

(FO-rewritings) can be exponentially (superpolynomially) more succinct than PE-rewritings. These prohibitively high lower bounds are caused by the fact that the chases (canonical models) for \mathcal{T}_n contain full binary trees of depth n and give rise to *exponentially-many* homomorphisms from \mathbf{q}_n to the labelled nulls of the chases, all of which have to be reflected in the rewritings of \mathbf{q}_n and \mathcal{T}_n .

In this paper, we investigate succinctness of CQ rewritings over ‘shallow’ ontologies whose (polynomial-size) chases are finite trees of depth 1 or 2 (which do not have chains of more than 1 or 2 labelled nulls). From the theoretical point of view, ontologies of depth 1 are important because their chases can only generate linearly-many homomorphisms of CQs to the labelled nulls; on the other hand, shallow ontologies are typical in the real-world OBDA applications. We obtain both positive and, unexpectedly, ‘negative’ results summarised below:

- (i) any CQ and ontology of depth 1 have a polynomial-size NDL-rewriting;
- (ii) PE-rewritings of some CQs and ontologies of depth 1 are of superpolynomial size;
- (iii) any tree-shaped CQ and ontology of depth 1 have a PE-rewriting of polynomial size;
- (iv) the existence of polynomial-size FO-rewritings for all CQs and ontologies of depth 1 is equivalent to an open problem ‘NL/poly \subseteq NC¹?’;
- (v) NDL- and PE-rewritings of some CQs and ontologies of depth 2 are of exponential size, while FO-rewritings are of superpolynomial size unless NP \subseteq P/poly.

We prove (i)–(v) by establishing a fundamental connection between FO-, PE- and NDL-rewritings, on the one hand, and, respectively, formulas, monotone formulas and monotone circuits computing certain monotone Boolean functions, on the other. These functions are associated with hypergraph representations of the tree-witness rewritings [22], reflecting possible homomorphisms of the given CQ to the labelled nulls of the chases for the given ontology. In particular, any hypergraph H of degree 2 (every vertex in which belongs to 2 hyperedges) corresponds to a CQ \mathbf{q}_H and an ontology \mathcal{T}_H of depth 1 such that answering \mathbf{q}_H over \mathcal{T}_H and single-individual data instances amounts to computing the hypergraph function for H . We show that representing Boolean functions as hypergraphs of degree 2 is polynomially equivalent to representing their duals as nondeterministic branching programs (NBPs) [18]. This correspondence and known results on NBPs [31, 19] give (i), (ii) and (iv) above. To prove (v), we observe that hypergraphs of degree 3 are computationally as powerful as nondeterministic Boolean circuits (NP/poly) and encode the function CLIQUE _{n,k} (\vec{e}) (graph \vec{e} with n vertices has a k -clique) as CQs over ontologies of depth 2. It also follows that there exist polynomial-size FO-rewritings for all CQs and ontologies (of depth 2) with polynomially-many tree witnesses iff all functions in NP/poly are computed by polynomial-size formulas, that is, iff NP/poly \subseteq NC¹ (which is a well-known open problem). Finally, we show that any tree-shaped CQ \mathbf{q} and ontology \mathcal{T} have a PE-rewriting of size $O(|\mathcal{T}|^2 \cdot |\mathbf{q}|^{1+\log d})$, where d is a parameter related to the number of tree witnesses sharing a common variable. This gives (iii) since $d = 2$ for ontologies of depth 1. We also note that the problem ‘ $\mathcal{T}, \mathcal{A} \models \mathbf{q}$?’, for tree-shaped Boolean CQs and any \mathcal{T} , is fixed-parameter tractable (recall that the problem ‘ $\mathcal{A} \models \mathbf{q}$?’, for tree-shaped \mathbf{q} , is known to be tractable [35], while ‘ $\mathcal{T}, \mathcal{A} \models \mathbf{q}$?’ is NP-hard [21]).

As shown in [16], exponential rewritings can be made polynomial at the expense of polynomially-many additional existential quantifiers over a domain with *two constants* not necessarily occurring in the CQs; cf. [6]. Intuitively, given \mathbf{q} , \mathcal{T} and \mathcal{A} , the extra quantifiers guess a homomorphism from \mathbf{q} to the chase for $(\mathcal{T}, \mathcal{A})$, whereas the standard rewritings (without extra constants) represent such homomorphisms explicitly (likewise NFAs are exponentially more succinct than DFAs, and \exists -QBFs are exponentially more succinct than SAT). A more practical utilisation of additional constants was suggested in the combined approach to OBDA [24], where they are used to construct a polynomial-size encoding of the chase for the given ontology and data over which the original CQ is evaluated. This encoding may introduce (exponentially-many in the worst case) spurious answers that are eliminated by a special polynomial-time filtering procedure.

2 The Tree-Witness Rewriting

In this paper, we assume that an *ontology*, \mathcal{T} , is a finite set of *tuple-generating dependencies* (tgds) of the form

$$\forall \vec{x} (\varphi(\vec{x}) \rightarrow \exists \vec{y} \bigwedge \psi_i(\vec{x}, \vec{y})), \quad (2)$$

where φ and the ψ_i are unary or binary atoms without constants and $|\vec{x} \cup \vec{y}| \leq 2$. These tgds are expressible via tgds in (1) using fresh binary predicates, whereas disjointness constraints in (1) do not contribute to the size of rewritings. Although the language given by (1) is slightly different from *OWL 2 QL*, all the results obtained here are applicable to *OWL 2 QL* ontologies as well. When writing tgds, we will omit the universal quantifiers. The size, $|\mathcal{T}|$, of \mathcal{T} is the number of predicate occurrences in \mathcal{T} . A *data instance*, \mathcal{A} , is a finite set of ground atoms. The set of individual constants in \mathcal{A} is denoted by $\text{ind}(\mathcal{A})$. Taken together, \mathcal{T} and \mathcal{A} form the *knowledge base* (KB) $(\mathcal{T}, \mathcal{A})$. To simplify notation, we will assume that the data instances in all KBs are *complete* in the following sense: for any ground atom $S(\vec{a})$ with $\vec{a} \subseteq \text{ind}(\mathcal{A})$, if $\mathcal{T}, \mathcal{A} \models S(\vec{a})$ then $S(\vec{a}) \in \mathcal{A}$ (see Lemma 1 below).

A *conjunctive query* (CQ) $\mathbf{q}(\vec{x})$ is a formula $\exists \vec{y} \varphi(\vec{x}, \vec{y})$, where φ is a conjunction of unary or binary atoms $S(\vec{z})$ with $\vec{z} \subseteq \vec{x} \cup \vec{y}$ (without loss of generality, we assume that CQs do not contain constants). A tuple $\vec{a} \subseteq \text{ind}(\mathcal{A})$ is a *certain answer to $\mathbf{q}(\vec{x})$ over $(\mathcal{T}, \mathcal{A})$* if $\mathcal{I} \models \mathbf{q}(\vec{a})$ for all models \mathcal{I} of \mathcal{T} and \mathcal{A} ; in this case we write $\mathcal{T}, \mathcal{A} \models \mathbf{q}(\vec{a})$. If $\vec{x} = \emptyset$ then the CQ \mathbf{q} is called *Boolean*; a certain answer to such a \mathbf{q} over $(\mathcal{T}, \mathcal{A})$ is ‘yes’ if $\mathcal{T}, \mathcal{A} \models \mathbf{q}$ and ‘no’ otherwise. Where convenient, we regard a CQ as the set of its atoms.

Given a CQ $\mathbf{q}(\vec{x})$ and an ontology \mathcal{T} , an FO-formula $\mathbf{q}'(\vec{x})$ *without constants* is called an *FO-rewriting of $\mathbf{q}(\vec{x})$ and \mathcal{T}* if, for any (complete) data instance \mathcal{A} and any $\vec{a} \subseteq \text{ind}(\mathcal{A})$, we have $(\mathcal{T}, \mathcal{A}) \models \mathbf{q}(\vec{a})$ iff $\mathcal{A} \models \mathbf{q}'(\vec{a})$? If \mathbf{q}' is a positive existential formula, we call it a *PE-rewriting of \mathbf{q} and \mathcal{T}* . We also consider rewritings in the form of nonrecursive datalog queries. Recall [1] that a *datalog program*, Π , is a finite set of Horn clauses $\forall \vec{x} (\gamma_1 \wedge \dots \wedge \gamma_m \rightarrow \gamma_0)$, where each γ_i is an atom of the form $P(x_1, \dots, x_l)$ with $x_i \in \vec{x}$. The atom γ_0 is the *head* of the clause, and $\gamma_1, \dots, \gamma_m$ its *body*. All variables in the head must also occur in the body. A predicate P *depends* on Q in Π if Π has a clause with P in the head and Q in the body; Π is *nonrecursive* if this dependence relation is acyclic. For a nonrecursive program Π and an atom $\mathbf{q}'(\vec{x})$, (Π, \mathbf{q}') is called an *NDL-rewriting of $\mathbf{q}(\vec{x})$ and \mathcal{T}* in case $\mathcal{T}, \mathcal{A} \models \mathbf{q}(\vec{a})$ iff $\Pi, \mathcal{A} \models \mathbf{q}'(\vec{a})$, for any (complete) \mathcal{A} and $\vec{a} \subseteq \text{ind}(\mathcal{A})$. Rewritings over *arbitrary* data are defined without stipulating that the data instances in KBs are complete.

Lemma 1. (i) *For any (PE-) FO-rewriting \mathbf{q}' of \mathbf{q} and \mathcal{T} over complete data, there is a (PE-) FO-rewriting \mathbf{q}'' over arbitrary data with $|\mathbf{q}''| \leq O(|\mathbf{q}'| \cdot |\mathcal{T}|)$.*

(ii) *For any NDL-rewriting (Π, \mathbf{q}') of \mathbf{q} and \mathcal{T} over complete data, there is an NDL-rewriting (Π', \mathbf{q}') over arbitrary data with $|\Pi'| \leq |\Pi| + O(|\mathcal{T}|)$.*

Proof. Given \mathcal{T} and \mathbf{q} , we define a partial order \leq on the formulas $\varrho(x)$ of the form $S(x)$, $S(x, x)$, $\exists y S(x, y)$ and $\exists y S(y, x)$, where S is a predicate in \mathcal{T} or \mathbf{q} , to be the transitive and reflexive closure of the following relation \leq_1 :

$$\varrho_1(x) \leq_1 \varrho_2(x) \quad \text{iff} \quad \varrho_1(x) \rightarrow \varrho_2(x) \in \mathcal{T}.$$

This partial order \leq is extended to the formulas $\varrho(x_1, x_2)$ of the form $S(x_1, x_2)$ and $S(x_2, x_1)$, where S is a binary predicate in \mathcal{T} or \mathbf{q} , by adding the transitive and reflexive closure of the following relation \leq_2 :

$$\varrho_1(\vec{x}) \leq_2 \varrho_2(\vec{x}) \quad \text{iff} \quad \varrho_1(\vec{x}) \rightarrow \varrho_2(\vec{x}) \in \mathcal{T}.$$

We also define an equivalence relation on such $\varrho(\vec{x})$ by taking $\varrho_1(\vec{x}) \equiv \varrho_2(\vec{x})$ iff $\varrho_1(\vec{x}) \leq \varrho_2(\vec{x})$ and $\varrho_2(\vec{x}) \leq \varrho_1(\vec{x})$. In each equivalence class, we fix a representative, denoted $\varrho(\vec{x})/\equiv$.

(i) For PE- and FO-rewritings, we replace each $S(\vec{x})$ in \mathbf{q}' with a disjunction of all $\varrho(\vec{x})$ such that $\varrho(\vec{x}) \leq S(\vec{x})$. The size of the resulting rewriting \mathbf{q}'' increases linearly in $|\mathcal{T}|$.

²Thus, we do not allow the rewriting from [16] since it contains constants.

(ii) We assume without loss of generality that q' is not a predicate name in \mathcal{T} . Let Π^* be the result of replacing each predicate name S in Π that occurs in \mathcal{T} with a fresh predicate name S^* . Define Π' to be the union of Π^* and the following clauses:

$$\begin{aligned} S^*(\vec{x}) &\leftarrow \varrho(\vec{x}), & \text{for all } \varrho(\vec{x}) \text{ with } \varrho(\vec{x}) \equiv S(\vec{x}), \\ S^*(\vec{x}) &\leftarrow \varrho^*(\vec{x}), & \text{for all } \varrho(\vec{x})/\equiv \text{ such that } \varrho(\vec{x}) \text{ is an immediate predecessor of } S(\vec{x}) \text{ in } \leq, \end{aligned}$$

where $\varrho^*(\vec{x})$ is the result of replacing the predicate, S_1 , in ϱ with S_1^* . It should be clear that (Π', q') is an NDL-rewriting of q and \mathcal{T} over arbitrary data and that the size of the additional clauses in Π' is linear in $|\mathcal{T}|$ (and does not depend on q). \square

We now define an improved version of the tree-witness PE-rewriting [22] that will be used to establish links with formulas and circuits computing certain monotone Boolean functions.

