

Finite Entailment of UCRPQs over \mathcal{ALC} Ontologies

Víctor Gutiérrez-Basulto¹, Albert Gutowski², Yazmín Ibáñez-García¹, Filip Murlak²

¹Cardiff University, UK

²University of Warsaw, Poland

{a.gutowski, f.murlak}@mimuw.edu.pl, {gutierrezbasultov, ibanezgarcia}@cardiff.ac.uk

Abstract

We investigate the problem of *finite* entailment of ontology-mediated queries. We consider the expressive query language, unions of conjunctive regular path queries (UCRPQs), extending the well-known class of union of conjunctive queries, with regular expressions over roles. We look at ontologies formulated using the description logic \mathcal{ALC} , and show a tight 2EXPTIME upper bound for entailment of UCRPQs. At the core of our decision procedure, there is a novel automata-based technique introducing a stratification of interpretations induced by the deterministic finite automaton underlying the input UCRPQ.

1 Introduction

At the intersection of knowledge representation and database theory lies the fundamental problem of *ontology-mediated query entailment (OMQE)*, where the background knowledge provided by an ontology is used to enrich the answers to queries posed to databases. In this context, description logics (DLs) are a widely accepted family of logics used to formulate ontologies. By now, the OMQE problem under the unrestricted semantics (reasoning over arbitrary models) is well understood for various query languages and DLs (Schneider and Simkus 2020). In contrast, for the *finite* OMQE problem, where one is interested in reasoning over finite models only, the overall landscape is rather incomplete. However, in recent years, the study of finite OMQE has been gaining traction, considering both lightweight and expressive DLs and (mostly) unions of conjunctive queries (Rosati 2008; Ibáñez-García, Lutz, and Schneider 2014; Rudolph 2016; Gogacz, Ibáñez-García, and Murlak 2018; Gogacz et al. 2019; Danielski and Kieronski 2019; Gogacz et al. 2020; Bednarczyk and Kieronski 2022).

In this paper we consider the problem of finite OMQE with unions of conjunctive regular path queries (UCRPQs) as the query language. UCRPQs (Florescu, Levy, and Suciu 1998; Calvanese et al. 2000) are a powerful navigational query language for graph databases in which one can express that two entities are related by a path of edges that can be specified by a regular language over binary relations. So, UCRPQs extend unions of conjunctive queries (UCQs) with atoms that might contain regular expressions that traverse the edges of the database. Indeed, path navigation is included in the query language XPath 2.0 for XML data, and

it is also present in the SPARQL 1.1 query language for RDF data through the property path feature. Given the resemblance of instance data stored in ABoxes in DLs to graph-like data, several investigations on unrestricted entailment of various types of navigational query languages mediated by DL ontologies have been carried out (Stefanoni et al. 2014; Calvanese, Eiter, and Ortiz 2014; Bienvenu, Ortiz, and Simkus 2015; Gutiérrez-Basulto, Ibáñez-García, and Jung 2018; Gogacz et al. 2019; Bednarczyk and Rudolph 2019), yielding algorithmic approaches and optimal complexity bounds. For finite entailment of regular path queries mediated by DL ontologies, there are only undecidability results available (Rudolph 2016). The most relevant positive news are the decidability and computational complexity results by Danielski and Kieronski (2019) and Gogacz et al. (2020) on finite entailment of conjunctive queries with transitive closure over roles mediated by expressive DL ontologies.

We focus on ontologies formulated using the description logic \mathcal{ALC} . Note that entailment of UCRPQs over \mathcal{ALC} ontologies is not *finitely controllable*, i.e. finite and unrestricted entailment do not coincide as it is *not* the case that for any \mathcal{ALC} knowledge base \mathcal{K} and any UCRPQ φ , it holds that \mathcal{K} entails φ over all (unrestricted) models iff \mathcal{K} entails φ over all finite models. By assuming that the represented world is finite, we can therefore not reuse existing complexity bounds or algorithmic approaches to UCRPQ entailment. From a usability perspective, the suitability of this assumption depends on the potential applications. A particular interest for navigational queries comes from bioinformatics and cheminformatics (Lysenko et al. 2016; Cook et al. 2016; Galgonek et al. 2016; Hu, Qiu, and Dumontier 2015; Rajabi and Sanchez-Alonso 2021; Chen et al. 2020). For instance, experts often need to find associations between entities in protein, cellular, drug, and disease networks (represented as graph databases), so that e.g. gene-disease-drug associations (corresponding to paths in the database) can be discovered for developing new treatment methods. In this type of applications, databases and the models they represent are clearly meant to be finite. Importantly, biochemical networks contain complex motifs involving e.g. *cycles* or *cliques*. This type of patterns can be described using UCRPQs, however, without the finiteness assumption these patterns could be disregarded as the associated query might not be entailed when reasoning over all models (including infinite ones).

Contribution

The main technical contribution of our investigation is the development of a dedicated automata-based method for entailment of UCRPQs over \mathcal{ALC} ontologies, providing an optimal upper bound. More precisely, we obtain the following result, where the matching lower bound is inherited from (Ortiz and Simkus 2014).

Theorem 1. *Finite entailment of UCRPQs over \mathcal{ALC} ontologies is 2EXPTIME-complete.*

In prior work, Rudolph (2016) showed that finite entailment of 2RPQs in \mathcal{ALCIOF} is undecidable. Theorem 1 thus provides a key step towards delimiting the decidability boundary of finite OMQE with navigational queries.

At the heart of our approach to finite entailment of UCRPQs in \mathcal{ALC} there is a stratification of interpretations induced by the deterministic finite automaton underlying the UCRPQ. This stratification builds upon the so-called *tape construction*, previously used to efficiently evaluate queries in the extension of XPath 1.0 where arbitrary regular expressions may appear as path expressions (Bojańczyk and Parys 2011). To realize the tape construction, our method represents UCRPQs by means of a semiautomaton \mathcal{B} (Ginzburg 1968) and defines an expansion of \mathcal{B} , allowing to trace runs of \mathcal{B} that begin in all possible states, on all infixes of the input word. We make interpretations \mathcal{I} knowledgeable of the expansion by enriching paths of \mathcal{I} with its possible runs and by associating edges of \mathcal{I} with levels ℓ induced by the transitions of the expansion. In a similar fashion we also make CRPQs sensible of levels. With this at hand, we tackle finite entailment by eliminating the lowest level from a query and from an interpretation, and then recursively solving the simpler problem. At each step of this process, we should be able to arrange solutions to simpler problems in a hierarchical way so that we can reason over them. To this aim, we consider a variant of entailment that includes an *environment*, which will provide the necessary information to position the arranged solutions to simpler problems in the context of larger interpretations. To better keep track of the complexity of our recursive method, we introduce a modification of the entailment problem modulo environment in which we look at a particular type of finite models: (ℓ, ℓ') -models, which are models with edges of levels ℓ or higher that are ‘consistent’ w.r.t. queries referring to edges of level ℓ' or higher. We solve the problem of finding (ℓ, ℓ') -models recursively by increasing ℓ and ℓ' in an alternating way, until both reach the maximum level $n + 1$, with n the number of states of \mathcal{B} . This will mean solving finite entailment modulo environment, and thus standard finite entailment as well.

Missing proofs can be found in the technical report available at <https://arxiv.org/pdf/2204.14261.pdf>.

Related Work

We next discuss some existing work relevant to our study.

OMQE of Navigational Queries. As previously discussed, there exist various works on unrestricted entailment of navigational query languages mediated by DL ontologies. Most of them concentrate on extensions of regular path queries

(RPQs), such as UCRPQs, and consider both Horn (Bienvenu, Ortiz, and Simkus 2015) and expressive DLs (Calvanese, Eiter, and Ortiz 2014; Gutiérrez-Basulto, Ibáñez-García, and Jung 2018; Gogacz et al. 2019; Bednarczyk and Rudolph 2019). There have been also some studies on entailment of graph XPath queries (Stefanoni et al. 2014; Bienvenu et al. 2014; Kostylev, Reutter, and Vrgoc 2014).

Finite OMQE. There exist various decidability results and optimal complexity bounds for finite entailment of union of conjunctive queries in Horn DLs (Rosati 2008; Ibáñez-García, Lutz, and Schneider 2014) and in expressive DLs from the \mathcal{S} family (Gogacz, Ibáñez-García, and Murlak 2018; Gogacz et al. 2019; Danielski and Kieronski 2019). In most cases, the computational complexity coincides with that of the unrestricted case, but the algorithmic approaches are completely different. On the negative side, undecidability of finite entailment of UCQs in the more expressive DL \mathcal{SHOIF} was shown by (Rudolph 2016), as well as the undecidability result for finite entailment of 2RPQs in \mathcal{ALCIOF} . Closer to our work are the positive results on finite entailment of UCQs with transitive closure over roles in expressive DLs allowing for transitivity or transitive closure over roles (Danielski and Kieronski 2019; Gogacz et al. 2020). These results close the distance to the undecidability frontier for finite entailment from a different angle by considering ontology languages more expressive than \mathcal{ALC} , but a subclass of UCRPQs as query language. In the context of database theory research, finite OMQE (also called open-world query entailment) has also been investigated; for instance, Amarilli and Benedikt (2020) study finite OMQE for inclusion dependencies and functional dependencies over relations of arbitrary arity, and Pratt-Hartmann (2009) looks at finite OMQE in the two-variable fragment of FOL with counting quantifiers.

Finite Controllability. There have been also a few works on finite controllability in the context of DLs. For instance, Bednarczyk and Kieroński (2022) recently showed that the \mathcal{ZOI} and \mathcal{ZOO} members of the \mathcal{Z} family are finitely controllable for UCQs. Beyond DLs, there have been several works on UCQ-finite controllability: for the guarded fragment of FOL (Bárány, Gottlob, and Otto 2014) or for various fragments of existential rules (Civili and Rosati 2012; Gogacz and Marcinkowski 2013; Baget et al. 2011; Amendola, Leone, and Manna 2018; Gottlob, Manna, and Pieris 2018). Closer to our study, is the work by Figueira, Figueira, and Baque (2020) on the classification of finitely and non-finitely controllable subclasses of CRPQs over ontologies formulated in the guarded-negation fragment of FOL or in the frontier fragment of existential rules. However, no complexity results or algorithms for finite entailment are provided for the non-finitely controllable cases.

2 Preliminaries

2.1 Description Logics

We consider a vocabulary consisting of countably infinite disjoint sets of *concept names* N_C , *role names* N_R , and *individual names* N_I . \mathcal{ALC} -concepts C, D are defined by the

grammar

$$C, D ::= A \mid \neg C \mid C \sqcap D \mid \exists r.C$$

where $A \in \mathbf{N}_C$ and $r \in \mathbf{N}_R$. We use standard abbreviations \perp , \top , $C \sqcup D$ and $\forall r.C$.

An \mathcal{ALC} -TBox \mathcal{T} is a finite set of *concept inclusions* (CIs) $C \sqsubseteq D$, where C, D are \mathcal{ALC} -concepts. An ABox \mathcal{A} is a finite non-empty set of *concept and role assertions* of the form $A(a)$, $r(a, b)$, where $A \in \mathbf{N}_C$, $r \in \mathbf{N}_R$ and $\{a, b\} \subseteq \mathbf{N}_I$. We write $\text{ind}(\mathcal{A})$ for the *set of individual names* occurring in \mathcal{A} . A *knowledge base* (KB) is a pair $\mathcal{K} = (\mathcal{T}, \mathcal{A})$. We write $\text{CN}(\mathcal{K})$ and $\text{rol}(\mathcal{K})$ for the *sets of all concept and role names* occurring in \mathcal{K} . We let $\|\mathcal{K}\|$ be the total size of the representation of \mathcal{K} .

Without loss of generality, we assume throughout the paper that all CIs are in one of the following *normal forms*:

$$\bigcap_i A_i \sqsubseteq \bigcup_j B_j, \quad A \sqsubseteq \exists r.B, \quad A \sqsubseteq \forall r.B,$$

where $A, A_i, B, B_j \in \mathbf{N}_C$, $r \in \mathbf{N}_R$, and empty disjunction and conjunction are equivalent to \perp and \top , respectively. Additionally, for each $A \in \text{CN}(\mathcal{K})$ there is a complementary $\bar{A} \in \text{CN}(\mathcal{K})$ axiomatized with $\top \sqsubseteq A \sqcup \bar{A}$ and $A \sqcap \bar{A} \sqsubseteq \perp$.

2.2 Interpretations

The semantics is given as usual via *interpretations* $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ consisting of a non-empty *domain* $\Delta^{\mathcal{I}}$ and an *interpretation function* $\cdot^{\mathcal{I}}$ mapping concept names to subsets of the domain and role names to binary relations over the domain, and individual names to elements of the domain. The interpretation of complex concepts C is defined in the usual way (Baader et al. 2017). An interpretation \mathcal{I} is a *model of a TBox* \mathcal{T} , written $\mathcal{I} \models \mathcal{T}$ if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ for all CIs $C \sqsubseteq D \in \mathcal{T}$. It is a *model of an ABox* \mathcal{A} , written $\mathcal{I} \models \mathcal{A}$, if $\text{ind}(\mathcal{A}) \subseteq \Delta^{\mathcal{I}}$, $a^{\mathcal{I}} = a$ for each $a \in \text{ind}(\mathcal{A})$, $(a, b) \in r^{\mathcal{I}}$ for all $r(a, b) \in \mathcal{A}$, and $a \in A^{\mathcal{I}}$ for all $A(a) \in \mathcal{A}$. The first two conditions constitute the so-called *standard name assumption*. Finally, \mathcal{I} is a *model of a KB* $\mathcal{K} = (\mathcal{T}, \mathcal{A})$, written $\mathcal{I} \models \mathcal{K}$, if $\mathcal{I} \models \mathcal{T}$ and $\mathcal{I} \models \mathcal{A}$.

An interpretation \mathcal{I} is *finite* if $\Delta^{\mathcal{I}}$ is finite. An interpretation \mathcal{I}' is a *sub-interpretation* of \mathcal{I} , written as $\mathcal{I}' \subseteq \mathcal{I}$, if $\Delta^{\mathcal{I}'} \subseteq \Delta^{\mathcal{I}}$, $A^{\mathcal{I}'} \subseteq A^{\mathcal{I}}$, and $r^{\mathcal{I}'} \subseteq r^{\mathcal{I}}$ for all $A \in \mathbf{N}_C$ and $r \in \mathbf{N}_R$. For $\Sigma \subseteq \mathbf{N}_C \cup \mathbf{N}_R$, \mathcal{I} is an interpretation *over signature* Σ if $A^{\mathcal{I}} = \emptyset$ and $r^{\mathcal{I}} = \emptyset$ for all $A \in \mathbf{N}_C \setminus \Sigma$ and $r \in \mathbf{N}_R \setminus \Sigma$. The union $\mathcal{I} \cup \mathcal{J}$ of \mathcal{I} and \mathcal{J} is an interpretation such that $\Delta^{\mathcal{I} \cup \mathcal{J}} = \Delta^{\mathcal{I}} \cup \Delta^{\mathcal{J}}$, $A^{\mathcal{I} \cup \mathcal{J}} = A^{\mathcal{I}} \cup A^{\mathcal{J}}$, and $r^{\mathcal{I} \cup \mathcal{J}} = r^{\mathcal{I}} \cup r^{\mathcal{J}}$ for all $A \in \mathbf{N}_C$ and $r \in \mathbf{N}_R$.

A *unary K-type* is a subset of $\text{CN}(\mathcal{K})$ including either A or \bar{A} for each $A \in \text{CN}(\mathcal{K})$. For an interpretation \mathcal{I} and an element $d \in \Delta^{\mathcal{I}}$, the *unary K-type of d in \mathcal{I}* is $\text{tp}^{\mathcal{I}}(d) = \{A \in \text{CN}(\mathcal{K}) \mid d \in A^{\mathcal{I}}\}$. We say that \mathcal{I} *realizes* a unary K-type τ if $\tau = \text{tp}^{\mathcal{I}}(d)$ for some $d \in \Delta^{\mathcal{I}}$.

