Spectra of Cardinality Queries over Description Logic Knowledge Bases

Quentin Manière[1,2], **Marcin Przybyłko**[1,3]
[1]Department of Computer Science, Leipzig University, Germany
[2]Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI), Dresden/Leipzig, Germany
[3]Institute of Informatics, University of Warsaw, Poland
quentin.maniere@uni-leipzig.de, m.przybylko@uw.edu.pl

Abstract

Recent works have explored the use of counting queries coupled with Description Logic ontologies. The answer to such a query in a model of a knowledge base is either an integer or ∞, and its spectrum is the set of its answers over all models. While it is unclear how to compute and manipulate such a set in general, we identify a class of counting queries whose spectra can be effectively represented. Focusing on atomic counting queries, we pinpoint the possible shapes of a spectrum over $\mathcal{ALCIF}$ ontologies: they are essentially the subsets of $\mathbb{N} \cup \{\infty\}$ closed under addition. For most sublogics of $\mathcal{ALCIF}$, we show that possible spectra enjoy simpler shapes, being $[\![m, \infty]\!]$ or variations thereof. To obtain our results, we refine constructions used for finite model reasoning and notably rely on a cycle-reversion technique for the Horn fragment of $\mathcal{ALCIF}$. We also study the data complexity of computing the proposed effective representation and establish the $\mathsf{FP}^{\mathsf{NP}[\log]}$-completeness of this task under several settings.

1 Introduction

Ontology-mediated query answering (OMQA) uses ontologies to offer a user-friendly vocabulary for formulating queries or to encapsulate domain knowledge that can be utilized to retrieve more comprehensive answers (Poggi et al. 2008; Xiao et al. 2018). Ontologies expressed in Description Logics (DLs), a family of knowledge representation languages underpinning the OWL Web Ontology Language, have received special attention (Artale et al. 2009; Baader et al. 2017), and the core reasoning task of OMQA, the query answering task, has been extensively studied for conjunctive queries (CQs) and unions thereof. Under the OMQA framework, answering CQs is addressed by considering every possible model of the knowledge base (KB), that is every extension of the data that satisfies the ontology, and returning so-called *certain answers*, *i.e.* answers true in every model.

A recent line of research has explored ways of leveraging OMQA to support counting queries, a well-known class of aggregate queries that allows to perform analytics on data. Several semantics for such queries have been investigated, differing on how the possibility of multiple models is taken into account. In (Feier, Lutz, and Przybylko 2021), this has been addressed by returning the number of certain answers to a query, while in (Calvanese et al. 2008) an epistemic semantics was adopted – enforcing the counting operator to only involve known data values and making it possible to use the usual notion of certain answers.

In this paper we adopt the semantic of (Kostylev and Reutter 2015; Bienvenu, Manière, and Thomazo 2020) that defines a counting query as a CQ in which some variables have been designated as *counting variables*. The answer to a counting query in a model of the KB is then obtained as the number of different assignments for the counting variables when considering every possible homomorphism of the CQ into the model. Finding uniform bounds on those answers, *i.e.* model-independent bounds, has been viewed as a notion of certain answers and is now well-understood for a variety of DLs (Calvanese et al. 2020; Bienvenu, Manière, and Thomazo 2022; Manière 2022). The following example highlights that even the tightest uniform bounds give, in general, a poor over-approximation of the set of answers.

Example 1. *Consider an empty KB $\mathcal{K}$ and a counting query q asking for the number of pairs (z_1, z_2) such that z_1 and z_2 are friends of* alice, *that is $q = \exists z_1\, \exists z_2\, \mathit{friendOf}(z_1, \mathsf{alice}) \wedge \mathit{friendOf}(z_2, \mathsf{alice})$. Clearly, the set of possible answers to q across models of $\mathcal{K}$ is $\{n^2 \mid n \in \mathbb{N}\} \cup \{\infty\}$. The tightest uniform bounds on this set are given by the interval $[\![0, \infty]\!]$.*

Rather than aiming for an over-approximation of the set of possible answers, we intent to give a comprehensive description of this subset of $\mathbb{N}^\infty = \{0, 1, 2, \ldots, \infty\}$ that we call the *spectrum of the counting query*, inspired by the notion of spectrum of a formula that refers to the possible cardinalities of its models (Fagin 1974; Durand, Fagin, and Loescher 1997). We investigate the possible shapes of these spectra for *counting conjunctive queries (CCQs)* mentioned above and for ontologies expressed in the $\mathcal{ALCIF}$ DL. This expressive DL is contained in $\mathcal{SHIQ}$, in which traditional CQ answering is well-understood (Glimm et al. 2008; Lutz 2008), and supports functionality constraints whose interactions with counting queries have never been studied to the best of our knowledge (those proposed in (Calvanese et al. 2020) and denoted $\mathcal{N}^-$ being much more restricted).

One of the challenges encountered in our work is to clarify how to represent spectra. Indeed, the set of possible answers of a CCQ across models of a KB might, *a priori*, be an arbitrary set of natural numbers, and thus hard to describe by means other than providing the CCQ-KB couple. We aim to identify classes of ontology-mediated queries (OMQs) whose spectra admit an *effective representation*. By effective, we intend a representation that is (i) finite, ideally with a size that can be bounded by the size of the OMQ, (ii) independent from the specific ontology language and (iii) spectrum membership can be efficiently tested, *i.e.* in polynomial time w.r.t. the size of the integer and of the representation.

Contributions. We introduce the notion of a spectrum for a CCQ and show that connected and individual-free CCQs evaluated on $\mathcal{ALCIF}$ KBs always admit well-behaved spectra, as those are subsets of $\mathbb{N}^\infty$ closed under addition. We then propose an effective representation of such spectra.

This motivates a focus on cardinality queries, *i.e.* Boolean atomic CCQs (Bienvenu, Manière, and Thomazo 2021), that fall in the above class. First, we fully characterize possible spectra shapes for concept cardinality queries on $\mathcal{ALCIF}$ KBs, showing that every subset of $\mathbb{N}^\infty$ closed under addition is realizable. We also study several sublogics of $\mathcal{ALCIF}$, extending $\mathcal{EL}$ and DL-Lite$_{\mathsf{core}}$, for most presenting full characterizations. For some, only simpler shapes, such as $[\![m, \infty]\!]$, are possible. For $\mathcal{ELIF}_\bot$, the Horn fragment of $\mathcal{ALCIF}$, we notably use variations of the cycle-reversion techniques introduced to tackle finite model reasoning in such DLs (Cosmadakis, Kanellakis, and Vardi 1990; Rosati 2008; Ibáñez-García, Lutz, and Schneider 2014).

We further study the data complexity of computing the proposed effective representations of spectra. For many settings, such as concept cardinality queries on $\mathcal{ALC}$ KBs, we are able to establish $\mathsf{FP}^{\mathsf{NP}[\log]}$-completeness of this problem. Several of our upper bounds notably rely on existing results regarding DLs equipped with closed predicates.

Via connections with the concept cardinality case and refinements of the corresponding constructions, we also investigate the case of role cardinality queries. This latter class of OMQs features challenging shapes of spectra already for $\mathcal{EL}_\bot$ KBs, and we prove that computing effective representations of those is $\mathsf{FP}^{\mathsf{NP}[\log]}$-complete already for $\mathcal{EL}$ KBs.

An appendix with full proofs can be found in the long version of this paper, see (Manière and Przybyłko 2024).

2 Preliminaries

With $\mathbb{N}$ we denote the set of natural numbers $\mathbb{N} = \{0, 1, \dots\}$ and by $(\mathbb{N}^\infty, +)$ the semigroup of natural numbers with infinity ∞ and the usual definition of addition $+$. In particular, $a + \infty = \infty + a = \infty$ for all elements $a \in \mathbb{N}^\infty$. We recall that every subsemigroup of $(\mathbb{N}^\infty, +)$, *i.e.* every subset closed under addition, is *ultimately periodic* (see *e.g.* (Grillet 2001), Chapter 2, Proposition 4.1), which ensures that every subsemigroup of $\mathbb{N}^\infty$ takes the following shape.

Lemma 1. *Let V be a subsemigroup of $(\mathbb{N}^\infty, +)$. Then $V = S \cup \{M + \alpha \cdot n \mid n \in \mathbb{N}\}$ where S is a finite subset of $\mathbb{N}^\infty$ and $M, \alpha \in \mathbb{N}^\infty$.*

If $\alpha = 1$, we write $S \cup [\![M, \infty[\![$ for $S \cup \{M + n \mid n \in \mathbb{N}\}$.

2.1 $\mathcal{ALCIF}$ and Other Description Logics

Let $\mathsf{N_C}$, $\mathsf{N_R}$, and $\mathsf{N_I}$ be countably infinite and mutually disjoint sets of *concept names*, *role names*, and *individual names*. An *inverse role* takes the form r^- with r a role name, and a *role* is a role name or an inverse role. We denote $\mathsf{N}_\mathsf{R}^\pm$ the set of roles. If $\mathrm{r} = \mathrm{s}^-$ is an inverse role, then r^- denotes s. An $\mathcal{ALCI}$ *concept* is built according to the rule $\mathrm{C}, \mathrm{D} ::= \top \mid \mathrm{A} \mid \neg\mathrm{C} \mid \mathrm{C} \sqcap \mathrm{D} \mid \exists\mathrm{r}.\mathrm{D}$ where $\mathrm{A} \in \mathsf{N_C}$ and $\mathrm{r} \in \mathsf{N}_\mathsf{R}^\pm$. We use $\bot$ as an abbreviation for $\neg\top$, $\mathrm{C} \sqcup \mathrm{D}$ for $\neg(\neg\mathrm{C} \sqcap \neg\mathrm{D})$, $\forall\mathrm{r}.\mathrm{C}$ for $\neg\exists\mathrm{r}.\neg\mathrm{C}$ and $\exists\mathrm{r}$ for $\exists\mathrm{r}.\top$.

