Reasonable Highly Expressive Query Languages: Extended Technical Report

Pierre Bourhis[†] and **Markus Krötzsch**[‡] and **Sebastian Rudolph**[‡]

[†]CNRS CRIStAL UMR 9189   [‡]Technische Universität Dresden, Germany

pierre.bourhis@inria.fr   {markus.kroetzsch,sebastian.rudolph}@tu-dresden.de

Abstract

Expressive query languages are gaining relevance in knowledge representation (KR), and new reasoning problems come to the fore. Especially query containment is interesting in this context. The problem is known to be decidable for many expressive query languages, but exact complexities are often missing. We introduce a new query language, *guarded queries* (GQ), which generalizes most known languages where query containment is decidable. GQs can be nested (more expressive), or restricted to linear recursion (less expressive). Our comprehensive analysis of the computational properties and expressiveness of (linear/nested) GQs also yields insights on many previous languages.

1 Introduction

The significance of query languages in KR is twofold. On the one hand, evaluating queries in the presence of a background ontology allows us to express more complex information needs, leading to the notion of *ontology-based query answering*. This topic has been studied for a wide range of ontology languages and many different query languages, including conjunctive queries [Calvanese *et al.*, 2007b; Eiter *et al.*, 2009] and (many variants of) regular path queries [Calvanese *et al.*, 2007a; Calvanese *et al.*, 2009; Bienvenu *et al.*, 2014]. On the other hand, recursive queries can be used to "implement" reasoning, such that the query language plays the role of a logical calculus that computes subsumptions. Examples include [Xiao *et al.*, 2010; Krötzsch, 2011; Bischoff *et al.*, 2014], and the references given therein.

In both application areas, we can see a tendency towards more and more powerful recursive queries. Recent works introduced several highly expressive query languages related to applications in KR: *Monadically Defined Queries* (MQs) [Rudolph and Krötzsch, 2013] and *Monadic Disjunctive SNP* queries (coMMSNP) [Bienvenu *et al.*, 2013]. Both can be viewed as fragments of (disjunctive) Datalog.

The proliferation of query languages and their uses in KR brings new questions to the fore. The complexity of ontology-based query answering has been studied from its inception, whereas the equally important question of relative expressiveness was studied only recently [Bienvenu *et al.*, 2013]. An-

other important question is the problem of *query containment*, where we consider two queries Q_1 and Q_2, and asks if every answer to Q_1 is also an answer to Q_2 over all possible inputs. Deciding query containment is relevant for query rewriting algorithms, where it needs to be checked if new queries are subsumed by previous ones to ensure termination. Further relevant applications are query optimization (finding a simpler yet equivalent query), and verification (checking that a query contains specific test cases). In addition, query containment has a range of applications in databases, e.g. in information integration and database integrity checking.

Although Datalog provides a useful framework for studying many recursive query languages, it does, unfortunately, not have a decidable query containment problem [Shmueli, 1987]. In contrast, the containment is known to be decidable for regular path queries, MQs, and coMMSNP queries. In the latter two cases, however, no upper complexity bound is known. Nevertheless, numerous results exist for various smaller query languages. For the following overview, recall that a predicate in a Datalog program is *intensional (IDB)* if it occurs in some rule head, and *extensional (EDB)* otherwise.

Non-recursive Datalog and unions of conjunctive queries A non-recursive Datalog program is equivalent to a (possibly exponential) union of conjunctive queries (UCQ), and thus expressible in first-order logic. Containment of Datalog (in the following referred to as Dlog) in UCQ is 2ExpTime-complete, while containment of Dlog in non-recursive Datalog is 3ExpTime-complete [Chaudhuri and Vardi, 1997]. Some restrictions for decreasing the complexity of these problems have been considered. Deciding if a *linear* Datalog program (LinDlog, where rule bodies contain at most one recursive predicate) is contained in a UCQ is ExpSpace-complete; complexity further decreases to PSpace when the linear Datalog program is monadic (LinMDlog, see below) [Chaudhuri and Vardi, 1994; Chaudhuri and Vardi, 1997].

Monadic Datalog A monadic Datalog (MDlog) program is one with only unary IDB predicates. Containment between two MDlog programs is 2ExpTime-complete. The upper bound is well known since the 80's [Cosmadakis *et al.*, 1988], while the lower bound has been established only recently [Benedikt *et al.*, 2012]. Moreover, the containment of Dlog in MDlog is also decidable by a straightforward appli-

Figure 1: Query languages and complexities; languages higher up in the graph are more expressive

cation of [Courcelle, 1991, Theorem 5.5].[1] So far, however, tight bounds have not been known for this case.

Guarded Datalog Guarded Datalog (GDlog) requires that, for each rule, the variables of the head should appear in a single EDB atom in the rule body. This notion of (frontier-)guarded rules has been known for a while [Calì *et al.*, 2008; Baget *et al.*, 2011], but its first use as a query language is recent [Bárány *et al.*, 2012]. GDlog is a proper extension of MDlog, since monadic rules can always be rewritten into guarded rules [Bárány *et al.*, 2012]. Query containment for GDlog is 2EXPTIME-complete, as it corresponds to a satisfiability problem for guarded negation fixed point logic [Bárány *et al.*, 2011].

Navigational Queries Conjunctive two-way regular path queries (C2RPQs) generalize conjunctive queries (CQs) by regular expressions over binary predicates [Florescu *et al.*, 1998; Calvanese *et al.*, 2003]. Variants of this idea are used in the RDF query language SPARQL 1.1 and the XML query language XPath. Roughly, C2RPQ is a conjunction of atoms of the form xLy where L is a two-way regular expression. A pair of nodes $\langle n_1, n_2 \rangle$ is a valuation of the pair $\langle x, y \rangle$ if and only if there exists a path between n_1 and n_2 matching L. Containment of such queries is EXPSPACE-complete [Florescu *et al.*, 1998; Calvanese *et al.*, 2003; Abiteboul and Vianu, 1999; Deutsch and Tannen, 2002], while containment of Dlog in C2RPQ is 2EXPTIME-complete [Calvanese *et al.*, 2005].

Fragments of Monadic Second-Order Logic More recently, Monadically Defined Queries (MQs) and their nested version (MQ$^+$s) have been introduced [Rudolph and Krötzsch, 2013] as a proper generalization of MDlog that also captures (unions of) C2RPQs. MQs are expressible in both Dlog and monadic second-order logic, but (in contrast to these languages), feature a decidable query containment problem. The most general recent query language for which containment is known to be decidable is coMMSNP [Bienvenu *et al.*, 2013],

a fragment of monadic second-order logic motivated by descriptive complexity. As opposed to the above languages, coMMSNP is a non-deterministic query language, closely related to disjunctive Datalog. A simple inspection of the definitions shows that the deterministic (disjunction-free) fragment, i.e., "Horn-coMMSNP", agrees with MQ.

In this paper, we further extend the known recursive query languages and at the same time settle all major questions related to the complexity of their query containment problems. Figure 1 gives an overview of all languages we consider, together with their respective query-answering complexities.

The main new query language we consider is called *guarded queries* (GQ), and is based on the use of frontier-guarded Datalog rules. GQ can be viewed as an extension of MQ, and is indeed inspired by a similar extension for coMMSNP [Bienvenu *et al.*, 2013]. GQ thus also generalizes frontier-guarded Datalog. We further introduce the nested and linear variant of GQ, and establish complexity results for query answering in all cases.

We then turn towards query containment. We obtain tight complexity bounds for (nested) GQs and many other query languages, which are summarized in Table 1. To show the upper bounds, we extend known automata-based approaches by a number of new techniques. Lower bounds are obtained by simulating space-bounded alternating Turing machines in a way that allows for an exponential increase in space with each nesting level. Finally, we also sketch how our results transfer to the case of linear Datalog, where many complexities can be slightly reduced.

In summary, our results settle open problems for (nested) MQs, and they paint a comprehensive and detailed picture of the state of the art in Datalog query containment.

2 Preliminaries

We consider a standard language of first-order predicate logic, based on an infinite set **C** of *constant symbols*, an infinite set **P** of *predicate symbols*, and an infinite set **V** of first-

---

order *variables*. Each predicate $p \in \mathbf{P}$ is associated with a natural number $\mathsf{ar}(p)$ called the *arity* of p. The list of predicates and constants forms the language's *signature* $\mathscr{S} = \langle \mathbf{P}, \mathbf{C} \rangle$. We generally assume $\mathscr{S} = \langle \mathbf{P}, \mathbf{C} \rangle$ to be fixed, and only refer to it explicitly if needed.

Formulae, Rules, and Queries A *term* is a variable $x \in \mathbf{V}$ or a constant $c \in \mathbf{C}$. We use symbols s, t to denote terms, x, y, z, v, w to denote variables, a, b, c to denote constants. Expressions like $\boldsymbol{t}, \boldsymbol{x}, \boldsymbol{c}$ denote finite lists of such entities. We use the standard predicate logic definitions of *atom* and *formula*, using symbols φ, ψ for the latter.

Datalog queries are defined over an extended signature with additional predicate symbols, called *IDB predicates*; all other predicates are called *EDB predicates*. A *Datalog rule* is a formula of the form $\forall \boldsymbol{x}, \boldsymbol{y}.\varphi[\boldsymbol{x}, \boldsymbol{y}] \to \psi[\boldsymbol{x}]$ where φ and ψ are conjunctions of atoms, called the *body* and *head* of the rule, respectively, and where ψ only contains IDB predicates. We usually omit universal quantifiers when writing rules. Sets of Datalog rules will be denoted by symbols $\mathbb{P}, \mathbb{R}, \mathbb{S}$. A set of Datalog rules $\mathbb{P}$ is

- *monadic* if all IDB predicates are of arity one;
- *frontier-guarded* if the body of every rule contains an atom $p(\boldsymbol{t})$ such that p is an EDB predicate and $\boldsymbol{t}$ contains all variables that occur in the rule's head;
- *linear* if every rule contains at most one IDB predicate in its body.

A *conjunctive query* (CQ) is a formula $Q[\boldsymbol{x}] = \exists \boldsymbol{y}.\psi[\boldsymbol{x}, \boldsymbol{y}]$ where $\psi[\boldsymbol{x}, \boldsymbol{y}]$ is a conjunction of atoms; a *union of conjunctive queries* (UCQ) is a disjunction of such formulae. A *Datalog query* $\langle \mathbb{P}, Q \rangle$ consists of a set of Datalog rules $\mathbb{P}$ and a conjunctive query Q over IDB or EDB predicates (Q could be expressed as a rule in Datalog, but not in all restrictions of Datalog we consider). We write Dlog for the language of Datalog queries. A monadic Datalog query is one where $\mathbb{P}$ is monadic, and similarly for other restrictions. We use the query languages MDlog (monadic), GDlog (frontier-guarded), LinDlog (linear), and LinMDlog (linear, monadic).

Databases and Semantics We use the standard semantics of first-order logic (FOL). A *database instance* $\mathcal{I}$ over a signature $\mathscr{S} = \langle \mathbf{P}, \mathbf{C} \rangle$ consists of a set $\Delta^{\mathcal{I}}$ called *domain* and a function $\cdot^{\mathcal{I}}$ that maps constants $c \in \mathbf{C}$ to domain elements $c^{\mathcal{I}} \in \Delta^{\mathcal{I}}$ and predicate symbols $p \in \mathbf{P}$ to relations $p^{\mathcal{I}} \subseteq (\Delta^{\mathcal{I}})^{\mathsf{ar}(p)}$, where $p^{\mathcal{I}}$ is the *extension* of p.

Given a database instance $\mathcal{I}$ and a formula $\varphi[\boldsymbol{x}]$ with free variables $\boldsymbol{x} = \langle x_1, \ldots, x_m \rangle$, the *extension* of $\varphi[\boldsymbol{x}]$ is the subset of $(\Delta^{\mathcal{I}})^m$ containing all those tuples $\langle \delta_1, \ldots, \delta_m \rangle$ for which $\mathcal{I}, \{x_i \mapsto \delta_i \mid 1 \leq i \leq m\} \models \varphi[\boldsymbol{x}]$. We denote this by $\langle \delta_1, \ldots, \delta_m \rangle \in \varphi^{\mathcal{I}}$ or by $\mathcal{I} \models \varphi(\delta_1, \ldots, \delta_m)$; a similar notation is used for all other types of query languages. Two formulae $\varphi[\boldsymbol{x}]$ and $\psi[\boldsymbol{x}]$ are called *equivalent* if their extensions coincide for every database instance $\mathcal{I}$.

The set of answers of a UCQ $Q[\boldsymbol{x}]$ over $\mathcal{I}$ is its extension. A Datalog program $\mathbb{P}$ is *satisfied* by database instance $\mathcal{I}'$ over the extended signature of EDB and IDB predicates, if all rules of $\mathbb{P}$ are satisfied by $\mathcal{I}'$ in the usual sense. The set of answers of a Datalog query $\langle \mathbb{P}, Q \rangle$ over $\mathcal{I}$ is the intersection of the extensions of Q over all extended database instances $\mathcal{I}'$ that satisfy $\mathbb{P}$ and agree with $\mathcal{I}$ on constants and EDB predicates. Datalog can also be defined as the least fixpoint of the inflationary evaluation of Q on $\mathcal{I}$ [Abiteboul *et al.*, 1994].

Note that we do not require database instances to have a finite domain, since all of our results are valid in either case. This is due to the fact that every entailment of a Datalog program has a finite witness, and that all of our query languages are positive, i.e., that their answers are preserved under homomorphisms of database instances.

One of the main tasks on reasoning on queries is to determine if a query contains another. In particular, for two Datalog queries $\langle \mathbb{P}, Q \rangle$ and $\langle \mathbb{P}', Q' \rangle$, we say that $\langle \mathbb{P}, Q \rangle$ *is contained in* $\langle \mathbb{P}', Q' \rangle$, denoted $\langle \mathbb{P}, Q \rangle \sqsubseteq \langle \mathbb{P}', Q' \rangle$ iff for each database instance $\mathcal{I}$ over the signature of EDB predicates and constants, the set of answers of $\langle \mathbb{P}, Q \rangle$ over $\mathcal{I}$ is included in the set of answers of $\langle \mathbb{P}', Q' \rangle$ over $\mathcal{I}$.

3 Guarded Queries

Rudolph and Krötzsch [2013] introduced *monadically defined queries* (MQs[2]) as a generalization of conjunctive two-way regular path queries (C2RPQs) and monadic Datalog (MDlog) for which query containment is still decidable. The idea underlying this approach is that candidate query answers are checked by evaluating a monadic Datalog program, i.e., in contrast to the usual evaluation of Datalog queries, we start with a "guessed" answer that is the input to a Datalog program. To implement this, the candidate answer is represented by special constants λ that the Datalog program can refer to. This mechanism was called *flag & check*, since the special constants act as flags to indicate the answer to be checked.

Example 1. *A query that computes the transitive closure over a relation p can be defined as follows.*

$$p(\lambda_1, y) \to \mathsf{U}(y)$$
$$\mathsf{U}(y) \wedge p(y, z) \to \mathsf{U}(z)$$
$$\mathsf{U}(\lambda_2) \to \mathsf{hit}$$

One defines the answer of the query to contain all pairs $\langle \delta_1, \delta_2 \rangle$ for which the rules entail hit *when interpreting λ_1 as δ_1 and λ_2 as δ_2.*

The original approach used monadic Datalog for its close relationship to monadic second-order logic, which was the basis for showing decidability of query containment. In this work, however, we develop new techniques for showing the decidability (and exact complexity) of this problem directly. It is therefore suggestive to consider other types of Datalog programs for the "check" part. The next definition introduces the general approach for arbitrary Datalog programs, and defines interesting fragments by imposing further restrictions.

Definition 1. *Consider a signature $\mathscr{S}$. An FCP ("flag & check program") of arity m is a set of Datalog rules $\mathbb{P}$ with $k \geq 0$ IDB predicates $\mathsf{U}_1, \ldots, \mathsf{U}_k$, that may use the additional constant symbols $\lambda_1, \ldots, \lambda_m \notin \mathscr{S}$ and an additional nullary predicate symbol* hit. *An FCQ ("flag & check query") P is*

---

[2]Here we shorten the original acronym MODEQ to MQ.

of the form $\exists \mathbf{y}.\mathbb{P}(z)$, where $\mathbb{P}$ is an FCP of arity $|z|$ and all variables in $\mathbf{y}$ occur in z. The variables $\mathbf{x}$ that occur in z but not in $\mathbf{y}$ are the *free variables* of P.