2.3 Queries and Finite Entailment

We next introduce the query language. We concentrate on Boolean queries, that is, queries without answer variables. The extension to queries with answer variables is standard;

see, for example, (Glimm et al. 2008). A *conjunctive regular path query* (CRPQ) is a first-order formula

$$\varphi = \exists \mathbf{x} \psi(\mathbf{x})$$

such that $\psi(\mathbf{x})$ is constructed using \wedge over atoms of the form $A(t)$ or $\mathcal{E}(t, t')$ where $A \in \mathbf{N}_C$, t, t' are variables from \mathbf{x} or individual names from \mathbf{N}_I , and \mathcal{E} is a *path expression* defined by the grammar

$$\mathcal{E}, \mathcal{E}' ::= r \mid \mathcal{E}^* \mid \mathcal{E} \cup \mathcal{E}' \mid \mathcal{E} \circ \mathcal{E}'$$

where $r \in \mathbf{N}_R$. Thus, \mathcal{E} is essentially a regular expression over the (infinite) alphabet $\{r \mid r \in \mathbf{N}_R\}$. The set of individual names in φ is denoted with $\text{ind}(\varphi)$. A *conjunctive query* (CQ) is a CRPQ that does not use the operators $*$, \cup and \circ in path expressions, and a *regular path query* (RPQ) consists of a single atom of the form $\mathcal{E}(t, t')$.

The semantics of CRPQs is defined via matches. Let us fix a CRPQ $\varphi = \exists \mathbf{x} \psi(\mathbf{x})$ and an interpretation \mathcal{I} . A *match* for φ in \mathcal{I} is a function

$$\pi : \mathbf{x} \cup \text{ind}(\varphi) \rightarrow \Delta^{\mathcal{I}}$$

such that $\pi(a) = a$, for all $a \in \text{ind}(\varphi)$, and $\mathcal{I}, \pi \models \psi(\mathbf{x})$ under the standard semantics of first-order logic extended with a rule for atoms of the form $\mathcal{E}(t, t')$. More formally, we define:

- $\mathcal{I}, \pi \models \psi_1 \wedge \psi_2$ iff $\mathcal{I}, \pi \models \psi_1$ and $\mathcal{I}, \pi \models \psi_2$;
- $\mathcal{I}, \pi \models A(t)$ iff $\pi(t) \in A^{\mathcal{I}}$;
- $\mathcal{I}, \pi \models \mathcal{E}(t, t')$ iff $(\pi(t), \pi(t')) \in \mathcal{E}^{\mathcal{I}}$, where $\mathcal{E}^{\mathcal{I}}$ is defined inductively as $(\mathcal{E}^*)^{\mathcal{I}} = (\mathcal{E}^{\mathcal{I}})^*$, $(\mathcal{E}_1 \cup \mathcal{E}_2)^{\mathcal{I}} = \mathcal{E}_1^{\mathcal{I}} \cup \mathcal{E}_2^{\mathcal{I}}$, $(\mathcal{E}_1 \circ \mathcal{E}_2)^{\mathcal{I}} = \mathcal{E}_1^{\mathcal{I}} \circ \mathcal{E}_2^{\mathcal{I}}$.

An interpretation \mathcal{I} *satisfies* φ , written $\mathcal{I} \models \varphi$, if there exists a match for φ in \mathcal{I} . A *union of CRPQs* (UCRPQ) is a finite set of CRPQs and a *union of CQs* (UCQ) is a finite set of CQs. An interpretation \mathcal{I} satisfies an UCRPQ Φ , written as $\mathcal{I} \models \Phi$, if $\mathcal{I} \models \varphi$ for some $\varphi \in \Phi$. We say that \mathcal{K} *finitely entails* Φ , written $\mathcal{K} \models_{\text{fin}} \Phi$, if each finite model of \mathcal{K} satisfies Φ . A model of \mathcal{K} that does not satisfy Φ is a *counter-model*. The *finite entailment problem* asks if a given KB \mathcal{K} finitely entails a given query Φ .

2.4 UCRPQs via Semiautomata

We work with UCRPQs represented by means of a *semiautomaton* (Ginzburg 1968) $\mathcal{B} = (Q, \Gamma, \delta)$ where Q is a finite set of states, $\Gamma \subseteq \{r \mid r \in \mathbf{N}_R\}$ is a finite alphabet—throughout the paper we assume $\Gamma = \text{rol}(\mathcal{K})$, and $\delta : Q \times \Gamma \rightarrow Q$ is the transition function. A semiautomaton is essentially a deterministic finite automaton without initial and final states; a run of a semiautomaton \mathcal{B} over a word w is defined just like for a finite automaton, except that it can begin in any state and there is no notion of accepting runs. Under this representation, an RPQ is an atom over a binary predicate of the form $\mathcal{B}_{q, q'}$ where $q, q' \in Q$ are states of \mathcal{B} . We let $\mathcal{I}, \pi \models \mathcal{B}_{q, q'}(t, t')$ iff $(\pi(t), \pi(t')) \in \mathcal{B}_{q, q'}^{\mathcal{I}}$, where $\mathcal{B}_{q, q'}^{\mathcal{I}}$ is the set of pairs (e, e') such that for some $n \in \mathbb{N}$ there exist $r_1, \dots, r_n \in \Gamma$ and $e_0, \dots, e_n \in \Delta^{\mathcal{I}}$ such that

- $e_0 = e$ and $e_n = e'$;

- $(e_{i-1}, e_i) \in (r_i)^{\mathcal{I}}$ for all $i \in \{1, \dots, n\}$;
- there exists a run of \mathcal{B} on the word $r_1 \dots r_n$ that begins in state q and ends in state q' .

We also allow *edge atoms* of the form $r(x, x')$ for $r \in \Gamma$.

Each UCRPQ Φ can be effectively rewritten into a UCRPQ Φ' expressed by means of a semiautomaton \mathcal{B} of size $k \cdot 2^{O(m)}$ where k is the number of path expressions in Φ and m is their maximal size. The size of CRPQs in Φ' is bounded by the size of CRPQs in Φ and $|\Phi'| = 2^{\text{poly}(\|\Phi\|)}$, where $\|\Phi\|$ is the total size of Φ .

For simplicity we work with KBs $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ where the ABox \mathcal{A} is *trivial*; that is, $\text{ind}(\mathcal{A}) = \{a\}$ for some $a \in \mathbb{N}_1$ and \mathcal{A} contains only concept assertions. The general finite entailment problem can be reduced to this special case using the following lemma.

Lemma 1. *Given an oracle for finite entailment for trivial ABoxes, the general finite entailment $(\mathcal{T}, \mathcal{A}) \models_{\text{fin}} \Phi$ can be decided in time $2^{\text{poly}(\|(\mathcal{T}, \mathcal{A})\|)} \cdot 2^{\text{poly}(\|\Phi\|)}$ using calls to the oracle for $\mathcal{K}' = (\mathcal{T}, \mathcal{A}')$ and Φ' consisting of $2^{\text{poly}(\|\Phi\|)}$ CRPQs of linear size over the same semiautomaton as Φ .*

2.5 Entailment Modulo Environment

We solve the entailment problem using a divide-and-conquer approach in which counter-models are decomposed into simpler ones, whose existence is easier to decide. Each level of this recursive procedure will involve certain modifications to the TBox. For complexity reasons we need to pay close attention to these changes, making sure that no blow-up is involved. To make it easier, we generalize the entailment problem by turning the modifications into a separate part of the input, which allows fixing the TBox for the duration of the whole procedure. At every level of the recursion, we will need to reason ‘externally’ about the way simpler pieces are put together to form the larger counter-model, and ‘internally’ about how to specify the required properties of a piece depending on what is happening outside. We will think of the models as induced subinterpretations of a larger interpretation. Dually, the remaining part of the larger interpretation can be seen as an external context, in which our models live. The relevant features of this context will be represented by environments, which we now define.

An *environment* $\mathcal{E} = (\Theta, \varepsilon)$ consists of a set Θ of unary types and a function $\varepsilon : \Theta \rightarrow 2^{\text{rol}(\mathcal{K}) \times \text{CN}(\mathcal{K})}$. The intended meaning is that only types from Θ are allowed and each element of an allowed unary type τ has an r -edge to an element in the extension of B in the external context for each $(r, B) \in \varepsilon(\tau)$. Accordingly, we say that \mathcal{I} is a *model of \mathcal{K} modulo \mathcal{E}* and write $\mathcal{I} \models_{\mathcal{E}}^{\mathcal{K}}$ if \mathcal{I} realizes only unary types from Θ and it is a model of \mathcal{K} under the following *relaxed semantics of existential restrictions*:

- for every existential restriction $\exists r.B$ in \mathcal{K} and every element $d \in \Delta^{\mathcal{I}}$, $d \in (\exists r.B)^{\mathcal{I}}$ iff either there is an r -edge in \mathcal{I} from d to an element $e \in B^{\mathcal{I}}$ or $(r, B) \in \varepsilon(\text{tp}^{\mathcal{I}}(d))$.

(The semantics of universal restrictions is not altered and it is the environment’s responsibility to account for them.) Correspondingly, a query Φ is *finitely entailed by \mathcal{K} modulo \mathcal{E}* , written $\mathcal{K} \models_{\text{fin}}^{\mathcal{E}} \Phi$, if for each finite interpretation \mathcal{I} , if

$\mathcal{I} \models^{\mathcal{E}} \mathcal{K}$ then $\mathcal{I} \models \Phi$. The problem of *finite entailment modulo environment* is to decide for a given KB \mathcal{K} , environment \mathcal{E} , and query Φ if $\mathcal{K} \models_{\text{fin}}^{\mathcal{E}} \Phi$.

Note that finite entailment modulo environment and ordinary finite entailment are interreducible. In one direction, it is enough to take the set of all unary \mathcal{K} -types for Θ and set $\varepsilon(\tau) = \emptyset$ for all $\tau \in \Theta$. In the other direction, the conditions imposed on unary types and the relaxed semantics of existential restrictions can be expressed easily in the TBox. The latter reduction, however, might significantly increase the size of the TBox. It is easier to control the size of the input at different levels of the recursion when these conditions are explicitly represented in the environment.

3 Expansion and Decorations

In order to handle UCRPQs expressed by means of a semiautomaton \mathcal{B} we need to be able to trace runs of \mathcal{B} that begin in all possible states, on all infixes of the input word. We achieve this using the following construction.

Let us fix an arbitrary linear order on the set Q of the states of \mathcal{B} . The *expansion* of \mathcal{B} is a semiautomaton $\hat{\mathcal{B}}$ whose set of states is the set \hat{Q} of all permutations of Q . Thus, an element of \hat{Q} can be seen as a tuple $\mathbf{p} = (p_1, p_2, \dots, p_n)$ such that p_i is the image of the i th state of \mathcal{B} under the respective permutation. We refer to positions in this tuple as *levels*. In particular, the *level of $q \in Q$ in \mathbf{p}* is the unique i such that $q = p_i$. Assuming $\delta : Q \times \text{rol}(\mathcal{K}) \rightarrow Q$ is the transition function of \mathcal{B} , we define the transition function

$$\hat{\delta} : \hat{Q} \times \text{rol}(\mathcal{K}) \rightarrow \hat{Q}$$

of $\hat{\mathcal{B}}$ by letting $\hat{\delta}(\mathbf{p}, r)$ be the permutation \mathbf{p}' obtained by listing all states appearing in the sequence

$$\delta(\mathbf{p}, r) = (\delta(p_1, r), \delta(p_2, r), \dots, \delta(p_n, r))$$

in the order of their first appearances, followed by all remaining states of \mathcal{B} ordered as in Q . Note that the level of $\delta(p_i, r)$ in \mathbf{p}' is at most i . Consider the set $P \subseteq \{1, 2, \dots, n\}$ of levels i such that the level of $\delta(p_i, r)$ in \mathbf{p}' is equal to i . It follows from the definition of \mathbf{p}' that $P = \{1, 2, \dots, \ell\}$ for some $\ell \in \{1, 2, \dots, n\}$. We call this number ℓ the *level of transition $\mathbf{p} \xrightarrow{r} \mathbf{p}'$* .

From each run of $\hat{\mathcal{B}}$ on a word w we can reconstruct all runs of \mathcal{B} on w . Let $\mathbf{p}_0, \mathbf{p}_1, \dots, \mathbf{p}_m$ be a run of $\hat{\mathcal{B}}$ on w . Consider a run q_0, q_1, \dots, q_m of \mathcal{B} on w . For $i = 0, 1, \dots, m$, let ℓ_i be the level of q_i in \mathbf{p}_i . Any sequence $\ell_0, \ell_1, \dots, \ell_m$ associated like this with a run of \mathcal{B} will be called a *thread* in the run of $\hat{\mathcal{B}}$ (see Fig. 1). Notice that two threads that begin at different levels can meet at the same level somewhere along the run; if this happens they remain equal until the end of the run. Also, threads can be born in the middle of a run of $\hat{\mathcal{B}}$, but they never disappear. A crucial property of threads is that they are non-increasing sequences: the level of q_{i+1} in \mathbf{p}_{i+1} is bounded by the level of q_i in \mathbf{p}_i .

Lemma 2. *Let $\mathbf{p}_0, \mathbf{p}_1, \dots, \mathbf{p}_m$ be a run of $\hat{\mathcal{B}}$ on w , and let q, q' be states of \mathcal{B} . There is a run of \mathcal{B} on w from q to q' iff there exist positions $0 \leq j_1 < j_2 < \dots < j_k = m$, levels $n \geq \ell_1 > \ell_2 > \dots > \ell_k \geq 1$, and states q_0, q_1, \dots, q_k with $1 \leq k \leq n$ such that $q_0 = q$, $q_k = q'$, and*




Figure 1: A thread in a run of the expansion of a semiautomaton.

- the level of q_0 in \mathbf{p}_0 is ℓ_1 and the level of q_k in \mathbf{p}_m is ℓ_k ;
- for all $i \in \{1, 2, \dots, k-1\}$, the level of q_i in \mathbf{p}_{j_i} is ℓ_i and the level of $\delta(q_i, w[j_i+1])$ in $\mathbf{p}_{j_{i+1}}$ is ℓ_{i+1} ;
- for all $i \in \{1, 2, \dots, k\}$, each transition taken in the segment of the run from $\mathbf{p}_{j_{i-1}+1}$ (or \mathbf{p}_0 for $i=1$) to \mathbf{p}_{j_i} has level at least ℓ_i .

As an illustration of Lemma 2, consider the run of the expanded semiautomaton shown in Fig. 1. Tracing the run of the original semiautomaton on the same word, starting in state q_4 , we discover the positions $j_1 = 3$ and $j_2 = 7$ where the corresponding thread drops to a lower level. Between these positions, the thread stays at the same level, beginning with $\ell_1 = 4$ (taking transitions of levels $5, 4, 5 \geq \ell_1$), followed by $\ell_2 = 3$ (taking transitions of levels $5, 3, 5 \geq \ell_2$), and $\ell_3 = 1$ (taking a transition of level $5 \geq \ell_3$).

The next step is to account for the possible runs of $\hat{\mathcal{B}}$ over paths in the interpretation. Towards this goal, we decorate elements of the interpretation with states of $\hat{\mathcal{B}}$. To avoid additional blow-up, we represent states of $\hat{\mathcal{B}}$ using combinations of fresh concept names $C_{q,\ell}$ where q is a state of \mathcal{B} and $\ell \in \{1, 2, \dots, n\}$ is a level; we write $\text{CN}(\hat{\mathcal{B}})$ for the set of all $C_{q,\ell}$. For a state $\mathbf{p} = (p_1, p_2, \dots, p_n)$ of $\hat{\mathcal{B}}$, by $C_{\mathbf{p}}$ we mean the concept $C_{p_1,1} \sqcap C_{p_2,2} \sqcap \dots \sqcap C_{p_n,n}$. We say that an element $e \in \Delta^{\mathcal{I}}$ is *decorated* with state \mathbf{p} if $e \in C_{\mathbf{p}}^{\mathcal{I}}$. An interpretation \mathcal{I} is *$\hat{\mathcal{B}}$ -decorated* if no element has incoming edges over different roles from $\text{rol}(\mathcal{K})$ and \mathcal{I} satisfies the CIs

$$C_{\mathbf{p}} \sqsubseteq \forall r. C_{\delta(\mathbf{p},r)}, \quad C_{\mathbf{p}} \sqcap C_{\mathbf{p}'} \sqsubseteq \perp, \quad \top \sqsubseteq \bigsqcup_{\mathbf{p} \in \hat{Q}} C_{\mathbf{p}}$$

for all states \mathbf{p}, \mathbf{p}' of $\hat{\mathcal{B}}$ such that $\mathbf{p} \neq \mathbf{p}'$. The axiomatization above is exponential in the size of \mathcal{B} , but we can do better.

