

Answering FO+MOD Queries Under Updates on Bounded Degree Databases*

Christoph Berkholz¹, Jens Keppeler², and Nicole Schweikardt³

- 1 Humboldt-Universität zu Berlin, Berlin, Germany
berkholz@informatik.hu-berlin.de
- 2 Humboldt-Universität zu Berlin, Berlin, Germany
keppelej@informatik.hu-berlin.de
- 3 Humboldt-Universität zu Berlin, Berlin, Germany
schweikn@informatik.hu-berlin.de

Abstract

We investigate the query evaluation problem for fixed queries over fully dynamic databases, where tuples can be inserted or deleted. The task is to design a dynamic algorithm that immediately reports the new result of a fixed query after every database update.

We consider queries in first-order logic (FO) and its extension with modulo-counting quantifiers (FO+MOD), and show that they can be efficiently evaluated under updates, provided that the dynamic database does not exceed a certain degree bound.

In particular, we construct a data structure that allows to answer a Boolean FO+MOD query and to compute the size of the query result within constant time after every database update. Furthermore, after every update we are able to immediately enumerate the new query result with constant delay between the output tuples. The time needed to build the data structure is linear in the size of the database.

Our results extend earlier work on the evaluation of first-order queries on static databases of bounded degree and rely on an effective Hanf normal form for FO+MOD recently obtained by Heimberg, Kuske, and Schweikardt (LICS 2016).

1998 ACM Subject Classification H.2.4 [Systems] Relational Databases, Query Processing, H.2.3 [Languages] Query Languages, F.1.2 [Modes of Computation] Interactive and Reactive Computation

Keywords and phrases dynamic databases, query enumeration, counting problem, first-order logic with modulo-counting quantifiers, Hanf locality

Digital Object Identifier 10.4230/LIPIcs.ICDT.2017.8

1 Introduction

Query evaluation is a fundamental task in databases, and a vast amount of literature is devoted to the complexity of this problem. In this paper we study query evaluation on relational databases in the “dynamic setting”, where the database may be updated by inserting or deleting tuples. In this setting, an evaluation algorithm receives a query φ and an initial database D and starts with a preprocessing phase that computes a suitable data structure to represent the result of evaluating φ on D . After every database update, the data structure is updated so that it represents the result of evaluating φ on the updated

* We acknowledge the financial support by the German Research Foundation DFG under grant SCHW 837/5-1.



database. The data structure shall be designed in such a way that it quickly provides the query result, preferably in constant time (i.e., independent of the database size). We focus on the following flavours of query evaluation.

- *Testing*: Decide whether a given tuple \bar{a} is contained in $\varphi(D)$.
- *Counting*: Compute $|\varphi(D)|$ (i.e., the number of tuples that belong to $\varphi(D)$).
- *Enumeration*: Enumerate $\varphi(D)$ with a bounded delay between the output tuples.

Here, as usual, $\varphi(D)$ denotes the k -ary relation obtained by evaluating a k -ary query φ on a relational database D . For *Boolean* queries, all three tasks boil down to

- *Answering*: Decide if $\varphi(D) \neq \emptyset$.

Compared to the *dynamic descriptive complexity* framework introduced by Patnaik and Immerman [17], which focuses on the *expressive power* of first-order logic on dynamic databases and has led to a rich body of literature (see [18] for a survey), we are interested in the *complexity* of query evaluation. The query language studied in this paper is FO+MOD, the extension of first-order logic FO with modulo-counting quantifiers of the form $\exists^{i \bmod m} x \psi$, expressing that the number of witnesses x that satisfy ψ is congruent to i modulo m . FO+MOD can be viewed as a subclass of SQL that properly extends the relational algebra.

Following [2], we say that a query evaluation algorithm is efficient if the update time is either constant or at most polylogarithmic ($\log^c n$) in the size of the database. As a consequence, efficient query evaluation in the dynamic setting is only possible if the static problem (i.e., the setting without database updates) can be solved in linear or pseudo-linear ($n^{1+\varepsilon}$) time. Since this is not always possible, we provide a short overview on known results about first-order query evaluation on static databases and then proceed by discussing our results in the dynamic setting.

First-order query evaluation on static databases. The problem of deciding whether a given database D satisfies a FO-sentence φ is AW[*]-complete (parameterised by $\|\varphi\|$) and it is therefore generally believed that the evaluation problem cannot be solved in time $f(\|\varphi\|)\|D\|^c$ for any computable f and constant c (here, $\|\varphi\|$ and $\|D\|$ denote the size of the query and the database, respectively). For this reason, a long line of research focused on increasing classes of sparse instances ranging from databases of *bounded degree* [19] (where every domain element occurs only in a constant number of tuples in the database) to classes that are *nowhere dense* [9]. In particular, Boolean first-order queries can be evaluated on classes of databases of bounded degree in linear time $f(\|\varphi\|)\|D\|$, where the constant factor $f(\|\varphi\|)$ is 3-fold exponential in $\|\varphi\|$ [19, 7]. As a matter of fact, Frick and Grohe [7] showed that the 3-fold exponential blow-up in terms of the query size is unavoidable assuming $\text{FPT} \neq \text{AW}[*]$.

Durand and Grandjean [5] and Kazana and Segoufin [11] considered the task of enumerating the result of a k -ary first-order query on bounded degree databases and showed that after a linear time preprocessing phase the query result can be enumerated with constant delay. This result was later extended to classes of databases of bounded expansion [12]. Kazana and Segoufin [12] also showed that counting the number of result tuples of a k -ary first-order query on databases of bounded expansion (and hence also on databases of bounded degree) can be done in time $f(\|\varphi\|)\|D\|$. In [6] an analogous result was obtained for classes of databases of low degree (i.e., degree at most $\|D\|^{o(1)}$) and pseudo-linear time $f(\|\varphi\|)\|D\|^{1+\varepsilon}$; the paper also presented an algorithm for enumerating the query results with constant delay after pseudo-linear time preprocessing.

Our contribution. We extend the known linear time algorithms for first-order logic on classes of databases of bounded degree to the more expressive query language FO+MOD.

Moreover, and more importantly, we lift the tractability to the dynamic setting and show that the result of FO and FO+MOD-queries can be maintained with constant update time. In particular, we obtain the following results. Let φ be a fixed k -ary FO+MOD-query and d a fixed degree bound on the databases under consideration. Given an initial database D , we construct in linear time $f(\|\varphi\|, d)\|D\|$ a data structure that can be updated in constant time $f(\|\varphi\|, d)$ when a tuple is inserted into or deleted from a relation of D . After each update the data structure allows to

- immediately answer φ on D if φ is a Boolean query (Theorem 4.1),
- test for a given tuple \bar{a} whether $\bar{a} \in \varphi(D)$ in time $\mathcal{O}(k^2)$ (Theorem 6.1),
- immediately output the number of result tuples $|\varphi(D)|$ (Theorem 8.1), and
- enumerate all tuples $(a_1, \dots, a_k) \in \varphi(D)$ with $\mathcal{O}(k^3)$ delay (Theorem 9.4).

For fixed d , the parameter function $f(\|\varphi\|, d)$ is 3-fold exponential in terms of the query size, which is (by Frick and Grohe [7]) optimal assuming $\text{FPT} \neq \text{AW}[*]$.

Outline. Our dynamic query evaluation algorithm crucially relies on the locality of FO+MOD and in particular an effective Hanf normal form for FO+MOD on databases of bounded degree recently obtained by Heimberg, Kuske, and Schweikardt [10]. After some basic definitions in Section 2 we briefly state their result in Section 3 and obtain a dynamic algorithm for Boolean FO+MOD-queries in Section 4. After some preparations for non-Boolean queries in Section 5, we present the algorithm for testing in Section 6. In Section 7 we reduce the task of counting and enumerating FO+MOD-queries in the dynamic setting to the problem of counting and enumerating independent sets in graphs of bounded degree. We use this reduction to provide efficient dynamic counting and enumeration algorithms in Section 8 and 9, respectively, and we conclude in Section 10. Due to space constraints some technical proofs are deferred to the full version of the paper [3].

2 Preliminaries

We write \mathbb{N} for the set of non-negative integers and let $\mathbb{N}_{\geq 1} := \mathbb{N} \setminus \{0\}$ and $[n] := \{1, \dots, n\}$ for all $n \in \mathbb{N}_{\geq 1}$. By 2^M we denote the power set of a set M . For a partial function f we write $\text{dom}(f)$ and $\text{codom}(f)$ for the domain and the codomain of f , respectively.

Databases. We fix a countably infinite set **dom**, the *domain* of potential database entries. Elements in **dom** are called *constants*. A *schema* is a finite set σ of relation symbols, where each $R \in \sigma$ is equipped with a fixed *arity* $\text{ar}(R) \in \mathbb{N}_{\geq 1}$. Let us fix a schema $\sigma = \{R_1, \dots, R_{|\sigma|}\}$. A *database* D of schema σ (σ -db, for short), is of the form $D = (R_1^D, \dots, R_{|\sigma|}^D)$, where each R_i^D is a finite subset of **dom** ^{$\text{ar}(R_i)$} . The *active domain* $\text{adom}(D)$ of D is the smallest subset A of **dom** such that $R_i^D \subseteq A^{\text{ar}(R_i)}$ for each R_i in σ .

The *Gaifman graph* of a σ -db D is the undirected simple graph $G^D = (V, E)$ with vertex set $V := \text{adom}(D)$, where there is an edge between vertices u and v whenever $u \neq v$ and there are $R \in \sigma$ and $(a_1, \dots, a_{\text{ar}(R)}) \in R^D$ such that $u, v \in \{a_1, \dots, a_{\text{ar}(R)}\}$. A σ -db D is called *connected* if its Gaifman graph G^D is connected; the *connected components* of D are the connected components of G^D . The *degree* of a database D is the degree of its Gaifman graph G^D , i.e., the maximum number of neighbours ^{D} of a node of G^D . Throughout this paper we fix a number $d \in \mathbb{N}$ and restrict attention to databases of degree at most d .