An $\mathcal{ALCIF}$ *TBox* is a finite set of *concept inclusions (CIs)* $\mathrm{C} \sqsubseteq \mathrm{D}$ and of *functionality restrictions* $\mathrm{C} \sqsubseteq \leq 1\ \mathrm{r}.\mathrm{D}$ where C, D are $\mathcal{ALCI}$ concepts and r is a role. An *ABox* is a finite set of *concept assertions* $\mathrm{A}(\mathsf{a})$ and *role assertions* $\mathrm{r}(\mathsf{a}, \mathsf{b})$ where $\mathrm{A} \in \mathsf{N_C}$, $\mathrm{r} \in \mathsf{N_R}$ and $\mathsf{a}, \mathsf{b} \in \mathsf{N_I}$. The set of individual names used in the ABox $\mathcal{A}$ is denoted $\mathsf{Ind}(\mathcal{A})$. An $\mathcal{ALCIF}$ *knowledge base (KB)* takes the form $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ with $\mathcal{T}$ an $\mathcal{ALCIF}$ TBox and $\mathcal{A}$ an ABox.

We also investigate restrictions of $\mathcal{ALCIF}$. Each fragment is obtained by disallowing concepts, roles constructors, or axiom shapes in the standard way. An $\mathcal{EL}$ *concept* is an $\mathcal{ALCI}$ concept that uses neither negation nor inverse roles and an $\mathcal{EL}$ TBox only supports CIs of $\mathcal{EL}$ concepts. Allowing inverse roles is indicated by $\mathcal{I}$; *concept disjointness* axioms of shape $\mathrm{C} \sqcap \mathrm{D} \sqsubseteq \bot$ by subscript $\bot$; unrestricted use of negation by replacing $\mathcal{EL}$ with $\mathcal{ALC}$; and functionality restrictions by $\mathcal{F}$. A DL-Lite concept has shape $\mathrm{A} \mid \exists\mathrm{r}$ for $\mathrm{A} \in \mathsf{N_C}$ and $\mathrm{r} \in \mathsf{N}_\mathsf{R}^\pm$. A DL-Lite$_{\mathsf{core}}$ TBox only supports CIs and CDs of DL-Lite concepts. DL-Lite$_\mathcal{F}$ extends DL-Lite$_{\mathsf{core}}$ with unqualified functionality restrictions $\top \sqsubseteq \leq 1\ \mathrm{r}.\top$ (Calvanese et al. 2006).

The semantics is defined in terms of interpretations $\mathcal{I} = (\Delta^\mathcal{I}, \cdot^\mathcal{I})$ in the standard way. $\Delta^\mathcal{I}$ is a non-empty *domain* and $\cdot^\mathcal{I}$ an *interpretation function*. We refer to (Baader et al. 2017) for details. An interpretation $\mathcal{I}$ satisfies a CI $\mathrm{C} \sqsubseteq \mathrm{D}$ if $\mathrm{C}^\mathcal{I} \subseteq \mathrm{D}^\mathcal{I}$ and a functionality restriction $\mathrm{C} \sqsubseteq \leq 1\ \mathrm{r}.\mathrm{D}$ if for each $d \in \mathrm{C}^\mathcal{I}$, there is at most one $e \in \mathrm{D}^\mathcal{I}$ such that $(d, e) \in \mathrm{r}^\mathcal{I}$. It satisfies an assertion $\mathrm{A}(\mathsf{a})$ if $\mathsf{a} \in \mathrm{A}^\mathcal{I}$ and $\mathrm{r}(\mathsf{a}, \mathsf{b})$ if $(\mathsf{a}, \mathsf{b}) \in \mathrm{r}^\mathcal{I}$. We make the *standard names assumption*: in every interpretation $\mathcal{I}$, we assume $\mathsf{a}^\mathcal{I} = \mathsf{a}$ for every $\mathsf{a} \in \mathsf{Ind}(\mathcal{A})$. An interpretation $\mathcal{I}$ is a *model* of a TBox $\mathcal{T}$, denoted $\mathcal{I} \models \mathcal{T}$, if it satisfies all its axioms. Models of ABoxes and KBs are defined likewise. A TBox $\mathcal{T}$ (resp. a KB $\mathcal{K}$) entails an assertion, a CI or a functionality restriction α, denoted $\mathcal{T} \models \alpha$ (resp. $\mathcal{K} \models \alpha$) if all its models satisfy α.

2.2 Spectra of Counting Queries

We consider two countably infinite and mutually disjoint sets: a set of *variables* and a set of *counting variables*. A *counting conjunctive query (CCQ)* takes the form $q(\bar{x}) = \exists\bar{y}\ \exists\bar{z}\ \psi(\bar{x}, \bar{y}, \bar{z})$, where $\bar{x}$ and $\bar{y}$ are tuples of distinct variables, $\bar{z}$ is a tuple of distinct counting variables and ψ is a conjunction of concept and role atoms whose terms are drawn from $\mathsf{N_I} \cup \bar{x} \cup \bar{y} \cup \bar{z}$. We call $\bar{x}$ (resp. $\bar{y}$, resp. $\bar{z}$) the *answer* (resp. *existential*, resp. *counting*) variables of q. A CCQ is *Boolean* if $\bar{x} = \emptyset$.

For a tuple $\bar{a} \in \mathsf{N_I}^{|\bar{a}|}$ of individuals and a model $\mathcal{I}$ of a KB $\mathcal{K}$, we define $\#q(\bar{a})^\mathcal{I}$ the answer of $q(\bar{a})$ on $\mathcal{I}$ as:

$$\#\{\pi_{|\bar{z}} \mid \pi : q \to \mathcal{I} \text{ homomorphism s.t. } \pi(\bar{x}) = \bar{a}\}.$$

The spectrum of $q(\bar{a})$ on $\mathcal{K}$ is further defined as:

$$\mathsf{Sp}_\mathcal{K}(q(\bar{a})) := \{\#q(\bar{a})^\mathcal{I} \mid \mathcal{I} \models \mathcal{K}\}.$$

Note that $\mathsf{Sp}_{\mathcal{K}}(q(\bar{a})) = \mathsf{Sp}_{\mathcal{K}}(q[\bar{a}/\bar{x}](\bar{a}_\emptyset))$, where $q[\bar{a}/\bar{x}]$ denotes the Boolean CCQ obtained from q by substituting every answer variable $x_i \in \bar{x}$ by the corresponding $a_i \in \bar{a}$, and $\bar{a}_\emptyset$ the empty tuple. We thus focus w.l.o.g. on Boolean CCQs $q(\bar{x}_\emptyset)$, denoted simply q for readability.

The main interest of this paper is to compute representations of spectra that are effective in the sense of Points (i)–(iii) in the introduction. While we do not know whether all spectra can be effectively represented, we identify a class of OMQs, namely connected and individual-free CCQs on $\mathcal{ALCIF}$ KBs, whose spectra admit such a representation. We recall that q is *connected* if its Gaifman graph is, and is *individual-free* if none of its atom involves a term from $\mathsf{N_I}$.

Lemma 2. *If $\mathcal{K}$ is an $\mathcal{ALCIF}$ KB and q a connected and individual-free CCQ, then $\mathsf{Sp}_{\mathcal{K}}(q)$ is closed under addition. Furthermore, if q is satisfiable w.r.t. $\mathcal{K}$, then $\infty \in \mathsf{Sp}_{\mathcal{K}}(q)$.*

In other words, spectra of connected and individual-free CCQs are subsemigroups of $\mathbb{N}^\infty$ and, by Lemma 1, are of form $S \cup \{M + \alpha \cdot n \mid n \in \mathbb{N}\}$. Thus, for this class, computing representations of spectra can be defined as follows.

Problem 1. *Given a KB $\mathcal{K}$ and a CCQ q, compute a special value $\emptyset$ if $\mathsf{Sp}_{\mathcal{K}}(q) = \emptyset$, otherwise a finite set $S \subseteq \mathbb{N}^\infty$ and numbers $M, \alpha \in \mathbb{N}^\infty$ s.t. $\mathsf{Sp}_{\mathcal{K}}(q) = S \cup \{M + \alpha \cdot n \mid n \in \mathbb{N}\}$.*

It can be verified that such representations as triples (S, M, α) comply with Points (i)–(iii) from the introduction and are, in this sense, effective.

Remark 1. *Notice $\mathsf{Sp}_{\mathcal{K}}(q) = \emptyset$ iff $\mathcal{K}$ is unsatisfiable; and, likewise, $\mathsf{Sp}_{\mathcal{K}}(q) = \{0\}$ iff $\mathcal{K}$ is satisfiable but q is unsatisfiable w.r.t. $\mathcal{K}$. In particular, if $\mathcal{K}$ is an $\mathcal{ELI}$ KB, then $\mathsf{Sp}_{\mathcal{K}}(q)$ cannot be $\emptyset$ nor $\{0\}$. Similarly, if $\mathcal{K}$ is an $\mathcal{ELIF}$ KB, then $\mathsf{Sp}_{\mathcal{K}}(q)$ cannot be $\{0\}$. An effective representation of $\{0\}$ in the sense of Problem 1 is $(S, M, \alpha) = (\emptyset, 0, 0)$.*

A subset of $\mathbb{N}^\infty$ is *trivial* if it is either $\emptyset$ or $\{0\}$.

