

Querying the Unary Negation Fragment with Regular Path Expressions*

Jean Christoph Jung, Carsten Lutz, Mauricio Martel, and Thomas Schneider

University of Bremen, Bremen, Germany
{jeanjung,clu,martel,ts}@informatik.uni-bremen.de

Abstract

The unary negation fragment of first-order logic (UNFO) has recently been proposed as a generalization of modal logic that shares many of its good computational and model-theoretic properties. It is attractive from the perspective of database theory because it can express conjunctive queries (CQs) and ontologies formulated in many description logics (DLs). Both are relevant for ontology-mediated querying and, in fact, CQ evaluation under UNFO ontologies (and thus also under DL ontologies) can be ‘expressed’ in UNFO as a satisfiability problem. In this paper, we consider the natural extension of UNFO with regular expressions on binary relations. The resulting logic UNFO^{reg} can express (unions of) conjunctive two-way regular path queries (C2RPQs) and ontologies formulated in DLs that include transitive roles and regular expressions on roles. Our main results are that evaluating C2RPQs under UNFO^{reg} ontologies is decidable, 2EXPTIME -complete in combined complexity, and CONP -complete in data complexity, and that satisfiability in UNFO^{reg} is 2EXPTIME -complete, thus not harder than in UNFO.

1998 ACM Subject Classification F.4.1 [Mathematical Logic]: Computational Logic

Keywords and phrases Query Answering, Regular Path Queries, Decidable Fragments of First-Order Logic, Computational Complexity

Digital Object Identifier 10.4230/LIPIcs.ICDT.2018.15

1 Introduction

In ontology-mediated querying, queries against incomplete and heterogeneous data are supported by an ontology that provides domain knowledge and assigns a semantics to the data [15, 19, 37, 41]. The ontologies are often formulated in a specialized language such as a description logic [4, 5] or an existential rule language [6, 7, 22, 33] while the actual query is typically a conjunctive query (CQ) or a mild extension thereof such as a union of CQs (UCQ). However, it can also be useful to consider more expressive decidable fragments of first-order logic (FO) as an ontology language as this serves to explore the limits of the ontology-mediated querying approach, to provide maximum expressive power for ontology formulation, and to put ontology-mediated querying into a more general logical perspective. Notably, this has been done in [8, 9, 19], where the guarded fragment (GF), the unary negation fragment (UNFO), and the guarded negation fragment (GN) of FO have been used as ontology languages. These fragments originate from the attempt to explain the good computational behaviour of modal and description logics and to extend their expressive power in a natural way. While GF and UNFO are orthogonal in expressive power, GN

* A full version of this paper with an appendix that contains detailed proofs is available at <http://www.informatik.uni-bremen.de/tdki/research/papers.html>



subsumes both of these fragments [9] and all of them subsume many common modal and description logics. It is an important result that, for all these fragments, ontology-mediated querying with UCQs remains decidable and that the complexity stays within the expected, namely 2EXPTIME combined and CONP data.

From the perspective of database theory, it is an attractive property of both UNFO and GN (but not of GF) that they can express CQs and UCQs. In ontology-mediated querying, this allows to ‘express’ the evaluation of ontology-mediated queries in terms of satisfiability in a natural way. It is easiest to state this for Boolean queries: if (\mathcal{O}, Σ, q) is an ontology-mediated query (OMQ) where \mathcal{O} is an ontology, Σ a set of predicate symbols (that is, relation names) that may occur in the data, and q a UCQ, and D is a Σ -database, then $D \models (\mathcal{O}, \Sigma, q)$ iff $\mathcal{O} \wedge D \wedge \neg q$ is unsatisfiable. When \mathcal{O} is formulated in UNFO or in GN, then so is $\mathcal{O} \wedge D \wedge \neg q$. What is more, the containment of OMQs can also be ‘expressed’ as a satisfiability problem in the natural case where both OMQs contain the same ontology and Σ is the set of all predicates symbols; from now on, we generally mean this case when speaking of OMQ containment. But also beyond ontology-mediated querying, we believe that the ability to express UCQs makes UNFO and GN attractive as an expressive logical backdrop for database theory.

In this paper, we study the natural extension UNFO^{reg} of UNFO with regular path expressions on binary relations. The resulting logic has the attractive property that it allows to express regular path queries [29] and conjunctive two-way regular path queries (C2RPQs) [25] as well as unions thereof (UC2RPQs). Such queries play a central role in the area of graph databases [2, 10] and they have also received considerable attention in ontology-mediated querying [12, 17, 18, 26, 27, 28, 40]. An additional reason to consider UNFO^{reg} is provided by the observation that transitive roles are an important feature of many common description logics (a role is a binary relation), but that transitive roles cannot be expressed in UNFO. In UNFO^{reg} , even transitive closure of roles and regular expressions on roles are expressible, two features that are provided by several expressive description logics [3, 24]. As a concrete example, every ontology formulated in $\mathcal{ALCT}^{\text{reg}}$, the extension of the common description logic \mathcal{ALCT} with regular expressions on roles [44], can be expressed in UNFO^{reg} and thus the evaluation of ontology-mediated queries (\mathcal{O}, Σ, q) where \mathcal{O} is formulated in $\mathcal{ALCT}^{\text{reg}}$ and q is a UC2RPQ can be ‘expressed’ as a satisfiability problem in UNFO^{reg} ; of course, the same is true when \mathcal{O} is formulated in UNFO^{reg} itself. We remark that transitive roles cannot be expressed in GF and GNF either, and that adding transitive relations to GF without losing decidability requires to impose rather strong syntactical restrictions [46], especially so in an ontology-mediated querying context [34]. Adding transitive relations to GNF has, to the best of our knowledge, not yet been studied.

The main problem that we are interested in is evaluating OMQs in which the ontology is formulated in UNFO^{reg} and the actual query is a UC2RPQ. We show that this problem is decidable, 2EXPTIME-complete in combined complexity and CONP-complete in data complexity. We further consider the OMQ containment problem and show that it is 2EXPTIME-complete as well. We additionally show that the complexity of model checking in UNFO^{reg} is the same as in UNFO, namely complete for $\text{P}^{\text{NP}[O(\log^2 n)]}$.

As explained above, both OMQ evaluation and OMQ containment can be reduced to satisfiability in polynomial time. For studying the combined complexity of the former and the complexity of the latter, we thus concentrate on the satisfiability problem and prove a 2EXPTIME upper bound. Note that the addition of regular expressions does thus not increase the complexity of this problem as satisfiability in UNFO is also 2EXPTIME-complete [47] and that the lower bound holds already when the arity of predicates is bounded by two,

as a consequence of the results in [39]. Our proof proceeds by first showing that every satisfiable UNFO^{reg} formula φ has a model whose treewidth is bounded by the size of φ , then establishing a characterization of the satisfaction of C2RPQs (that occur as a building block in φ) in such models in terms of certain witness trees, and finally showing that this infrastructure gives rise to a decision procedure based on two-way alternating tree automata. This ‘direct approach’ is in contrast to the reduction to satisfiability in the μ -calculus used for UNFO in [47] which seems unwieldy in the presence of regular path expressions. Note in particular that an important reason for the relative simplicity of the reduction in [47] is that there is always a model of bounded treewidth in which any two bags overlap in at most one element; this is no longer true in UNFO^{reg}. To establish the CONP upper bound on data complexity, we first observe that it suffices to consider a database satisfiability problem (given a database D , is there a model of the fixed UNFO^{reg} sentence φ that extends D ?) and then establish a certain kind of decoration of D as a witness for D being a positive instance, in a way such that witnesses can be guessed and verified in polynomial time.

Related work. For general background on ontology-mediated querying, we refer to [15, 19, 37, 41] and the references therein. OMQ containment was considered in [11, 13, 14, 21]. UNFO was introduced and studied by ten Cate and Segoufin in [47] and it was considered as an ontology language for OMQs in [19]. Regular path queries, C2RPQs, and variations thereof emerge from graph databases, see the surveys [2, 10] and references therein. We use C2RPQs that admit nesting via node tests, as considered in [16], see also [20]. Sometimes, this is referred to as ‘nested’ C2RPQs. There are several further extensions of C2RPQs that are not considered in this paper. We still mention two of them. A more powerful form of nesting is obtained by allowing to use C2RPQs with two answer variables in place of binary predicates in regular expressions, giving rise to regular queries [42]. Another expressive extension of C2RPQs is defined by monadically defined queries, which implement a certain ‘flag and check’ paradigm [43]. OMQs in which the actual query is some form of regular path queries are considered in [12, 17, 18, 26, 27, 28, 40]. As discussed in more detail later, UNFO^{reg} is also related (but orthogonal in expressive power) to propositional dynamic logic with intersection and converse (ICPDL) [31] and to UNFO extended with fixed points [47].

2 Preliminaries

We assume that a countably infinite supply of predicate symbols of each arity is available. In the *unary negation fragment of first-order logic extended with regular path expressions* (UNFO^{reg}), formulas φ are formed according to the following grammar:

$$\begin{aligned} \varphi & ::= P(\mathbf{x}) \mid E(x, y) \mid x = y \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \exists x \varphi \mid \neg \varphi(x) \\ E & ::= R \mid R^- \mid E \cup E \mid E \cdot E \mid E^* \mid \varphi(x)? \end{aligned}$$

where P ranges over predicate symbols, R over binary predicate symbols, and, in the $\neg \varphi(x)$ clause, φ has no free variables besides (possibly) x . Expressions E formed according to the second line are called (*regular*) *path expressions* and expressions $\varphi(x)?$ according to the last clause in that line are called *tests*. Tests are similar to the test operator in propositional dynamic logic (PDL) [30] and to node tests in XPath [32] and in some versions of regular path queries [16, 20]. When we write $\varphi(\mathbf{x})$, we generally mean that the free variables of φ are among \mathbf{x} , but not all variables from \mathbf{x} need actually be free in φ . For a UNFO^{reg} formula $\varphi(x)$, we use $\forall x \varphi$ to abbreviate $\neg \exists x \neg \varphi(x)$.

► **Example 1.** The following are UNFO^{reg} formulas: $\forall x (\exists y R(x, y) \wedge \neg (R \cdot R^*)(x, x))$ and $\exists y (R^*(x, y) \wedge S^*(x, y))$.

A *structure* \mathfrak{A} takes the form $(A, R_1^{\mathfrak{A}}, R_2^{\mathfrak{A}}, \dots)$ where A is a non-empty set called the *domain* and R_i is an n_i -ary relation over A if R_i is a predicate symbol of arity n_i . Whenever convenient, we use $\text{dom}(\mathfrak{A})$ to refer to A . Every path expression E is interpreted as a binary relation $E^{\mathfrak{A}}$ over A : $R^{\mathfrak{A}}$ is part of \mathfrak{A} , $(R^-)^{\mathfrak{A}}$ is the converse of $R^{\mathfrak{A}}$, $(E_1 \cup E_2)^{\mathfrak{A}} = E_1^{\mathfrak{A}} \cup E_2^{\mathfrak{A}}$, $(E_1 \circ E_2)^{\mathfrak{A}} = E_1^{\mathfrak{A}} \circ E_2^{\mathfrak{A}}$, $(E^*)^{\mathfrak{A}}$ is the reflexive-transitive closure of $E^{\mathfrak{A}}$, and $(\varphi(x))^{\mathfrak{A}} = \{(a, a) \mid \mathfrak{A} \models \varphi(a)\}$. UNFO^{reg} formulas are then interpreted under the standard first-order semantics with path expressions being treated in the same way as binary predicates. A UNFO^{reg} sentence $\varphi(\mathbf{x})$ is *satisfiable* if there is a structure \mathfrak{A} such that $\mathfrak{A} \models \varphi$. Such an \mathfrak{A} is called a *model* of $\varphi(\mathbf{x})$.

► **Example 2.** Reconsider the UNFO^{reg} formulas from Example 1. It can be verified that the first sentence is satisfiable, but not in a finite model. Thus, in contrast to UNFO (and to propositional dynamic logic), UNFO^{reg} lacks the finite model property. The second sentence expresses a property that cannot be expressed in UNFO extended with fixed points, as studied in [47], which can formally be shown using UN-bisimulations, also defined in [47]. In fact, UNFO^{reg} and UNFO with fixed points are orthogonal in expressive power. Another related logic is ICPDL, that is, PDL extended with intersection and converse [31]. This logic, too, is orthogonal in expressive power to UNFO^{reg}. For example, the existence of a 4-clique can be expressed as a UNFO sentence, but not in ICPDL since every satisfiable ICPDL formula is satisfiable in a structure of tree width two.

The expressive power of UNFO^{reg} is closely related to that of conjunctive 2-way regular path queries. A *database* D is a finite structure such that for every $a \in \text{dom}(D)$, there is an $\mathbf{a} \subseteq \text{dom}(D)$ and a predicate symbol P such that $a \in \mathbf{a} \in P^D$. Since a database is a syntactic object, we refer to the elements of $\text{dom}(D)$ as *constants* whereas we speak about *elements* in the context of (semantic) structures. A *conjunctive 2-way regular path query* (C2RPQ) is a formula of the form $q(\mathbf{x}) = \exists \mathbf{y} \varphi(\mathbf{x}, \mathbf{y})$ where $\varphi(\mathbf{x}, \mathbf{y})$ is a conjunction of *atoms* of the form $R(\mathbf{z})$ and $E(z_1, z_2)$, R a predicate symbol and E a two-way regular path query, that is, an expression formed according to the second line of the syntax definition of UNFO^{reg}, but allowing only formulas $\varphi(x)$ that are C2RPQs in tests. The variables \mathbf{x} are the *answer variables* of $q(\mathbf{x})$ and $q(\mathbf{x})$ is *Boolean* if $\mathbf{x} = \emptyset$. A *union of C2RPQs* (UC2RPQ) is a disjunction of C2RPQs that all have the same answer variables. A *conjunctive query* (CQ) is a C2RPQ that does not use atoms of the form $E(z_1, z_2)$. The *answers* to a UC2RPQ $q(\mathbf{x})$ on a database D , denoted $\text{ans}(q, D)$, are defined in the standard way, see for example [42]. Note that every UC2RPQ is a UNFO^{reg} formula.

► **Example 3.** Consider the following database about family relationships, using binary predicates *Child* and *Spouse*, and written as a set of facts.

$$D = \left\{ \begin{array}{l} \text{Child}(\text{Nívea}, \text{Clara}), \text{Child}(\text{Clara}, \text{Blanca}), \text{Child}(\text{Blanca}, \text{Alba}), \\ \text{Spouse}(\text{Nívea}, \text{Severo}), \text{Spouse}(\text{Esteban}, \text{Clara}) \end{array} \right\}$$

The following C2RPQ asks for all pairs (x, y) such that x is an ancestor of y in a line of only married ancestors (using the shorthand $R^+ = R \cdot R^*$).

$$q(x, y) = (m(z)^? \cdot \text{Child})^+(x, y) \quad \text{where} \quad m(z) = \exists z' (\text{Spouse} \cup \text{Spouse}^-)(z, z')$$

We have $\text{ans}(q, D) = \{(\text{Nívea}, \text{Clara}), (\text{Nívea}, \text{Blanca}), (\text{Clara}, \text{Blanca})\}$.

Let $q(\mathbf{x}) = \exists \mathbf{y} \varphi(\mathbf{x}, \mathbf{y})$ be a C2RPQ. We use $\text{var}(q)$ to denote the variables that occur in q outside of tests, that is, $\mathbf{x} \cup \mathbf{y}$. We do not distinguish between $q(\mathbf{x})$ and the set of all atoms in φ , writing e.g. $R(x, y, z) \in q(\mathbf{x})$ to mean that $R(x, y, z)$ is an atom in φ . For simplicity,

we treat an atom $E(x, x)$ in a C2RPQ $q(\mathbf{x})$ where E is the test $\varphi(y)$? as an atom of the form $\varphi(x)$; that is, w.l.o.g. we use tests not only in path expressions but also directly as atoms of a C2RPQ. A C2RPQ $q(\mathbf{x})$ can be viewed as a finite hypergraph in the expected way, that is, every atom $R(\mathbf{z})$ and $E(z_1, z_2)$ is viewed as a hyperedge. We say that $q(\mathbf{x})$ is *connected* if the Gaifman graph of this hypergraph is connected. It is interesting to observe that foundational problems concerning UC2RPQs can be phrased as (un)satisfiability problems in UNFO^{reg}.

► **Example 4.**

1. The problem whether a Boolean UC2RPQ $q()$ evaluates to true on a database D (i.e., whether the empty tuple is in $\text{ans}(q, D)$) corresponds to the unsatisfiability of $\varphi_D() \wedge \neg q()$ where $\varphi_D()$ is D viewed as a Boolean CQ in the obvious way.
2. The problem whether a Boolean UC2RPQ $q_1()$ is contained in a Boolean UC2RPQ $q_2()$ (defined in the usual way) corresponds to the unsatisfiability of $q_1() \wedge \neg q_2()$.

Both reductions extend to the case of non-Boolean queries by simulating answer variables using fresh unary predicates, see the proof of Lemma 6 below.

An *ontology-mediated query (OMQ)* is a triple (\mathcal{O}, Σ, q) where \mathcal{O} is a logical sentence called the *ontology*, Σ is a set of predicate symbols called the *data signature*, and q is a query. In this paper, we shall primarily be interested in the case where \mathcal{O} is an UNFO^{reg} sentence and q is a UC2RPQ. We use $(\text{UNFO}^{\text{reg}}, \text{UC2RPQ})$ to denote the set of OMQs of this form and similarly for other ontology languages and query languages. Let $Q = (\mathcal{O}, \Sigma, q)$ be from $(\text{UNFO}^{\text{reg}}, \text{UC2RPQ})$ and D a database that uses only symbols from Σ . We call $\mathbf{a} \subseteq \text{dom}(D)$ a *certain answer to Q on D* if $\mathbf{a} \in \text{ans}(q, \mathfrak{A})$ for every structure \mathfrak{A} that extends D and is a model of \mathcal{O} , where \mathfrak{A} *extends D* if $\text{dom}(D) \subseteq \text{dom}(\mathfrak{A})$ and $P^D \subseteq P^{\mathfrak{A}}$ for all predicate symbols P . Note that this semantics embodies a ‘standard names assumption’, that is, constants in D are interpreted as themselves. The set of all certain answers to Q on D is denoted $\text{cert}(Q, D)$. We say that Q is *Boolean* if q is. For a Boolean OMQ Q , we write $D \models Q$ to indicate that Q is true on D , meaning that the empty tuple is in $\text{cert}(Q, D)$.