Updates. We allow to update a given database of schema σ by inserting or deleting tuples as follows (note that both types of commands may change the database's active domain and

the database's degree). A *deletion* command is of the form $\text{delete } R(a_1, \dots, a_r)$ for $R \in \sigma$, $r = \text{ar}(R)$, and $a_1, \dots, a_r \in \mathbf{dom}$. When applied to a σ -db D , it results in the updated σ -db D' with $R^{D'} = R^D \setminus \{(a_1, \dots, a_r)\}$ and $S^{D'} = S^D$ for all $S \in \sigma \setminus \{R\}$.

An *insertion* command is of the form $\text{insert } R(a_1, \dots, a_r)$ for $R \in \sigma$, $r = \text{ar}(R)$, and $a_1, \dots, a_r \in \mathbf{dom}$. When applied to a σ -db D in the unrestricted setting, it results in the updated σ -db D' with $R^{D'} = R^D \cup \{(a_1, \dots, a_r)\}$ and $S^{D'} = S^D$ for all $S \in \sigma \setminus \{R\}$. In this paper, we restrict attention to databases of degree at most d . Therefore, when applying an insertion command to a σ -db D of degree $\leq d$, the command is carried out only if the resulting database D' still has degree $\leq d$; otherwise D remains unchanged and instead of carrying out the insertion command, an error message is returned.

Queries. We fix a countably infinite set \mathbf{var} of *variables*. We consider the extension FO+MOD of first-order logic FO with modulo-counting quantifiers. For a fixed schema σ , the set FO+MOD[σ] is built from atomic formulas of the form $x_1 = x_2$ and $R(x_1, \dots, x_{\text{ar}(R)})$, for $R \in \sigma$ and variables $x_1, x_2, \dots, x_{\text{ar}(R)} \in \mathbf{var}$, and is closed under Boolean connectives \neg, \wedge , existential first-order quantifiers $\exists x$, and modulo-counting quantifiers $\exists^{i \bmod m} x$, for a variable $x \in \mathbf{var}$ and integers $i, m \in \mathbb{N}$ with $m \geq 2$ and $i < m$. The intuitive meaning of a formula of the form $\exists^{i \bmod m} x \psi$ is that the number of witnesses x that satisfy ψ is congruent i modulo m . As usual, $\forall x, \vee, \rightarrow, \leftrightarrow$ will be used as abbreviations when constructing formulas. It will be convenient to add the quantifier $\exists^{\geq m} x$, for $m \in \mathbb{N}_{\geq 1}$; a formula of the form $\exists^{\geq m} x \psi$ expresses that the number of witnesses x which satisfy ψ is $\geq m$. This quantifier is just syntactic sugar and does not increase the expressive power of FO+MOD.

The *quantifier rank* $\text{qr}(\varphi)$ of a FO+MOD-formula φ is the maximum nesting depth of quantifiers that occur in φ . By $\text{free}(\varphi)$ we denote the set of all *free variables* of φ , i.e., all variables x that have at least one occurrence in φ that is not within a quantifier of the form $\exists x, \exists^{\geq m} x$, or $\exists^{i \bmod m} x$. A *sentence* is a formula φ with $\text{free}(\varphi) = \emptyset$.

An *assignment* for φ in a σ -db D is a partial mapping α from \mathbf{var} to $\text{adom}(D)$, where $\text{free}(\varphi) \subseteq \text{dom}(\alpha)$. We write $(D, \alpha) \models \varphi$ to indicate that φ is satisfied when evaluated in D with respect to *active domain semantics* while interpreting every free occurrence of a variable x with the constant $\alpha(x)$. Recall from [1] that “active domain semantics” means that quantifiers are evaluated with respect to the database's active domain. In particular, $(D, \alpha) \models \exists x \psi$ iff there exists an $a \in \text{adom}(D)$ such that $(D, \alpha \frac{a}{x}) \models \psi$, where $\alpha \frac{a}{x}$ is the assignment α' with $\alpha'(x) = a$ and $\alpha'(y) = \alpha(y)$ for all $y \in \text{dom}(\alpha) \setminus \{x\}$. Accordingly, $(D, \alpha) \models \exists^{\geq m} x \psi$ iff $|\{a \in \text{adom}(D) : (D, \alpha \frac{a}{x}) \models \psi\}| \geq m$, and $(D, \alpha) \models \exists^{i \bmod m} x \psi$ iff $|\{a \in \text{adom}(D) : (D, \alpha \frac{a}{x}) \models \psi\}| \equiv i \pmod{m}$.

A *k-ary FO+MOD query of schema σ* is of the form $\varphi(x_1, \dots, x_k)$ where $k \in \mathbb{N}$, $\varphi \in \text{FO+MOD}[\sigma]$, and $\text{free}(\varphi) \subseteq \{x_1, \dots, x_k\}$. We will often assume that the tuple (x_1, \dots, x_k) is clear from the context and simply write φ instead of $\varphi(x_1, \dots, x_k)$ and $(D, (a_1, \dots, a_k)) \models \varphi$ instead of $(D, \frac{a_1, \dots, a_k}{x_1, \dots, x_k}) \models \varphi$, where $\frac{a_1, \dots, a_k}{x_1, \dots, x_k}$ denotes the assignment α with $\alpha(x_i) = a_i$ for all $i \in [k]$. When evaluated in a σ -db D , the k -ary query $\varphi(x_1, \dots, x_k)$ yields the k -ary relation

$$\varphi(D) := \left\{ (a_1, \dots, a_k) \in \text{adom}(D)^k : (D, \frac{a_1, \dots, a_k}{x_1, \dots, x_k}) \models \varphi \right\}.$$

Boolean queries are k -ary queries with $k = 0$. As usual, for Boolean queries we will write $\varphi(D) = \mathbf{no}$ instead of $\varphi(D) = \emptyset$, and $\varphi(D) = \mathbf{yes}$ instead of $\varphi(D) \neq \emptyset$; and we write $D \models \varphi$ to indicate that $(D, \alpha) \models \varphi$ for any assignment α .

Sizes and Cardinalities. The *size* $\|\sigma\|$ of a schema σ is the sum of the arities of its relation symbols. The size $\|\varphi\|$ of an FO+MOD query φ of schema σ is the length of φ when

viewed as a word over the alphabet $\sigma \cup \mathbf{var} \cup \mathbb{N} \cup \{=, \wedge, \neg, \exists, \text{mod}, (,)\}$. For a k -ary query $\varphi(x_1, \dots, x_k)$ and a σ -db D , the *cardinality of the query result* is the number $|\varphi(D)|$ of tuples in $\varphi(D)$. The *cardinality* $|D|$ of a σ -db D is defined as the number of tuples stored in D , i.e., $|D| := \sum_{R \in \sigma} |R^D|$. The *size* $\|D\|$ of D is defined as $\|\sigma\| + |\text{adom}(D)| + \sum_{R \in \sigma} \text{ar}(R) \cdot |R^D|$ and corresponds to the size of a reasonable encoding of D .

Dynamic Algorithms for Query Evaluation. We adopt the framework for dynamic algorithms for query evaluation of [2]; the next paragraphs are taken almost verbatim from [2]. Following [4], we use Random Access Machines (RAMs) with $\mathcal{O}(\log n)$ word-size and a uniform cost measure to analyse our algorithms. We will assume that the RAM's memory is initialised to 0. In particular, if an algorithm uses an array, we will assume that all array entries are initialised to 0, and this initialisation comes at no cost (in real-world computers this can be achieved by using the *lazy array initialisation technique*, cf. e.g. [16]). A further assumption is that for every fixed dimension $k \in \mathbb{N}_{\geq 1}$ we have available an unbounded number of k -ary arrays \mathbf{A} such that for given $(n_1, \dots, n_k) \in \mathbb{N}^k$ the entry $\mathbf{A}[n_1, \dots, n_k]$ at position (n_1, \dots, n_k) can be accessed in constant time.¹

Our algorithms will take as input a k -ary FO+MOD-query $\varphi(x_1, \dots, x_k)$, a parameter d , and a σ -db D_0 of degree $\leq d$. For all query evaluation problems considered in this paper, we aim at routines **preprocess** and **update** which achieve the following.

Upon input of $\varphi(x_1, \dots, x_k)$ and D_0 , **preprocess** builds a data structure \mathbf{D} which represents D_0 (and which is designed in such a way that it supports the evaluation of φ on D_0). Upon input of a command **update** $R(a_1, \dots, a_r)$ (with **update** $\in \{\text{insert}, \text{delete}\}$), calling **update** modifies the data structure \mathbf{D} such that it represents the updated database D . The *preprocessing time* t_p is the time used for performing **preprocess**; the *update time* t_u is the time used for performing an **update**. In this paper, t_u will be independent of the size of the current database D . By **init** we denote the particular case of the routine **preprocess** upon input of a query $\varphi(x_1, \dots, x_k)$ and the *empty* database D_\emptyset (where $R^{D_\emptyset} = \emptyset$ for all $R \in \sigma$). The *initialisation time* t_i is the time used for performing **init**. In all dynamic algorithms presented in this paper, the **preprocess** routine for input of $\varphi(x_1, \dots, x_k)$ and D_0 will carry out the **init** routine for $\varphi(x_1, \dots, x_k)$ and then perform a sequence of $|D_0|$ update operations to insert all the tuples of D_0 into the data structure. Consequently, $t_p = t_i + |D_0| \cdot t_u$.

In the following, D will always denote the database that is currently represented by the data structure \mathbf{D} .

To solve the *enumeration problem under updates*, apart from the routines **preprocess** and **update**, we aim at a routine **enumerate** such that calling **enumerate** invokes an enumeration of all tuples (without repetition) that belong to the query result $\varphi(D)$. The *delay* t_d is the maximum time used during a call of **enumerate**

- until the output of the first tuple (or the end-of-enumeration message **EOE**, if $\varphi(D) = \emptyset$),
- between the output of two consecutive tuples, and
- between the output of the last tuple and the end-of-enumeration message **EOE**.

To *test* if a given tuple belongs to the query result, instead of **enumerate** we aim at a routine **test** which upon input of a tuple $\bar{a} \in \mathbf{dom}^k$ checks whether $\bar{a} \in \varphi(D)$. The *testing time* t_t is the time used for performing a **test**. To solve the *counting problem under updates*, instead of **enumerate** or **test** we aim at a routine **count** which outputs the cardinality $|\varphi(D)|$ of the query result. The *counting time* t_c is the time used for performing a **count**.