Lemma 3. *Given \mathcal{B} one can compute in polynomial time a TBox $\hat{\mathcal{T}}_{\mathcal{B}}$ such that $\mathcal{I} \models \hat{\mathcal{T}}_{\mathcal{B}}$ iff \mathcal{I} is $\hat{\mathcal{B}}$ -decorated.*

To every edge in a $\hat{\mathcal{B}}$ -decorated interpretation \mathcal{I} we can assign a level as follows. Consider elements $e, e' \in \Delta^{\mathcal{I}}$ such that $(e, e') \in r^{\mathcal{I}}$ for some $r \in \text{rol}(\mathcal{K})$. Note that $(e, e') \notin s^{\mathcal{I}}$ for every $s \in \text{rol}(\mathcal{K}) \setminus \{r\}$. Let \mathbf{p} and \mathbf{p}' be the states decorating e and e' , respectively. It holds that $\mathbf{p} \xrightarrow{r} \mathbf{p}'$. By the level of the edge (e, e') we shall understand the level of this

transition. A *level- ℓ interpretation* is a $\hat{\mathcal{B}}$ -decorated interpretation that does not contain edges of level strictly below ℓ ; if $\ell > n$, this means that there are no edges at all. The following lemma is the key to our algorithm.

Lemma 4. *Consider a level- ℓ interpretation \mathcal{I} and elements $e \in C_{q,\ell}^{\mathcal{I}}$ and $e' \in C_{q',\ell}^{\mathcal{I}}$. Then, $(e, e') \in \mathcal{B}_{q,q'}^{\mathcal{I}}$ iff there is a path from e to e' in \mathcal{I} .*

We make use of Lemma 4 by decomposing RPQs into segments corresponding to different levels, as was done for the runs of $\hat{\mathcal{B}}$ in Lemma 2. To facilitate this, we make our queries aware of levels. A *$\hat{\mathcal{B}}$ -decorated CRPQ* is a CRPQ φ represented by means of semiautomaton \mathcal{B} that contains exactly one atom of the form $C_{q,\ell}(x)$ and exactly one atom of the form $C_{q',\ell'}(x')$ for each atom $\mathcal{B}_{q,q'}(x, x')$ in φ . We call ℓ and ℓ' the *begin level* and the *end level* of atom $\mathcal{B}_{q,q'}(x, x')$ in φ , respectively. Because levels never increase in a thread of a run of $\hat{\mathcal{B}}$, we can assume without loss of generality that $\ell \geq \ell'$ always holds. A *level- ℓ CRPQ* is a $\hat{\mathcal{B}}$ -decorated CRPQ that contains no RPQ atoms of end level strictly below ℓ . As all end levels are at most n , a level- ℓ CRPQ for $\ell > n$ contains no RPQ atoms; that is, it is a CQ. To *complete* a CRPQ φ means to turn it into a $\hat{\mathcal{B}}$ -decorated CRPQ φ' by adding unary atoms over concepts $C_{q,\ell}$ in an arbitrary minimal way. Each resulting φ' is called a *completion* of φ . Over $\hat{\mathcal{B}}$ -decorated interpretations, φ is equivalent to the union of its completions. The *completion* of a UCRPQ Φ is the union of all completions of all CRPQs in Φ .

We conclude this section by showing how to turn any counterexample to $\mathcal{K} \models_{\text{fin}}^{\varepsilon} \Phi$ into a $\hat{\mathcal{B}}$ -decorated one. Let \mathcal{I} be an interpretation over $\text{CN}(\mathcal{K}) \cup \text{rol}(\mathcal{K})$. The *product* of \mathcal{I} and $\hat{\mathcal{B}}$ is the interpretation $\mathcal{I} \times \hat{\mathcal{B}}$ over $\text{CN}(\hat{\mathcal{B}}) \cup \text{CN}(\mathcal{K}) \cup \text{rol}(\mathcal{K})$ such that

- $\Delta^{\mathcal{I} \times \hat{\mathcal{B}}} = \Delta^{\mathcal{I}} \times \text{rol}(\mathcal{K}) \times \hat{Q}$,
- $C^{\mathcal{I} \times \hat{\mathcal{B}}} = C^{\mathcal{I}} \times \text{rol}(\mathcal{K}) \times \hat{Q}$ for all $C \in \text{CN}(\mathcal{K})$,
- $C_{q,\ell}^{\mathcal{I} \times \hat{\mathcal{B}}} = \Delta^{\mathcal{I}} \times \text{rol}(\mathcal{K}) \times \{(p_1, p_2, \dots, p_n) \in \hat{Q} : p_{\ell} = q\}$ for all $q \in Q$ and $\ell \in \{1, 2, \dots, n\}$,
- $r^{\mathcal{I} \times \hat{\mathcal{B}}} = \{(e, s, \mathbf{p}), (e', r, \mathbf{p}') : (e, e') \in r^{\mathcal{I}}, \mathbf{p} \xrightarrow{r} \mathbf{p}', s \in \text{rol}(\mathcal{K})\}$ for $r \in \text{rol}(\mathcal{K})$.

Note that if \mathcal{I} is finite, so is $\mathcal{I} \times \hat{\mathcal{B}}$.

Lemma 5. *Let Φ be a UCRPQ, \mathcal{K} an \mathcal{ALC} KB with a trivial ABox, and \mathcal{E} an environment.*

- $\mathcal{I} \times \hat{\mathcal{B}}$ is a $\hat{\mathcal{B}}$ -decorated interpretation.
- If $\mathcal{I} \not\models \Phi$ then $\mathcal{I} \times \hat{\mathcal{B}} \not\models \Phi$.
- If $\mathcal{I} \models^{\mathcal{E}} \mathcal{K}$ then $\mathcal{I} \times \hat{\mathcal{B}} \models^{\mathcal{E}} \mathcal{K}$ up to identifying the unique individual a in \mathcal{K} with some $(a, r, \mathbf{p}) \in \Delta^{\mathcal{I} \times \hat{\mathcal{B}}}$.

4 Core Computational Problem

To solve the entailment problem we eliminate the lowest level from the query and from the interpretation, and solve the problem with fewer levels recursively. Eliminating each level will involve interpretations built from pieces that are solutions for the simplified problem. Evaluating CRPQs over such interpretations requires breaking them down into fragments and it must accommodate single RPQs witnessed across multiple pieces.

For a UCRPQ Φ let $\tilde{\Phi}$ be the completion of an equivalent UCRPQ represented by means of a semiautomaton \mathcal{B} . A *fragment* of $\varphi \in \tilde{\Phi}$ is either of the following:

- a $\hat{\mathcal{B}}$ -decorated CRPQ of the form $C_{q_1, \ell_1}(y_1) \wedge \mathcal{B}_{q_1, q_2}(y_1, y_2) \wedge C_{q_2, \ell_2}(y_2)$ or $C_{q_1, \ell_1}(y_1) \wedge \mathcal{B}_{q_1, q_2}(y_1, y_2) \wedge C_{q_2, \ell_2}(y_2) \wedge r(y_2, y_3) \wedge C_{q_3, \ell_3}(y_3)$ where y_1, y_2, y_3 are fresh variables and $r \in \text{rol}(\mathcal{K})$,
- a connected $\hat{\mathcal{B}}$ -decorated CRPQ that can be obtained from φ by dropping selected atoms, replacing selected RPQ atoms $\mathcal{B}_{q, q'}(x, x')$ by a subset of $\mathcal{B}_{q, q_1}(x, y_1), r(y_1, y_2), \mathcal{B}_{q_3, q'}(y_3, x')$ for some fresh variables y_1, y_2, y_3 and $r \in \text{rol}(\mathcal{K})$, and completing the resulting CRPQ.

A fragment of Φ is a fragment of any of the CRPQs in $\tilde{\Phi}$. Importantly, a fragment of a fragment of Φ is also a fragment of Φ , and each $\varphi \in \tilde{\Phi}$ is a fragment of Φ . Up to renaming fresh variables, Φ has $2^{\text{poly}(\|\Phi\|)}$ different fragments, despite \mathcal{B} being exponential in $\|\Phi\|$.

We now enrich interpretations again by including information about matched fragments of $\tilde{\Phi}$. For each fragment φ of $\tilde{\Phi}$ and each $\emptyset \neq V \subseteq \text{var}(\varphi)$ we choose a fresh concept name $A_{\varphi, V}$. We call an interpretation \mathcal{I} *correct* (wrt. Φ) if $e \in A_{\varphi, V}^{\mathcal{I}}$ iff $\pi(V) = \{e\}$ for some match π for φ in \mathcal{I} . Assuming \mathcal{I} is correct, $\mathcal{I} \models \Phi$ iff $A_{\varphi, V}^{\mathcal{I}} \neq \emptyset$ for some $\varphi \in \tilde{\Phi}$ and $\emptyset \neq V \subseteq \text{var}(\varphi)$. Correctness is not compositional: the union of two correct interpretations sharing a single element need not be correct. As our method of eliminating levels relies on such decompositions of interpretations, we replace correctness with a notion that is weaker, but compositional.

We first abstract the decomposition of a $\hat{\mathcal{B}}$ -decorated CRPQ induced by a match in a union of disjoint ‘peripheric’ interpretations, each sharing a single element with a single ‘core’ interpretation (Fig. 2 shows three ‘peripheric’ interpretations connected to the ‘core’ by single edges, included in the ‘peripheric’ interpretations). A *partition* of a $\hat{\mathcal{B}}$ -decorated CRPQ φ into $\varphi', \varphi_1, \dots, \varphi_k$ is obtained as follows. Choose $X', X_1, \dots, X_k \subseteq \text{var}(\varphi)$ such that

- $X_i \cap X_j = \emptyset$ for all $i \neq j$;




Figure 2: CRPQ φ is distributed over the bags constituting \mathcal{I} .

- for each atom of the form $r(x, x')$ in φ there exists i such that either $\{x, x'\} \subseteq X_i$ or $\{x, x'\} \subseteq X'$.

Based on X', X_1, \dots, X_k define $\varphi', \varphi_1, \dots, \varphi_k$ as follows. For each atom of the form $r(x, x')$ in φ choose i such that $\{x, x'\} \subseteq X_i$ and add $r(x, x')$ to φ_i or add $r(x, x')$ to φ' provided that $\{x, x'\} \subseteq X'$. For each RPQ atom $\mathcal{B}_{q, q'}(x, x')$ of begin level ℓ and end level ℓ' in φ do one of the following:

- provided that $\{x, x'\} \subseteq X'$, add $\mathcal{B}_{q, q'}(x, x')$ to φ' ;
- choose i such that $\{x, x'\} \subseteq X_i$ but $\{x, x'\} \not\subseteq X'$, and add $\mathcal{B}_{q, q'}(x, x')$ to φ_i (light green RPQ in Fig. 2);
- choose i such that $x \in X' \setminus X_i$ and $x' \in X_i \setminus X'$, a level m such that $\ell \geq m \geq \ell'$, a state p of \mathcal{B} , and a fresh variable y , and add $\mathcal{B}_{q, p}(x, y) \wedge C_{p, m}(y)$ to φ' and $C_{p, m}(y) \wedge \mathcal{B}_{p, q'}(y, x')$ to φ_i (blue and orange in Fig. 2);
- choose i such that $x \in X_i \setminus X'$ and $x' \in X' \setminus X_i$, a level m such that $\ell \geq m \geq \ell'$, a state p of \mathcal{B} , and a fresh variable y , and add $\mathcal{B}_{q, p}(x, y) \wedge C_{p, m}(y)$ to φ_i and $C_{p, m}(y) \wedge \mathcal{B}_{p, q'}(y, x')$ to φ' (dark green in Fig. 2);
- choose $i \neq j$ such that $x \in X_i \setminus X'$ and $x' \in X_j \setminus X'$, levels m, m' such that $\ell \geq m \geq m' \geq \ell'$, states p, p' of \mathcal{B} , and fresh variables y, y' , add $\mathcal{B}_{q, p}(x, y) \wedge C_{p, m}(y)$ to φ_i , $C_{p, m}(y) \wedge \mathcal{B}_{p, p'}(y, y') \wedge C_{p', m'}(y')$ to φ' , and $C_{p', m'}(y') \wedge \mathcal{B}_{p', q'}(y', x')$ to φ_j (purple in Fig. 2).

Note that for each $\mathcal{B}_{q, q'}(x, x')$ exactly one of the above actions can be performed and the choice of i and j is unique. To complete the construction, add to φ' all unary atoms of φ over variables already used in φ' , and similarly for each φ_i . Observe that for each $X' \subseteq \text{var}(\varphi)$ there is exactly one choice of X_1, X_2, \dots, X_k (up to a permutation) such that the resulting $\varphi_1, \varphi_2, \dots, \varphi_k$ are connected (regardless of the choice of p, p' and m, m'). Assuming that φ is a fragment of Φ , it then holds that so are $\varphi_1, \varphi_2, \dots, \varphi_k$.

We call \mathcal{I} *consistent* (wrt. Φ) if for each partition of a fragment φ of Φ into a CRPQ φ' and fragments $\varphi_1, \varphi_2, \dots, \varphi_k$ with $\text{var}(\varphi_i) \cap \text{var}(\varphi_j) = \emptyset$ for $i \neq j$, $V_i = \text{var}(\varphi_i) \cap \text{var}(\varphi')$, and $\emptyset \neq V \subseteq \text{var}(\varphi) \cap \text{var}(\varphi')$, there is no match π for φ' in \mathcal{I} such that $\pi(V_i) = \{e_i\} \subseteq (A_{\varphi_i, V_i})^{\mathcal{I}}$ for all i but $\pi(V) = \{e\} \not\subseteq (A_{\varphi, V})^{\mathcal{I}}$. Clearly, all correct interpretations

are consistent. The converse is not true in general, but the following key property is preserved.

Lemma 6. *For every UCRPQ Φ and every consistent \widehat{B} -decorated interpretation \mathcal{I} , if $A_{\varphi, V}^{\mathcal{I}} = \emptyset$ for each $\varphi \in \tilde{\Phi}$ and $\emptyset \neq V \subseteq \text{var}(\varphi)$, then $\mathcal{I} \not\models \Phi$.*

Consistency is sufficient to express entailment, but it does not yield well to the recursive elimination of levels. We generalize it by refining the information about matched fragments of Φ . We introduce fresh concepts $A_{\varphi, V}^{\kappa}$ where φ is a fragment of Φ , $\emptyset \neq V \subseteq \text{var}(\varphi)$,

$$\kappa : \text{var}(\varphi) \rightarrow \{1, 2, \dots, \ell\},$$

and $\kappa(V) = \{\ell\}$ for some $\ell \in \{1, 2, \dots, n+1\}$. We write CN_{ℓ}^{Φ} for the set of $A_{\psi, V}^{\kappa}$ such that $\kappa(V) = \{\ell\}$. Intuitively, κ is a synopsis of when specific fragments of ψ were matched during the recursive search for the model. Specifically, $\kappa(x) = \ell$ indicates that x was matched after all levels strictly below ℓ had been eliminated from the query, but while level ℓ was still present. Accordingly, ℓ -consistency, defined below, ensures that the synopses built so far are consistently updated while level ℓ is being handled.

We call \mathcal{I} ℓ -consistent (wrt. Φ) if for each partition of a fragment φ of Φ into a CRPQ φ' of level ℓ and fragments $\varphi_1, \varphi_2, \dots, \varphi_k$ with $\text{var}(\varphi_i) \cap \text{var}(\varphi_j) = \emptyset$ for $i \neq j$, $V_i = \text{var}(\varphi_i) \cap \text{var}(\varphi')$, and $\emptyset \neq V \subseteq \text{var}(\varphi) \cap \text{var}(\varphi')$, there is no match π for φ' in \mathcal{I} such that $\pi(V_i) = \{e_i\} \subseteq (A_{\varphi_i, V_i}^{\kappa_i})^{\mathcal{I}}$ for all i but $\pi(V) = \{e\} \not\subseteq (A_{\varphi, V}^{\kappa})^{\mathcal{I}}$ where

- $\kappa_i(x) \leq \ell$ for all $x \in \text{var}(\varphi_i)$,
- $\kappa(x) = \kappa_i(x)$ for all $x \in \text{var}(\varphi_i) \setminus V_i$,
- $\kappa(x) = \ell$ for all $x \in \text{var}(\varphi) \cap \text{var}(\varphi')$.

We stress that while φ' has level ℓ , fragments $\varphi, \varphi_1, \dots, \varphi_k$ can have any level. Note also that ℓ -consistency speaks only of concept names in $\text{CN}_1^{\Phi} \cup \text{CN}_2^{\Phi} \cup \dots \cup \text{CN}_{\ell}^{\Phi}$. Identifying $A_{\varphi, V}$ with $A_{\varphi, V}^{\kappa}$ for κ constantly equal to 1, we get that consistency and 1-consistency are equivalent.