Let $\mathcal{I}$ be a database instance over $\mathscr{S}$. The extension $\mathbb{P}^{\mathcal{I}}$ of $\mathbb{P}$ is the set of all tuples $\langle \delta_1, \ldots, \delta_m \rangle \in (\Delta^{\mathcal{I}})^m$ such that every database instance $\mathcal{I}'$ that extends $\mathcal{I}$ to the signature of $\mathbb{P}$ and that satisfies $\langle \lambda_1^{\mathcal{I}'}, \ldots, \lambda_m^{\mathcal{I}'} \rangle = \langle \delta_1, \ldots, \delta_m \rangle$ also entails hit. The semantics of FCQs is defined in the obvious way based on the extension of FCPs.

A GQ is an FCQ $\exists \mathbf{y}.\mathbb{P}(z)$ such that $\mathbb{P}$ is *frontier-guarded*. Similarly, we define MQ *(monadic)*, LinMQ *(linear, monadic)*, and LinGQ *(linear, frontier-guarded)* queries.

In contrast to Rudolph and Krötzsch [2013], we do not define monadic queries as conjunctive queries of FCPs, but we merely allow existential quantification to project some of the FCP variables. Proposition 1 below shows that this does not reduce expressiveness.

We generally consider monadic Datalog as a special case of frontier-guarded Datalog. Monadic Datalog rules do not have to be frontier-guarded. A direct way to obtain a suitable guard is to assume that there is a unary domain predicate that contains all (relevant) elements of the domain of the database instance. However, it already suffices to require *safety* of Datalog rules, i.e., that the variable in the head of a rule must also occur in the body. Then every element that is inferred to belong to an IDB relation must also occur in some EDB relation. We can therefore add single EDB guard atoms to each rule in all possible ways without modifying the semantics. This is a polynomial operation, since all variables in the guards are fresh, other than the single head variable that we want to guard. We therefore find, in particular, that GQ captures the expressiveness of MQ. The converse is not true, as the following example illustrates.

Example 2. *The following 4-ary LinGQ generalizes Example 1 by checking for the existence of two parallel p-chains of arbitrary length, where each pair of elements along the chains is connected by a relation q, like the steps of a ladder.*

$$q(\lambda_1, \lambda_2) \rightarrow \mathrm{U}_q(\lambda_1, \lambda_2)$$
$$\mathrm{U}_q(x, y) \wedge p(x, x') \wedge p(y, y'), q(x', y') \rightarrow \mathrm{U}_q(x', y')$$
$$\mathrm{U}_q(\lambda_3, \lambda_4) \rightarrow \mathrm{hit}$$

One might assume that the following MQ *is equivalent:*

$$q(\lambda_1, \lambda_2) \rightarrow \mathrm{U}_1(\lambda_1)$$
$$q(\lambda_1, \lambda_2) \rightarrow \mathrm{U}_2(\lambda_2)$$
$$\mathrm{U}_1(x) \wedge \mathrm{U}_2(y) \wedge p(x, x') \wedge p(y, y'), q(x', y') \rightarrow \mathrm{U}_1(x')$$
$$\mathrm{U}_1(x) \wedge \mathrm{U}_2(y) \wedge p(x, x') \wedge p(y, y'), q(x', y') \rightarrow \mathrm{U}_2(y')$$
$$\mathrm{U}_1(\lambda_3) \wedge \mathrm{U}_2(\lambda_4) \rightarrow \mathrm{hit}$$

However, the latter query also matches structures that are not ladders. For example, the following database yields the answer $\langle a, b, c, d \rangle$, although there is no corresponding ladder structure: $\{q(a, b), p(a, c), p(b, e), q(c, e), p(a, e'), p(b, d), q(e', d)\}$. One can extend the MQ *to avoid this case, but any such fix is "local" in the sense that a sufficiently large ladder-like structure can trick the query.*

Rudolph and Krötzsch [2013] showed that monadically defined queries can be expressed both in Datalog and in monadic second-order logic. While we lose the connection to monadic second-order logic with GQs, the expressibility in Datalog remains. The encoding is based on the intuition that the choice of the candidate answers for λ "contextualizes" the inferences of the Datalog program. To express this without special constants, we can store this context information in predicates of suitably increased arity.

Example 3. *The 4-ary LinGQ of Example 2 can be expressed with the following Datalog query. For brevity, let $\mathbf{y}$ be the variable list $\langle y_1, y_2, y_3, y_4 \rangle$, which provides the context for the IDB facts we derive.*

$$q(y_1, y_2) \rightarrow \mathrm{U}_q^+(y_1, y_2, \mathbf{y})$$
$$\mathrm{U}_q(x, y, \mathbf{y}) \wedge p(x, x') \wedge p(y, y'), q(x', y') \rightarrow \mathrm{U}_q^+(x', y', \mathbf{y})$$
$$\mathrm{U}_q(y_3, y_4, \mathbf{y}) \rightarrow \mathsf{goal}(\mathbf{y})$$

This result is obtained by a straightforward extension of the translation algorithm for MQ*s [Rudolph and Krötzsch, 2013], which may not produce the most concise representation. Also note that the first rule in this program is not safe, since y_3 and y_4 occur in the head but not in the body. According to the semantics we defined, such variables can be bound to any element in the active domain of the given database instance (i.e., they behave as if bound by a unary domain predicate).*

This observation justifies that we consider MQs, GQs, etc. as Datalog fragments. It is worth noting that the translation does not change the number of IDB predicates in the body of rules, and thus preserves linearity. The relation to (linear) Datalog also yields some complexity results for query answering; we will discuss these at the end of the next section, after introducing nested variants our query languages.

4 Nested Queries

Every query language gives rise to a nested language, where we allow for nested queries to be used as if they were predicates. Sometimes, this does not lead to a new query language (like for CQ and Dlog), but often it affects complexities and/or expressiveness. It has been shown that both are increased when moving from MQs to their nested variants [Rudolph and Krötzsch, 2013]. We will see that nesting also has strong effects on the complexity of query containment.

Definition 2. *We define k-nested FCPs inductively. A 1-nested FCP is an FCP. A k+1-nested FCP is an FCP that may use k-nested FCPs of arity m instead of predicate symbols of arity m in rule bodies. The semantics of nested FCPs is immediate based on the extension of FCPs. A k-nested FCQ P is of the form $\exists \mathbf{y}.\mathbb{P}(z)$, where $\mathbb{P}$ is a k-nested FCP of arity $|z|$ and all variables in $\mathbf{y}$ occur in z.*

A k-nested GQ *query is a k-nested frontier-guarded FCQ. For the definition of* frontier-guarded*, we still require EDB predicates in guards: subqueries cannot be guards. The language of k-nested* GQ *queries is denoted GQ^k; the language of arbitrarily nested* GQ *queries is denoted GQ^+. Similarly, we define languages MQ^k and MQ^+ (monadic), LinMQ^k and LinMQ^+ (linear, monadic), and LinGQ^k and LinGQ^+ (linear, frontier-guarded).*

Note that nested queries can use the same additional symbols (predicates and constants); this does not lead to any semantic interactions, however, as the interpretation of the special symbols is "private" to each query. To simplify notation, we assume that distinct (sub)queries always contain distinct special symbols. The relationships of the query languages we introduced here are summarized in Figure 1, where upwards links denote increased expressiveness. An interesting observation that is represented in this figure is that linear Datalog is closed under nesting:

Theorem 3. $\mathsf{LinDlog} = \mathsf{LinDlog}^+$.

Proof. We will prove that any $\mathsf{LinDlog}^+$ query can be rewritten into a $\mathsf{LinDlog}$ query of polynomial size. We make simplifying assumptions on the structure of the nested query which can be easily obtained by polynomial transformations and make the presentation easier: we assume that every rule body of any query occurring at any nesting depth contains at most one subquery atom (using, e.g., Proposition 1). Second, we assume that all variables and IDB predicates that are not in the same scope are appropriately renamed apart.

In order to proof our claim, we will first show that any $\mathsf{LinDlog}^2$ can be rewritten into an equivalent $\mathsf{LinDlog}$ query. Applying the rewriting iteratively inside-out (and observing that even manyfold application can be done in polynomial total time) then allows to conclude that there is a polynomial rewriting of any $\mathsf{LinDlog}^+$ query of arbitrary depth into a $\mathsf{LinDlog}$ query.

Consider a $\mathsf{LinDlog}^2$ query $P = \langle \mathbb{P}, p \rangle$ and assume w.l.o.g. that every rule body of the rules contains at most one $\mathsf{LinDlog}^1$ subquery. Now, going through all rules of $\mathbb{P}$ we produce the rules $\mathbb{P}'$ of the unnested but equivalent version.

Consider a rule $\rho \in \mathbb{P}$ having the shape

$$Q(x_1, \ldots, x_n) \wedge p(y_1, \ldots y_\ell) \wedge B_1 \wedge \ldots \wedge B_k \rightarrow H$$

where p is the body IDB predicate and where $Q = \langle \mathbb{Q}, q \rangle$ is a $\mathsf{LinDlog}^1$ query. For any k-ary IDB predicate r inside $\mathbb{Q}$ we increase its arity by ℓ and let $\mathbb{P}'$ contain all rules of $\mathbb{Q}'$ which is obtained from the rules ρ' of $\mathbb{Q}$ by

- replacing any (head or body) IDB atom $r(z_1, \ldots, z_k)$ of ρ' by $r(z_1, \ldots, z_k, y_1, \ldots y_\ell)$ and
- in case ρ' does not contain any IDB body atom, add $p(y_1, \ldots y_\ell)$ to the body.

Further we let $\mathbb{P}'$ contain the rule

$$q(x_1, \ldots, x_n, y_1, \ldots y_\ell) \wedge \wedge B_1 \wedge \ldots \wedge B_k \rightarrow H.$$

In case of a rule $\rho \in \mathbb{P}$ having the shape

$$Q(x_1, \ldots, x_n) \wedge B_1 \wedge \ldots \wedge B_k \rightarrow H$$

we add $\mathbb{Q}$ to $\mathbb{P}'$ without change and let $\mathbb{P}'$ contain the rule

$$q(x_1, \ldots, x_n) \wedge B_1 \wedge \ldots \wedge B_k \rightarrow H.$$

In case a rule $\rho \in \mathbb{P}$ does not contain a subquery atom we simply add ρ to $\mathbb{P}'$.

It can now easily verified that $\langle \mathbb{P}, p \rangle$ and $\langle \mathbb{P}', p \rangle$ are equivalent: first it is straightforward, that $\langle \mathbb{P}, p \rangle$ is equivalent to $\langle \mathbb{P}^\flat, p \rangle$ where $\mathbb{P}^\flat$ is obtained from $\mathbb{P}$ by replacing every

$Q(x_1, \ldots, x_n)$ by $q(x_1, \ldots, x_n)$ (that is, the according goal predicate) and then adding all rules from $\mathbb{Q}$ with no changes made to them. Second one can show that there is a direct correspondence between proof trees of $\langle \mathbb{P}^\flat, p \rangle$ and linearized proof trees of $\langle \mathbb{P}', p \rangle$ which yields the desired result. □

Another kind of nesting that does not add expressiveness is the nesting of FCQs in UCQs. Indeed, it turns out that (nested) FCQs can internalize arbitrary conjunctions and disjunctions of FCQs (of the same nesting level). This even holds when restricting to linear rules.

Proposition 1. *Let P be a positive query, i.e., a Boolean expression of disjunctions and conjunctions, of LinMQ^k queries with $k \geq 1$. Then there is a LinMQ^k query P' of size polynomial in P that is equivalent to P. Analogous results hold when replacing LinMQ^k by MQ^k, GQ^k, or LinMQ^k queries.*

Proof. We show the claim by induction, by expressing the innermost disjunctions and conjunctions of P with equivalent LinMQ^k queries of linear size. We consider positive queries without existential quantifiers (i.e., where all variables are answer variables), but the inner LinMQ^k may use existential quantifiers.

Let $P[\boldsymbol{x}] = P_1[\boldsymbol{x_1}] \vee \ldots \vee P_n[\boldsymbol{x_n}]$ be a disjunction of LinMQ^k queries. Each query P_i is of the form $\exists z_i.P'_i[\boldsymbol{x'_i}]$, where $\boldsymbol{x'_i}$ is the list of free variables of P'_i (corresponding to constants λ), and z_i contains exactly those variables of $\boldsymbol{x'_i}$ that do not occur in $\boldsymbol{x_i}$. We assume without loss of generality that z_i is disjoint from z_j if $i \neq j$, and that each P'_i uses a unique set of IDBs that does not occur in other queries. We consider queries $\bar{P}_i$ obtained by replacing the special constant that represents a variable $x_j \in \boldsymbol{x}$ by the special constant λ_j (assumed to not occur in P yet). Thus, the queries $\bar{P}_i$ share special constants exactly where queries P_1 share variables. We can now define the LinMQ^k P' as $\exists z_1 \ldots z_n.\bar{P}_1 \cup \ldots \cup \bar{P}_n$, where we assume that the correspondence of special constants to free variables is such that the existential quantifiers refer to the same variables as before.

Let $P[\boldsymbol{x}] = P_1[\boldsymbol{x_1}] \wedge \ldots \wedge P_n[\boldsymbol{x_n}]$ be a conjunction of LinMQ^k queries. Let $P_i = \exists z_i.P'_i[\boldsymbol{x'_i}]$ as before, and let U_i for $i \in \{1, \ldots, n-1\}$ be fresh IDB predicates. The queries $\bar{P}_i$ are defined as before by renaming special constants to reflect shared variables. For each $i \in \{1, \ldots, n\}$, the set of rules $\hat{P}_i$ is obtained from $\bar{P}_i$ as follows: if $i < n$, then every rule $\varphi \rightarrow \mathsf{hit} \in \bar{P}_i$ is replaced by the rule $\varphi \rightarrow \mathsf{U}_i(\lambda_1)$, where λ_1 is a fixed special constant in the queries; if $i > 1$, then every rule $\varphi \rightarrow \psi \in \bar{P}_i$ where φ does not contain an IDB predicate is replaced by the rule $\varphi \wedge \mathsf{U}_{i-1}(\lambda_1) \rightarrow \psi$, where λ_1 is as before. The LinMQ^k P' is defined as $\exists z_1 \ldots z_n.\hat{P}_1 \cup \ldots \cup \hat{P}_n$.

These constructions lead to equivalent LinMQ^k queries of linear size, so the claim follows by inductions. The cases for MQ^k, GQ^k, and LinMQ^k follow from the same constructions (note that, without the requirement of linearity, a simpler construction is possible in the case of conjunctions). □

Query answering for MQs has been shown to be NP-complete (combined complexity) and P-complete (data complexity). For MQ^+, the combined complexity increases to PSPACE while the data complexity remains the same. These

results can be extended to GQs. We also note the complexity for frontier-guarded Datalog, for which we are not aware of any published result.

Theorem 4. *The combined complexity of evaluating* GQ *queries over a database instance is* NP-*complete. The same holds for* GDlog *queries. The combined complexity of evaluating* GQ$^+$ *queries is* PSPACE-*complete. The data complexity is* P-*complete for* GDlog, GQ, *and* GQ$^+$.

Proof. The lower bounds are immediate from the matching complexities for MQ and MQ$^+$ queries, respectively [Rudolph and Krötzsch, 2013].

First, we prove that checking if a tuple is an answer of a GQ over a database instance $\mathcal{I}$ is in NP for combined complexity. Let $\mathcal{I}$ be an instance, let P be a GQ with frontier guarded rules $\mathbb{P}$, and let δ be be a candidate answer for P as in Definition 1.

Since each rule in $\mathbb{P}$ is frontier-guarded, each intentional fact that is derived when checking the answer follows from the application of one particular rule, instantiated to match one particular (guard) EDB fact in the body. Therefore, the number of IDB facts that can be derived is polynomially bounded in the size of $\mathcal{I}$ and $\mathbb{P}$.

Thus, for every derivation of $\mathbb{P}$, only a polynomial number of rule applications are necessary, since it is enough to derive each IDB fact once. It is clear that one can guess such a derivation, where we guess, for each derivable IDB fact, one specific rule instance by which it is derived. The correctness of this guess can be checked in polynomial time, showing that the problem can be solved in NP.

We now show that checking an answer of a GQ$^+$ over an instance $\mathcal{I}$ is in PSPACE. Let $\mathcal{I}$ be an instance, let P be a GQ^k with frontier guarded rules $\mathbb{P}$ (that may contain subqueries), and let δ be a candidate answer for P as in Definition 1. We demonstrate by induction on k that checking if δ is a solution for P w.r.t. $\mathcal{I}$ is in NPSPACE. For the induction base, the claim follows from the above result for GQs.

