhoods.

Definition 3. *Consider an interpretation \mathcal{M} and an element $c \in \Delta^{\mathcal{M}}$. Its n -neighbourhood $\mathcal{N}_n^{\mathcal{M}, \mathcal{D}}(c)$ w.r.t. a subdomain $\mathcal{D} \subseteq \Delta^{\mathcal{M}}$ is defined inductively as:*

$$\begin{cases} \mathcal{N}_0^{\mathcal{M}, \mathcal{D}}(c) := \{c\} \\ \mathcal{N}_{n+1}^{\mathcal{M}, \mathcal{D}}(c) := \mathcal{N}_n^{\mathcal{M}, \mathcal{D}}(c) \cup \{e \mid \exists d \in \mathcal{N}_n^{\mathcal{M}, \mathcal{D}}(c) \setminus \mathcal{D}, \exists R \in \mathbf{N}_R^{\pm}, (d, e) \in R^{\mathcal{M}}\} \end{cases}$$

Observe that we stop adding successors when we reach an element from \mathcal{D} . In particular, for $c \in \mathcal{D}$, we have $\mathcal{N}_n^{\mathcal{M}, \mathcal{D}}(c) = \{c\}$ for every value of n . It follows that the statement ' $c_1 \in \mathcal{N}_n^{\mathcal{M}, \mathcal{D}}(c_2)$ iff $c_2 \in \mathcal{N}_n^{\mathcal{M}, \mathcal{D}}(c_1)$ ' does not hold in general.

Example 5. *In Figure 1d, neighbourhoods $\mathcal{N}_2^{\mathcal{I}'_e, \Delta^*}(\gamma)$ and $\mathcal{N}_2^{\mathcal{I}'_e, \Delta^*}(\delta)$ are depicted (in green, resp. blue). In particular, notice $\mathbf{a} \notin \mathcal{N}_2^{\mathcal{I}'_e, \Delta^*}(\delta)$ since $\alpha \in \Delta^*$.*

Recall that the definition of $\Delta^{\mathcal{I}'}$ ensures that any $c \in \Delta^{\mathcal{I}'} \setminus \Delta^*$ is actually an element of Δ° and therefore we have $c = \mathbf{a}w$ for some individual name \mathbf{a} and word $w \in \Omega^*$. The tree-shaped structure of Δ° ensures that for all n , there exists a unique prefix $r_{n,c}$ of $\mathbf{a}w$ such that (i) $f'(r_{n,c}) \in \mathcal{N}_n^{\mathcal{I}', \Delta^*}(c)$ and (ii) for any $d \in \mathcal{N}_n^{\mathcal{I}', \Delta^*}(c)$, there exists a unique word $w_{n,c}^d$ such that $d = f'(r_{n,c} \cdot w_{n,c}^d)$.

This leads us to characterize the n -neighbourhood of an element c via the following function $\chi_{n,c}$, whose domain Ω_n is the set of words over Ω with length $\leq 2n$. Notice that, departing from [10], we keep track of *sets* of satisfied concepts, in order to handle conjunction of concepts on the LHS of axioms.

$$\chi_{n,c} : \Omega_n \rightarrow \Delta^* \cup 2^{\text{sig}(\mathcal{T})} \cup \{\emptyset\}$$

$$w \mapsto \begin{cases} f'(r_{n,c}w) & \text{if } r_{n,c}w \in \Delta^{\mathcal{C}\kappa} \wedge f'(r_{n,c}w) \in \Delta^* \\ \{\mathbf{A} \mid \mathbf{A} \in \text{sig}(\mathcal{T}), f'(r_{n,c}w) \in \mathbf{A}^{\mathcal{I}'}\} & \text{if } r_{n,c}w \in \Delta^{\mathcal{C}\kappa} \wedge f'(r_{n,c}w) \notin \Delta^* \\ \emptyset & \text{otherwise} \end{cases}$$

We can now introduce the equivalence relation we use to merge elements:

Definition 4. The equivalence relation \sim_n on $\Delta^{\mathcal{I}'}$ is defined as follows: an element from Δ^* is \sim_n -equivalent only to itself; two elements c_1, c_2 from $\Delta^{\mathcal{I}'} \setminus \Delta^*$ are \sim_n -equivalent iff $w_{n,c_1}^{c_1} = w_{n,c_2}^{c_2}$, $\chi_{n,c_1} = \chi_{n,c_2}$, and $|c_1| = |c_2| \pmod{2|q|+3}$.