We highlight that Example 1 illustrates a situation in which the individual-freeness condition is not met.

3 Spectrum of a Concept Cardinality Query

In this section, we focus on concept cardinality queries $q_{\mathrm{C}} := \exists z\ \mathrm{C}(z)$, where $\mathrm{C} \in \mathsf{N_C}$ is a concept name and z a counting variable. Computing the spectrum of q_{C} over a KB $\mathcal{K}$ thus corresponds to the natural task of deciding the possible values of $\left|\mathrm{C}^{\mathcal{I}}\right|$ across the models $\mathcal{I}$ of $\mathcal{K}$. Every concept cardinality query satisfies preconditions of Lemma 2 and thus its spectrum can be represented as in Problem 1. Conversely, one can ask which sets are spectra of such queries. We say a set V is *$\mathcal{L}$-concept realizable* if there is a concept C and a $\mathcal{L}$ KB $\mathcal{K}$ s.t. $\mathsf{Sp}_{\mathcal{K}}(q_{\mathrm{C}}) = V$. We begin with $\mathcal{ALCIF}$ KBs and prove they can realize all subsemigroups of $\mathbb{N}^\infty$.

Theorem 1. *A non-trivial subset of $\mathbb{N}^\infty$ is $\mathcal{ALCIF}$-concept realizable iff it is a subsemigroup of $\mathbb{N}^\infty$ containing ∞.*

Notice Lemma 2 already ensures that being a subsemigroup of $\mathbb{N}^\infty$ containing ∞ is necessary. The other direction is a generalization of the following example that illustrates how to realize a shape of spectrum with $\alpha = 2$.

Example 2. *Consider the $\mathcal{ALCIF}$ TBox $\mathcal{T} = \{\ \top \sqsubseteq \mathrm{C},\ \mathrm{A} \sqsubseteq \exists \mathrm{r}.\neg\mathrm{A},\ \neg\mathrm{A} \sqsubseteq \exists \mathrm{r}.\mathrm{A},\ \top \sqsubseteq\ \leq 1\, \mathrm{r}.\top,\ \top \sqsubseteq\ \leq 1\, \mathrm{r}^-.\top\ \}$. Then, $\mathsf{Sp}_{(\mathcal{T},\emptyset)}(q_{\mathrm{C}}) = 2\mathbb{N} \cup \{\infty\}$.*

Notice that to achieve the non-trivial period of $\alpha = 2$ in the spectrum $2\mathbb{N} \cup \{\infty\}$, we rely on a role that is both functional and inverse functional. In fact, limiting one of these two features forces a trivial periodic behavior, *i.e.* $\alpha = 1$, and further allows for easier computation of the spectra.

3.1 Limiting Inverse Functional Roles

We now move towards $\mathcal{ALCF}$ and $\mathcal{ALCI}$, in which the functionality of an inverse role cannot be expressed. As a consequence, spectra of a concept cardinality query over such KBs enjoy the following well-behaved shapes.

Theorem 2. *A non-trivial subset of $\mathbb{N}^\infty$ is $\mathcal{ALCI}$- (resp. $\mathcal{ALCF}$-) concept realizable iff it has shape $\{0\} \cup [\![M, \infty]\!]$ or shape $[\![M, \infty]\!]$ for some $M \in \mathbb{N}$.*

The main ingredient for the 'only-if' part of Theorem 2 is a technique that extends any model $\mathcal{I}$ in which $\left|\mathrm{C}^{\mathcal{I}}\right| > 0$ into a model $\mathcal{J}$ with $\left|\mathrm{C}^{\mathcal{J}}\right| = \left|\mathrm{C}^{\mathcal{I}}\right| + 1$, as used in (Baader, Bednarczyk, and Rudolph 2020) for $\mathcal{ALCF}$ KBs. Conversely, it is not difficult to find KBs, already in DL-Lite$_{\mathsf{core}}$ or $\mathcal{EL}_\bot$, that realize these shapes notably relying on CD axioms for the shape $\{0\} \cup [\![M, \infty]\!]$.

Moreover, if we focus on negation-free DLs, the situation becomes even more favorable:

Theorem 3. *A non-trivial subset of $\mathbb{N}^\infty$ is $\mathcal{ELI}$- (resp. $\mathcal{ELF}$-) concept realizable iff it has shape $[\![M, \infty]\!]$ for some $M \in \mathbb{N}$. For $\mathcal{ELIF}$, the shape $\{\infty\}$ is also permitted.*

For the shape $\{\infty\}$, we use the following well-known example of an $\mathcal{ELIF}$ KB (notice it is also a DL-Lite$_{\mathcal{F}}$ KB).

Example 3. *Consider the $\mathcal{ELIF}$ KB $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ with $\mathcal{A} = \{\mathrm{r}(\mathsf{a},\mathsf{a}), \mathrm{r}(\mathsf{a},\mathsf{b})\}$ and $\mathcal{T} = \{\ \mathrm{C} \sqsubseteq \exists\mathrm{r}.\top,\ \exists\mathrm{r}^-.\top \sqsubseteq \mathrm{C},\ \top \sqsubseteq\ \leq 1\, \mathrm{r}^-.\top\ \}$. It can be verified that $\mathsf{Sp}_{\mathcal{K}}(q_{\mathrm{C}}) = \{\infty\}$.*

3.2 $\mathcal{ELIF}_\bot$ KBs and Cycles Reversion

We now turn to the two remaining DLs, namely $\mathcal{ELIF}_\bot$ and DL-Lite$_{\mathcal{F}}$, in which inverse functional roles and negation are supported. We begin with an example illustrating that, compared to the previously investigated restrictions of $\mathcal{ALCIF}$, new spectrum shapes can be realized.

Example 4. *We construct an $\mathcal{ELIF}_\bot$ KB $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ s.t. $\mathsf{Sp}_{\mathcal{K}}(q_{\mathrm{C}}) = \{4\} \cup [\![6, \infty]\!]$. The TBox $\mathcal{T}$ contains the axioms:*

$$\top \sqsubseteq \mathrm{C} \sqcap \exists\mathrm{r}.\mathrm{A}_1 \sqcap \exists\mathrm{r}.\mathrm{A}_2 \qquad \top \sqsubseteq\ \leq 1\, \mathrm{s}^-.\top$$
$$\exists\mathrm{r}.\mathrm{X} \sqsubseteq \mathrm{Y} \qquad \exists\mathrm{r}.\mathrm{Y} \sqsubseteq \mathrm{X} \qquad \mathrm{Y} \sqsubseteq \exists\mathrm{s}.\mathrm{X} \qquad \mathrm{X} \sqsubseteq \exists\mathrm{s}.\mathrm{Y}$$
$$\mathrm{A}_1 \sqcap \mathrm{A}_2 \sqsubseteq \bot \qquad \mathrm{X} \sqcap \mathrm{Y} \sqsubseteq \bot$$

The ABox $\mathcal{A}$ contains the 8 concept assertions $\mathrm{A}_1(\mathsf{x}_1)$, $\mathrm{A}_2(\mathsf{x}_2)$, $\mathrm{A}_1(\mathsf{y}_1)$, $\mathrm{A}_2(\mathsf{y}_2)$, $\mathrm{X}(\mathsf{x}_1)$, $\mathrm{X}(\mathsf{x}_2)$, $\mathrm{Y}(\mathsf{y}_1)$, $\mathrm{Y}(\mathsf{y}_2)$, and the 12 role assertions $\mathrm{s}(u_1, v_1)$, $\mathrm{s}(u_2, v_2)$, $\mathrm{r}(u_i, v_j)$ for each $i, j \in \{1, 2\}$ and each $u, v \in \{\mathsf{x}, \mathsf{y}\}$ with $u \neq v$.

A representation of this spectrum according to Problem 1 is $(S, M, \alpha) := (\{4\}, 6, 1)$. Such possibly non-trivial part S of the spectrum make a full characterization of realizable sets hard to reach. Interestingly, however, every $\mathcal{ELIF}_\bot$-concept realizable set can be represented with $\alpha = 1$.

Theorem 4. *If a non-trivial subset of $\mathbb{N}^\infty$ is $\mathcal{ELIF}_\bot$-concept realizable, then it has shape $\{\infty\}$, $\{0, \infty\}$, or $S \cup [\![M, \infty]\!]$ for some $M \in \mathbb{N}$ and $S \subseteq [\![0, M]\!]$.*

The remainder of this section is devoted to the proof of this theorem. Let us first eliminate the easy cases, proving $\{\infty\}$ and $\{0,\infty\}$ are already DL-Lite$_\mathcal{F}$-concept realizable. The $\{\infty\}$ shape has been obtained in Example 3. To realize $\{0,\infty\}$, we rely on concept disjointness as follows:

Example 5. *Consider the DL-Lite$_\mathcal{F}$ TBox $\mathcal{T}$ containing:*

$$\begin{array}{llll} \mathrm{C} \sqsubseteq \exists \mathrm{r} & \exists \mathrm{r}^- \sqsubseteq \mathrm{C} & \top \sqsubseteq\ \leq 1\ \mathrm{r}^-.\top & \multirow{2}{*}{$\exists \mathrm{r}^- \sqcap \exists \mathrm{s}^- \sqsubseteq \bot$} \\ \mathrm{C} \sqsubseteq \exists \mathrm{s} & \exists \mathrm{s}^- \sqsubseteq \mathrm{C} & \top \sqsubseteq\ \leq 1\ \mathrm{s}^-.\top & \end{array}$$

It is immediate to verify that $\mathsf{Sp}_{(\mathcal{T},\emptyset)}(q_\mathrm{C}) = \{0,\infty\}$.