► **Example 5.** Consider the OMQ $Q = (\mathcal{O}, \Sigma, q')$ based on an extension of the C2RPQ q from Example 3, where \mathcal{O} defines a single mother as an unmarried woman who has a child, using additional unary predicates *Female* and *SingleMother*, and q' has an additional conjunct requiring that y is a single mother, that is:

$$\begin{aligned} \mathcal{O} &= \forall x (\text{SingleMother}(x) \leftrightarrow \text{Female}(x) \wedge \text{Single}(x) \wedge \exists y \text{Child}(x, y)) \\ q'(x, y) &= q(x, y) \wedge \text{SingleMother}(y) \\ \Sigma &= \{\text{Child}, \text{Spouse}, \text{Female}, \text{Single}\} \end{aligned}$$

Note that \mathcal{O} is equivalent to a UNFO^{reg} (even plain UNFO) formula obtained by eliminating \leftrightarrow in the usual way. Let $D' = D \cup \{\text{Female}(\text{Blanca}), \text{Single}(\text{Blanca})\}$, where D is the database from Example 3. Then $\text{cert}(Q, D') = \{(\text{Nívea}, \text{Blanca}), (\text{Clara}, \text{Blanca})\}$, but $\text{cert}(Q, D) = \emptyset$.

OMQ evaluation in $(\text{UNFO}^{\text{reg}}, \text{UC2RPQ})$ is the problem to decide, given an OMQ Q from $(\text{UNFO}^{\text{reg}}, \text{UC2RPQ})$, a database D , and an $\mathbf{a} \subseteq \text{dom}(D)$, whether $\mathbf{a} \in \text{cert}(Q, D)$. This is a relevant problem since ontologies formulated in many logics used as ontology languages can be translated into an equivalent UNFO^{reg} sentence in polynomial time. In particular, this is the case for the basic description logics \mathcal{ALC} and \mathcal{ALCI} [19] and for their extensions with transitive closure of roles [3] and with regular expressions over roles [24]. For any of these logics \mathcal{L} , this of course also yields a polynomial time reduction of OMQ evaluation in $(\mathcal{L}, \text{UC2RPQ})$ to OMQ

evaluation in $(\text{UNFO}^{\text{reg}}, \text{UC2RPQ})$. Even UNFO itself has occasionally been considered as an ontology language [19].

For the rather common extension of the description logic \mathcal{ALC} with transitive roles [5], an equivalence preserving translation of ontologies into UNFO^{reg} sentences is not possible since UNFO^{reg} cannot enforce that a binary predicate is transitive. However, a transitive role can be simulated using the transitive closure of a binary predicate R (and never using R without transitive closure). In this way, one still obtains the desired polynomial time reduction of OMQ evaluation. The same reduction can be applied even to UNFO^{reg} extended with transitive relations. We use $\text{UNFO}_{\text{trans}}^{\text{reg}}$ to denote the extension of UNFO^{reg} where sentences take the form $\varphi_{\text{trans}} \wedge \varphi$ with φ_{trans} a conjunction of atoms of the form $\text{trans}(R)$, R a binary predicate symbol, and φ a UNFO^{reg} sentence. An atom $\text{trans}(R)$ is satisfied in a structure \mathfrak{A} if $R^{\mathfrak{A}}$ is transitive.

Evaluation of Boolean OMQs in $(\text{UNFO}_{\text{trans}}^{\text{reg}}, \text{UC2RPQ})$ reduces in polynomial time to satisfiability in $\text{UNFO}_{\text{trans}}^{\text{reg}}$ since $D \models (\mathcal{O}, \Sigma, q)$ iff $\varphi_D() \wedge \mathcal{O} \wedge \neg q()$ is unsatisfiable. The reduction can be extended to non-Boolean queries by simulating answer variables using fresh unary predicates. Because of this observation, in the main body of the paper we concentrate on deciding satisfiability rather than OMQ evaluation.

► **Lemma 6.** *OMQ evaluation in $(\text{UNFO}_{\text{trans}}^{\text{reg}}, \text{UC2RPQ})$ reduces in polynomial time to satisfiability in UNFO^{reg} , and so does satisfiability in $\text{UNFO}_{\text{trans}}^{\text{reg}}$.*

Together with Theorem 14 it thus follows that UNFO can be extended with transitive relations without losing decidability or affecting the complexity of satisfiability and of OMQ evaluation. This is in contrast to the guarded fragment, where in both cases decidability can only be obtained by adopting additional syntactic restrictions. While for satisfiability it suffices to assume that transitive relations are only used in guard positions, even stronger restrictions are necessary for OMQ evaluation [35, 46, 34].

There are also other interesting reasoning problems that can be reduced to satisfiability in UNFO^{reg} . Here we consider OMQ containment, leaving out transitive roles for simplicity. Let $Q_1 = (\mathcal{O}, \Sigma_{\text{full}}, q_1)$ and $Q_2 = (\mathcal{O}, \Sigma_{\text{full}}, q_2)$ be OMQs from $(\text{UNFO}^{\text{reg}}, \text{UC2RPQ})$ with the same number of answer variables and where Σ_{full} is the *full* data signature, that is, the set of all predicate symbols. We say that Q_1 is contained in Q_2 and write $Q_1 \subseteq Q_2$ if for every database D , $\text{cert}(Q_1, D) \subseteq \text{cert}(Q_2, D)$. We observe that OMQ containment can also be reduced to satisfiability in polynomial time.

► **Lemma 7.** *OMQ containment in $(\text{UNFO}^{\text{reg}}, \text{UC2RPQ})$ reduces in polynomial time to satisfiability in UNFO^{reg} .*

There are also versions of OMQ containment that admit different ontologies in the two involved OMQs and more restricted data signatures in place of Σ_{full} [11, 13, 14, 21]. These are computationally harder and a polynomial time reduction to satisfiability cannot be expected. In fact, it follows from results in [21] that this more general form of OMQ containment is 2NEXPTIME-hard already when the ontologies are formulated in the description logic \mathcal{ALCI} , a fragment of UNFO, and when the actual queries are CQs. Decidability remains an open problem. We remark that when the actual queries in OMQs are CQs, then OMQ containment under the full data signature can be reduced to query evaluation in a straightforward way, essentially by viewing the query from the left-hand OMQ as a database. In the presence of regular path queries, however, this does not seem to be easily possible.

We next introduce a normal form for UNFO^{reg} sentences, similar but not identical to the normal form used for UNFO in [47]. For a set \mathcal{L} of UNFO^{reg} formulas with one free variable,

a *C2RPQ extended with \mathcal{L} -formulas* is a C2RPQ in which all tests $\varphi(x)?$ in atoms $E(z_1, z_2)$ have been replaced with tests $\psi(x)?$, $\psi(x)$ a formula from \mathcal{L} . The set of *normal UNFO^{reg} formulas* is the smallest set of formulas such that

1. every connected C2RPQ with exactly one free variable, extended with normal UNFO^{reg} formulas, is a normal UNFO^{reg} formula;
2. if $\varphi(x)$ and $\psi(x)$ are normal UNFO^{reg} formulas, then $\neg\varphi(x)$, $\varphi(x) \vee \psi(x)$, and $\exists x \varphi(x)$ are normal UNFO^{reg} formulas.

Observe that Item 1 serves as an induction start since every connected C2RPQ without tests (and with one free variable) is a normal UNFO^{reg} formula. Note that normal formulas are closed under conjunction in the sense that the conjunction of normal formulas $\varphi_1(x)$ and $\varphi_2(x)$ is a C2RPQ extended with normal UNFO^{reg} formulas and thus a normal formula. Thus, unary disjunction could be eliminated, but for our purposes it is more convenient to keep it. We note in passing that using this normal form, it is easy to observe that UNFO^{reg} has the same expressive power as C2RPQs that admit both tests and negated tests.

The *width* of a normal UNFO^{reg} formula is the maximal number of variables in a C2RPQ that occurs in it (not counting the variables that occur in the C2RPQ only inside tests). The *atom width* is defined analogously, but referring to the number of atoms instead of the number of variables. In the context of normal UNFO^{reg} formulas, for brevity we speak of C2RPQs when meaning C2RPQs extended with normal UNFO^{reg} formulas. The *size* of a UNFO^{reg} formula is the number of symbols needed to write it, with variable symbols and predicate symbols being counted as a single symbol.

► **Lemma 8.** *Every UNFO^{reg} sentence φ can be transformed into an equivalent normal UNFO^{reg} sentence φ' in single exponential time. Moreover, the width and the atom width of φ' are at most polynomial in the size of φ and the path expressions that occur in φ' are exactly those in φ .*

In the following sections, we replace atoms $E(z_1, z_2)$ in the C2RPQs that occur in a normal UNFO^{reg} formula with atoms of the form $\mathcal{A}(z_1, z_2)$ where \mathcal{A} is a *nondeterministic automaton on finite words (NFA)* over a suitable alphabet; we call such atoms *NFA atoms*. Formally, an NFA is a tuple $(Q, \Sigma, \Delta, q_0, F)$ where Q is a finite set of states, Σ a finite alphabet, $\Delta \subseteq Q \times \Sigma \times Q$ a transition relation, $q_0 \in Q$ an initial state and $F \subseteq Q$ a set of final states. When deciding the satisfiability of a UNFO^{reg} sentence φ_0 , we will generally take Σ to be $\{R, R^- \mid R \text{ a binary predicate in } \varphi_0\} \cup \{\varphi(x)? \mid \varphi(x)? \text{ a test in } \varphi_0\}$. Clearly, all path expressions in φ_0 are regular expressions over this alphabet. Since every regular expression can be converted into an equivalent NFA in polynomial time, we can thus w.l.o.g. assume the NFA-based presentation. Let $\mathcal{A} = (Q, \Sigma, \Delta, q_0, F)$ be an NFA. Then we use $\mathcal{A}[F/F']$ to denote the NFA obtained from \mathcal{A} by replacing F with $F' \subseteq Q$ and $\mathcal{A}[q_0/q]$ for the NFA obtained from \mathcal{A} by replacing q_0 with $q \in Q$. For a structure \mathfrak{A} , an NFA \mathcal{A} , and $a, b \in A$, we write $\mathfrak{A} \models \mathcal{A}(a, b)$ if there are $a_1, \dots, a_n \in A$ and a word $w \in L(\mathcal{A})$ of length $n - 1$ such that $a = a_1$, $a_n = b$, $(a_i, a_{i+1}) \in R^{\mathfrak{A}}$ if the i th symbol in w is R , $(a_{i+1}, a_i) \in R^{\mathfrak{A}}$ if the i th symbol in w is R^- , and $a_i = a_{i+1}$ and $\mathfrak{A} \models \varphi(a_i)$ if the i th symbol in w is $\varphi(x)?$. This gives a semantics to NFA atoms. The *size* of a normal UNFO^{reg} formula with NFA atoms is defined in the same way as the size of a UNFO formula, where every NFA $\mathcal{A} = (Q, \Sigma, \Delta, q_0, F)$ contributes the cardinality of Q plus the cardinality of Δ plus the cardinality of F .

3 Tree-like Structures and Witness Trees

We give a characterization of satisfiability in UNFO^{reg} that is tailored towards implementation by tree automata. In particular, we show that every satisfiable UNFO^{reg} formula φ has

a model whose treewidth is bounded by the width of φ , introduce a representation of such models in terms of labeled trees, and characterize the satisfaction of C2RPQs in models represented in this way in terms of tree-shaped witnesses. To simplify the technical development, in this section and the subsequent one we disallow predicates of arity zero. Note that an atom $P()$ can be simulated by the formula $\exists x P(x)$, so this assumption is w.l.o.g. We work with normal UNFO^{reg} sentences throughout the section.

A (*directed*) tree is a prefix-closed subset $T \subseteq (\mathbb{N} \setminus \{0\})^*$. A node $w \in T$ is a *successor* of $v \in T$ and v is a *predecessor* of w if $w = v \cdot i$ for some $i \in \mathbb{N}$. Moreover, w is a *neighbor* of v if it is a successor or predecessor of v . A *tree-like structure* is a pair (T, \mathbf{bag}) where T is a tree and \mathbf{bag} a function that assigns to every $w \in T$ a finite structure $\mathbf{bag}(w)$ such that

$$\text{the set of nodes } \{w \in T \mid a \in \text{dom}(w)\} \text{ is connected in } T, \text{ for each } a \in \bigcup_{w \in T} \text{dom}(w)$$

where, here and in the remainder of the paper, $\text{dom}(w)$ is a shorthand for $\text{dom}(\mathbf{bag}(w))$. The *width* of (T, \mathbf{bag}) is the maximum domain size of structures that occur in the range of \mathbf{bag} . Its *outdegree* is the outdegree of T . A tree-like structure (T, \mathbf{bag}) defines the associated structure $\mathfrak{A}_{(T, \mathbf{bag})}$ which is the (non-disjoint) union of all structures $\mathbf{bag}(w)$, $w \in T$. We use $\text{dom}(T, \mathbf{bag})$ as a shorthand for $\text{dom}(\mathfrak{A}_{(T, \mathbf{bag})})$. As witnessed by its representation (T, \mathbf{bag}) , the treewidth of the structure $\mathfrak{A}_{(T, \mathbf{bag})}$ is bounded by the maximum cardinality of $\text{dom}(\mathbf{bag}(w))$, $w \in T$. We will show that every satisfiable UNFO^{reg} sentence φ_0 is satisfiable in a tree-like structure whose width is bounded by the width of φ_0 . In UNFO, it suffices to consider structures of this form in which bags overlap in at most one element; this is not the case in UNFO^{reg}.

Let φ_0 be a normal UNFO^{reg} sentence. We use $\text{sub}(\varphi_0)$ to denote the subformulas of φ_0 with at most one free variable, and where the free variable is renamed to x . Then $\text{cl}(\varphi_0)$ denotes the smallest set of normal UNFO^{reg} formulas that contains $\text{sub}(\varphi_0)$ and is closed under single negation. A *1-type for φ_0* is a subset $t \subseteq \text{cl}(\varphi_0)$ that satisfies the following conditions:

1. $\varphi \in t$ iff $\neg\varphi \notin t$ for all $\neg\varphi \in \text{cl}(\varphi_0)$;
2. $\varphi \vee \psi \in t$ iff $\varphi \in t$ or $\psi \in t$ for all $\varphi \vee \psi \in \text{cl}(\varphi_0)$.

We use $\text{TP}(\varphi_0)$ to denote the set of all 1-types for φ_0 .

A *type decorated tree-like structure for φ_0* is a triple (T, \mathbf{bag}, τ) with (T, \mathbf{bag}) a tree-like structure such that only predicates from φ_0 occur in the range of \mathbf{bag} and $\tau : \text{dom}(T, \mathbf{bag}) \rightarrow \text{TP}(\varphi_0)$. Let (T, \mathbf{bag}, τ) be such a structure, \mathcal{A} an NFA, and $a, b \in \text{dom}(T, \mathbf{bag})$. We write $\mathfrak{A}_{(T, \mathbf{bag}), \tau} \models \mathcal{A}(a, b)$ if $\mathfrak{A}_{(T, \mathbf{bag})} \models \mathcal{A}(a, b)$ with the semantics of tests reinterpreted: instead of demanding that $\mathfrak{A} \models \varphi(a')$ for a test $\varphi_0(x)?$ to hold at an element a' , we now require that $\varphi \in \tau(a')$. Let $\varphi(x) = \exists \mathbf{y} \psi(x, \mathbf{y})$ be a C2RPQ and $a \in \text{dom}(T, \mathbf{bag})$. A *homomorphism from $\varphi(x)$ to (T, \mathbf{bag}, τ)* is a function $h : \{x\} \cup \mathbf{y} \rightarrow \text{dom}(T, \mathbf{bag})$ such that the following conditions are satisfied:

- $h(\mathbf{x}) \in R^{\mathfrak{A}_{(T, \mathbf{bag})}}$ for each $R(\mathbf{x}) \in \varphi(x)$;
- $\mathfrak{A}_{(T, \mathbf{bag}), \tau} \models \mathcal{A}(h(y), h(z))$ for each $\mathcal{A}(y, z) \in \varphi(x)$.

A type decorated tree-like structure (T, \mathbf{bag}, τ) for φ_0 is *proper* if:

1. for all $\exists x \varphi(x) \in \text{cl}(\varphi_0)$, $\exists x \varphi(x) \in \tau(a)$ iff there is a $b \in \text{dom}(T, \mathbf{bag})$ with $\varphi(x) \in \tau(b)$;
2. for all C2RPQs $\varphi(x) \in \text{cl}(\varphi_0)$, $\varphi(x) \in \tau(a)$ iff there is a homomorphism h from $\varphi(x)$ to (T, \mathbf{bag}, τ) such that $h(x) = a$.

The following lemma establishes proper type decorated tree-like structures for φ_0 as witnesses for the satisfiability of φ_0 . The proof of the ‘only if’ direction is via an unraveling procedure that constructs a type decorated tree-like structure in a top-down manner, introducing

fresh bags to satisfy C2RPQs and to implement a step-by-step chase of paths that witness satisfaction of NFA atoms in C2RPQs.