For the induction step, using the same argument as before, we can see that the number of IDB facts that can be derived by $\mathbb{P}$ is still polynomial. Therefore, we can again guess a polynomial derivation as before, though the rule instances now may refer to subqueries of smaller nesting depth. By the induction hypothesis, whenever we need to verify the applicability of such a rule, we can use an NPSPACE algorithm for the nested query. The overall number of such checks is polynomial, yielding the overall NPSPACE algorithm. The result follows since NPSPACE =PSPACE [Savitch, 1970].

The fact that query evaluation is in P for data complexity is immediate from the fact our queries can be expressed in Datalog, which is known to have this data complexity. A direct proof is also obtained by observing that the number of possible derivation sequences that the above algorithms need to consider is in itself polynomial in $\mathcal{I}$ if P is fixed, so that the algorithms themselves are already in P for data complexity. □

The lower bounds in the previous case are immediate from know results for MQs. In particular, the hardness proof for nested MQs also shows that queries of a particular fixed nesting level can encode the validity problem for quantified boolean formulae with a certain number of quantifier alternations; this explains why we show the combined complexity of MQ^k to be in the Polynomial Hierarchy in Figure 1. A modification of this hardness proof of Rudolph and Krötzsch [2013] allows us to obtain the same results for the combined complexities in the linear cases; matching upper bounds follow from Theorem 4.

Theorem 5. *The combined complexity of evaluating* LinMQ, LinGDlog, *or* LinGQ *queries over a database instance is* NP-*complete. The combined complexity of evaluating* LinMQ$^+$ *or* LinGQ$^+$ *queries is* PSPACE-*complete. The data complexity is* NLOGSPACE-*complete for all of these query languages.*

Proof. The claimed NP-completeness is immediate. Hardness follows from the hardness of CQ query answering. Membership follows from the membership of GQ.

The claimed membership in PSPACE follows from the PSPACE-membership of LinDlog; note that this uses Theorem 3. Hardness for LinGQ$^+$ follows from the hardness for LinMQ$^+$, which we show by modifying the PSPACE-hardness proof for monadically defined queries from [Rudolph and Krötzsch, 2013].

We show the result by providing a reduction from the validity problem of quantified Boolean formulae (QBFs). We recap that for any QBF, it is possible to construct in polynomial time an equivalent QBF that has the specific shape

$$Q_1 x_1 Q_2 x_2 \ldots Q_n x_n \bigvee_{L \in \mathcal{L}} \bigwedge_{\ell \in L} \ell,$$

with $Q_1, \ldots Q_n \in \{\exists, \forall\}$ and $\mathcal{L}$ being a set of sets of literals over the propositional variables $x_1, \ldots, x_n$. In words, we assume our QBF to be in prenex form with the propositional part of the formula in disjunctive normal form. For every literal set $L = \{x_{k_1}, \ldots, x_{k_i}, \neg x_{k_{i+1}}, \ldots, \neg x_{k_j}\}$, we now define the n-ary FCP $p_L = \{t(\lambda_{k_1}) \wedge \ldots \wedge t(\lambda_{k_i}) \wedge f(\lambda_{k_{i+1}}) \wedge \ldots \wedge f(\lambda_{k_j}) \to \text{hit}\}$. Moreover, we define the n-ary FCP $p_{\mathcal{L}} = \{p_L(\lambda_1, \ldots, \lambda_n) \to \text{hit} \mid L \in \mathcal{L}\}$. Letting $p_{\mathcal{L}} = p_n$ we now define FCPs $p_{n-1} \ldots p_0$ in descending order. If $Q_i = \exists$, then the $i-1$-ary FCP p_{i-1} is defined as the singleton rule set $\{p_i(\lambda_1, \ldots, \lambda_{i-1}, y) \to \text{hit}\}$. In case $Q_i = \forall$, we let p_{i-1} contain the rules

$$f(x) \to \text{U}_?(x)$$
$$\text{U}_!(x) \wedge f(x) \wedge t(y) \to \text{U}_?(y)$$
$$\text{U}_!(x) \wedge t(x) \to \text{hit}$$
$$\text{U}_?(x) \wedge p_i(\lambda_1, \ldots, \lambda_{i-1}, x) \to \text{U}_!(x)$$

Note that p_0 is a Boolean LinMQ$^+$ query the size of which is polynomial in the size of the input QBF.

Now, let D be the database containing the two individuals 0 and 1 as well as the facts $f(0)$ and $t(1)$. We now show that the considered QBF is true exactly if $D \models p_0()$. To this end, we first note that by construction the extension of p_L contains exactly those n-tuples $\langle \delta_1, \ldots, \delta_n \rangle$ for which the corresponding truth value assignment *val*, sending x_i to **true** iff $\delta_i = 1$, makes the formula $\bigwedge_{\ell \in L} \ell$ true. In the same way,

the extension of $p_{\mathcal{L}}$ represents the set of truth value assignments satisfying $\bigvee_{L \in \mathcal{L}} \bigwedge_{\ell \in L} \ell$. Then, by descending induction, we can show that the extensions of p_i encode the assignments to free propositional variables of the subformula $Q_{i+1} x_{i+1} \ldots Q_n x_n \bigvee_{L \in \mathcal{L}} \bigwedge_{\ell \in L} \ell$ that make this formula true. Consequently, p_0 has a nonempty extension if the entire considered QBF is true.

Finally, the NLogSpace-completeness for data complexity is again immediate, where the upper bound is obtained from LinDlog, and the lower bound follows from the well-known hardness of reachability queries, which can be expressed in LinMDlog. □

5 Complexity of Query Subsumption

In this section, we first discuss an automata-based way to decide query containment, which will give rise to upper complexity bounds. To obtain matching lower bounds, we provide direct encodings of Alternating Turing Machines (ATMs) with a fixed space bound [Chandra *et al.*, 1981]. Finally we show that containment checking is often slightly simpler in fragments of linear Datalog.

5.1 Deciding Query Containment with Automata

We first recall a general technique of reducing query containment to the containment problem for (tree) automata [Chaudhuri and Vardi, 1997], which we build our proofs on. An introduction to tree automata is included in the appendix.

A common way to describe the answers of a Dlog query $P = \langle \mathbb{P}, p \rangle$ is to consider its *expansion trees*. Intuitively speaking, the goal atom $p(\boldsymbol{x})$ can be rewritten by applying rules of $\mathbb{P}$ in a backward-chaining manner until all IDB predicates have been eliminated, resulting in a CQ. The answers of P coincide with the (infinite) union of answers to the CQs obtained in this fashion. The rewriting itself gives rise to a tree structure, where each node is labeled by the instance of the rule that was used in the rewriting, and the leaves are instances of rules that contain only EDB predicates in their body. The set of all expansion trees provides a regular description of P that we exploit to decide containment.

To formalize this approach, we describe the set of all expansion trees as a tree language, i.e., as a set of trees with node labels from a finite alphabet. The number of possible labels of nodes in expansion trees is unbounded, since rules are instantiated using fresh variables. To obtain a finite alphabet of labels, one limits the number of variables and thus the overall number of possible rule instantiations [Chaudhuri and Vardi, 1997].

Definition 6. *Given a* Dlog *query* $P = \langle \mathbb{P}, p \rangle$, $\mathcal{R}_{\mathbb{P}}$ *is the set of all instantiations of rules of* $\mathbb{P}$ *using only the variables* $\mathcal{V}_{\mathbb{P}} = \{v_1, \ldots, v_n\}$, *where* n *is twice the maximal number of variables occurring in any rule of* $\mathbb{P}$.

A proof tree *for* P *is a tree with labels from* $\mathcal{R}_{\mathbb{P}}$, *such that (a) the root is labeled by a rule with* p *as its head predicate; (b) if a node is labeled by a rule* ρ *with an IDB atom* B *in its body, then it has a child node that is labeled by* ρ' *with head atom* B. *The label of a node* e *is denoted* $\pi(e)$.

Consider two nodes e_1 *and* e_2 *in a proof tree with lowest common ancestor* e. *Two occurrences of a variable* v *in* $\pi(e_1)$

and $\pi(e_2)$ *are connected if* v *occurs in the head of* $\pi(f)$ *for all nodes* f *on the shortest path between* e_1 *and* e_2, *with the possible exception of* e.

A proof tree encodes an expansion tree where we replace every set of mutually connected variable occurrences by a fresh variable. Conversely, every expansion tree is represented by a proof tree that replaces fresh body variables by variables that do not occur in the head; this is always possible since proof trees can use twice as many variables as any rule of $\mathbb{P}$. The set of proof trees is a regular tree language that can be described by an automaton.

Proposition 2 (Proposition 5.9 [Chaudhuri and Vardi, 1997])**.** *For a* Dlog *query* $P = \langle \mathbb{P}, p \rangle$, *there is a tree automaton* $\mathcal{A}_P$ *of size exponential in* P *that accepts exactly the set of all proof trees of* P.

In order to use $\mathcal{A}_P$ to decide containment of P in another query P', we construct an automaton $\mathcal{A}_{P \sqsubseteq P'}$ that accepts all proof trees of P that are "matched" by P'. Indeed, every proof tree induces a *witness*, i.e., a minimal matching database instance, and one can check whether or not P' can produce the same query answer on this instance. If this is the case for all proof trees of P, then containment is shown.

5.2 Deciding Guarded Query Containment

Our first result provides the upper bound for deciding containment of GQ queries. In fact, the result extends to arbitrary Dlog queries on the left-hand side.

Theorem 7. *Containment of* Dlog *queries in* GQ *queries can be decided in* 3ExpTime.

To prove this, we need to construct the tree automaton $\mathcal{A}_{P \sqsubseteq P'}$ for an arbitrary GQ P'. As a first step, we construct an alternating 2-way tree automaton $\mathcal{A}_{P \sqsubseteq P'}^+$ that accepts the proof trees that we would like $\mathcal{A}_{P \sqsubseteq P'}$ to accept, but with nodes additionally being annotated with information about the choice of λ values to guide the verification.

We first construct automata to verify the match of a single, non-recursive rule that may refer to λ constants. The rule does not have to be monadic or frontier-guarded. Our construction is inspired by a similar construction for CQs by Chaudhuri and Vardi [Chaudhuri and Vardi, 1997], with the main difference that the answer variables in our case are not taken from the root of the tree but rather from one arbitrary node that is marked accordingly.

To define this formally, we introduce trees with additional annotations besides their node labels. Clearly, such trees can be viewed as regular labelled trees by considering annotations to be components of one label; our approach, however, leads to a more readable presentation.

Definition 8. *Consider a Datalog program* $\mathbb{P}$, *a rule* $\rho = \varphi \to p(\boldsymbol{x})$, *and* $n \geq 0$ *special constants* $\boldsymbol{\lambda} = \lambda_1, \ldots, \lambda_n$. *The proof-tree variables* $\mathcal{V}_{\mathbb{P}}$ *used in* $\mathcal{R}_{\mathbb{P}}$ *are as in Definition 6.*

A proof tree for $\mathbb{P}$ *is* $\boldsymbol{\lambda}$*-annotated if every node has an additional* $\boldsymbol{\lambda}$*-label that is a partial mapping* $\{\lambda_1, \ldots, \lambda_n\} \to \mathcal{V}_{\mathbb{P}}$, *such that: every special constant* λ_i *occurs in at least one* $\boldsymbol{\lambda}$*-label, and whenever a constant* λ_i *occurs in two* $\boldsymbol{\lambda}$*-labels, it is mapped to the same variable and both variable occurrences are connected.*

A proof tree for $\mathbb{P}$ is p-annotated if exactly one node has an additional p-label of the form $p(\boldsymbol{v})$, where $\boldsymbol{v}$ is a list of variables from $\mathcal{V}_{\mathbb{P}}$.

A matching tree T for ρ and $\mathbb{P}$ is a λ-annotated and p-annotated proof tree for $\mathbb{P}$ for which there is a mapping $\nu : \text{Var}(\rho) \cup \{\lambda_1, \ldots, \lambda_n\} \to \mathcal{V}_{\mathbb{P}}$ such that

1. *$\nu(p(\boldsymbol{x})) = p(\boldsymbol{v})$;*

2. *for every atom α of φ, there is a node e_α in T such that the rule instance that e_α is labeled with contains the EDB atom $\nu(\alpha)$ in its body;*

3. *if λ_i occurs in α, then the λ-label maps λ_i to the occurrence of $\nu(\lambda_i)$ in e_α;*

4. *if $\alpha, \alpha' \in \varphi$ share a variable x, then the occurrences of $\nu(x)$ in e_α and $e_{\alpha'}$ are connected.*

Proposition 3. *There is an automaton $\mathcal{A}_{P,\rho}$ that accepts exactly the annotated matching trees for ρ and $\mathbb{P}$, and which is exponential in the size of ρ and $\mathbb{P}$.*

Proof. We first construct an automaton $\mathcal{A}'_{P,\rho}$ that accepts matching trees where each node is additionally annotated by a partial mapping of the form $\text{Var}(\rho) \to \mathcal{V}_{\mathbb{P}}$ (called $\text{Var}(\rho)$-*label*), such that: every special variable $x \in \text{Var}(\rho)$ occurs in at least one $\text{Var}(\rho)$-label, and whenever a variable $x \in \text{Var}(\rho)$ occurs in two, it is mapped to the same variable and both variable occurrences are connected. Note that this is essentially the same condition that we imposed for λ-annotations.

The intersection of tree automata can be computed in polynomial time. We can therefore construct automata to check part of the conditions for (annotated) matching trees to simplify the definitions. We first construct an automaton $\mathcal{A}_x$ for checking the condition on $\text{Var}(\rho)$-labels for one variable $x \in \text{Var}(\rho)$. We define $\mathcal{A}_x = \langle \Sigma, Q_x, Q_x^s, \delta_x, Q_x^e \rangle$, where the alphabet Σ consists of quadruples of proof-tree labels (from $\mathcal{R}_{\mathbb{P}}$), λ-labels, p-labels, and $\text{Var}(\rho)$-labels. The state set Q_x is $\{a, b, \text{accept}\} \cup \{q_v \mid v \in \mathcal{V}_{\mathbb{P}}\}$, signifying that the current node is *a*bove the first node annotated with a mapping for x, *b*elow or *b*esides any nodes that were annotated with a mapping for x, or at a node where x is mapped to a variable v. That start-state set is $Q_x^s = \{a\} \cup \{q_v \mid v \in \mathcal{V}_{\mathbb{P}}\}$; the end-state set if $Q_x^e = \{\text{accept}\}$.

Consider a rule $\rho' \in \mathcal{R}_{\mathbb{P}}$ of the form $r_1(\boldsymbol{v_1}) \wedge \ldots \wedge r_n(\boldsymbol{v_n}) \wedge h_1(\boldsymbol{w_1}) \wedge \ldots \wedge h_m(\boldsymbol{w_m}) \to h(\boldsymbol{v})$, where r_i are EDB predicates and $h_{(i)}$ are IDB predicates. For the case that $m > 0$, there is a transition $\langle q_1, \ldots, q_m \rangle \in \delta(q, \langle \rho', \_, \_, v \rangle)$ exactly if the following conditions are satisfied:

- if $q = a$ and $\nu(x)$ is undefined, then $q_i = a$ for one $1 \le i \le m$ and $q_j = b$ for all $1 \le j \le m$ with $i \ne j$;

- if $q = q_v$ and $\nu(x) = v$, then $q_i = q_v$ for all $1 \le i \le m$ such that v occurs in $\boldsymbol{w_i}$ and $q_i = b$ for all other i;

- if $q = b$ and $\nu(x)$ is undefined, then $q_i = b$ for all $1 \le i \le m$.

For the case $m = 0$, there is a transition $\langle \text{accept} \rangle \in \delta(q, \langle \rho', \_, \_, v \rangle)$ exactly if:

- if $q = q_v$ and $\nu(x) = v$;

- if $q = b$ and $\nu(x)$ is undefined.

It is easy to check that the automaton $\mathcal{A}_x$ satisfies the required condition. Now an automaton for checking the condition on $\text{Var}(\rho)$-labels can be constructed as the intersection $\mathcal{A}'_{\text{Var}(\rho)} = \bigcap_{x \in \text{Var}(\rho)} \mathcal{A}_x$. The automaton $\mathcal{A}'_\lambda$ for checking the condition on λ-labels is constructed in a similar fashion. Likewise, an automaton $\mathcal{A}'_p$ for checking the condition on p-labels is easy to define.