It remains to verify that every other non-trivial subset V of $\mathbb{N}^\infty$ that is $\mathcal{ELIF}_\bot$-concept realizable has shape $S \cup [\![M,\infty]\!]$ for some $M \in \mathbb{N}$ and $S \subseteq [\![0,M]\!]$. Let V be such a set and $\mathcal{K}$ an $\mathcal{ELIF}_\bot$ KB s.t. $V = \mathsf{Sp}_\mathcal{K}(q_\mathrm{C})$ for some concept name C. We prove that $\mathsf{Sp}_\mathcal{K}(q_\mathrm{C})$ actually contains two consecutive non-zero integers n and $n+1$, which guarantees, from closure under addition, that every integer greater than $n(n+1)$ is also in $\mathsf{Sp}_\mathcal{K}(q_\mathrm{C})$. Setting $M = n(n+1)$ and $S = \mathsf{Sp}_\mathcal{K}(q_\mathrm{C}) \cap [\![0,M]\!]$ will then conclude the proof.

Since V is non-trivial and neither $\{0,\infty\}$ nor $\{\infty\}$, it contains a non-zero integer. In other words, the concept C admits a finite interpretation in some (potentially infinite) model ($\star$). To exploit this fact, we refine cycle-reversion techniques which have been developed to study finite reasoning in similar logics (Cosmadakis, Kanellakis, and Vardi 1990; Rosati 2008; Ibáñez-García, Lutz, and Schneider 2014). More precisely, we tailor the notion of cycles to characterize under which conditions the interpretation of C may be finite. By ($\star$), those conditions are satisfied and we adapt a construction from the latter reference to produce models $\mathcal{I}$ and $\mathcal{J}$ of $\mathcal{K}$ s.t $\left|\mathrm{C}^\mathcal{J}\right| = \left|\mathrm{C}^\mathcal{I}\right| + 1 < \infty$ as desired. Henceforth, we assume $\mathcal{ELIF}_\bot$ KBs to be in normal form, that is every axiom in the TBox has one of the following shapes:

$$K \sqsubseteq \mathrm{A} \quad K \sqsubseteq \exists \mathrm{r}.K' \quad \exists \mathrm{r}.K \sqsubseteq K' \quad K \sqsubseteq\ \leq 1\ \mathrm{r}.K'$$

where $\mathrm{A} \in \mathsf{N_C} \cup \{\bot\}$, $\mathrm{r} \in \mathsf{N_R^\pm}$ and K, K' are conjunctions of concepts names. This is a reformulation of the normal form used in (Ibáñez-García, Lutz, and Schneider 2014) and it can be verified that putting a KB in such a normal form does not affect spectra of queries on this KB.

We now present our refined notion of cycles which itself relies on the following definition of inverse functional paths.

Definition 1. *An* inverse functional path (IFP) in $\mathcal{T}$ *is a sequence $K_0, \mathrm{r}_1, K_1, \ldots, \mathrm{r}_n, K_n$ where $n \geq 1$, $K_0, \ldots, K_n$ are conjunctions of concept names and $\mathrm{r}_1, \ldots, \mathrm{r}_n$ are (potentially inverse) roles s.t. for all $0 \leq i < n$:*

$$\mathcal{T} \models K_i \sqsubseteq \exists \mathrm{r}_{i+1}.K_{i+1} \quad \textit{and} \quad \mathcal{T} \models K_{i+1} \sqsubseteq\ \leq 1\ \mathrm{r}_{i+1}^-.K_i.$$

The interesting cycles for a concept C are the IFPs looping on themselves and forcing the presence of (at least) one instance of C "per instance of the cycle". This latter property can also be expressed in terms of IFPs.

Definition 2. *An IFP $K_0, \mathrm{r}_1, K_1, \ldots, \mathrm{r}_n, K_n$ is a* C-generating cycle *in $\mathcal{T}$ if $\mathcal{T} \models K_n \sqsubseteq K_0$ and there exists an IFP $L_0, \mathrm{s}_1, L_1, \ldots, \mathrm{s}_m, L_m$ such that $\mathcal{T} \models K_i \sqsubseteq L_0$ for some $0 \leq i \leq n$ and $\mathcal{T} \models L_m \sqsubseteq \mathrm{C}$.*

We now reconcile with existing cycle reversion techniques by considering a completion of the original TBox containing reversed versions of each C-generating cycle.

Definition 3. *We denote $\mathcal{T}_\mathrm{C}$ the $\mathcal{ELIF}_\bot$ TBox obtained from $\mathcal{T}$ by adding the following axioms, for each* C*-generating cycle $K_0, \mathrm{r}_1, K_1, \ldots, \mathrm{r}_n, K_n$ in $\mathcal{T}$ and each $0 \leq i < n$:*

$$K_{i+1} \sqsubseteq \exists \mathrm{r}_i^-.K_i \quad \textit{and} \quad K_i \sqsubseteq\ \leq 1\ \mathrm{r}_{i+1}.K_{i+1}$$

The key result regarding this cycle reversion technique focused on a single concept is the following lemma:

Lemma 3. *Let $(\mathcal{T},\mathcal{A})$ be an $\mathcal{ELIF}_\bot$ KB and C a concept name. There exists a model $\mathcal{I}$ of $\mathcal{K}$ s.t. $\left|\mathrm{C}^\mathcal{I}\right| < \infty$ iff the KB $(\mathcal{T}_\mathrm{C},\mathcal{A})$ is satisfiable. Furthermore, every such model $\mathcal{I}$ is a model of $(\mathcal{T}_\mathrm{C},\mathcal{A})$.*

The 'only-if' direction of the above is the easy one: the IFPs in Definition 2 enforce that for every K_i on a C-generating cycle, there is an injection from $K_i^\mathcal{I}$ to $\mathrm{C}^\mathcal{I}$. Since $\mathrm{C}^\mathcal{I}$ is finite, so are all these $K_i^\mathcal{I}$. It follows that the injective function from $K_i^\mathcal{I}$ to $K_{i+1}^\mathcal{I}$ defined by $\mathrm{r}_{i+1}^\mathcal{I}$ is actually a bijection. From there, it is readily checked that $\mathcal{I}$ is a model of $(\mathcal{T}_\mathrm{C},\mathcal{A})$ as claimed, and thus $(\mathcal{T}_\mathrm{C},\mathcal{A})$ is satisfiable.

For the 'if' direction of Lemma 3, assume $\mathcal{K}_\mathrm{C} = (\mathcal{T}_\mathrm{C},\mathcal{A})$ is satisfiable. We adapt a construction from (Ibáñez-García, Lutz, and Schneider 2014) to assemble a model $\mathcal{I}$ of $\mathcal{K}_\mathrm{C}$ (thus, of $\mathcal{K}$) in which $\mathrm{C}^\mathcal{I}$ is finite. Our construction actually takes as input any $\mathcal{ELIF}_\bot$ KB $\mathcal{K} = (\mathcal{T},\mathcal{A})$ and guarantees the above finiteness condition for all "safe" concepts of $\mathcal{T}$. A concept C is a *safe concept* of $\mathcal{T}$ if every axiom from $\mathcal{T}_\mathrm{C}$ is already entailed by $\mathcal{T}$. In particular, C is safe in $\mathcal{T}_\mathrm{C}$.

We introduce some relevant preliminaries. Let $\mathsf{N_C}(\mathcal{T})$ be the set of concept names used in $\mathcal{T}$. A type for $\mathcal{T}$ is a subset $t \subseteq \mathsf{N_C}(\mathcal{T})$ s.t. there is a model $\mathcal{I}$ of $\mathcal{T}$ and a $d \in \Delta^\mathcal{I}$ s.t. $\mathsf{tp}_\mathcal{I}(d) = t$, where $\mathsf{tp}_\mathcal{I}(d)$ is the type realized at d in $\mathcal{I}$, *i.e.*:

$$\mathsf{tp}_\mathcal{I}(d) := \{\mathrm{A} \in \mathsf{N_C}(\mathcal{T}) \mid d \in \mathrm{A}^\mathcal{I}\}$$

We use $\mathsf{TP}(\mathcal{T})$ to denote the set of all types of $\mathcal{T}$. A type is *critical* in $\mathcal{T}$ if it occurs on a C-generating cycle for some safe concept C of $\mathcal{T}$. Otherwise it is a *free* type in $\mathcal{T}$. For $t, t' \in \mathsf{TP}(\mathcal{T})$ and r a role, we write:

- $t \to_\mathrm{r} t'$ if $\mathcal{T} \models t \sqsubseteq \exists r.t'$ and t' is maximal for this property;
- $t \to_\mathrm{r}^1 t'$ if $t \to_\mathrm{r} t'$ and $\mathcal{T} \models t' \sqsubseteq\ \leq 1\ \mathrm{r}^-.t$;
- $t\ {}^1{\leftrightarrow}_\mathrm{r}^1\ t'$ if $t \to_\mathrm{r}^1 t'$ and $t' \to_{\mathrm{r}^-}^1 t$.

A type class is a non-empty set $P \subseteq \mathsf{TP}(\mathcal{T})$ such that $t \in P$ and $t\ {}^1{\leftrightarrow}_\mathrm{r}^1\ t'$ implies $t' \in P$, and P is minimal with this condition. Note that the set of all type classes is a partition of $\mathsf{TP}(\mathcal{T})$. We set $P \prec P'$ if there are $t \in P$ and $t' \in P'$ with $t' \subsetneq t$. Let $\prec^+$ be the transitive closure of $\prec$. It is known from (Ibáñez-García, Lutz, and Schneider 2014) that $\prec^+$ is a strict partial order.