As is well-known [1], for any KB $(\mathcal{T}, \mathcal{A})$, there is a *canonical model* (or *chase*) $\mathcal{C}_{\mathcal{T}, \mathcal{A}}$ such that $\mathcal{T}, \mathcal{A} \models q(\vec{a})$ iff $\mathcal{C}_{\mathcal{T}, \mathcal{A}} \models q(\vec{a})$, for all CQs $q(\vec{x})$ and $\vec{a} \subseteq \text{ind}(\mathcal{A})$. The domain of $\mathcal{C}_{\mathcal{T}, \mathcal{A}}$ consists of $\text{ind}(\mathcal{A})$ and the witnesses, or *labelled nulls*, introduced by the existential quantifiers in \mathcal{T} .

For any formula $\varrho(x)$ of the form $S(x)$, $S(x, x)$, $\exists y S(x, y)$ or $\exists y S(y, x)$, where S is a predicate in \mathcal{T} , we denote by $\mathcal{C}_{\mathcal{T}}^{\varrho(a)}$ the canonical model of the KB $(\mathcal{T} \cup \{A(x) \rightarrow \varrho(x)\}, \{A(a)\})$, where A is a fresh unary predicate. We say that \mathcal{T} is of *depth* k , $1 \leq k < \omega$, if one of the $\mathcal{C}_{\mathcal{T}}^{\varrho(a)}$ contains a chain of the form $R_0(w_0, w_1) \dots R_{k-1}(w_{k-1}, w_k)$, with not necessarily distinct w_i , but none of the $\mathcal{C}_{\mathcal{T}}^{\varrho(a)}$ has such a chain of greater length.

Suppose we are given a CQ $q(\vec{x}) = \exists \vec{y} \varphi(\vec{x}, \vec{y})$ and an ontology \mathcal{T} . For a pair $\mathfrak{t} = (\mathfrak{t}_r, \mathfrak{t}_i)$ of disjoint sets of variables in q , with $\mathfrak{t}_i \subseteq \vec{y}$ and $\mathfrak{t}_i \neq \emptyset$ (\mathfrak{t}_r can be empty), set

$$\mathbf{q}_{\mathfrak{t}} = \{ S(\vec{z}) \in q \mid \vec{z} \subseteq \mathfrak{t}_r \cup \mathfrak{t}_i \text{ and } \vec{z} \not\subseteq \mathfrak{t}_i \}.$$

We call $\mathfrak{t} = (\mathfrak{t}_r, \mathfrak{t}_i)$ a *tree witness for q and \mathcal{T} generated by ϱ* if $\mathbf{q}_{\mathfrak{t}}$ is a minimal subset of q for which there exists a homomorphism $h: \mathbf{q}_{\mathfrak{t}} \rightarrow \mathcal{C}_{\mathcal{T}}^{\varrho(a)}$ such that $\mathfrak{t}_r = h^{-1}(a)$ and $\mathbf{q}_{\mathfrak{t}}$ contains all atoms of q with at least one variable from \mathfrak{t}_i (cf. aggregated unifiers from [23]). Note that the same tree witness $\mathfrak{t} = (\mathfrak{t}_r, \mathfrak{t}_i)$ can be generated by different ϱ . Now, we set

$$\text{tw}_{\mathfrak{t}}(\mathfrak{t}_r) = \bigvee_{\mathfrak{t} \text{ generated by } \varrho} \exists z (\varrho(z) \wedge \bigwedge_{x \in \mathfrak{t}_r} (x = z)). \quad (3)$$

The variables in \mathfrak{t}_i do not occur in $\text{tw}_{\mathfrak{t}}$ and are called *internal*. The length, $|\text{tw}_{\mathfrak{t}}|$, of $\text{tw}_{\mathfrak{t}}$ is $O(|q| \cdot |\mathcal{T}|)$. Tree witnesses \mathfrak{t} and \mathfrak{t}' are *conflicting* if $\mathbf{q}_{\mathfrak{t}} \cap \mathbf{q}_{\mathfrak{t}'} \neq \emptyset$. Denote by $\Theta_{\mathcal{T}}^q$ the set of tree witnesses for q and \mathcal{T} . A subset $\Theta \subseteq \Theta_{\mathcal{T}}^q$ is *independent* if no pair of distinct tree witnesses in it is conflicting. Let $\mathbf{q}_{\Theta} = \bigcup_{\mathfrak{t} \in \Theta} \mathbf{q}_{\mathfrak{t}}$. The following PE-formula \mathbf{q}_{tw} is called the *tree-witness rewriting of q and \mathcal{T}* :

$$\mathbf{q}_{\text{tw}}(\vec{x}) = \bigvee_{\Theta \subseteq \Theta_{\mathcal{T}}^q \text{ independent}} \exists \vec{y} \left(\bigwedge_{S(\vec{z}) \in \mathbf{q}_{\Theta}} S(\vec{z}) \wedge \bigwedge_{\mathfrak{t} \in \Theta} \text{tw}_{\mathfrak{t}}(\mathfrak{t}_r) \right). \quad (4)$$

Example 2. Consider the following ontology and CQ:

$$\begin{aligned} \mathcal{T} &= \{ A_1(x) \rightarrow \exists y (R_1(x, y) \wedge Q(x, y)), A_2(x) \rightarrow \exists y (R_2(x, y) \wedge Q(y, x)) \}, \\ \mathbf{q}(x_1, x_2) &= \exists y_1 y_2 (R_1(x_1, y_1) \wedge Q(y_2, y_1) \wedge R_2(x_2, y_2)). \end{aligned}$$

The CQ q is shown in Fig. 1 alongside the $\mathcal{C}_{\mathcal{T}}^{A_k(a)}$, $k = 1, 2$. There are two tree witnesses, \mathfrak{t}^1 and \mathfrak{t}^2 , for q and \mathcal{T} with

$$\mathbf{q}_{\mathfrak{t}^1} = \{ R_1(x_1, y_1), Q(y_2, y_1) \} \quad \text{and} \quad \mathbf{q}_{\mathfrak{t}^2} = \{ Q(y_2, y_1), R_2(x_2, y_2) \}$$

(they are shown as shaded rectangles in Fig. 1); the tree witness $\mathfrak{t}^k = (\mathfrak{t}_r^k, \mathfrak{t}_i^k)$, for $k = 1, 2$, is generated by $A_k(x)$ with $\mathfrak{t}_r^k = \{x_k, y_{3-k}\}$ and $\mathfrak{t}_i^k = \{y_k\}$, which gives

$$\text{tw}_{\mathfrak{t}^k}(x_k, y_{3-k}) = \exists z (A_k(z) \wedge (x_k = z) \wedge (y_{3-k} = z)).$$




Figure 1: Query $q(x_1, x_2)$ and canonical models $\mathcal{C}_{\mathcal{T}}^{A_1(a)}$ and $\mathcal{C}_{\mathcal{T}}^{A_2(a)}$ from Example 2.

As t^1 and t^2 are conflicting, we obtain the following rewriting:

$$\exists y_1 y_2 [(R_1(x_1, y_1) \wedge Q(y_2, y_1) \wedge R_2(x_2, y_2)) \vee (R_2(x_2, y_2) \wedge \text{tw}_{t^1}) \vee (R_1(x_1, y_1) \wedge \text{tw}_{t^2})].$$

Theorem 3 ([22]). *For any complete data instance \mathcal{A} and any $\vec{a} \subseteq \text{ind}(\mathcal{A})$, we have $\mathcal{T}, \mathcal{A} \models q(\vec{a})$ iff $\mathcal{A} \models q_{\text{tw}}(\vec{a})$.*

The number of tree witnesses, $|\Theta_{\mathcal{T}}^q|$, is bounded by $3^{|q|}$. On the other hand, there is a sequence of queries q_n and ontologies \mathcal{T}_n with exponentially many (in $|q_n|$) tree witnesses [22]. The length of q_{tw} is $O(2^{|\Theta_{\mathcal{T}}^q|} \cdot |q| \cdot |\mathcal{T}|)$. If any two tree-witnesses for q and \mathcal{T} are *compatible*—that is, they are either non-conflicting or one is included in the other—then q_{tw} can be equivalently transformed into the PE-rewriting

$$q'_{\text{tw}}(\vec{x}) = \exists \vec{y} \bigwedge_{S(\vec{z}) \in \mathbf{q}} (S(\vec{z}) \vee \bigvee_{t \in \Theta_{\mathcal{T}}^q \text{ with } S(\vec{z}) \in \mathbf{q}_t} \text{tw}_t(t_r))$$

of size $O(|\Theta_{\mathcal{T}}^q| \cdot |q|^2 \cdot |\mathcal{T}|)$. Our aim now is to investigate transformations of this kind in the more abstract setting of Boolean functions. In Section 5, we shall see an example of q and \mathcal{T} with only $|q|$ -many tree witnesses any PE-rewriting of which is of superpolynomial size because of multiple combinations of incompatible tree witnesses.

3 Hypergraph Functions

The rewriting q_{tw} gives rise to monotone Boolean functions we call hypergraph functions. For the complexity theory of monotone Boolean functions, the reader is referred to [3, 18]. Let $H = (V, E)$ be a hypergraph with *vertices* $v \in V$ and *hyperedges* $e \in E$, $E \subseteq 2^V$. A subset $X \subseteq E$ is *independent* if $e \cap e' = \emptyset$, for any distinct $e, e' \in X$. Denote by V_X the set of vertices occurring in the hyperedges of X . With each $v \in V$ and $e \in E$ we associate propositional variables p_v and p_e , respectively. The *hypergraph function* f_H for H is given by the Boolean formula

$$f_H = \bigvee_{X \subseteq E \text{ independent}} \left(\bigwedge_{v \in V \setminus V_X} p_v \wedge \bigwedge_{e \in X} p_e \right). \quad (5)$$

The rewriting q_{tw} of q and \mathcal{T} defines a hypergraph $H_{\mathcal{T}}^q$ whose vertices are the atoms of q and hyperedges are the sets \mathbf{q}_t , for $t \in \Theta_{\mathcal{T}}^q$. Formula (5) for $H_{\mathcal{T}}^q$ is the same as rewriting (4) with the atoms $S(\vec{z}) \in \mathbf{q}$ and the tree witness formulas tw_t , for $t \in \Theta_{\mathcal{T}}^q$, treated as propositional variables, $p_{S(\vec{z})}$ and p_t , respectively.

Example 4. For q and \mathcal{T} from Example 2, the hypergraph $H_{\mathcal{T}}^q$ is shown in Fig. 2 and

$$f_{H_{\mathcal{T}}^q} = (p_{R_1(x_1, y_1)} \wedge p_{Q(y_2, y_1)} \wedge p_{R_2(x_2, y_2)}) \vee (p_{R_2(x_2, y_2)} \wedge p_{t^1}) \vee (p_{R_1(x_1, y_1)} \wedge p_{t^2}).$$

Suppose the function $f_{H_{\mathcal{T}}^q}$ is computed by some Boolean formula χ . Consider the FO-formula obtained by adding the prefix $\exists \vec{y}$ to χ and replacing each $p_{S(\vec{z})}$ in it with $S(\vec{z})$ and each p_t with the formula $\text{tw}_t(t_r)$ of length $O(|q| \cdot |\mathcal{T}|)$. By comparing (5) and (4), we see that the resulting FO-formula is a rewriting of q and \mathcal{T} . This gives the first claim in the following theorem; the second one requires some basic skills in datalog programming. (Recall [3] that *monotone* Boolean formulas and circuits contain only \wedge and \vee .)




Figure 2: Hypergraph H_T^q for q and \mathcal{T} from Example 2.

Theorem 5. *If $f_{H_T^q}$ is computed by a (monotone) Boolean formula χ then there is a (PE-) FO-rewriting of q and \mathcal{T} of size $O(|\chi| \cdot |q| \cdot |\mathcal{T}|)$.*

If $f_{H_T^q}$ is computed by a monotone Boolean circuit C then there is an NDL-rewriting of q and \mathcal{T} of size $O(|C| \cdot |q| \cdot |\mathcal{T}|)$.