¹ While this can be accomplished easily in the RAM-model, for an implementation on real-world computers one would probably have to resort to replacing our use of arrays by using suitably designed hash functions.

To *answer* a *Boolean* query under updates, instead of **enumerate**, **test**, or **count** we aim at a routine **answer** that produces the answer **yes** or **no** of φ on D . The *answer time* t_a is the time used for performing **answer**. Whenever speaking of a *dynamic algorithm*, we mean an algorithm that has routines **preprocess** and **update** and, depending on the problem at hand, at least one of the routines **answer**, **test**, **count**, and **enumerate**.

Throughout the paper, we often adopt the view of *data complexity* and suppress factors that may depend on the query φ or the degree bound d , but not on the database D . E.g., “linear preprocessing time” means $t_p \leq f(\varphi, d) \cdot \|D_0\|$ and “constant update time” means $t_u \leq f(\varphi, d)$, for a function f with codomain \mathbb{N} . When writing *poly*(n) we mean $n^{\mathcal{O}(1)}$.

3 Hanf Normal Form for FO+MOD

Our algorithms for evaluating FO+MOD queries rely on a decomposition of FO+MOD queries into *Hanf normal form*. To describe this normal form, we need some more notation.

Two formulas φ and ψ of schema σ are called *d-equivalent* (in symbols: $\varphi \equiv_d \psi$) if $\varphi(D) = \psi(D)$ for all σ -dbs D of degree $\leq d$. For a σ -db D and a set $A \subseteq \text{adom}(D)$ we write $D[A]$ to denote the restriction of D to the domain A , i.e., $R^{D[A]} = \{\bar{a} \in R^D : \bar{a} \in A^{\text{ar}(R)}\}$, for all $R \in \sigma$. For two σ -dbs D and D' and two k -tuples $\bar{a} = (a_1, \dots, a_k)$ and $\bar{a}' = (a'_1, \dots, a'_k)$ of elements in $\text{adom}(D)$ and $\text{adom}(D')$, resp., we write $(D, \bar{a}) \cong (D', \bar{a}')$ to indicate that there is an isomorphism² π from D to D' that maps a_i to a'_i for all $i \in [k]$.

The *distance* $\text{dist}^D(a, b)$ between two elements $a, b \in \text{adom}(D)$ is the minimal length (i.e., the number of edges) of a path from a to b in D 's Gaifman graph G^D (if no such path exists, we let $\text{dist}^D(a, b) = \infty$; note that $\text{dist}^D(a, a) = 0$). For $r \geq 0$ and $a \in \text{adom}(D)$, the *r-ball* around a in D is the set $N_r^D(a) := \{b \in \text{adom}(D) : \text{dist}^D(a, b) \leq r\}$. For a σ -db D and a tuple $\bar{a} = (a_1, \dots, a_k)$ we let $N_r^D(\bar{a}) := \bigcup_{i \in [k]} N_r^D(a_i)$. The *r-neighbourhood* around \bar{a} in D is defined as the σ -db $\mathcal{N}_r^D(\bar{a}) := D[N_r^D(\bar{a})]$.

For $r \geq 0$ and $k \geq 1$, a *type* τ (over σ) with k centres and radius r (for short: *r-type with k centres*) is of the form (T, \bar{t}) , where T is a σ -db, $\bar{t} \in \text{adom}(T)^k$, and $\text{adom}(T) = N_r^T(\bar{t})$. The elements in \bar{t} are called the *centres* of τ . For a tuple $\bar{a} \in \text{adom}(D)^k$, the *r-type of \bar{a} in D* is defined as the *r-type with k centres* $(\mathcal{N}_r^D(\bar{a}), \bar{a})$.

For a given *r-type with k centres* $\tau = (T, \bar{t})$ it is straightforward to construct a first-order formula $\text{sph}_\tau(\bar{x})$ (depending on r and τ) with k free variables $\bar{x} = (x_1, \dots, x_k)$ which expresses that the *r-type of \bar{x}* is isomorphic to τ , i.e., for every σ -db D and all $\bar{a} = (a_1, \dots, a_k) \in \text{adom}(D)^k$ we have $(D, \bar{a}) \models \text{sph}_\tau(\bar{x}) \iff (\mathcal{N}_r^D(\bar{a}), \bar{a}) \cong (T, \bar{t})$. The formula $\text{sph}_\tau(\bar{x})$ is called a *sphere-formula* (over σ and \bar{x}); the numbers r and k are called *locality radius* and *arity*, resp., of the sphere-formula.

A *Hanf-sentence* (over σ) is a sentence of the form $\exists^{\geq m} x \text{sph}_\tau(x)$ or $\exists^{i \bmod m} x \text{sph}_\tau(x)$, where τ is an *r-type* (over σ) with 1 centre, for some $r \geq 0$. The number r is called *locality radius* of the Hanf-sentence. A formula in *Hanf normal form* (over σ) is a Boolean combination³ of sphere-formulas and Hanf-sentences (over σ). The *locality radius* of a formula ψ in Hanf normal form is the maximum of the locality radii of the Hanf-sentences and the sphere-formulas that occur in ψ . The formula is *d-bounded* if all types τ that occur in sphere-formulas or Hanf-sentences of ψ are *d-bounded*, i.e., T is of degree $\leq d$, where $\tau = (T, \bar{t})$. Our query evaluation algorithms for FO+MOD rely on the following result by Heimberg, Kuske, and Schweikardt [10].

² An *isomorphism* $\pi: D \rightarrow D'$ is a bijection from $\text{adom}(D)$ to $\text{adom}(D')$ with $(b_1, \dots, b_r) \in R^D \iff (\pi(b_1), \dots, \pi(b_r)) \in R^{D'}$ for all $R \in \sigma$, for $r := \text{ar}(R)$, and for all $b_1, \dots, b_r \in \text{adom}(D)$.

³ Throughout this paper, whenever we speak of *Boolean combinations* we mean *finite* Boolean combinations.

► **Theorem 3.1** ([10]). *There is an algorithm which receives as input a degree bound $d \in \mathbb{N}$ and a FO+MOD[σ]-formula φ , and constructs a d -equivalent formula ψ in Hanf normal form (over σ) with the same free variables as φ . For any $d \geq 2$, the formula ψ is d -bounded and has locality radius $\leq 4^{\text{qr}(\varphi)}$, and the algorithm's runtime is $2^{d^{2^{\mathcal{O}(\|\varphi\| + \|\sigma\|)}}$.*

The first step of all our query evaluation algorithms is to use Theorem 3.1 to transform a given query $\varphi(\bar{x})$ into a d -equivalent query $\psi(\bar{x})$ in Hanf normal form. The following lemma summarises easy facts that are useful for evaluating the sphere-formulas that occur in ψ .

► **Lemma 3.2.** *Let $d \geq 2$ and let D be a σ -db of degree $\leq d$. Let $r \geq 0$, $k \geq 1$, and $\bar{a} = (a_1, \dots, a_k) \in \text{adom}(D)$.*

- (a) $|N_r^D(\bar{a})| \leq k \sum_{i=0}^r d^i \leq kd^{r+1}$.
- (b) *Given D and \bar{a} , the r -neighbourhood $N_r^D(\bar{a})$ can be computed in time $(kd^{r+1})^{\mathcal{O}(\|\sigma\|)}$.*
- (c) $N_r^D(a_1, a_2)$ is connected if and only if $\text{dist}^D(a_1, a_2) \leq 2r + 1$.
- (d) *If $N_r^D(\bar{a})$ is connected, then $N_r^D(\bar{a}) \subseteq N_{r+(k-1)(2r+1)}^D(a_i)$, for all $i \in [k]$.*
- (e) *Let D' be a σ -db of degree $\leq d$ and let $\bar{b} = (b_1, \dots, b_k) \in \text{adom}(D')$. It can be tested in time $(kd^{r+1})^{\mathcal{O}(\|\sigma\| + kd^{r+1})} \leq 2^{\mathcal{O}(\|\sigma\|k^2d^{2r+2})}$ whether $(N_r^D(\bar{a}), \bar{a}) \cong (N_r^{D'}(\bar{b}), \bar{b})$.*

The time bound stated in part (e) of Lemma 3.2 is obtained by a brute-force approach. When using Luks' polynomial time isomorphism test for bounded degree graphs [15], the time bound of Lemma 3.2(e) can be improved to $(kd^{r+1})^{\text{poly}(d\|\sigma\|)}$. However, the asymptotic overall runtime of our algorithms for evaluating FO+MOD-queries won't improve when using Luks algorithm instead of the brute-force isomorphism test of Lemma 3.2(e).

4 Answering Boolean FO+MOD Queries Under Updates

In [7], Frick and Grohe showed that in the static setting (i.e., without database updates), Boolean FO-queries φ can be answered on databases D of degree $\leq d$ in time $2^{d^{2^{\mathcal{O}(\|\varphi\|)}}} \cdot \|D\|$. Our first main theorem extends their result to FO+MOD-queries and the dynamic setting.

► **Theorem 4.1.** *There is a dynamic algorithm that receives a schema σ , a degree bound $d \geq 2$, a Boolean FO+MOD[σ]-query φ , and a σ -db D_0 of degree $\leq d$, and computes within $t_p = f(\varphi, d) \cdot \|D_0\|$ preprocessing time a data structure that can be updated in time $t_u = f(\varphi, d)$ and allows to return the query result $\varphi(D)$ with answer time $t_a = \mathcal{O}(1)$. The function $f(\varphi, d)$ is of the form $2^{d^{2^{\mathcal{O}(\|\varphi\|)}}$.*

If φ is a d -bounded Hanf-sentence of locality radius r , then $f(\varphi, d) = 2^{\mathcal{O}(\|\sigma\|d^{2r+2})}$, and the initialisation time is $t_i = \mathcal{O}(\|\varphi\|)$.