► **Lemma 9.** *A normal UNFO^{reg} sentence φ_0 of size n and width m is satisfiable iff there is a proper type decorated tree-like structure (T, \mathbf{bag}, τ) for φ_0 of width at most m and outdegree at most $n^2 + n$ such that $\varphi_0 \in \tau(a)$ for some a .*

As the next step, we take a closer look at Point 2 of properness, that is, we characterize carefully the existence of a homomorphism h from $\varphi(x)$ to (T, \mathbf{bag}, τ) such that $h(x) = a$ in a way that is tailored towards implementation by tree automata. This gives rise to the notion of a witness tree below. We start with introducing the notions of subdivisions and splittings which shall help us to take care of the fact that the homomorphic image of a query $q(\mathbf{x})$ may be spread over several bags of a tree-like structure, and in fact this might even be the case for a single NFA atom.

An *instantiated C2RPQ* is a C2RPQ in which all free variables have been replaced with constants. We write $\varphi(\mathbf{a})$ to indicate that the constants in the instantiated C2RPQ are exactly \mathbf{a} . When working with instantiated C2RPQs, we drop existential quantifiers, assuming that all variables are implicitly existentially quantified. For brevity, we often omit the word ‘instantiated’ and only speak of C2RPQs. We speak of *terms* to mean both variables and constants, and we denote terms with t .

Let $\varphi(\mathbf{a})$ be a connected C2RPQ, Δ be a domain, and $s \geq 1$. A (Δ, s) -*subdivision* of of an atom $\mathcal{A}(t, t') \in \varphi(\mathbf{a})$ is a set of atoms

$$\mathcal{A}[F/\{q_1\}](t, b_1), \mathcal{A}[q_0/q_1, F/\{q_2\}](b_1, b_2), \dots, \mathcal{A}[q_0/q_{k-1}, F/\{q_k\}](b_{k-1}, b_k), \mathcal{A}[q_0/q_k](b_k, t')$$

where q_1, \dots, q_k are states of \mathcal{A} , $k \leq s$, and b_1, \dots, b_k are constants from Δ . A C2RPQ $\psi(\mathbf{a}')$ is a (Δ, s) -*subdivision* of $\varphi(\mathbf{a})$ if it is obtained from $\varphi(\mathbf{a})$ by replacing zero or more NFA atoms with (Δ, s) -subdivisions. Let $\psi(\mathbf{a}')$ be a (Δ, s) -subdivision of $\varphi(\mathbf{a})$. A *splitting* of $\psi(\mathbf{a}')$ is a sequence $\psi_0(\mathbf{a}_0), \dots, \psi_\ell(\mathbf{a}_\ell)$, $\ell \geq 0$, of C2RPQs that is a partition of $\psi(\mathbf{a}')$ (viewed as a set of atoms) where we also allow the special case that $\psi_0(\mathbf{a}_0)$ is empty (and thus $\psi_1(\mathbf{a}_1), \dots, \psi_\ell(\mathbf{a}_\ell)$ is the actual partition). We require that the following conditions are satisfied:

1. $\psi_1(\mathbf{a}_1), \dots, \psi_\ell(\mathbf{a}_\ell)$ are connected;
2. $\text{var}(\psi_i(\mathbf{a}_i)) \cap \text{var}(\psi_j(\mathbf{a}_j)) \subseteq \text{var}(\psi_0(\mathbf{a}_0))$ for $1 \leq i < j \leq \ell$;
3. each of $\psi_1(\mathbf{a}_1), \dots, \psi_\ell(\mathbf{a}_\ell)$ contains at most one atom from each subdivision of an atom in $\varphi(\mathbf{a})$.

Intuitively, the $\vartheta_0(\mathbf{a})$ component of a splitting is the part of $\psi(\mathbf{a}')$ that maps into a bag that we are currently focussing on while the other components are pushed to neighboring bags.

► **Example 10.** Consider $q(a) = \{\mathcal{A}(a, y), T(a, z), Q(a, y, z)\}$ with $\mathcal{A} = \rightarrow \textcircled{0} \xrightarrow{R} \textcircled{1} \curvearrowright R, S$.

Let $\Delta = \{a, b, c\}$ and $\mathcal{A}_{ij} = \mathcal{A}[0/i, F/j]$. An example for a $(\Delta, 2)$ -subdivision of $\mathcal{A}(a, y)$ is $\{\mathcal{A}_{01}(a, b), \mathcal{A}_{11}(b, b), \mathcal{A}_{11}(b, y)\}$, which yields the following $(\Delta, 2)$ -subdivision of $\varphi(a)$: $\psi(a, b) = \{\mathcal{A}_{01}(a, b), \mathcal{A}_{11}(b, b), \mathcal{A}_{11}(b, y), T(a, z), Q(a, y, z)\}$. $\psi(a, b)$ admits a splitting into ψ_0, ψ_1 as follows: $\psi_0(a, b) = \{\mathcal{A}_{01}(a, b), \mathcal{A}_{11}(b, b), T(a, z)\}$ and $\psi_1(a, b) = \{\mathcal{A}_{11}(b, y), Q(y, z, a)\}$.

The *query closure* $\text{qcl}(\varphi_0, \Delta, s)$ is defined as the smallest set such that the following conditions are satisfied:

- if $\varphi(x) \in \text{cl}(\varphi_0)$ is a C2RPQ and $a \in \Delta$, then $\varphi(a) \in \text{qcl}(\varphi_0, \Delta, s)$;

15:10 Querying the Unary Negation Fragment with Regular Path Expressions

- if $\varphi(\mathbf{a}) \in \text{qcl}(\varphi_0, \Delta, s)$, $\psi(\mathbf{a}')$ is a (Δ, s) -subdivision of $\varphi(\mathbf{a})$, $\psi_0(\mathbf{a}_0), \dots, \psi_\ell(\mathbf{a}_\ell)$ is a splitting of $\psi(\mathbf{a}')$, $1 \leq i \leq \ell$, and $\psi'_i(\mathbf{a}'_i)$ is obtained from $\psi_i(\mathbf{a}_i)$ by consistently replacing zero or more variables with constants from Δ , then $\psi'_i(\mathbf{a}'_i) \in \text{qcl}(\varphi_0, \Delta, s)$.

► **Lemma 11.** *The cardinality of $\text{qcl}(\varphi_0, \Delta, s)$ is bounded by $p \cdot (a^2 d^m)^{m'}$, where p is the number of C2RPQs in φ_0 , a the maximal number of states in an NFA in φ_0 , d the cardinality of Δ , m the width of φ_0 , and m' the atom width of φ_0 .*

We are almost ready to define witness trees. The following notion of a homomorphism is more local than the ones used so far as it only concerns a single bag rather than the entire tree-like structure. Let (T, bag, τ) be a type decorated tree-like structure, $w \in T$, \mathcal{A} an NFA, and $a, b \in \text{dom}(w)$. We write $\text{bag}(w), \tau \models \mathcal{A}(a, b)$ if $\text{bag}(w) \models \mathcal{A}(a, b)$ with the semantics of tests reinterpreted: instead of demanding $\text{bag}(w) \models \varphi(a')$ for a test $\varphi(x)?$ to hold at an element a' , we now require that $\varphi(x) \in \tau(a')$. Let $\varphi(\mathbf{a})$ be a C2RPQ. A *homomorphism from $\varphi(\mathbf{a})$ to $\text{bag}(w)$ given τ* is a function $h : \mathbf{a} \cup \text{var}(\varphi) \rightarrow \text{dom}(w)$ such that the following conditions are satisfied:

- $h(a) = a$ for each $a \in \mathbf{a}$;
- $h(\mathbf{t}) \in R^{\text{bag}(w)}$ for each $R(\mathbf{t}) \in \varphi(\mathbf{a})$;
- $\text{bag}(w), \tau \models \mathcal{A}(h(t), h(t'))$ for each $\mathcal{A}(t, t') \in \varphi(a)$.

Let n be the size of φ_0 , $a \in \text{dom}(T, \text{bag})$, and $\varphi(x) \in \text{cl}(\varphi_0)$ a C2RPQ. A *witness tree for $\varphi(a)$ in (T, bag, τ)* is a finite labeled tree (W, σ) with $\sigma : W \rightarrow T \times \text{qcl}(\varphi_0, \text{dom}(T, \text{bag}), n^2)$ such that the root is labeled with $\sigma(\varepsilon) = (w, \varphi(a))$ for some $w \in T$ with $a \in \text{dom}(w)$ and the following conditions are satisfied for all $u \in W$:

- (*) if $\sigma(u) = (w, \psi(\mathbf{a}))$, then there is a $(\text{dom}(w), n^2)$ -subdivision $\vartheta'(\mathbf{a})$ of $\psi(\mathbf{a})$, a splitting $\vartheta_0(\mathbf{a}_0), \dots, \vartheta_\ell(\mathbf{a}_\ell)$ of $\vartheta'(\mathbf{a})$, a homomorphism h from $\vartheta_0(\mathbf{a}_0)$ to $\text{bag}(w)$ given τ , and successors u_1, \dots, u_ℓ of u such that $\sigma(u_i) = (w_i, \vartheta'_i(\mathbf{a}'_i))$ for $1 \leq i \leq \ell$, where each w_i is a neighbor of w in T with $\mathbf{a}'_i \subseteq \text{dom}(w_i)$ and $\vartheta'_i(\mathbf{a}'_i)$ is obtained from $\vartheta_i(\mathbf{a}_i)$ by replacing each variable x in the domain of h with the constant $h(x)$.

Informally, a witness tree decomposes a homomorphism h from $\varphi(x)$ to $\mathfrak{A}_{T, \text{bag}}$ into local ‘chunks’, each of which concerns only a single bag. In particular, the splitting $\vartheta_0(\mathbf{a}_0), \dots, \vartheta_k(\mathbf{a}_k)$ in (*) breaks the current C2RPQ down into components that are satisfied in different parts of the tree-like structure. We need to first subdivide since satisfaction of NFA atoms is witnessed by an entire path, and this path can pass through the current node several times. Fortunately, the number of points introduced in a subdivision can be bounded: we can w.l.o.g. choose a shortest path and such a path can pass through w at most once for each element in $\text{dom}(w)$ and each state of the automaton \mathcal{A} , thus we need at most n^2 points in subdivisions.

► **Lemma 12.** *Let (T, bag, τ) be a type decorated tree-like structure, $\varphi(x) \in \text{cl}(\varphi_0)$ a C2RPQ, and $a \in \text{dom}(T, \text{bag})$. Then there is a homomorphism h from $\varphi(x)$ to (T, bag, τ) with $h(x) = a$ iff there is a witness tree for $\varphi(a)$ in (T, bag, τ) .*

4 Automata-Based Decision Procedure

We now reduce satisfiability of UNFO^{reg} sentences to the nonemptiness problem of two-way alternating tree automata. We start with recalling this automata model and discuss the encoding of tree-like structures as an input to automata.

Two-way alternating tree automata. A tree is *k-ary* if each node has *exactly* k successors. As a convention, we set $w \cdot 0 = w$ and $wc \cdot (-1) = w$, leave $\varepsilon \cdot (-1)$ undefined, and for any $k \in \mathbb{N}$, set $[k] = \{-1, 0, \dots, k\}$. Let Σ be a finite alphabet. A Σ -labeled tree is a pair (T, L) with T a tree and $L : T \rightarrow \Sigma$ a node labeling function.

An *alternating 2-way tree automaton (2ATA)* over Σ -labeled k -ary trees is a tuple $\mathcal{A} = (Q, \Sigma, q_0, \delta, F)$ where Q is a finite set of *states*, $q_0 \in Q$ is an *initial state*, δ is the *transition function*, and F is the (*parity*) *acceptance condition*, that is, a finite sequence G_1, \dots, G_k with $G_1 \subseteq G_2 \subseteq \dots \subseteq G_k = Q$. The transition function maps a state q and an input letter $a \in \Sigma$ to a positive Boolean formula over the constants `true` and `false`, and variables from $[k] \times Q$. The semantics is given in terms of runs in the appendix of the long version. As usual, $L(\mathcal{A})$ denotes the set of trees accepted by \mathcal{A} . The *nonemptiness problem* for 2ATAs is the problem to decide, given a 2ATA \mathcal{A} , whether $L(\mathcal{A})$ is nonempty. It can be solved in time single exponential in the number of states and the number of sets in the parity condition, and linear in the size of the transition function [48].

Encoding of tree-like structures. Let φ_0 be a normal UNFO^{reg} sentence whose satisfiability we want to decide. By Lemma 9, this corresponds to deciding the existence of a proper type decorated tree-like structure for φ_0 (of certain dimensions) and thus our aim is to build a 2ATA \mathcal{A} such that $L(\mathcal{A}) \neq \emptyset$ if and only if there is such a structure. 2ATAs cannot run directly on tree-like structures because the labeling of the underlying trees is not finite: we have already shown that UNFO^{reg} does not have the finite model property and thus it might be necessary that infinitely many elements occur in the bags. We therefore use an appropriate encoding that ‘reuses’ element names so that we can make do with finitely many element names overall, similar to what has been done, for example, in [36, 1].

Let R_1, \dots, R_ℓ be the predicate symbols that occur in φ_0 and let m be the width of φ_0 . Fix a finite set Δ with $2m$ elements and define Σ to be the set of all pairs (\mathbf{bag}, τ) such that $\mathbf{bag} = (A, R_1^{\mathbf{bag}}, \dots, R_\ell^{\mathbf{bag}})$ is a structure that satisfies $A \subseteq \Delta$ and $|A| \leq m$, and $\tau : A \rightarrow \text{TP}(\varphi_0)$ is a map that assigns a 1-type to every element in \mathbf{bag} .

Let (T, L) be a Σ -labeled tree. For convenience, we use \mathbf{bag}_w to refer to the first component of $L(w)$ and τ_w to refer to the second component, that is, $L(w) = (\mathbf{bag}_w, \tau_w)$. Moreover, dom_w is shorthand for $\text{dom}(\mathbf{bag}_w)$. For an element $d \in \Delta$, we say that $v, w \in T$ are *d-equivalent* if $d \in \text{dom}_u$ for all u on the unique shortest path from v to w . Informally, this means that d represents the same element in \mathbf{bag}_v and in \mathbf{bag}_w . In case that $d \in \text{dom}_w$, we use $[w]_d$ to denote the set of all v that are d -equivalent to w . We say that (T, L) is *type consistent* if, for all $d \in \Delta$ and all d -equivalent $v, w \in T$, $\tau_v(d) = \tau_w(d)$. Each type consistent (T, L) represents a type decorated tree-like structure $(T, \mathbf{bag}', \tau')$ of width at most m as follows. The domain of $\mathfrak{A}_{(T', \mathbf{bag}'')}$ is the set of all equivalence classes $[w]_d$ with $w \in T$ and $d \in \text{dom}_w$. The function τ' maps each domain element $[w]_d$ to $\tau_w(d)$, which is well-defined since (T, L) is type consistent. Finally, for every $w \in T$, the structure $\mathbf{bag}'(w) = (A(w), R_1^{\mathbf{bag}'(w)}, \dots, R_\ell^{\mathbf{bag}'(w)})$ is defined by:

$$\begin{aligned} A(w) &= \{[w]_d \mid d \in \text{dom}_w\}, \\ R_i^{\mathbf{bag}'(w)} &= \{([w]_{d_1}, \dots, [w]_{d_j}) \mid (d_1, \dots, d_j) \in R_i^{\mathbf{bag}_w}\} \quad \text{for } 1 \leq i \leq \ell. \end{aligned}$$

Conversely, for every type decorated tree-like structure (T, \mathbf{bag}, τ) of width m , there is a Σ -labeled tree (T, L) that represents a type decorated tree-like structure $(T, \mathbf{bag}', \tau')$ such that there is an isomorphism π between $\mathfrak{A}_{(T, \mathbf{bag})}$ and $\mathfrak{A}_{(T, \mathbf{bag}'')}$ that satisfies $\tau(d) = \tau'(\pi(d))$, for all $d \in \text{dom}(T, \mathbf{bag})$. In fact, since Δ is of size $2m$, it is possible to select a mapping $\pi : \text{dom}(T, \mathbf{bag}) \rightarrow \Delta$ such that for each $w \in T \setminus \{\varepsilon\}$ and each $d \in \text{dom}(w) \setminus \text{dom}(w \cdot -1)$, we have $\pi(d) \notin \{\pi(e) \mid e \in \text{dom}(w \cdot -1)\}$. Define the Σ -labeled tree (T, L) by setting, for all $w \in T$, \mathbf{bag}_w to the image of $\mathbf{bag}(w)$ under π and τ_w to the map defined by $\tau_w(h(d)) = \tau(d)$, for all $d \in \text{dom}_w$. Clearly, π satisfies the desired properties.

The notion of a witness tree carries over straightforwardly from type decorated tree-like structures to type consistent Σ -labeled trees. In fact, one only needs to replace τ with τ_w in Condition (*). Then, there is a witness tree for $\varphi(a)$ in a type consistent (T, L) iff there

is a witness tree for $\varphi(a)$ in the type decorated tree-like structure (T, bag', τ') represented by (T, L) . The notion of properness also carries over straightforwardly. For easier reference, we spell it out explicitly below, and also replace the homomorphisms from the original formulation by witness trees as suggested by Lemma 12. A type consistent Σ -labeled tree (T, L) is *proper* if for all $w \in T$ and $a \in \text{dom}_w$,

- 1'. for all $\exists x \varphi(x) \in \text{cl}(\varphi_0)$, $\exists x \varphi(x) \in \tau_w(a)$ iff there is a $v \in T$, $b \in \text{dom}_v$ with $\varphi(x) \in \tau_v(b)$;
 2'. for all C2RPQs $\varphi(x) \in \text{cl}(\varphi_0)$, $\varphi(x) \in \tau_w(a)$ iff there is a witness tree for $\varphi(a)$ in (T, L) .