Proof. The first claim was explained above, so we only prove the second one. Let $q(\vec{x}) = \exists \vec{y} \varphi(\vec{x}, \vec{y})$. First, we define a unary predicate D_0 by taking the following clauses

$$D_0(z) \leftarrow \varrho(z), \quad (6)$$

for every formula $\varrho(z)$ of the form $S(z)$, $\exists y S(z, y)$ and $\exists y S(y, z)$, where S is a predicate occurring in \mathcal{T} or q (we say that such a ϱ is *in the signature of \mathcal{T} and q*). Intuitively, the interpretation of D_0 contains all the individual constants of the given data instance. We then set $\vec{z} = \vec{x} \cup \vec{y}$ and define a $|\vec{z}|$ -ary predicate D by the following clause

$$D(\vec{z}) \leftarrow \bigwedge_{z \in \vec{z}} D_0(z). \quad (7)$$

We need the predicate D to ensure that all the clauses in our NDL-program are safe (that is, every variable in the head of a clause also occurs in the body).

Suppose H_T^q has m vertices and l hyperedges. Let g_1, \dots, g_n be the nodes of C ordered in such a way that g_1, \dots, g_m correspond to the atoms $S_1(\vec{z}_1), \dots, S_m(\vec{z}_m)$ of q , g_{m+1}, \dots, g_{m+l} correspond to the tree witnesses t^1, \dots, t^l and g_{m+l+1}, \dots, g_n correspond to the gates of C with g_n its output. For $1 \leq i \leq m$, we take the clauses

$$G_i(\vec{z}) \leftarrow S_i(\vec{z}_i) \wedge D(\vec{z}). \quad (8)$$

For $m < i \leq m + l$, take the clauses

$$G_i(\vec{z}) \leftarrow \varrho(z_0) \wedge \bigwedge_{y \in t_i^{i-m}} (z_0 = y) \wedge D(\vec{z}), \quad \text{for all } \varrho(z) \text{ with } t_i^{i-m} \text{ is generated by } \varrho(z). \quad (9)$$

where z_0 is a fresh variable. For $i > m + l$, we take the clauses

$$G_i(\vec{z}) \leftarrow G_j(\vec{z}) \wedge G_{j'}(\vec{z}) \wedge D(\vec{z}), \quad \text{if } g_i = g_j \wedge g_{j'}, \quad (10)$$

$$\left. \begin{array}{l} G_i(\vec{z}) \leftarrow G_j(\vec{z}) \wedge D(\vec{z}), \\ G_i(\vec{z}) \leftarrow G_{j'}(\vec{z}) \wedge D(\vec{z}), \end{array} \right\} \quad \text{if } g_i = g_j \vee g_{j'}. \quad (11)$$

Denote the resulting set of clauses (6)–(11) by Π . We claim that (Π, G_n) is an NDL-rewriting of q and \mathcal{T} over complete data. To see this, we can transform (Π, G_n) to a PE-formula of the form

$$\exists \vec{y} \left[\psi(\vec{x}, \vec{y}) \wedge \bigwedge_{z \in \vec{x} \cup \vec{y}} \left(\bigvee_{\varrho \text{ in signature of } \mathcal{T}, q} \varrho(z) \right) \right],$$

where $\exists \vec{y} \psi(\vec{x}, \vec{y})$ can be constructed by taking the Boolean formula representing C and replacing $p_S(\vec{z})$ with $S(\vec{z})$ and p_t with $\text{tw}_t(t)$. It follows from the first claim of the theorem that $\exists \vec{y} \psi(\vec{x}, \vec{y})$ is a rewriting of q and \mathcal{T} over complete data. It should be clear that the big conjunction does not change this fact. \square

Thus, the problem of constructing short rewritings is reducible to the problem of finding short (monotone) Boolean formulas or circuits computing the hypergraph functions.

In the next section, we consider hypergraphs as programs for computing Boolean functions and compare them with the well-known formalisms of nondeterministic branching programs (NBPs) and nondeterministic Boolean circuits [3, 18].

4 Hypergraphs, NBPs and Nondeterministic Boolean Circuits

Let p_1, \dots, p_n be propositional variables. An *input* to a hypergraph program or an NBP is a vector $\vec{\alpha} \in \{0, 1\}^n$ assigning the truth-value $\vec{\alpha}(p_i)$ to each of the p_i . We extend this notation to negated variables and constants by setting $\vec{\alpha}(\neg p_i) = \neg \vec{\alpha}(p_i)$, $\vec{\alpha}(0) = 0$ and $\vec{\alpha}(1) = 1$.

A *hypergraph program* (HGP) is a hypergraph $H = (V, E)$ in which every vertex is labelled with 0, 1, p_i or $\neg p_i$. We say that the hypergraph program H *computes* a Boolean function f in case, for any input $\vec{\alpha}$, we have $f(\vec{\alpha}) = 1$ iff there is an independent subset in E that *covers all zeros*—that is, contains all the vertices in V labelled with 0 under $\vec{\alpha}$. A hypergraph program is *monotone* if there are no negated variables among its vertex labels. The *size*, $|H|$, of a hypergraph program H is the number of hyperedges in it.

We say that a hypergraph (program) H is of *degree* $\leq n$ if every vertex in it belongs to at most n hyperedges; H is of *degree* n if every vertex in it belongs to exactly n hyperedges. We denote by $\text{HGP}(f)$ ($\text{HGP}^n(f)$) the minimal size of hypergraph programs (of degree $\leq n$) computing f ; $\text{HGP}_+(f)$ and $\text{HGP}_+^n(f)$ are used for the size of monotone programs.

We show first that monotone hypergraph programs of degree ≤ 2 capture the computational power of hypergraph functions for hypergraphs of degree ≤ 2 . On the one hand, a monotone hypergraph program H computes the subfunction of f_H obtained by setting $p_e = 1$, for all $e \in E$, and setting p_v to be equal to the label of v . On the other hand, any hypergraph function f_H can be computed by a monotone hypergraph program of degree 2 and size $O(|H|)$.

Lemma 6. *For any hypergraph H of degree $\leq n$, there is a monotone HGP of degree $\leq \max(2, n)$ and size $2|H|$ computing the function f_H .*

Proof. Given a hypergraph $H = (V, E)$, we label each $v \in V$ by a variable p_v . For each $e \in E$, we add a fresh vertex a_e labelled with 1 and a fresh vertex b_e labelled with p_e ; then we create a new hyperedge $e' = \{a_e, b_e\}$ and add a_e to the hyperedge e . We claim that the resulting hypergraph program H' computes f_H . Indeed, for any input $\vec{\alpha}$ with $\vec{\alpha}(p_e) = 0$, we have to include the edge e' into the cover, and so cannot include the edge e itself. Thus, the program returns 1 iff there is an independent set X of hyperedges with $\vec{\alpha}(p_e) = 1$, for all $e \in X$, covering all zeros of the variables p_v . It follows that H' computes f_H . \square

Lemma 7. *If f is computable by a (monotone) HGP H of degree ≤ 2 , then it can also be computed by a (monotone) HGP of degree 2 and size $|H| + 3$.*

Proof. Let v_1, \dots, v_k be vertices of degree 0 and v_{k+1}, \dots, v_l vertices of degree 1 in H . It suffices to extend H with vertices, x, y, z labelled with 1, 0, 0, respectively, and hyperedges $e_1 = \{v_1, \dots, v_l, x, y\}$, $e_2 = \{v_1, \dots, v_k, x, z\}$ and $e_3 = \{y, z\}$. It is easy to see that each cover should contain e_3 but cannot contain e_1, e_2 . Indeed, y and z should both be covered. However, e_1 and e_2 intersect and cannot be both in the same cover. Thus, y and z should be covered by e_3 , while e_1 and e_2 , intersecting e_3 , are not in the cover. After these choices we are left to deal with the original hypergraph. Clearly, this construction preserves monotonicity. \square

Our next result in this section establishes a link between hypergraph programs of degree ≤ 2 and NBPs. Recall [18] that an NBP is a directed multigraph with two distinguished vertices, s and t , and the arcs labelled with 0, 1, p_i or $\neg p_i$ (the arcs of the first type have no effect, the arcs of the second type are called *rectifiers*, and those of the third and fourth types *contacts*). We assume that s has no incoming and t no outgoing arcs, and note that NBPs may have multiple parallel arcs (with distinct labels) connecting two nodes. We write $v \rightarrow_{\vec{\alpha}} v'$ if there is a directed path from v to v' labelled with 1 under $\vec{\alpha}$. An NBP *computes* a Boolean function f if $f(\vec{\alpha}) = 1$ just in case $s \rightarrow_{\vec{\alpha}} t$. The *size* of an NBP is the number of arcs in it. An NBP is *monotone* if it has no negated variables among its labels. We denote by $\text{NBP}(f)$ (respectively, $\text{NBP}_+(f)$) the minimal size of (monotone) NBPs computing f . $\text{NBP}(\text{poly})$ is the class of Boolean functions computable by polynomial-size NBPs. As usual, f^* is the Boolean function dual to f .

Theorem 8. (i) *For any Boolean function f , $\text{HGP}^2(f)$ and $\text{NBP}(\neg f)$ are polynomially related.*
(ii) *For any monotone Boolean function f , $\text{HGP}_+^2(f)$ and $\text{NBP}_+(f^*)$ are polynomially related.*

Proof. We only prove (i); (ii) is proved by the same argument. Suppose $\neg f$ is computed by an NBP G . We construct a hypergraph program H of degree ≤ 2 as follows. For each arc e in G , H has two vertices e^0 and e^1 , which represent the beginning and the end of e . The vertex e^0 is labelled with the *negated* label of e in G and e^1 with 1. We also add to H a vertex t labelled with 0. For each arc e in G , H has an e -hyperedge $\{e^0, e^1\}$. For each vertex v in G but s and t , H has a v -hyperedge that consists of all vertices e^1 , for the arcs e leading to v , and all vertices e^0 , for the arcs e leaving v . For the vertex t , H contains a hyperedge that consists of t and all vertices e^1 , for the arcs e leading to t . We claim that the constructed hypergraph program H computes f . Indeed, if $s \not\rightarrow_{\vec{\alpha}} t$ in G then the following subset of hyperedges is independent and covers all zeros: all e -hyperedges, for the arcs e reachable from s and labelled with 1 under $\vec{\alpha}$, and all v -hyperedges with $s \not\rightarrow_{\vec{\alpha}} v$. Conversely, if $s \rightarrow_{\vec{\alpha}} t$ then it can be shown by induction that, for each arc e_i of the path, the e_i -hyperedge must be in the cover of all zeros. Thus, no independent set can cover t , which is labelled with 0.

Suppose f is computed by a hypergraph program H of degree 2 with hyperedges e_1, \dots, e_k . We first provide a graph-theoretic characterisation of independent sets covering all zeros based on the implication graph [5] (or the chain criterion of Lemma 8.3.1 [9]). With any hyperedge e_i we associate a propositional variable p_{e_i} and with an input $\vec{\alpha}$ we associate the following set $\Phi_{\vec{\alpha}}$ of binary clauses:

- $\neg p_{e_i} \vee \neg p_{e_j}$, if $e_i \cap e_j \neq \emptyset$ (informally: intersecting hyperedges cannot be chosen at the same time),
- $p_{e_i} \vee p_{e_j}$, if there is $v \in e_i \cap e_j$ such that $\vec{\alpha}(v) = 0$ (informally: all zeros must be covered; note that all vertices have at most two incident edges).

By definition, X is an independent set covering all zeros iff $X = \{e_i \mid \vec{\beta}(p_{e_i}) = 1\}$, for some assignment $\vec{\beta}$ satisfying $\Phi_{\vec{\alpha}}$. Let $B_{\vec{\alpha}} = (V, E_{\vec{\alpha}})$ be a directed graph with

$$\begin{aligned} V &= \{e_i^+, e_i^- \mid 1 \leq i \leq k\}, \\ E_{\vec{\alpha}} &= \{(e_i^+, e_j^-) \mid e_i \cap e_j \neq \emptyset\} \cup \{(e_i^-, e_j^+) \mid v \in e_i \cap e_j \text{ and } \vec{\alpha}(v) = 0\}. \end{aligned}$$

(V is the set of all ‘literals’ for the variables of $\Phi_{\vec{\alpha}}$ and $E_{\vec{\alpha}}$ is the arcs for the implicational form of the clauses of $\Phi_{\vec{\alpha}}$; note that $\neg p_{e_i} \vee \neg p_{e_j}$ gives rise to two implications, $p_{e_i} \rightarrow \neg p_{e_j}$ and $p_{e_j} \rightarrow \neg p_{e_i}$, and so to two arcs in the graph). By Lemma 8.3.1 in [9], $\Phi_{\vec{\alpha}}$ is satisfiable iff there is no e_i with a (directed) cycle going through e_i^+ and e_i^- . It will be convenient for us to regard the $B_{\vec{\alpha}}$, for assignments $\vec{\alpha}$, as a single labelled directed graph B with arcs of the form (e_i^+, e_j^-) labelled with 1 and arcs of the form (e_i^-, e_j^+) labelled with $\neg v$, for $v \in e_i \cap e_j$. It should be clear that $B_{\vec{\alpha}}$ has a cycle going through e_i^+ and e_i^- iff $e_i^- \rightarrow_{\vec{\alpha}} e_i^+$ and $e_i^+ \rightarrow_{\vec{\alpha}} e_i^-$ in B .