Proof. W.l.o.g. we assume that all the symbols of σ occur in φ (otherwise, we remove from σ all symbols that do not occur in φ). In the preprocessing routine, we first use Theorem 3.1 to transform φ into a d -equivalent sentence ψ in Hanf normal form; this takes time $2^{d^{2^{\mathcal{O}(\|\varphi\|)}}$. The sentence ψ is a Boolean combination of d -bounded Hanf-sentences (over σ) of locality radius at most $r := 4^{\text{qr}(\varphi)}$. Let ρ_1, \dots, ρ_s be the list of all types that occur in ψ . Thus, every Hanf-sentence in ψ is of the form $\exists^{\geq k} x \text{ sph}_{\rho_j}(x)$ or $\exists^{i \bmod m} x \text{ sph}_{\rho_j}(x)$ for some $j \in [s]$ and $k, i, m \in \mathbb{N}$ with $k \geq 1$, $m \geq 2$, and $i < m$. For each $j \in [s]$ let r_j be the radius of ρ_j . Thus, ρ_j is an r_j -type with 1 centre (over σ).

For each $j \in [s]$ our data structure will store the number $A[j]$ of all elements $a \in \text{adom}(D)$ whose r_j -type is isomorphic to ρ_j , i.e., $(N_{r_j}^D(a), a) \cong \rho_j$. The initialisation for the empty

database D_\emptyset lets $A[j] = 0$ for all $j \in [s]$. In addition to the array A , our data structure stores a Boolean value \mathbf{Ans} where $\mathbf{Ans} = \varphi(D)$ is the answer of the Boolean query φ on the current database D . This way, the query can be answered in time $\mathcal{O}(1)$ by simply outputting \mathbf{Ans} . The initialisation for the empty database D_\emptyset computes \mathbf{Ans} as follows. Every Hanf-sentence of the form $\exists^{\geq k} x \text{ sph}_{\rho_j}(x)$ in ψ is replaced by the Boolean constant **false**. Every Hanf-sentence of the form $\exists^{i \bmod m} x \text{ sph}_{\rho_j}(x)$ is replaced by **true** if $i = 0$ and by **false** otherwise. The resulting formula, a Boolean combination of the Boolean constants **true** and **false**, then is evaluated, and we let \mathbf{Ans} be the obtained result. The entire initialisation takes time at most $t_i = f(\varphi, d) = 2^{d^{2^{\mathcal{O}(\|\varphi\|)}}}$. If φ is a Hanf-sentence, we even have $t_i = \mathcal{O}(\|\varphi\|)$.

To update our data structure upon a command **update** $R(a_1, \dots, a_k)$, for $k = \text{ar}(R)$ and **update** $\in \{\text{insert}, \text{delete}\}$, we proceed as follows. The idea is to remove from the data structure the information on all the database elements whose r_j -neighbourhood (for some $j \in [s]$) is affected by the update, and then to recompute the information concerning all these elements on the updated database.

Let D_{old} be the database before the update is received and let D_{new} be the database after the update has been performed. We consider each $j \in [s]$. All elements whose r_j -neighbourhood might have changed, belong to the set $U_j := N_{r_j}^{D'}(\bar{a})$, where $D' := D_{new}$ if the update command is **insert** $R(\bar{a})$, and $D' := D_{old}$ if the update command is **delete** $R(\bar{a})$.

To remove the old information from $A[j]$, we compute for each $a \in U_j$ the neighbourhood $T_a := N_{r_j}^{D_{old}}(a)$, check whether $(T_a, a) \cong \rho_j$, and if so, decrement the value $A[j]$.

To recompute the new information for $A[j]$, we compute for all $a \in U_j$ the neighbourhood $T'_a := N_{r_j}^{D_{new}}(a)$, check whether $(T'_a, a) \cong \rho_j$, and if so, increment the value $A[j]$.

Using Lemma 3.2 we obtain for each $j \in [s]$ that $|U_j| \leq kd^{r_j+1}$. For each $a \in U_j$, the neighbourhoods T_a and T'_a can be computed in time $(d^{r_j+1})^{\mathcal{O}(\|\sigma\|)}$, and testing for isomorphism with ρ_j can be done in time $(d^{r_j+1})^{\mathcal{O}(\|\sigma\|+d^{r_j+1})}$. Thus, the update of $A[j]$ is done in time $k \cdot (d^{r_j+1})^{\mathcal{O}(\|\sigma\|+d^{r_j+1})} \leq 2^{d^{2^{\mathcal{O}(\|\varphi\|)}}}$ (note that $k \leq \|\sigma\| \leq \|\varphi\|$ and $r_j \leq 4^{\text{qr}(\varphi)} \leq 2^{\mathcal{O}(\|\varphi\|)}$).

After having updated $A[j]$ for each $j \in [s]$, we recompute the query answer \mathbf{Ans} as follows. Every Hanf-sentence of the form $\exists^{\geq k} x \text{ sph}_{\rho_j}(x)$ in ψ is replaced by the Boolean constant **true** if $A[j] \geq k$, and by the Boolean constant **false** otherwise. Every Hanf-sentence of the form $\exists^{i \bmod m} x \text{ sph}_{\rho_j}(x)$ is replaced by **true** if $A[j] \equiv i \bmod m$, and by **false** otherwise. The resulting formula, a Boolean combination of the Boolean constants **true** and **false**, then is evaluated, and we let \mathbf{Ans} be the obtained result. Thus, recomputing \mathbf{Ans} takes time $\text{poly}(\|\psi\|)$.

In summary, the entire update time is $t_u = f(\varphi, d) = 2^{d^{2^{\mathcal{O}(\|\varphi\|)}}}$. In case that φ is a d -bounded Hanf-sentence of locality radius r , we even have $t_u = k \cdot (d^{r+1})^{\mathcal{O}(\|\sigma\|+d^{r+1})} \leq 2^{\mathcal{O}(\|\sigma\|d^{2r+2})}$. This completes the proof of Theorem 4.1. \blacktriangleleft

In [7], Frick and Grohe obtained a matching lower bound for answering Boolean FO-queries of schema $\sigma = \{E\}$ on databases of degree at most $d := 3$ in the static setting. They used the (reasonable) complexity theoretic assumption $\text{FPT} \neq \text{AW}[*]$ and showed that if this assumption is correct, then there is no algorithm that answers Boolean FO-queries φ on σ -dbs D of degree ≤ 3 in time $2^{2^{2^{\mathcal{O}(\|\varphi\|)}}} \cdot \text{poly}(\|D\|)$ in the static setting (see Theorem 2 in [7]). As a consequence, the same lower bound holds in the dynamic setting and shows that in Theorem 4.1, the 3-fold exponential dependency on the query size $\|\varphi\|$ cannot be substantially lowered (unless $\text{FPT} = \text{AW}[*]$):

► **Corollary 4.2.** *Let $\sigma := \{E\}$ and let $d := 3$. If $\text{FPT} \neq \text{AW}[*]$, then there is no dynamic algorithm that receives a Boolean $\text{FO}[\sigma]$ -query φ and a σ -db D_0 , and computes within $t_p \leq f(\varphi) \cdot \text{poly}(\|D_0\|)$ preprocessing time a data structure that can be updated in time $t_u \leq f(\varphi)$ and allows to return the query result $\varphi(D)$ with answer time $t_a \leq f(\varphi)$, for a function f with $f(\varphi) = 2^{2^{2^{\mathcal{O}(\|\varphi\|)}}}$.*

5 Technical Lemmas on Types and Spheres Useful for Handling Non-Boolean Queries

For our algorithms for evaluating non-Boolean queries it will be convenient to work with a fixed list of representatives of d -bounded r -types, provided by the following straightforward lemma.

► **Lemma 5.1.** *There is an algorithm which upon input of a schema σ , a degree bound $d \geq 2$, a radius $r \geq 0$, and a number $k \geq 1$, computes a list $\mathcal{L}_r^{\sigma,d}(k) = \tau_1, \dots, \tau_\ell$ (for a suitable $\ell \geq 1$) of d -bounded r -types with k centres (over σ), such that for every d -bounded r -type τ with k centres (over σ) there is exactly one $i \in [\ell]$ such that $\tau \cong \tau_i$. The algorithm's runtime is $2^{(kd^{r+1})^{\mathcal{O}(\|\sigma\|)}}$. Furthermore, upon input of a d -bounded r -type τ with k centres (over σ), the particular $i \in [\ell]$ with $\tau \cong \tau_i$ can be computed in time $2^{(kd^{r+1})^{\mathcal{O}(\|\sigma\|)}}$.*

Throughout the remainder of this paper, $\mathcal{L}_r^{\sigma,d}(k)$ will always denote the list provided by Lemma 5.1. The following lemma will be useful for evaluating Boolean combinations of sphere-formulas.

► **Lemma 5.2.** *Let σ be a schema, let $r \geq 0$, $k \geq 1$, $d \geq 2$, and let $\mathcal{L}_r^{\sigma,d}(k) = \tau_1, \dots, \tau_\ell$. Let $\bar{x} = (x_1, \dots, x_k)$ be a list of k pairwise distinct variables. For every Boolean combination $\psi(\bar{x})$ of d -bounded sphere-formulas of radius at most r (over σ), there is an $I \subseteq [\ell]$ such that $\psi(\bar{x}) \equiv_d \bigvee_{i \in I} \text{sph}_{\tau_i}(\bar{x})$. Furthermore, given $\psi(\bar{x})$, the set I can be computed in time $\text{poly}(\|\psi\|) \cdot 2^{(kd^{r+1})^{\mathcal{O}(\|\sigma\|)}}$.*

The lemma's proof is based on the following observations. Negations can be eliminated by the equivalence $\neg \text{sph}_{\tau_j}(\bar{x}) \equiv_d \bigvee_{i \in [\ell] \setminus \{j\}} \text{sph}_{\tau_i}(\bar{x})$. To eliminate conjunctions, observe that for $i \neq i'$ the formula $\text{sph}_{\tau_i}(\bar{x}) \wedge \text{sph}_{\tau_{i'}}(\bar{x})$ is unsatisfiable. Thus, by the distributive law we obtain for all $m \geq 1$ and all $I_1, \dots, I_m \subseteq [\ell]$ that

$$\bigwedge_{j \in [m]} \left(\bigvee_{i \in I_j} \text{sph}_{\tau_i}(\bar{x}) \right) \equiv_d \bigvee_{i_1 \in I_1} \cdots \bigvee_{i_m \in I_m} \left(\text{sph}_{\tau_{i_1}}(\bar{x}) \wedge \cdots \wedge \text{sph}_{\tau_{i_m}}(\bar{x}) \right) \equiv_d \bigvee_{i \in I} \text{sph}_{\tau_i}(\bar{x})$$

for $I := I_1 \cap \cdots \cap I_m$.