It is straightforward to verify that (T, L) is proper iff the type decorated tree-like structure (T, bag', τ') represented by (T, L) is proper. Thus, our aim is to build a 2ATA \mathcal{A} that accepts exactly the proper type consistent Σ -labeled trees (T, L) such that $\varphi_0 \in \tau_w(a)$ for some $w \in T$ and $a \in \text{dom}_w$.

Automata construction. Let n be the size of φ_0 , $k = n^2 + n$ the bound on the outdegree from Lemma 9, and assume from now on that the automata run over k -ary Σ -labeled trees. It is straightforward to construct a 2ATA \mathcal{A}_0 that accepts (T, L) iff it is type consistent and satisfies Condition 1' of properness and the condition that $\varphi_0 \in \tau_w(a)$ for some $w \in T$ and $a \in \text{dom}_w$. The number of states of the automaton is linear in the size of φ_0 ; details are omitted. We next show how to construct a 2ATA $\mathcal{A}_1 = (Q, \Sigma, q_0, \delta, F)$ that accepts a type consistent (T, L) iff Condition 2' is satisfied. The automaton uses the set of states

$$Q = \{q_0\} \cup \{\varphi(\mathbf{a}), \bar{\varphi}(\mathbf{a}) \mid \varphi(\mathbf{a}) \in \text{qcl}(\varphi_0, \Delta, n^2)\}.$$

where states of the form $\varphi(\mathbf{a})$ are used to verify the ‘only if’ part of Condition 2' while states of the form $\bar{\varphi}(\mathbf{a})$ are used to verify the contrapositive of the ‘if’ part, that is, whenever a C2RPQ $\varphi(x) \in \text{cl}(\varphi_0)$ is not in $\tau_w(a)$, then there is no witness tree for $\varphi(a)$ in (T, L) .

Starting from the initial state, \mathcal{A}_1 loops over all nodes and domain elements using the following transitions, for all $(\text{bag}, \tau) \in \Sigma$:

$$\delta(q_0, (\text{bag}, \tau)) = \bigwedge_{1 \leq i \leq k} (i, q_0) \wedge \bigwedge_{a \in \text{dom}(\text{bag})} \left(\bigwedge_{\substack{\varphi(x) \in \tau(a), \\ \varphi(x) \text{ a C2RPQ}}} \varphi(a) \wedge \bigwedge_{\substack{\neg \varphi(x) \in \tau(a), \\ \varphi(x) \text{ a C2RPQ}}} \bar{\varphi}(a) \right)$$

We next give transitions for states of the form $\varphi(\mathbf{a}) \in \text{qcl}(\varphi_0, \Delta, n^2)$. Informally, if the automaton visits a node w in state $\varphi(\mathbf{a})$, then this is an obligation to show that there is a witness tree whose root is labeled with $(w, \varphi(\mathbf{a}))$. In particular, the automaton has to demonstrate that there are suitable successors for the root of the witness tree, implementing Condition (*). For a more concise definition of the transitions, we first establish a suitable notation. Let $\varphi(\mathbf{a}), \vartheta_1(\mathbf{a}_1), \dots, \vartheta_\ell(\mathbf{a}_\ell) \in \text{qcl}(\varphi_0, \Delta, n^2)$ and $(\text{bag}, \tau) \in \Sigma$. We write $\varphi(\mathbf{a}) \rightarrow_{(\text{bag}, \tau)} \vartheta_1(\mathbf{a}_1), \dots, \vartheta_\ell(\mathbf{a}_\ell)$ if there is a (Δ, n^2) -subdivision $\vartheta(\mathbf{a}')$ of $\varphi(\mathbf{a})$, a splitting $\vartheta'_0(\mathbf{a}'_0), \dots, \vartheta'_\ell(\mathbf{a}'_\ell)$ of $\vartheta(\mathbf{a}')$, a homomorphism h from $\vartheta'_0(\mathbf{a}'_0)$ to bag given τ , and $\vartheta_i(\mathbf{a}_i)$ is obtained from $\vartheta'_i(\mathbf{a}'_i)$ by replacing each variable x in the domain of h with the constant $h(x)$; please note that this is an essential part of Condition (*). Then, we include for each $\varphi(\mathbf{a}) \in \text{qcl}(\varphi_0, \Delta, n^2)$ and each $(\text{bag}, \tau) \in \Sigma$ the transition

$$\delta(\varphi(\mathbf{a}), (\text{bag}, \tau)) = \bigvee_{\varphi(\mathbf{a}) \rightarrow_{(\text{bag}, \tau)} \vartheta_1(\mathbf{a}_1), \dots, \vartheta_\ell(\mathbf{a}_\ell)} \bigwedge_{1 \leq i \leq \ell} \bigvee_{j \in [k] \setminus \{0\}} (j, \vartheta_i(\mathbf{a}_i))$$

if $\mathbf{a} \subseteq \text{dom}(\text{bag})$ and set $\delta(\varphi(\mathbf{a}), (\text{bag}, \tau)) = \text{false}$ otherwise. States of the form $\bar{\varphi}(\mathbf{a})$ are treated dually, that is, using the transitions

$$\delta(\bar{\varphi}(\mathbf{a}), (\text{bag}, \tau)) = \bigwedge_{\varphi(\mathbf{a}) \rightarrow_{(\text{bag}, \tau)} \vartheta_1(\mathbf{a}_1), \dots, \vartheta_\ell(\mathbf{a}_\ell)} \bigvee_{1 \leq i \leq \ell} \bigwedge_{j \in [k] \setminus \{0\}} (j, \bar{\vartheta}_i(\mathbf{a}_i))$$

if $\mathbf{a} \subseteq \text{dom}(\text{bag})$ and setting $\delta(\overline{\varphi}(\mathbf{a}), (\text{bag}, \tau)) = \text{true}$ otherwise.

To ensure that the witness trees constructed by the states of the form $\varphi(\mathbf{a})$ are finite, we use the parity condition $F = G_1, G_2$ with $G_1 = \text{qcl}(\varphi_0, \Delta, n^2)$ and $G_2 = Q$. From an accepting run of \mathcal{A}_1 on an input tree (T, L) , one can extract the witness trees that are required to show that the ‘only if’ direction of Condition 2’ is satisfied. Moreover, the run demonstrates that the witness trees forbidden by the ‘if’ direction do not exist. We thus obtain the following.

► **Lemma 13.** *The UNFO^{reg} sentence φ_0 is satisfiable iff $L(\mathcal{A}_0) \cap L(\mathcal{A}_1)$ is not empty.*

Putting together Lemmas 8, 11, and 13, it follows that satisfiability in UNFO^{reg} is in 2EXPTIME . The corresponding lower bound is inherited from UNFO [47].

► **Theorem 14.** *In UNFO^{reg} , satisfiability is 2EXPTIME -complete.*

5 OMQ Evaluation and Containment

We study the complexity of OMQ evaluation and OMQ containment in $(\text{UNFO}^{\text{reg}}, \text{UC2RPQ})$. Recall that the complexity of OMQ evaluation can be measured in different ways. In *combined complexity*, both the OMQ and the database on which it is evaluated are considered to be an input. In *data complexity*, the OMQ is fixed and the database is the only input. We first state our main result regarding the combined complexity of OMQ evaluation and the complexity of OMQ containment.

► **Theorem 15.** *In $(\text{UNFO}^{\text{reg}}, \text{UC2RPQ})$,*

1. *OMQ evaluation is 2EXPTIME -complete in combined complexity and*
2. *OMQ containment is 2EXPTIME -complete.*

The upper bounds in Theorem 15 are a consequence of Lemmas 6 and 7 and Theorem 14. The lower bounds hold already when predicates are at most binary. For Point 1 this follows from the fact that OMQ evaluation is 2EXPTIME -hard even for OMQs from the class $(\mathcal{ALCT}, \text{CQ})$ where the ontology is formulated in the description logic \mathcal{ALCT} , a fragment of UNFO with only unary and binary predicates, and the actual query is a CQ [39]. The same is true for Point 2 since in $(\mathcal{ALCT}, \text{CQ})$, OMQ evaluation can be reduced in polynomial time to OMQ containment in a straightforward way.

We next study the data complexity of $(\text{UNFO}^{\text{reg}}, \text{UC2RPQ})$. A coNP lower bound is again inherited from (rather small) fragments of $(\text{UNFO}^{\text{reg}}, \text{C2RPQ})$ [38, 23]. We give a coNP upper bound, thus establishing the following.

► **Theorem 16.** *OMQ evaluation in $(\text{UNFO}^{\text{reg}}, \text{UC2RPQ})$ is coNP -complete in data complexity.*

Instead of directly considering OMQ evaluation, we work with a problem that we call database satisfiability. A database D is *satisfiable* with an UNFO^{reg} sentence φ if there is a model of φ that extends D . Let φ be an UNFO^{reg} sentence and Σ a set of predicate symbols. The *database satisfiability problem associated with φ and Σ* is to decide, given a Σ -database D , whether D is satisfiable with φ . Note that OMQ evaluation can be reduced in polynomial time to Boolean OMQ evaluation as in the proof of Lemma 6. Moreover, for a Boolean OMQ $Q = (\mathcal{O}, \Sigma, q)$ and a Σ -database D , $D \not\models Q$ iff D is satisfiable with $\mathcal{O} \wedge \neg q$. Consequently, a coNP upper bound for OMQ evaluation in $(\text{UNFO}^{\text{reg}}, \text{UC2RPQ})$ can be proved by establishing an NP upper bound for database satisfiability in UNFO^{reg} .

Let φ_0 be an UNFO^{reg} formula and Σ a set of predicate symbols. We may assume w.l.o.g. that φ_0 is normal and that every symbol from Σ occurs in φ_0 . Subdivisions and splittings, defined as in Section 3, shall again play an important role. However, instead of subdividing an atom $\mathcal{A}(t, t')$ into at most n^2 many atoms, we use at most *two* intermediary points. Informally, this splits a witnessing path for $\mathcal{A}(t, t')$ into three parts: the first part is from t to the first element from D that appears on the path, the third subdivision atom represents the part from the last element from D that appears on the path to t' , and the second atom represents the remaining middle part of the path.

We use $\text{ecl}(\varphi_0)$ to denote the union of $\text{cl}(\varphi_0)$ and $\text{qcl}(\varphi_0)$, closed under single negation, where $\text{qcl}(\varphi_0)$ is $\text{qcl}(\varphi_0, \{x\}, 2)$ extended with the set of all $\mathcal{A}[q_0/s, F/\{s'\}](x, x)$ such that \mathcal{A} is an NFA that occurs in φ_0 and s, s' are states in \mathcal{A} . An *extended 1-type* for φ_0 is a subset $t \subseteq \text{ecl}(\varphi_0)$ such that t satisfies the conditions for being a 1-type from Section 3. We denote with $\text{eTP}(\varphi_0)$ the set of all extended 1-types for φ_0 .

Let D be a Σ -database. A *type decoration* for D is a mapping $\tau : \text{dom}(D) \rightarrow \text{eTP}(\varphi_0)$. We write $D, \tau \models \mathcal{A}(a, b)$ if $D \models \mathcal{A}(a, b)$ with the semantics of tests reinterpreted: instead of demanding $D \models \varphi(a')$ for a test $\varphi(x)?$ to hold at an element a' , we now require that $\varphi(x) \in \tau(a')$. Let $\varphi(\mathbf{a})$ be an (instantiated) C2RPQ. A *homomorphism from $\varphi(\mathbf{a})$ to D given τ* is a function $h : \mathbf{a} \cup \text{var}(\varphi) \rightarrow \text{dom}(D)$ such that the following conditions are satisfied: $h(\mathbf{a}) = \mathbf{a}$, $h(\mathbf{t}) \in R^D$ for each $R(\mathbf{t}) \in \varphi(\mathbf{a})$, and for each $\mathcal{A}(t, t') \in \varphi(\mathbf{a})$, there are $a_1, \dots, a_n \in \text{dom}(D)$ and states s_0, \dots, s_n from \mathcal{A} , and a word $\nu_1 \cdots \nu_{n-1}$ from the alphabet of \mathcal{A} such that

- (a) $a_1 = h(t)$, $a_n = h(t')$, $s_0 = q_0$, and $s_n \in F$,
- (b) $(a_i, a_{i+1}) \in R^D$ if $\nu_i = R$, $(a_{i+1}, a_i) \in R^D$ if $\nu_i = R^-$, and $\theta(x) \in \tau(a_i)$ and $a_{i+1} = a_i$ if $\nu_i = \theta(x)?$, for $1 \leq i < n$, and
- (c) $(s, \nu_i, s_{i+1}) \in \Delta$ for some s with $\mathcal{A}[q_0/s_i, F/\{s\}](x, x) \in \tau(a_{i+1})$, for $0 \leq i < n$.

Note that Condition (c) admits the spontaneous change from state s_i to state s at a_{i+1} , without reading any of the ν_j symbols, when the atom $\mathcal{A}[q_0/s_i, F/\{s\}](x, x)$ is contained in $\tau(a_{i+1})$, asserting that we can indeed get from s_i to s starting at a_{i+1} and cycling back there while reading some unknown subword.

A type decoration τ is called *proper* if for all $a \in \text{dom}(D)$, the following hold:

1. $\bigwedge_{\psi(x) \in \tau(a)} \psi(a)$ is satisfiable;
2. $\exists x \varphi(x) \in \tau(a)$ iff $\exists x \varphi(x) \in \tau(b)$, for all $a, b \in \text{dom}(D)$ and all $\exists x \varphi(x) \in \text{cl}(\varphi_0)$;
3. if $\neg \psi(x) \in \tau(a)$ for some $\psi(x) \in \text{qcl}(\varphi_0)$, then for each $(\text{dom}(D), 2)$ -subdivision $\vartheta(\mathbf{a})$ of $\psi(a)$ and each splitting $\vartheta_0(\mathbf{a}_0), \vartheta_1(a_1), \dots, \vartheta_\ell(a_\ell)$ of $\vartheta(\mathbf{a})$ such that there is a homomorphism h from $\vartheta_0(\mathbf{a}_0)$ to D given τ , there is an $i \in \{1, \dots, \ell\}$ such that $\neg \vartheta_i(x) \in \tau(a_i)$.

Our NP procedure for database satisfiability is, given a Σ -database D , to guess a type decoration τ for D and to then verify in deterministic polynomial time that D is proper. Note that the size of a type decoration is $\mathcal{O}(c \cdot |D|)$ for some constant c . The satisfiability checks in Point 1 of properness concern sentences whose size is independent of D , thus they need only constant time. Point 2 can be checked in time quadratic in the size of D . For Point 3, note that there are only polynomially many $(\text{dom}(D), 2)$ -subdivisions and splittings (in the size of D). To check the existence of the required homomorphism h , we can go through all candidates, directly verifying the homomorphism condition for relational atoms and proceedings as follows for NFA atoms: first extend D by exhaustively adding ‘implied facts’ of the form $\mathcal{A}(a, b)$, also taking into account assertions of the form $\mathcal{A}[q_0/s_i, F/\{s\}](x, x)$ that occur in τ -labels, as in Condition (c) above, and then treat NFA atoms like relational atoms. The following lemma finishes the proof of Theorem 16.

► **Lemma 17.** *D is satisfiable with φ_0 iff D has a proper type decoration τ such that $\varphi_0 \in \tau(a_0)$ for some $a_0 \in \text{dom}(D)$.*

6 Model Checking

We show that model checking in UNFO^{reg} is complete for $\text{P}^{\text{NP}[O(\log^2 n)]}$, the class of problems that can be solved in polynomial time given access to an NP oracle, but with only $O(\log^2 n)$ many oracle calls admitted. It thus has the same complexity as model checking in UNFO. Formally, the model checking problem for UNFO^{reg} is as follows: given a finite structure \mathfrak{A} and a UNFO^{reg} sentence φ , does $\mathfrak{A} \models \varphi$ hold? Without tests in path expressions, UNFO^{reg} model checking can easily be reduced to model checking in UNFO: simply extend the input structure by exhaustively adding ‘implied facts’ of the form $\mathcal{A}(a, b)$ and then replace every \mathcal{A} with a fresh binary relation symbol in both φ and \mathfrak{A} , obtaining an instance of UNFO model checking. With tests, this does not work. We would need multiple calls to UNFO model checking, essentially one call for every subformula inside a test in the input formula, but this brings us outside of $\text{P}^{\text{NP}[O(\log^2 n)]}$. We thus resort to expanding the $\text{P}^{\text{NP}[O(\log^2 n)]}$ upper bound proof from [47], which is by reduction to a $\text{P}^{\text{NP}[O(\log^2 n)]}$ -complete circuit value problem.

► **Theorem 18.** *The UNFO^{reg} model checking problem is $\text{P}^{\text{NP}[O(\log^2 n)]}$ -complete.*

7 Conclusion

We have proved that OMQ evaluation in $(\text{UNFO}^{\text{reg}}, \text{UC2RPQ})$ is decidable, 2EXPTIME -complete in combined complexity, and CONP -complete in data complexity, and that OMQ containment and satisfiability are also 2EXPTIME -complete. There are several interesting topics for future work. First, in contrast to UNFO, UNFO^{reg} does not have the finite model property and thus it would be interesting to study OMQ evaluation over finite models as well as finite satisfiability. Second, there are various natural directions for further increasing the expressive power. For example, one could allow any UNFO^{reg} formula with two free variables as a base case in regular path expressions instead of only atomic formulas. Such a logic would be strictly more expressive than propositional dynamic logic (PDL) with converse and intersection [31] and it would push the expressive power of UNFO^{reg} into the direction of regular queries, which have recently been proposed as an extension of C2RPQs [42]. Another natural extension was proposed by a reviewer of this paper: replace C2RPQs with linear Datalog to remove the asymmetry between binary relations and relations of higher arity in UNFO^{reg} . Additional relevant extensions could arise from the aim to capture additional description logics. From this perspective, it would for example be natural to extend UNFO^{reg} with constants, with fixed points, and with so-called role inclusions, please see [5]. Since functional relations and similar forms of counting play an important role in description logics, we remark that it is implicit in [47] that satisfiability (and thus OMQ evaluation) is undecidable in UNFO extended with two functional relations. Finally, it would be interesting to investigate the complexity of OMQ containment in $(\text{UNFO}^{\text{reg}}, \text{C2RPQ})$ without the restriction to a single ontology and to the full data signature. For (UNFO, CQ) , a 2NEXPTIME upper bound can be proved by a slight adaptation of the technique in [21], also using (a slightly refined version of) the translation from (UNFO, CQ) to monadic disjunctive Datalog from [19]. However, accommodating C2RPQs in this approach seems nontrivial.