The initial interpretation $\mathcal{I}_\mathcal{K}^0$ is defined by introducing an element for every ABox individual and an element d_t for each $t \in \mathsf{TP}(\mathcal{T}_f)$. Formally, we define:

$$\begin{aligned} \Delta^{\mathcal{I}_\mathcal{K}^0} &= \mathsf{Ind}(\mathcal{A}) \cup \{d_t \mid t \in \mathsf{TP}(\mathcal{T}_f)\} \\ \mathrm{A}^{\mathcal{I}_\mathcal{K}^0} &= \{\mathsf{a} \in \mathsf{Ind}(\mathcal{A}) \mid \mathrm{A} \in \mathsf{tp}_\mathcal{K}(\mathsf{a})\} \cup \{d_t \mid \mathrm{A} \in d_t\} \\ \mathrm{r}^{\mathcal{I}_\mathcal{K}^0} &= \{(\mathsf{a},\mathsf{b}) \mid \mathrm{r}(\mathsf{a},\mathsf{b}) \in \mathcal{A}\} \end{aligned}$$

where $\mathsf{tp}_\mathcal{K}(\mathsf{a}) := \{\mathrm{A} \in \mathsf{N_C} \mid \mathcal{K} \models \mathrm{A}(\mathsf{a})\}$.

We describe three completion rules $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3$ applicable to an interpretation $\mathcal{I}$. Informally, whenever an existing d

with type t' needs a r.t-successor for some t, then $\mathcal{C}_3$ connects d to d_t if the chosen witness may be used by several such elements d (that is $t' \not\rightarrow^1_{\mathrm{r}} t$). If, on the other hand, the witness cannot be reused, then $\mathcal{C}_1$ simply introduces a dedicated fresh element e if t is either free or not in the type class of t'. Otherwise t is critical and in the type class P of t'. Then $\mathcal{C}_2$ introduces or reuses existing elements to instantiate the whole type class P at once. This requires only finitely many fresh instances of each type in P, in particular, critical types in P are instantiated only finitely many times.

$\mathcal{C}_1$. For each $d \in \Delta^{\mathcal{I}}$, each $t \in \mathsf{TP}(\mathcal{T})$ and each $\mathrm{r} \in \mathsf{N}^{\pm}_{\mathsf{R}}$ such that $\mathsf{tp}_{\mathcal{I}}(d) \rightarrow^1_{\mathrm{r}} t$, $d \notin (\exists \mathrm{r}.t)^{\mathcal{I}}$, and either $t \not\rightarrow^1_{\mathrm{r}^-} \mathsf{tp}_{\mathcal{I}}(d)$ or t is a free type in $\mathcal{T}$, add a fresh domain element e and modify the interpretation of concept names such that $\mathsf{tp}_{\mathcal{I}}(e) = t$ and $(d, e) \in \mathrm{r}^{\mathcal{I}}$.

$\mathcal{C}_2$. Choose a type class P that is minimal w.r.t. the order $\prec^+$, a $\lambda = s \;{}^1\!\!\leftrightarrow^1_{\mathrm{r}}\; s'$ with $s \in P$, and an element $d \in s^{\mathcal{I}} \setminus (\exists \mathrm{r}.s')^{\mathcal{I}}$. If such a choice is not possible, then the application of $\mathcal{C}_2$ just returns the original model $\mathcal{I}$. Otherwise, for each $\lambda = s \;{}^1\!\!\leftrightarrow^1_{\mathrm{r}}\; s'$ with $s \in P$, set:

$$X^{\mathcal{I}}_{\lambda,1} = s^{\mathcal{I}} \setminus (\exists \mathrm{r}.s')^{\mathcal{I}} \qquad X^{\mathcal{I}}_{\lambda,2} = s'^{\mathcal{I}} \setminus (\exists \mathrm{r}^-.s)^{\mathcal{I}}.$$

Take (i) a fresh set Δ_s for each $s \in P$ such that $|\Delta_s| \leq \max\{|t^{\mathcal{I}}| \mid t \in P\}$ and (ii) a bijection π_λ from $X^{\mathcal{I}}_{\lambda,1} \cup \Delta_s$ to $X^{\mathcal{I}}_{\lambda,2} \cup \Delta_{s'}$ for each $\lambda = s \;{}^1\!\!\leftrightarrow^1_{\mathrm{r}}\; s'$ with $s, s' \in P$ and $\mathrm{r} \in \mathsf{N}_{\mathsf{R}}$. A concrete construction of such sets and bijections can follow the one detailed in (Ibáñez-García, Lutz, and Schneider 2014). We additionally require the above to minimize $|\biguplus_{s \in P} \Delta_s|$. Now extend $\mathcal{I}$ as follows:

- add all domain elements in $\biguplus_{s \in P} \Delta_s$;
- extend $\mathrm{r}^{\mathcal{I}}$ with π_λ, for each $\lambda = s \;{}^1\!\!\leftrightarrow^1_{\mathrm{r}}\; s'$ with $s, s' \in P$ and r a role name;
- interpret concept names so that $\mathsf{tp}_{\mathcal{I}}(d) = s$ for all $d \in \Delta_s$, $s \in P$.

$\mathcal{C}_3$. For each $d \in \Delta^{\mathcal{I}}$, each $t \in \mathsf{TP}(\mathcal{T})$ and each $\mathrm{r} \in \mathsf{N}^{\pm}_{\mathsf{R}}$ such that $\mathsf{tp}_{\mathcal{I}}(d) \rightarrow_{\mathrm{r}} t$, $\mathsf{tp}_{\mathcal{I}}(d) \not\rightarrow^1_{\mathrm{r}} t$, and $d \notin (\exists \mathrm{r}.t)^{\mathcal{I}}$, add the edge (d, d_t) to $\mathrm{r}^{\mathcal{I}}$.

We denote $\mathcal{C}_k(\mathcal{I})$ the application of $\mathcal{C}_k$ to interpretation $\mathcal{I}$. For $\mathcal{C}_2$, this is ambiguous since its application may depend on several choices (a minimal type class, etc). This does not matter for our construction and we simply assume a fixed choice. While it is easily verified that $\mathcal{C}_3$ is idempotent, that is $\mathcal{C}_3(\mathcal{C}_3(\mathcal{I})) = \mathcal{C}_3(\mathcal{I})$, it is not the case for $\mathcal{C}_1$ in general. However, since applying $\mathcal{C}_1$ on $\mathcal{I}$ does not alter the interpretation of concept and roles names on the original domain $\Delta^{\mathcal{I}}$, we can safely define $\mathcal{C}_1^\infty(\mathcal{I}) = \bigcup_{n=1}^\infty \mathcal{C}_1^n(\mathcal{I})$, where $\mathcal{C}_1^n$ denotes n successive applications of $\mathcal{C}_1$. We now view $\mathcal{C}_1^\infty$ as a completion rule, which is clearly idempotent.

Starting from the initial interpretation $\mathcal{I}^0_{\mathcal{K}}$ previously defined, we complete it as follows:

$$\mathcal{I}^{n+1}_{\mathcal{K}} = \mathcal{C}_3(\mathcal{C}_2(\mathcal{C}_1^\infty(\mathcal{I}^n_{\mathcal{K}}))) \qquad \text{and} \qquad \mathcal{I}_{\mathcal{K}} = \textstyle\bigcup_{n=0}^\infty \mathcal{I}^n_{\mathcal{K}}.$$

Here again, notice that each rule application on $\mathcal{I}$ preserves the interpretation of concept names on $\Delta^{\mathcal{I}}$ and can only extend those of role names, so $\mathcal{I}_{\mathcal{K}}$ is well-defined. In fact, we prove that $\mathcal{I}_{\mathcal{K}}$ is obtained after finitely many steps: there exists $N \in \mathbb{N}$ such that $\mathcal{I}^{N+1}_{\mathcal{K}} = \mathcal{I}^N_{\mathcal{K}}$. Crucially, this guarantees that only finitely many instances of every critical type and safe concepts are introduced. This culminates in the following lemma, which also concludes the proof of Lemma 3.

Lemma 4. *If $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ is a satisfiable $\mathcal{ELIF}_\bot$ KB, then $\mathcal{I}_{\mathcal{K}}$ is a model of $\mathcal{K}$ and $\mathrm{C}^{\mathcal{I}_{\mathcal{K}}}$ is finite for all safe C of $\mathcal{T}$.*

Now, to finish the proof of Theorem 4, we build a model with exactly one more instance of C than in the model $\mathcal{I}_{\mathcal{K}_\mathrm{C}}$ obtained by the above procedure on $\mathcal{K}_\mathrm{C} := (\mathcal{T}_\mathrm{C}, \mathcal{A})$. To do so, we essentially relaunch this procedure on a simpler KB $\mathcal{K}' = (\mathcal{T}_\mathrm{C}, \{\mathrm{C}(\mathsf{a})\})$, where a is a fresh individual name, and then form the disjoint union of $\mathcal{I}_{\mathcal{K}'}$ with $\mathcal{I}_{\mathcal{K}_\mathrm{C}}$. However, this approach is too naive as the model $\mathcal{I}_{\mathcal{K}'}$ might contain several instances of the concept C, due to the initial elements d_t for each type $t \in \mathsf{TP}(\mathcal{T})$. This cannot easily be solved by identifying the respective d_t elements from $\mathcal{I}_{\mathcal{K}'}$ and $\mathcal{I}_{\mathcal{K}_\mathrm{C}}$, as such an operation may violate some functionality constraints.