For evaluating a Boolean combination $\psi(\bar{x})$ of sphere-formulas and Hanf-sentences on a given σ -db D , an obvious approach is to first consider every Hanf-sentence χ that occurs in ψ , to check if $D \models \chi$, and replace every occurrence of χ in ψ with **true** (resp., **false**) if $D \models \chi$ (resp., $D \not\models \chi$). The resulting formula $\psi'(\bar{x})$ is then transformed into a disjunction $\psi''(\bar{x}) := \bigvee_{i \in I} \text{sph}_{\tau_i}(\bar{x})$ by Lemma 5.2, and the query result $\psi(D) = \psi''(D)$ is obtained as the union of the query results $\text{sph}_{\tau_i}(D)$ for all $i \in I$.

While this works well in the static setting (i.e., without database updates), in the dynamic setting we have to take care of the fact that database updates might change the status of a Hanf-sentence χ in ψ , i.e., an update operation might turn a database D with $D \models \chi$ into a database D' with $D' \not\models \chi$ (and vice versa). Consequently, the formula $\psi''(\bar{x})$ that is equivalent to $\psi(\bar{x})$ on D might be inequivalent to $\psi(\bar{x})$ on D' .

To handle the dynamic setting correctly, at the end of each update step we will use the following lemma (the lemma's proof is an easy consequence of Lemma 5.2).

► **Lemma 5.3.** *Let σ be a schema. Let $s \geq 0$ and let χ_1, \dots, χ_s be FO+MOD[σ]-sentences. Let $r \geq 0, k \geq 1, d \geq 2$, and let $\mathcal{L}_r^{\sigma,d}(k) = \tau_1, \dots, \tau_\ell$. Let $\bar{x} = (x_1, \dots, x_k)$ be a list of k pairwise distinct variables. For every Boolean combination $\psi(\bar{x})$ of the sentences χ_1, \dots, χ_s and of d -bounded sphere-formulas of radius at most r (over σ), and for every $J \subseteq [s]$ there is a set $I \subseteq [\ell]$ such that*

$$\psi_J(\bar{x}) \equiv_d \bigvee_{i \in I} \text{sph}_{\tau_i}(\bar{x}),$$

where ψ_J is the formula obtained from ψ by replacing every occurrence of a sentence χ_j with **true** if $j \in J$ and with **false** if $j \notin J$ (for every $j \in [s]$).

Given ψ and J , the set I can be computed in time $\text{poly}(\|\psi\|) \cdot 2^{(kd^{r+1})^{\mathcal{O}(\|\sigma\|)}}$.

To evaluate a single sphere-formula $\text{sph}_\tau(\bar{x})$ for a given r -type τ with k centres (over σ), it will be useful to decompose τ into its connected components as follows. Let $\tau = (T, \bar{t})$ with $\bar{t} = (t_1, \dots, t_k)$. Consider the Gaifman graph G^T of T and let C_1, \dots, C_c be the vertex sets of the c connected components of G^T . For each connected component C_j of G^T , let \bar{t}_j be the subsequence of \bar{t} consisting of all elements of \bar{t} that belong to C_j , and let k_j be the length of \bar{t}_j . Since (T, \bar{t}) is an r -type with the k centres, we have $T = \mathcal{N}_r^T(\bar{t})$, and thus $c \leq k$ and $k_j \geq 1$ for all $j \in [c]$. To avoid ambiguity, we make sure that the list C_1, \dots, C_c is sorted in such a way that for all $j < j'$ we have $i < i'$ for the smallest i with $t_i \in C_j$ and the smallest i' with $t_{i'} \in C_{j'}$.

For each C_j consider the r -type with k_j centres $\rho_j = (T[C_j], \bar{t}_j)$. Let ν_j be the unique integer such that ρ_j is isomorphic to the ν_j -th element in the list $\mathcal{L}_r^{\sigma,d}(k_j)$, and let τ_{j,ν_j} be the ν_j -th element in this list.

It is straightforward to see that the formula $\text{sph}_\tau(\bar{x})$ is d -equivalent to the formula

$$\text{conn-sph}_\tau(\bar{x}) := \bigwedge_{j \in [c]} \text{sph}_{\tau_{j,\nu_j}}(\bar{x}_j) \wedge \bigwedge_{j \neq j'} \neg \text{dist}_{\leq 2r+1}^{k_j, k_{j'}}(\bar{x}_j, \bar{x}_{j'}), \quad (1)$$

where \bar{x}_j is the subsequence of \bar{x} obtained from \bar{x} in the same way as \bar{t}_j is obtained from \bar{t} , and $\text{dist}_{\leq 2r+1}^{k_j, k_{j'}}(\bar{x}_j, \bar{x}_{j'})$ is a formula of schema σ which expresses that for some variable y in \bar{x}_j and some variable y' in $\bar{x}_{j'}$, the distance between y and y' is $\leq 2r+1$. I.e., for $\bar{a} = (a_1, \dots, a_{k_j})$ and $\bar{b} = (b_1, \dots, b_{k_{j'}})$ we have $(\bar{a}, \bar{b}) \in \text{dist}_{\leq 2r+1}^{k_j, k_{j'}}(D) \iff \text{dist}^D(\bar{a}, \bar{b}) \leq 2r+1$, where

$$\text{dist}^D(\bar{a}, \bar{b}) \leq 2r+1 \text{ means that } \text{dist}^D(a_i, b_{i'}) \leq 2r+1 \text{ for some } i \in [k_j] \text{ and } i' \in [k_{j'}]. \quad (2)$$

Using the Lemmas 3.2 and 5.1, the following lemma is straightforward.

► **Lemma 5.4.** *There is an algorithm which upon input of a schema σ , numbers $r \geq 0, k \geq 1$, and $d \geq 2$, and an r -type τ with k centres (over σ) computes the formula $\text{conn-sph}_\tau(\bar{x})$, along with the corresponding parameters c and $k_j, \nu_j, \bar{x}_j, \tau_{j,\nu_j}$ for all $j \in [c]$.*

The algorithm's runtime is $2^{(kd^{r+1})^{\mathcal{O}(\|\sigma\|)}}$.

We define the *signature* of τ to be the tuple $\text{sgn}(\tau)$ built from the parameters c and $(k_j, \nu_j, \{\mu \in [k] : x_\mu \text{ belongs to } \bar{x}_j\})_{j \in [c]}$ obtained from the above lemma. The signature $\text{sgn}^D(\bar{a})$ of a tuple \bar{a} in a database D (w.r.t. radius r) is defined as $\text{sgn}(\rho)$ for $\rho := (\mathcal{N}_r^D(\bar{a}), \bar{a})$. Note that $\bar{a} \in \text{sph}_\tau(D) \iff \text{sgn}^D(\bar{a}) = \text{sgn}(\tau)$.

6 Testing Non-Boolean FO+MOD Queries Under Updates

This section is devoted to the proof of the following theorem.

► **Theorem 6.1.** *There is a dynamic algorithm that receives a schema σ , a degree bound $d \geq 2$, a k -ary FO+MOD[σ]-query $\varphi(\bar{x})$ (for some $k \in \mathbb{N}$), and a σ -db D_0 of degree $\leq d$, and computes within $t_p = f(\varphi, d) \cdot \|D_0\|$ preprocessing time a data structure that can be updated in time $t_u = f(\varphi, d)$ and allows to test for any input tuple $\bar{a} \in \mathbf{dom}^k$ whether $\bar{a} \in \varphi(D)$ within testing time $t_t = \mathcal{O}(k^2)$. The function $f(\varphi, d)$ is of the form $2^{d^{2^{\mathcal{O}(\|\varphi\|)}}$.*

For the proof, we use the lemmas provided in Section 5 and the following lemma.

► **Lemma 6.2.** *There is a dynamic algorithm that receives a schema σ , a degree bound $d \geq 2$, numbers $r \geq 0$ and $k \geq 1$, an r -type τ with k centres (over σ), and a σ -db D_0 of degree $\leq d$, and computes within $t_p = 2^{(kd^{r+1})^{\mathcal{O}(\|\sigma\|)}} \cdot \|D_0\|$ preprocessing time a data structure that can be updated in time $t_u = 2^{(kd^{r+1})^{\mathcal{O}(\|\sigma\|)}}$ and allows to test for any input tuple $\bar{a} \in \mathbf{dom}^k$ whether $\bar{a} \in \text{sph}_\tau(D)$ within testing time $t_t = \mathcal{O}(k^2)$.*

Proof sketch. The preprocessing routine starts by using Lemma 5.4 to compute the formula $\text{conn-sph}_\tau(\bar{x})$, along with the according parameters c and $k_j, \nu_j, \bar{x}_j, \tau_{j, \nu_j}$ for each $j \in [c]$. This is done in time $2^{(kd^{r+1})^{\mathcal{O}(\|\sigma\|)}}$. We let $\text{sgn}(\tau)$ be the signature of τ (defined directly after Lemma 5.4). Recall that $\text{conn-sph}_\tau(\bar{x}) \equiv_d \text{sph}_\tau(\bar{x})$, and recall from equation (1) the precise definition of the formula $\text{conn-sph}_\tau(\bar{x})$. Our data structure will store the following information on the database D :

- the set Γ of all tuples $\bar{b} \in \text{adom}(D)^{k'}$ where $k' \leq k$ and $\mathcal{N}_r^D(\bar{b})$ is connected, and
- for every $j \in [c]$ and every k_j -tuple $\bar{b} \in \Gamma$, the unique number $\nu_{\bar{b}}$ such that $\rho_{\bar{b}} := (\mathcal{N}_r^D(\bar{b}), \bar{b})$ is isomorphic to the $\nu_{\bar{b}}$ -th element in the list $\mathcal{L}_r^{\sigma, d}(k_j)$.

We want to store this information in such a way that for any given tuple $\bar{b} \in \mathbf{dom}^{k'}$ it can be checked in time $\mathcal{O}(k)$ whether $\bar{b} \in \Gamma$. To ensure this, we use a k' -ary array $\Gamma_{k'}$ that is initialised to 0, and where during update operations the entry $\Gamma_{k'}[\bar{b}]$ is set to 1 for all $\bar{b} \in \Gamma$ of arity k' . In a similar way we can ensure that for any given $j \in [c]$ and any $\bar{b} \in \Gamma$ of arity k_j , the number $\nu_{\bar{b}}$ can be looked up in time $\mathcal{O}(k)$.