The required NBP will contain two distinguished vertices, s and t , and, for each hyperedge e_i , two copies, B_i^+ and B_i^- , of B with arcs from s to the e_i^- vertex of B_i^+ , from the e_i^+ vertex of B_i^+ to the e_i^+ vertex of B_i^- and from the e_i^- vertex of B_i^- to t . This construction guarantees that $s \rightarrow_{\vec{\alpha}} t$ iff there is e_i such that $B_{\vec{\alpha}}$ contains a cycle going through e_i^+ and e_i^- . \square

In terms of expressive power, polynomial-size NBPs are a nonuniform analogue of the class NL; in symbols: $\text{NBP}(\text{poly}) = \text{NL}/\text{poly}$. Compared to other nonuniform computational models, (monotone) NBPs sit between (monotone) Boolean formulas and Boolean circuits [31]. As shown above, a (monotone) Boolean function f is computable by a polynomial-size (monotone) HGP of degree ≤ 2 iff its dual f^* is computable by a polynomial-size (monotone) NBP. (The problem whether f^* can be replaced with f is open; a negative solution would give a solution to the open problem 5 from [31].) Thus, (monotone) HGPs of degree ≤ 2 also sit between (monotone) Boolean formulas and Boolean circuits. However, (monotone) hypergraphs of degree ≤ 3 turn out to be much more powerful than (monotone) hypergraphs of degree ≤ 2 : we show now that polynomial-size (monotone) HGPs of degree ≤ 3 can compute NP-hard Boolean functions.

A function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ is computed by a *nondeterministic Boolean circuit* $C(\vec{x}, \vec{y})$, with $|\vec{x}| = n$, if for any $\vec{\alpha} \in \{0, 1\}^n$, we have $f(\vec{\alpha}) = 1$ iff there is $\vec{\beta} \in \{0, 1\}^m$ with $C(\vec{\alpha}, \vec{\beta}) = 1$. The variables in \vec{y} are called *advice variables*. We say that a nondeterministic circuit $C(\vec{x}, \vec{y})$ is *monotone* if the negations in C are only applied to variables in \vec{y} . Denote by $\text{NBC}(f)$ (respectively, $\text{NBC}_+(f)$) the minimal size of (monotone) nondeterministic Boolean circuits computing f .

Theorem 9. (i) For any Boolean function f , $\text{HGP}(f)$, $\text{HGP}^3(f)$ and $\text{NBC}(f)$ are polynomially related.
(ii) For any monotone Boolean function f , $\text{HGP}_+(f)$, $\text{HGP}_+^3(f)$ and $\text{NBC}_+(f)$ are polynomially related.

Proof. Clearly, $\text{HGP}(f) \leq \text{HGP}^3(f)$.

Now, given a (monotone) HGP of size m , we construct a (monotone) nondeterministic circuit $C(\vec{x}, \vec{y})$ of size $\text{poly}(m)$. Its \vec{x} -variables are the variables of the program, and its advice variables correspond to the edges of the program. The circuit C will return 1 on $(\vec{\alpha}, \vec{\beta})$ iff the family $\{e_i \mid \vec{\beta}(e_i) = 1\}$ of edges of the hypergraph forms an independent set covering all zeros under $\vec{\alpha}$. It is easy to construct a polynomial-size circuit checking this property. Indeed, for each pair of intersecting edges e_i, e_j , it is enough to take disjunction $\neg e_i \vee \neg e_j$, and for each vertex of the hypergraph labelled with p and adjacent to edges e_{i_1}, \dots, e_{i_k} to take disjunction $p \vee e_{i_1} \vee \dots \vee e_{i_k}$. (Note that applications of \neg to advice variables in the monotone case are allowed.) It then remains to take a conjunction of these disjunctions. Finally, it is easy to see that the resulting nondeterministic circuit is monotone if the hypergraph program is monotone.

Conversely, suppose f is computed by a nondeterministic circuit $C(\vec{x}, \vec{y})$. Let g_1, \dots, g_n be the nodes of C (including the inputs \vec{x} and \vec{y}). We construct an HGP of degree ≤ 3 computing f by taking, for each i , a vertex g_i labelled with 0 and a pair of hyperedges \bar{e}_{g_i} and e_{g_i} , both containing g_i . No other edge contains g_i , and so either \bar{e}_{g_i} or e_{g_i} should be present in any cover of zeros. (Intuitively, if the node g_i is positive then e_{g_i} belongs to the cover; otherwise, \bar{e}_{g_i} is there.) To ensure this property, for each input variable x_i , we add a vertex labelled with $\neg x_i$ to e_{x_i} and a fresh vertex labelled with x_i to \bar{e}_{x_i} . For each gate g_i , we consider three cases.

- If $g_i = \neg g_j$ then we add a vertex labelled with 1 to e_{g_i} and \bar{e}_{g_j} , and a vertex labelled with 1 to \bar{e}_{g_i} and e_{g_j} .
- If $g_i = g_j \vee g_{j'}$ then we add a vertex labelled with 1 to e_{g_j} and \bar{e}_{g_i} , add a vertex labelled with 1 to $e_{g_{j'}}$ and \bar{e}_{g_i} ; then, we add vertices h_j and $h_{j'}$ labelled with 1 to \bar{e}_{g_j} and $\bar{e}_{g_{j'}}$, respectively, and a vertex u_i labeled with 0 to \bar{e}_{g_i} ; finally, we add hyperedges $\{h_j, u_i\}$ and $\{h_{j'}, u_i\}$.
- If $g_i = g_j \wedge g_{j'}$ then we use the dual construction.

It is not hard to see that e_{g_i} is in the cover iff it contains \bar{e}_{g_j} in the first case, and e_{g_i} is in the cover iff it contains at least one of e_{g_j} and $e_{g_{j'}}$ in the second one. Indeed, in the second case if, say, the cover contains e_{g_j} then it cannot contain \bar{e}_{g_i} , and so it contains e_{g_i} . The vertex u_i in this case can be covered by the hyperedge $\{h_j, u_i\}$ since \bar{e}_{g_j} is not in the cover. Conversely, if neither e_{g_j} nor $e_{g_{j'}}$ is in the cover, then it must contain both \bar{e}_{g_j} and $\bar{e}_{g_{j'}}$, and so, neither $\{h_j, u_i\}$ nor $\{h_{j'}, u_i\}$ can belong to the cover and we will have to include \bar{e}_{g_i} to the cover. Finally, we add one more vertex labelled with 0 to e_g for the output gate g of C . By induction on the structure of C one can show that, for each $\vec{\alpha}$, there is $\vec{\beta}$ such that $C(\vec{\alpha}, \vec{\beta}) = 1$ iff the constructed HGP returns 1 on $\vec{\alpha}$.

If C is monotone, we remove all vertices labelled with $\neg x_i$. Then, for an input $\vec{\alpha}$, there is a cover of zeros in the resulting HGP iff there are $\vec{\beta}$ and $\vec{\alpha}' \leq \vec{\alpha}$ with $C(\vec{\alpha}', \vec{\beta}) = 1$. \square

Now, we use the developed machinery to investigate the size of rewritings over ontologies of depth 1 and 2.

5 Rewritings over Ontologies of Depth 1

Theorem 10. For any ontology \mathcal{T} of depth 1 and any CQ q , the hypergraph $H_{\mathcal{T}}^q$ is of degree ≤ 2 and $|\Theta_{\mathcal{T}}^q| \leq |q|$.

Proof. We have to show that every atom in q belongs to at most two q_t , $t \in \Theta_{\mathcal{T}}^q$. Suppose $t = (t_r, t_i)$ is a tree witness and $y \in t_i$. Since \mathcal{T} is of depth 1, $t_i = \{y\}$ and t_r consists of all those variables z in q for which $S(y, z) \in q$ or $S(z, y) \in q$, for some S . Thus, different tree witnesses have different internal variables y . An atom of the form $A(u) \in q$ is in q_t iff $u = y$. An atom of the form $P(u, v) \in q$ is in q_t iff either $u = y$ or $v = y$. Therefore, $P(u, v) \in q$ can only be covered by the tree witness with internal u and by the tree witness with internal v . \square

The canonical model $\mathcal{C}_{\mathcal{T}_H}^{A_{e_1}(a)}$ is shown on the right-hand side of the picture above. Note that each z_e determines the tree witness \mathfrak{t}^e with $\mathbf{q}_{\mathfrak{t}^e} = \{R_v(z_{i_1(v)}, z_{i_2(v)}) \mid v \in e\}$; \mathfrak{t}^e and $\mathfrak{t}^{e'}$ are conflicting iff $e \cap e' \neq \emptyset$. It follows that H is isomorphic to $H_{\mathcal{T}_H}^{\mathbf{q}_H}$. In fact, this example generalises to the following:

Theorem 13. *Any hypergraph H of degree 2 is isomorphic to $H_{\mathcal{T}_H}^{\mathbf{q}_H}$, with \mathcal{T}_H being an ontology of depth 1.*

Proof. We show that the map $h: v \mapsto R_v(z_{i_1(v)}, z_{i_2(v)})$ is an isomorphism between H and $H_{\mathcal{T}_H}^{\mathbf{q}_H}$. By the definition of \mathbf{q}_H , h is a bijection between V and the atoms of \mathbf{q}_H . For any $e \in E$, there is a tree witness $\mathfrak{t}^e = (\mathfrak{t}_i^e, \mathfrak{t}_e^e)$ generated by $A_e(x)$ with

$$\mathfrak{t}_i^e = \{z_e\} \quad \text{and} \quad \mathfrak{t}_e^e = \{z_{e'} \mid e' \cap e \neq \emptyset\},$$

and $\mathbf{q}_{\mathfrak{t}^e}$ consists of the $h(v)$, for $v \in e$. Conversely, every tree witness \mathfrak{t} for \mathbf{q}_H and \mathcal{T}_H contains $z_e \in \mathfrak{t}_i$, for some $e \in E$, and so $\mathbf{q}_{\mathfrak{t}} = \{h(v) \mid v \in e\}$. \square

We now show that answering \mathbf{q}_H over \mathcal{T}_H and certain single-individual data instances amounts to computing the Boolean function f_H . Let $H = (V, E)$ be a hypergraph of degree 2 with $V = \{v_1, \dots, v_n\}$ and $E = \{e_1, \dots, e_m\}$. We denote by $\vec{\alpha}(v_i)$ the i -th component of $\vec{\alpha} \in \{0, 1\}^n$, by $\vec{\beta}(e_j)$ the j -th component of $\vec{\beta} \in \{0, 1\}^m$, and set

$$\mathcal{A}_{\vec{\alpha}, \vec{\beta}} = \{R_{v_i}(a, a) \mid \vec{\alpha}(v_i) = 1\} \cup \{A_{e_j}(a) \mid \vec{\beta}(e_j) = 1\}.$$

Theorem 14. *Let $H = (V, E)$ be a hypergraph of degree 2. Then $\mathcal{T}_H, \mathcal{A}_{\vec{\alpha}, \vec{\beta}} \models \mathbf{q}_H$ iff $f_H(\vec{\alpha}, \vec{\beta}) = 1$, for any $\vec{\alpha} \in \{0, 1\}^{|V|}$ and $\vec{\beta} \in \{0, 1\}^{|E|}$.*

Proof. (\Leftarrow) Let X be an independent subset of E such that $\bigwedge_{v \in V \setminus V_X} p_v \wedge \bigwedge_{e \in X} p_e$ is true on $\vec{\alpha}$ (for the p_v) and $\vec{\beta}$ (for the p_e). Define $h: \mathbf{q}_H \rightarrow \mathcal{C}_{\mathcal{T}_H, \mathcal{A}_{\vec{\alpha}, \vec{\beta}}}$ by taking $h(z_e) = a$ if $e \notin X$ and $h(z_e) = w_e$ otherwise, where w_e is the labelled null in the canonical model $\mathcal{C}_{\mathcal{T}_H, \mathcal{A}_{\vec{\alpha}, \vec{\beta}}}$ introduced to witness the existential quantifier in (12). One can check that h is a homomorphism, and so $\mathcal{T}_H, \mathcal{A}_{\vec{\alpha}, \vec{\beta}} \models \mathbf{q}_H$.