Remark 3. The set of concepts from $\text{sig}(\mathcal{T})$ satisfied by $c \in \Delta^{\mathcal{I}'}$ is exactly $\chi_{n,c}(w_{n,c}^c)$. Therefore, if $c \sim_n c'$, then c and c' satisfy the same concept names.

Remark 4. If $c \sim_n c'$, then $c \sim_m c'$ for any $m \leq n$.

We obtain a smaller countermodel for our CCQ q by merging elements with respect to $\sim_{|q|+1}$. We will use \bar{e} for the equivalence class of e w.r.t. $\sim_{|q|+1}$ and let $\mathbf{p} : e \mapsto \bar{e}$ denote the canonical projection from \mathcal{I}' to \mathcal{J} . To improve the readability of later material, we introduce notation $\overline{\Delta^*}$ for the set $\{\bar{\sigma} \mid \sigma \in \Delta^*\}$.

Definition 5. The reduced interlacing \mathcal{J} is the interpretation with domain $\Delta^{\mathcal{I}'} / \sim_{|q|+1}$ and interpretation function $\cdot^{\mathcal{J}} := \mathbf{p} \circ \cdot^{\mathcal{I}'}$.

Example 6. The reduced interlacing \mathcal{J}_e , together with two 2-neighbourhoods $\mathcal{N}_2^{\mathcal{J}_e, \overline{\Delta^*}}(\bar{\gamma})$ and $\mathcal{N}_2^{\mathcal{J}_e, \overline{\Delta^*}}(\bar{\delta})$, are displayed in Figure 1e. Notice \mathcal{J}_e remains a countermodel for q_e and candidate integer 3.

It remains to show that the reduced interlacing \mathcal{J} is a countermodel. As elements were merged using a local criteria, it is not possible in general to exhibit a homomorphism from the whole \mathcal{J} to \mathcal{I}' , injecting matches as in Lemma 3. However, local solutions are possible: a match of q in \mathcal{J} maps each connected component C of q into a $|q|$ -neighbourhood $\mathcal{N}_{|q|}^{\mathcal{J}, \overline{\Delta^*}}(\bar{c})$. By exhibiting a homomorphism $\rho_c : \mathcal{N}_{|q|}^{\mathcal{J}, \overline{\Delta^*}}(\bar{c}) \rightarrow \mathcal{N}_{|q|}^{\mathcal{I}', \Delta^*}(c)$, we can find a match of C in \mathcal{I}' . Such matches for q 's connected components together form a match of the full q in \mathcal{I}' .

We define such homomorphisms ρ_c inductively on $\mathcal{N}_k^{\mathcal{J}, \overline{\Delta^*}}(\bar{c})$ with k increasing from 0 to $|q|$. Starting from the element $\bar{c} \in \mathcal{N}_0^{\mathcal{J}, \overline{\Delta^*}}(\bar{c})$, we can naturally carry it back as $\rho_c(\bar{c}) = c \in \mathcal{N}_0^{\mathcal{I}', \Delta^*}(c)$. Assume now that we have defined $\rho_c(\bar{d})$ for some $\bar{d} \in \mathcal{N}_n^{\mathcal{J}, \overline{\Delta^*}}(\bar{c})$ and that we are moving further to an element $\bar{e} \in \mathcal{N}_{n+1}^{\mathcal{J}, \overline{\Delta^*}}(\bar{c})$ along an edge (\bar{d}, \bar{e}) in \mathcal{J} . In the case of $\bar{e} \notin \overline{\Delta^*}$, the following lemma produces a candidate $\rho_c(\bar{e})$, namely e' , which is to $\rho_c(\bar{d})$, namely d' , what \bar{e} is to \bar{d} .