Acknowledgements. ERC Consolidator Grant 647289 CODA and DFG grant SCHN 1234/3.

References

- 1 Antoine Amarilli, Michael Benedikt, Pierre Bourhis, and Michael Vanden Boom. Query answering with transitive and linear-ordered data. In *Proc. IJCAI*, pages 893–899, 2016.
- 2 Renzo Angles and Claudio Gutiérrez. Survey of graph database models. *ACM Comput. Surv.*, 40(1):1:1–1:39, 2008.
- 3 Franz Baader. Augmenting concept languages by transitive closure of roles: An alternative to terminological cycles. In *Proc. IJCAI*, pages 446–451, 1991.
- 4 Franz Baader, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, and Peter F. Patel-Schneider. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge Univ. Press, 2nd edition, 2007.
- 5 Franz Baader, Ian Horrocks, Carsten Lutz, and Ulrike Sattler. *An Introduction to Description Logic*. Cambridge University Press, 2017.
- 6 Jean-François Baget, Michel Leclère, Marie-Laure Mugnier, and Eric Salvat. On rules with existential variables: Walking the decidability line. *Artif. Intell.*, 175(9-10):1620–1654, 2011.
- 7 Jean-François Baget, Marie-Laure Mugnier, Sebastian Rudolph, and Michaël Thomazo. Walking the complexity lines for generalized guarded existential rules. In *Proc. IJCAI*, pages 712–717, 2011.
- 8 Vince Bárány, Georg Gottlob, and Martin Otto. Querying the guarded fragment. *Logical Methods in Computer Science*, 10(2), 2014.
- 9 Vince Bárány, Balder ten Cate, and Luc Segoufin. Guarded negation. *J. ACM*, 62(3):22:1–22:26, 2015.
- 10 Pablo Barceló. Querying graph databases. In *Proc. PODS*, pages 175–188, 2013.
- 11 Pablo Barceló, Gerald Berger, and Andreas Pieris. Containment for rule-based ontology-mediated queries. In *Proc. PODS*, 2018.
- 12 Meghyn Bienvenu, Diego Calvanese, Magdalena Ortiz, and Mantas Šimkus. Nested regular path queries in description logics. In *Proc. KR*, 2014.
- 13 Meghyn Bienvenu, Peter Hansen, Carsten Lutz, and Frank Wolter. First order-rewritability and containment of conjunctive queries in Horn description logics. In *Proc. IJCAI*, pages 965–971. IJCAI/AAAI Press, 2016.
- 14 Meghyn Bienvenu, Carsten Lutz, and Frank Wolter. Query containment in description logics reconsidered. In *Proc. KR*. AAAI Press, 2012.
- 15 Meghyn Bienvenu and Magdalena Ortiz. Ontology-mediated query answering with data-tractable description logics. In *Proc. Reasoning Web*, volume 9203 of *LNCS*, pages 218–307. Springer, 2015.
- 16 Meghyn Bienvenu, Magdalena Ortiz, and Mantas Šimkus. Conjunctive regular path queries in lightweight description logics. In *Proc. IJCAI*, pages 761–767. IJCAI/AAAI, 2013.
- 17 Meghyn Bienvenu, Magdalena Ortiz, and Mantas Šimkus. Navigational queries based on frontier-guarded datalog: Preliminary results. In *Proc. AMW*, volume 1378 of *CEUR Workshop Proceedings*, 2015.
- 18 Meghyn Bienvenu, Magdalena Ortiz, and Mantas Šimkus. Regular path queries in lightweight description logics: Complexity and algorithms. *J. Artif. Intell. Res.*, 53:315–374, 2015.
- 19 Meghyn Bienvenu, Balder ten Cate, Carsten Lutz, and Frank Wolter. Ontology-based data access: A study through disjunctive datalog, CSP, and MMSNP. *ACM Trans. Database Syst.*, 39(4):33:1–33:44, 2014.
- 20 Pierre Bourhis, Markus Krötzsch, and Sebastian Rudolph. How to best nest regular path queries. In *Proc. DL*, volume 1193 of *CEUR Workshop Proceedings*, pages 404–415, 2014.
- 21 Pierre Bourhis and Carsten Lutz. Containment in monadic disjunctive datalog, MMSNP, and expressive description logics. In *Proc. KR*, pages 207–216. AAAI Press, 2016.

- 22 Andrea Cali, Georg Gottlob, and Michael Kifer. Taming the infinite chase: Query answering under expressive relational constraints. In *Proc. KR*, pages 70–80, 2008.
- 23 Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. Data complexity of query answering in description logics. *Artif. Intell.*, 195:335–360, 2013.
- 24 Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, and Daniele Nardi. Reasoning in expressive description logics. In *Handbook of Automated Reasoning*, pages 1581–1634. Elsevier and MIT Press, 2001.
- 25 Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, and Moshe Y. Vardi. Containment of conjunctive regular path queries with inverse. In *Proc. KR*, pages 176–185, 2000.
- 26 Diego Calvanese, Thomas Eiter, and Magdalena Ortiz. Answering regular path queries in expressive description logics: An automata-theoretic approach. In *Proc. AAAI*, pages 391–396, 2007.
- 27 Diego Calvanese, Thomas Eiter, and Magdalena Ortiz. Regular path queries in expressive description logics with nominals. In *Proc. IJCAI*, pages 714–720, 2009.
- 28 Diego Calvanese, Thomas Eiter, and Magdalena Ortiz. Answering regular path queries in expressive description logics via alternating tree-automata. *Inf. Comput.*, 237:12–55, 2014.
- 29 Isabel F. Cruz, Alberto O. Mendelzon, and Peter T. Wood. A graphical query language supporting recursion. In *Proc. SIGMOD*, pages 323–330, 1987.
- 30 Michael J. Fischer and Richard E. Ladner. Propositional dynamic logic of regular programs. *J. Comput. Syst. Sci.*, 18(2):194–211, 1979.
- 31 Stefan Göller, Markus Lohrey, and Carsten Lutz. PDL with intersection and converse: satisfiability and infinite-state model checking. *J. Symb. Log.*, 74(1):279–314, 2009.
- 32 Georg Gottlob, Christoph Koch, and Reinhard Pichler. Efficient algorithms for processing XPath queries. In *Proc. VLDB*, pages 95–106, 2002.
- 33 Georg Gottlob, Giorgio Orsi, Andreas Pieris, and Mantas Šimkus. Datalog and its extensions for semantic web databases. In *Proc. Reasoning Web*, volume 7487 of *LNCS*, pages 54–77. Springer, 2012.
- 34 Georg Gottlob, Andreas Pieris, and Lidia Tendera. Querying the guarded fragment with transitivity. In *Proc. ICALP II*, volume 7966 of *LNCS*, pages 287–298. Springer, 2013.
- 35 Erich Grädel. On the restraining power of guards. *J. Symb. Log.*, 64(4):1719–1742, 1999.
- 36 Erich Grädel and Igor Walukiewicz. Guarded fixed point logic. In *Proc. LICS-99*, pages 45–54, 1999.
- 37 Roman Kontchakov and Michael Zakharyashev. An introduction to description logics and query rewriting. In *Proc. Reasoning Web*, volume 8714 of *LNCS*, pages 195–244. Springer, 2014.
- 38 Adila Krisnadhi and Carsten Lutz. Data complexity in the \mathcal{EL} family of description logics. In *Proc. LPAR*, volume 4790 of *LNCS*, pages 333–347. Springer, 2007.
- 39 Carsten Lutz. The complexity of conjunctive query answering in expressive description logics. In *Proc. IJCAR*, volume 5195 of *LNCS*, pages 179–193. Springer, 2008.
- 40 Magdalena Ortiz, Sebastian Rudolph, and Mantas Šimkus. Query answering in the Horn fragments of the description logics \mathcal{SHOIQ} and \mathcal{SROIQ} . In *Proc. IJCAI*, pages 1039–1044, 2011.
- 41 Antonella Poggi, Domenico Lembo, Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, and Riccardo Rosati. Linking data to ontologies. *J. Data Semantics*, 10:133–173, 2008.
- 42 Juan L. Reutter, Miguel Romero, and Moshe Y. Vardi. Regular queries on graph databases. *Theory Comput. Syst.*, 61(1):31–83, 2017.

- 43 Sebastian Rudolph and Markus Krötzsch. Flag & check: data access with monadically defined queries. In *Proc. PODS*, pages 151–162, 2013.
- 44 Klaus Schild. A correspondence theory for terminological logics: Preliminary report. In *Proc. IJCAI*, pages 466–471, 1991.
- 45 Philippe Schnoebelen. Oracle circuits for branching-time model checking. In *Proc. ICALP*, volume 2719 of *LNCS*, pages 790–801. Springer, 2003.
- 46 Wiesław Szwast and Lidia Tendera. The guarded fragment with transitive guards. *Ann. Pure Appl. Logic*, 128(1-3):227–276, 2004.
- 47 Balder ten Cate and Luc Segoufin. Unary negation. *Logical Methods in Computer Science*, 9(3), 2013.
- 48 Moshe Y. Vardi. Reasoning about the past with two-way automata. In *Proc. ICALP-98*, pages 628–641, 1998.

A

 Proofs for Section 2

► **Lemma 6.** *OMQ evaluation in $(\text{UNFO}_{\text{trans}}^{\text{reg}}, \text{UC2RPQ})$ reduces in polynomial time to satisfiability in UNFO^{reg} , and so does satisfiability in $\text{UNFO}_{\text{trans}}^{\text{reg}}$.*

Proof. We proceed in three steps: first, we reduce evaluation of *Boolean* OMQs in $(\text{UNFO}_{\text{trans}}^{\text{reg}}, \text{UC2RPQ})$ to satisfiability in $\text{UNFO}_{\text{trans}}^{\text{reg}}$; second, we reduce satisfiability in $\text{UNFO}_{\text{trans}}^{\text{reg}}$ to satisfiability in UNFO^{reg} ; third, we reduce evaluation of (general) OMQs to the Boolean case. For the sake of a convenient notation, we denote structures (and thus databases) as sets of facts.

(1) For the reduction from Boolean OMQ evaluation to satisfiability, let D be a database and $Q = (\mathcal{O}, \Sigma, q)$ an OMQ with \mathcal{O} a $\text{UNFO}_{\text{trans}}^{\text{reg}}$ formula and q a Boolean UC2RPQ. We show that

$$D \models Q \quad \text{iff} \quad \varphi_D \wedge \mathcal{O} \wedge \neg q \text{ is unsatisfiable,}$$

where φ_D is D viewed as a Boolean CQ in the obvious way.

The ‘if’ direction is immediate. To prove ‘only if’, assume $\mathfrak{A} \models \varphi_D() \wedge \mathcal{O} \wedge \neg q$ for some structure \mathfrak{A} . In particular, $\mathfrak{A} \models \varphi_D()$ is witnessed by a homomorphism h from D to \mathfrak{A} . If h is injective, then \mathfrak{A} is an extension of D and thus witnesses $D \not\models Q$ as desired. If h is not injective, we have to extend \mathfrak{A} iteratively, in each step taking an element a that has two distinct preimages b_1, b_2 under h , creating a fresh copy a' , duplicating all tuples in which a participates, and updating h such that it maps b_1 to a and b_2 to a' . After exhaustive application of this step, we obtain a structure \mathfrak{A}' that homomorphically embeds into \mathfrak{A} and is UN-bisimilar [47] to \mathfrak{A} , plus an injective homomorphism h' from D to \mathfrak{A}' . Together with the assumption, this implies that $\mathfrak{A}' \models \mathcal{O} \wedge \neg q$ and \mathfrak{A}' extends D ; hence $D \not\models Q$ as desired.

(2) Let φ be $\text{UNFO}_{\text{trans}}^{\text{reg}}$ sentence with transitivity atoms $\text{trans}(R_1), \dots, \text{trans}(R_n)$. Transform φ into a UNFO^{reg} sentence φ' by dropping the transitivity atoms and replacing each atom $R_i(x, y)$ with the (regular expression) atom $R_i^+(x, y)$. It is straightforward to show that φ is satisfiable if and only if φ' is.

(3) For the reduction from (general) OMQ evaluation to the Boolean case, let $Q = (\mathcal{O}, \Sigma, q(\mathbf{x}))$ be an OMQ in $(\text{UNFO}_{\text{trans}}^{\text{reg}}, \text{UC2RPQ})$ with $q(\mathbf{x}) = q_1(\mathbf{x}) \vee \dots \vee q_n(\mathbf{x})$ such that $q_i(\mathbf{x}) = \exists \mathbf{y}_i \varphi_i(\mathbf{x}, \mathbf{y}_i)$ and $\mathbf{x} = x_1, \dots, x_n$, D a database, and $\mathbf{a} \subseteq \text{dom}(D)$ with $\mathbf{a} = a_1, \dots, a_n$. We construct a new OMQ Q' and database D' by taking fresh unary predicates P_1, \dots, P_n and setting:

$$\begin{aligned} D' &= D \cup \{P_1(a_1), \dots, P_n(a_n)\} \\ Q' &= (\mathcal{O}, \Sigma, q'()), \quad \text{where} \\ q'() &= q'_1() \vee \dots \vee q'_n(), \quad \text{with} \\ q'_i() &= \exists \mathbf{x} \mathbf{y}_i (\varphi_i(\mathbf{x}, \mathbf{y}_i) \wedge P_1(x_1) \wedge \dots \wedge P_n(x_n)) \end{aligned}$$

It suffices to prove the following claim, which is nearly straightforward.

Claim. $\mathbf{a} \in \text{cert}(Q, D) \Leftrightarrow () \in \text{cert}(Q', D')$

Proof of Claim. We proceed via contraposition in both directions.

‘ \Rightarrow ’ Assume $() \notin \text{cert}(Q', D')$. Then there is a structure \mathfrak{A} that extends D' and is a model of \mathcal{O} with $() \notin \text{ans}(q', \mathfrak{A})$. Since \mathfrak{A} extends D' it also holds that $\mathbf{a} \notin \text{ans}(q, \mathfrak{A})$ (because otherwise the witnessing homomorphism would also witness $() \in \text{ans}(q', \mathfrak{A})$, a contradiction). Hence $\mathbf{a} \notin \text{cert}(Q, D)$.

‘ \Leftarrow ’ Suppose $\mathbf{a} \notin \text{cert}(Q, D)$. Then there is a structure \mathfrak{A} that extends D and is a model of \mathcal{O} with $\mathbf{a} \notin \text{ans}(q, \mathfrak{A})$. Consider the structure $\mathfrak{A}' = \mathfrak{A} \cup \{P_1(a_1), \dots, P_n(a_n)\}$, which extends \mathfrak{A} and is a model of \mathcal{O} (since the P_i were fresh). In addition, it holds that $() \notin \text{ans}(q', \mathfrak{A}')$ (because otherwise the witnessing homomorphism would also witness $\mathbf{a} \notin \text{ans}(q, \mathfrak{A})$, a contradiction). Hence $() \notin \text{cert}(Q', D')$. \blacktriangleleft

► **Lemma 7.** *OMQ containment in $(\text{UNFO}^{\text{reg}}, \text{UC2RPQ})$ reduces in polynomial time to satisfiability in UNFO^{reg} .*

Proof. It suffices to reduce containment between *Boolean* OMQs to satisfiability; the case of general OMQs can be reduced to the Boolean case by applying the construction described in the proof of Lemma 6, Step 3, to both input OMQs.

Let $Q_1 = (\mathcal{O}, \Sigma_{\text{full}}, q_1)$ and $Q_2 = (\mathcal{O}, \Sigma_{\text{full}}, q_2)$ be OMQs with \mathcal{O} a UNFO^{reg} sentence and q_1, q_2 Boolean UC2RPQs. Then the following holds:

$$Q_1 \subseteq Q_2 \quad \text{iff} \quad \mathcal{O} \wedge q_1 \wedge \neg q_2 \text{ is unsatisfiable.}$$

For the ‘if’ direction, assume $Q_1 \not\subseteq Q_2$. Then there is a database D with $\text{cert}(Q_1, D) \not\subseteq \text{cert}(Q_2, D)$, i.e., there is a structure \mathfrak{A} extending D such that $\mathfrak{A} \models \mathcal{O}$, $\mathfrak{A} \models q_1$, and $\mathfrak{A} \not\models q_2$. Hence $\mathcal{O} \wedge q_1 \wedge \neg q_2$ is satisfiable.