(\Rightarrow) Suppose $h: \mathbf{q}_H \rightarrow \mathcal{C}_{\mathcal{T}_H, \mathcal{A}_{\vec{\alpha}, \vec{\beta}}}$ is a homomorphism. We show that the set $X = \{e \in E \mid h(z_e) \neq a\}$ is independent. Indeed, if $e, e' \in X$ and $v \in e \cap e'$, then h sends one variable of the R_v -atom to the labelled null w_e and the other end to $w_{e'}$, which is impossible. We claim that $f_H(\vec{\alpha}, \vec{\beta}) = 1$. Indeed, for each $v \in V \setminus V_X$, h sends both ends of the R_v -atom to a , and so $\vec{\alpha}(v) = 1$. For each $e \in X$, we must have $h(z_e) = w_e$ because $h(z_e) \neq a$, and so $\vec{\beta}(e) = 1$. It follows that $f_H(\vec{\alpha}, \vec{\beta}) = 1$. \square

We are now fully equipped to show that there exist CQs and ontologies of depth 1 without polynomial-size PE-rewritings:

Theorem 15. *There is a sequence of CQs \mathbf{q}_n and ontologies \mathcal{T}_n of depth 1, both of polynomial size in n , such that any PE-rewriting of \mathbf{q}_n and \mathcal{T}_n is of size $n^{\Omega(\log n)}$.*

Proof. As shown in [19], there is a sequence f_n of monotone Boolean functions that are computable by polynomial-size monotone NBPs, but any monotone Boolean formulas computing f_n are of size $n^{\Omega(\log n)}$. In fact, f_n from [19] checks whether two given vertices are connected by a path in a given undirected graph (alternatively, one could use the functions from [17]).

By Theorem 8 (ii) and Lemma 7, there is a sequence of polynomial-size monotone HGPs H'_n of degree 2 computing f_n^* . By applying Theorem 13 to the hypergraph H_n of H'_n , we obtain a sequence of CQs \mathbf{q}_n and ontologies \mathcal{T}_n of depth 1 such that H_n is isomorphic to $H_{\mathcal{T}_n}^{\mathbf{q}_n}$. We show now that any PE-rewriting \mathbf{q}'_n of \mathbf{q}_n and \mathcal{T}_n can be transformed to a monotone Boolean formula computing f_n and having size $\leq |\mathbf{q}'_n|$.

To define such a formula, we eliminate the quantifiers in \mathbf{q}'_n in the following way: take a constant a and replace every subformula of the form $\exists x \psi(x)$ in \mathbf{q}'_n with $\psi(a)$, repeating this operation as many times as possible. The resulting formula \mathbf{q}''_n is built from atoms of the form $A_e(a)$, $R_v(a, a)$ and $S_e(a, a)$ using \wedge and \vee . For every data instance \mathcal{A} with a single individual a , we have $\mathcal{T}_n, \mathcal{A} \models \mathbf{q}_n$ iff $\mathcal{A} \models \mathbf{q}''_n$. Let χ_n be

the result of replacing $S_e(a, a)$ in q_n'' with \perp , $A_e(a)$ with p_e and $R_v(a, a)$ with p_v . Clearly, $|\chi_n| \leq |q_n'|$. By the definition of $\mathcal{A}_{\vec{\alpha}, \vec{\beta}}$ and Theorem 14, we obtain:

$$\chi_n(\vec{\alpha}, \vec{\beta}) = 1 \quad \text{iff} \quad \mathcal{A}_{\vec{\alpha}, \vec{\beta}} \models q_n'' \quad \text{iff} \quad \mathcal{T}_n, \mathcal{A}_{\vec{\alpha}, \vec{\beta}} \models q_n \quad \text{iff} \quad f_{H_n}(\vec{\alpha}, \vec{\beta}) = 1.$$

As H_n' computes f_n^* , we can obtain f_n^* from f_{H_n} by replacing each p_e with 1 and each p_v with the label of v in H_n' . The same substitution in χ_n (with \top and \perp in place of 1 and 0) gives a monotone formula that computes f_n^* . By swapping \vee and \wedge in it, we obtain a monotone formula χ_n' computing f_n . It remains to recall that $|q_n'| \geq |\chi_n'| = n^{\Omega(\log n)}$. \square

It may be of interest to note that the function f_n in the proof above is in the complexity class L. The algorithm computing f_n by querying the NDL-rewriting of Theorem 11 over single-individual data instances runs in polynomial time; the algorithm querying any PE-rewriting to compute f_n requires, by Theorem 15, superpolynomial time.

We note further that instead of reachability in undirected graphs in Theorem 15 we could use reachability in directed graphs. Indeed, since the undirected case reduces to the directed one, we have the same lower bound for computing directed reachability by monotone formulas. On the other hand, it is known that directed reachability also can be computed by polynomial-size monotone circuits. As reachability in directed graphs is NL/poly-complete under NC^1 -reductions, the argument in the proof of Theorem 15 shows that the existence of short FO-rewritings of CQs and ontologies of depth 1 is equivalent to a well-known open problem in computational complexity:

Theorem 16. *There exist polynomial-size FO-rewritings for all CQs and ontologies of depth 1 iff all functions in NL/poly are computed by polynomial-size Boolean formulas, that is, iff $\text{NL/poly} \subseteq \text{NC}^1$.*

Proof. (\Leftarrow) Suppose $\text{NL/poly} \subseteq \text{NC}^1$. Take an arbitrary CQ and an ontology of depth 1. By Theorem 10, its hypergraph H_T^q is of degree ≤ 2 and polynomial size. By Lemma 6, there is a polynomial-size HGP H computing $f_{H_T^q}$, whence, by Theorem 8 (i), there is a polynomial-size NBP computing $\neg f_{H_T^q}$, and so $f_{H_T^q}$ is in $\text{CONL/poly} = \text{NL/poly}$. Therefore, by our assumption, it can be computed by a polynomial-size Boolean formula. By Theorem 5, the latter translates into a polynomial-size FO-rewriting of q and \mathcal{T} .

(\Rightarrow) Suppose that there exist polynomial-size FO-rewritings for all CQs and ontologies of depth 1. Consider a sequence of functions f_n that compute the connectivity function in *directed* graphs. Since $f_n \in \text{NL}$ and $\text{NL} = \text{CONL}$, the functions $\neg f_n$ are computable by a sequence of polynomial-size NBPs. Now we use an argument similar to the one in the proof of Theorem 15. We apply Theorem 8 (i) and Lemma 7 to $\neg f_n$ and obtain a polynomial-size HGP H_n' of degree 2 that computes f_n . By Theorem 13, there are sequences of CQs q_n and ontologies \mathcal{T}_n of depth 1 such that f_n is a subfunction of $f_{H_n'^{q_n}}$ in the sense that f_n is the result of replacing each p_e with 1 and each p_v with the label of v in H_n' . By our assumption, there is a polynomial-size FO-rewriting q_n' of q_n and \mathcal{T}_n . We eliminate the quantifiers in q_n' and apply to the result the substitution giving f_n from $f_{H_n'^{q_n}}$ to obtain a polynomial-size propositional Boolean formula that computes f_n . Since f_n is NL/poly-complete under NC^1 reductions, we then must have $\text{NL/poly} \subseteq \text{NC}^1$. \square

As we shall see in Section 7, *tree-shaped* CQs and ontologies of depth 1 always have polynomial-size PE-rewritings.

6 Rewritings over Ontologies of Depth 2

Our next aim is to show that CQs and ontologies of depth 2 can compute the NP-complete function checking whether a graph with n vertices has a k -clique. We remind the reader (see, e.g., [3] for details) that the monotone Boolean function $\text{CLIQUE}_{n,k}(\vec{e})$ of $n(n-1)/2$ variables $e_{jj'}$, $1 \leq j < j' \leq n$, returns 1 iff the graph with vertices $\{1, \dots, n\}$ and edges $\{\{j, j'\} \mid e_{jj'} = 1\}$ contains a k -clique. A series of papers, started by Razborov's [30], gave an exponential lower bound for the size of monotone circuits computing $\text{CLIQUE}_{n,k}$: $2^{\Omega(\sqrt{k})}$ for $k \leq \frac{1}{4}(n/\log n)^{2/3}$ [2]. For monotone formulas, an even better lower bound is known: $2^{\Omega(k)}$ for $k = 2n/3$ [29].

We first construct a monotone HGP computing $\text{CLIQUE}_{n,k}$ and then use the intuition behind the construction to encode $\text{CLIQUE}_{n,k}$ by means of a Boolean CQ $q_{n,k}$ and an ontology $\mathcal{T}_{n,k}$ of depth 2 and polynomial size. As a consequence, any PE- or NDL-rewriting of $q_{n,k}$ and $\mathcal{T}_{n,k}$ is of exponential size, while any FO-rewriting is of superpolynomial size unless $\text{NP} \subseteq \text{P/poly}$.

Given n and k , let $H_{n,k}$ be a monotone HGP with vertices

$$\begin{aligned} w_{jj'} &\text{ labelled with } e_{jj'}, & (1 \leq j < j' \leq n), \\ u_{jj'} \text{ and } u_{j'j} &\text{ labelled with } 1, & (1 \leq j < j' \leq n), \\ v_i &\text{ labelled with } 0 & (1 \leq i \leq k), \end{aligned}$$

and hyperedges

$$\begin{aligned} h^{jj'} &= \{w_{jj'}, u_{jj'}\} \quad \text{and} \quad h^{j'j} = \{w_{jj'}, u_{j'j}\} & (1 \leq j < j' \leq n), \\ f^{ij} &= \{v_i\} \cup \{u_{jj'} \mid j' \neq j\} & (1 \leq i \leq k, 1 \leq j \leq n). \end{aligned}$$

Informally, the $w_{jj'}$ represent the edges of the complete graph with n vertices; they can be turned ‘on’ or ‘off’ by means of the variables $e_{jj'}$. The vertex $u_{jj'}$ together with the hyperedge $h^{jj'}$ represent the ‘half’ of the edge connecting j and j' that is adjacent to j ; the other ‘half’ is represented by $u_{j'j}$ and $h^{j'j}$. The vertices v_i represent a k -clique and the edge f^{ij} corresponds to the choice of the vertex j of the graph as the i th element of the clique. The hypergraph of $H_{4,2}$ is shown in Fig. 4.




Figure 4: The hypergraph of $H_{4,2}$.

Theorem 17. *The HGP $H_{n,k}$ computes $\text{CLIQUE}_{n,k}$.*

Proof. We show that, for each $\vec{e} \in \{0, 1\}^{n(n-1)/2}$, there is an independent set X of hyperedges covering all zeros in $H_{n,k}$ under \vec{e} iff $\text{CLIQUE}_{n,k}(\vec{e}) = 1$.

(\Leftarrow) Let $\lambda: \{1, \dots, k\} \rightarrow \{1, \dots, n\}$ be such that $C = \{\lambda(i) \mid 1 \leq i \leq k\}$ is a k -clique in the graph G given by \vec{e} . Then

$$X = \{f^{i\lambda(i)} \mid 1 \leq i \leq k\} \cup \{h^{jj'} \mid j \notin C, j' \in C\} \cup \{h^{jj'} \mid j, j' \notin C \text{ and } j < j'\}$$

is independent and covers all zeros in $H_{n,k}$ under \vec{e} . Indeed, X is independent because, in every $h^{jj'} \in X$, the index j does not belong to C . By definition, each $f^{i\lambda(i)}$ covers v_i , for $1 \leq i \leq k$. Thus, it remains to show that any $w_{jj'}$ with $e_{jj'} = 0$ (that is, the edge $\{j, j'\}$ belongs to the complement of G) is covered by some hyperedge. All edges of the complement of G can be divided into two groups: those that are adjacent to C , and those that are not. The $w_{jj'}$ that correspond to the edges of the former group are covered by the

$h^{jj'}$ from the middle disjunct of X , where j corresponds to the end of the edge $\{j, j'\}$ that is not C . To cover $w_{jj'}$ of the latter group, take $h^{jj'}$ from the last disjunct of X .