Instead, we produce an incomplete version $\mathcal{J}_{\mathcal{K}'}$ of $\mathcal{I}_{\mathcal{K}'}$ in which elements d_t are absent and applications of rule $\mathcal{C}_3$ are ignored. Formally, for an $\mathcal{ELIF}_\bot$ KB $\mathcal{K}$, the interpretation $\mathcal{J}^0_{\mathcal{K}}$ is defined as $\mathcal{I}^0_{\mathcal{K}}$, but without the d_t elements, and we further define, for all $n \geq 0$:

$$\mathcal{J}^{n+1}_{\mathcal{K}} = \mathcal{C}_2(\mathcal{C}_1^\infty(\mathcal{J}^n_{\mathcal{K}})) \qquad \text{and} \qquad \mathcal{J}_{\mathcal{K}} = \textstyle\bigcup_{n=0}^\infty \mathcal{J}^n_{\mathcal{K}}.$$

The resulting interpretation $\mathcal{J}_{\mathcal{K}}$ is in general not a model of $\mathcal{K}$ due to the non-applied $\mathcal{C}_3$ rules. It is however possible to reuse d_t elements of another model, *e.g.* those from $\mathcal{I}_{\mathcal{K}}$.

Lemma 5. *If $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ is a satisfiable $\mathcal{ELIF}_\bot$ KB and C is safe in $\mathcal{T}$, then $\mathcal{J} = \mathcal{C}_3(\mathcal{I}_{\mathcal{K}} \cup \mathcal{J}_{\mathcal{K}'})$ is a model of $\mathcal{K}$, where $\mathcal{K}' = (\mathcal{T}, \{\mathrm{C}(\mathsf{a})\})$ with $\mathsf{a} \notin \mathsf{Ind}(\mathcal{A})$. Furthermore, $\mathrm{C}^{\mathcal{J}} = \mathrm{C}^{\mathcal{I}_{\mathcal{K}}} \cup \{\mathsf{a}\}$.*

This concludes the proof of Theorem 4 as we obtain two models $\mathcal{I}_{\mathcal{K}_\mathrm{C}}$ and $\mathcal{C}_3(\mathcal{I}_{\mathcal{K}_\mathrm{C}} \cup \mathcal{J}_{(\mathcal{T}_\mathrm{C}, \{\mathrm{C}(\mathsf{a})\})})$ with respectively n and $n+1$ instances of C, both n and $n+1$ being finite.

3.3 The Case of DL-Lite$_{\mathcal{F}}$

We now build upon the above technique to obtain the following complete characterization for DL-Lite$_{\mathcal{F}}$ KBs.

Theorem 5. *A non-trivial subset of $\mathbb{N}^\infty$ is DL-Lite$_{\mathcal{F}}$-concept realizable iff it has shape $\{\infty\}$, $\{0, \infty\}$, $\{0\} \cup [\![M, \infty]\!]$ or $[\![M, \infty]\!]$ for some $M \in \mathbb{N}$.*

The 'if' direction follows notably from Examples 3 and 5. The converse is a consequence of the following lemma.

Lemma 6. *Let $\mathcal{K}$ be a DL-Lite$_{\mathcal{F}}$ KB and C a concept name. If there exists a model $\mathcal{I}$ of $\mathcal{K}$ with $1 \leq |\mathrm{C}^{\mathcal{I}}| < \infty$, then there exists a model $\mathcal{J}$ of $\mathcal{K}$ with $|\mathrm{C}^{\mathcal{J}}| = |\mathrm{C}^{\mathcal{I}}| + 1$.*

Proof sketch. As for $\mathcal{ELIF}_\bot$, we consider the interpretation $\mathcal{J}_{\mathcal{K}'}$ where $\mathcal{K}' = (\mathcal{T}_\mathrm{C}, \{\mathrm{C}(\mathsf{a})\})$. We can then connect $\mathcal{J}_{\mathcal{K}'}$ to any model $\mathcal{I}$ of $\mathcal{K}$ in which there is at least one instance of C by finding appropriate witnesses for the pending roles. This is possible as DL-Lite$_{\mathcal{F}}$ does not support qualified existential restrictions, thus imposing very little constraints on the types of the required witnesses. □

	DL-Lite$_{\text{core}}$	DL-Lite$_{\mathcal{F}}$	$\mathcal{EL}$, $\mathcal{ELI}_\bot$, $\mathcal{ELF}_\bot$	$\mathcal{ALC}$, $\mathcal{ELIF}$, $\mathcal{ELIF}_\bot$, $\mathcal{ALCI}$, $\mathcal{ALCF}^*$	$\mathcal{ALCF}$	$\mathcal{ALCIF}$
Concept	in FP	in $\mathsf{FP}^{\mathsf{NP}[\log]}$	in $\mathsf{FP}^{\mathsf{NP}[1]}$	$\mathsf{FP}^{\mathsf{NP}[\log]}$-c	$\mathsf{FP}^{\mathsf{NP}[\log]}$-c	$\mathsf{FP}^{\mathsf{NP}[\log]}$-h
Role	in FP	in $\mathsf{FP}^{\mathsf{NP}[\log]}$	$\mathsf{FP}^{\mathsf{NP}[\log]}$-c	$\mathsf{FP}^{\mathsf{NP}[\log]}$-c	$\mathsf{FP}^{\mathsf{NP}[\log]}$-h	$\mathsf{FP}^{\mathsf{NP}[\log]}$-h

Table 1: Worst-case complexity of $\mathsf{Spectrum}(q, \mathcal{T})$ if q is a concept (resp. role) cardinality query and $\mathcal{T}$ is expressed in one of the DLs on the first row. -h stands for -hard and -c for -complete.

4 Complexity of Computing Spectra

We now tackle the problem of computing the proposed effective representation of spectra, helped by our knowledge of their possible shapes. We focus on data complexity: for a fixed cardinality query q and a fixed TBox $\mathcal{T}$, we study the complexity of the $\mathsf{Spectrum}(q, \mathcal{T})$ problem, which, given an ABox $\mathcal{A}$ as input, computes the output of Problem 1 from Section 2.

We use functional complexity classes: FP is the class of functions computable in polynomial time by a Turing machine; $\mathsf{FP}^{\mathsf{NP}[1]}$ is FP with $\mathcal{O}(1)$ many queries to an NP oracle; and $\mathsf{FP}^{\mathsf{NP}[\log]}$ is allowed $\mathcal{O}(\log(n))$ many queries to NP where n is the size of the input and p a polynomial. We refer to (Krentel 1988; Jenner and Torán 1995) for details and recall the following inclusions: $\mathsf{FP} \subseteq \mathsf{FP}^{\mathsf{NP}[1]} \subseteq \mathsf{FP}^{\mathsf{NP}[\log]}$.

Our complexity results are summarized in Table 1. We start with upper bounds.

Theorem 6. $\mathsf{Spectrum}(q_{\mathrm{C}}, \mathcal{T})$ *is in:*

- $\mathsf{FP}^{\mathsf{NP}[\log]}$ *if* $\mathcal{T}$ *is in* $\mathcal{ELIF}_\bot$, $\mathcal{ALCI}$ *or* $\mathcal{ALCF}$.
- $\mathsf{FP}^{\mathsf{NP}[1]}$ *if* $\mathcal{T}$ *is in* $\mathcal{ELI}_\bot$ *or* $\mathcal{ELF}_\bot$.
- FP *if* $\mathcal{T}$ *is in DL-Lite*$_{\text{core}}$.

The backbone of the above complexity results relies on the fact that, in all the concerned cases, the possible spectra are of the form $S \cup [\![M, \infty]\!]$ where $S \subseteq [\![0, M[\![$ and M is either ∞ or M is polynomial w.r.t. data complexity.

Lemma 7. *Let* $\mathcal{T}$ *be an* $\mathcal{ELIF}_\bot$*-,* $\mathcal{ALCI}$*-, or* $\mathcal{ALCF}$*-TBox and* C *a concept name. There exists a polynomial* $p(x)$*, with coefficients computable from* $\mathcal{T}$ *and* C*, such that* $p(x) \geq 1$ *for every* $x \geq 0$ *and, for every KB* $\mathcal{K} = (\mathcal{T}, \mathcal{A})$*, either*

- $p(|\mathcal{A}|) \notin \mathsf{Sp}_{\mathcal{K}}(q_{\mathrm{C}})$ *and* $\mathsf{Sp}_{\mathcal{K}}(q_{\mathrm{C}}) \subseteq \{0, \infty\}$*; or*
- $p(|\mathcal{A}|) \in \mathsf{Sp}_{\mathcal{K}}(q_{\mathrm{C}})$ *and* $\mathsf{Sp}_{\mathcal{K}}(q_{\mathrm{C}}) = S \cup [\![M, \infty]\!]$ *where* $S \subseteq [\![0, M]\!]$ *and* $M < p(|\mathcal{A}|)$.

Moreover, in the latter case, $M \leq \min(\mathsf{Sp}_{\mathcal{K}}(q_{\mathrm{C}})) + p(0)$.

This allows us to present a simple and uniform description of an algorithm computing spectra. Let $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ be the input KB, q the input cardinality query, and $p(x)$ the polynomial from Lemma 7.

First the algorithm tests whether $p(|\mathcal{A}|) \in \mathsf{Sp}_{\mathcal{K}}(q_{\mathrm{C}})$. If not then we are in the first case of Lemma 7, that is $\mathsf{Sp}_{\mathcal{K}}(q_{\mathrm{C}}) \subseteq \{0, \infty\}$. The algorithm now tests whether $\mathcal{K}$ is satisfiable: if not, it returns $\emptyset$. Then, it checks whether q_{C} is satisfiable with respect to $\mathcal{K}$: if not, it returns $\{0\}$ (represented in the sense of Problem 1 by the triple $(\emptyset, 0, 0)$). Finally, the algorithm checks whether $0 \in \mathsf{Sp}_{\mathcal{K}}(q_{\mathrm{C}})$. If yes, it returns $\{0, \infty\}$ and $\{\infty\}$ otherwise (respectively represented in the sense of Problem 1 by triples $(\{0\}, \infty, 0)$ and $(\emptyset, \infty, 0)$).