For the ‘only if’ direction, assume that $\mathcal{O} \wedge q_1 \wedge \neg q_2$ is satisfiable. Then $\mathfrak{A} \models \mathcal{O}$, $\mathfrak{A} \models q_1$, and $\mathfrak{A} \not\models q_2$ for some structure \mathfrak{A} . Then \mathfrak{A} contains some (finite) database D witnessing $Q_1 \not\subseteq Q_2$. \blacktriangleleft

► **Lemma 8.** *Every UNFO^{reg} sentence φ can be transformed into an equivalent normal UNFO^{reg} sentence φ' in single exponential time. Moreover, the width and the atom width of φ' are at most polynomial in the size of φ and the path expressions that occur in φ' are exactly those in φ .*

Proof. By Lemma 4.1 of [47], we can convert any UNFO sentence φ in single exponential time into an equivalent UNFO sentence φ' generated by the following grammar:

$$\varphi(x) ::= \exists \mathbf{y} \psi(x, \mathbf{y}) \mid \neg \varphi(x) \mid \varphi(x) \vee \varphi(x)$$

where $\exists \mathbf{y} \psi(x, \mathbf{y})$ is a CQ that might contain equality atoms. The transformation steps are rather straightforward and also work for UNFO^{reg} with the only difference that $\exists \mathbf{y} \psi(x, \mathbf{y})$ is then a C2RPQ that might contain equality atoms. The transformation may cause an at most single exponential blowup in formula size and it satisfies the requirements regarding parameters formulated in Lemma 8.

We can easily eliminate equality atoms in C2RPQs by identifying variables; when a free variable is identified with a quantified variable, we use the name of the free variable.

It remains to make C2RPQs connected. Let $\exists \mathbf{y} \psi(x, \mathbf{y})$ be a C2RPQ subformula such that $\psi(x, \mathbf{y})$ has the connected components $\psi(x, \mathbf{y}_0), \psi(\mathbf{y}_1), \dots, \psi(\mathbf{y}_k)$, $k \geq 1$. We replace it with the conjunction of $\varphi_0 = \exists \mathbf{y}_0 \psi(x, \mathbf{y}_0)$ and $\varphi_1 = \exists \mathbf{y}_1 \psi(\mathbf{y}_1), \dots, \varphi_k = \exists \mathbf{y}_k \psi(\mathbf{y}_k)$, that is, with the C2RPQ $\varphi_0?(x), \varphi_1?(x), \dots, \varphi_k?(k)$. Note that this C2RPQ and all C2RPQs inside the tests are connected. In fact, in can be verified that the resulting UNFO^{reg} sentence is normal according to our definition. \blacktriangleleft

B Proofs for Section 3

► **Lemma 9.** *A normal UNFO^{reg} sentence φ_0 of size n and width m is satisfiable iff there is a proper type decorated tree-like structure (T, bag, τ) for φ_0 of width at most m and outdegree at most $n^2 + n$ such that $\varphi_0 \in \tau(a)$ for some a .*

Proof. (\Leftarrow) Assume a proper type decorated tree-like structure (T, bag, τ) for φ_0 . It is easily verified by induction on the structure of formulas that, for all $\psi(x) \in \text{cl}(\varphi_0)$ and all $a \in \text{dom}(T, \text{bag})$, we have $\mathfrak{A}_{(T, \text{bag})} \models \psi(a)$ if, and only if, $\psi(x) \in \tau(a)$. Since $\varphi_0 \in \tau(a)$ for some a , we get $\mathfrak{A}_{(T, \text{bag})} \models \varphi_0$.

(\Rightarrow) Let $\mathfrak{A} \models \varphi_0$ for a UNFO^{reg} sentence φ_0 of size n . In order to construct a tree-like structure (T, bag) , we use an unraveling technique during which we maintain a mapping $h : \text{dom}(T, \text{bag}) \rightarrow A$ and an additional labeling $E(w)$ containing expressions of the form $\mathcal{A}(b, b')$, for each $w \in T$. Throughout the construction, we preserve as an invariant that h is a homomorphism and that for each $\mathcal{A}(b, b') \in E(w)$, we have $b, b' \in \text{dom}(w)$ and $\mathfrak{A} \models \mathcal{A}(h(b), h(b'))$.

Start with choosing elements $a_1, \dots, a_k \in A$ such that, for every formula of the form $\exists x \psi(x) \in \text{cl}(\varphi_0)$ with $\mathfrak{A} \models \exists x \psi(x)$, there is some $i \in \{1, \dots, k\}$ and $\mathfrak{A} \models \psi(a_i)$. Initialize (T, bag) by setting $T = \{\varepsilon, 1, \dots, k\}$, $\text{bag}(\varepsilon)$ to the empty structure, and $\text{bag}(i) = \mathfrak{A}|_{\{a_i\}}$, for every $i \in \{1, \dots, k\}$, where $\mathfrak{A}|_X$ denotes the restriction of \mathfrak{A} to domain X . Moreover, set $h(a_i) = a_i$ and $E(\varepsilon) = E(i) = \emptyset$, for all $i \in \{1, \dots, k\}$. Clearly, the invariants are satisfied.

Then, apply the following steps exhaustively and in a fair way:

- Choose a node $w \in T$, an element $a \in \text{dom}(w)$, and a C2RPQ $\psi(x) \in \text{cl}(\varphi_0)$ with $\mathfrak{A} \models \psi(h(a))$. There is a mapping $\beta : \text{var}(\psi) \rightarrow A$ such that $\beta(x) = h(a)$, and for each atom $R(\mathbf{z}) \in \psi$, we have $\beta(\mathbf{z}) \in R^{\mathfrak{A}}$, and for each $\mathcal{A}(z, z') \in \psi$, \mathcal{A} an NFA, we have $\mathfrak{A} \models \mathcal{A}(\beta(z), \beta(z'))$. Let $B = \{\beta(z) \mid z \in \text{var}(\psi)\}$, and create a successor v of w where the associated structure $\text{bag}(v)$ is obtained from $\mathfrak{A}|_B$ by replacing each $b \in B \setminus \{h(a)\}$ with a fresh b' , and $h(a)$ with a . Extend h by setting $h(b') = b$, for all introduced b' . Finally, set $E(w) = \{\mathcal{A}(\beta(z), \beta(z')) \mid \mathcal{A}(z, z') \in \psi\}$ where \bar{b} is b' for all $b \in B \setminus \{h(a)\}$ and $\bar{h}(a)$ is a .
- Choose a node $w \in T$ and a label $\mathcal{A}(b, b') \in E(w)$. By the invariant for $E(w)$, we know that $\mathfrak{A} \models \mathcal{A}(h(b), h(b'))$. Take a shortest sequence $a_1, \dots, a_n \in A$ and word $\hat{w} \in L(\mathcal{A})$ of length $n - 1$ such that $a_1 = h(b)$, $a_n = h(b')$, and $(a_i, a_{i+1}) \in R^{\mathfrak{A}}$ if the i th symbol in \hat{w} is R , $(a_{i+1}, a_i) \in R^{\mathfrak{A}}$ if the i th symbol in \hat{w} is R^- , and $a_i = a_{i+1}$ and $\mathfrak{A} \models \varphi(a_i)$ if the i th symbol in \hat{w} is $\varphi(x)$?. If $n \leq 2$, then, by construction, $\text{bag}(w) \models \mathcal{A}(b, b')$, so we can stop. If $n > 2$, let $\mathcal{A} = (Q, \Sigma, q_0, \Delta, F)$ and q_0, \dots, q_n be a sequence of states such that $(q_i, a_{i+1}, q_{i+1}) \in \Delta$ for all i with $0 \leq i < n$ and $q_n \in F$. Define $B = \{a_1, a_2, a_n, a_{n-1}\}$ and $B' = B \setminus \{a_1, a_n\}$. Then create a successor v of w , and set $\text{bag}(v)$ to the structure obtained from $\mathfrak{A}|_B$ by replacing every $d \in B'$ with a fresh d' , a_1 with b , and a_n with b' . Finally, extend h by setting $h(d') = d$ for all $d \in B'$, and set $E(v)$ to the singleton set containing $\mathcal{A}[q_0/q_1, F/\{q_{n-1}\}](\bar{a}_2, \bar{a}_{n-1})$, where \bar{d} is d' , for all $d \in B'$, and $\bar{a}_1 = b$ and $\bar{a}_n = b'$.

By definition of the rules, the invariants regarding h and E are preserved.

Let (T^*, bag^*) and h^* be the tree-like structure and homomorphism, respectively, obtained in the limit of the unraveling. We define a type decoration τ by taking $\tau(a) = \{\psi(x) \in \text{cl}(\varphi_0) \mid \mathfrak{A} \models \psi(h^*(a))\}$, for all $a \in \text{dom}(T^*, \text{bag}^*)$. It is not difficult to verify that $\mathfrak{A}_{(T^*, \text{bag}^*)} \models \psi(a)$ iff $\psi(x) \in \tau(a)$, for all $a \in \text{dom}(T^*, \text{bag}^*)$ and all $\psi(x) \in \text{cl}(\varphi_0)$. In particular, we have:

- if $\psi(x) \in \tau(a)$ is a C2RPQ, then $\mathfrak{A} \models \psi(h^*(a))$. Since steps 1 and 2 were applied exhaustively, we know that $\mathfrak{A}_{(T^*, \text{bag}^*)} \models \psi(a)$.
- Let $\neg\psi(x) \in \tau(a)$ for a C2RPQ $\psi(x)$ and assume $\mathfrak{A}_{(T^*, \text{bag}^*)} \models \psi(a)$. Since h^* is a homomorphism from $\mathfrak{A}_{(T^*, \text{bag}^*)}$ to \mathfrak{A} , we also have $\mathfrak{A} \models \psi(h^*(a))$, a contradiction to the definition of τ .

Consequently, $(T^*, \text{bag}^*, \tau)$ is proper. Finally, observe that the size of the bag created is bounded by m in the first step and by 4 in the second step. For the outdegree, observe that a bag created in the second step has outdegree 1, while a bag created in the first step has

outdegree at most $n^2 + n$. \blacktriangleleft

► **Lemma 11.** *The cardinality of $\text{qcl}(\varphi_0, \Delta, s)$ is bounded by $p \cdot (a^2 d^m)^{m'}$, where p is the number of C2RPQs in φ_0 , a the maximal number of states in an NFA in φ_0 , d the cardinality of Δ , m the width of φ_0 , and m' the atom width of φ_0 .*

Proof. Each query in $\text{qcl}(\varphi_0, \Delta, s)$ can be obtained from a C2RPQ $\varphi(x) \in \text{cl}(\varphi_0)$ by first dropping atoms, then replacing NFA atoms $\mathcal{A}(t, t')$ with an atom $\mathcal{A}'(\hat{t}, \hat{t}')$ where \mathcal{A}' is obtained from \mathcal{A} by replacing the initial state and replacing the set of final states with a single state, \hat{t} is t or a fresh constant, and \hat{t}' is t' or a fresh constant, and finally replacing existentially quantified variables with constants from Δ ; to see this, it is important to consider Condition 3 of splittings and to note that the ψ_0 -component of splittings is not included in $\text{qcl}(\varphi_0, \Delta, s)$. Thus the number of atoms in each query in $\text{qcl}(\varphi_0, \Delta, s)$ is at most m' . Furthermore, each atom can take on $a^2 d^m$ variations by replacing NFA states as described and/or replacing variables with constants. Thus each of the p C2RPQs $\varphi(x) \in \text{cl}(\varphi_0)$ contributes at most $(a^2 d^m)^{m'}$ C2RPQs to $\text{qcl}(\varphi_0, \Delta, s)$. \blacktriangleleft

► **Lemma 12.** *Let (T, bag, τ) be a type decorated tree-like structure, $\varphi(x) \in \text{cl}(\varphi_0)$ a C2RPQ, and $a \in \text{dom}(T, \text{bag})$. Then there is a homomorphism h from $\varphi(x)$ to (T, bag, τ) with $h(x) = a$ iff there is a witness tree for $\varphi(a)$ in (T, bag, τ) .*

Proof. (\Rightarrow) Let h be a homomorphism from $\varphi(x)$ to (T, bag, τ) with $h(x) = a$. We inductively construct a witness tree (W, σ) for $\varphi(a)$ in (T, bag, τ) . During the construction, we maintain the following invariants for all nodes $u \in W$ with $\sigma(u) = (w, \psi(\mathbf{a}))$:

- (i) for all $R(\mathbf{t}) \in \psi(\mathbf{a})$, we have $h(\mathbf{t}) \in R^{\mathfrak{A}_{(T, \text{bag})}}$;
- (ii) for all $\mathcal{A}(t, t') \in \psi(\mathbf{a})$, we have $\mathfrak{A}_{(T, \text{bag})}, \tau \models \mathcal{A}(h(t), h(t'))$.

We start the construction with setting $\sigma(\varepsilon) = (w, \varphi(a))$ for some $w \in T$ with $a \in \text{dom}(w)$, obviously satisfying (i) and (ii).

Then, apply the following step exhaustively. Let $u \in W$ be an unprocessed node in the witness tree constructed so far, and assume that $\sigma(u) = (w, \psi(\mathbf{a}))$. First, assign to each constant $a \in \text{dom}(T, \text{bag}) \setminus \text{dom}(w)$ the (uniquely defined) value $\kappa(a) \in \{-1, 1, \dots, m\}$, m the outdegree of T , such that $w \cdot \kappa(a)$ lies on the shortest path from w to the unique world where a appears for the first time in (T, bag) .

We use the mapping κ to assign atoms occurring in (a subdivision of) $\psi(\mathbf{a})$ to neighboring nodes of w , intuitively, to reflect where h maps different parts of $\psi(\mathbf{a})$. Formally, we use queries $\psi_{-1}(\mathbf{b}_{-1}), \dots, \psi_\ell(\mathbf{b}_\ell)$, initialized with \emptyset , where $\psi_i(\mathbf{b}_i)$ collects the parts of $\psi(\mathbf{a})$ which are sent to $w \cdot i$, for all i . We process all atoms in $\psi(\mathbf{a})$ as follows:

1. For each atom $R(\mathbf{t}) \in \psi(\mathbf{a})$, by invariant (i), we can fix a $v \in T$ with $h(\mathbf{t}) \in R^{\text{bag}(v)}$ which has minimal distance to w . We distinguish three cases:

- a. if $h(\mathbf{t}) \in R^{\text{bag}(w)}$, then add $R(\mathbf{t})$ to ψ_0 ;
- b. if $h(\mathbf{t}) \subseteq \text{dom}(w)$ and $h(\mathbf{t}) \notin R^{\text{bag}(w)}$, then add $R(\mathbf{t})$ to ψ_i where i is the (unique) number such that $w \cdot i$ is on the shortest path to v ;
- c. if $h(\mathbf{t}) \not\subseteq \text{dom}(w)$, add $R(\mathbf{t})$ to ψ_i where $i = \kappa(h(x))$ for some $x \in \mathbf{t}$ with $h(x) \notin \text{dom}(w)$. It is important to note that i is uniquely defined in this way. Indeed, assume two variables $x, y \in \mathbf{t}$ with $h(x), h(y) \notin \text{dom}(w)$ and $\kappa(h(x)) \neq \kappa(h(y))$. Then, one of $w \cdot \kappa(h(x))$ and $w \cdot \kappa(h(y))$ does not lie on the shortest path from w to v , say the latter. However, by the connectedness property of tree decompositions, we know that then $h(y)$ appears in $\text{dom}(w)$, a contradiction.

2. For an atom $\mathcal{A}(t, t') \in \psi(\mathbf{a})$, by invariant (ii), we can fix sequences a_1, \dots, a_n and s_0, \dots, s_n , and a word $\nu_1 \cdots \nu_{n-1} \in L(\mathcal{A})$ of minimal length such that $a_1 = h(t)$, $a_n = h(t')$, $s_0 = q_0$, $s_n \in F$, $(s_i, \nu_i, s_{i+1}) \in \Delta$, for all $i \in \{1, \dots, n-1\}$, $(a_i, a_{i+1}) \in R^{\mathfrak{A}(T, \text{bag}, \tau)}$ if $\nu_i = R$, $(a_{i+1}, a_i) \in R^{\mathfrak{A}(T, \text{bag}, \tau)}$ if $\nu_i = R^-$, and $\theta(x) \in \tau(a_i)$ and $a_{i+1} = a_i$ if $\nu_i = \theta(x)$?. Let I be the set of all $i \in \{1, \dots, n\}$ such that $a_i \in \text{dom}(w)$.

If $I = \emptyset$, then t, t' are variables with $\kappa(h(t)) = \kappa(h(t'))$. Add $\mathcal{A}(t, t')$ to $\psi_{\kappa(h(t))}$.

Otherwise, that is, $I \neq \emptyset$, let $i_1 < \dots < i_k$ be a linear order of the elements in I . If $a_1 \notin \text{dom}(w)$ and thus $1 \notin I$, then add $\mathcal{A}[F/\{s_{i_1+1}\}](t, a_{i_1})$ to $\psi_{\kappa(a_1)}$, and, if $a_n \notin \text{dom}(w)$ and thus $n \notin I$, then add $\mathcal{A}[q_0/s_{i_k}](a_{i_k}, t')$ to $\psi_{\kappa(a_n)}$. Moreover, for each $j \in \{1, \dots, k\}$ such that ν_{i_j} is not a test $\theta(x)$? do:

- if $\nu_{i_j} = R$ and $(a_{i_j}, a_{i_{j+1}}) \in R^{\text{bag}(w)}$, then add $\mathcal{A}[q_0/s_{i_j}, F/\{s_{i_{j+1}}\}](a_{i_j}, a_{i_{j+1}})$ to ψ_0 ;
- if $\nu_{i_j} = R$ and $(a_{i_j}, a_{i_{j+1}}) \notin R^{\text{bag}(w)}$, then add $\mathcal{A}[q_0/s_{i_j}, F/\{s_{i_{j+1}}\}](a_{i_j}, a_{i_{j+1}})$ to ψ_i , where $i = \kappa(a_{i_{j+1}})$;
- if $\nu_{i_j} = R^-$ and $(a_{i_{j+1}}, a_{i_j}) \in R^{\text{bag}(w)}$, then add $\mathcal{A}[q_0/s_{i_j}, F/\{s_{i_{j+1}}\}](a_{i_j}, a_{i_{j+1}})$ to ψ_0 ;
- if $\nu_{i_j} = R^-$ and $(a_{i_{j+1}}, a_{i_j}) \notin R^{\text{bag}(w)}$, then add $\mathcal{A}[q_0/s_{i_j}, F/\{s_{i_{j+1}}\}](a_{i_j}, a_{i_{j+1}})$ to ψ_i , where $i = \kappa(a_{i_{j+1}})$.