(\Rightarrow) Suppose X is an independent set covering all zeros labelling the vertices of $H_{n,k}$, for an input \vec{e} . The vertex v_i , $1 \leq i \leq k$, is labelled with 0, and so there is $\lambda(i)$ such that $f^{i\lambda(i)} \in X$. We claim that $C = \{\lambda(i) \mid 1 \leq i \leq k\}$ is a k -clique in the graph given by \vec{e} . Indeed, suppose that the graph has no edge between some vertices $j, j' \in C$, that is, $e_{jj'} = 0$ for $j < j'$. Since $w_{jj'}$ is labelled with 0, it must be covered by a hyperedge in X , which can only be either $h^{jj'}$ or $h^{j'j}$ (see the picture above). But $h^{jj'}$ intersects $f^{\lambda^{-1}(j)j}$ and $h^{j'j}$ intersects $f^{\lambda^{-1}(j')j'}$, which is a contradiction. \square




Figure 5: The CQ $q_{4,2}$ for $H_{4,2}$.

We are now in a position to define $\mathcal{T}_{n,k}$ of depth 2 and $q_{n,k}$, both of polynomial size in n , that can compute $\text{CLIQUE}_{n,k}$. Let $q_{n,k}$ contain the following atoms (all variables are quantified):

$$\begin{aligned} T_{ij}(v_i, z_{ij}) & & (1 \leq i \leq k, 1 \leq j \leq n), \\ P_{jj'}(w_{jj'}, x_{jj'}), P_{j'j}(w_{jj'}, x_{j'j}) & & (1 \leq j < j' \leq n), \\ Q(u_{jj'}, x_{jj'}), U(u_{jj'}, z_{ij}) & & (1 \leq j \neq j' \leq n, 1 \leq i \leq k). \end{aligned}$$

Figs. 5 and 6 show two different views of the CQ $q_{4,2}$ for $H_{4,2}$. Fig. 7 illustrates the fragments of $q_{n,k}$ centred in each variable of the form z_{ij} and $x_{j'j}$ (the fragment centred in $x_{j'j}$ is similar to that of $x_{jj'}$ except the index of the $w_{jj'}$).

The ontology $\mathcal{T}_{n,k}$ mimics the arrangement of atoms in the layers depicted in Fig. 7 and contains the following tgds, where $1 \leq i \leq k$ and $1 \leq j \neq j' \leq n$,

$$\begin{aligned} A_{ij}(x) &\rightarrow \exists y \left[\bigwedge_{j'' \neq j} T_{ij''}(y, x) \wedge U(y, x) \wedge Q(y, x) \wedge A'_{ij}(y) \right], \\ A'_{ij}(x) &\rightarrow \exists y [T_{ij}(x, y) \wedge U(x, y)], \\ B_{jj'}(x) &\rightarrow \exists y [P_{j'j}(y, x) \wedge U(y, x) \wedge B'_{j'j}(y)], \\ B'_{j'j}(x) &\rightarrow \exists y [P_{jj'}(x, y) \wedge Q(x, y)]. \end{aligned}$$




Figure 6: The CQ $q_{4,2}$ for $H_{4,2}$.




Figure 7: Fragments of $q_{n,k}$ and the canonical models $C_{T_{n,k}}^{A_{ij}(a)}$ and $C_{T_{n,k}}^{B_{jj'}(a)}$.

The canonical models $C_{T_{n,k}}^{A_{ij}(a)}$ and $C_{T_{n,k}}^{B_{jj'}(a)}$ are also illustrated in Fig. 7 with the horizontal dashed lines showing possible ways of embedding the fragments of $q_{n,k}$ into them. These embeddings give rise to the following tree witnesses:

– $\mathfrak{t}^{ij} = (\mathfrak{t}_r^{ij}, \mathfrak{t}_i^{ij})$ generated by $A_{ij}(x)$, for $1 \leq i \leq k$ and $1 \leq j \leq n$, where

$$\begin{aligned} \mathfrak{t}_r^{ij} &= \{z_{ij'}, x_{jj'} \mid 1 \leq j' \leq n, j' \neq j\} \cup \{z_{i'j} \mid 1 \leq i' \leq k, i' \neq i\}, \\ \mathfrak{t}_i^{ij} &= \{v_i, z_{ij}\} \cup \{u_{jj'} \mid 1 \leq j' \leq n, j' \neq j\}; \end{aligned}$$

– $\mathfrak{s}^{jj'} = (\mathfrak{s}_r^{jj'}, \mathfrak{s}_i^{jj'})$ and $\mathfrak{s}^{j'j} = (\mathfrak{s}_r^{j'j}, \mathfrak{s}_i^{j'j})$, generated by $B_{jj'}(x)$ and $B_{j'j}(x)$, respectively, for $1 \leq j < j' \leq n$, where

$$\begin{aligned} \mathfrak{s}_r^{jj'} &= \{x_{j'j}\} \cup \{z_{ij} \mid 1 \leq i \leq k\}, & \mathfrak{s}_i^{jj'} &= \{w_{jj'}, u_{jj'}, x_{jj'}\}, \\ \mathfrak{s}_r^{j'j} &= \{x_{jj'}\} \cup \{z_{ij'} \mid 1 \leq i \leq k\}, & \mathfrak{s}_i^{j'j} &= \{w_{jj'}, u_{j'j}, x_{j'j}\}. \end{aligned}$$

The tree witnesses \mathfrak{t}^{ij} , $\mathfrak{s}^{jj'}$ and $\mathfrak{s}^{j'j}$ are uniquely determined by their most remote (from the root) variables, z_{ij} , $x_{jj'}$ and $x_{j'j}$, respectively, and correspond to the hyperedges f^{ij} , $h^{jj'}$, $h^{j'j}$ of $H_{n,k}$; their internal variables of the form v_i , $w_{jj'}$ and $u_{jj'}$ correspond to the vertices in the respective hyperedge (see Fig. 5).

Given a vector \vec{e} representing a graph with n vertices, we construct a data instance $\mathcal{A}_{\vec{e}}$ with a single individual a by taking the following atoms:

$$Q(a, a), \quad U(a, a), \quad A_{ij}(a), \text{ for } 1 \leq i \leq k \text{ and } 1 \leq j \leq n, \\ P_{jj'}(a, a) \text{ and } P_{j'j}(a, a), \text{ for } 1 \leq j < j' \leq n \text{ with } e_{jj'} = 1.$$

Lemma 18. $\mathcal{T}_{n,k}, \mathcal{A}_{\vec{e}} \models \mathbf{q}_{n,k}$ iff $\text{CLIQUE}_{n,k}(\vec{e}) = 1$.

Proof. (\Rightarrow) Suppose $\mathcal{T}_{n,k}, \mathcal{A}_{\vec{e}} \models \mathbf{q}_{n,k}$. Then there is a homomorphism g from $\mathbf{q}_{n,k}$ to the canonical model \mathcal{C} of $(\mathcal{T}_{n,k}, \mathcal{A}_{\vec{e}})$. Since the only points of \mathcal{C} that belong to $\exists y T_{ij}(x, y)$ are of the form c_{ij} (see Fig. 7) and $\mathbf{q}_{n,k}$ contains atoms of the form $T_{ij}(v_i, z_{ij})$, there is $\lambda: \{1, \dots, k\} \rightarrow \{1, \dots, n\}$ such that $g(v_i) = c_{i\lambda(i)}$. We claim that $C = \{\lambda(i) \mid 1 \leq i \leq k\}$ is a k -clique in the graph given by \vec{e} .

We first show that λ is injective. Suppose to the contrary that $\lambda(i) = \lambda(i') = j$, for $i \neq i'$. Since $\mathbf{q}_{n,k}$ contains $T_{ij}(v_i, z_{ij})$ and $T_{i'j}(v_{i'}, z_{i'j})$, we have $g(z_{ij}) = c'_{ij}$ and $g(z_{i'j}) = c'_{i'j}$. Take $j' \neq j$. Since $U(u_{jj'}, z_{ij}), U(u_{jj'}, z_{i'j}) \in \mathbf{q}_{n,k}$, we obtain $g(u_{jj'}) = c_{ij}$ and $g(u_{jj'}) = c_{i'j}$, contrary to $i \neq i'$.

Next, we show that $e_{jj'} = 1$, for all $j, j' \in C$ with $j < j'$. Since $U(u_{jj'}, z_{ij})$ is in $\mathbf{q}_{n,k}$, we have $g(u_{jj'}) = c_{ij}$, and so $g(x_{jj'}) = a$. Similarly, we also have $g(u_{j'j}) = c_{i'j'}$ and $g(x_{j'j}) = a$. Then, since $\mathbf{q}_{n,k}$ contains both $P_{jj'}(w_{jj'}, x_{jj'})$ and $P_{j'j}(w_{j'j}, x_{j'j})$ and \mathcal{C} contains no pair of points in both $P_{jj'}$ and $P_{j'j}$ apart from (a, a) , we obtain $e_{jj'} = 1$ whenever $g(x_{jj'}) = g(x_{j'j}) = a$, as shown in Fig. 8.




Figure 8: Proof of Lemma 18.

(\Leftarrow) Suppose $\lambda: \{1, \dots, k\} \rightarrow \{1, \dots, n\}$ is a k -clique. We construct a homomorphism g from $\mathbf{q}_{n,k}$ to the canonical model of $(\mathcal{T}_{n,k}, \mathcal{A}_{\vec{e}})$ by taking (see Fig. 7), for $1 \leq i \leq k$ and $1 \leq j < j' \leq n$,

$$g(v_i) = c_{i\lambda(i)}, \\ g(z_{ij}) = \begin{cases} c'_{ij}, & \text{if } j = \lambda(i), \\ a & \text{otherwise,} \end{cases} \quad g(w_{jj'}) = \begin{cases} a, & \text{if } j, j' \in C, \\ d'_{jj'}, & \text{if } j' \notin C \text{ and } j \in C, \\ d_{jj'}, & \text{otherwise,} \end{cases}$$

and, for $1 \leq j \neq j' \leq n$,

$$g(u_{jj'}) = \begin{cases} c_{\lambda^{-1}(j)j}, & \text{if } j \in C, \\ d_{jj'}, & \text{if } j \notin C, j' \in C, \\ d'_{jj'}, & \text{if } j, j' \notin C, j < j', \\ a, & \text{if } j, j' \notin C, j' < j, \end{cases} \quad g(x_{jj'}) = \begin{cases} a, & \text{if } j \in C, \\ d'_{jj'}, & \text{if } j \notin C, j' \in C, \\ d_{jj'}, & \text{if } j, j' \notin C, j < j', \\ a, & \text{if } j, j' \notin C, j' < j. \end{cases}$$

This homomorphism mimics the cover X constructed for $H_{n,k}$ in the proof of Theorem 17. The internal variables of the tree witnesses from X are sent to labelled nulls, and all other points are sent to a . For example, in the definition of $g(u_{jj'})$, the first case corresponds to $u_{jj'} \in f^{\lambda^{-1}(j)j} \in X$; the second and third cases to $u_{jj'} \in h^{jj'} \in X$; and in the fourth case, $u_{jj'}$ is not covered by X . It follows that $\mathcal{T}_{n,k}, \mathcal{A}_{\vec{e}} \models \mathbf{q}_{n,k}$. \square

Theorem 19. *There exists a sequence of CQs \mathbf{q}_n and ontologies \mathcal{T}_n of depth 2 any PE- and NDL-rewritings of which are of exponential size, while any FO-rewriting is of superpolynomial size unless $\text{NP} \subseteq \text{P/poly}$.*

Proof. Given a PE-, FO- or NDL-rewriting $q'_{n,k}$ of $q_{n,k}$ and $\mathcal{T}_{n,k}$, we show how to construct, respectively, a monotone Boolean formula, a Boolean formula or a monotone Boolean circuit for the function $\text{CLIQUE}_{n,k}$ of size $|q'_{n,k}|$.

Suppose $q'_{n,k}$ is a PE-rewriting of $q_{n,k}$ and $\mathcal{T}_{n,k}$. We eliminate the quantifiers in $q'_{n,k}$ by replacing first every subformula of the form $\exists x \psi(x)$ in $q'_{n,k}$ with $\psi(a)$, and then replacing each $P_{jj'}(a, a)$ and $P_{j'j}(a, a)$ with $e_{jj'}$, each $T_{ij}(a, a)$, $A'_{ij}(a)$ and $B'_{jj'}(a)$ with 0 and each $U(a, a)$, $Q(a, a)$, $A_{ij}(a)$ and $B_{jj'}(a)$ with 1. One can check that the resulting propositional monotone Boolean formula computes $\text{CLIQUE}_{n,k}$.