The **test** routine upon input of a tuple \bar{a} computes the signature $\text{sgn}^D(\bar{a})$ of \bar{a} in D , tests whether $\text{sgn}^D(\bar{a}) = \text{sgn}(\tau)$ and outputs “yes” if this is the case and “no” otherwise. Using the information stored in our data structure, all this can be done in time $\mathcal{O}(k^2)$. The bound on the update time follows from the fact that the insertion or deletion of a tuple affects only a small number of entries in the data structure. ◀

Theorem 6.1 is now obtained by combining Theorem 3.1, Lemma 6.2, Theorem 4.1, and Lemma 5.3.

Proof of Theorem 6.1. For $k = 0$, the theorem immediately follows from Theorem 4.1. Consider the case where $k \geq 1$. As in the proof of Theorem 4.1, we assume w.l.o.g. that all the symbols of σ occur in φ . We start the preprocessing routine by using Theorem 3.1 to transform $\varphi(\bar{x})$ into a d -equivalent query $\psi(\bar{x})$ in Hanf normal form; this takes time $2^{d^{2^{\mathcal{O}(\|\varphi\|)}}$. The formula ψ is a Boolean combination of d -bounded Hanf-sentences and sphere-formulas (over σ) of locality radius at most $r := 4^{\text{qr}(\varphi)}$, and each sphere-formula is of arity at most k . Let χ_1, \dots, χ_s be the list of all Hanf-sentences that occur in ψ .

We use Lemma 5.1 to compute the list $\mathcal{L}_r^{\sigma, d}(k) = \tau_1, \dots, \tau_\ell$. In parallel for each $i \in [\ell]$, we use the algorithm provided by Lemma 6.2 for $\tau := \tau_i$. Furthermore, for each $j \in [s]$, we

use the algorithm provided by Theorem 4.1 upon input of the Hanf-sentence $\varphi := \chi_j$. In addition to the components used by these dynamic algorithms, our data structure also stores

- the set $J := \{j \in [s] : D \models \chi_j\}$,
- the particular set $I \subseteq [\ell]$ provided by Lemma 5.3 for $\psi(\bar{x})$ and J , and
- the set $K = \{\text{sgn}(\tau_i) : i \in I\}$, where for each type τ , $\text{sgn}(\tau)$ is the signature of τ defined directly after Lemma 5.4.

The **test** routine upon input of a tuple $\bar{a} = (a_1, \dots, a_k)$ proceeds in the same way as in the proof of Lemma 6.2 to compute in time $\mathcal{O}(k^2)$ the signature $\text{sgn}^D(\bar{a})$ of the tuple \bar{a} . For every $i \in [\ell]$ we have $\bar{a} \in \text{sph}_{\tau_i}(D) \iff \text{sgn}^D(\bar{a}) = \text{sgn}(\tau_i)$. Thus, $\bar{a} \in \varphi(D) \iff \text{sgn}^D(\bar{a}) \in K$. Therefore, the **test** routine checks whether $\text{sgn}^D(\bar{a}) \in K$ and outputs “yes” if this is the case and “no” otherwise. To ensure that this test can be done in time $\mathcal{O}(k^2)$, we use an array construction for storing K (similar to the one for storing Γ in the proof of Lemma 6.2).

The **update** routine runs in parallel the update routines for all the used dynamic data structures. Afterwards, it recomputes J by calling the **answer** routine for χ_j for all $j \in [s]$. Then, it uses Lemma 5.3 to recompute I . The set K is then recomputed by applying Lemma 5.4 for $\tau := \tau_i$ for all $i \in I$. It is straightforward to see that the overall runtime of the **update** routine is $t_u = 2^{d^2 \mathcal{O}(\|\sigma\|)}$. This completes the proof of Theorem 6.1. ◀

7 Representing Databases by Coloured Graphs

To obtain dynamic algorithms for counting and enumerating query results, it will be convenient to work with a representation of databases by coloured graphs that is similar to the representation used in [6]. For defining this representation, let us consider a fixed d -bounded r -type τ with k centres (over a schema σ). Use Lemma 5.4 to compute the formula $\text{conn-sph}_\tau(\bar{x})$ (for $\bar{x} = (x_1, \dots, x_k)$) and the according parameters c and $k_j, \nu_j, \bar{x}_j, \tau_j, \nu_j$, and let $\text{sgn}(\tau)$ be the signature of τ . To keep the notation simple, we assume w.l.o.g. that $\bar{x}_1 = x_1, \dots, x_{k_1}$, $\bar{x}_2 = x_{k_1+1}, \dots, x_{k_1+k_2}$ etc.

Recall that $\text{sph}_\tau(\bar{x})$ is d -equivalent to the formula

$$\text{conn-sph}_\tau(\bar{x}) := \bigwedge_{j \in [c]} \text{sph}_{\tau_j, \nu_j}(\bar{x}_j) \wedge \bigwedge_{j \neq j'} \neg \text{dist}_{\leq 2r+1}^{k_j, k_{j'}}(\bar{x}_j, \bar{x}_{j'}).$$

To count or enumerate the results of the formula $\text{sph}_\tau(\bar{x})$ we represent the database D by a c -coloured graph \mathcal{G}_D . Here, a c -coloured graph \mathcal{G} is a database of the particular schema

$$\sigma_c := \{E, C_1, \dots, C_c\},$$

where E is a binary relation symbol and C_1, \dots, C_c are unary relation symbols. We define \mathcal{G}_D in such a way that the task of counting or enumerating the results of the query $\text{sph}_\tau(\bar{x})$ on the database D can be reduced to counting or enumerating the results of the query

$$\varphi_c(z_1, \dots, z_c) := \bigwedge_{j \in [c]} C_j(z_j) \wedge \bigwedge_{j \neq j'} \neg E(z_j, z_{j'}) \quad (3)$$

on the c -coloured graph \mathcal{G}_D . The vertices of \mathcal{G}_D correspond to tuples over $\text{adom}(D)$ whose r -neighbourhood is connected; a vertex has colour C_j if its associated tuple \bar{a} is in $\text{sph}_{\tau_j, \nu_j}(D)$; and an edge between two vertices indicates that $\text{dist}^D(\bar{a}; \bar{b}) \leq 2r+1$, for their associated tuples \bar{a} and \bar{b} . The following lemma allows to translate a dynamic algorithm for counting or enumerating the results of the query $\varphi_c(z_1, \dots, z_c)$ on c -coloured graphs into a dynamic algorithm for counting or enumerating the result of the query $\text{sph}_\tau(\bar{x})$ on D .

► **Lemma 7.1.** *Suppose that the counting problem (the enumeration problem) for $\varphi_c(\bar{z})$ on σ_c -dbs of degree at most d' can be solved by a dynamic algorithm with initialisation time $t_i(c, d')$, update time $t_u(c, d')$, and counting time $t_c(c, d')$ (delay $t_d(c, d')$). Then for every schema σ and every $d \geq 2$ the following holds.*

(1) *Let $r \geq 0$, $k \geq 1$, τ a d -bounded r -type with k centres, and fix $d' := d^{2k^2(2r+1)}$ and $\hat{t}_x := \max_{c=1}^k t_x(c, d')$ for $t_x \in \{t_i, t_u, t_c, t_d\}$. The counting problem (the enumeration problem) for $\text{sph}_\tau(\bar{x})$ on σ -dbs of degree at most d can be solved by a dynamic algorithm with counting time \hat{t}_c (delay $\mathcal{O}(\hat{t}_d k)$), update time $t'_u \leq \hat{t}_u d^{\mathcal{O}(k^2 r + k \|\sigma\|)} + 2^{\mathcal{O}(\|\sigma\| k^2 d^{2r+2})}$, and initialisation time \hat{t}_i .*

(2) *The counting problem (the enumeration problem) for k -ary FO+MOD-queries $\varphi(\bar{x})$ on σ -dbs of degree at most d can be solved with counting time $\mathcal{O}(1)$ (delay $\mathcal{O}(\hat{t}_d k)$), update time $(\hat{t}_u + \hat{t}_c) 2^{d^{2^{\mathcal{O}(\|\varphi\|)}}$, and initialisation time $\hat{t}_i 2^{d^{2^{\mathcal{O}(\|\varphi\|)}}$ where $\hat{t}_x = \max_{c=1}^k t_x(c, d^{2^{\mathcal{O}(\|\varphi\|)}}$ for $t_x \in \{t_i, t_u, t_c, t_d\}$.*

Proof sketch. The first part is a simple reduction from $\text{conn-sph}_\tau(\bar{x})$ to φ_c and can be found in the full version of the paper. The second part for $k = 0$ follows immediately from Theorem 4.1. Consider the case where $k \geq 1$. W.l.o.g. we assume that all the symbols of σ occur in φ (otherwise, we remove from σ all symbols that do not occur in φ). We start the preprocessing routine by using Theorem 3.1 to transform $\varphi(\bar{x})$ into a d -equivalent query $\psi(\bar{x})$ in Hanf normal form; this takes time $2^{d^{2^{\mathcal{O}(\|\varphi\|)}}$. The formula ψ is a Boolean combination of d -bounded Hanf-sentences and sphere-formulas (over σ) of locality radius at most $r := 4^{\text{qr}(\varphi)}$, and each sphere-formula is of arity at most k . Note that for $d' := d^{2k^2(2r+1)}$ as used in the first part it holds that $d' = d^{2^{\mathcal{O}(\|\varphi\|)}}$. Let χ_1, \dots, χ_s be the list of all Hanf-sentences that occur in ψ (recall that $s \leq 2^{d^{2^{\mathcal{O}(\|\varphi\|)}}$).