Lemma 4. Given two elements $\bar{d}, \bar{e} \in \Delta^{\mathcal{J}} \setminus \overline{\Delta^*}$, if there exists a role P from \mathbf{N}_R^\pm such that $(\bar{d}, \bar{e}) \in P^{\mathcal{J}}$, then there exists a unique element $R.B \in \Omega$ such that one of the two following conditions is satisfied:

edge⁺. $|e| = |d| + 1 \pmod{2|q|+3}$, $w_{|q|+1,e}^e = w_{|q|+1-1,d}^d \cdot R.B$ and $\mathcal{T} \models R \sqsubseteq P$.

Furthermore, for all $d' \sim_k d$, the element $e' := d' \cdot R.B$ belongs to $\Delta^{\mathcal{I}'}$ and satisfies $e' \sim_{k-1} e$.

edge⁻. $|d| = |e| + 1 \pmod{2|q|+3}$, $w_{|q|+1,d}^d = w_{|q|+1-1,e}^e \cdot R.B$ and $\mathcal{T} \models R^- \sqsubseteq P$.

Furthermore, for all $d' \sim_k d$, we have e' such that $d' = e' \cdot R.B$ and the prefix e' satisfies $e' \sim_{k-1} e$.

Notice the “strength” of the equivalence relation \sim_k between \bar{e} and $\rho_c(\bar{e})$ decreases as we move further in the neighbourhood of \bar{c} . For example in Figure 1e: building $\rho_\gamma(\bar{\gamma}_2)$ from $\rho_\gamma(\bar{\gamma}) := \gamma$ we obtain γ_1 , with $\gamma_1 \sim_1 \gamma_2$ but $\gamma_1 \not\sim_2 \gamma_2$. However, since we start from $\rho_c(\bar{c}) := c \sim_{|q|+1} c$ and explore a $|q|$ -neighbourhood, the index remains at least 1. This is essential as \sim_1 encodes relations to elements of $\overline{\Delta^*}$ as the next lemma shows. It allows in particular to treat the case of $\bar{e} \in \overline{\Delta^*}$.

Lemma 5. *If $(\bar{d}, \bar{e}) \in \mathbb{R}^{\mathcal{J}}$ for some $e \in \Delta^*$, and if $d' \sim_1 d$, then $(d', e) \in \mathbb{R}^{\mathcal{I}'}$.*

It remains to free ourselves from the particular choice of \bar{d} , which is likely not to be the only element of $\mathcal{N}_n^{\mathcal{J}, \overline{\Delta^*}}(\bar{c})$ connected to \bar{e} . This cannot be observed on our running example as $|q_e| = 1$, but the possibility of cycles in the reduced interlacing should still be clear from Figure 1e. Taking a closer look at Lemma 4, we observe that $\rho_c(\bar{e})$, that is e' , is obtained either by adding a letter to $\rho_c(\bar{d})$, that is d' , or by removing the last letter of $\rho_c(\bar{d})$, and that these letters coincide with those in the suffixes of elements d and e . Therefore, when moving from \bar{c} to \bar{e} and ignoring self-cancelling steps, each added letter must appear in the suffix of e and, similarly, each removed letter must appear in the suffix of c . The challenge is therefore to quantify the number of additions and removals to build $\rho_c(\bar{e})$ directly from c and \bar{e} . The next definition captures the relative difference of letters between \bar{c} and \bar{e} , encoded in $|c|$ and $|e| \pmod{2|q| + 3}$.

Definition 6. *Let $\bar{c} \in \Delta^{\mathcal{J}}$ and $n \leq |q|$. The relative depth of $\bar{e} \in \mathcal{N}_n^{\mathcal{J}, \overline{\Delta^*}}(\bar{c})$ from \bar{c} is the integer $\delta_{\bar{c}}(\bar{e}) \in [-n, n]$ such that $|e| = |c| + \delta_{\bar{c}}(\bar{e}) \pmod{2|q| + 3}$.*

Remark 5. *By induction on $n \leq |q|$, it is straightforward to see that $\delta_{\bar{c}}(\bar{e})$ is well defined. Unicity is ensured by $\delta_{\bar{c}}(\bar{e}) \leq n \leq |q|$. A consequence of Lemma 4 is that for the smallest $n \leq |q|$ such that $\bar{e} \in \mathcal{N}_n^{\mathcal{J}, \overline{\Delta^*}}(\bar{c})$ we have $\delta_{\bar{c}}(\bar{e}) = n \pmod{2}$.*