In what follows, by an (ℓ, ℓ') -interpretation we mean an ℓ' -consistent level- ℓ interpretation. By an (ℓ, ℓ') -model of \mathcal{K} modulo \mathcal{E} we mean an (ℓ, ℓ') -interpretation that is model of \mathcal{K} modulo \mathcal{E} . The actual problem we will be solving is the following (ℓ, ℓ') -model problem for $\ell \leq \ell'$: Given a KB \mathcal{K} with a trivial ABox, an environment \mathcal{E} , and a UCRPQ Φ decide if there exists a finite (ℓ, ℓ') -model of \mathcal{K} modulo \mathcal{E} .

By Lemma 5, entailment modulo environment (with trivial ABox) can be reduced to the $(1, 1)$ -model problem by modifying the environment to forbid all unary types containing $A_{\varphi, V}^{\kappa}$ for any $\varphi \in \tilde{\Phi}$, $\emptyset \neq V \subseteq \text{var}(\varphi)$, and κ constantly equal 1. Note that the reduction does not affect the query Φ , nor the KB \mathcal{K} . However, it introduces up to $2^{\text{poly}(\|\Phi\|)}$ new concept names $A_{\varphi, V}$ and $C_{q, \ell}$. Consequently, the number of unary types is at most $2^{|\text{CN}(\mathcal{K})| + 2^{\text{poly}(\|\Phi\|)}}$. It follows that the size of the environment is bounded by $2^{\|\mathcal{K}\| + 2^{\text{poly}(\|\Phi\|)}}$.

To solve the $(1, 1)$ -model problem we will proceed recursively, incrementing ℓ and ℓ' in an alternating fashion, until $\ell = \ell' = n+1$. At each level of the recursion we will be making multiple recursive calls. During the recursion

the UCRPQ Φ and the TBox \mathcal{T} will remain unchanged, but the ABox and the environment will evolve. Importantly, we will not introduce any new concepts, so the size of the environment will always be bounded by $2^{\|\mathcal{K}\| + 2^{\text{poly}(\|\Phi\|)}}$. The size of the ABox will be bounded by $\|\mathcal{K}\| + 2^{\text{poly}(\|\Phi\|)}$ and the number of individuals will never grow. In consequence, the total cost of the algorithm can be computed as the cost of a single recursion step times the number of steps. In the following sections we will show that each recursion step can be carried out in time $2^{O(\|\mathcal{K}\|) + 2^{\text{poly}(\|\Phi\|)}}$, excluding the cost of the recursive calls. The depth of the recursion is $O(n) = 2^{\text{poly}(\|\Phi\|)}$. The number of recursive calls within a single recursion step is also bounded by $2^{O(\|\mathcal{K}\|) + 2^{\text{poly}(\|\Phi\|)}}$. This means that the total number of recursion steps is $2^{\|\mathcal{K}\| \cdot 2^{\text{poly}(\|\Phi\|)}}$ and so is the overall complexity of the recursive algorithm for the $(1, 1)$ -model problem.

5 Incrementing the Level of Queries

The main goal of this section is to solve the (ℓ, ℓ) -model problem by reduction to multiple instances of the $(\ell, \ell+1)$ -model problem for $\ell \leq n$. The $(n+1, n+1)$ -model problem is discussed briefly at the end of the section.

As a first step, we observe that it is enough to consider interpretations whose DAG of strongly connected components is a tree. For this purpose we define *tree-like* interpretations as those that can be decomposed into multiple finite subinterpretations, called *bags*, arranged into a (possibly infinite) tree such that: (1) all bags are pairwise disjoint; (2) between each parent and child bag there is a single edge, pointing from an element of the parent bag to an element of the child bag; (3) all other edges are between elements of the same bag. We think of edges between bags as 2-element interpretations, called *edge-bags*, sharing the origin with the parent bag and the target with the child bag. A tree-like interpretation is then a union of all its bags and edge-bags. Fig. 2 shows a tree-like interpretation with 4 bags and 3 edge-bags. In tree-like interpretations ℓ -consistency is a local property.

Lemma 7. *A tree-like interpretation is ℓ -consistent iff each of its bags and edge-bags is ℓ -consistent.*

Lemma 8. *There is a finite (ℓ, ℓ') -model of \mathcal{K} modulo \mathcal{E} iff there is a finite tree-like (ℓ, ℓ') -model of \mathcal{K} modulo \mathcal{E} whose bags are strongly connected.*

The next step is to eliminate the lowest level from the queries. An $(\ell+1)$ -reduct of a level- ℓ CRPQ φ is any CRPQ that can be obtained from φ by first splitting each RPQ atom $\mathcal{B}_{q_1, q_2}(x_1, x_2)$ of begin level $\ell_1 > \ell$ and end level ℓ into $\mathcal{B}_{q_1, q'_1}(x_1, x'_1) \wedge C_{q'_1, \ell'_1}(x'_1) \wedge r(x'_1, x'_2) \wedge C_{q'_2, \ell}(x'_2) \wedge \mathcal{B}_{q'_2, q_2}(x'_2, x_2)$ where $\ell_1 \geq \ell'_1 \geq \ell+1$, and then dropping from the resulting CRPQ all atoms whose begin and end level is ℓ (all unary atoms are kept). Note that each $(\ell+1)$ -reduct φ' of φ is a conjunction of at most $|\varphi|$ disjoint fragments of φ and that $\text{var}(\varphi) \subseteq \text{var}(\varphi')$.

Lemma 9. *Over \widehat{B} -decorated interpretations, each level- ℓ CRPQ implies the union of its $(\ell+1)$ -reducts. Over strongly-connected level- ℓ interpretations, each level- ℓ CRPQ is equivalent to the union of its $(\ell+1)$ -reducts.*

Because Lemma 8 guarantees tree-like solutions with strongly connected bags, we can replace ℓ -consistency with *strong ℓ -consistency*: the only difference is that π ranges over matches of all possible $(\ell+1)$ -reducts of φ' , rather than over matches of φ' itself. We restate Lemma 8 as follows.

Lemma 10. *There is a finite (ℓ, ℓ) -model of \mathcal{K} modulo \mathcal{E} iff there is a finite tree-like level- ℓ model of \mathcal{K} modulo \mathcal{E} whose edge bags are ℓ -consistent and bags strongly ℓ -consistent.*

It remains to show how to find models of the latter form. Let us first see how to find one consisting of a single bag; that is, how to find a finite strongly ℓ -consistent level- ℓ model of \mathcal{K} modulo \mathcal{E} . We will show that this amounts to finding a finite $(\ell, \ell+1)$ -model of \mathcal{K} modulo \mathcal{E}' for one of the $(\ell+1)$ -reducts \mathcal{E}' of \mathcal{E} described below.

Consider a fragment φ , a non-empty set $V \subseteq \text{var}(\varphi)$, a partition of φ into a CRPQ φ' of level ℓ , and fragments $\varphi_1, \varphi_2, \dots, \varphi_k$, as in the definition of (strong) ℓ -consistency. Let $\kappa : \text{var}(\varphi) \rightarrow \{1, 2, \dots, \ell\}$ be such that $\kappa(\text{var}(\varphi) \cap \text{var}(\varphi')) = \{\ell\}$. Let ψ' be an $(\ell+1)$ -reduct of φ' . Consider CRPQs ψ with $\text{var}(\varphi) \subseteq \text{var}(\psi)$ that can be partitioned into ψ' and $\varphi_1, \varphi_2, \dots, \varphi_k$. Choose the one with minimal $\text{var}(\psi)$. This amounts to merging back all RPQ atoms split during the partition of φ , provided that their segments were not affected by replacing φ' with ψ' . The CRPQ ψ is not a fragment, because it need not be connected: Figure 2 right illustrates passing from φ to ψ consisting of two disconnected fragments. Let $\psi_1, \psi_2, \dots, \psi_m$ be the fragments constituting ψ and let $U_i = V \cap \text{var}(\psi_i)$. An $(\ell+1)$ -reduct \mathcal{E}' of \mathcal{E} is constructed by iterating over all possible choices of $\varphi, V, \varphi', \varphi_1, \varphi_2, \dots, \varphi_k, \psi', \kappa$, as above, and pruning \mathcal{E} for each choice in one of the following ways:

- either pick i such that $U_i = \emptyset$ and remove all unary types that contain any $A_{\psi_i, W_i}^{\kappa_i}$ with $W_i \subseteq \text{var}(\psi_i) \cap \text{var}(\psi')$, $\kappa_i(\text{var}(\psi_i) \cap \text{var}(\psi')) = \{\ell+1\}$, and $\kappa_i(x) = \kappa(x)$ for all $x \in \text{var}(\psi_i) \setminus \text{var}(\psi')$;
- or remove all unary types that contain some $A_{\psi_i, U_i}^{\kappa_i}$ with $\kappa_i(\text{var}(\psi_i) \cap \text{var}(\psi')) = \{\ell+1\}$ and $\kappa_i(x) = \kappa_i(x)$ for all $x \in \text{var}(\psi_i) \setminus \text{var}(\psi')$, for each i such that $U_i \neq \emptyset$, but do not contain $A_{\varphi, V}^{\kappa}$.

Lemma 11. *\mathcal{I} is a strongly ℓ -consistent level- ℓ model of \mathcal{K} modulo \mathcal{E} iff some interpretation that agrees with \mathcal{I} over all role names and all concept names except $\text{CN}_{\ell+1}^{\Phi}$ is an $(\ell+1)$ -consistent level- ℓ model of \mathcal{K} modulo \mathcal{E}' for some $(\ell+1)$ -reduct \mathcal{E}' of \mathcal{E} .*

Finite models consisting of multiple bags can be constructed bottom-up by a least fixed point procedure, using Lemma 11 to find each bag.

Lemma 12. *The (ℓ, ℓ) -model problem for an ALC KB \mathcal{K} , a UCRPQ Φ , and an environment \mathcal{E} can be solved in time*

$$2^{O(\|\mathcal{K}\|)} + 2^{\text{poly}(\|\Phi\|)}$$

given an oracle for the $(\ell, \ell+1)$ -model problem (with the same UCRPQ and TBox).

At the bottom of the recursion we need to check if there exists a $(n+1, n+1)$ -model for \mathcal{K} modulo \mathcal{E} . Now, a $\hat{\mathcal{B}}$ -decorated interpretation is level- $(n+1)$ iff it is *discrete*; that is, it has no edges at all. This allows solving the problem by a direct inspection. Because the ABox is trivial and ℓ -consistency is preserved under restrictions of the domain, it is enough to go through all singleton interpretations.

Lemma 13. *The $(n+1, n+1)$ -model problem for an ALC KB \mathcal{K} , a UCRPQ Φ , and an environment \mathcal{E} can be solved in time $2^{O(\|\mathcal{K}\|)} + 2^{\text{poly}(\|\Phi\|)}$.*

6 Incrementing the Level of Models

In this section we solve the (ℓ, ℓ') -model problem by reduction to multiple instances of the (ℓ', ℓ') -model problem for $\ell < \ell'$; that is, we eliminate level- ℓ edges from the interpretations. Like in Section 5, we rely on tree-like models of a special form; this time, however, they may be infinite and an additional step is needed to turn them into finite ones.

A $\hat{\mathcal{B}}$ -decorated interpretation is ℓ' -flat if it is a tree-like interpretation where all edges between bags have level strictly below ℓ' , whereas all edges inside bags have level at least ℓ' .

Lemma 14. *If there exists a finite (ℓ, ℓ') -model of \mathcal{K} modulo \mathcal{E} then there exists an ℓ' -flat (ℓ, ℓ') -model of \mathcal{K} modulo \mathcal{E} with bounded degree and bag size.*

In contrast to Lemma 8, the above only shows that the reformulated condition is necessary. We show that it is sufficient, by turning an arbitrary ℓ' -flat (ℓ, ℓ') -model of \mathcal{K} modulo \mathcal{E} with bounded degree and bag size into a finite (ℓ, ℓ') -model of \mathcal{K} modulo \mathcal{E} . For this we use *coloured blocking*. For $d \in \Delta^{\mathcal{I}}$, the m -neighbourhood $N_m^{\mathcal{I}}(d)$ of d is the interpretation obtained by restricting \mathcal{I} to elements $e \in \Delta^{\mathcal{I}}$ within distance m from d in \mathcal{I} , enriched with a fresh concept interpreted as $\{d\}$. A *colouring* of \mathcal{I} with k colours is an extension \mathcal{I}' of \mathcal{I} to k fresh concept names B_1, \dots, B_k such that $B_1^{\mathcal{I}'}, \dots, B_k^{\mathcal{I}'}$ is a partition of $\Delta^{\mathcal{I}'} = \Delta^{\mathcal{I}}$. We say that $d \in B_i^{\mathcal{I}'}$ has colour B_i . We call \mathcal{I}' *m-proper* if for each $d \in \Delta^{\mathcal{I}'}$ all elements of $N_m^{\mathcal{I}'}(d)$ have different colours.

Fact 1 (Gogacz, Ibáñez-García, and Murlak 2018). *If \mathcal{I} has bounded degree, then for all $m \geq 0$ there exists an m -proper colouring \mathcal{I}' of \mathcal{I} with finitely many colours. Consider interpretation \mathcal{J} obtained from \mathcal{I}' by redirecting some edges such that the old target and the new target have isomorphic m -neighbourhoods in \mathcal{I}' . Then, for each conjunctive query φ with at most \sqrt{m} binary atoms, if $\mathcal{I} \models \varphi$, then $\mathcal{J} \models \varphi$.*

Let \mathcal{I} be an ℓ' -flat (ℓ, ℓ') -model of \mathcal{K} modulo \mathcal{E} of bounded degree with bags of size at most M . In order to make Fact 1 applicable, we need to express the ℓ' -consistency condition over \mathcal{I} by means of a finite set of conjunctive queries, rather than CRPQs. Towards this end, we show that over \mathcal{I} each level- ℓ' CRPQ is equivalent to a UCQ. We rely on the following observation.

Lemma 15. *In a match of a $\hat{\mathcal{B}}$ -decorated CRPQ in a $\hat{\mathcal{B}}$ -decorated interpretation, each path witnessing an RPQ atom of end level at least ℓ' uses at most $n - \ell'$ edges of level strictly below ℓ' .*

We say that a CRPQ φ is *bounded* by K over an interpretation \mathcal{J} if for each match of φ in \mathcal{J} each RPQ atom of φ can be witnessed by a path of length at most K .

Lemma 16. *Let \mathcal{J} be $\widehat{\mathcal{B}}$ -decorated interpretation made up of disjoint level- ℓ' interpretations of size at most M connected by edges of level strictly below ℓ' . Assuming $\ell' \leq n$, each level- ℓ' CRPQ is bounded by $M(n - \ell' + 1)^2$ over \mathcal{J} .*

For a $\widehat{\mathcal{B}}$ -decorated CRPQ φ , let $\varphi^{(K)}$ be the UCQ obtained by taking the union of all CQs that can be obtained from φ by eliminating each RPQ atom $\mathcal{B}_{q,q'}(x, x')$ in one of the following ways: either remove the atom and equate variables x and x' , or replace the atom with a CQ of the form

$$r_1(x, y_1) \wedge r_2(y_1, y_2) \wedge \dots \wedge r_N(y_{N-1}, x')$$

where $N \leq K$, y_1, \dots, y_{N-1} are fresh variables, and there is a run of \mathcal{B} on $r_1 \dots r_N$ that begins in q and ends in q' .