It is crucial that in Step 2 the cardinality of I is bounded by m^2 , m the size of φ_0 . More precisely, in sequences a_1, \dots, a_n and s_0, \dots, s_n of minimal length, there are no $i < j$ such that $a_i = a_j$ and $s_i = s_j$, as otherwise we can obtain shorter sequences by dropping a_{i+1}, \dots, a_j and s_{i+1}, \dots, s_j . As both $\text{dom}(w)$ and the number of states in \mathcal{A} is bounded by m , the claimed bound follows. It should thus be clear that $\vartheta(\mathbf{a}) = \bigcup_i \psi_i(\mathbf{b}_i)$ is a subdivision of $\psi(\mathbf{a})$.

Note that the $\psi_i(\mathbf{b}_i)$ need not be connected. Define a splitting $\vartheta_0(\mathbf{a}_0), \dots, \vartheta_\ell(\mathbf{a}_\ell)$ of $\vartheta(\mathbf{a})$ by setting $\vartheta_0(\mathbf{a}_0) = \psi_0(\mathbf{b}_0)$, and including, for each $i \in \{-1, 1, \dots, m\}$, each connected component $\psi'(\mathbf{a}')$ of $\psi_i(\mathbf{b}_i)$ in the sequence. By construction, h is a homomorphism from $\vartheta_0(\mathbf{a}_0)$ to $\text{bag}(w)$ given τ . Finally, extend the witness tree by adding, for each $\vartheta_i(\mathbf{a}_i)$, $1 \leq i \leq \ell$, a successor u_i of u with $\sigma(u_i) = (w \cdot j, \vartheta'_i(\mathbf{a}'_i))$ where j is such that $\vartheta_i(\mathbf{a}_i) \subseteq \psi_j(\mathbf{b}_j)$ and $\vartheta'_i(\mathbf{a}'_i)$ is obtained from $\vartheta_i(\mathbf{a}_i)$ by replacing each variable x in the domain of h with $h(x)$. By construction, u satisfies (*). Moreover, it is routine to verify that invariants (i) and (ii) are preserved.

It remains to argue that the described process results in a finite tree. Clearly, the constructed tree is finitely branching. For finite depth, consider first atoms of the form $R(\mathbf{t}) \in \varphi(a)$. In Step 1, these atoms are always ‘sent’ to a closest v such that $h(\mathbf{t}) \in R^{\text{bag}(v)}$ (Items 1b and 1c). This v is reached after finitely many steps. Consider now atoms of the form $\mathcal{A}(t, t')$. Assume some atom $\mathcal{A}(t_0, t'_0)$ is obtained in Step 2 applied to $\mathcal{A}(t, t')$ and that $\mathcal{A}(t_0, t_1) \in \vartheta_i(\mathbf{a}_i)$. Let $\mathcal{A}(t'_0, t'_1)$ be the corresponding atom in $\vartheta'_i(\mathbf{a}'_i)$. Then, by the minimality condition, the witnessing sequences selected in Step 2 when applied to $\mathcal{A}(t'_0, t'_1)$ while processing $\vartheta'_i(\mathbf{a}'_i)$ is strictly shorter than the witnessing sequence for $\mathcal{A}(t, t')$. Thus, finite depth follows.

(\Leftarrow) Let (W, σ) be a witness tree for $\varphi(a)$ in (T, bag, τ) . We inductively construct a homomorphism h from $\varphi(x)$ to (T, bag, τ) such that $h(x) = a$.

Start with $h(x) = a$. Then apply the following rule exhaustively. Let $u \in W$ be a node in the witness tree with $\sigma(u) = (w, \psi(\mathbf{a}))$ such that all predecessor nodes have been processed. Let g be the homomorphism witnessing Condition (*) for u , and define $h(z) = g(z)$ for all z in the domain of g but not in the domain of h .

It is a consequence of Condition 2 of splittings and (*), that h is well-defined. We show that the result h of this process is a homomorphism from $\varphi(x)$ to (T, bag, τ) .

- Consider first atoms of the form $R(\mathbf{t})$. We prove by induction that, if $R(\mathbf{t}) \in \psi(\mathbf{a})$ for

some $u \in W$ with $\sigma(u) = (w, \psi(\mathbf{a}))$, then $h(\mathbf{t}) \in R^{\mathfrak{A}(T, \text{bag})}$. The induction base is the case when $R(\mathbf{t})$ is put into $\vartheta_0(\mathbf{a})$ by (*). By definition of h , we know $h(\mathbf{t}) \in R^{\text{bag}(w)}$, thus also $h(\mathbf{t}) \in R^{\mathfrak{A}(T, \text{bag})}$. In the induction step, let $R(\mathbf{t}) \in \psi(\mathbf{a})$, but in (*) the atom $R(\mathbf{t})$ is put into $\vartheta_i(\mathbf{a}_i)$ for some $i > 0$. By the definition of, $R(\mathbf{t}') \in \vartheta'_i(\mathbf{a}')$ where \mathbf{t}' is obtained from \mathbf{t} by instantiating some variables $z \in \mathbf{t}$ with $h(z)$. By induction, we know that $h(\mathbf{t}') \in R^{\mathfrak{A}(T, \text{bag})}$, thus also $h(\mathbf{t}) \in R^{\mathfrak{A}(T, \text{bag})}$.

- Consider now atoms of the form $\mathcal{A}(t, t')$. We prove by induction that, if $\mathcal{A}(t, t') \in \psi(\mathbf{a})$ for some $u \in W$ with $\sigma(u) = (w, \psi(\mathbf{a}))$, then $\mathfrak{A}_{(T, \text{bag})}, \tau \models \mathcal{A}(h(t), h(t'))$. The induction base is the case when $\mathcal{A}(t, t')$ is put into $\vartheta_0(\mathbf{a}_0)$ by (*). By (*) and the definition of h , we know that h is a homomorphism from $\mathcal{A}(t, t')$ to $\text{bag}(w)$, hence $\text{bag}(w), \tau \models \mathcal{A}(h(t), h(t'))$ and thus $\mathfrak{A}_{(T, \text{bag})}, \tau \models \mathcal{A}(h(t), h(t'))$. In the induction step, the atom $\mathcal{A}(t, t')$ is subdivided into atoms, say $\alpha_1, \dots, \alpha_k$. However, by definition of subdivisions and by (*), it is straightforward to prove that $\mathfrak{A}_{(T, \text{bag})}, \tau \models \mathcal{A}(h(t), h(t'))$ given that $\mathfrak{A}_{(T, \text{bag})}, \tau \models \alpha_i$ for all i , by induction hypothesis. ◀

C Proofs for Section 4

We give the semantics of 2ATAs. A *run* of \mathcal{A} on a labeled tree (T, L) is a $(T_r \times Q)$ -labeled tree (T_r, r) such that $r(\varepsilon) = (\varepsilon, q_0)$ and whenever $x \in T_r, r(x) = (w, q)$, and $\delta(q, L(w)) = \theta$, then there is a set $\mathcal{V} = \{(m_1, q_1), \dots, (m_n, q_n)\} \subseteq [k] \times Q$ such that \mathcal{V} satisfies θ and for $1 \leq i \leq n$, we have $x \cdot i \in T_r, w \cdot m_i$ is defined, and $r(x \cdot i) = (w \cdot m_i, q_i)$. A run is *accepting* if every infinite path π satisfies the parity condition F , that is, if there is an even i such that $\text{inf}(\pi) \cap G_i \neq \emptyset$ and $\text{inf}(\pi) \cap G_{i-1} = \emptyset$, where $\text{inf}(\pi) \subseteq Q$ denotes the set of states that occur infinitely often in π . The automaton accepts an input tree if there is an accepting run for it.

► **Theorem 14.** *In UNFO^{reg} , satisfiability is 2EXPTIME-complete.*

Proof. The lower bound follows from that for UNFO [47]. For the upper bound, let φ be a UNFO^{reg} sentence of size n . By Lemma 8, we can transform φ into an equivalent normal UNFO^{reg} sentence φ_0 whose width m and atom width m' are polynomial in n and whose size is single exponential in n . Note that the number of 1-types for φ_0 is double exponential in n and that, by the bounds stated in Lemma 8 and by Lemma 11, the cardinality of $\text{qcl}(\varphi_0, \Delta, n^2)$ is single exponential in n .

We then build \mathcal{A}_0 and \mathcal{A}_1 for φ_0 as described above. The number of states of \mathcal{A}_0 is exponential in n and, by the bound on $\text{qcl}(\varphi_0, \Delta, n^2)$ stated above, the same is true for the number of states of \mathcal{A}_1 . The alphabet Σ is of cardinality double exponential in n . The transition functions of \mathcal{A}_0 and \mathcal{A}_1 can be computed in time double exponential in n . Constructing the intersection 2ATA does not increase the number of states. In summary, the final 2ATA \mathcal{A} can be constructed in time double exponential in n and has single exponentially many states in n . The number of sets in the parity condition is a constant. Consequently, nonemptiness of \mathcal{A} can be decided in time double exponential in φ . ◀

D Proofs for Section 5

► **Lemma 17.** *D is satisfiable with φ_0 iff D has a proper type decoration τ such that $\varphi_0 \in \tau(a_0)$ for some $a_0 \in \text{dom}(D)$.*

Proof. The ‘only if’ direction is rather straightforward. Let \mathfrak{A} be a model of D and φ_0 .

Then we can define, for every $a \in \text{dom}(D)$,

$$\tau(a) = \{\varphi(x) \in \text{ecl}(\varphi_0) \mid \mathfrak{A} \models \varphi(a)\}.$$

We aim to show that τ is a proper type decoration of D and that $\varphi_0 \in \tau(a_0)$, for some $a_0 \in D$. The latter is clear as, by assumption, $\mathfrak{A} \models \varphi_0$. We verify Point 1 to 3 of properness. Points 1 and 2 are clear as $\tau(a)$ is read off from a model of $\tau(a)$. Assume to the contrary of what we aim to show that Point 3 is violated. Then there is an $a \in \text{dom}(D)$, a $\neg\psi(x) \in \tau(a)$ for some $\psi(x) \in \text{qcl}(\varphi_0)$, a subdivision $\vartheta(\mathbf{a})$ of $\psi(a)$, a splitting $\vartheta_0(\mathbf{a}_0), \vartheta_1(a_1), \dots, \vartheta_k(k_k)$ of $\vartheta(\mathbf{a})$ and a homomorphism h from $\vartheta_0(\mathbf{a}_0)$ to D given τ , such that $\neg\vartheta_i(x) \notin \tau(a_i)$, for all $1 \leq i \leq k$. By definition of τ , we get that $\mathfrak{A} \models \vartheta_i(a_i)$, for all i . It can be verified that this implies $\mathfrak{A} \models \psi(a)$, in contradiction to $\neg\psi(x) \in \tau(a)$. Thus, τ is proper.

Now for the ‘if’ direction. Assume that D has a proper type decoration τ . Then we find, for each $a \in \text{dom}(D)$, a model \mathfrak{A}_a of the formula in Point 1 of the definition of properness. We assume w.l.o.g. that the domains of \mathfrak{A}_a and \mathfrak{A}_b are disjoint when $a \neq b$ and that each \mathfrak{A}_a shares with D only the constant a . Let \mathfrak{A} be obtained by taking the union of D and all models \mathfrak{A}_a . Clearly, \mathfrak{A} is a model of D . It thus remains to show that it is also a model of φ_0 . We start with an auxiliary claim.

Claim. For all $a \in \text{dom}(D)$, $b \in \text{dom}(\mathfrak{A}_a)$, and for all $\varphi(x) \in \text{ecl}(\varphi_0)$ with a free variable x , we have $\mathfrak{A}_a \models \varphi(b)$ iff $\mathfrak{A} \models \varphi(b)$.

Proof of the Claim. The proof is by induction on the structure of $\varphi(x)$. Since φ_0 is normal, there are three cases: negation $\neg\varphi(x)$, unary disjunction $\varphi(x) \vee \psi(x)$, and C2RPQs $\varphi(x)$. Negation and unary disjunction are immediate; we consider only C2RPQs.

Let $\varphi(x) \in \text{ecl}(\varphi_0)$ be a C2RPQ. For (\Rightarrow) , assume that $\mathfrak{A}_a \models \varphi(b)$, that is, there is a mapping $h : \text{var}(\varphi) \rightarrow \text{dom}(\mathfrak{A}_a)$ such that $h(x) = b$ and

- $h(\mathbf{x}) \in R^{\mathfrak{A}_a}$ for all $R(\mathbf{x}) \in \varphi(x)$, and
- $\mathfrak{A}_a \models \mathcal{A}(h(z), h(z'))$ for all $\mathcal{A}(z, z') \in \varphi(x)$.

By definition of \mathfrak{A} and the induction hypothesis applied to tests in $\mathcal{A}(z, z')$, we also have

- $h(\mathbf{x}) \in R^{\mathfrak{A}}$ for all $R(\mathbf{x}) \in \varphi(x)$, and
- $\mathfrak{A} \models \mathcal{A}(h(z), h(z'))$ for all $\mathcal{A}(z, z') \in \varphi(x)$.

Thus, we obtain $\mathfrak{A} \models \varphi(b)$.

For (\Leftarrow) , assume that $\mathfrak{A} \models \varphi(b)$, that is, there is a mapping $h : \text{var}(\varphi) \rightarrow \text{dom}(\mathfrak{A})$ such that $h(x) = b$ and

- $h(\mathbf{x}) \in R^{\mathfrak{A}}$ for all $R(\mathbf{x}) \in \varphi(x)$, and
- $\mathfrak{A} \models \mathcal{A}(h(z), h(z'))$ for all $\mathcal{A}(z, z') \in \varphi(x)$.

Fix $a \in \text{dom}(D)$ such that $b \in \text{dom}(\mathfrak{A}_a)$, and let $U = \{a\} \cup (\text{dom}(\mathfrak{A}) \setminus \text{dom}(\mathfrak{A}_a))$. We first define a query $\psi_0(x)$, which intuitively contains those parts of φ that are mapped to U by h (that is, outside \mathfrak{A}_a); the free variable x ‘represents’ the domain element a . Formally, we process $\varphi(x)$ as follows:

1. for all $R(\mathbf{x}) \in \varphi(x)$ with $h(\mathbf{x}) \not\subseteq \text{dom}(\mathfrak{A}_a)$, we have, by construction of \mathfrak{A} , that $h(\mathbf{x}) \subseteq U$. In this case, we include $R(\mathbf{x}')$ in $\psi_0(x)$, where \mathbf{x}' is obtained from \mathbf{x} by renaming every $y \in \mathbf{x}$ satisfying $h(y) = a$ with x .
2. for all $\mathcal{A}(z, z') \in \varphi(x)$, there are sequences $a_1, \dots, a_n, s_0, \dots, s_n$, and a word $\nu_1 \cdots \nu_{n-1} \in L(\mathcal{A})$ such that $a_1 = h(z)$, $a_n = h(z')$, $s_0 = q_0$, $s_n \in F$, $(s_i, \nu_i, s_{i+1}) \in \Delta$, for all $i \in \{1, \dots, n-1\}$, $(a_i, a_{i+1}) \in R^{\mathfrak{A}}$ if $\nu_i = R$, $(a_{i+1}, a_i) \in R^{\mathfrak{A}}$ if $\nu_i = R^-$, and $\theta(x) \in \tau(a_i)$ and $a_{i+1} = a_i$ if $\nu_i = \theta(x)$?

For all i, j with $1 \leq i \leq j \leq n$ such that a_i, \dots, a_j is a subsequence of a_1, \dots, a_n maximal with $a_k \in U$ for all $i \leq k \leq j$, we distinguish four cases:

- if $i = 1$ and $j = n$, then add $\mathcal{A}[F/\{s_n\}](z, z')$ to $\psi_0(x)$,
- if $i = 1$ and $j \neq n$, then add $\mathcal{A}[F/\{s_{j-1}\}](z, x)$ to $\psi_0(x)$,
- if $i \neq 1$ and $j = n$, then add $\mathcal{A}[q_0/s_{i-1}, F/\{s_n\}](x, z')$ to $\psi_0(x)$,
- if $i \neq 1$ and $j \neq n$, then add $\mathcal{A}[q_0/s_{i-1}, F/\{s_{j-1}\}](x, x)$ to ψ_0 .

Note that $a_j = a$ in the first, $a_i = a$, in the second, and $a_i = a_j = a$ in the last case.

Let $\psi_1(x), \dots, \psi_k(x)$ be all connected components of $\psi_0(x)$ with x considered as ‘constant’. It is not hard to see that $\psi_i(x) \in \text{qcl}(\varphi_0)$, for every $i \in \{1, \dots, k\}$.