We use Lemma 5.1 to compute the list $\mathcal{L}_r^{\sigma, d}(k) = \tau_1, \dots, \tau_\ell$ (note that $\ell \leq 2^{d^{2^{\mathcal{O}(\|\varphi\|)}}$). In parallel for each $i \in [\ell]$, we use the dynamic algorithm for $\text{sph}_{\tau_i}(\bar{x})$ provided from the lemma's part (1). Furthermore, for each $j \in [s]$, we use the dynamic algorithm provided by Theorem 4.1 upon input of the Hanf-sentence $\varphi := \chi_j$. In addition to the components used by these dynamic algorithms, our data structure also stores

- the set $J := \{j \in [s] : D \models \chi_j\}$,
- the particular set $I \subseteq [\ell]$ provided by Lemma 5.3 for $\psi(\bar{x})$ and J , and
- the cardinality $n = |\varphi(D)|$ of the query result.

The **count** routine simply outputs the value n in time $\mathcal{O}(1)$. The **enumerate** routine runs the **enumerate** routine on $\text{sph}_{\tau_i}(D)$ for every $i \in I$. Note that this enumerates, without repetition, all tuples in $\varphi(D)$, because by Lemma 5.3, $\varphi(D)$ is the union of the sets $\text{sph}_{\tau_i}(D)$ for all $i \in I$, and this is a union of pairwise disjoint sets. The **update** routine runs in parallel the update routines for all used dynamic data structures. Afterwards, it recomputes J by calling the **answer** routine for χ_j for all $j \in [s]$. Then, it uses Lemma 5.3 to recompute I . The number n is then recomputed by letting $n = \sum_{i \in I} n_i$, where n_i is the result of the **count** routine for τ_i . It is straightforward to verify that the overall runtime of the **update** routine is bounded by $(\hat{t}_u + \hat{t}_c) 2^{d^{2^{\mathcal{O}(\|\varphi\|)}}$. ◀

8

 Counting Results of FO+MOD Queries Under Updates

This section is devoted to the proof of the following theorem.

► **Theorem 8.1.** *There is a dynamic algorithm that receives a schema σ , a degree bound $d \geq 2$, a k -ary FO+MOD[σ]-query $\varphi(\bar{x})$ (for some $k \in \mathbb{N}$), and a σ -db D_0 of degree $\leq d$, and computes within $t_p = f(\varphi, d) \cdot \|D_0\|$ preprocessing time a data structure that can be updated in time $t_u = f(\varphi, d)$ and allows to return the cardinality $|\varphi(D)|$ of the query result within time $\mathcal{O}(1)$. The function $f(\varphi, d)$ is of the form $2^{d^{2^{\mathcal{O}(\|\varphi\|)}}$.*

The theorem follows immediately from Lemma 7.1 and the following dynamic counting algorithm for the query $\varphi_c(\bar{z})$.

► **Lemma 8.2.** *There is a dynamic algorithm that receives a number $c \geq 1$, a degree bound $d \geq 2$, and a σ_c -db \mathcal{G}_0 of degree $\leq d$, and computes $|\varphi_c(\mathcal{G})|$ with $d^{\mathcal{O}(c^2)}$ initialisation time, $\mathcal{O}(1)$ counting time, and $d^{\mathcal{O}(c^2)}$ update time.*

Proof. Recall that $\varphi_c(z_1, \dots, z_c) = \bigwedge_{i \in [c]} C_i(z_i) \wedge \bigwedge_{j \neq j'} \neg E(z_j, z_{j'})$. For all $j, j' \in [c]$ with $j \neq j'$ consider the formula $\theta_{j,j'}(z_1, \dots, z_c) := E(z_j, z_{j'}) \wedge \bigwedge_{i \in [c]} C_i(z_i)$. Furthermore, let $\alpha(z_1, \dots, z_c) := \bigwedge_{i \in [c]} C_i(z_i)$. Clearly, for every σ_c -db \mathcal{G} we have

$$\begin{aligned} \alpha(\mathcal{G}) &= C_1^{\mathcal{G}} \times \dots \times C_c^{\mathcal{G}}, \\ \varphi_c(\mathcal{G}) &= \alpha(\mathcal{G}) \setminus \left(\bigcup_{j \neq j'} \theta_{j,j'}(\mathcal{G}) \right), \quad \text{and hence,} \quad |\varphi_c(\mathcal{G})| = |\alpha(\mathcal{G})| - \left| \bigcup_{j \neq j'} \theta_{j,j'}(\mathcal{G}) \right|. \end{aligned}$$

By the *inclusion-exclusion principle* we obtain for $J := \{(j, j') : j, j' \in [c], j \neq j'\}$ that

$$\left| \bigcup_{j \neq j'} \theta_{j,j'}(\mathcal{G}) \right| = \sum_{\emptyset \neq K \subseteq J} (-1)^{|K|-1} \left| \bigcap_{(j,j') \in K} \theta_{j,j'}(\mathcal{G}) \right| = \sum_{\emptyset \neq K \subseteq J} (-1)^{|K|-1} |\varphi_K(\mathcal{G})|$$

for the formula $\varphi_K(z_1, \dots, z_c) := \bigwedge_{i \in [c]} C_i(z_i) \wedge \bigwedge_{(j,j') \in K} E(z_j, z_{j'})$.

Our data structure stores the following values:

- $|C_i^{\mathcal{G}}|$, for each $i \in [c]$, and $n_1 := |\alpha(\mathcal{G})| = \prod_{i \in [c]} |C_i^{\mathcal{G}}|$,
- $|\varphi_K(\mathcal{G})|$, for each $K \subseteq J$ with $K \neq \emptyset$, and
- $n_2 := \sum_{\emptyset \neq K \subseteq J} (-1)^{|K|-1} |\varphi_K(\mathcal{G})|$ and $n_3 := n_1 - n_2$.

Note that $n_3 = |\varphi_c(\mathcal{G})|$ is the desired size of the query result. Therefore, the **count** routine can answer in time $\mathcal{O}(1)$ by just outputting the number n_3 .

It remains to show how these values can be initialised and updated during updates of \mathcal{G} . The initialisation for the empty graph initialises all the values to 0. In the **update** routine, the values for $|C_i^{\mathcal{G}}|$ and n_1 can be updated in a straightforward way (using time $\mathcal{O}(c)$). For each $K \subseteq J$, the update of $|\varphi_K(\mathcal{G})|$ is provided within time $d^{\mathcal{O}(c^2)}$ by the following Claim 8.3, whose proof can be found in the full version of the paper.

► **Claim 8.3.** *For every $K \subseteq J$, the cardinality $|\varphi_K(\mathcal{G})|$ of a σ_c -db \mathcal{G} of degree at most d can be updated within time $d^{\mathcal{O}(c^2)}$ after $d^{\mathcal{O}(c^2)} \cdot |\mathcal{G}_0|$ preprocessing time.*

Once we have available the updated numbers $|\varphi_K(\mathcal{G})|$ for all $K \subseteq J$, the value n_2 can be computed in time $\mathcal{O}(|2^J|) \leq 2^{\mathcal{O}(c^2)}$. And n_3 is then obtained in time $\mathcal{O}(1)$. Altogether, performing the **update** routine takes time at most $d^{\mathcal{O}(c^2)}$. The **preprocess** routine initialises all values for the empty graph and then uses $|\mathcal{G}_0|$ update steps to insert all the tuples of \mathcal{G}_0 into the data structure. This completes the proof of Lemma 8.2. ◀

9 Enumerating Results of FO+MOD Queries Under Updates

In this section we prove (and afterwards, improve) the following theorem.

► **Theorem 9.1.** *There is a dynamic algorithm that receives a schema σ , a degree bound $d \geq 2$, a k -ary FO+MOD[σ]-query $\varphi(\bar{x})$ (for some $k \in \mathbb{N}$), and a σ -db D_0 of degree $\leq d$, and computes within $t_p = f(\varphi, d) \cdot \|D_0\|$ preprocessing time a data structure that can be updated in time $t_u = f(\varphi, d)$ and allows to enumerate $\varphi(D)$ with $d^{2^{\mathcal{O}(\|\varphi\|)}}$ delay.*

The function $f(\varphi, d)$ is of the form $2^{d^{2^{\mathcal{O}(\|\varphi\|)}}$.

The theorem follows immediately from Lemma 7.1 and the following dynamic enumeration algorithm for the query $\varphi_c(\bar{z})$.

► **Lemma 9.2.** *There is a dynamic algorithm that receives a number $c \geq 1$, a degree bound $d \geq 2$, and a σ_c -db \mathcal{G} of degree $\leq d$, and computes within $t_p = d^{\text{poly}(c)} \cdot |\mathcal{G}_0|$ preprocessing time a data structure that can be updated in time $d^{\text{poly}(c)}$ and allows to enumerate the query result $\varphi_c(\mathcal{G})$ with $\mathcal{O}(c^3 d)$ delay.*

Proof. For a σ_c -db \mathcal{G} and a vertex $v \in \text{adom}(\mathcal{G})$ we let $N^{\mathcal{G}}(v)$ be the set of all neighbours of v in \mathcal{G} . I.e., $N^{\mathcal{G}}(v)$ is the set of all $w \in \text{adom}(\mathcal{G})$ such that (v, w) or (w, v) belongs to $E^{\mathcal{G}}$.

The underlying idea of the enumeration procedure is the following greedy strategy. We cycle through all vertices $u_1 \in C_1^{\mathcal{G}}$, $u_2 \in C_2^{\mathcal{G}} \setminus N^{\mathcal{G}}(u_1)$, $u_3 \in C_3^{\mathcal{G}} \setminus (N^{\mathcal{G}}(u_1) \cup N^{\mathcal{G}}(u_2))$, \dots , $u_c \in C_c^{\mathcal{G}} \setminus \bigcup_{i \leq c-1} N^{\mathcal{G}}(u_i)$ and output (u_1, \dots, u_c) . This strategy does not yet lead to a constant delay enumeration, as there might be vertex tuples (u_1, \dots, u_i) (for $i < c$) that do extend to an output tuple (u_1, \dots, u_c) , but where many possible extensions are checked before this output tuple is encountered. We now show how to overcome this problem and describe an enumeration procedure with $\mathcal{O}(c^3 d)$ delay and update time $d^{\text{poly}(c)}$.