We can now identify how many additions and removals cancelled each other. Indeed, if it takes n steps to reach \bar{e} from \bar{c} , with relative difference of $\delta := \delta_{\bar{c}}(\bar{e})$, then $n - |\delta|$ is the length of the self-cancelling path, hence: $\frac{n-|\delta|}{2}$ cancelled additions and $\frac{n-|\delta|}{2}$ cancelled removals. Therefore, the actual amount of additions is $\frac{n-|\delta|}{2} + \delta$ if $\delta \geq 0$, or $\frac{n-|\delta|}{2}$ if $\delta \leq 0$, that is in both cases $\frac{n+\delta}{2}$. Similarly we obtain $\frac{n-\delta}{2}$ for the actual amount of removals. The next theorem formalizes all these intuitions: $\rho_{n,c}(\bar{e})$ (in non-trivial cases) is obtained by removing the $\frac{n-\delta}{2}$ last letters of c and keeping the $\frac{n+\delta}{2}$ last letters from the suffix of e . For example on Figures 1d, element $\rho_\gamma(\bar{\gamma}_2) = \gamma_1$ is obtained by removing $\frac{1-(-1)}{2} = 1$ letter from γ and keeping $\frac{1+(-1)}{2} = 0$ letter from the suffix of γ_2 . It is then a technicality to verify these syntactical operations on words make sense in the domain of \mathcal{I}' .

Theorem 3. *For all $c \in \Delta^{\mathcal{I}'}$ and all $n \leq |q|$, the following mapping:*

$$\rho_{n,c}(\bar{e}) : \mathcal{N}_n^{\mathcal{J}, \overline{\Delta^*}}(\bar{c}) \rightarrow \mathcal{N}_n^{\mathcal{I}', \Delta^*}(c) \quad \bar{e} \mapsto \begin{cases} \rho_{n-1,c}(\bar{e}) & \text{if } \bar{e} \in \mathcal{N}_{n-1}^{\mathcal{J}, \overline{\Delta^*}}(\bar{c}) \\ e & \text{if } \bar{e} \in \overline{\Delta^*} \\ r_{\frac{n-\delta_{\bar{c}}(\bar{e})}{2}, c} \cdot w_{\frac{n+\delta_{\bar{c}}(\bar{e})}{2}, e} & \text{otherwise} \end{cases}$$

is a homomorphism satisfying $\rho_{n,c}(\bar{e}) \sim_{|q|+1-n} e$ and $\rho_{n,c}^{-1}(\overline{\Delta^}) \subseteq \overline{\Delta^*}$.*

Let us clarify how Theorem 3 concludes our proof with the following lemma.

Lemma 6. *If \mathcal{I} is a countermodel, then so is its reduced interlacing \mathcal{J} .*

Proof (sketch). It is mostly routine work of definition chaining to show that \mathcal{J} is a model, except for negative role inclusions, where Theorem 3 is needed to move violations of $R_1 \sqcap R_2 \sqsubseteq \perp$ in \mathcal{J} back into \mathcal{I}' . As sketched earlier, we can employ the local homomorphisms between neighbourhoods to transform a match of q in \mathcal{J} into a match in \mathcal{I}' . Matches in \mathcal{I}' being contained in Δ^* , our original match in \mathcal{J} must be contained in $\overline{\Delta^*}$. Thus, the mapping of matches in \mathcal{J} to matches in \mathcal{I}' is essentially the identity (if we conflate Δ^* and $\overline{\Delta^*}$), hence injective.

Finally, we obtain Theorem 1 by analyzing the size of \mathcal{J} (i.e. counting the equivalence classes in $\Delta^{\mathcal{J}}$), keeping in mind that due to Lemma 1 and our assumption on m , we have $|\Delta^*| \leq |\text{Ind}(\mathcal{A})| + |q| (|\text{Ind}(\mathcal{A})| + 3 |\mathcal{T}| 2^{|\mathcal{T}|})^{|q|}$.

4 Lower bounds for \mathcal{EL}

We now consider the simpler setting of \mathcal{EL} , incomparable with DL-Lite $_{\mathcal{R}}$, and show that the same lower bounds hold, both in data and combined complexity.