It remains to construct an automaton for checking the conditions (a)–(d) of Definition 8. To do this, we interpret the $\text{Var}(\rho)$-labels and λ-labels as partial specifications of the required mapping ν. Condition (a) further requires that $\nu(\boldsymbol{x}) = \boldsymbol{v}$, i.e., that the $\text{Var}(\rho)$-label at the unique node annotated with $p(\boldsymbol{v})$ contains this mapping. It is easy to verify this with an automaton $\mathcal{A}'_{(a)}$. Together, $\mathcal{A}'_{(a)}$, $\mathcal{A}'_\lambda$, and $\mathcal{A}'_{\text{Var}(\rho)}$ provide a consistent variable mapping that respects the p-label (a) and the connectedness of variable occurrences, i.e., (c) and (d). To check the remaining condition (b), we use an automaton $\mathcal{A}'_{(b)}$.

The automaton for (b) will use auxiliary markers to record which atoms have been matched in the current node and how exactly this was done. We record such a match as a partial function from atoms $q(\boldsymbol{z}) \in \varphi$ to instances $q(\boldsymbol{w})$ of such atoms using variables $\boldsymbol{w} \subseteq \mathcal{V}_{\mathbb{P}}$. The set of all such partial functions is denoted $\text{Match}_{\varphi, \mathbb{P}}$. Note that this set is exponential (not double exponential).

We now define $\mathcal{A}'_{(b)} = \langle \Sigma, Q, Q_s, \delta, Q_e \rangle$ where Σ is as for $\mathcal{A}_x$ above. The set of states Q is $\{\text{accept}\} \cup (2^\varphi \times \text{Match}_{\varphi, \mathbb{P}})$, where elements from 2^φ encode the subset of φ that should be witnessed at or below the current node, and the elements from $\text{Match}_{\varphi, \mathbb{P}}$ encode atoms that must be matched at the current node with their respective instantiations. The start-state set Q_s is $\{\langle \varphi, \mu \rangle \mid \mu \in \text{Match}_{\varphi, \mathbb{P}}\}$; the end-state set Q_e is $\{\text{accept}\}$. The transition function δ is defined as follows. Consider a rule $\rho' \in \mathcal{R}_{\mathbb{P}}$ of the form $r_1(\boldsymbol{v_1}) \wedge \ldots \wedge r_n(\boldsymbol{v_n}) \wedge h_1(\boldsymbol{w_1}) \wedge \ldots \wedge h_m(\boldsymbol{w_m}) \to h(\boldsymbol{v})$, where r_i are EDB predicates and $h_{(i)}$ are IDB predicates. For the case $m > 0$, there is a transition $\langle \langle \beta_1, \mu_1 \rangle, \ldots, \langle \beta_m, \mu_n \rangle \rangle \in \delta(\langle \beta, \mu \rangle, \langle \rho', \nu_\lambda, \_, \nu_{\text{Var}(\rho)} \rangle)$ exactly if the set $\beta \subseteq \varphi$ can be partitioned into sets $\beta', \beta_1, \ldots, \beta_m$ such that $(\nu_\lambda \cup \nu_{\text{Var}(\rho)})(\beta') = \mu(\beta')$ and $\mu(\beta') \subseteq \{r_1(\boldsymbol{v_1}), \ldots, r_n(\boldsymbol{v_n})\}$. The element μ_i of successor states can be chosen freely; the validity of the choice will be checked later. For the case $m = 0$, there is a transition $\langle \text{accept} \rangle \in \delta(\langle \beta, \mu \rangle, \langle \rho', \nu_\lambda, \_, \nu_{\text{Var}(\rho)} \rangle)$ exactly if $(\nu_\lambda \cup \nu_{\text{Var}(\rho)})(\beta) = \mu(\beta)$ and $\mu(\beta) \subseteq \{r_1(\boldsymbol{v_1}), \ldots, r_n(\boldsymbol{v_n})\}$. In fact, the information from $\text{Match}_{\varphi, \mathbb{P}}$ is not strictly necessary to define the transition, since the relevant elements μ are always determined by other choices in the transition. However, having this information explicit will be important in later proofs.

The automaton $\mathcal{A}'_{P,\rho}$ is obtained as the intersection $\mathcal{A}'_{\text{Var}(\rho)} \cap \mathcal{A}'_\lambda \cap \mathcal{A}'_p \cap \mathcal{A}'_{(a)} \cap \mathcal{A}'_{(b)}$. It is easy to verify that it accepts exactly the $\text{Var}(\rho)$-annotated matching trees. Note that $\mathcal{A}'_{P,\rho}$ is exponential in size, already due to the exponentially large alphabet Σ. Now the required automaton $\mathcal{A}_{P,\rho}$ is obtained by "forgetting" the $\text{Var}(\rho)$-label in transitions of $\mathcal{A}'_{P,\rho}$. This projection operation for tree automata is possible with a polynomial increase in size: every state of $\mathcal{A}_{P,\rho}$ is a pair of a state of $\mathcal{A}'_{P,\rho}$ and a $\text{Var}(\rho)$-label; transitions of $\mathcal{A}_{P,\rho}$ are defined as

for $\mathcal{A}'_{P,\rho}$, but keeping $\mathsf{Var}(\rho)$-label information in states and introducing transitions for all possible $\mathsf{Var}(\rho)$-labels in child nodes. □

We want to use the automata $\mathcal{A}_{P,\rho}$ to verify the entailment of a single rule within a Datalog derivation. We would like an automaton to check whether a whole derivation is possible. Unfortunately, we cannot check these derivations using automata of the form $\mathcal{A}_{P,\rho}$, which each need to be run on a p-annotated tree which has the unique entailment of the rule marked. The length of a derivation is unbounded, and we would not be able to distinguish an unbounded amount of p-markers. To overcome this problem, we create a modified automaton $\mathcal{A}^+_{P,\rho,v}$ that simulates the behavior of $\mathcal{A}_{P,\rho}$ on a tree with annotation $p(v)$. For $\mathcal{A}^+_{P,\rho,v}$ to know which node the annotation $p(v)$ refers to, it has to be started at this node. This is a non-standard notion of run, where we do not start at the root of the tree. Moreover, starting in the middle of the tree makes it necessary to consider both nodes below and above the current position, and $\mathcal{A}^+_{P,\rho,v}$ therefore needs to be an *alternating 2-way tree automaton*.

Proposition 4. *There is an alternating 2-way tree automaton $\mathcal{A}^+_{P,\rho,v}$ that is polynomial in the size of $\mathcal{A}_{P,\rho}$ such that, whenever $\mathcal{A}_{P,\rho}$ accepts a matching tree T that has the p-annotation $p(v)$ on node e, then $\mathcal{A}^+_{P,\rho,v}$ has an accepting run that starts from the corresponding node e' on the tree T' that is obtained by removing the p-annotation from T.*

Proof. Using alternating 2-way automata, we can traverse a tree starting from any node, visiting each node once. To control the direction of the traversal, we create multiple copies of each state q: states q_{down} are processed like normal states in $\mathcal{A}_{P,\rho}$, states q_{up} use an inverted transition of $\mathcal{A}_{P,\rho}$ to move up the tree into a state $q_{\sigma,i}$; these auxiliary states are used to check that the label of the upper node is actually σ and to start new downwards processes for all child nodes other than the one (i) that we came from.

To ensure that the constructed automaton $\mathcal{A}^+_{P,\rho,v}$ simulates the behavior of $\mathcal{A}_{P,\rho}$ in case the annotation $p(v)$ is found, we eliminate all transitions that mention other p-annotations. Moreover, we assume without loss of generality that the states of $\mathcal{A}_{P,\rho}$ that allow a transition mentioning $p(v)$ cannot be left through any other transition; this can always be ensured by duplicating states and using them exclusively for one kind of transition. Let Q_p be the set of states of $\mathcal{A}_{P,\rho}$ that admit (only) transitions mentioning $p(v)$. Let $\mathcal{A}'_{P,\rho} = \langle \Sigma', Q, Q_s, \delta', Q_e \rangle$ denote the automaton over the alphabet Σ' of λ-annotated proof trees (without p-annotations), with the same (start/end) states as $\mathcal{A}_{P,\rho}$, and where δ' is defined based on the transition function δ of $\mathcal{A}_{P,\rho}$ as follows: $\delta'(\langle \rho', M \rangle)$ is the union of all sets of the form $\delta(\langle \rho', \lambda\text{-label}, p\text{-label} \rangle)$ where p-label is either $p(v)$ or empty. By this construction, there is a correspondence between the accepting runs of $\mathcal{A}_{P,\rho}$ over trees where one node e is annotated with $p(v)$ and accepting runs of $\mathcal{A}'_{P,\rho}$ (on trees without p-annotations) for which the node e is visited in some state of Q_p.

Let s be the maximal out-degree of proof trees for P, i.e., the maximal number of IDB atoms in bodies of P. The state

set Q^+ of $\mathcal{A}^+_{P,\rho,v}$ is given by the disjoint union $\{q_{\mathsf{up}} \mid q \in Q\} \cup \{q_{\sigma,i} \mid q \in Q, \sigma \in \Sigma, 1 \leq i \leq s\} \cup \{q_{\mathsf{down}} \mid q \in Q\} \cup \{\mathsf{start}, \mathsf{accept}\}$. The start-state set is $Q^+_s = \{\mathsf{start}\}$ and the end-state set is $Q^+_e = \{\mathsf{accept}\} \cup \{q_{\mathsf{down}} \mid q \in Q_e\}$.

Transitions of $\mathcal{A}^+_{P,\rho,v}$ are defined as follows:

- For all $\sigma \in \Sigma$, let $\delta^+(\mathsf{start}, \sigma)$ be the disjunction of all formulae $\langle 0, q_{\mathsf{up}} \rangle \wedge \langle 0, q_{\mathsf{down}} \rangle$ where $q \in Q_p$.

- For states q_{down} and $\sigma \in \Sigma$, let $\delta^+(q_{\mathsf{down}}, \sigma)$ be the disjunction of all formulae $\langle 1, q^1_{\mathsf{down}} \rangle \wedge \ldots \wedge \langle m, q^m_{\mathsf{down}} \rangle$ for which $\mathcal{A}'_{P,\rho}$ has a transition $\langle q^1, \ldots, q^m \rangle \in \delta'(q, \sigma)$.

- For states q_{up} and $\sigma \in \Sigma$, let $\delta^+(q_{\mathsf{up}}, \sigma)$ be the disjunction of all formulae $\langle -1, q'_{\sigma',i} \rangle$ for which $\mathcal{A}'_{P,\rho}$ has a transition $\langle q^1, \ldots, q^{i-1}, q, q^{i+1}, \ldots, q^m \rangle \in \delta'(q', \sigma')$ and the current node is the ith child of its parent (we can assume that this information is encoded in the labels σ, even for basic proof trees, which increases the alphabet only linearly; we omit this in our definitions since it would clutter all other parts of our proof without need).

- For states $q_{\sigma,i,q'}$, let $\delta^+(q_{\sigma,i,q'}, \sigma)$ be the disjunction of all formulae $\langle 0, q_{\mathsf{up}} \rangle \wedge \langle 1, q^1_{\mathsf{down}} \rangle \wedge \ldots \wedge \langle i-1, q^{i-1}_{\mathsf{down}} \rangle \wedge \langle i+1, q^{i+1}_{\mathsf{down}} \rangle \wedge \ldots \wedge \langle m, q^m_{\mathsf{down}} \rangle$ for which $\mathcal{A}'_{P,\rho}$ has a transition $\langle q^1, \ldots, q^{i-1}, q', q^{i+1}, q^m \rangle \in \delta'(q, \sigma)$.

- For all starting states $q \in Q_s$ of $\mathcal{A}'_{P,\rho}$ and $\sigma \in \Sigma$, let $\delta(q_{\mathsf{up}}, \sigma) = \langle 0, \mathsf{accept} \rangle$.

It is not hard to verify that $\mathcal{A}^+_{P,\rho,v}$ has the required properties. □

Using the automata $\mathcal{A}^+_{P,\rho,v}$, we can now obtain the claimed alternating 2-way automaton $\mathcal{A}^+_{P \sqsubseteq P'}$ for a GQ P'. Intuitively speaking, $\mathcal{A}^+_{P \sqsubseteq P'}$ concatenates the automata $\mathcal{A}^+_{P,\rho,v}$ using alternation: whenever a derivation requires a (recursive) IDB atom, a suitable process $\mathcal{A}^+_{P,\rho,v}$ is initiated, starting from a node in the middle of the tree. The construction relies on guardedness, which ensures that we can always find a suitable start node (corresponding to the node that was p-annotated earlier), by finding a suitable guard EDB atom in the tree.

Proposition 5. *For a Dlog query P and a GQ query P' with special constants λ, there is an alternating 2-way automaton $\mathcal{A}^+_{P \sqsubseteq P'}$ of exponential size that accepts the λ-annotated proof trees of P that encode expansion trees with λ assignments for which P' has a match.*

Proof. Let P' be the set $\{\rho_1, \ldots, \rho_\ell\}$. For every IDB predicate p, let P'_p denote the set of rules in P' with head predicate p (possibly hit). Without loss of generality, we assume that distinct rules use distinct sets of variables. For every frontier-guarded rule ρ', let $\mathsf{guard}(\rho')$ be a fixed EDB atom that acts as a guard in this rule, i.e., an atom that refers to all variables in the head of ρ'.

Consider a rule $\rho' \in P'$ with IDB atoms $q_1(t_1), \ldots, q_m(t_m)$ in its body. We construct new rules from ρ' by replacing each atom $q_i(t_i)$ with a guard atom $\mathsf{guard}(\rho'_i)$, suitably unified. Formally, assume that there are rules $\rho'_i \in P'_{q_i}$ with head $q_i(s_i)$ and a substitution θ that is a most general unifier for the problems $t_i\theta = s_i\theta$, for all $i \in \{1, \ldots, m\}$, and that maps

every variable in ρ'_i that does not occur in the head to a globally fresh variable. Then the *guard expansion* of ρ' for $(\rho'_i)_{i=1}^m$ and θ is the rule that is obtained from $\rho'\theta$ by replacing each body atom $q_i(\boldsymbol{t}_i)\theta$ by $\mathsf{guard}(\rho'_i)\theta$. By construction, two distinct atoms $\mathsf{guard}(\rho'_i)\theta$ and $\mathsf{guard}(\rho'_j)\theta$ do not share variables, unless at positions that correspond to head variables in rules ρ'_i and ρ'_j. The atoms $\mathsf{guard}(\rho'_i)\theta$ in a guard expansion are called *replacement guards*. We consider two guard expansions to be equivalent if they only differ in the choice of the most general unifier. Let $\mathsf{Guard}(\rho')$ be the set of all guard expansions of $\rho' \in P'$, i.e., a set containing one representative of each class of equivalent guard expansions. $\mathsf{Guard}(\rho')$ is exponential since there are up to $|P'|^m$ non-equivalent guard expansions for a rule with m IDB atoms.

The automaton $\mathcal{A}_{P\sqsubseteq P'}^+$ is constructed as follows. For every guard expansion $\rho_g \in \bigcup_{\rho'\in P'} \mathsf{Guard}(\rho')$ and every list $\boldsymbol{v}$ of proof-tree variables of the arity of the head of ρ_g, consider the alternating 2-way tree automaton $\mathcal{A}_{P,\rho_g,\boldsymbol{v}}^+$ of Proposition 4. We assume w.l.o.g. that the state sets of these automata are mutually disjoint. Let $\mathcal{A}_{P\sqsubseteq P'}^+ = \langle \Sigma, Q, Q_s, \delta, Q_e \rangle$. As before, Σ consists of pairs of a rule instance from $\mathcal{R}_{\mathbb{P}}$ and a partial mapping of λ to $\mathcal{V}_{\mathbb{P}}$. The state set Q is the disjoint union of all state sets of the automata of form $\mathcal{A}_{P,\rho_g,\boldsymbol{v}}^+$. The start-state set Q_s is the disjoint union of all start-state sets of automata $\mathcal{A}_{P,\rho_g,\boldsymbol{v}}^+$ for which ρ_g is a guard expansion of a rule with head hit (and $\boldsymbol{v}$ is the empty list). The end-state set Q_e is the disjoint union of all end-state sets of automata $\mathcal{A}_{P,\rho_g,\boldsymbol{v}}^+$.

The transition function δ is defined as follows. By the construction in Proposition 3, each state q in the automaton $\mathcal{A}_{P,\rho}$ encodes a partial mapping $\mathsf{match}(q)$ from body atoms of ρ to instantiated atoms that use variables from $\mathcal{V}_{\mathbb{P}}$, which are matched at the current tree node. This information is preserved through alphabet projections, intersections, and even through the construction in Proposition 4. We can therefore assume that each state q of $\mathcal{A}_{P\sqsubseteq P'}^+$ is associated with a partial mapping $\mathsf{match}(q)$.