Assume first that $\mathfrak{A}_a \models \psi_i(a)$, for all $i \in \{1, \dots, k\}$. In this case, we can modify h and the witnessing sequences to ‘live’ completely in \mathfrak{A}_a , and thus obtain $\mathfrak{A}_a \models \varphi(b)$. Assume now that $\mathfrak{A}_a \not\models \psi_i(a)$, for some $i \in \{1, \dots, k\}$. By Point 1, we have $\neg\psi_i(x) \in \tau(a)$. We now derive a contradiction to Point 3 using the mapping h from above. For doing so, we assume that $\text{dom}(D) = \{b_1, \dots, b_\ell\}$ and define queries $\vartheta_0(\mathbf{a}_0), \vartheta'_1(b_1), \dots, \vartheta'_\ell(b_\ell)$, where intuitively $\vartheta_0(\mathbf{a}_0)$ contains those atoms from $\psi_i(x)$ which are mapped to $\text{dom}(D)$ by h while $\vartheta'_i(b_i)$ contains all atoms which are mapped to \mathfrak{A}_{b_i} by h . Formally, we proceed as follows:

- for all $R(\mathbf{x}) \in \psi_i(x)$, we either have $h(\mathbf{x}) \in R^D$ or $h(\mathbf{x}) \in R^{\mathfrak{A}_{b_j}}$, for some j . In the former case, add $R(\mathbf{x})$ to $\vartheta_0(\mathbf{a})$; in the latter case, add $R(\mathbf{x})$ to $\vartheta'_j(b_j)$, where \mathbf{x}' is obtained from \mathbf{x} by replacing every $y \in \mathbf{x}$ satisfying $h(y) = b_j$ with b_j ;
- Let $\mathcal{A}^*(z, z')$ be an atom that was added in Step 2 above, and let a_i, \dots, a_j and s_{i-1}, \dots, s_j be the sub-sequences corresponding to this atom which were assumed there. Depending on where the sequence a_i, \dots, a_j lies with respect to $\text{dom}(D)$, we distinguish five cases.
 - If $\{a_i, \dots, a_j\}$ is disjoint from $\text{dom}(D)$, then there is a k such that $\{a_i, \dots, a_j\} \subseteq \text{dom}(\mathfrak{A}_{b_k})$. Add $\mathcal{A}^*(z, z')$ to $\vartheta'_k(b_k)$ in this case.
 - If $\{a_i, \dots, a_j\}$ is not disjoint from $\text{dom}(D)$ and $a_i, a_j \notin \text{dom}(D)$, then let l, u be the unique numbers with $i \leq l \leq u \leq j$ such that $a_l, a_u \in \text{dom}(D)$ but $a_k \notin \text{dom}(D)$ for all $i \leq k < l$ and all $u < k \leq j$. Moreover, fix k, k' such that $a_i \in \text{dom}(\mathfrak{A}_{b_k})$ and $a_j \in \text{dom}(\mathfrak{A}_{b_{k'}})$. Note that $a_l = b_k$ and $a_u = b_{k'}$ in this case. Then add $\mathcal{A}^*[F/\{s_{l-1}\}](z, a_l)$ to $\vartheta'_k(a_l)$, $\mathcal{A}^*[q_0/s_{l-1}, F/\{s_{u-1}\}](a_l, a_u)$ to $\vartheta_0(\mathbf{a}_0)$, and $\mathcal{A}^*[q_0/s_{u-1}](a_u, z')$ to $\vartheta'_{k'}(a_u)$.
 - If $a_i \in \text{dom}(D)$ and $a_j \notin \text{dom}(D)$, then let u be the unique number with $i \leq u \leq j$ such that $a_u \in \text{dom}(D)$, but $a_k \notin \text{dom}(D)$ for all $u < k \leq j$, and fix k such that $a_j \in \text{dom}(\mathfrak{A}_{b_k})$. Note that $a_u = b_k$ in this case. Then add $\mathcal{A}^*[F/\{s_{u-1}\}](a_i, a_u)$ to $\vartheta_0(\mathbf{a}_0)$ and $\mathcal{A}^*[q_0/s_{u-1}](a_u, z')$ to $\vartheta'_k(a_u)$.
 - If $a_j \in \text{dom}(D)$ and $a_i \notin \text{dom}(D)$, then let l be the unique number with $i \leq l \leq j$ such that $a_l \in \text{dom}(D)$, but $a_k \notin \text{dom}(D)$ for all $i \leq k < l$, and fix k such that $a_i \in \text{dom}(\mathfrak{A}_{b_k})$. Note that $a_l = b_k$ in this case. Then add $\mathcal{A}^*[F/\{s_{l-1}\}](z, a_l)$ to $\vartheta'_k(a_l)$ and $\mathcal{A}^*[q_0/s_{l-1}](a_l, a_i)$ to $\vartheta_0(\mathbf{a}_0)$.
 - If $a_i, a_j \in \text{dom}(D)$, then add $\mathcal{A}^*(a_i, a_j)$ to $\vartheta_0(\mathbf{a}_0)$

Now obtain a sequence $\vartheta_1(a_1), \dots, \vartheta_k(a_k)$ by replacing each $\vartheta'_i(b_i)$ with its connected components. It should be clear that $\vartheta_0(\mathbf{a}_0), \vartheta_1(a_1), \dots, \vartheta_k(a_k)$ is a splitting of a subdivision of $\psi_i(a)$. Moreover, h is a homomorphism from $\vartheta_0(\mathbf{a}_0)$ to D given τ . However, by construction of $\vartheta_i(a_i)$, it should be clear that $\mathfrak{A}_{a_i} \models \vartheta_i(a_i)$, for all i . Thus, $\vartheta_i(x) \in \tau(a_i)$ for all i , a contradiction to Condition 3. This finishes the proof of the Claim.

Based on the previous claim, we can establish by structural induction that, for all sentences $\varphi \in \text{ecl}(\varphi_0)$, we have:

$$\varphi \in \tau(a) \text{ for all } a \in \text{dom}(D) \quad \text{iff} \quad \mathfrak{A} \models \varphi.$$

Note that Condition 2 is used to prove the induction base. As $\varphi_0 \in \tau(a_0)$ for some $a_0 \in \text{dom}(D)$, this yields the desired $\mathfrak{A} \models \varphi_0$. \blacktriangleleft

E Proofs for Section 6

► **Theorem 18.** *The UNFO^{reg} model checking problem is P^{NP}[O(log² n)]-complete.*


Proof. The lower bound follows from that for UNFO [47].

For the upper bound, we give a polynomial-time reduction of the model checking problem for UNFO^{reg} to a restricted version of the problem ‘Tree Block Satisfaction’, which was shown to be P^{NP}[O(log² n)]-complete in [45]. This version, called TB(SAT) in [47], is defined as follows.

A *TB-tree* of *width* $k \geq 1$ is a tree consisting of *blocks*, where each block is a kind of Boolean circuit that has k outputs and, for each of its n children, has k inputs,¹ see Figure 1. The i -th output of a block is determined by the values of its inputs in a way defined by an existentially quantified Boolean formula (\exists QBF) χ_i of the form

$$\chi_i = \exists \mathbf{b}_1 c_1 \dots \mathbf{b}_m c_m \mathbf{d} (c_1 = \text{input}_{i_1}(\mathbf{b}_1) \wedge \dots \wedge c_m = \text{input}_{i_m}(\mathbf{b}_m) \wedge \psi), \quad \text{where}$$

- $i_1, \dots, i_m \leq n$;
- each \mathbf{b}_j is a tuple of $\log k$ variables, encoding a number $\#\mathbf{b}_j \leq k$;
- $\text{input}_{i_j}(\mathbf{b}_j)$ represents the value of the $\#\mathbf{b}_j$ -th output bit of the i_j -th child block (e.g., if $\#\mathbf{b}_j = 5$, then $\text{input}_2(\mathbf{b}_j) = y_5^{(2)}$ in Figure 1);
- ψ is a Boolean formula using any of the existentially quantified variables.



■ **Figure 1** A block in a TB-tree of width k with n children.

TB(SAT) is the following problem: given a TB-tree of width k , does the first output bit of the root block have value 1? For the reduction, we show how to construct, for a given UNFO^{reg} sentence φ and structure \mathfrak{A} , a TB-tree $T_{\mathfrak{A},\varphi}$ such that

$$T_{\mathfrak{A},\varphi} \text{ is a yes-instance of TB(SAT) iff } \mathfrak{A} \models \varphi. \tag{*}$$

Let $|\text{dom}(\mathfrak{A})| = k$ and assume a linear order on the elements of \mathfrak{A} from 1 to k . For a given $a \in \text{dom}(\mathfrak{A})$, we use $\#a$ to denote the position of a in this order. We construct $T_{\mathfrak{A},\varphi}$ of width k via induction on the structure of φ . The construction satisfies the following invariant: For every subformula $\psi(x)$ of φ with *at most one* free variable, and every $a \in \text{dom}(\mathfrak{A})$ with $\#a = i$,

$$\text{The } i\text{-th output gate of } T_{\mathfrak{A},\psi(x)} \text{ is true iff } \mathfrak{A} \models \psi(a). \tag{**}$$

In case ψ has no free variable, $\mathfrak{A} \models \psi(a)$ stands for $\mathfrak{A} \models \psi$.

¹ In the general case [45], a block may have additional inputs, which do not connect to children and are thus inputs of the TB-tree. Our version does not allow this; i.e., we restrict ourselves to TB-trees without inputs.

It is easy to see that $(**)$ implies $(*)$: just set $\psi(x) = \varphi$ and $i = 1$. Furthermore, $(**)$ is readily checked in every step of the induction.

The shape of each $T_{\mathfrak{A},\psi(x)}$ will roughly reflect that of the syntactic tree of $\psi(x)$. When we construct the \exists QBFs χ_i of each block, we will use \mathbf{b} to denote a vector of $\log k$ variables and $\#\mathbf{b} = i$ as a shorthand for the Boolean formula expressing that \mathbf{b} represents the binary encoding of i .

Let $\psi(x)$ be a subformula of φ .

Case 1: $\psi(x) = \neg\vartheta(x)$. Construct $T_{\mathfrak{A},\psi(x)}$ from $T_{\mathfrak{A},\vartheta(x)}$ by adding a new root block whose i -th output is defined by the formula that negates the i -th input of the single child:

$$\chi_i := \exists \mathbf{b} c (c = \text{input}_1(\mathbf{b}) \wedge \#\mathbf{b} = i \wedge c = 0)$$

Case 2: $\psi(x)$ is built from atomic formulas and UNFO^{reg} formulas in one free variable using conjunction, disjunction, and existential quantification. Let y_1, \dots, y_n be the variables in $\psi(x)$ that are quantified on the ‘top level’, i.e., outside the scope of any test in $\psi(x)$. Let $\vartheta_1(z_1), \dots, \vartheta_m(z_m)$ the maximal strict subformulas in at most one free variable, where $z_i \in \{x, y_1, \dots, y_n\}$ for all $i \leq m$. We can assume w.l.o.g. that the z_i are distinct and coincide with y_1, \dots, y_m and that x occurs in φ only in atoms of the form $x = y_j$. These assumptions can always be satisfied by introducing additional quantified variables and equality atoms. We construct $T_{\mathfrak{A},\psi(x)}$ from the $T_{\mathfrak{A},\vartheta_i(x)}$ by adding a new root block whose children are the roots of the $T_{\mathfrak{A},\vartheta_i(x)}$, and whose i -th output is defined by the formula

$$\chi_i := \exists \mathbf{b}_1 c_1 \dots \mathbf{b}_m c_m \mathbf{b}_{m+1} \dots \mathbf{b}_n \left(\bigwedge_{j \leq m} c_j = \text{input}_j(\mathbf{b}_j) \right) \wedge \chi_{\mathfrak{A}},$$

where the \mathbf{b}_j, c_j are used to refer to the values of the subformulas $\vartheta_j(y_j)$, the $\mathbf{b}_{m+1}, \dots, \mathbf{b}_n$ correspond to the additional y_j and $\chi_{\mathfrak{A}}$ is obtained from $\psi(x)$ as follows:

- Every subformula $\vartheta_j(y_j)$ is replaced by c_j .
- Every equational atom $x = y_j$ is replaced by $\#\mathbf{b}_j = i$ and $y_j = y_\ell$ by $\#\mathbf{b}_j = \#\mathbf{b}_\ell$.
- Every relational atom $R(y_{j_1}, \dots, y_{j_\ell})$ is replaced by a Boolean formula enumerating all tuples in $R^{\mathfrak{A}}$:

$$\bigvee_{(a_1, \dots, a_\ell) \in R^{\mathfrak{A}}} \left(\#\mathbf{b}_{j_1} = \#a_1 \wedge \dots \wedge \#\mathbf{b}_{j_\ell} = \#a_\ell \right)$$

- Every regular atom $\mathcal{A}(y_j, y_n)$ is replaced with the Boolean formula

$$\bigvee_{a, b \in \text{dom}(\mathfrak{A})} \left(\#\mathbf{b}_j = \#a \wedge \#\mathbf{b}_n = \#b \wedge \alpha_{\mathcal{A}, a, b} \right),$$

where $\alpha_{\mathcal{A}, a, b}$ is an \exists QBF that evaluates to true iff there is a path from element a to b in \mathfrak{A} that is accepted by \mathcal{A} . After bound renaming, the quantifiers from $\alpha_{\mathcal{A}, a, b}$ can be moved forward such that χ_i becomes a well-formed \exists QBF. To encode \mathcal{A} ’s behavior in $\alpha_{\mathcal{A}, a, b}$, we assume that $\mathcal{A} = (Q, \Sigma, q_0, \Delta, F)$, where

- Q with $|Q| = t$ is the set of states;
- $\Sigma = \{R, R^- \mid R \text{ a binary predicate in } \varphi\} \cup \{\vartheta(x)? \mid \vartheta(x)? \text{ a test in } \varphi\}$ is the input alphabet;
- q_0 is the initial state;
- $\Delta \subseteq Q \times \Sigma \times Q$ is the transition relation;
- $F \subseteq Q$ is the set of accepting states.

For every $p, q \in Q$, denote with $\mathcal{A}[p, q]$ the NFA obtained from \mathcal{A} by setting p to be the initial state and q to be the only accepting state.

The encoding uses fresh Boolean variables $x_{p,q,a,b}^\ell$ with $\ell \leq t \cdot k$, $p, q \in Q$, and $a, b \in \text{dom}(\mathfrak{A})$. The truth value of the variable $x_{p,q,a,b}^\ell$ indicates whether there is a path of length ℓ from a to b in \mathfrak{A} that is accepted by $\mathcal{A}[p, q]$. It is clear that, whenever there is some path from a to b accepted by $\mathcal{A}[p, q]$, then there is always a path of length $\leq t \cdot k$ because one can always omit loops between two positions in a path that agree in state and element visited. Therefore the above restriction $\ell \leq t \cdot k$ suffices for a correct modeling of \mathcal{A} 's behavior, and the number of variables needed is polynomial.

The formula $\alpha_{\mathcal{A},a,b}$ enforces the correct truth values of these variables via induction on ℓ and requires that some $x_{q_0,q_f,a,b}^\ell$ with $q_f \in F$ be true:

$$\begin{aligned} \alpha_{\mathcal{A},a,b} = & \exists \mathbf{b}'_{1,1} c'_{1,1} \dots \mathbf{b}'_{m,k} c'_{m,k} \quad \exists_{\substack{h=0,\dots,|Q| \\ p,q \in Q \\ a,b \in \text{dom}(\mathfrak{A})}} x_{p,q,a,b}^h \\ & \left(\bigwedge_{j \leq m, \ell \leq k} c'_{j,\ell} = \text{input}_j(\mathbf{b}'_{j,\ell}) \wedge \# \mathbf{b}'_{j,\ell} = \ell \right) \\ & \wedge \beta_{\mathcal{A}} \wedge \gamma_{\mathcal{A}} \wedge \delta_{\mathcal{A}} \wedge \bigvee_{\ell \leq t \cdot k} \bigvee_{q_f \in F} x_{q_0,q_f,a,b}^\ell, \end{aligned}$$

where the $\mathbf{b}'_{j,\ell}, c'_{j,\ell}$ make the values of all $\vartheta_j(y_j)$ in all elements of the structure accessible for evaluating tests in regular atoms, and the conjuncts $\beta_{\mathcal{A}}, \gamma_{\mathcal{A}}, \delta_{\mathcal{A}}$ have the purpose to set the $x_{\cdot,\cdot,\cdot,\cdot}^\ell$, for $\ell = 0, \ell = 1$, and $\ell \geq 2$, respectively. They are defined as follows.

$$\begin{aligned} \beta_{\mathcal{A}} = & \bigwedge_{\substack{q \in Q \\ a \in \text{dom}(\mathfrak{A})}} x_{q,q,a,a}^0 \wedge \bigwedge_{\substack{p,q \in Q \\ a,b \in \text{dom}(\mathfrak{A}) \\ p \neq q \text{ or } a \neq b}} \neg x_{p,q,a,b}^0 \\ \gamma_{\mathcal{A}} = & \bigwedge_{\substack{(p,R,q) \in Q \\ (a,b) \in R^{\mathfrak{A}}}} x_{p,q,a,b}^1 \wedge \bigwedge_{\substack{(p,R^-,q) \in Q \\ (a,b) \in R^{\mathfrak{A}}}} x_{p,q,b,a}^1 \wedge \bigwedge_{\substack{(p,\vartheta_j^?,q) \in Q \\ a \in \text{dom}(\mathfrak{A})}} (x_{p,q,a,a}^1 \leftrightarrow c'_{\#a}) \wedge \gamma'_{\mathcal{A}}, \end{aligned}$$

where $\gamma'_{\mathcal{A}}$ is the conjunction of $\neg x_{\cdot,\cdot,\cdot,\cdot}^1$ for all $x_{\cdot,\cdot,\cdot,\cdot}^1$ that do not occur in the preceding conjuncts of $\gamma_{\mathcal{A}}$. Finally,

$$\delta_{\mathcal{A}} = \bigwedge_{\substack{p,q \in Q \\ a,b \in \text{dom}(\mathfrak{A}) \\ 1 \leq \ell < t \cdot k}} \left(x_{p,q,a,b}^{\ell+1} \leftrightarrow \bigvee_{\substack{r \in Q \\ c \in \text{dom}(\mathfrak{A})}} (x_{p,r,a,c}^1 \wedge x_{r,q,c,b}^{\ell+1}) \right)$$

Case 2 also covers the case where $\psi(x)$ has no strict subformula with at most one free variable. It is easy to check that the invariant (***) holds and that $T_{\mathfrak{A},\varphi}$ can be constructed in polynomial time. \blacktriangleleft