If $p(|\mathcal{A}|) \in \mathsf{Sp}_{\mathcal{K}}(q_{\mathrm{C}})$ then we are in the second case of Lemma 7, that is $M < p(|\mathcal{A}|)$ and $S \subseteq [\![0, M]\!]$. The algorithm first performs a binary search on the interval $[\![0, p(|\mathcal{A}|)]\!]$ to find $\min(\mathsf{Sp}_{\mathcal{K}}(q_{\mathrm{C}}))$. In each step of the search, the algorithm is given a number $n \in [\![0, p(|\mathcal{A}|)]\!]$ and performs a minimality test that verifies whether $n > \min(\mathsf{Sp}_{\mathcal{K}}(q_{\mathrm{C}}))$. Note that once the value of $\min(\mathsf{Sp}_{\mathcal{K}}(q_{\mathrm{C}}))$ is found, the additional remark in Lemma 7 guarantees that $M \leq \min(\mathsf{Sp}_{\mathcal{K}}(q_{\mathrm{C}})) + p(0)$ and thus $S \subseteq [\![\min(\mathsf{Sp}_{\mathcal{K}}(q_{\mathrm{C}})), \min(\mathsf{Sp}_{\mathcal{K}}(q_{\mathrm{C}}))+p(0)]\!]$. The algorithm performs membership tests on this latter interval to compute M and S. Finally, the algorithm returns the set $S \cup [\![M, \infty]\!]$ (represented in the sense of Problem 1 by the triple $(S, M, 1)$).

The correctness of the algorithm follows directly from Lemma 7. The exact computational complexity depends on the number and cost of satisfiability checks, membership tests, and minimality tests. For instance, for $\mathcal{ELIF}_\bot$ the two initial satisfiability checks can by performed by NP oracles (Glimm et al. 2008). Similarly, the membership tests can be resolved by an NP oracle as they can be seen as instances of *closed predicates* problem, see (Lukumbuzya and Šimkus 2021). The minimality tests can be performed by an NP oracle that guesses $n' < n$ and performs a membership test. Since the algorithm uses logarithmically many minimality tests to compute $\min(\mathsf{Sp}_{\mathcal{K}}(q_{\mathrm{C}}))$ and no more than $p(0)+1$ membership tests to compute S, the desired upper bound holds.

In DL-Lite$_{\text{core}}$, we can perform the satisfiability checks, the membership tests, and the minimality tests in polynomial time, see (Calvanese et al. 2006) and (Manière 2022) [Theorem 51] respectively, resulting in overall polynomial running time.

The following theorem provides two lower bounds, notably establishing $\mathsf{FP}^{\mathsf{NP}[\log]}$-completeness in several cases.

Theorem 7. *There exists an* $\mathcal{ELIF}$ *(resp.* $\mathcal{ALC}$*) TBox* $\mathcal{T}$ *such that* $\mathsf{Spectrum}(q_{\mathrm{C}}, \mathcal{T})$ *is* $\mathsf{FP}^{\mathsf{NP}[\log]}$*-hard.*

We reduce from the problem of computing the maximal size of an independent set in a graph, known to be $\mathsf{FP}^{\mathsf{NP}[\log]}$-hard (Krentel 1988). We briefly sketch the proof for $\mathcal{ALC}$: consider the TBox $\mathcal{T} := \{\neg\mathrm{C} \sqsubseteq \forall \mathrm{r}.\mathrm{C}\}$. Given a graph $G = \langle V, E \rangle$, we construct an ABox $\mathcal{A}$ consisting in $\mathrm{r}(u, v)$ for every $\{u, v\} \in E$. Intuitively, $\neg\mathrm{C}$ describes an independent set and we prove that k is the maximal size of an independent set in G iff $\mathsf{Sp}_{(\mathcal{T},\mathcal{A})}(q_{\mathrm{C}}) = [\![|V| - k, \infty]\!]$, that is the triple $(\emptyset, |V| - k, 1)$ is our representation of $\mathsf{Sp}_{(\mathcal{T},\mathcal{A})}(q_{\mathrm{C}})$.

5 The Case of Role Cardinality Queries

In this section, we briefly mention the similar results we obtain regarding role cardinality queries, *i.e.* CCQs with form $q_{\mathrm{r}} := \exists z_1 \exists z_2\ \mathrm{r}(z_1, z_2)$, where $\mathrm{r} \in \mathsf{N_R}$ is a role name and z_1, z_2 are counting variables. Computing the spectrum of q_{r} on a KB $\mathcal{K}$ thus corresponds to deciding the possible values of $\left|\mathrm{r}^{\mathcal{I}}\right|$ across models $\mathcal{I}$ of $\mathcal{K}$. Every such query satisfies preconditions of Lemma 2 and thus its spectrum can be represented as in Problem 1. We say a set V is *$\mathcal{L}$-role realizable* if there is a role r and a $\mathcal{L}$ KB $\mathcal{K}$ s.t. $\mathsf{Sp}_{\mathcal{K}}(q_{\mathrm{r}}) = V$.

We first highlight that, if a DL of interest can express that a role is functional, then there is a strong connection between role-realizable and concept-realizable sets. Indeed, if $\mathcal{T} \models \top \sqsubseteq\, \leq 1\ \mathrm{r}.\top$, then in every model $\mathcal{I}$ of $\mathcal{T}$, we have $\left|\mathrm{r}^{\mathcal{I}}\right| = \left|(\exists \mathrm{r})^{\mathcal{I}}\right|$. The following is an immediate consequence.

Lemma 8. *Let $\mathcal{L}$ be a fragment of $\mathcal{ALCIF}$ and $V \subseteq \mathbb{N}^\infty$. If V is $\mathcal{L}$-concept realizable and axioms $\top \sqsubseteq\, \leq 1\ \mathrm{r}.\top$, $\exists \mathrm{r} \sqsubseteq \mathrm{C}$ and $\mathrm{C} \sqsubseteq \exists \mathrm{r}$ are permitted in $\mathcal{L}$, then V is $\mathcal{L}$-role realizable.*

The above notably applies to DL-Lite$_{\mathcal{F}}$ and $\mathcal{ALCF}$ KBs. For $\mathcal{ALCIF}$ KBs and joint with Theorem 1, we obtain:

Corollary 1. *A non-trivial subset of $\mathbb{N}^\infty$ is $\mathcal{ALCIF}$-role realizable iff it is a subsemigroup of $\mathbb{N}^\infty$ containing ∞.*

In the case of concept names, we established identical results for $\mathcal{ALCI}$ and $\mathcal{ALCF}$ KBs (Theorem 2). This does not hold with a role name. In fact, we prove that $\mathcal{ALCF}$ KBs realize the same sets as $\mathcal{ALCIF}$, except for $\{\infty\}$ and $\{0, \infty\}$.

Theorem 8. *A non-trivial subset of $\mathbb{N}^\infty$ is $\mathcal{ALCF}$-role realizable iff it is a subsemigroup of $\mathbb{N}^\infty$ containing ∞ and at least a non-zero natural.*

To establish the above, we strongly rely on qualified functional dependencies, *i.e.* axioms $\mathrm{B}_1 \sqsubseteq\, \leq 1\ \mathrm{r}.\mathrm{B}_2$ where $\mathrm{B}_1, \mathrm{B}_2$ are not just $\top$. As $\mathcal{ALCIF}$ is sometimes defined to only support unqualified functionality, *i.e.* only $\mathrm{B}_1 = \mathrm{B}_2 = \top$, we also treat this fragment, here denoted $\mathcal{ALCIF}^*$.

Theorem 9. *If a non-trivial subset of $\mathbb{N}^\infty$ is $\mathcal{ALCI}$- (resp. $\mathcal{ALCF}^*$-) role realizable, then it has shape $S \cup [\![M, \infty]\!]$ for some $M \in \mathbb{N}$ and $S \subseteq [\![0, M]\!]$. In the case of $\mathcal{ELIF}_\perp$, shapes $\{\infty\}$ and $\{0, \infty\}$ are also permitted.*

A key ingredient is to reuse techniques from Section 3 on concept cardinality queries, notably for $\mathcal{ELIF}_\perp$ KBs, is the observation, simply following from the semantics of concepts, that in every interpretation $\mathcal{I}$, we have:

$$\max\left(\left|(\exists \mathrm{r})^{\mathcal{I}}\right|, \left|(\exists \mathrm{r}^-)^{\mathcal{I}}\right|\right) \leq \left|\mathrm{r}^{\mathcal{I}}\right| \leq \left|(\exists \mathrm{r})^{\mathcal{I}}\right| \cdot \left|(\exists \mathrm{r}^-)^{\mathcal{I}}\right|.$$

Moreover, note that Theorem 9 is not a complete characterization. Indeed, as for concept cardinality queries on $\mathcal{ELIF}_\perp$ KBs (see Example 4), we exhibit a simple setting in which the spectrum contains a non-trivial part S, already for an $\mathcal{EL}_\perp$ TBox and the empty ABox.