For every state $q \in Q_{P,\rho_g,\boldsymbol{v}}$ and every $\sigma \in \Sigma$, we define $\delta(q,\sigma) = \delta_{P,\rho_g,\boldsymbol{v}}(q,\sigma) \wedge \psi$, where ψ defined as follows. For every replacement guard atom α of ρ_g for which $\mathsf{match}(q)(\alpha)$ is defined, we consider the formula $\psi_\alpha = \langle 0, q_1 \rangle \vee \ldots \vee \langle 0, q_\ell \rangle$, where

- $\alpha = \mathsf{guard}(\rho')\theta$ for some rule ρ' and substitution θ;
- $\mathsf{match}(q)(\alpha) = \alpha\theta'$ for some substitution θ';
- $q_1, \ldots q_\ell$ are the start states of the automaton $\mathcal{A}_{P,\rho',z\theta\theta'}$ where $p(\boldsymbol{z})$ is the head of ρ'.

Now ψ is the conjunction of all formulae ψ_α thus defined. $\quad\square$

We are now ready to prove Theorem 7. The automaton $\mathcal{A}_{P\sqsubseteq P'}^+$ allows us to check the answers of P' on a proof tree that is λ-annotated to assign values for answer constants. We can transform this alternating 2-way automaton into a tree automaton $\mathcal{A}'_{P\sqsubseteq P'}$ that is exponentially larger, i.e., doubly exponential in the size of the input. To remove the need for λ-labels, we modify the automaton $\mathcal{A}'_{P\sqsubseteq P'}$ so that it can only perform a transition from its start state if it finds that the constants in λ are assigned to the answer variables of P in the root. Finally, we obtain $\mathcal{A}_{P\sqsubseteq P'}$ by projecting to the alphabet

$\mathcal{R}_{\mathbb{P}}$ without λ-annotations; this is again possible in polynomial effort. The containment problem $P \sqsubseteq P'$ is equivalent by deciding the containment of $\mathcal{A}_P$ in $\mathcal{A}_{P\sqsubseteq P'}$, which is possible in exponential time w.r.t. the size of the automata. Since $\mathcal{A}_P$ is exponential and $\mathcal{A}_{P\sqsubseteq P'}$ is double exponential, we obtain the claimed triple exponential bound.

Our proof of Theorem 7 can be used to obtain another interesting result for the case of frontier-guarded Datalog. If P is a GDlog query, which does not use any special constants λ, then the λ-annotations are not relevant and $\mathcal{A}_{P\sqsubseteq P'}^+$ can be constructed as an alternating 2-way automaton on proof trees. For this, we merely need to modify the construction in Proposition 5 to start in start states of automata for rules that entail the goal predicate of P' with the expected binding of variables to answer variables of P. We can then omit the projection step, which required us to convert $\mathcal{A}_{P\sqsubseteq P'}^+$ into a tree automaton earlier. Instead, we can construct from $\mathcal{A}_{P\sqsubseteq P'}^+$ a complement tree automaton $\bar{\mathcal{A}}_{P\sqsubseteq P'}$ that is only exponentially larger than $\mathcal{A}_{P\sqsubseteq P'}^+$, i.e., doubly exponential overall [Cosmadakis *et al.*, 1988][Theorem A.1]. Containment can then be checked by checking the non-emptiness of $\mathcal{A}_P \cap \bar{\mathcal{A}}_{P\sqsubseteq P'}$, which is possible in polynomial time, leading to a 2ExpTime algorithm overall.

Theorem 9. *Containment of* Dlog *queries in* GDlog *queries can be decided in* 2ExpTime.

This generalizes an earlier result of Cosmadakis *et al.* [1988] for monadic Datalog, using an alternative, direct proof.

Finally, we can lift our results to the case of nested queries. Using Proposition 1, we can make the simplifying assumption that rules with some nested query in their body contain only one nested query and a guard atom as the only other atom. Thus all rules with nested queries have the form $g(\boldsymbol{s}) \wedge Q(\boldsymbol{t}) \rightarrow p(\boldsymbol{u})$, where g is an EDB predicate, Q is a nested query, and the variables $\boldsymbol{u}$ occur in $\boldsymbol{s}$.

In Proposition 4, we constructed alternating 2-way automata $\mathcal{A}_{P,\rho,\boldsymbol{v}}^+$ that can check the entailment of a particular atom $p(\boldsymbol{v})$ starting from a node within the tree. Analogously, we now construct automata $\mathcal{A}_{P,Q,\theta}^+$ that check that the nested query Q matches partially, where θ is a substitution that interprets query variables in terms of proof-tree variables on the current node of the tree. Only the variables that occur in $g(\boldsymbol{s})$ and $Q(\boldsymbol{t})$ are mapped by θ; the remaining variables can be interpreted arbitrarily, possibly in distant parts of the proof tree.

To construct $\mathcal{A}_{P,Q,\theta}^+$, we use the alternating 2-way automaton $\mathcal{A}_{P\sqsubseteq Q}^+$, constructed in Proposition 5 (assuming, for a start, that Q is not nested). This automaton is extended to an alternating 2-way automaton $\mathcal{A}_{P,Q}^+$ that accepts trees with a unique annotation of the form $\langle Q, \theta \rangle$, for which we check that it is consistent with the λ-annotation (i.e., for each query variable x mapped by θ, the corresponding constant λ is assigned to $\theta(x)$ at the node that is annotated with $\langle Q, \theta \rangle$). We then obtain a (top-down) tree automaton $\mathcal{A}_{P,Q}$ by transforming $\mathcal{A}_{P,Q}^+$ into a tree automaton (exponential), and projecting away the λ-annotations (polynomial). The automaton $\mathcal{A}_{P,Q}$ is analogous to the tree automaton $\mathcal{A}_{P,\rho}$ of Proposition 3. Using the same

transformation as in Proposition 4, we obtain an alternating 2-way automaton $\mathcal{A}_{P,Q,\theta}^{+}$ for each θ.

The automaton $\mathcal{A}_{P\sqsubseteq P'}^{+}$ for a nested query P' is constructed as in Proposition 5, but using the automata $\mathcal{A}_{P,Q,\theta}^{+}$ instead of automata $\mathcal{A}_{P,\rho,v}^{+}$ to check the entailment of a subquery Q. The size of $\mathcal{A}_{P\sqsubseteq P'}^{+}$ is increased by one exponential since the size of $\mathcal{A}_{P,Q,\theta}^{+}$ is exponentially increased when projecting out λ-labels for Q. Applying this construction inductively, we obtain the following result.

Theorem 10. *Containment of* Dlog *queries in* GQ^{k} *queries can be decided in* $(k+2)$ExpTime.

5.3 Simulating Alternating Turing Machines

To show the hardness of query containment problems, we generally provide direct encodings of Alternating Turing Machines (ATMs) with a fixed space bound [Chandra *et al.*, 1981]. To simplify this encoding, we assume without loss of generality that every universal ATM configuration leads to exactly two successor configurations. The following definition defines ATM encodings formally. Rather than requiring concrete structures to encode ATMs, we abstract the encoding by means of queries that find suitable structures in a database instance; this allows us to apply the same definition for increasingly complex encodings. The following definition is illustrated in Figure 2.

Definition 11. *Consider an ATM* $\mathcal{M} = \langle Q, \Sigma, \Delta, q_s, q_e \rangle$ *and queries* FirstConf$[x,y]$, NextConf$_\delta[x,y]$ *for all* $\delta \in \Delta$, LastConf$[x]$, State$_q[x]$ *for all* $q \in Q$, Head$[x,y]$, ConfCell$[x,y]$, FirstCell$[x,y]$, NextCell$[x,y]$, LastCell$[x]$, *and* Symbol$[x,y]$. *To refer to tape symbols, we consider constants* c_σ *for all* $\sigma \in \Sigma$, *and to refer to positions of the head, we use constants* h *(here),* l *(left), and* r *(right).*

With respect to these queries, an element $c \in \mathsf{dom}(\mathcal{I})$ *in a database instance* $\mathcal{I}$ *encodes an* $\mathcal{M}$ *quasi-configuration of size* s *if* $\mathcal{I}$ *contains a structure*

State$_q(c)$, FirstCell(c,d_1),
ConfCell(c,d_1), Symbol(d_1,c_{σ_1}), Head(d_1,p_1), NextCell(d_1,d_2),
ConfCell(c,d_2), Symbol(d_2,c_{σ_2}), Head(d_2,p_2), $\ldots$, NextCell(d_{s-1},d_s),
ConfCell(c,d_s), Symbol(d_s,c_{σ_s}), Head(d_s,p_s), LastCell(d_s),

where $q \in Q$, $\sigma_i \in \Sigma$, *and* $p_i \in \{h,l,r\}$. *We say that* c *encodes an* $\mathcal{M}$ *configuration of size* s *if, in addition, the sequence* $(p_i)_{i=1}^{s}$ *has the form* $l, \ldots, l, h, r, \ldots, r$ *with zero or more occurrences of* r *and* l, *respectively.*

An element c *in* $\mathcal{I}$ *encodes a (quasi-)configuration tree of* $\mathcal{M}$ *in space* s *if*

- $\mathcal{I} \models$ FirstConf(c,d_1) *for some* d_1,

- d_1 *is the root of a tree with edges defined by* NextConf$_\delta$,

- *every node in this tree encodes an* $\mathcal{M}$ *(quasi-) configuration of size* s,

- *if there is a transition* $\mathcal{I} \models$ NextConf$_{\delta_1}(e,e_1)$, *where* $\delta_1 = \langle q, \sigma, q', \sigma', d \rangle$ *and* q *is a universal state, then there is also a transition* $\mathcal{I} \models$ NextConf$_{\delta_2}(e,e_2)$ *with* $\delta_1 \neq \delta_2$,

- *if* e *is a leaf node, then* $\mathcal{I} \models$ LastConf(e).

If the tree is an accepting run, then c *encodes an accepting run (of* $\mathcal{M}$ *in space* s).

A same-cell query *is a query* SameCell$[x,y]$ *such that, if* $c_1, c_2 \in \mathsf{dom}(\mathcal{I})$ *encode two quasi-configurations, and* $d_1, d_2 \in \mathsf{dom}(\mathcal{I})$ *represent the same tape cell in the encodings* c_1 *and* c_2, *respectively, then* $\langle d_1, d_2 \rangle \in$ SameCell$^{\mathcal{I}}$.

Two queries $P_1[x]$ *and* $P_2[x]$ *containment-encode accepting runs of* $\mathcal{M}$ *in space* s *if, for every database instance* $\mathcal{I}$ *and element* $c \in P_1^{\mathcal{I}} \setminus P_2^{\mathcal{I}}$, c *encodes an accepting run of* $\mathcal{M}$ *in space* s, *and every accepting run of* $\mathcal{M}$ *in space* s *is encoded by some* $c \in P_1^{\mathcal{I}} \setminus P_2^{\mathcal{I}}$ *for some* $\mathcal{I}$.

Note that elements c may encode more than one configuration (or configuration tree). This is not a problem in our arguments.

The conditions that ensure that a quasi-configuration tree is an accepting run can be expressed by a query, based on the queries given in Definition 11. More specifically, one can construct a query that accepts all elements that encode a quasi-configuration sequence that is *not* a run. Together with a query that accepts only encodings of quasi-configurations tree, this allows us to containment-encode accepting runs of an ATM. Only linear queries, possibly nested, will be needed to perform the required checks, even in the case of ATMs. To simplify the statements, we use LinMQ^0 as a synonym for UCQ.

Lemma 12. *Consider an ATM* $\mathcal{M}$, *and queries as in Definition 11, including* SameCell$[x,y]$, *that are* MQ^k *queries for some* $k \geq 0$. *There is a* MQ^k *query* $P[x]$, *polynomial in the size of* $\mathcal{M}$ *and the given queries, such that the following hold.*

- *For every accepting run of* $\mathcal{M}$ *in space* s, *there is some database instance* $\mathcal{I}$ *with some element* c *that encodes the run, such that* $c \notin P^{\mathcal{I}}$.

- *If an element* c *of* $\mathcal{I}$ *encodes a tree of quasi-configurations of* $\mathcal{M}$ *in space* s, *and if* $c \notin P^{\mathcal{I}}$, *then* c *encodes an accepting run of* $\mathcal{M}$ *in space* s.

Moreover, if all input queries are in LinMQ^k, *then so is P.*

Proof. We construct P from all (polynomially many) positive queries obtained by instantiating the query patterns in Figure 3. Since P needs to be a unary query with variable x, we extend every positive query that does not contain x with the atom FirstConf$[x,x']$ (omitted for space reasons in Figure 3). By Proposition 1 we can express the disjunctions of all the positive queries in Figure 3 as a LinMQ^k $P[x]$ of polynomial size (for $k=0$ it is a UCQ).

If an element c in a database instance $\mathcal{I}$ encodes an accepting run of $\mathcal{M}$ in space s, and $\mathcal{I}$ contains no other structures, then none of the queries in Figure 3 matches. Hence $c \notin P^{\mathcal{I}}$.

Conversely, assume that c encodes a tree of $\mathcal{M}$ quasi-configurations in space s and $c \notin P^{\mathcal{I}}$. If none of the queries in Figure 3 (1) match, the head positions of every configuration must form a sequence $l, \ldots, l, h, r, \ldots, r$; hence all quasi-configurations are actually configurations. Queries (2)–(4) ensure that the first and last configuration are in the start and end state, respectively, and that each transition is matched by suitable state and tape modifications. Queries (5) ensure that tape cells that are not at the head of the TM are not modified between configurations. Queries (6) ensure that the movement of the head is consistent with the transitions, and especially does not leave the prescribed space. Note that the
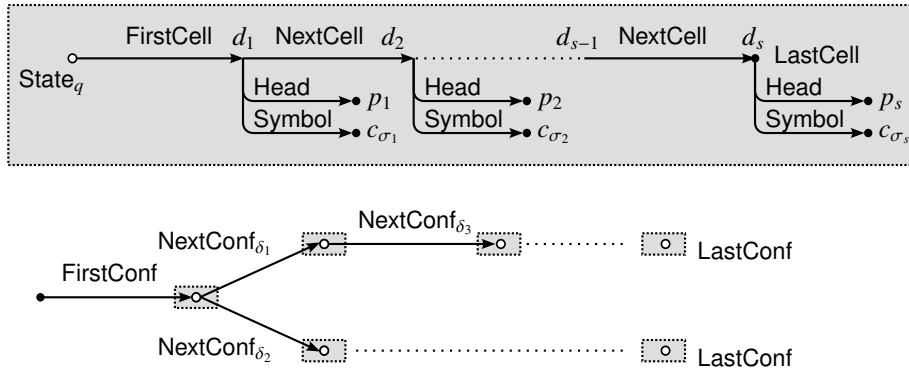
Figure 2: Illustration of the ATM encoding of Definition 11: shaded configurations (top) are used within the configuration tree (bottom); ConfCell queries are omitted for clarity

(1) Unique head marker and correct left/right head markers:

$$\mathsf{Head}(y, p_1) \wedge \mathsf{NextCell}(y, z) \wedge \mathsf{Head}(z, p_2) \quad \text{where } \langle p_1, p_2 \rangle \in \{\langle h, h \rangle, \langle h, l \rangle, \langle r, h \rangle, \langle r, l \rangle\}$$

$$\mathsf{Head}(y, h) \wedge \mathsf{Head}(y, p) \quad \text{where } p \in \{r, l\}$$

(2) Unique start configuration:

$$\mathsf{FirstConf}(x, y) \wedge \mathsf{State}_q(y) \quad \text{where } q \neq q_s$$

$$\mathsf{FirstConf}(x, y) \wedge \mathsf{FirstCell}(y, z) \wedge \mathsf{Head}(z, p) \quad \text{where } p \in \{l, r\}$$

$$\mathsf{FirstConf}(x, y) \wedge \mathsf{ConfCell}(y, z) \wedge \mathsf{Symbol}(z, c_\sigma) \quad \text{where } \sigma \neq \square$$