Note that for every $J \subseteq [c]$ we have $|\bigcup_{j \in J} N^{\mathcal{G}}(u_j)| \leq cd$. Hence, if a set $C_i^{\mathcal{G}}$ contains more than cd elements, we know that every considered tuple has an extension $u_i \in C_i^{\mathcal{G}}$ that is not a neighbour of any vertex in the tuple. Let $I := \{i \in [c] : |C_i^{\mathcal{G}}| \leq cd\}$ be the set of *small* colour classes in \mathcal{G} and to simplify the presentation we assume without loss of generality that $I = \{1, \dots, s\}$. In our data structure we store the current index set I and the set

$$\mathcal{S} := \left\{ (u_1, \dots, u_s) \in C_1^{\mathcal{G}} \times \dots \times C_s^{\mathcal{G}} : (u_j, u_{j'}) \notin E^{\mathcal{G}}, \text{ for all } j \neq j' \right\} \quad (4)$$

of tuples on the small colours. Note that a tuple $(u_1, \dots, u_s) \in C_1^{\mathcal{G}} \times \dots \times C_s^{\mathcal{G}}$ extends to an output tuple $(u_1, \dots, u_c) \in \varphi_c(\mathcal{G})$ if and only if it is contained in \mathcal{S} . As $|\mathcal{S}| \leq (cd)^c$, it is not hard to see that we can recompute the sets I and \mathcal{S} in time $d^{\text{poly}(c)}$ after every update. The enumeration procedure is given in Algorithm 1.

It is straightforward to see that this procedure enumerates $\varphi_c(\mathcal{G})$. Let us analyse the delay. Since for all $i > s$ we have that $|C_i^{\mathcal{G}}| > cd$, it follows that every call of $\text{ENUM}(u_1, \dots, u_i)$ leads to at least one recursive call of $\text{ENUM}(u_1, \dots, u_i, u_{i+1})$. Furthermore, there are at most cd iterations of the loop in line 7 that do *not* lead to a recursive call. As every test in line 8 can be done in time $\mathcal{O}(c)$, it follows that the time spans until the first recursive call, between the calls, and after the last call are bounded by $\mathcal{O}(c^2 d)$. As the recursion depth is c , the overall delay between two output tuples is bounded by $\mathcal{O}(c^3 d)$. ◀

By using similar techniques as in [6], we can obtain the following improved version of Lemma 9.2 where the delay is independent of the degree bound d .

Algorithm 1 Enumeration procedure.

```

1: for all  $(u_1, \dots, u_s) \in \mathcal{S}$  do ENUM( $u_1, \dots, u_s$ ).
2: Output the end-of-enumeration message EOE.
3:
4: function ENUM( $u_1, \dots, u_i$ )
5:   if  $i = c$  then output the tuple  $(u_1, \dots, u_c)$ .
6:   else
7:     for all  $u_{i+1} \in C_{i+1}^{\mathcal{G}}$  do
8:       if  $u_{i+1} \notin \bigcup_{j=1}^i N^{\mathcal{G}}(u_j)$  then ENUM( $u_1, \dots, u_i, u_{i+1}$ ).

```

► **Lemma 9.3.** *There is a dynamic algorithm that receives a number $c \geq 1$, a degree bound $d \geq 2$, and a σ_c -db \mathcal{G}_0 of degree $\leq d$, and computes within $t_p = d^{\text{poly}(c)} \cdot |\mathcal{G}_0|$ preprocessing time a data structure that can be updated in time $d^{\text{poly}(c)}$ and allows to enumerate the query result $\varphi_c(\mathcal{G})$ with $\mathcal{O}(c^2)$ delay.*

Proof idea. We proceed in a similar way as in the proof of Lemma 9.2. But in order to enumerate the tuples with only $\mathcal{O}(c^2)$ delay, we replace the loop in lines 7–8 of Algorithm 1 by a precomputed “skip” function that allows to iterate through all elements in $C_{i+1}^{\mathcal{G}} \setminus \bigcup_{j=1}^i N^{\mathcal{G}}(u_j)$ with $\mathcal{O}(c)$ delay. This technique has been introduced for static databases in [6]. It turns out that it is possible to maintain the additional information with $d^{\text{poly}(c)}$ update time. For details we refer the reader to the full version of the paper. ◀

By Lemma 7.1, this directly improves the delay in Theorem 9.1 from $d^{2^{\mathcal{O}(\|\varphi\|)}}$ to $\mathcal{O}(k^3)$ and leads to the following theorem.

► **Theorem 9.4.** *There is a dynamic algorithm that receives a schema σ , a degree bound $d \geq 2$, a k -ary FO+MOD[σ]-query $\varphi(\bar{x})$ (for some $k \in \mathbb{N}$), and a σ -db D_0 of degree $\leq d$, and computes within $t_p = f(\varphi, d) \cdot \|D_0\|$ preprocessing time a data structure that can be updated in time $t_u = f(\varphi, d)$ and allows to enumerate $\varphi(D)$ with $\mathcal{O}(k^3)$ delay. The function $f(\varphi, d)$ is of the form $2^{d^{2^{\mathcal{O}(\|\varphi\|)}}$.*

10 Conclusion

Our main results show that in the dynamic setting (i.e., allowing database updates), the results of k -ary FO+MOD-queries on bounded degree databases can be tested and counted in constant time and enumerated with constant delay, after linear time preprocessing and with constant update time. Here, “constant time” refers to data complexity and is of size $\text{poly}(k)$ concerning the delay and the time for testing and counting. The time for performing a database update is 3-fold exponential in the size of the query and the degree bound, and is worst-case optimal.

The starting point of our algorithms is to decompose the given query into a query in Hanf normal form, using a recent result of [10]. This normal form is only available for the setting with a fixed maximum degree bound d , i.e., the setting considered in this paper.

Recently, Kuske and Schweikardt [13] introduced a new kind of Hanf normal form for a variant of *first-order logic with counting* that contains and extends Libkin’s logic FO(Cnt) [14] and Grohe’s logic FO+C [8]. As an application it is shown in [13] that the present paper’s techniques can be lifted from FO+MOD to full first-order logic with counting.

An obvious future task is to investigate to which extent further query evaluation results that are known for the static setting can be lifted to the dynamic setting. More specifically: Are there efficient dynamic algorithms for evaluating (i.e., answering, testing, counting, or enumerating) results of first-order queries on other sparse classes of databases (e.g. planar, bounded treewidth, bounded expansion, nowhere dense) or databases of low degree, lifting the “static” results accumulated in [12, 9, 6] to the dynamic setting?

References

- 1 Serge Abiteboul, Richard Hull, and Victor Vianu. *Foundations of Databases*. Addison-Wesley, 1995. URL: <http://webdam.inria.fr/Alice/>.
- 2 Christoph Berkholz, Jens Keppeler, and Nicole Schweikardt. Answering conjunctive queries under updates. In *Proceedings of the 36th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS'17*, 2017. To appear.
- 3 Christoph Berkholz, Jens Keppeler, and Nicole Schweikardt. Answering FO+MOD queries under updates on bounded degree databases. *CoRR*, abs/1702.08764, 2017. URL: <http://arxiv.org/abs/1702.08764>.
- 4 Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms (3. ed.)*. MIT Press, 2009. URL: <http://mitpress.mit.edu/books/introduction-algorithms>.
- 5 Arnaud Durand and Etienne Grandjean. First-order queries on structures of bounded degree are computable with constant delay. *ACM Trans. Comput. Log.*, 8(4), 2007. doi:10.1145/1276920.1276923.
- 6 Arnaud Durand, Nicole Schweikardt, and Luc Segoufin. Enumerating answers to first-order queries over databases of low degree. In *Proceedings of the 33rd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS'14, Snowbird, UT, USA, June 22-27, 2014*, pages 121–131, 2014. doi:10.1145/2594538.2594539.
- 7 Markus Frick and Martin Grohe. The complexity of first-order and monadic second-order logic revisited. *Ann. Pure Appl. Logic*, 130(1-3):3–31, 2004. doi:10.1016/j.apal.2004.01.007.
- 8 Martin Grohe. *Descriptive Complexity, Canonisation, and Definable Graph Structure Theory*. Lecture Notes in Logic. Association for Symbolic Logic in conjunction with Cambridge University Press, to appear. Preliminary version available at <https://www.lii.rwth-aachen.de/de/13-mitarbeiter/professoren/39-book-descriptive-complexity.html>.
- 9 Martin Grohe, Stephan Kreutzer, and Sebastian Siebertz. Deciding first-order properties of nowhere dense graphs. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 89–98, 2014. doi:10.1145/2591796.2591851.
- 10 Lucas Heimberg, Dietrich Kuske, and Nicole Schweikardt. Hanf normal form for first-order logic with unary counting quantifiers. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS'16, New York, NY, USA, July 5-8, 2016*, pages 277–286, 2016. doi:10.1145/2933575.2934571.
- 11 Wojciech Kazana and Luc Segoufin. First-order query evaluation on structures of bounded degree. *Logical Methods in Computer Science*, 7(2), 2011. doi:10.2168/LMCS-7(2:20)2011.
- 12 Wojciech Kazana and Luc Segoufin. Enumeration of first-order queries on classes of structures with bounded expansion. In *Proceedings of the 32nd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS 2013, New York, NY, USA, June 22-27, 2013*, pages 297–308, 2013. doi:10.1145/2463664.2463667.

- 13 Dietrich Kuske and Nicole Schweikardt. First-order logic with counting: At least, *weak* Hanf normal forms always exist and can be computed! Unpublished manuscript, 2017.
- 14 Leonid Libkin. *Elements of Finite Model Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2004. doi:10.1007/978-3-662-07003-1.
- 15 Eugene M. Luks. Isomorphism of graphs of bounded valence can be tested in polynomial time. *J. Comput. Syst. Sci.*, 25(1):42–65, 1982. doi:10.1016/0022-0000(82)90009-5.
- 16 Bernard M. E. Moret and Henry D. Shapiro. *Algorithms from P to NP: Volume 1: Design & Efficiency*. Benjamin-Cummings, 1991.
- 17 Sushant Patnaik and Neil Immerman. Dyn-FO: A parallel, dynamic complexity class. *J. Comput. Syst. Sci.*, 55(2):199–209, 1997. doi:10.1006/jcss.1997.1520.
- 18 Thomas Schwentick and Thomas Zeume. Dynamic complexity: recent updates. *SIGLOG News*, 3(2):30–52, 2016. doi:10.1145/2948896.2948899.
- 19 Detlef Seese. Linear time computable problems and first-order descriptions. *Mathematical Structures in Computer Science*, 6(6):505–526, 1996.