Example 6. $\mathcal{T} = \{\top \sqsubseteq \exists \mathrm{r}.\mathrm{A}_1, \top \sqsubseteq \exists \mathrm{r}.\mathrm{A}_2, \mathrm{A}_1 \sqcap \mathrm{A}_2 \sqsubseteq \perp\}$ *is an $\mathcal{EL}_\perp$ TBox and we have $\mathsf{Sp}_{(\mathcal{T},\emptyset)}(q_{\mathrm{r}}) = \{4\} \cup [\![6, \infty]\!]$.*

In the remaining fragments of $\mathcal{ALCIF}$, our results mirror those of the concept cardinality case, as summarized by the following two theorems echoing Theorems 3 and 5.

Theorem 10. *A non-trivial subset of $\mathbb{N}^\infty$ is $\mathcal{ELI}$- (resp. $\mathcal{ELF}$-) role realizable iff it has shape $[\![M, \infty]\!]$ for some $M \in \mathbb{N}$. For $\mathcal{ELIF}$, the shape $\{\infty\}$ is also permitted.*

Theorem 11. *A non-trivial subset of $\mathbb{N}^\infty$ is DL-Lite$_{\mathcal{F}}$-role realizable iff it has shape $\{\infty\}$, $\{0, \infty\}$, $\{0\} \cup [\![M, \infty]\!]$ or $[\![M, \infty]\!]$ for some $M \in \mathbb{N}$. The same holds for DL-Lite$_{\mathsf{core}}$ but without shapes $\{\infty\}$ and $\{0, \infty\}$.*

Based on these results, we classify the complexity of computing the proposed representation. Our complexity results also appear in Table 1. For the upper bounds, we follow the same approach as presented in Section 4 for concept cardinality queries, and obtain the following:

Theorem 12. $\mathsf{Spectrum}(q_{\mathrm{r}}, \mathcal{T})$ *is in:*

- $\mathsf{FP}^{\mathsf{NP}[\log]}$ *if $\mathcal{T}$ is in $\mathcal{ELIF}_\perp$, $\mathcal{ALCI}$ or $\mathcal{ALCF}^*$.*
- FP *if $\mathcal{T}$ is in DL-Lite$_{\mathsf{core}}$.*

The following theorem provides a lower bound that applies already for $\mathcal{EL}$ KBs.

Theorem 13. *There exists an $\mathcal{EL}$ TBox $\mathcal{T}$ such that* $\mathsf{Spectrum}(q_{\mathrm{r}}, \mathcal{T})$ *is* $\mathsf{FP}^{\mathsf{NP}[\log]}$*-hard.*

6 Conclusion

We have characterized almost exhaustively the possible shapes of spectra for cardinality queries and proved that, in many settings, computing the proposed effective representation is $\mathsf{FP}^{\mathsf{NP}[\log]}$-complete w.r.t. data complexity. Whether an effective representation for the spectrum of a cardinality query over an $\mathcal{ALCIF}$ KB can be computed remains an open question, despite our work fully characterizing its possible shapes. For DL-Lite$_{\mathcal{F}}$ KBs, we conjecture FP membership as the use of NP oracles may not be necessary and might be replaced by direct checks in P as those employed for DL-Lite$_{\mathsf{core}}$ KBs, here used in a black-box manner.

Departing from data complexity, it is readily verified that the algorithm proposed in Section 4 provides a uniform procedure to compute the representation of a spectrum from an input KB and cardinality query, in all cases covered by Theorem 6 (resp. by Theorem 12 for role cardinality queries). This notably relies on the polynomial provided by Lemma 7 being computable given a TBox and a query. Hence, a careful inspection of the proof of the correctness of the algorithm could provide upper bounds for the combined complexity of spectrum computation. On the other hand, we believe that further obtaining meaningful lower bounds for such high-complexity functional classes is challenging.

We also emphasize that our investigation covers the 'standard' meaning of the spectrum for a logical formula, being the possible sizes of its models: it suffices to set $\mathrm{C} = \top$ in our results for concept cardinality queries.

We believe that it could be interesting to study the impact of our results on the closely related problem of answering (Boolean atomic) queries under the bag semantics. While the semantics adopted in the present paper does not coincide with bag semantics, as discussed for example in (Nikolaou et al. 2019; Calvanese et al. 2020), considerations regarding the spectra and some of the corresponding techniques might be adapted to this setting.

Acknowledgments

The authors acknowledge the financial support by the Federal Ministry of Education and Research of Germany and by the Sächsische Staatsministerium für Wissenschaft Kultur und Tourismus in the program Center of Excellence for AI-research “Center for Scalable Data Analytics and Artificial Intelligence Dresden/Leipzig”, project identification number: ScaDS.AI

Second author was supported by the DFG project LU1417/3-1 QTEC.

References

Artale, A.; Calvanese, D.; Kontchakov, R.; and Zakharyaschev, M. 2009. The DL-Lite Family and Relations. *Journal of Artificial Intelligence Research (JAIR)*, 36(1): 1–69.

Baader, F.; Bednarczyk, B.; and Rudolph, S. 2020. Satisfiability and Query Answering in Description Logics with Global and Local Cardinality Constraints. In *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI 2020)*, 616–623.

Baader, F.; Horrocks, I.; Lutz, C.; and Sattler, U. 2017. *An Introduction to Description Logic*. Cambridge University Press.

Bienvenu, M.; Manière, Q.; and Thomazo, M. 2020. Answering Counting Queries over DL-Lite Ontologies. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, 1608–1614.

Bienvenu, M.; Manière, Q.; and Thomazo, M. 2021. Cardinality Queries over DL-Lite Ontologies. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 1801–1807.

Bienvenu, M.; Manière, Q.; and Thomazo, M. 2022. Counting Queries over $\mathcal{ALCHI}$ Ontologies. In *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, 53–62.

Calvanese, D.; Corman, J.; Lanti, D.; and Razniewski, S. 2020. Counting Query Answers over a DL-Lite Knowledge Base. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, 1658–1666.

Calvanese, D.; Giacomo, G. D.; Lembo, D.; Lenzerini, M.; and Rosati, R. 2006. Data Complexity of Query Answering in Description Logics. In *Proceedings of the 10th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, 260–270.

Calvanese, D.; Kharlamov, E.; Nutt, W.; and Thorne, C. 2008. Aggregate queries over ontologies. In *Proceedings of the 2nd International Workshop on Ontologies and Information Systems for the Semantic Web (ONISW)*, 97–104.

Cosmadakis, S. S.; Kanellakis, P. C.; and Vardi, M. Y. 1990. Polynomial-Time Implication Problems for Unary Inclusion Dependencies. *Journal of the ACM*, 37(1): 15–46.

Durand, A.; Fagin, R.; and Loescher, B. 1997. Spectra with Only Unary Function Symbols. In *Proceedings of the 11th International Workshop on Computer Science Logic (CSL)*, 189–202.

Fagin, R. 1974. Generalized first-order spectra and polynomial-time recognizable sets. *Complexity of computation*, 7: 43–73.

Feier, C.; Lutz, C.; and Przybylko, M. 2021. Answer Counting under Guarded TGDs. In *Proceedings of the 24th International Conference on Database Theory (ICDT)*, 11:1–11:22.

Glimm, B.; Horrocks, I.; Lutz, C.; and Sattler, U. 2008. Conjunctive Query Answering for the Description Logic $\mathcal{SHIQ}$. *Journal of Artificial Intelligence Research (JAIR)*, 31(1): 157–204.

Grillet, P. A. 2001. *Commutative Semigroups*. Springer New York, NY. ISBN 978-0-7923-7067-3.

Ibáñez-García, Y. A.; Lutz, C.; and Schneider, T. 2014. Finite Model Reasoning in Horn Description Logics. In *Proceedings of the 14th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, 490–509.

Jenner, B.; and Torán, J. 1995. Computing Functions with Parallel Queries to NP. *Journal of Theoretical Computer Science (TCS)*, 141(1&2): 175–193.

Kostylev, E. V.; and Reutter, J. L. 2015. Complexity of answering counting aggregate queries over DL-Lite. *Journal of Web Semantics (JWS)*, 33: 94–111.

Krentel, M. W. 1988. The Complexity of Optimization Problems. *Journal of Computer and System Sciences*, 36(3): 490–509.

Lukumbuzya, S.; and Šimkus, M. 2021. Bounded Predicates in Description Logics with Counting. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 1966–1972.

Lutz, C. 2008. The Complexity of Conjunctive Query Answering in Expressive Description Logics. In *Proceedings of the 4th International Joint Conference on Automated Reasoning (IJCAR)*, 179–193.

Manière, Q. 2022. *Counting queries in ontology-based data access*. Ph.D. thesis, Université de Bordeaux.

Manière, Q.; and Przybyłko, M. 2024. Spectra of Cardinality Queries over Description Logic Knowledge Bases. arXiv:2412.12929.

Nikolaou, C.; Kostylev, E. V.; Konstantinidis, G.; Kaminski, M.; Cuenca Grau, B.; and Horrocks, I. 2019. Foundations of ontology-based data access under bag semantics. *Journal of Artificial Intelligence (AIJ)*, 91–132.

Poggi, A.; Lembo, D.; Calvanese, D.; De Giacomo, G.; Lenzerini, M.; and Rosati, R. 2008. Linking Data to Ontologies. *Journal on Data Semantics (JoDS)*, 10: 133–173.

Rosati, R. 2008. Finite model reasoning in DL-Lite. In *Proceedings of the 5th European Semantic Web Conference (ESWC)*, 215–229.

Xiao, G.; Calvanese, D.; Kontchakov, R.; Lembo, D.; Poggi, A.; Rosati, R.; and Zakharyaschev, M. 2018. Ontology-based data access: a survey. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 5511–5519.