Theorem 4. *CCQ answering in \mathcal{EL} is coNP-complete w.r.t. data complexity.*

Proof (sketch). We reduce the complement of the graph 3-colorability problem to answering the \mathcal{EL} OMQ (q, \mathcal{T}) , with $q = \exists z B(z)$ and \mathcal{T} containing $A \sqsubseteq \exists R.B$ and $\exists R.C_k \sqcap \exists E.(\exists R.C_k) \sqsubseteq B$ for $k \in \{1, 2, 3\}$.

Theorem 5. *CCQ answering in \mathcal{EL} is coNEXP-hard w.r.t. combined complexity.*

Proof (sketch). The proof adapts a reduction from the exponential grid tiling problem (Lemma 18 from [10]), the key difference being the use of existential restriction to replace role inclusions.

5 Outlook

We have initiated the study of CCQ answering for Horn DLs outside the DL-Lite family, establishing the same complexity bounds for $\mathcal{EL}(\mathcal{H}\mathcal{I}_\perp)$ as were known for DL-Lite $_{\mathcal{R}}$. There remain many questions to explore. Let us mention two potential research directions towards obtaining more practical algorithms. First, we can look for additional restrictions on the query or the ontology that ensure polynomial data complexity, as has been considered for DL-Lite [6,3,4]. Unfortunately, it appears that restrictions that work for DL-Lite are not sufficient to obtain tractability in \mathcal{EL} , so novel restrictions need to be identified. Second, it would be desirable, for \mathcal{EL} but also for DL-Lite $_{\mathcal{R}}$, to develop a more refined coNP procedure that is amenable to implementation, e.g. using SAT solvers.

References

1. Baader, F., Brandt, S., Lutz, C.: Pushing the EL envelope. In: Proceedings of the 19th international joint conference on Artificial intelligence, IJCAI. pp. 364–369 (2005)
2. Bienvenu, M., Calvanese, D., Ortiz, M., Simkus, M.: Nested regular path queries in description logics. In: Proc. of the 14th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR). pp. 218–227 (2014)
3. Bienvenu, M., Manière, Q., Thomazo, M.: Answering counting queries over DL-Lite ontologies. In: Proc. of the 29th International Joint Conference on Artificial Intelligence (IJCAI). pp. 1608–1614 (2020)
4. Bienvenu, M., Manière, Q., Thomazo, M.: Cardinality queries over DL-Lite ontologies. In: Proc. of the 30th International Joint Conference on Artificial Intelligence (IJCAI). pp. 1801–1807 (2021)
5. Bienvenu, M., Ortiz, M.: Ontology-mediated query answering with data-tractable description logics. In: Tutorial Lectures of the 11th Reasoning Web International Summer School. pp. 218–307 (2015)
6. Calvanese, D., Corman, J., Lanti, D., Razniewski, S.: Counting query answers over a DL-Lite knowledge base. In: Proc. of the 29th International Joint Conference on Artificial Intelligence (IJCAI). pp. 1658–1666 (2020)
7. Calvanese, D., Corman, J., Lanti, D., Razniewski, S.: Rewriting count queries over DL-Lite TBoxes with number restrictions. In: Proc. of the 33rd International Workshop on Description Logics (DL) (2020)
8. Calvanese, D., Kharlamov, E., Nutt, W., Thorne, C.: Aggregate queries over ontologies. In: Proc. of the 2nd International Workshop on Ontologies and Information Systems for the Semantic Web (ONISW). pp. 97–104 (2008)
9. Feier, C., Lutz, C., Przybylko, M.: Answer counting under guarded TGDs. In: Proc. of the 24th International Conference on Database Theory (ICDT) (2021)
10. Kostylev, E.V., Reutter, J.L.: Complexity of answering counting aggregate queries over DL-Lite. *Journal of Web Semantics (JWS)* pp. 94–111 (2015)
11. Nikolaou, C., Kostylev, E.V., Konstantinidis, G., Kaminski, M., Grau, B.C., Horrocks, I.: Foundations of ontology-based data access under bag semantics. *Artificial Intelligence (AIJ)* pp. 91 – 132 (2019)