If $q'_{n,k}$ is an FO-rewriting of $q_{n,k}$, then we eliminate the quantifiers by replacing $\exists x \psi(x)$ and $\forall x \psi(x)$ in $q'_{n,k}$ with $\psi(a)$, and then carry out the replacing procedure above, obtaining a propositional Boolean formula computing $\text{CLIQUE}_{n,k}$.

If $(\Pi, q'_{n,k})$ is an NDL-rewriting of $q_{n,k}$, we replace all the individual variables in Π with a and then perform the replacement described above. Denote the resulting propositional NDL-program by Π' . The program Π' can now be transformed into a monotone Boolean circuit computing $\text{CLIQUE}_{n,k}$: for every (propositional) variable p occurring in the head of a clause in Π' , we introduce an \vee -gate whose output is p and inputs are the bodies of the clauses with the head p ; and for each such body, we introduce an \wedge -gate whose inputs are the propositional variables in the body.

Now Theorem 19 follows from the lower bounds for monotone Boolean circuits and formulas computing $\text{CLIQUE}_{n,k}$ given at the beginning of this section. \square

As the function $\text{CLIQUE}_{n,k}$ is known to be NP/poly-complete with respect to NC^1 -reductions, we also obtain:

Theorem 20. *There exist polynomial-size FO-rewritings for all CQs and ontologies of depth 2 with polynomially-many tree witnesses iff all functions in NP/poly are computed by polynomial-size formulas, that is, iff $\text{NP/poly} \subseteq \text{NC}^1$.*

Proof. (\Leftarrow) Suppose $\text{NP/poly} \subseteq \text{NC}^1$. Consider an arbitrary CQ q and an ontology \mathcal{T} of depth 2 with polynomially-many tree witnesses. Then the hypergraph $H_{\mathcal{T}}^q$ is of polynomial size. The hypergraph function $f_{H_{\mathcal{T}}^q}$ is in the class NP/poly because the problem whether there exists an independent set of hyperedges in a hypergraph covering all zeros is in NP. Therefore, by our assumption, $f_{H_{\mathcal{T}}^q}$ can be computed by a polynomial-size formula, which translates into a polynomial-size FO-rewriting by Theorem 5.

(\Rightarrow) Conversely, suppose that there is a polynomial-size FO-rewriting for all CQs and ontologies of depth 2. In particular, there is a polynomial-size FO-rewriting for the CQs and ontologies of depth 2 encoding $\text{CLIQUE}_{n,k}$ defined above. These CQs and ontologies have polynomially-many tree witnesses. Our assumption and the construction in the proof of Theorem 19 provide us with a polynomial-size Boolean formula computing $\text{CLIQUE}_{n,k}$. Since $\text{CLIQUE}_{n,k}$ is NP/poly-complete under NC^1 reductions, this gives us $\text{NP/poly} \subseteq \text{NC}^1$. \square

7 Rewritings of Tree-Shaped CQs

A CQ is said to be *tree-shaped* if its Gaifman graph is a tree. It is well known [35, 12] that tree-shaped CQs (or, more generally, CQs of bounded treewidth) can be evaluated over plain data instances in polynomial time. In contrast, the evaluation of arbitrary CQs is NP-complete for combined complexity and $W[1]$ -complete for parameterised complexity. In this section, we consider tree-shaped CQs over ontologies.

At first sight, we do not gain much by focusing on tree-shaped CQs: answering such CQs over ontologies is NP-complete for combined complexity [21], while their PE- and NDL-rewritings can suffer an exponential blowup [20]. However, by examining the tree-witness rewriting (4), we see that the tw_t formula (3) defines a predicate over the data that can be computed in linear time. It follows that, for a tree-shaped q , every disjunct of (4) can also be regarded as a tree-shaped CQ of size $\leq |q|$. So, bearing in mind that $|\Theta_{\mathcal{T}}^q| \leq 3^{|q|}$, we obtain the following:

Theorem 21. *Given a tree-shaped CQ $q(\vec{x})$, an ontology \mathcal{T} , a data instance \mathcal{A} and a tuple $\vec{a} \subseteq \text{ind}(\mathcal{A})$, the problem of deciding whether $\mathcal{T}, \mathcal{A} \models q(\vec{a})$ is fixed-parameter tractable, with parameter $|q|$.*

Furthermore, if each variable in a tree-shaped CQ is covered by a ‘small’ number of tree witnesses then we can obtain polynomial-size PE- or NDL-rewritings.

Example 22. Consider the following ontology and CQ illustrated in Fig. 9:

$$\begin{aligned} \mathcal{T} &= \{ A_i(x) \rightarrow \exists y (R_i(x, y) \wedge R_{i+1}(y, x)) \mid 1 \leq i \leq 3 \}, \\ \mathbf{q} &= \exists y_1 \dots y_5 \bigwedge_{1 \leq i \leq 4} R_i(y_i, y_{i+1}). \end{aligned}$$




Figure 9: CQ and ontology from Example 22.

We construct a PE-rewriting \mathbf{q}^\dagger of \mathbf{q} and \mathcal{T} recursively by splitting \mathbf{q} into smaller subqueries. Suppose $\mathcal{T}, \mathcal{A} \models \mathbf{q}$, for some \mathcal{A} . Then there is a homomorphism $h: \mathbf{q} \rightarrow \mathcal{C}_{\mathcal{T}, \mathcal{A}}$. Consider the ‘central’ variable y_3 dividing \mathbf{q} in half. If $h(y_3)$ is in the data part of $\mathcal{C}_{\mathcal{T}, \mathcal{A}}$ then y_3 behaves like a free variable in \mathbf{q} . Since \mathbf{q} is tree-shaped, we can then proceed by constructing PE-rewritings, $\mathbf{q}_1^\dagger(y_3)$ and $\mathbf{q}_2^\dagger(y_3)$, for the subqueries

$$\begin{aligned} \mathbf{q}_1(y_3) &= \exists y_1 y_2 (R_1(y_1, y_2) \wedge R_2(y_2, y_3)), \\ \mathbf{q}_2(y_3) &= \exists y_4 y_5 (R_3(y_3, y_4) \wedge R_4(y_4, y_5)). \end{aligned}$$

If $h(y_3)$ is a labelled null, then y_3 must be an internal point of some tree witness for \mathbf{q} and \mathcal{T} . We have only one such tree witness, $\mathbf{t} = (\mathbf{t}_r, \mathbf{t}_i)$, generated by $A_2(x)$ with $\mathbf{t}_r = \{y_2, y_4\}$, $\mathbf{t}_i = \{y_3\}$ and $\mathbf{q}_t = \{R_2(y_2, y_3), R_3(y_3, y_4)\}$ (shaded in Fig. 9). But then $h(y_2) = h(y_4)$ and this element is in the data part of $\mathcal{C}_{\mathcal{T}, \mathcal{A}}$. So, we need PE-rewritings, $\mathbf{q}_3^\dagger(y_2)$ and $\mathbf{q}_4^\dagger(y_4)$, of the remaining fragments of \mathbf{q} :

$$\mathbf{q}_3(y_2) = \exists y_1 R_1(y_1, y_2), \quad \mathbf{q}_4(y_4) = \exists y_5 R_4(y_4, y_5).$$

If the required rewritings \mathbf{q}_i^\dagger , $1 \leq i \leq 4$, are constructed then we obtain a PE-rewriting \mathbf{q}^\dagger of \mathbf{q} and \mathcal{T} :

$$\mathbf{q}^\dagger = \exists y_3 (\mathbf{q}_1^\dagger(y_3) \wedge \mathbf{q}_2^\dagger(y_3)) \vee \exists y_2 y_4 (A_2(y_2) \wedge (y_2 = y_4) \wedge \mathbf{q}_3^\dagger(y_2) \wedge \mathbf{q}_4^\dagger(y_4)).$$

We analyse $\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3$ and \mathbf{q}_4 in the same way and obtain

$$\begin{aligned} \mathbf{q}_1^\dagger(y_3) &= \exists y_2 (\mathbf{q}_3^\dagger(y_2) \wedge R_2(y_2, y_3)) \vee \exists y_1 (A_1(y_3) \wedge (y_1 = y_3)), \\ \mathbf{q}_2^\dagger(y_3) &= \exists y_4 (R_3(y_3, y_4) \wedge \mathbf{q}_4^\dagger(y_4)) \vee \exists y_5 (A_3(y_3) \wedge (y_5 = y_3)), \end{aligned}$$

$\mathbf{q}_3^\dagger(y_2)$ and $\mathbf{q}_4^\dagger(y_4)$ equal to $\mathbf{q}_3(y_2)$ and $\mathbf{q}_4(y_4)$, respectively.

We now give a general definition of a PE-rewriting obtained by the strategy ‘divide and rewrite’ and applicable to any (not necessarily tree-shaped) CQ. Let $\mathbf{q}(\vec{x}) = \exists \vec{y} \varphi(\vec{x}, \vec{y})$ and an ontology \mathcal{T} be given. We recursively define a PE-query $\mathbf{q}^\dagger(\vec{x})$ as follows. Take the finest partition of $\exists \vec{y} \varphi(\vec{x}, \vec{y})$ into a conjunction $\bigwedge_j \exists \vec{y}_j \varphi_j(\vec{x}, \vec{y}_j)$ such that every atom containing some $y \in \vec{y}_j$ belongs to the same conjunct $\varphi_j(\vec{x}, \vec{y}_j)$. (Informally, the Gaifman graph of φ is cut along the answer variables \vec{x} .) By definition, the set of tree witnesses for $\exists \vec{y} \varphi(\vec{x}, \vec{y})$ and \mathcal{T} is the disjoint union of the sets of tree witnesses for the $\exists \vec{y}_j \varphi_j(\vec{x}, \vec{y}_j)$ and \mathcal{T} . Then we set $(\exists \vec{y} \varphi(\vec{x}, \vec{y}))^\dagger = \bigwedge_j \psi_j$, where ψ_j is $\varphi_j(\vec{x})$ in case \vec{y}_j is empty; otherwise, we choose a variable z in \vec{y}_j and define ψ_j to be the formula

$$\exists z (\exists [\vec{y}_j \setminus \{z\}] \varphi_j(\vec{x}, \vec{y}_j))^\dagger \vee \bigvee_{\substack{\mathbf{t} \text{ a tree witness for } \exists \vec{y}_j \varphi_j(\vec{x}, \vec{y}_j) \text{ and } \mathcal{T} \\ \text{such that } \mathbf{t} = (\mathbf{t}_r, \mathbf{t}_i) \text{ and } z \in \mathbf{t}_i}} \exists \vec{y}_{j, \mathbf{t}} ((\exists [\vec{y}_j \setminus \vec{y}_{j, \mathbf{t}}] \varphi_{j, \mathbf{t}}(\vec{x}, \vec{y}_j))^\dagger \wedge \text{tw}_{\mathbf{t}}(\mathbf{t}_r)),$$

where $\vec{y}_{j,t} = \vec{y}_j \cap t_r$ contains the variables in \vec{y}_j that occur among t_r , the quantifiers $\exists[\vec{y}_j \setminus \{z\}]$ and $\exists[\vec{y}_j \setminus \vec{y}_{j,t}]$ contain all variables in \vec{y}_j but z and $\vec{y}_{j,t}$, respectively, and $\varphi_{j,t}$ consists of all the atoms of φ_j except those in q_t . Note that the variables in t_i (in particular, z) do not occur in the disjunct for t (and so can be removed from the respective quantifier). Intuitively, the first disjunct represents the situation where z is mapped to a data individual and treated as a free variable in the rewriting of φ_j . The other disjuncts reflect the cases where z is mapped to a labelled null, and so z is an internal variable of a tree witness $t = (t_r, t_i)$ for $\exists\vec{y}_j \varphi_j(\vec{x}, \vec{y}_j)$ and \mathcal{T} . As the variables in t_r must be mapped to data individuals, this only leaves the set of atoms $\varphi_{j,t}$ with existentially quantified $\vec{y}_j \setminus \vec{y}_{j,t}$ for further rewriting. The existentially quantified variables in each of the disjuncts do not contain z , and so our recursion is well-founded. The proof of the following theorem is straightforward:

Theorem 23. *For any CQ $q(x)$ and ontology \mathcal{T} , $q^\dagger(\vec{x})$ is a PE-rewriting of q and \mathcal{T} (over complete data).*

The exact form of the rewriting q^\dagger depends on the choice of the variables z . We now consider two strategies for choosing these variables in the case of tree-shaped CQs. Let

$$d_{\mathcal{T}}^q = 1 + \max_{z \in \vec{y}} |\{t = (t_r, t_i) \in \Theta_{\mathcal{T}}^q \mid z \in t_i\}|.$$

We call $d_{\mathcal{T}}^q$ the *tree-witness degree* of q and \mathcal{T} . For example, the tree-witness degree of any CQ and ontology of depth 1 is at most 2, as observed in the proof of Theorem 10. In general, however, it can only be bounded by $1 + |\Theta_{\mathcal{T}}^q|$.