Fact 2. *If a $\widehat{\mathcal{B}}$ -decorated CRPQ φ is bounded by K on an interpretation \mathcal{J} , then $\mathcal{J} \models \varphi$ iff $\mathcal{J} \models \varphi^{(K)}$.*

The final step before we can apply Fact 1 is to express ℓ' -consistency as query evaluation. Consider a partition of a fragment φ of Φ into a CRPQ φ' of level ℓ' and fragments $\varphi_1, \varphi_2, \dots, \varphi_k$ with $\text{var}(\varphi_i) \cap \text{var}(\varphi_j) = \emptyset$ for $i \neq j$, $V_i = \text{var}(\varphi_i) \cap \text{var}(\varphi')$, and $\emptyset \neq V \subseteq \text{var}(\varphi) \cap \text{var}(\varphi')$. Let ψ be the CRPQ obtained from φ' as follows. Begin from a copy of φ' . For each $i \in \{1, \dots, k\}$, add to ψ an atom $A_{\varphi_i, V_i}^{\kappa_i}(u)$ for some κ_i satisfying $\kappa_i(x) \leq \ell$ for all $x \in \text{var}(\varphi_i)$ and $\kappa_i(x) = \ell$ for all $x \in V_i$, and some variable u in V_i (V_i is nonempty, because φ is connected), and equate all variables in V_i . Similarly, add to ψ the atom $\bar{A}_{\varphi', V}(u)$ for some κ satisfying $\kappa(x) = \ell$ for all $x \in \text{var}(\varphi) \cap \text{var}(\varphi')$ and $\kappa(x) = \kappa_i(x)$ for all $x \in \text{var}(\varphi) \cap \text{var}(\varphi_i)$, and some $u \in V$, and equate all variables in V . Let $\Phi_{\ell'}$ be the union of all CRPQs ψ obtained as above for different choices of φ , φ' , $\varphi_1, \varphi_2, \dots, \varphi_k$, V , and $\kappa_1, \kappa_2, \dots, \kappa_k$. Note that $\Phi_{\ell'}$ is a union of level- ℓ' CRPQs. If $\ell' > n$, $\Phi_{\ell'}$ is a UCQ.

Lemma 17. *If \mathcal{J} is a $\widehat{\mathcal{B}}$ -decorated interpretation, then \mathcal{J} is ℓ' -consistent iff $\mathcal{J} \not\models \Phi_{\ell'}$.*

Let $K = M(n - \ell' + 1)^2$. Let t be the maximal number of binary atoms in one CQ in $\Phi_{\ell'}^{(K)}$. (Note that if $\ell' > n$, the query $\Phi_{\ell'}$ is a UCQ and $\Phi_{\ell'}^{(K)}$ coincides with $\Phi_{\ell'}$.) Fix $m = t^2$ and let \mathcal{I}' be an m -proper colouring of \mathcal{I} . On each infinite branch, select the first bag \mathcal{M} such that for some bag \mathcal{M}' higher on this branch, the m -neighbourhood of the target element e of the edge from the parent of \mathcal{M} to \mathcal{M} is isomorphic to the m -neighbourhood of the target e' of the edge from the parent of \mathcal{M}' to \mathcal{M}' . Because the number of non-isomorphic m -neighbourhoods in a structure of bounded degree is bounded, the depth of the selected bags in the tree of bags is also bounded. The set of selected bags is finite and forms a maximal antichain. Let \mathcal{F} be the interpretation obtained by taking the union of all strict ancestors of the selected bags, and for each element e as above, redirect the edge coming from the parent of \mathcal{M} to e' .

Clearly, \mathcal{F} is a finite level- ℓ interpretation. It is routine to check that $\mathcal{F} \models^{\mathcal{E}} \mathcal{K}$. It remains to prove that \mathcal{F} is ℓ' -consistent. We know that \mathcal{I} is ℓ' -consistent. By Lemma 17,

$\mathcal{I} \not\models \Phi_{\ell'}$. By Lemma 16 and Fact 2, $\mathcal{I} \not\models \Phi_{\ell'}^{(K)}$. By Fact 1, $\mathcal{F} \not\models \Phi_{\ell'}^{(K)}$. By construction, \mathcal{F} satisfies the assumptions of Lemma 16. Hence, by Lemma 16 and Fact 2, $\mathcal{F} \not\models \Phi_{\ell'}$. We conclude that \mathcal{F} is ℓ' -consistent using Lemma 17.

Thus we have proved the converse of Lemma 14.

Lemma 18. *If there exists an ℓ' -flat (ℓ, ℓ') -model of \mathcal{K} modulo \mathcal{E} with bounded degree and bag size then there exists a finite (ℓ, ℓ') -model of \mathcal{K} modulo \mathcal{E} .*

Combining Lemmas 6, 14, and 18, we get that there is a finite (ℓ, ℓ') -model of \mathcal{K} modulo \mathcal{E} iff there is a bounded-degree ℓ' -flat model of \mathcal{K} modulo \mathcal{E} whose bags are ℓ' -consistent and have bounded size. As in an ℓ' -flat model each bag is a level- ℓ' interpretation, when we restrict our search to one-bag models the problem is an instance of the (ℓ', ℓ') -model problem. Models consisting of multiple bags can be built coinductively top-down by means of a greatest fixed point algorithm (similar to type elimination), using the (ℓ', ℓ') -model problem to check if each bag exists.

Lemma 19. *The (ℓ, ℓ') -model problem for an \mathcal{ALC} KB \mathcal{K} , a UCRPQ Φ , and an environment \mathcal{E} can be solved in time*

$$2^{O(\|\mathcal{K}\|)} + 2^{\text{poly}(\|\Phi\|)}$$

given an oracle for the (ℓ', ℓ') -model problem (with the same UCRPQ and TBox).

7 Looking Forward (and Back)

This paper provides first positive results on finite entailment of navigational queries over DLs ontologies. The main technical contribution is an optimal automata-based 2EXPTIME upper bound for finite entailment of UCRPQs in \mathcal{ALC} .

Let us take a look back at our journey. We devised an expansion of the semiautomaton used to represent UCRPQs to keep track of its runs that begin in all possible states, on all infixes of the input word. By making interpretations and CRPQs knowledgeable of the runs of this expansion, we are able to associate levels to them as dictated by the transitions of the expansion. To solve the entailment problem, we use a recursive method eliminating the lowest level from the query and from the interpretation, and solving then the simpler problem. In particular, we look at problem of finding (ℓ, ℓ') models, and solve it by recursively increasing ℓ and ℓ' in an alternating way, until both reach a maximum level: Section 5 and 6 respectively address the increment of the query level and of the model. We finally showed what to do when $\ell = \ell' = n + 1$, which as argued, is enough to solve the original finite entailment problem.

As for future work, the first immediate step is to extend our method to deal with test atoms of the form $A?$, which are usually available in UCRPQs. For the ontology language, we believe our method can be adapted to allow inverses, nominals or counting. Regarding more expressive query languages, the natural next step is to consider *two-way* CRPQs. Our current approach relies on the fact that information only flows forward, and it is not clear whether it can be adapted to deal with queries that can go back.

Acknowledgments

This work was supported by Poland's National Science Centre grant 2018/30/E/ST6/00042. It also benefited from inspiring discussions with Charles Paperman.

References

- Amarilli, A., and Benedikt, M. 2020. Finite open-world query answering with number restrictions. *ACM Trans. Comput. Log.* 21(4):27:1–27:73.
- Amendola, G.; Leone, N.; and Manna, M. 2018. Finite controllability of conjunctive query answering with existential rules: Two steps forward. In *IJCAI*, 5189–5193. ijcai.org.
- Baader, F.; Horrocks, I.; Lutz, C.; and Sattler, U. 2017. *An Introduction to Description Logic*. Cambridge University Press.
- Baget, J.-F.; Leclère, M.; Mugnier, M.-L.; and Salvat, E. 2011. On rules with existential variables: Walking the decidability line. *Artificial Intelligence* 175(9):1620–1654.
- Bárány, V.; Gottlob, G.; and Otto, M. 2014. Querying the guarded fragment. *Log. Methods Comput. Sci.* 10(2).
- Bednarczyk, B., and Kieronski, E. 2022. Finite entailment of local queries in the \mathcal{Z} family of description logics. In *Proc. of AAAI-2022*.
- Bednarczyk, B., and Rudolph, S. 2019. Worst-case optimal querying of very expressive description logics with path expressions and succinct counting. In *IJCAI*, 1530–1536. ijcai.org.
- Bienvenu, M.; Calvanese, D.; Ortiz, M.; and Simkus, M. 2014. Nested regular path queries in description logics. In *KR*.
- Bienvenu, M.; Ortiz, M.; and Simkus, M. 2015. Regular path queries in lightweight description logics: Complexity and algorithms. *J. Artif. Intell. Res.* 53:315–374.
- Bojańczyk, M., and Parys, P. 2011. XPath evaluation in linear time. *J. ACM* 58(4):17:1–17:33.
- Calvanese, D.; Giacomo, G. D.; Lenzerini, M.; and Vardi, M. Y. 2000. Containment of conjunctive regular path queries with inverse. In *KR*, 176–185. Morgan Kaufmann.
- Calvanese, D.; Eiter, T.; and Ortiz, M. 2014. Answering regular path queries in expressive description logics via alternating tree-automata. *Inf. Comput.* 237:12–55.
- Chen, C.; Huang, H.; Ross, K. E.; Cowart, J. E.; Arighi, C. N.; Wu, C. H.; and Natale, D. A. 2020. Protein ontology on the semantic web for knowledge discovery. *Scientific Data* 7(1):337.
- Civili, C., and Rosati, R. 2012. A broad class of first-order rewritable tuple-generating dependencies. In *Datalog*, volume 7494 of *Lecture Notes in Computer Science*, 68–80. Springer.
- Cook, C. E.; Bergman, M. T.; Finn, R. D.; Cochrane, G.; Birney, E.; and Apweiler, R. 2016. The European Bioinformatics Institute in 2016: Data growth and integration. *Nucleic acids research* 44:D20–D26.
- Danielski, D., and Kieronski, E. 2019. Finite satisfiability of unary negation fragment with transitivity. In *MFCS*, volume 138 of *LIPIcs*, 17:1–17:15.
- Figueira, D.; Figueira, S.; and Baque, E. P. 2020. Finite controllability for ontology-mediated query answering of CRPQ. In *KR*, 381–391.
- Florescu, D.; Levy, A. Y.; and Suciu, D. 1998. Query containment for conjunctive queries with regular expressions. In *PODS*, 139–148. ACM Press.
- Galgonek, J.; Hurt, T.; Michlíková, V.; Onderka, P.; Schwarz, J.; and Vondrášek, J. 2016. Advanced SPARQL querying in small molecule databases. *Journal of Cheminformatics* 8(1):31.
- Ginzburg, A. 1968. *Algebraic Theory of Automata*. ACM Monograph Series. Academic Press.
- Glimm, B.; Lutz, C.; Horrocks, I.; and Sattler, U. 2008. Conjunctive query answering for the description logic SHIQ. *J. Artif. Intell. Res.* 31:157–204.
- Gogacz, T., and Marcinkowski, J. 2013. Converging to the chase - A tool for finite controllability. In *LICS*, 540–549. IEEE Computer Society.
- Gogacz, T.; Gutiérrez-Basulto, V.; Ibáñez-García, Y.; Jung, J. C.; and Murlak, F. 2019. On finite and unrestricted query entailment beyond SQ with number restrictions on transitive roles. In *IJCAI*, 1719–1725. ijcai.org.
- Gogacz, T.; Gutiérrez-Basulto, V.; Gutowski, A.; Ibáñez-García, Y.; and Murlak, F. 2020. On finite entailment of non-local queries in description logics. In *KR*, 424–433.
- Gogacz, T.; Ibáñez-García, Y.; and Murlak, F. 2018. Finite query answering in expressive description logics with transitive roles. In *Proc. of KR-18*.
- Gottlob, G.; Manna, M.; and Pieris, A. 2018. Finite model reasoning in hybrid classes of existential rules. In *IJCAI*, 1831–1837. ijcai.org.
- Gutiérrez-Basulto, V.; Ibáñez-García, Y. A.; and Jung, J. C. 2018. Answering regular path queries over SQ ontologies. In *AAAI*, 1845–1852. AAAI Press.
- Hu, W.; Qiu, H.; and Dumontier, M. 2015. Link analysis of life science linked data. In *ISWC* (2).
- Ibáñez-García, Y. A.; Lutz, C.; and Schneider, T. 2014. Finite model reasoning in Horn description logics. In *KR*.
- Kostylev, E. V.; Reutter, J. L.; and Vrgoc, D. 2014. XPath for DL-Lite ontologies. In *Description Logics*.
- Lysenko, A.; Roznovát, I. A.; Saqi, M.; Mazein, A.; Rawlings, C. J.; and Auffray, C. 2016. Representing and querying disease networks using graph databases. *BioData Mining* 9(1):23.
- Ortiz, M., and Simkus, M. 2014. Revisiting the hardness of query answering in expressive description logics. In *RR*, volume 8741 of *Lecture Notes in Computer Science*, 216–223. Springer.
- Pratt-Hartmann, I. 2009. Data-complexity of the two-variable fragment with counting quantifiers. *Inf. Comput.* 207(8):867–888.

- Rajabi, E., and Sanchez-Alonso, S. 2021. Knowledge discovery using SPARQL property path: The case of disease data set. *Journal of Information Science* 47(5):677–687.
- Rosati, R. 2008. Finite model reasoning in DL-Lite. In *ESWC*.
- Rudolph, S. 2016. Undecidability results for database-inspired reasoning problems in very expressive description logics. In *KR*, 247–257. AAAI Press.
- Schneider, T., and Simkus, M. 2020. Ontologies and data management: A brief survey. *Künstliche Intell.* 34(3):329–353.
- Stefanoni, G.; Motik, B.; Krötzsch, M.; and Rudolph, S. 2014. The complexity of answering conjunctive and navigational queries over OWL 2 EL knowledge bases. *J. Artif. Intell. Res.* 51:645–705.

A Additional Definitions

Let us fix $\mathcal{K} = (\mathcal{T}, \mathcal{A})$. A *homomorphism* from interpretation \mathcal{I} to interpretation \mathcal{J} , written as $h : \mathcal{I} \rightarrow \mathcal{J}$ is a function $h : \Delta^{\mathcal{I}} \rightarrow \Delta^{\mathcal{J}}$ that preserves roles, concepts, and individual names: that is, for all $r \in \mathbf{N}_R$, $(h(d), h(e)) \in r^{\mathcal{J}}$ whenever $(d, e) \in r^{\mathcal{I}}$, for all $A \in \mathbf{N}_C$, $h(d) \in A^{\mathcal{J}}$ whenever $d \in A^{\mathcal{I}}$, and $h(a) = a$ for all $a \in \text{ind}(\mathcal{A})$.

B Proofs

Lemma 1. *Given an oracle for finite entailment for trivial ABoxes, the general finite entailment $(\mathcal{T}, \mathcal{A}) \models_{\text{fin}} \Phi$ can be decided in time $2^{\text{poly}(\|\mathcal{T}, \mathcal{A}\|)} \cdot 2^{\text{poly}(\|\Phi\|)}$ using calls to the oracle for $\mathcal{K}' = (\mathcal{T}, \mathcal{A}')$ and Φ' consisting of $2^{\text{poly}(\|\Phi\|)}$ CRPQs of linear size over the same semiautomaton as Φ .*

Proof. Let \mathcal{I} be a finite model of $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ such that $\mathcal{I} \models \Phi$. We can think of \mathcal{A} as an interpretation with domain $\text{ind}(\mathcal{A})$. Then, \mathcal{I} contains a subinterpretation \mathcal{I}' that is an isomorphic copy of \mathcal{A} , except that the extension of concepts over $\text{ind}(\mathcal{A})$ is kept as in \mathcal{I} . Let \mathcal{J} be the interpretation obtained by starting from \mathcal{I}' and for each $a \in \text{ind}(\mathcal{A})$, adding an isomorphic copy \mathcal{I}_a of \mathcal{I} sharing only a with \mathcal{I}' . Clearly, \mathcal{J} is a model of \mathcal{K} and $\mathcal{J} \models \Phi$, because \mathcal{I} is a homomorphic image of \mathcal{J} . Note also, that $\mathcal{I}_a \models \mathcal{T}$ for all $a \in \text{ind}(\mathcal{A})$. This shows that it suffices to look for counter models that are unions of a *core* interpretation \mathcal{J}' that is a copy of \mathcal{A} up to the interpretations of concept names, and a collection of disjoint *peripheric* models \mathcal{J}_a of \mathcal{T} such that $\Delta^{\mathcal{I}_a} \cap \Delta^{\mathcal{J}'} = \{a\}$ for $a \in \text{ind}(\mathcal{A})$.