(3) Valid, uniquely defined transitions:

$$\mathsf{State}_q(y) \wedge \mathsf{Head}(z, h) \wedge \mathsf{ConfCell}(y, z) \wedge \mathsf{Symbol}(z, c_\sigma) \wedge \mathsf{NextConf}_\delta(y, y') \wedge \quad \text{where } \delta = \langle q_1, \sigma_1, q_2, \sigma_2, d \rangle$$
$$\mathsf{State}_{q'}(y') \wedge \mathsf{ConfCell}(y', z') \wedge \mathsf{SameCell}(z', z) \wedge \mathsf{Symbol}(z', c_{\sigma'}) \quad \text{with } q_1 \neq q \text{ or } \sigma_1 \neq \sigma \text{ or } q_2 \neq q' \text{ or } \sigma_2 \neq \sigma'$$

(4) Unique end state:

$$\mathsf{LastConf}(y) \wedge \mathsf{State}_q(y) \quad \text{where } q \neq q_e$$

(5) Memory:

$$\mathsf{ConfCell}(y_1, x_1) \wedge \mathsf{Head}(x_1, r) \wedge \mathsf{Symbol}(x_1, c_\sigma) \wedge \mathsf{NextConf}_\delta(y_1, y_2) \wedge$$
$$\mathsf{ConfCell}(y_2, x_2) \wedge \mathsf{SameCell}(x_1, x_2) \wedge \mathsf{Symbol}(x_2, c_{\sigma'}) \quad \text{where } \sigma \neq \sigma'$$

$$\mathsf{ConfCell}(y_1, x_1) \wedge \mathsf{Head}(x_1, l) \wedge \mathsf{Symbol}(x_1, c_\sigma) \wedge \mathsf{NextConf}_\delta(y_1, y_2) \wedge$$
$$\mathsf{ConfCell}(y_2, x_2) \wedge \mathsf{SameCell}(x_1, x_2) \wedge \mathsf{Symbol}(x_2, c_{\sigma'}) \quad \text{where } \sigma \neq \sigma'$$

(6) Head movement:

$$\mathsf{ConfCell}(y_1, x_1) \wedge \mathsf{Head}(x_1, h) \wedge \mathsf{NextConf}_\delta(y_1, y_2) \wedge \quad \text{where } \delta = \langle q_1, \sigma_1, q_2, \sigma_2, \text{right} \rangle$$
$$\mathsf{ConfCell}(y_2, x_2) \wedge \mathsf{SameCell}(x_1, x_2) \wedge \mathsf{NextCell}(x_2, x_2') \wedge \mathsf{Head}(x_2', p) \quad \text{and } p \in \{r, l\}$$

$$\mathsf{ConfCell}(y_1, x_1) \wedge \mathsf{Head}(x_1, h) \wedge \mathsf{NextConf}_\delta(y_1, y_2) \wedge \quad \text{where } \delta = \langle q_1, \sigma_1, q_2, \sigma_2, \text{right} \rangle$$
$$\mathsf{ConfCell}(y_2, x_2) \wedge \mathsf{SameCell}(x_1, x_2) \wedge \mathsf{LastCell}(x_2) \wedge \mathsf{Head}(x_2, p) \quad \text{and } p \in \{r, l\}$$

$$\mathsf{ConfCell}(y_1, x_1) \wedge \mathsf{Head}(x_1, h) \wedge \mathsf{NextConf}_\delta(y_1, y_2) \wedge \quad \text{where } \delta = \langle q_1, \sigma_1, q_2, \sigma_2, \text{left} \rangle$$
$$\mathsf{ConfCell}(y_2, x_2) \wedge \mathsf{SameCell}(x_1, x_2) \wedge \mathsf{NextCell}(x_2', x_2) \wedge \mathsf{Head}(x_2', p) \quad \text{and } p \in \{r, l\}$$

$$\mathsf{ConfCell}(y_1, x_1) \wedge \mathsf{Head}(x_1, h) \wedge \mathsf{NextConf}_\delta(y_1, y_2) \wedge \quad \text{where } \delta = \langle q_1, \sigma_1, q_2, \sigma_2, \text{left} \rangle$$
$$\mathsf{ConfCell}(y_2, x_2) \wedge \mathsf{SameCell}(x_1, x_2) \wedge \mathsf{FirstCell}(z, x_2) \wedge \mathsf{Head}(x_2, p) \quad \text{and } p \in \{r, l\}$$

Figure 3: Queries to construct a containment encoding as in Lemma 12

queries allow transitions that try to move the head beyond the tape and require that the head stays in its current position in this case. This allows the ATM to recognize the end of the tape, which is important for the Turing machines that we consider below. With all these restrictions observed, c must encode a run of $\mathcal{M}$ in space s. $\qquad \square$

The previous result allows us to focus on the encoding of quasi-configuration trees and the definition of queries as required in Definition 11. Indeed, the main challenge below will be to enforce a sufficiently large tape for which we can still find a correct same-cell query.

5.4 Hardness of Monadic Query Containment

We can now prove our first major hardness result:

Theorem 13. *Deciding containment of* MDlog *queries in* MQ^k *queries is hard for* $(k + 2)$ExpTime.

Note that the statement includes the 3ExpTime-hardness for containment of MQs as a special case. To prove this result, we first construct an ExpSpace ATM that we then use to construct tapes of double exponential size.

Lemma 14. *For any ATM* $\mathcal{M}$, *there is an* MDlog *query* $P_1[x]$, *a* LinMQ $P_2[x]$, *queries as in Definition 11 that are* LinMQ*s, and a same-cell query that is a* UCQ, *such that* $P_1[x]$ *and* $P_2[x]$ *containment-encode accepting runs of* $\mathcal{M}$ *in exponential space.*

Proof. Let $\mathcal{M} = \langle Q, \Sigma, \Delta, q_s, q_e \rangle$ with Q partitioned into existential states $Q_\exists$ and universal states $Q_\forall$. In order to use Lemma 12, we first construct queries P_1' and P_2' that containment-encode quasi-configuration trees of $\mathcal{M}$ in space 2^ℓ for some ℓ that is linear in the size of the queries (w.r.t. to suitable queries as in Definition 11).

Our signature contains the binary predicates (distinguished from the queries of Definition 11 by using lower case letters) firstConf, nextConf$_\delta$ for all $\delta \in \Delta$, firstCell, nextCell, bit$_i$ for all $i \in \{1, \ldots, \ell\}$, symbol, head, as well as the unary predicates lastConf, and state$_q$ for all $q \in Q$.

We define P_1' to be the following MDlog query that has the goal predicate U_{goal} and uses two further constants 0 and 1:

$$\text{firstConf}(x, y) \wedge U_{conf}(y) \rightarrow U_{goal}(x)$$

$$\text{state}_q(x) \wedge \text{firstCell}(x, y) \wedge$$
$$U_{bit_1}(y) \rightarrow U_{conf}(x) \qquad \text{for } q \in Q$$
$$\text{bit}_{i-1}(x, 0) \wedge U_{bit_i}(x) \rightarrow U_{bit_{i-1}}(x) \qquad \text{for } i \in \{2, \ldots, \ell\}$$
$$\text{bit}_{i-1}(x, 1) \wedge U_{bit_i}(x) \rightarrow U_{bit_{i-1}}(x) \qquad \text{for } i \in \{2, \ldots, \ell\}$$
$$\text{symbol}(x, c_\sigma) \wedge U_{symbol}(x) \rightarrow U_{bit_\ell}(x) \qquad \text{for } \sigma \in \Sigma$$
$$\text{head}(x, h) \wedge U_{head}(x) \rightarrow U_{symbol}(x)$$
$$\text{head}(x, l) \wedge U_{head}(x) \rightarrow U_{symbol}(x)$$
$$\text{head}(x, r) \wedge U_{head}(x) \rightarrow U_{symbol}(x)$$
$$\text{nextCell}(x, y) \wedge U_{bit_1}(y) \rightarrow U_{head}(x)$$
$$\text{nextConf}_\delta(x, y) \wedge U_{conf}(y) \rightarrow U_{head}(x) \qquad \text{for } \delta = \langle q, \sigma, q', \sigma', d \rangle$$
$$\text{with } q \in Q_\exists$$
$$\text{nextConf}_{\delta_1}(x, y_1) \wedge U_{conf}(y_1) \wedge \qquad \text{for } \delta_1 = \langle q, \sigma, q', \sigma', d \rangle,$$
$$\text{nextConf}_{\delta_2}(x, y_2) \wedge U_{conf}(y_2) \rightarrow U_{head}(x) \qquad q \in Q_\forall, \text{ and } \delta_1 \neq \delta_2$$
$$\text{lastConf}(x) \rightarrow U_{head}(x)$$

P_1' encodes structures that resemble configuration trees, but with each configuration "tape" consisting of an arbitrary sequence of "cells" of the form $\text{bit}_1(x, v_1), \ldots, \text{bit}_\ell(x, v_\ell), \text{symbol}(x, c_\sigma), \text{head}(x, p)$, where each v_i is either 0 or 1. The values for the bit sequence encode a binary number of length ℓ. We provide a query P_2' which ensures that each sequence of cells encodes an ascending sequence of binary numbers from $00 \ldots 0$ to $11 \ldots 1$. More precisely, P_2' checks if there are any consecutive cells that violate this rule, i.e., the structures matched by P_1' but not by P_2' are those where each configuration contains 2^ℓ cells. The following query checks whether bit i is the rightmost bit

containing a 0 and bit i in the successor configuration also contains a 0, which is a situation that must not occur if the bit sequences encode a binary counter:

$$\text{bit}_i(y, 0) \wedge \text{bit}_{i+1}(y, 1) \wedge \ldots \wedge \text{bit}_\ell(y, 1) \wedge \text{nextCell}(y, z) \wedge \text{bit}_i(z, 0)$$

In a similar way, we can ensure that every bit to the right of the rightmost 0 is changed to 0, every bit that is left of a 0 remains unchanged, the first number is $0 \ldots 0$, and the last number is $1 \ldots 1$. The query P_2' is the union of all of these (polynomially many) conditions, each with new atom firstConf(x,y) added and all variables other than x existentially quantified; this ensures that we obtain a unary query that matches the same elements as P_1' if it matches at all.

We claim that the elements matching P_1' but not P_2' encode quasi-configuration trees of $\mathcal{M}$ in space 2^ℓ. Indeed, it is easy to specify the queries required by Definition 11. The most complicated query is ConfCell$[x, y]$, which can be defined by the following LinMQ:

$$\text{state}_q(\lambda_1) \wedge \text{nextCell}(\lambda_1, y) \rightarrow U(y) \qquad \text{for all } q \in Q$$
$$U(y) \wedge \text{nextCell}(y, z) \rightarrow U(y)$$
$$U(\lambda_2) \rightarrow \text{hit}$$

The remaining queries are now easy to specify, where we use ConfCell$[x, y]$, knowing that a conjunctive query over LinMQs can be transformed into a single LinMQ using Proposition 1:

$$\text{FirstConf}[x, y] := \text{firstConf}(x, y)$$
$$\text{NextConf}_\delta[x, y] := \exists z.\text{ConfCell}(x, z) \wedge \text{nextConf}_\delta(z, y)$$
$$\text{LastConf}[x] := \exists z.\text{ConfCell}(x, z) \wedge \text{lastConf}(z)$$
$$\text{State}_q[x] := \text{state}_q(x)$$
$$\text{Head}[x, y] := \text{head}(x, y)$$
$$\text{FirstCell}[x, y] := \text{firstCell}(x, y))$$
$$\text{NextCell}[x, y] := \text{nextCell}(x, y)$$
$$\text{LastCell}[x] := \text{lastConf}(x) \vee \exists z.\text{nextConf}(x, z)$$
$$\text{Symbol}[x, y] := \text{symbol}(x, y)$$
$$\text{SameCell}[x, y] := \exists v_1, \ldots, v_\ell.\text{bit}_1(x, v_1) \wedge \text{bit}_1(y, v_1) \wedge$$
$$\ldots \wedge \text{bit}_\ell(x, v_\ell) \wedge \text{bit}_\ell(y, v_\ell)$$

Using these queries, we can construct a LinMQ P as in Lemma 12 such that $P_1 = P_1'$ and $P_2 = P_2' \vee P$ containment-encode accepting runs of $\mathcal{M}$. □

Figure 4 illustrates the encoding that we use to prove Lemma 14. While it resembles the structure of Figure 2, the labels are now EDB predicates rather than (abstract) queries. The encoding of tapes attaches to each cell an ℓ-bit address (where bits are represented by constants 0 and 1). We can use these bits to count from 0 to 2^ℓ to construct tapes of this length. The query on the left-hand side can only enforce that there are cells with bit addresses, not that they actually count; even the exact length of the tape is unspecified. The query on the right-hand side of the containment then checks that consecutive cells (in all tapes that occur in the configuration tree) represent successor addresses, and that the first and last address is as expected.

Another difference from Figure 2 is that we now treat configurations as linear structures, with a beginning and an end.

In our representation of the configuration tree, we next configuration therefore connects to the last cell of the previous configuration's tape, rather than its start. We do this to ensure that the encoding works well even when restricting to linear queries. Indeed, the only non-linear rules in P_1 are used to enforce multiple successor configurations for universal states of an ATM. For normal TMs, even P_1 is in LinMDlog.

Note that we do not enforce any structure to define the query ConfCell; this query is implemented by a LinMQ that navigates over an arbitrary number of cells within one configuration. This is the main reason why we need LinMQs rather than UCQs here.

We now use the exponential space ATM of Lemma 14 to encode the tape of 2EXPSPACE ATM. The following result shows, that one can always obtain an exponentially larger tape by nesting linear queries on the right-hand side.

Lemma 15. *Assume that there is some space bound s such that, for every DTM $\mathcal{M}$, there is a MDlog query $P_1[x]$ and an MQ^{k+1} query $P_2[x]$ with $k \geq 0$, such that $P_1[x]$ and $P_2[x]$ containment-encode accepting runs of $\mathcal{M}$ in s, where the queries required by Definition 11 are MQ^{k+1} queries. Moreover, assume that there is a suitable same-cell query that is in MQ^k.*

Then, for every ATM $\mathcal{M}'$, there is a MDlog query $P_1'[x]$, an MQ^{k+1} $P_2'[x]$, and MQ^{k+1} queries as in Definition 11, such that $P_1'[x]$ and $P_2'[x]$ containment-encode an accepting run of $\mathcal{M}'$ in space $s' \geq 2^s$. Moreover, the size of the queries for this encoding is polynomial in the size of the queries for the original encoding.

Proof. There is a TM $\mathcal{M} = \langle Q, \Sigma, \Delta, q_s, q_e \rangle$ that counts from 0 to 2^s in binary (using space s) and then halts. $\mathcal{M}$ can be small (constant size) since our formalization of (A)TMs allows the TMs to recognize the last tape position to ensure that the maximal available space is used. The computation will necessarily take $s' > 2^s$ steps to complete since multiple steps are needed to increment the counter by 1. Let $P_1[x]$ and $P_2[x]$ be queries that containment-encode accepting runs of $\mathcal{M}$ in s, and let ConfCell, SameCell, etc. denote the respective LinMQ^k as in Definition 11.

Let $\mathcal{M}' = \langle Q', \Sigma', \Delta', q_s', q_e' \rangle$ be an arbitrary ATM. We use the signature of P_1, extended by additional binary predicates firstConf′, nextConf′$_\delta$ for all $\delta \in \Delta'$, symbol′, head′, as well as unary predicates lastConf′, and state′$_q$ for all $q \in Q'$. All of these are assumed to be distinct from predicates in P_1.

Let U_{goal} be the goal predicate of P_1, and let U_{tape} be a new unary IDB predicate. We construct the program $\bar{P}_1$ from P_1 as follows. For every rule of P_1 that does not contain an IDB atom in its body we add the atom $U_{tape}(x)$ to the body, where x is any variable that occurs in the rule. Intuitively speaking, the IDBs U_{tape} and U_{goal} mark the start and end of tapes of $\mathcal{M}'$, which are represented by runs of $\mathcal{M}$. Moreover, we modify $\bar{P}_1$ to "inject" additional state and head information for $\mathcal{M}'$ into configurations of $\mathcal{M}$, i.e., we extend P_1 to ensure that every element e with state$_q(e)$ also occurs in some symbol′$(e, c'_{\sigma'})$ and in some relation head′(e, p). This can always be achieved by adding a linear number of IDB predicates and rules.