Given a tree-shaped CQ $q(\vec{x}) = \exists\vec{y} \varphi(\vec{x}, \vec{y})$, we pick some variable as its root and define a partial order \preceq on the variables of q by taking $z \preceq z'$ iff z' occurs in the subtree of q rooted in z . The strategy used in [8] chooses the smallest z with respect to \preceq . Since the number of distinct subtrees of q is bounded by $|q|$ and NDL programs allow for structure sharing, this strategy yields an NDL-rewriting of size $|\mathcal{T}| \cdot |q| \cdot d_{\mathcal{T}}^q$:

Corollary 24 ([8]). *Any tree-shaped CQ and ontology with polynomially-many tree-witnesses have a polynomial-size NDL-rewriting.*

The depth of recursion in the rewiring process with the above strategy is $|q|$ in the worst case. Therefore, we can only obtain a PE-rewriting of exponential size in $|q|$. However, if we adopt the strategy of choosing z that splits the graph of each φ_j in half, then the depth of recursion does not exceed $\log |q|$, and so the resulting PE-rewriting is of polynomial size for q and \mathcal{T} of bounded tree-witness degree. This strategy is based on the following fact:

Proposition 25. *Any tree $T = (V, E)$ contains a vertex $v \in V$ such that each connected component obtained by removing v from T has at most $|V|/2$ vertices.*

As a consequence, we obtain:

Theorem 26. *For any tree-shaped CQ q and any ontology \mathcal{T} , there is a PE-rewriting of size $|\mathcal{T}| \cdot |q|^{1+\log d_{\mathcal{T}}^q}$ (over complete data).*

Proof. Denote by $F(n)$ the maximal size of p^\dagger for a subquery p of q with at most n atoms. We show by induction that $F(n) \leq |\mathcal{T}| \cdot n^{1+\log d}$, where $d = d_{\mathcal{T}}^q$. By definition, for each component p_j of the finest partition of p , the length of its contribution to p^\dagger does not exceed

$$F(n_j) + \sum_{i=1}^{d-1} (F(n_j - m_{ji}) + |\mathcal{T}| \cdot m_{ji}),$$

where n_j is the number of atoms in p_j and m_{ji} is the number of atoms in the i th tree witness with $z \in t_i$, $1 \leq m_{ji} \leq n_j$. By the induction hypothesis, the length of the contribution of p_j does not exceed

$$\begin{aligned} |\mathcal{T}| \cdot n_j^{1+\log d} + |\mathcal{T}| \cdot \sum_{i=1}^{d-1} ((n_j - m_{ji})^{1+\log d} + m_{ji}) &\leq \\ |\mathcal{T}| \cdot (n_j^{1+\log d} + (d-1) \cdot n_j^{1+\log d}) &= |\mathcal{T}| \cdot d \cdot n_j^{1+\log d}. \end{aligned}$$

By Proposition 25, we can choose z (at the preceding step) so that \mathbf{p} with n atoms is split into components $\mathbf{p}_1, \dots, \mathbf{p}_k$ each of which has $n_j \leq n/2$ atoms (by definition, $\sum_{j=1}^k n_j = n$). Then we obtain

$$F(n) \leq |\mathcal{T}| \cdot d \cdot \sum_{j=1}^k ((n/2)^{\log d} \cdot n_j) \leq |\mathcal{T}| \cdot n^{1+\log d},$$

as required. □

Corollary 27. *Any tree-shaped CQ q and ontology \mathcal{T} of depth 1 have a PE-rewriting of size $|\mathcal{T}| \cdot |q|^2$ (over complete data).*

8 Conclusions

We established a fundamental link between FO-rewritings of CQs over *OWL 2 QL* ontologies of depth 1 and 2 and—via the hypergraph functions and programs—classical computational models for Boolean functions. This link allowed us to apply the Boolean complexity theory and obtain both polynomial upper and exponential (or superpolynomial) lower bounds for the size of rewritings. It is to be noted that the high lower bounds were proved for CQs and ontologies with polynomially-many tree witnesses and polynomial-size chases.

A few challenging important questions remain open: (i) Are all hypergraphs representable as subgraphs of some tree-witness hypergraphs? (ii) Do all tree-shaped CQs have polynomial-size rewritings over ontologies of depth 2 (more generally, of bounded depth)? (iii) What is the size of CQ rewritings over a *fixed* ontology in the worst case? (The last question is related to the *non-uniform* approach to the complexity of query answering in OBDA on the level of individual ontologies [25].)

Acknowledgment This work was supported by the U.K. EPSRC project ‘ExODA: Integrating Description Logics and Database Technologies for Expressive Ontology-Based Data Access’ (EP/H05099X).

References

- [1] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
- [2] N. Alon and R. Boppana. The monotone circuit complexity of Boolean functions. *Combinatorica*, 7(1):1–22, 1987.
- [3] S. Arora and B. Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, New York, NY, USA, 1st edition, 2009.
- [4] A. Artale, D. Calvanese, R. Kontchakov, and M. Zakharyashev. The DL-Lite family and relations. *Journal of Artificial Intelligence Research (JAIR)*, 36:1–69, 2009.
- [5] B. Aspvall, M. Plass, and R. Tarjan. A linear-time algorithm for testing the truth of certain quantified boolean formulas. *Information Processing Letters*, 8(3):121–123, 1979.
- [6] J. Avigad. Eliminating definitions and Skolem functions in first-order logic. In *Proc. of the 16th Annual IEEE Symposium on Logic in Computer Science (LICS’01)*, pages 139–146. IEEE Computer Society, 2001.
- [7] J.-F. Baget, M. Leclère, M.-L. Mugnier, and E. Salvat. Extending decidable cases for rules with existential variables. In *Proc. of the 21st Int. Joint Conf. on Artificial Intelligence (IJCAI 2009)*, pages 677–682, 2009.
- [8] M. Bienvenu, M. Ortiz, M. Simkus, and G. Xiao. Tractable queries for lightweight description logics. In *Proc. of the 23rd Int. Joint Conf. on Artificial Intelligence (IJCAI 2013)*. IJCAI/AAAI, 2013.

- [9] E. Börger, E. Grädel, and Y. Gurevich. *The Classical Decision Problem*. Perspectives in Mathematical Logic. Springer, 1997.
- [10] A. Cali, G. Gottlob, and A. Pieris. Towards more expressive ontology languages: The query answering problem. *Artificial Intelligence*, 193:87–128, 2012.
- [11] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati. Tractable reasoning and efficient query answering in description logics: The *DL-Lite* family. *Journal of Automated Reasoning*, 39(3):385–429, 2007.
- [12] C. Chekuri and A. Rajaraman. Conjunctive query containment revisited. *Theoretical Computer Science*, 239(2):211–229, 2000.
- [13] A. Chortaras, D. Trivela, and G. Stamou. Optimized query rewriting for OWL 2 QL. In *Proc. of the 23rd Int. Conf. on Automated Deduction (CADE-23)*, volume 6803 of *Lecture Notes in Computer Science*, pages 192–206. Springer, 2011.
- [14] T. Eiter, M. Ortiz, M. Šimkus, T.-K. Tran, and G. Xiao. Query rewriting for Horn-SHIQ plus rules. In *Proc. of the 27th Nat. Conf. on Artificial Intelligence (AAAI 2012)*. AAAI Press, 2012.
- [15] G. Gottlob, G. Orsi, and A. Pieris. Ontological queries: Rewriting and optimization. In *Proc. of the 27th Int. Conf. on Data Engineering (ICDE 2011)*, pages 2–13. IEEE Computer Society, 2011.
- [16] G. Gottlob and T. Schwentick. Rewriting ontological queries into small nonrecursive datalog programs. In *Proc. of the 13th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR 2012)*. AAAI Press, 2012.
- [17] M. Grigni and M. Sipser. Monotone separation of logarithmic space from logarithmic depth. *Journal of Computer and System Sciences*, 50(3):433–437, 1995.
- [18] S. Jukna. *Boolean Function Complexity — Advances and Frontiers*, volume 27 of *Algorithms and Combinatorics*. Springer, 2012.
- [19] M. Karchmer and A. Wigderson. Monotone circuits for connectivity require super-logarithmic depth. In *Proc. of the 20th Annual ACM Symposium on Theory of Computing (STOC’88)*, pages 539–550. ACM, 1988.
- [20] S. Kikot, R. Kontchakov, V. Podolskii, and M. Zakharyashev. Exponential lower bounds and separation for query rewriting. In *Proc. of the 39th Int. Colloquium on Automata, Languages, and Programming (ICALP 2012), II*, volume 7392 of *Lecture Notes in Computer Science*, pages 263–274. Springer, 2012.
- [21] S. Kikot, R. Kontchakov, and M. Zakharyashev. On (In)Tractability of OBDA with OWL 2 QL. In *Proc. of DL*, volume 745 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2011.
- [22] S. Kikot, R. Kontchakov, and M. Zakharyashev. Conjunctive query answering with OWL 2 QL. In *Proc. of the 13th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR 2012)*, 2012.
- [23] M. König, M. Leclère, M.-L. Mugnier, and M. Thomazo. On the exploration of the query rewriting space with existential rules. In *Proc. of the 7th International Conference on Web Reasoning and Rule Systems (RR 2013)*, volume 7994 of *Lecture Notes in Computer Science*, pages 123–137. Springer, 2013.
- [24] C. Lutz, I. Seylan, D. Toman, and F. Wolter. The combined approach to OBDA: Taming role hierarchies using filters. In *Proc. of the 12th Int. Semantic Web Conf. (ISWC 2013)*, volume 8218 of *Lecture Notes in Computer Science*, pages 314–330. Springer, 2013.

- [25] C. Lutz and F. Wolter. Non-uniform data complexity of query answering in description logics. In *Proc. of the 13th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR 2012)*. AAAI, 2012.
- [26] H. Pérez-Urbina, B. Motik, and I. Horrocks. A comparison of query rewriting techniques for DL-Lite. In *Proc. of DL*, volume 477 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2009.
- [27] H. Pérez-Urbina, E. Rodríguez-Díaz, M. Grove, G. Konstantinidis, and E. Sirin. Evaluation of query rewriting approaches for OWL 2. In *Proc. of SSWS+HPCSW*, volume 943 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2012.
- [28] A. Poggi, D. Lembo, D. Calvanese, G. De Giacomo, M. Lenzerini, and R. Rosati. Linking data to ontologies. *Journal on Data Semantics*, X:133–173, 2008.
- [29] R. Raz and A. Wigderson. Monotone circuits for matching require linear depth. *Journal of the ACM*, 39(3):736–744, 1992.
- [30] A. Razborov. Lower bounds for the monotone complexity of some Boolean functions. *Dokl. Akad. Nauk SSSR*, 281(4):798–801, 1985.
- [31] A. Razborov. Lower bounds for deterministic and nondeterministic branching programs. In *Proc. of the 8th Int. Symposium on Fundamentals of Computation Theory (FCT'91)*, volume 529 of *Lecture Notes in Computer Science*, pages 47–60. Springer, 1991.
- [32] M. Rodríguez-Muro, R. Kontchakov, and M. Zakharyashev. Ontology-based data access: Ontop of databases. In *Proc. of the 12th Int. Semantic Web Conf. (ISWC 2013)*, volume 8218 of *Lecture Notes in Computer Science*, pages 558–573. Springer, 2013.
- [33] R. Rosati. Prexto: Query rewriting under extensional constraints in DL-Lite. In *Proc. of the 9th Extended Semantic Web Conf. (ESWC 2012)*, volume 7295 of *Lecture Notes in Computer Science*, pages 360–374. Springer, 2012.
- [34] R. Rosati and A. Almatelli. Improving query answering over DL-Lite ontologies. In *Proc. of the 12th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR 2010)*. AAAI Press, 2010.
- [35] M. Yannakakis. Algorithms for acyclic database schemes. In *Proc. of the 7th Int. Conf. on Very Large Data Bases (VLDB'81)*, pages 82–94. IEEE Computer Society, 1981.