The algorithm iterates through all possible core models \mathcal{J}' . For each \mathcal{J}' it needs to decide if there exist peripheric models \mathcal{J}_a for $a \in \text{ind}(\mathcal{A})$ such that no partial match π of $\varphi \in \Phi$ in \mathcal{J}' can be extended to a full match of φ in the whole \mathcal{J} . For this it is enough to know if $(\mathcal{T}, \mathcal{A}') \models_{\text{fin}} \Phi'$ where \mathcal{A}' ranges over trivial ABoxes using a fixed individual a and concept names from $\text{CN}(\mathcal{K})$, and Φ' ranges over sets of CRPQs φ_U for $\varphi \in \Phi$ and $U \subseteq \text{var}(\varphi)$ defined as follows. The CRPQ φ_U is obtained from φ by

- dropping all atoms that involve no variable from U , as well as all edge atoms involving a variable not in U ;
- replacing each $\mathcal{B}_{q,q'}(x, x')$ such that $x \in U$ and $x \notin U$ with $\mathcal{B}_{q,p}(x, a)$ for some p' , and each $\mathcal{B}_{q,q'}(x, x')$ such that $x \notin U$ and $x' \in U$ with $\mathcal{B}_{p,q'}(a, x')$ for some p .

Note that $|\Phi'| = 2^{\text{poly}(\|\Phi\|)}$ but all CRPQs in Φ' have size bounded by $\max_{\varphi \in \Phi} |\varphi|$ and the underlying semiautomaton \mathcal{B} is not altered.

The number of possible choices of \mathcal{A}' and Φ' is

$$2^{\text{CN}(\mathcal{K})} \cdot 2^{2^{\text{poly}(\|\Phi\|)}}.$$

The number of distinguishable choices for each peripheric model \mathcal{I}_a is

$$2^{2^{\text{poly}(\|\Phi\|)}}.$$

This gives up to

$$2^{|\text{CN}(\mathcal{K})| \cdot |\text{ind}(\mathcal{A})|} \cdot 2^{2^{\text{poly}(\|\Phi\|)} \cdot |\text{ind}(\mathcal{A})|} = 2^{\text{poly}(\|\mathcal{K}\|) \cdot 2^{\text{poly}(\|\Phi\|)}}$$

choices for the algorithm. For each choice there are $|\Phi| \cdot |\text{ind}(\mathcal{A})|^{O(m)}$ partial matches to consider. The cost of verifying a single match is polynomial in the size of \mathcal{J}' and the size of a single Φ' ; that is, $\text{poly}(\|\mathcal{K}\|, 2^{\text{poly}(\|\Phi\|)})$. Overall, the complexity of the algorithm is

$$2^{2^{\text{poly}(\|\Phi\|)} \cdot \text{poly}(\|\mathcal{K}\|)}.$$

□

Lemma 2. *Let $\mathbf{p}_0, \mathbf{p}_1, \dots, \mathbf{p}_m$ be a run of $\widehat{\mathcal{B}}$ on w , and let q, q' be states of \mathcal{B} . There is a run of \mathcal{B} on w from q to q' iff there exist positions $0 \leq j_1 < j_2 < \dots < j_k = m$, levels $n \geq \ell_1 > \ell_2 > \dots > \ell_k \geq 1$, and states q_0, q_1, \dots, q_k with $1 \leq k \leq n$ such that $q_0 = q$, $q_k = q'$, and*

- the level of q_0 in \mathbf{p}_0 is ℓ_1 and the level of q_k in \mathbf{p}_m is ℓ_k ;
- for all $i \in \{1, 2, \dots, k-1\}$, the level of q_i in \mathbf{p}_{j_i} is ℓ_i and the level of $\delta(q_i, w[j_i+1])$ in $\mathbf{p}_{j_{i+1}}$ is ℓ_{i+1} ;
- for all $i \in \{1, 2, \dots, k\}$, each transition taken in the segment of the run from $\mathbf{p}_{j_{i-1}+1}$ (or \mathbf{p}_0 for $i = 1$) to \mathbf{p}_{j_i} has level at least ℓ_i .

Proof. (\Rightarrow) The run of \mathcal{B} from q to q' induces a thread in the run of $\widehat{\mathcal{B}}$. We can split the thread into segments that stay at the same level, giving levels $\ell_1, \ell_2, \dots, \ell_k$, separated by transitions that decrease the level. Clearly, $1 \leq k \leq n$. The last positions on the subsequent levels give j_1, j_2, \dots, j_k . It is easy to check that the corresponding states satisfy the conditions specified in the lemma.

(\Leftarrow) The first and third condition, combined with the fact that threads are non-increasing, imply that between indexes $j_i + 1$ and j_{i+1} the thread—from $\delta(q_i, w[j_i + 1])$ in $\mathbf{p}_{j_{i+1}}$ to q_{i+1} (or q' for $i = m - 1$) in $\mathbf{p}_{j_{i+1}}$ —stays at the same level; similarly for the prefix and suffix. Combined with the transitions mentioned in the second condition they give a single thread witnessing a run of \mathcal{B} from q to q' on w . □

Lemma 3. *Given \mathcal{B} one can compute in polynomial time a TBox $\widehat{\mathcal{T}}_{\mathcal{B}}$ such that $\mathcal{I} \models \widehat{\mathcal{T}}_{\mathcal{B}}$ iff \mathcal{I} is $\widehat{\mathcal{B}}$ -decorated.*

Proof. It is straightforward to express the condition that no element has incoming edges over different roles from $\text{rol}(\mathcal{K})$. Pick a fresh concept name A_r for each $r \in \text{rol}(\mathcal{K})$ and include axioms $\top \sqsubseteq \forall r.A_r$ and $A_r \sqcap A_s \sqsubseteq \perp$ for all $r, s \in \text{rol}(\mathcal{K})$ with $r \neq s$.

We next provide an alternative axiomatization of

$$C_{\mathbf{p}} \sqsubseteq \forall r.C_{\delta(\mathbf{p}, r)} \quad (1)$$

$$C_{\mathbf{p}} \sqcap C_{\mathbf{p}'} \sqsubseteq \perp \quad (2)$$

$$\top \sqsubseteq \bigsqcup_{\mathbf{p} \in \widehat{Q}} C_{\mathbf{p}} \quad (3)$$

To encode conditions (3) we include the following axioms for every $\ell \in \{1, \dots, n\}$.

$$\top \sqsubseteq \bigsqcup_{q \in Q} C_{q, \ell}$$

and

$$C_{q, \ell} \sqcap C_{q', \ell} \sqsubseteq \perp$$

with $q, q' \in Q$ such that $q \neq q'$. These, together with the following will enforce condition (2).

$$C_{q,\ell} \sqcap C_{q,\ell'} \sqsubseteq \perp$$

for each pair ℓ, ℓ' with $\ell \neq \ell'$ and each $q \in Q$.

To ensure that the transitions of $\hat{\mathcal{B}}$ are faithfully represented, we will use auxiliary concepts $A_{i,j}^r, D_{i,j}^r, B_i^r$ with $i, j \in \{1, \dots, n\}$, and r a role name.

Let $\mathbf{p} = (p_1, \dots, p_n) \in \hat{Q}$, and let $\delta(\mathbf{p}, r) = (p'_1, \dots, p'_n)$, for some arbitrary (but fixed) role name r . We will use concept $A_{i,j}^r$ to indicate that $\delta(qp_i, r) = p'_k$ for some $k \in \{1, 2, \dots, i\}$. Further, $D_{i,j}^r$ will indicate that $\delta(p_i, r) = p'_j$. Finally, B_ℓ^r will help to indicate that the level of \mathbf{p}' is equal to ℓ . More precisely, if an element d in the domain encodes the state \mathbf{p} , and the level of $\delta(\mathbf{p}, r)$ is ℓ , then every r -successor of d must satisfy B_ℓ^r (see (9)) below.

We have the following axioms:

$$\top \sqsubseteq A_{1,1}^r \quad (4)$$

For every triple q, ℓ, ℓ' with $q \in Q$ and $\ell, \ell' \in \{1, \dots, n\}$

$$C_{q,\ell} \sqcap A_{\ell,\ell'}^r \sqsubseteq \forall r. \left(\bigcup_{1 \leq k \leq \ell'} C_{\delta(q,r),k} \right) \quad (5)$$

$$C_{q,\ell} \sqcap \exists r. C_{\delta(q,r),\ell'} \sqsubseteq D_{\ell,\ell'}^r \quad (6)$$

For every ℓ, k with $1 \leq \ell, k < n$:

$$D_{\ell,k}^r \sqcap A_{\ell,k}^r \sqsubseteq A_{\ell+1,k+1}^r \quad (7)$$

$$D_{\ell,k'}^r \sqcap A_{\ell,k}^r \sqsubseteq A_{\ell+1,k}^r \quad \text{for every } k' < k \quad (8)$$

For every $k \in \{1, \dots, n\}$, and every $\ell < k$,

$$D_{n,\ell}^r \sqcap A_{n,k}^r \sqsubseteq \forall r. B_k^r \quad (9)$$

And for every $\ell < n$,

$$B_\ell^r \sqsubseteq \left(\bigcup_{i < j} (C_{q_i,\ell} \sqcap C_{q_j,\ell+1}) \right) \quad (10)$$

Finally, we require that for every $\ell, k, k' \in \{1, \dots, n\}$ such that $k \neq k'$:

$$A_{\ell,k}^r \sqcap A_{\ell,k'}^r \sqsubseteq \perp \quad (11)$$

$$D_{\ell,k}^r \sqcap D_{\ell,k'}^r \sqsubseteq \perp \quad (12)$$

$$B_k^r \sqcap B_{k'}^r \sqsubseteq \perp \quad (13)$$

Intuitively, the axioms encode the listing order of $(\delta(p_1), \dots, \delta(p_n))$ in $\hat{\delta}(\mathbf{p}) = (p'_1, \dots, p'_n)$ as follows. (4) encodes that the first position of the tuple is (the only) available for $\delta(p_1, r)$. Further, by (5) we have that if $q = p_\ell$ and the next available position for p_ℓ is ℓ' , then $\delta(p_\ell, r) = p'_k$ for some $1 \leq k \leq \ell'$, which means in particular that $\delta(p_1, r) = p'_1$.

As mentioned above, (6) is used for “recording” the level of $\delta(p_\ell, r)$ using the concept $D_{\ell,\ell'}^r$. That is, $D_{\ell,\ell'}^r$ holds whenever $\delta(p_\ell, r) = p'_{\ell'}$.

Clearly, if $\delta(p_\ell, r) = p'_k$, and its next available position was k , then the next available position for $\delta(p_{\ell+1})$ is $k+1$. This situation is captured by (7). On the other hand,

if $\delta(p_\ell, r)$ does not take position k (which is only possible if $\delta(p_\ell, r) = p'_{k'}$, with $k' < k$) then k is available for $\delta(p_{\ell+1}, r)$, as captured by (8).

Now, we need to account for the positions not taken by any $\delta(p_i, r)$. By the way the positions are taken, it is enough to record the smallest position unused. This information is encoded using the concept B_k^r . Thus, if $A_{n,k}^r$ and $D_{n,\ell}^r$, for some $\ell < k$, are both satisfied then k is the next available position for listing the remaining states ordered as in $Q = q_1, \dots, q_n$. This is captured by (9) and (10).

Finally, the role of CIs (11)–(13) is to ensure the consistency of the information encoded.

With this intuition in mind, it is not difficult to see (1) is faithfully encoded. \square

Lemma 4. Consider a level- ℓ interpretation \mathcal{I} and elements $e \in C_{q,\ell}^{\mathcal{I}}$ and $e' \in C_{q',\ell}^{\mathcal{I}}$. Then, $(e, e') \in \mathcal{B}_{q,q'}^{\mathcal{I}}$ iff there is a path from e to e' in \mathcal{I} .

Proof. (\Rightarrow) This is obvious because the definition of $\mathcal{B}_{q,q'}^{\mathcal{I}}$ requires a path from e to e' in \mathcal{I} .

(\Leftarrow) Consider a path from e to e' in \mathcal{I} and the corresponding run of the automaton $\hat{\mathcal{B}}$. We will focus on the thread starting in the state q . It starts on level ℓ because $e \in C_{q,\ell}^{\mathcal{I}}$ and cannot drop below level ℓ because \mathcal{I} is a level- ℓ interpretation. Hence, the thread ends in e' on level ℓ . But the state in e' on level ℓ is q' . Thus, this thread corresponds to a correct run of \mathcal{B} from q to q' . \square

Lemma 5. Let Φ be a UCRPQ, \mathcal{K} an ALC KB with a trivial ABox, and \mathcal{E} an environment.

- $\mathcal{I} \times \hat{\mathcal{B}}$ is a $\hat{\mathcal{B}}$ -decorated interpretation.
- If $\mathcal{I} \not\models \Phi$ then $\mathcal{I} \times \hat{\mathcal{B}} \not\models \Phi$.
- If $\mathcal{I} \models^{\mathcal{E}} \mathcal{K}$ then $\mathcal{I} \times \hat{\mathcal{B}} \models^{\mathcal{E}} \mathcal{K}$ up to identifying the unique individual a in \mathcal{K} with some $(a, r, \mathbf{p}) \in \Delta^{\mathcal{I} \times \hat{\mathcal{B}}}$.

Proof. $\mathcal{I} \times \hat{\mathcal{B}}$ is a $\hat{\mathcal{B}}$ -decorated interpretation by construction.

The mapping $(e, r, \mathbf{p}) \mapsto e$ from $\mathcal{I} \times \hat{\mathcal{B}}$ to \mathcal{I} is a homomorphism. Should some $\phi \in \Phi$ be matched in $\mathcal{I} \times \hat{\mathcal{B}}$, one could compose the match with the homomorphism above to obtain a match in \mathcal{I} .

Assume that $\mathcal{I} \models \mathcal{K}$. Satisfaction of the ABox transfers directly to $\mathcal{I} \times \hat{\mathcal{B}}$. Let us see that $\mathcal{I} \times \hat{\mathcal{B}}$ is a model of the TBox of \mathcal{K} . For CIs of the forms $\bigcap_i A_i \sqsubseteq \bigcup_j B_j$ and $A \sqsubseteq \forall r. B$ this follows from the existence of a homomorphic mapping from $\mathcal{I} \times \hat{\mathcal{B}}$ to \mathcal{I} , described above. For CIs of the form $A \sqsubseteq \forall r. B$ the reason is that the transition function of $\hat{\mathcal{B}}$ is defined for each state \mathbf{p} of $\hat{\mathcal{B}}$ and each $r \in \text{rol}(\mathcal{K})$, thus each r -edge originating in e will have its counterpart originating in (e, s, \mathbf{p}) for each s and \mathbf{p} . \square

Lemma 6. For every UCRPQ Φ and every consistent $\hat{\mathcal{B}}$ -decorated interpretation \mathcal{I} , if $A_{\varphi,V}^{\mathcal{I}} = \emptyset$ for each $\varphi \in \tilde{\Phi}$ and $\emptyset \neq V \subseteq \text{var}(\varphi)$, then $\mathcal{I} \not\models \Phi$.

Proof. By contradiction, suppose that $\mathcal{I} \models \Phi$. Then there exists a match π of some $\phi \in \tilde{\Phi}$ in \mathcal{I} . Take any element e in the image of π . The definition of consistency applied for

the trivial partition with $\varphi = \varphi'$ implies that $e \in A_{\varphi, V}^{\mathcal{I}}$ for $V = \pi^{-1}(a)$. \square

Lemma 7. *A tree-like interpretation is ℓ -consistent iff each of its bags and edge-bags is ℓ -consistent.*

Proof. Left to right implication is obvious. For right to left implication assume the contrary, that all bags and edge-bags are ℓ -consistent, but there is:

- a fragment φ of Φ ,
- a partition of φ into a CRPQ φ' of level ℓ and fragments $\varphi_1, \varphi_2, \dots, \varphi_k$, and sets V, V_1, V_2, \dots, V_k such that:
 - $\text{var}(\varphi_i) \cap \text{var}(\varphi_j) = \emptyset$ for $i \neq j$,
 - $V_i = \text{var}(\varphi_i) \cap \text{var}(\varphi')$,
 - $\emptyset \neq V \subseteq \text{var}(\varphi) \cap \text{var}(\varphi')$;
- a match π for φ' in \mathcal{I} and functions $\kappa, \kappa_1, \dots, \kappa_k$ such that:
 - $\pi(V_i) = \{e_i\} \subseteq (A_{\varphi_i, V_i}^{\kappa_i})^{\mathcal{I}}$ for all i ,
 - $\pi(V) = \{e\} \not\subseteq (A_{\varphi, V}^{\kappa})^{\mathcal{I}}$,
 - $\kappa_i(x) \leq \ell$ for all $x \in \text{var}(\varphi_i)$,
 - $\kappa(x) = \kappa_i(x)$ for all $x \in \text{var}(\varphi_i) \setminus V_i$,
 - $\kappa(x) = \ell$ for all $x \in \text{var}(\varphi) \cap \text{var}(\varphi')$.