Now P_1' is defined to be a MDlog query with goal predicate U'_{goal} (assumed, like all IDB predicates of form U' below, to

be distinct from any IDB predicate in $\bar{P}_1$), which is obtained as the union of $\bar{P}_1$ with the following rules:

$$\text{firstConf}'(x, y) \wedge U'_{conf}(y) \to U'_{goal}(x)$$
$$\text{state}'_q(x) \wedge U_{goal}(x) \to U'_{conf}(x) \qquad \text{for } q \in Q$$
$$\text{nextCell}'(x, y) \wedge U_{goal}(y) \to U_{tape}(x) \qquad \text{for } q \in Q$$
$$\text{nextConf}'_\delta(x, y) \wedge U'_{conf}(y) \to U_{tape}(x) \qquad \text{for } \delta = \langle q, \sigma, q', \sigma', d \rangle$$
$$\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{with } q \in Q_\exists$$
$$\text{nextConf}'_{\delta_1}(x, y_1) \wedge U'_{conf}(y_1) \wedge \qquad \text{for } \delta_1 = \langle q, \sigma, q', \sigma', d \rangle,$$
$$\text{nextConf}'_{\delta_2}(x, y_2) \wedge U'_{conf}(y_2) \to U_{tape}(x) \qquad q \in Q_\forall, \text{ and } \delta_1 \neq \delta_2$$
$$\text{lastConf}'(x) \to U_{tape}(x)$$

P_1' encodes trees of trees of $\mathcal{M}$ quasi-configurations in space s. The structures matched by P_1' but not by P_2 encode trees of accepting runs of $\mathcal{M}$ in space s (note that these runs are linear, since $\mathcal{M}$ is not alternating). Every such run consists of the same number $s' \geq 2^s$ of configurations; these configurations represent the tape cells of our encoding of $\mathcal{M}'$ sequences. This encoding is formalized by queries as follows. The queries FirstConf′$[x, y]$, State′$_q[x]$, Head′$[x, y]$, and Symbol′$[x, y]$ are directly expressed by singleton CQs that use the eponymous predicates firstConf′(x, y), etc. To access cells of $\mathcal{M}'$, we can use the analogous queries to access configurations of $\mathcal{M}$: FirstCell′$[x, y] = $ FirstConf(x, y), NextCell′$[x, y] = $ NextConf(x, y), and LastCell′$[x] = $ LastConf(x).

The remaining queries can be expressed as LinMQ queries. To present these queries in a more readable way, we specify them in regular expression syntax rather than giving many rules for each. It is clear that regular expressions over unary and binary predicates can be expressed in LinMQ (it was already shown that MQs can express regular path queries, which is closely related [Rudolph and Krötzsch, 2013]). We use abbreviation P1SYMBOL to express the regular expression that is a disjunction of all predicate symbols that occur in P_1 (this allows us to skip over any structures generated by P_1; with the specific forms of P_1 that can occur in our proofs, one could make this more specific to use only certain binary predicates, but our formulation does not depend on internals of P_1). Moreover, let STATE be the disjunction of all atoms state′$_q(x)$ and $\exists y.$head′(x, y) (both unary).

$$\text{NextConf}'_\delta[x, y] := \text{STATE P1SYMBOL}^* \text{ nextConf}'_\delta$$
$$\text{LastConf}'[x] := \text{STATE P1SYMBOL}^* \text{ lastConf}'$$
$$\text{ConfCell}'[x, y] := \text{STATE P1SYMBOL}^* \text{ HEAD}$$

The unary query LastConf′$[x]$ uses the variable at the beginning of the expression as its answer. It is easy to verify that the elements accepted by P_1' but not by P_2 encode sequences of quasi-configurations of $\mathcal{M}'$ in space s' with respect to these queries. To apply Lemma 12, we need to specify an additional SameCell′ query for this encoding.

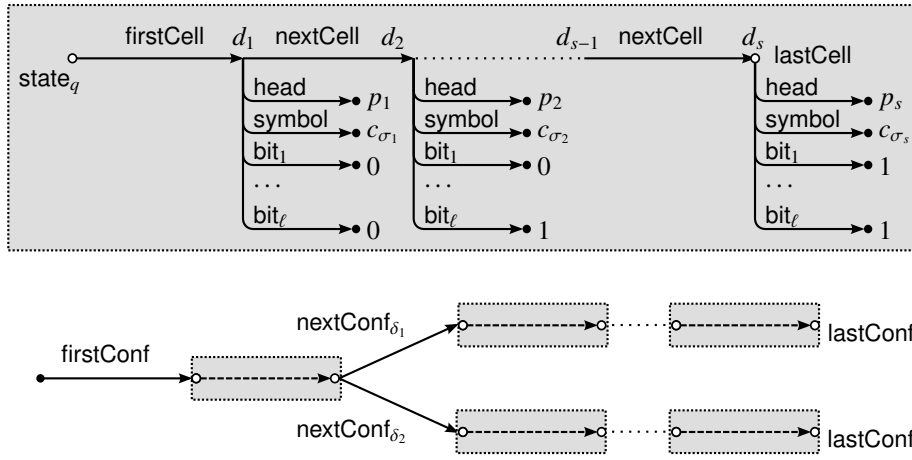SameCell′ is expressed by an MQ^{k+1} query that can in general not be expressed by a MQ^k query:

Figure 4: Illustration of the ATM encoding of Lemma 14: shaded configurations (top) are used within the configuration tree (bottom)

$$\text{FirstCell}(\lambda_1, x) \rightarrow U_1(x)$$

$$U_1(x) \wedge \text{NextCell}(x, x') \rightarrow U_1(x')$$

$$\text{State}_q(\lambda_1) \wedge \text{FirstCell}(\lambda_1, x) \wedge \text{Symbol}(x, z) \wedge \text{Head}(x, v) \wedge$$
$$\text{State}_q(\lambda_2) \wedge \text{FirstCell}(\lambda_2, y) \wedge \text{Symbol}(y, z) \wedge \text{Head}(y, v) \rightarrow U_2(y)$$
$$\text{for all } q \in Q$$

$$U_1(x) \wedge U_2(y) \wedge \text{SameCell}(x, y) \wedge$$
$$\text{NextCell}(x, x') \wedge \text{Symbol}(x', z) \wedge \text{Head}(x', v) \wedge$$
$$\text{NextCell}(y, y') \wedge \text{Symbol}(y', z) \wedge \text{Head}(y', v) \rightarrow U_2(y')$$

$$U_2(y) \wedge \text{LastCell}(y) \rightarrow \text{hit}$$

where FirstCell, Symbol, SameCell, and LastCell are the queries for which P_1 and P_2 containment-encode runs of $\mathcal{M}$. Note that our constructions already ensure that the sequences of $\mathcal{M}$-cells compared by SameCell$'$ are of the same length.

To complete the proof, we apply Lemma 12 to construct an MQ^{k+1} $\bar{P}_2$. The MQ^{k+1} P'_2 is obtained by expressing the disjunction of P_2 and $\bar{P}_2$ as an MQ^{k+1} using Proposition 1. Then P'_1 and P'_2 containment-encode accepting runs of $\mathcal{M}'$ in space s'. □

The most difficult to express in this construction is the new same-cell query SameCell$'$. The first two rules simply mark the tape starting at λ_1 with U_1. The next two rules then compare the two (potentially very long) tapes from configurations of $\mathcal{M}$ to check if they contain exactly the same symbols at each position, and the last rule finishes. Since the tapes are not connected in any known way, we have to be careful to ensure never to loose the connection to either of the tapes, to avoid comparing random cells from other parts of the database. Indeed, the last two rules do not mention λ_1 or λ_2 at all. We need two IDB predicates to achieve this, which carefully mark the two tapes cell by cell.

Another important thing to note is that the query SameCell is only used exactly once in exactly one rule. Indeed, if we were using it twice, then the length of our queries would grow exponentially when applying the construction inductively. This is the reason why we encode symbols and head positions with constants, rather than using unary predicates like for states. In the latter case, we need many rules, one

for each predicate, as can be seen in the third rule above. One could try to avoid the use of constants by more complex encodings that encode information using paths of different lengths as done by Björklund et al. [Björklund *et al.*, 2008]. However, some additional device is needed to ensure that database instances are sufficiently closely connected in this case, which may again require constants, IDBs of higher arity, or a greater nesting level of LinMQ queries to navigate larger distances.

With the previous results, Theorem 13 can be proved by an easy induction: for the base case $k = 1$ we apply Lemma 15 to the result of Lemma 14; for the induction step we use Lemma 15 again.

Proof. The claim is shown by induction on k. For the base case, we show that deciding containment of MQ queries is 3ExpTime-hard. By Lemma 14, for any DTM $\mathcal{M}^0$, there is a MDlog query P_1^0, a LinMQ P_2^0, LinMQs as in Definition 11, and a same-cell query that is a UCQ with respect to which P_1^0 and P_2^0 containment-encode accepting runs of $\mathcal{M}^0$ in exponential space s. By applying Lemma 15, we obtain, for an arbitrary ATM $\mathcal{M}^1$, a MDlog query P_1^1, an MQ P_2^1, and MQ queries as in Definition 11 (including a same-cell query), that containment-encode accepting runs of $\mathcal{M}^1$ in space $s' \geq 2^s$.

The induction step for $k > 1$ is immediate from Lemma 15.
□

5.5 Linear Datalog

Not only query answering, but also containment checking is often slightly simpler in fragments of linear Datalog. Intuitively, this is so because derivations can be represented as words rather than as trees. Thus, the automata theoretic techniques that we have used in Section 5.2 can be applied with automata on words where some operations are easier. In particular, containment of (nondeterministic) automata on words can be checked in polynomial space rather than in exponential time. This allows us to establish the following theorems, which reduce the 2ExpTime upper bound of Theorem 9 to Ex-

PSPACE and the $(k+2)$EXPTIME upper bound of Theorem 10 to $(k+1)$EXPSPACE.

Theorem 16. *Containment of* LinDlog *queries in* GDlog *queries can be decided in* EXPSPACE.

Theorem 17. *Containment of* LinDlog *queries in* GQ^k *queries can be decided in* $(k+1)$EXPSPACE.

Establishing matching lower bounds for the complexity turns out to be more difficult. In general, we loose the power of alternation, which explains the reduction in complexity. The general approach of encoding (non-alternating) Turing machines is the same as in Section 5.3, where Definition 11 is slightly simplified since we do not need to consider universal states, so that configuration trees turn into configuration sequences. Moreover, Lemma 12 applies to this case as well, since it only requires linear queries. Likewise, our general inductive step in Lemma 15 uses deterministic (non-alternating) TMs to construct exponentially long tapes. Moreover, it turns out that the construction of an initial exponential space TM in Lemma 14 leads to linear queries if the TM has no universal states.

Yet it is challenging to lift the exact encodings of Lemma 14 and Lemma 15. The same-cell query that we constructed in Lemma 15 for our inductive argument is non-linear. As explained in Section 5.4, the use of two IDBs to mark both sequences of tape cells is essential there to ensure correctness. The main problem is that we must not loose connection to either of the sequences during our checks. As an alternative to using IDBs on both sequences, one could use the ConfCell query to ensure that the compared cells belong to the right configurations. This leads to the following same-cell query:

$$\text{State}_q(\lambda_1) \wedge \text{FirstCell}(\lambda_1, x) \wedge \text{Symbol}(x, z) \wedge \text{Head}(x, v) \wedge$$
$$\text{State}_q(\lambda_2) \wedge \text{FirstCell}(\lambda_2, y) \wedge \text{Symbol}(y, z) \wedge \text{Head}(y, v) \rightarrow \text{U}(y)$$
$$\text{for all } q \in Q$$

$$\text{U}(y) \wedge \text{ConfCell}(\lambda_1, x) \wedge \text{SameCell}(x, y) \wedge$$
$$\text{NextCell}(x, x') \wedge \text{Symbol}(x', z) \wedge \text{Head}(x', v) \wedge$$
$$\text{NextCell}(y, y') \wedge \text{Symbol}(y', z) \wedge \text{Head}(y', v) \rightarrow \text{U}(y')$$

$$\text{U}(y) \wedge \text{LastCell}(y) \rightarrow \text{hit}$$

While this works in principle, it has the problem that the ConfCell query of Lemma 14 is a LinMQ, not a UCQ. Therefore, if we construct a same-cell query for the 2EXPSPACE case, we obtain LinMQ^2 queries, which yields the following result:

Theorem 18. *Deciding containment of* LinMDlog *queries in* LinMQ^k *queries is hard for* kEXPSPACE.

In order to do better, one can try to express ConfCell as a UCQ. In general, this is not possible on the database instances that the left-hand query in Lemma 14 recognizes, since cells may have an exponential distance to their configuration while UCQs can only recognize local structures. To make ConfCell local, we can modify the left-hand query to ensure that every cell is linked directly to its configuration with a binary predicate inConf. Using binary IDB predicates, we can do this with the following set of frontier-guarded rules (for all $q \in Q$, $i \in \{2, \ldots, \ell\}$, $\sigma \in \Sigma$, and $\delta \in \Delta$, respectively):

$$\text{firstConf}(x, y) \wedge \text{U}_{conf}(y) \rightarrow \text{U}_{goal}(x)$$
$$\text{state}_q(x) \wedge \text{nextCell}(x, y) \wedge \text{inConf}(y, x) \wedge \text{U}_{bit_1}(y, x) \rightarrow \text{U}_{conf}(x)$$
$$\text{bit}_{i-1}(x, 0) \wedge \text{U}_{bit_i}(y, z) \wedge \text{inConf}(x, z) \rightarrow \text{U}_{bit_{i-1}}(x, z)$$

$$\text{bit}_{i-1}(x, 1) \wedge \text{U}_{bit_i}(y, z) \wedge \text{inConf}(x, z) \rightarrow \text{U}_{bit_{i-1}}(x, z)$$
$$\text{symbol}(x, c_\sigma) \wedge \text{U}_{symbol}(x, z) \wedge \text{inConf}(x, z) \rightarrow \text{U}_{bit_\ell}(x, z)$$
$$\text{head}(x, h) \wedge \text{U}_{head}(x, z) \wedge \text{inConf}(x, z) \rightarrow \text{U}_{symbol}(x, z)$$
$$\text{head}(x, l) \wedge \text{U}_{head}(x, z) \wedge \text{inConf}(x, z) \rightarrow \text{U}_{symbol}(x, z)$$
$$\text{head}(x, r) \wedge \text{U}_{head}(x, z) \wedge \text{inConf}(x, z) \rightarrow \text{U}_{symbol}(x, z)$$
$$\text{nextCell}(x, y) \wedge \text{U}_{bit_1}(y, z) \wedge \text{inConf}(x, z) \rightarrow \text{U}_{head}(x, z)$$
$$\text{nextConf}_\delta(x, y) \wedge \text{U}_{conf}(y) \wedge \text{inConf}(x, z) \rightarrow \text{U}_{head}(x, z)$$
$$\text{lastConf}(x) \wedge \text{inConf}(x, z) \rightarrow \text{U}_{head}(x, z)$$

Structures matched by this query provide direct links from each element to their configuration element, and we can thus formulate ConfCell as a UCQ and obtain the following.

Theorem 19. *Deciding containment of* LinGDlog *queries in* LinMQ^k *queries is hard for* $(k+1)$EXPSPACE.

It is not clear if this result can be extended to containments of LinMQ in LinMQ^k; the above approach does not suggest any suitable modification. In particular, the propagation of inConf in the style of a transitive closure does not work, since elements may participate in many inConf relations. On the other hand, the special constants λ in LinMQs cannot be used to refer to the current configuration, since there can be an unbounded number of configurations but only a bounded number of special constants. It is possible, however, to formulate a LinMQ Config[x] that generates the required structure for a single configuration, since one can then represents the configuration by λ. We can generate arbitrary sequences of such structures by using Config[x] as a nested query to that matches a regular expression firstConf (Config NextConf)* Config lastConf, where we use NextConf to express the disjunction of all nextConf$_\delta$ relations. This proves the following statement.

Theorem 20. *Deciding containment of* LinMQ^2 *queries in* LinMQ^k *queries is hard for* $(k+1)$EXPSPACE.

Finally, we can also continue to use the same approach for encoding SameCell as in Section 5.4, without using ConfCell, while still restricting to linear Datalog (and thus to non-alternating TMs) on the left-hand side. This leads us to the following result.

Theorem 21. *Deciding containment of* LinMDlog *queries in* MQ^k *queries is hard for* $(k+1)$EXPSPACE.

We have thus established tight complexity bounds for the containment of nested GQs, while there remains a gap (of one exponential or one nesting level) for MQs.

6 Conclusions

We have studied the most expressive fragments of Datalog for which query containment is known to be decidable, and we provided exact complexities for query answering and containment in most cases. Our results are summarized in Table 1. While containment tends to be non-elementary for nested queries, we have identified tight exponential complexity hierarchies depending on nesting depth. Our results settle several open problems for known query languages: the complexity of query containment for MQ and MQ$^+$, the complexity of query containment of Dlog in GDlog, and the expressivity of nested LinDlog.