For each RPQ atom $\mathcal{B}_{q, q'}(x, y)$ in φ' , choose a path from $\pi(x)$ to $\pi(y)$ witnessing that the atom is satisfied. Pick the parameters above and the witnessing paths for which π spans through the smallest number of bags and edge-bags (we count a bag if some edge atom in φ' is mapped by π to an edge of this bag, or if some witnessing path shares an edge with this bag). Note that the match of φ' given by π and the witnessing paths is connected. This is because the whole query φ is connected and because $\pi(V_i)$ consists of just one element for each i – the query φ' itself might not be connected, although it would be if we equated all variables in each V_i . The number of bags π spans through must be at least two: were it contained in one bag, this bag would be inconsistent. Essentially, we will show that we can derive the fact that $e \in (A_{\varphi, V}^{\kappa})^{\mathcal{I}}$ from ℓ -consistency conditions for some matches spanning through smaller number of bags.

Let b be the bag of e (not edge-bag, so it is unique). The match π necessarily spans through the bag b : otherwise no edge or RPQ is matched inside b , so φ' consists only of unary atoms, which means that $k = 1$, $\varphi_1 = \varphi$ and $e = e_1$, which easily leads to contradiction. Let $\psi', \psi_1, \dots, \psi_m$ be a partition of φ taking into account the bag b , match π and the chosen witnessing paths, where ψ_1, \dots, ψ_m are fragments. That is, in the definition of a partition:

- the initial set X' is the set of variables of φ which are mapped by π to the bag b (note that π is defined only on $\text{var}(\varphi')$);
- each RPQ is split (or not) in an appropriate way, depending on whether the corresponding witnessing path has zero, one, or two endpoints in the bag b , and the fresh variables are assigned level and state according to the last (or first) elements in b on the corresponding witnessing paths;

- the sets X_i are chosen in the way which results in ψ_i being fragments.

Let $U_i = \text{var}(\psi_i) \cap \text{var}(\psi')$ and π' be a match agreeing with π on $\text{var}(\psi') \cap \text{var}(\varphi')$, extended with fresh variables mapped to the appropriate elements on witnessing paths. The set $\pi'(U_i)$ consists of exactly one element, call it e'_i . Indeed, for some i there is j such that $\psi_i = \varphi_j$ and $U_i = V_j$, and so $\pi'(U_i) = \pi(V_j) = \{e_j\}$; and for other i , ψ_i consists of some part of φ' , possibly merged with some φ_j , with all the variables shared with ψ' being matched to an endpoint of an edge leaving or entering bag b .

We claim that for each i , $e'_i \in A_{\psi_i, U_i}^{\lambda_i}$ for some appropriate λ_i , such that we can use ℓ -consistency for b to show that $e \in (A_{\varphi, V}^{\kappa})^{\mathcal{I}}$.

To show that, we need to relate all φ_j to some ψ_i . Specifically, for each $i \in \{1, \dots, m\}$, consider all $j \in \{1, \dots, k\}$ such that $\text{var}(\psi_i) \cap \text{var}(\varphi_j) \neq \emptyset$. Note that variables in all $\psi', \psi_1, \dots, \psi_m$ are exactly the variables of φ , along with some fresh variables splitting some RPQs, and analogously for $\varphi', \varphi_1, \dots, \varphi_k$ (and each of the fragments φ_j shares at least one variable with φ). Since ψ' can be seen as a part of φ' (in particular, $\text{var}(\varphi) \cap \text{var}(\psi') \subseteq \text{var}(\varphi) \cap \text{var}(\varphi')$), it is easy to see that each j will be assigned to some i . For each i , ψ_i can be (yet again) partitioned into ψ'_i (intuitively being the common part of φ' and ψ_i) and φ_j for all j assigned to this i . There is also a match π_i agreeing with π on $\text{var}(\psi'_i) \cap \text{var}(\varphi')$, as usual, extended with fresh variables from U_i mapped to the appropriate elements on witnessing paths. Since this match spans through fewer bags than π (it does not span through the bag b , as all edges and parts of RPQs that were mapped inside b are in ψ'), the ℓ -consistency condition must be satisfied for this choice of parameters; since $\pi(V_j) = \{e_j\} \subseteq (A_{\varphi_j, V_j}^{\kappa_j})^{\mathcal{I}}$ for all j , we get that $e'_i \in A_{\psi_i, U_i}^{\lambda_i}$ for λ_i which agrees with κ_j on all $\text{var}(\varphi_j) \setminus V_j$, and is equal to ℓ on all other variables. Using this fact for all i and using ℓ -consistency in the bag b for the match π' , we get that $e \in (A_{\varphi, V}^{\kappa})^{\mathcal{I}}$ for κ which agrees with each κ_i on all variables from φ_i apart from V_i and is equal to ℓ on all other variables, which is exactly what needed to be shown. \square

Lemma 8. *There is a finite (ℓ, ℓ') -model of \mathcal{K} modulo \mathcal{E} iff there is a finite tree-like (ℓ, ℓ') -model of \mathcal{K} modulo \mathcal{E} whose bags are strongly connected.*

Proof. This is proved by routine unravelling. Let \mathcal{I} be a finite (ℓ, ℓ') -model of \mathcal{K} modulo \mathcal{E} . For each element d in \mathcal{I} define \mathcal{I}_d as the subinterpretation of \mathcal{I} obtained by restricting the domain of \mathcal{I} to the elements in the maximal strongly connected subset of $\Delta^{\mathcal{I}}$ that contains d . Recall that the ABox \mathcal{A} of \mathcal{K} is trivial. Let a be the unique element of $\text{ind}(\mathcal{A})$. Begin the construction of a tree-like model \mathcal{J} by taking a copy of \mathcal{I}_a . Then, as long as there exists an element e in \mathcal{J} and a CI $A \sqsubseteq \exists r.B$ in the TBox of \mathcal{K} such that $e \in A^{\mathcal{J}}$ but there is yet no $e' \in B^{\mathcal{J}}$ such that $(e, e') \in r^{\mathcal{J}}$, find the original d of e in \mathcal{I} and an element $d' \in B^{\mathcal{I}}$ such that $(d, d') \in r^{\mathcal{I}}$. Add to \mathcal{J} a copy of $\mathcal{I}_{d'}$ as a new bag, with an r -edge from e to the copy of d' . This construction gives a finite interpretation: the height of the tree of bags associated

to \mathcal{I} is bounded by the height of the DAG of strongly connected components of \mathcal{I} . It is straightforward to check that \mathcal{J} is a level- ℓ model of \mathcal{K} modulo \mathcal{E} . It is also clear that \mathcal{J} can be mapped homomorphically to \mathcal{I} by mapping each element of \mathcal{J} to its original in \mathcal{I} . Because ℓ' -consistency is defined in terms of forbidden matches it follows immediately that ℓ' -consistency of \mathcal{I} implies ℓ' -consistency of \mathcal{J} . \square

Lemma 9. *Over $\widehat{\mathcal{B}}$ -decorated interpretations, each level- ℓ CRPQ implies the union of its $(\ell+1)$ -reducts. Over strongly-connected level- ℓ interpretations, each level- ℓ CRPQ is equivalent to the union of its $(\ell+1)$ -reducts.*

Proof. Use Lemma 4 and the definition of reducts. \square

Lemma 11. *\mathcal{I} is a strongly ℓ -consistent level- ℓ model of \mathcal{K} modulo \mathcal{E} iff some interpretation that agrees with \mathcal{I} over all role names and all concept names except $\text{CN}_{\ell+1}^\Phi$ is an $(\ell+1)$ -consistent level- ℓ model of \mathcal{K} modulo \mathcal{E}' for some $(\ell+1)$ -reduct \mathcal{E}' of \mathcal{E} .*

Proof. The proof will use constructions that are very similar to the ones used in the proof of Lemma 7. Let us start with right to left implication. Let \mathcal{I} be some $(\ell+1)$ -consistent model of \mathcal{K} modulo \mathcal{E}' for some $(\ell+1)$ -reduct \mathcal{E}' of \mathcal{E} . We will show that the exact same model \mathcal{I} is a strongly ℓ -consistent model of \mathcal{K} modulo \mathcal{E} (note that strong ℓ -consistency does not mention concept names from $\text{CN}_{\ell+1}^\Phi$, so adjusting them is not necessary for this implication). Assume the contrary; so there is:

- a fragment φ ,
- a partition of φ into a CRPQ φ' of level ℓ and fragments $\varphi_1, \varphi_2, \dots, \varphi_k$, and sets V, V_1, V_2, \dots, V_k such that:
 - $\text{var}(\varphi_i) \cap \text{var}(\varphi_j) = \emptyset$ for $i \neq j$,
 - $V_i = \text{var}(\varphi_i) \cap \text{var}(\varphi')$,
 - $\emptyset \neq V \subseteq \text{var}(\varphi) \cap \text{var}(\varphi')$;
- a match π for some $(\ell+1)$ -reduct ψ' of φ' in \mathcal{I} and functions $\kappa, \kappa_1, \dots, \kappa_k$ such that:
 - $\pi(V_i) = \{e_i\} \subseteq (A_{\varphi_i, V_i}^{\kappa_i})^{\mathcal{I}}$ for all i ,
 - $\pi(V) = \{e\} \not\subseteq (A_{\varphi, V}^{\kappa})^{\mathcal{I}}$,
 - $\kappa_i(x) \leq \ell$ for all $x \in \text{var}(\varphi_i)$,
 - $\kappa(x) = \kappa_i(x)$ for all $x \in \text{var}(\varphi_i) \setminus V_i$,
 - $\kappa(x) = \ell$ for all $x \in \text{var}(\varphi) \cap \text{var}(\varphi')$.

For each RPQ atom $\mathcal{B}_{q, q'}(x, y)$ in ψ' , choose a path from $\pi(x)$ to $\pi(y)$ witnessing that the atom is satisfied.

When constructing \mathcal{E}' , the exact same parameters were considered (φ , its partition into $\varphi', \varphi_1, \varphi_2, \dots, \varphi_k$, $(\ell+1)$ -reduct ψ' of φ' , the set V and the function κ), sets U_1, \dots, U_m and fragments ψ_1, \dots, ψ_m were defined (note that for a fixed match π and witnessing paths, partial matches and partial witnessing paths for ψ_i can also be obtained, for the parts common with ψ'), and a choice was made:

- either pick i such that $U_i = \emptyset$ and remove all unary types that contain any $A_{\psi_i, W_i}^{\lambda_i}$ with $W_i \subseteq \text{var}(\psi_i) \cap \text{var}(\psi')$, $\lambda_i(\text{var}(\psi_i) \cap \text{var}(\psi')) = \{\ell+1\}$, and $\lambda_i(x) = \kappa(x)$ for all $x \in \text{var}(\psi_i) \setminus \text{var}(\psi')$;
- or remove all unary types that contain some $A_{\psi_i, U_i}^{\lambda_i}$ with $\lambda_i(\text{var}(\psi_i) \cap \text{var}(\psi')) = \{\ell+1\}$ and $\kappa(x) = \lambda_i(x)$ for all $x \in \text{var}(\psi_i) \setminus \text{var}(\psi')$, for each i such that $U_i \neq \emptyset$, but do not contain $A_{\varphi, V}^{\kappa}$.

For each $i \in \{1, \dots, m\}$ consider a partition of ψ_i into ψ'_i (intuitively being the common part of ψ_i and ψ') and fragments φ_j for all j such that $\text{var}(\varphi_j) \cap \text{var}(\psi_i) \neq \emptyset$, and a match π_i of ψ'_i agreeing with π and the choice of witnessing paths.

If the first choice was made, take the chosen i , choose any element e' in the image of $\text{var}(\psi_i) \cap \text{var}(\psi')$ under π_i , and let W_i be the preimage of e' under π_i . By $(\ell+1)$ -consistency, we see that $e' \in A_{\psi_i, W_i}^{\lambda_i}$, where λ_i agrees with κ_j (and κ) on $\text{var}(\varphi_j) \setminus V_j$ for all j such that $\text{var}(\varphi_j) \cap \text{var}(\psi_i) \neq \emptyset$, and is equal to $\ell+1$ on all other variables. However, all unary types containing this concept were forbidden in \mathcal{E}' , which is a contradiction with \mathcal{I} being a model of \mathcal{K} modulo \mathcal{E}' .

If the second choice was made, use $(\ell+1)$ -consistency for each ψ_i with $U_i \neq \emptyset$ and its partition as described above, which proves that $e \in A_{\psi_i, U_i}^{\lambda_i}$ for some λ_i agreeing with κ_j (and κ) on $\text{var}(\varphi_j) \setminus V_j$ for all j such that $\text{var}(\varphi_j) \cap \text{var}(\psi_i) \neq \emptyset$, and is equal to $\ell+1$ on all other variables. Thus, because of the environment \mathcal{E}' , $A_{\varphi, V}^{\kappa}$ must also be satisfied. (Note that the union of $\text{var}(\psi_i) \setminus \text{var}(\psi')$ for all i with $U_i \neq \emptyset$ is equal to the union of $\text{var}(\varphi_j) \setminus V_j$ for all φ_j which are not disjoint with all ψ_i considered here).

Now we will prove left to right implication. Assume that \mathcal{I} is a strongly ℓ -consistent model of \mathcal{K} modulo \mathcal{E} . Let \mathcal{I}' be an interpretation that agrees with \mathcal{I} over all role and concept names except $\text{CN}_{\ell+1}^\Phi$, and in which the interpretation of these concept names is *correct* in the following sense. First, for each $e \in (A_{\varphi, V}^{\kappa})^{\mathcal{I}'}$ with CN_ℓ^Φ we let $e \in (A_{\varphi, V}^{\kappa'})^{\mathcal{I}'}$ where $\kappa'(x) = \ell+1$ for $x \in V$ and $\kappa'(x) = \kappa(x)$ for $x \in \text{var}(\varphi) \setminus V$. Then, we add element e to $(A_{\varphi, V}^{\kappa})^{\mathcal{I}'}$ with $(A_{\varphi, V}^{\kappa})^{\mathcal{I}'} \in \text{CN}_{\ell+1}^\Phi$ if and only if there exists some partition of φ into $\varphi', \varphi_1, \dots, \varphi_k$, sets V_1, \dots, V_k and a match π for φ' in \mathcal{I}' with all requirements exactly as in the definition of $(\ell+1)$ -consistency, in which additionally $\kappa_i(x) \leq \ell$ for all $x \in \text{var}(\varphi_i) \setminus V_i$ and $\kappa_i(x) = \ell+1$ for all $x \in V_i$, for $i \in \{1, \dots, k\}$. Note that this does not leave many choices regarding the partition: all variables x such that $\kappa(x) = \ell+1$ must be in φ' , all other variables of φ must be outside φ' , so it is even known which RPQ atoms are split; only the levels and states of the fresh variables might differ between different partitions.

The interpretation \mathcal{I}' is $(\ell+1)$ -consistent, since any partition of some φ and a corresponding match (as in the definition of $(\ell+1)$ -consistency) for which $\kappa_i(V_i) = \{\ell+1\}$ for some i , can be “unwrapped” to ones satisfying the correctness condition, as follows. If $e_i \in A_{\psi_i, V_i}^{\kappa_i}$ and $\kappa_i(V_i) =$

$\{\ell + 1\}$, by the correctness condition, there is a partition of φ_i and an appropriate match witnessing that. This partition and match of φ_i can be “merged” into the partition and match of φ . Applying this procedure for all i such that $\kappa_i(V_i) = \{\ell + 1\}$ results in a partition and a match of φ to which the correctness condition can be applied.

Now we will construct an appropriate $(\ell + 1)$ -reduct \mathcal{E}' of \mathcal{E} . Consider some parameters $\kappa, \varphi, \varphi', \varphi_1, \dots, \varphi_k, \psi', V$, define $\psi_1, \dots, \psi_m, U_1, \dots, U_m$ as when constructing an $(\ell + 1)$ -reduct of the environment; recall that κ uses levels at most ℓ , and $\kappa(\text{var}(\varphi) \cap \text{var}(\varphi')) = \{\ell\}$. We need to choose one of the two options mentioned in the construction. If for some i such that $U_i = \emptyset$, the interpretation \mathcal{I}' does not violate the restrictions imposed by the first choice (that is, there are no elements in concept $A_{\psi_i, W_i}^{\kappa_i}$ in \mathcal{I}' for all W_i and κ_i as defined during the construction of a reduct), choose this option and this i . Otherwise, choose the second option, knowing that in \mathcal{I}' for each i such that $U_i = \emptyset$ there is some element $e_i \in A_{\psi_i, W_i}^{\kappa_i}$ for some W_i and κ_i such that $W_i \subseteq \text{var}(\psi_i) \cap \text{var}(\psi')$, $\kappa_i(\text{var}(\psi_i) \cap \text{var}(\psi')) = \{\ell + 1\}$, and $\kappa_i(x) = \kappa(x)$ for all $x \in \text{var}(\psi_i) \setminus \text{var}(\psi')$.

We claim that \mathcal{I}' is a model of \mathcal{K} modulo the environment \mathcal{E}' constructed as above. Suppose this is not the case. By the construction of \mathcal{E}' , the interpretation \mathcal{I}' satisfies all the requirements imposed by the first choice. Hence, it must be the case that \mathcal{I}' does not satisfy the requirements imposed by the second choice; that is, in \mathcal{I}' there is an element e in concepts $A_{\psi_i, U_i}^{\kappa_i}$ where $\kappa_i(\text{var}(\psi_i) \cap \text{var}(\psi')) = \{\ell + 1\}$ and $\kappa(x) = \kappa_i(x)$ for all $x \in \text{var}(\psi_i) \setminus \text{var}(\psi')$, for all i such that $U_i \neq \emptyset$, but not in $A_{\varphi, V}^{\kappa}$.

For each $i \in \{1, \dots, m\}$ and the witnessing element $e_i \in \left(A_{\psi_i, W_i}^{\kappa_i}\right)^{\mathcal{I}'}$ (if $U_i \neq \emptyset$, we let $e_i = e$ and $W_i = U_i$), take the partition and match π_i guaranteed by the correctness condition; all variables from $\text{var}(\psi_i) \cap \text{var}(\psi')$ are guaranteed to be in the domain of π_i . Merge all these matches together into one match π' . It is nearly a match of ψ' ; since κ_i determines for each i which RPQ atoms in ψ_i will be split in the partition, and it agrees with κ on this matter, we can identify the splitting variables in ψ' and in the domain of π_i ; but these variables ($\text{var}(\psi') \setminus \bigcup_i \text{var}(\psi_i)$) might have different states and levels assigned in π' ; since \mathcal{I} is a level- ℓ interpretation, the assigned levels are at most ℓ . Let $\tilde{\psi}'$ be an $(\ell + 1)$ -reduct of ψ' with states and levels of these variables adjusted to match the ones in π' . Let $\tilde{\pi}$ be π' adjusted to $\tilde{\psi}'$ (it is easy to modify a match for a CRPQ to a match of its $(\ell + 1)$ -reduct). We get that there is a partition of φ into $\tilde{\varphi}'$, $\tilde{\varphi}_1, \dots, \tilde{\varphi}_h$, sets $\tilde{V}_1, \dots, \tilde{V}_h$ and the match $\tilde{\pi}$ of an $(\ell + 1)$ -reduct $\tilde{\psi}'$ of $\tilde{\varphi}'$ to \mathcal{I}' , where $\tilde{\pi}(\tilde{V}_i) = \{\tilde{e}_i\}$, $\tilde{e}_i \in \left(A_{\tilde{\varphi}_i, \tilde{V}_i}^{\lambda_i}\right)^{\mathcal{I}'}$, and λ_i do not use level $\ell + 1$.

Using strong ℓ -consistency of \mathcal{I} with these parameters gives us that $e \in \left(A_{\varphi, V}^{\lambda}\right)^{\mathcal{I}}$ for some λ . One just needs to show that $\lambda(x) = \kappa(x)$ for all $x \in \text{var}(\varphi)$ to arrive at a contradiction. The matches π_i obtained from correctness condition for $e_i \in \left(A_{\psi_i, W_i}^{\kappa_i}\right)^{\mathcal{I}'}$ guarantee that $\kappa_i(x) =$

$\kappa(x)$ for all $x \in \text{var}(\psi_i) \setminus \text{var}(\psi')$. We also know that all $x \in \text{var}(\psi_i) \cap \text{var}(\psi')$ are in the domain of π_i , so (by the above use of strong ℓ -consistency) $\lambda(x) = \ell$ for $x \in \text{var}(\psi_i) \cap \text{var}(\psi')$, and $\lambda(x) = \kappa(x)$ for all other variables of φ , which is exactly what is needed. \square

Lemma 12. *The (ℓ, ℓ) -model problem for an \mathcal{ALC} KB \mathcal{K} , a UCRPQ Φ , and an environment \mathcal{E} can be solved in time*

$$2^{O(\|\mathcal{K}\|) + 2^{\text{poly}(\|\Phi\|)}}$$

given an oracle for the $(\ell, \ell + 1)$ -model problem (with the same UCRPQ and TBox).

Proof. By Lemma 10 it suffices to decide if there is a finite tree-like level- ℓ model of $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ modulo $\mathcal{E} = (\Theta, \varepsilon)$ whose edge-bags are ℓ -consistent and bags are strongly ℓ -consistent. Our algorithm will compute the set of unary types that are realizable in such interpretations of increasing height. Here, by the height of a tree-like interpretation we mean the number of edges on the longest path from the root bag to a leaf bag.

The algorithm begins from the empty set of types $\Phi_0 = \emptyset$. In round $h = 1, 2, \dots$, based on the set Φ_{h-1} it computes the set Φ_h of types that can be realized in models of height $h - 1$. The type τ is added to Φ_h iff there exists a finite strongly ℓ -consistent level- ℓ model of $\mathcal{K}_\tau = (\mathcal{T}, \mathcal{A}_\tau)$ modulo \mathcal{E}_h where

$$\mathcal{A}_\tau = \{A(a) \mid A \in \tau\}$$

for a designated $a \in \mathbb{N}_1$ and \mathcal{E}_h is defined based on \mathcal{E} and Φ_{h-1} as explained below; for the existence test we use Lemma 11.

For unary types τ_1, τ_2 and a role name r let $\mathcal{J}_{(\tau_1, r, \tau_2)}$ be the edge-bag built from an element e_1 of type τ_1 and an element e_2 of type τ_2 connected by an r -edge. We let $\mathcal{E}_h = (\Theta, \varepsilon_h)$ and include (r, B) in $\varepsilon_h(\tau_1)$ for $\tau_1 \in \Theta$ iff either $(r, B) \in \varepsilon(\tau_1)$ or there exists $\tau_2 \in \Theta_{h-1}$ such that $B \in \tau_2$ and $\mathcal{I}_{(\tau_1, r, \tau_2)}$ is an (ℓ, ℓ) -interpretation and satisfies all CIs of the form $A' \sqsubseteq \forall r.B'$ in \mathcal{T} . Recall that the first condition amounts to checking that $\mathcal{I}_{(\tau_1, r, \tau_2)}$ is $\hat{\mathcal{B}}$ -decorated ($\mathcal{I}_{(\tau_1, r, \tau_2)} \models \hat{\mathcal{T}}_{\mathcal{B}}$), level- ℓ , and ℓ -consistent.

Because the computed sets satisfy

$$\Phi_0 \subseteq \Phi_1 \subseteq \dots \subseteq \Phi_{h-1} \subseteq \Phi_h \subseteq \dots,$$

after at most $2^{\|\mathcal{CN}(\mathcal{K})\| + 2^{\text{poly}(\|\mathcal{K}\|)}}$ rounds the sets Φ_h stabilize. The algorithm should return yes iff the last Φ_h contains a unary type compatible with the assertions made by the ABox \mathcal{A} of \mathcal{K} about the unique individual it mentions. It is not hard to check that each round can be performed in time $2^{O(\|\mathcal{K}\|) + 2^{\text{poly}(\|\mathcal{K}\|)}}$, yielding the desired complexity upper bound. \square

Lemma 13. *The $(n + 1, n + 1)$ -model problem for an \mathcal{ALC} KB \mathcal{K} , a UCRPQ Φ , and an environment \mathcal{E} can be solved in time $2^{O(\|\mathcal{K}\|) + 2^{\text{poly}(\|\Phi\|)}}$.*

Proof. Let us first see when a discrete interpretation is $(n + 1)$ -consistent. Consider a partition of a fragment φ into $\varphi', \varphi_1, \dots, \varphi_k$ like in the definition of $(n + 1)$ -consistency.

Because φ' has level $n + 1$, it must be a UCQ. If φ' contains a binary atom, it cannot be matched in a discrete interpretation. Hence, we can assume that φ' contains no binary atoms. Fragments $\varphi_1, \dots, \varphi_k$ share no variables, but φ is connected, so $k = 1$. It follows that $\varphi_1 = \varphi$ and $V \subseteq \text{var}(\varphi') \subseteq V_1$. Then, $(n + 1)$ -consistency reduces to the condition $A_{\varphi, V_1}^\kappa \sqsubseteq A_{\varphi, V}^\kappa$ for all κ such that $\kappa(x) = n + 1$ for all $x \in V_1$. We can capture $(n + 1)$ -consistency of the model by replacing \mathcal{E} with the environment $\mathcal{E}' = (\Theta', \epsilon')$ obtained from \mathcal{E} by filtering out unary types that violate this condition. This can be done in time polynomial in the size of \mathcal{E} .

It remains to decide if there is a discrete model of \mathcal{K} modulo \mathcal{E}' . This is the case iff for the individual a mentioned in \mathcal{A} there is a type $\tau \in \Theta'$ compatible with the assertions on a in \mathcal{A} such that for each concept inclusion $A \sqsubseteq \exists r.B$ in \mathcal{T} if $A \in \tau$ then $(r, B) \in \epsilon'(\tau)$. This can be checked in time polynomial in the size of \mathcal{K} and \mathcal{E}' .

Overall, the existence of an $(n + 1, n + 1)$ -model of \mathcal{K} modulo \mathcal{E} can be decided in time polynomial in the size of \mathcal{K} and \mathcal{E} ; that is, in time $2^{O(\|\mathcal{K}\|) + 2^{\text{poly}(\|\Phi\|)}}$. \square

Lemma 14. *If there exists a finite (ℓ, ℓ') -model of \mathcal{K} modulo \mathcal{E} then there exists an ℓ' -flat (ℓ, ℓ') -model of \mathcal{K} modulo \mathcal{E} with bounded degree and bag size.*

Proof. This is also proved by routine unravelling, much like Lemma 8. The difference is that this time for \mathcal{I}_d we take the interpretation \mathcal{I} with all edges of level strictly below ℓ' removed. This unravelling procedure may pass through the same element multiple times on the same branch, so the resulting tree-like structure \mathcal{J} may be infinite. Because new bags are added to \mathcal{J} only when a witness is missing in the parent bag, it follows that all edges between bags have level strictly below ℓ' (all edges of level at least ℓ' are already copied in the parent bag, together with their targets). Hence, \mathcal{J} is ℓ' -flat. The size of each bag is equal to the size of \mathcal{I} and the degree within each bag is bounded by the maximal degree in \mathcal{I} . The number of child bags connected to the same element in the parent bag is bounded by the size of the TBox. Hence, the degree in \mathcal{J} is bounded. Checking that \mathcal{J} is a (ℓ, ℓ') -model of \mathcal{K} modulo \mathcal{E} is straightforward, just like in Lemma 8. \square

Lemma 15. *In a match of a $\widehat{\mathcal{B}}$ -decorated CRPQ in a $\widehat{\mathcal{B}}$ -decorated interpretation, each path witnessing an RPQ atom of end level at least ℓ' uses at most $n - \ell'$ edges of level strictly below ℓ' .*

Proof. Suppose a $\widehat{\mathcal{B}}$ -decorated CRPQ φ is matched in a $\widehat{\mathcal{B}}$ -decorated interpretation \mathcal{J} . Each path in \mathcal{J} corresponds to the run ρ of \widehat{B} obtained by reading the states decorating the elements on the path. Such a path witnesses an RPQ atom $\mathcal{B}_{q, q'}(x, x')$ iff the thread of ρ beginning in q ends in q' . If the atom has end level at least ℓ' , then the level of q' in the last state of ρ must be at least ℓ' . Observe however that each edge of level strictly below ℓ' brings all threads from levels ℓ' and higher at least one level down. Consequently, the path may use at most $n - \ell'$ edges of level strictly below ℓ' . \square

Lemma 16. *Let \mathcal{J} be $\widehat{\mathcal{B}}$ -decorated interpretation made up of disjoint level- ℓ' interpretations of size at most M connected by edges of level strictly below ℓ' . Assuming $\ell' \leq n$, each level- ℓ' CRPQ is bounded by $M(n - \ell' + 1)^2$ over \mathcal{J} .*

Proof. Consider a level- ℓ' CRPQ φ matched in \mathcal{J} . Consider a witnessing path e_0, e_1, \dots, e_k in \mathcal{J} . Let $\mathbf{p}_0, \mathbf{p}_1, \dots, \mathbf{p}_k$ be the run of $\widehat{\mathcal{B}}$ corresponding to the witnessing path and let $\ell_0, \ell_1, \dots, \ell_k$ be the thread in ρ that corresponds to the witnessing run q_0, q_1, \dots, q_k of \mathcal{B} . We have $\ell_0 \geq \ell_1 \geq \dots \geq \ell_k \geq \ell'$. Suppose that for some $i < j$ we have $e_i = e_j$ and $\ell_i = \ell_j = \ell'$. It follows immediately that $\mathbf{p}_i = \mathbf{p}_j$ and $q_i = q_j$. Thus, we can choose a shorter witnessing path by skipping $e_{i+1}, e_{i+2}, \dots, e_j$. Consequently, it is enough to look at witnessing paths that visit each element at most $(n - \ell' + 1)$ times. From the assumption on the structure of \mathcal{J} and from Lemma 15 it follows that every simple witnessing path has length at most $M(n - \ell' + 1)^2$. \square

Lemma 19. *The (ℓ, ℓ') -model problem for an \mathcal{ALC} KB \mathcal{K} , a UCRPQ Φ , and an environment \mathcal{E} can be solved in time*

$$2^{O(\|\mathcal{K}\|) + 2^{\text{poly}(\|\Phi\|)}}$$

given an oracle for the (ℓ', ℓ') -model problem (with the same UCRPQ and TBox).

Proof. By Lemma 14 it suffices to decide if there exists an ℓ' -flat (ℓ, ℓ') -model of \mathcal{K} modulo \mathcal{E} with bounded degree and bag size. Using Lemma 7 and the definition of ℓ' -flatness this amounts to deciding if there exists a (possibly infinite) tree-like model of \mathcal{K} modulo \mathcal{E} such that each bag is a finite (ℓ', ℓ') -interpretation, each edge-bag is an (ℓ, ℓ') -interpretation but not level- ℓ' , and the size of bags and the degree of elements is bounded.

The algorithm is similar to the one in Lemma 12. The main difference is that the model can now be infinite. Suppose for a while, however, that we are interested in computing only finite models. Then we can proceed just like in Lemma 12, computing sets

$$\emptyset = \Phi_0 \subseteq \Phi_1 \subseteq \Phi_2 \subseteq \dots$$

but as we are after ℓ' -consistent bags, rather than strongly ℓ' -consistent, we reduce directly to the (ℓ', ℓ') -model problem for \mathcal{K}_τ defined like before, and \mathcal{E}_h define almost like before, the difference being that we additionally require that the edge in $\mathcal{J}_{(\tau_1, r, \tau_2)}$ has level strictly below ℓ' . We do not need to do anything about the size of the bags and the degree, because in a finite interpretation these are always bounded.

In order to take into account also infinite models, we replace induction by co-induction. The algorithm proceeds just like described above but it starts from Φ_0 containing all unary types over $\text{CN}(\mathcal{K})$ and concepts $C_{q, k}$ and $A_{\psi, V}^\kappa$. It follows that

$$\Phi_0 \supseteq \Phi_1 \supseteq \dots \supseteq \Phi_{h-1} \supseteq \Phi_h \supseteq \dots$$

Like before, the sequence must stabilize after at most $2^{|\text{CN}(\mathcal{K})| + 2^{\text{poly}(\|\mathcal{K}\|)}}$ steps. We claim that the algorithm can answer yes iff the last computed Φ_h contains a type compatible with the ABox of \mathcal{K} . This is because one can build the

potentially infinite model top down plugging in as bags the models witnessing the addition of τ to Φ_h in the last round. Importantly, the number of these witnesses is finite, because the number of invoked instances of the (ℓ', ℓ') -model problem is finite. Consequently, the size of bags in the constructed model is bounded. The number of child bags attached to each element is bounded by the number of existential restrictions in the TBox of \mathcal{K} , so the degree in the constructed model is also bounded. The complexity bound follows like in Lemma 12. \square