	UCQ, LinMDlog, MDlog, LinGDlog, GDlog	LinMQ^k, LinGQ^k	MQ^k, GQ^k	$\text{LinMQ}^+, \text{MQ}^+$, $\text{LinGQ}^+, \text{GQ}^+$	Dlog
LinMQ	PSPACE-h [Chaudhuri and Vardi, 1994] EXPSPACE [Th.16]	kEXPSPACE-h [Th.18] $(k+1)$EXPSPACE [Th.17]	$(k+1)$EXPSPACE-c [Th.21]\[Th.17]	Nonelementary [Th.18]	Undecidable [Abiteboul et al., 1994]
LinGDlog, LinMQ^n $(n \geq 2)$, $\text{LinMQ}^+, \text{LinGQ}^+$, LinGQ^n, LinDlog	EXPSPACE-c [Th.20]\[Th.16]	$(k+1)$EXPSPACE-c [Th.19,20]\[Th.17]	$(k+1)$EXPSPACE-c [Th.19,20]\[Th.17]	Nonelementary [Th.19,20]	Undecidable [Abiteboul et al., 1994]
MDlog, GDlog, MQ^n, GQ^n, MQ^+, GQ^+, Dlog	2EXPTIME-c [Benedikt et al., 2012], [Chaudhuri and Vardi, 1997]\ [Cosmadakis et al., 1988], [Th.9]	$(k+2)$EXPTIME-c [Th.13]\[Th.10]	$(k+2)$EXPTIME-c [Th.13]\[Th.10]	Nonelementary [Th.13]	Undecidable [Shmueli, 1987]

Table 1: Summary of the known complexities of query containment for several Datalog fragments; sources for each claim are shown in square brackets, using \ to separate sources for lower and upper complexity bounds, respectively

Moreover, we have built on the recent "flag & check" approach of monadically defined queries to derive various natural extensions, which lead to new query languages with interesting complexity results. In most cases, we observed that the extension from monadic to frontier-guarded Datalog does not affect any of the complexities, whereas it might have an impact on expressivity. In contrast, the restriction to linear Datalog has the expected effects, both for query answering and for containment.

The only case for which our results for containment complexity are not tight is when we restrict rules to be both linear and monadic: while small variations in the involved query languages lead to the expected tight bounds, this particular combination eludes our analysis. This case is closely related to conjunctive regular path queries, and inspiration might be drawn from recent results in this field [Reutter, 2013]. This case could be studied as part of a future program for analyzing the behavior of (nested) conjunctive regular path queries, which are also a special form of monadic, linear Datalog.

Another interesting open question is the role of constants. Our hardness proofs, especially in the nested case, rely on the use of constants to perform certain checks more efficiently. Without this, it is not clear how an exponential blow-up of our encoding (or the use of additional nesting levels) could be avoided. Of course, constants can be simulated if we have either predicates of higher arity or special constants as in "flag & check" queries. However, for the case of (linear) monadic Datalog without constants, we conjecture that containment complexities are reduced by one exponential each.

Promising directions for future research include the study of practical containment algorithms, since the automata-based techniques we use do not lend themselves for implementation yet. Another interesting topic is to study problems where we ask for the *existence* of a subsuming query of a certain type rather than merely check containment of two given queries. The most prominent instance of this scenario is the *boundedness* problem, which asks whether a given Datalog program can be expressed by some (yet unknown) UCQ. It has been shown that this problem can be studied using tree-automata-based techniques as for query containment [Cosmadakis et al., 1988], though other approaches have been applied as well [Bárány et al., 2012]. Besides boundedness, one can also ask more general questions of *rewritability*, e.g., whether some Datalog program can be expressed in monadic Datalog or in a regular path query.

References

[Abiteboul and Vianu, 1999] Serge Abiteboul and Victor Vianu. Regular path queries with constraints. *J. Comput. Syst. Sci.*, 58(3):428–452, 1999.

[Abiteboul et al., 1994] Serge Abiteboul, Richard Hull, and Victor Vianu. *Foundations of Databases*. Addison Wesley, 1994.

[Baget et al., 2011] Jean-François Baget, Michel Leclère, Marie-Laure Mugnier, and Eric Salvat. On rules with existential variables: Walking the decidability line. *Artificial Intelligence*, 175(9–10):1620–1654, 2011.

[Bárány et al., 2011] Vince Bárány, Balder ten Cate, and Luc Segoufin. Guarded negation. In Luca Aceto, Monika Henzinger, and Jiri Sgall, editors, *ICALP (2)*, volume 6756 of *LNCS*, pages 356–367. Springer, 2011.

[Bárány et al., 2012] Vince Bárány, Balder ten Cate, and Martin Otto. Queries with guarded negation. *PVLDB*, 5(11):1328–1339, 2012.

[Benedikt et al., 2012] Michael Benedikt, Pierre Bourhis, and Pierre Senellart. Monadic datalog containment. In *Proc. 39th Int. Coll. on Automata, Languages, and Programming (ICALP'12)*, pages 79–91, 2012.

[Bienvenu et al., 2013] Meghyn Bienvenu, Balder ten Cate, Carsten Lutz, and Frank Wolter. Ontology-based data access: A study through disjunctive datalog, CSP, and MMSNP. In Hull and Fan [2013], pages 213–224.

[Bienvenu et al., 2014] Meghyn Bienvenu, Diego Calvanese, Magdalena Ortiz, and Mantas Simkus. Nested regular path queries in description logics. In Chitta Baral, Giuseppe De Giacomo, and Thomas Eiter, editors, *Proc. 14th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'14)*. AAAI Press, 2014.

[Bischoff *et al.*, 2014] Stefan Bischoff, Markus Krötzsch, Axel Polleres, and Sebastian Rudolph. Schema-agnostic query rewriting for SPARQL 1.1. In Peter Mika, Tania Tudorache, Abraham Bernstein, Chris Welty, Craig A. Knoblock, Denny Vrandečić, Paul T. Groth, Natasha F. Noy, Krzysztof Janowicz, and Carole A. Goble, editors, *Proc. 13th Int. Semantic Web Conf. (ISWC'14)*, volume 8796 of *LNCS*, pages 584–600. Springer, 2014.

[Björklund *et al.*, 2008] Henrik Björklund, Wim Martens, and Thomas Schwentick. Optimizing conjunctive queries over trees using schema information. In Edward Ochmanski and Jerzy Tyszkiewicz, editors, *Proc. 3rdInt. Symposium on Mathematical Foundations of Computer Science*, volume 5162 of *LNCS*, pages 132–143. Springer, 2008.

[Boutilier, 2009] Craig Boutilier, editor. *Proc. 21st Int. Joint Conf. on Artificial Intelligence (IJCAI'09)*. IJCAI, 2009.

[Calì *et al.*, 2008] Andrea Calì, Georg Gottlob, and Michael Kifer. Taming the infinite chase: Query answering under expressive relational constraints. In Gerhard Brewka and Jérôme Lang, editors, *Proc. 11th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'08)*, pages 70–80. AAAI Press, 2008.

[Calvanese *et al.*, 2003] Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, and Moshe Y. Vardi. Reasoning on regular path queries. *SIGMOD Record*, 32(4):83–92, 2003.

[Calvanese *et al.*, 2005] Diego Calvanese, Giuseppe De Giacomo, and Moshe Y. Vardi. Decidable containment of recursive queries. *Theor. Comput. Sci.*, 336(1):33–56, 2005.

[Calvanese *et al.*, 2007a] Diego Calvanese, Thomas Eiter, and Magdalena Ortiz. Answering regular path queries in expressive description logics: An automata-theoretic approach. In *Proc. 22nd AAAI Conf. on Artificial Intelligence (AAAI'07)*, pages 391–396. AAAI Press, 2007.

[Calvanese *et al.*, 2007b] Diego Calvanese, Guiseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. Tractable reasoning and efficient query answering in description logics: The DL-Lite family. *J. of Automated Reasoning*, 39(3):385–429, 2007.

[Calvanese *et al.*, 2009] Diego Calvanese, Thomas Eiter, and Magdalena Ortiz. Regular path queries in expressive description logics with nominals. In Boutilier [2009], pages 714–720.

[Chandra *et al.*, 1981] Ashok K. Chandra, Dexter C. Kozen, and Larry J. Stockmeyer. Alternation. *J. of the ACM*, 28(1):114–133, 1981.

[Chaudhuri and Vardi, 1994] Surajit Chaudhuri and Moshe Y. Vardi. On the complexity of equivalence between recursive and nonrecursive Datalog programs. In *Proc. 13th Symp. on Principles of Database Systems (PODS'93)*, pages 107–116, 1994.

[Chaudhuri and Vardi, 1997] Surajit Chaudhuri and Moshe Y. Vardi. On the equivalence of recursive and nonrecursive Datalog programs. *J. of Comput. Syst. Sci.*, 54(1):61–78, 1997.

[Cosmadakis *et al.*, 1988] Stavros Cosmadakis, Haim Gaifman, Paris Kanellakis, and Moshe Vardi. Decidable optimization problems for database logic programs. In *Proc. 20th Annual ACM Symp. on Theory of Computing (STOC'88)*, pages 477–490. ACM, 1988.

[Courcelle, 1991] Bruno Courcelle. Recursive queries and context-free graph grammars. *Theor. Comput. Sci.*, 78(1):217–244, 1991.

[Deutsch and Tannen, 2002] Alin Deutsch and Val Tannen. Optimization properties for classes of conjunctive regular path queries. In *Revised Papers from the 8th Int. Workshop on Database Programming Languages (DBPL'01)*, pages 21–39. Springer, 2002.

[Eiter *et al.*, 2009] Thomas Eiter, Carsten Lutz, Magdalena Ortiz, and Mantas Simkus. Query answering in description logics with transitive roles. In Boutilier [2009], pages 759–764.

[Florescu *et al.*, 1998] Daniela Florescu, Alon Levy, and Dan Suciu. Query containment for conjunctive queries with regular expressions. In Alberto O. Mendelzon and Jan Paredaens, editors, *Proc. 17th Symp. on Principles of Database Systems (PODS'98)*, pages 139–148. ACM, 1998.

[Hull and Fan, 2013] Richard Hull and Wenfei Fan, editors. *Proc. 32nd Symp. on Principles of Database Systems (PODS'13)*. ACM, 2013.

[Krötzsch, 2011] Markus Krötzsch. Efficient rule-based inferencing for OWL EL. In Toby Walsh, editor, *Proc. 22nd Int. Joint Conf. on Artificial Intelligence (IJCAI'11)*, pages 2668–2673. AAAI Press/IJCAI, 2011.

[Reutter, 2013] Juan L. Reutter. Containment of nested regular expressions. *CoRR*, abs/1304.2637, 2013.

[Rudolph and Krötzsch, 2013] Sebastian Rudolph and Markus Krötzsch. Flag & check: Data access with monadically defined queries. In Hull and Fan [2013], pages 151–162.

[Savitch, 1970] Walter J. Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of Computer and System Sciences*, 4(2):177 – 192, 1970.

[Shmueli, 1987] O. Shmueli. Decidability and expressiveness aspects of logic queries. In *Proc. 6th Symp. on Principles of Database Systems (PODS'87)*, pages 237–249. ACM, 1987.

[Xiao *et al.*, 2010] Guohui Xiao, Stijn Heymans, and Thomas Eiter. DReW: a reasoner for Datalog-rewritable description logics and dl-programs. In Thomas Eiter, Adil El Ghali, Sergio Fernández, Stijn Heymans, Thomas Krennwallner, and François Lévy, editors, *Proc. 1st Int. Workshop on Business Models, Business Rules and Ontologies (BuRO'10)*, pages 1–14. ONTORULE Project, 2010.

A   Tree Automata

We use standard definitions for two-way alternating tree automata as introduced in [Cosmadakis *et al.*, 1988]. A regular (one-way, non-alternating) tree automaton is obtained by restricting this definition.

Tree automata run over ranked, labelled trees of some maximal arity (out-degree) f. A ranked tree can be seen a function t mapping sequences of positive natural numbers (encoding nodes in the tree) to symbols from a fixed finite alphabet (the labels of each node). Each letter of the alphabet is ranked, i.e., associated with an arity that defines how many child nodes a node labeled with this symbol should have. The domain of t, denoted Nodes(t), satisfies the following closure property: if $i \cdot j \in$ Nodes(t), then $i \in$ Nodes(t) and $i \cdot k \in$ Nodes(t) for all $1 \leq k \leq j$. Given a ranked tree t, we write $i \in$ Nodes(t) to denote an arbitrary node of t and $t(i)$ to denote the label of i in t. We denote by Trees(Σ) the set of trees over the alphabet Σ.

A two-way alternating tree automaton $\mathcal{A}$ is a tuple $\langle \Sigma, Q, Q_s, \delta, Q_e \rangle$ where

- Σ is a tree alphabet;
- Q is a set of states;
- $Q_s \subseteq Q$ is the set of initial states;
- $Q_e \subseteq Q$ is the set of accepting states;
- δ is a transition function from $Q \times \Sigma$: let $q \in Q$ be a state and $\sigma \in \Sigma$ be a letter of arity ℓ; then $\delta(q, \sigma)$ is a positive boolean combination of elements in $\{-1, 0, 1, \cdots, \ell\} \times Q$.

The numbers used in transitions encode directions, where -1 is up and 0 is stay. For example $\delta(q, \sigma) = (\langle 1, s_1 \rangle \wedge \wedge 1, s_2) \vee (\langle -1, t_3 \rangle \wedge \langle 2, t_4 \rangle)$ is an example of transition for a state q and a node labeled σ: a node labeled by σ can be in the state q iff its first child can be in the states s_1 and s_2, or its parent and its second child can be in the states s_3 and s_4, respectively.

Let t be a tree over Σ. A run τ of $\mathcal{A}$ over t is a tree labeled by elements of $Q \times \{-1, 0, 1, \cdots, f\} \times$ Nodes(t) $\cup \{-1\}$. τ satisfies the following properties:

- τ is finite.
- The root of τ is labelled by (q_0, i, n), where q_0 is in Q_s.
- If a node v is labelled by (q, i, n) and n is not a node of t, then v is a leaf of τ.
- If a node v is labelled by (q, i, n'), n is a node of τ labelled by σ of arity l and v' is labelled by (q_1, j, n') then
  - if $j = -1$, then there exists $u \leq k$ such that $n = n'.u$
  - if $j = 0$, then $n = n'$
  - if $j \leq k$, then $n' = n.j$.
- if a node v is labelled by (q, i, n), $n \in t$ labelled by σ and the children of v are labelled by $(q_1, j_1, n_1) \cdots (q_k, j_k, n_k)$ then $\delta(q, \sigma)$ is satisfied when interpreting the sybmols $\{\langle j_1, q_1 \rangle, \cdots, \langle j_k, q_k \rangle\}$ as *true* and all other symbols as *false*.

τ is *valid* iff, for each leaf of τ labelled by (q, i, n), q is in Q_e. $\mathcal{A}$ *accepts* a tree t if there exists a valid run of t over $\mathcal{A}$. We denote by Trees($\mathcal{A}$). The set of trees accepted by $\mathcal{A}$.

A regular (one-way, non-alternating) tree automaton is a 2-way alternating tree automaton where all transitions for a symbol σ of rank ℓ are boolean formulae of the form $(\langle 1, q_{11} \rangle \wedge \ldots \wedge \langle \ell, q_{\ell 1} \rangle) \vee \ldots \vee (\langle 1, q_{1n} \rangle \wedge \ldots \wedge \langle \ell, q_{\ell n} \rangle)$ for some $n \geq 0$. In particular, directions 0 and -1 do not occur. In this case, we can represent transitions as sets of lists of states $\{\langle q_{11}, \ldots, q_{\ell 1} \rangle, \ldots, \langle q_{1n}, \ldots, q_{\ell n} \rangle\}$.

Finally, we recall two useful theorems from [Cosmadakis *et al.*, 1988].

Theorem 22 (Theorem A.1 of [Cosmadakis *et al.*, 1988])**.** *Let $\mathcal{A}$ be a two-way alternating automaton. Then there exists a tree automaton $\mathcal{A}$ whose size is exponential in the size of $\mathcal{A}$ such that* Trees($\mathcal{A}$) = Trees(Σ) \ Trees($\mathcal{A}$).

Theorem 23 (Theorem A.2 of [Cosmadakis *et al.*, 1988])**.** *Let $\mathcal{A}$ be a two-way alternating automaton. Then there exists a tree automaton $\mathcal{A}$ whose size is exponential in the size of $\mathcal{A}$ such that* Trees($\mathcal{A}$) = Trees($\mathcal{A}$).