

A STRATEGY FOR DYNAMIC PROGRAMS: START OVER AND MUDDLE THROUGH

SAMIR DATTA ^a, ANISH MUKHERJEE ^b, THOMAS SCHWENTICK ^c, NILS VORTMEIER ^c,
AND THOMAS ZEUME ^c

^a Chennai Mathematical Institute & UMI ReLaX, Chennai, India
e-mail address: sdatta@cmi.ac.in

^b Chennai Mathematical Institute, Chennai, India
e-mail address: anish@cmi.ac.in

^c TU Dortmund University, Dortmund, Germany
e-mail address: {thomas.schwentick, nils.vortmeier, thomas.zeume}@tu-dortmund.de

ABSTRACT. In the setting of DYNFO, dynamic programs update the stored result of a query whenever the underlying data changes. This update is expressed in terms of first-order logic. We introduce a strategy for constructing dynamic programs that utilises periodic computation of auxiliary data from scratch and the ability to maintain a query for a limited number of change steps. We show that if some program can maintain a query for $\log n$ change steps after an AC^1 -computable initialisation, it can be maintained by a first-order dynamic program as well, i.e., in DYNFO. As an application, it is shown that decision and optimisation problems defined by monadic second-order (MSO) formulas are in DYNFO, if only change sequences that produce graphs of bounded treewidth are allowed. To establish this result, a Feferman-Vaught-type composition theorem for MSO is established that might be useful in its own right.

1. INTRODUCTION

Each time a database is changed, any previously computed and stored result of a fixed query might become outdated. However, when the change is small, it is plausible that the new query result is highly related to the old result. In that case it might be more efficient to use previously computed information for obtaining the new answer to the query instead of recomputing the query result from scratch. A theoretical framework for studying when the result of a query over relational databases can be updated in a declarative fashion was formalised by Patnaik and Immerman [PI97], and Dong, Su, and Topor [DST95]. In their

Key words and phrases: dynamic complexity, treewidth, monadic second order logic.

This article is the full version of [DMS⁺17]. The authors acknowledge the financial support by the DAAD-DST grant “Exploration of New Frontiers in Dynamic Complexity”. The first and the second authors were partially funded by a grant from Infosys foundation. The second author was partially supported by a TCS PhD fellowship. The last three authors acknowledge the financial support by DFG grant SCHW 678/6-2 on “Dynamic Expressiveness of Logics”.

formalisation, a dynamic program consists of a set of logical formulas that update a query result after the insertion or deletion of a tuple. The formulas may use additional auxiliary relations that, of course, need to be updated as well. The queries maintainable in this way via first-order formulas constitute the dynamic complexity class DYNFO.

Recent work has confirmed that DYNFO is a quite powerful class, since it captures, e.g., the reachability query for directed graphs [DKM⁺18], and even allows for more complex change operations than single-tuple changes [SVZ18, DMVZ18].

In this article we introduce a general strategy for dynamic programs that further underscores the expressive power of DYNFO. All prior results for DYNFO yield dynamic programs that are able to maintain a query for arbitrary long sequences of changes. Even if this is not (known to be) possible for a given query, in a practical scenario it might still be favourable to maintain the query result dynamically for a bounded number of changes, then to apply a more complex algorithm that recomputes certain auxiliary information from scratch, such that afterwards the query can again be maintained for some time, and so on.

Here we formalise this approach. Let \mathcal{C} be a complexity class and $f: \mathbb{N} \rightarrow \mathbb{N}$ a function. A query q is called (\mathcal{C}, f) -maintainable, if there is a dynamic program (with first-order definable updates) that, starting from some input structure \mathcal{A} and auxiliary relations computed in \mathcal{C} from \mathcal{A} , can answer q for $f(|\mathcal{A}|)$ many steps, where $|\mathcal{A}|$ denotes the size of the universe of \mathcal{A} .

We feel that this notion might be interesting in its own right. However, in this article we concentrate on the case where \mathcal{C} is (uniform) AC^1 and $f(n) = \log n$. The class AC^1 contains all queries that can be computed by a (uniform) circuit of depth $\mathcal{O}(\log n)$ that uses polynomially many \wedge -, \vee -, and \neg -gates, where \wedge - and \vee -gates may have unbounded fan-in. We show that $(AC^1, \log n)$ -maintainable queries are actually in DYNFO, and thus can be maintained for arbitrary long change sequences.

We apply this insight to show that all queries and optimisation problems definable in monadic second-order logic (MSO) are in DYNFO for (classes of) structures of bounded treewidth, by proving that they are $(AC^1, \log n)$ -maintainable. The same can be said about guarded second-order logic (GSO), simply because it is expressively equivalent to MSO on such classes [Cou94]. This implies that decision problems like 3-COLOURABILITY or HAMILTONCYCLE as well as optimisation problems like VERTEXCOVER and DOMINATINGSET are in DYNFO, for such classes of structures. This result is therefore a dynamic version of Courcelle’s Theorem which states that all problems definable in (certain extensions of) MSO can statically be solved in linear time for graphs of bounded treewidth [Cou90].

The proof that MSO-definable queries are $(AC^1, \log n)$ -maintainable on structures of bounded treewidth makes use of a Feferman–Vaught-type composition theorem for MSO which might be useful in its own right.

The result that $(AC^1, \log n)$ -maintainable queries are in DYNFO comes with a technical restriction: in a nutshell, it holds for queries that are invariant under insertion of (many) isolated elements. We call such queries *almost domain independent* and refer to Section 3 for a precise definition.

We emphasise that the main technical challenge in maintaining MSO-queries on graphs of bounded treewidth is that tree decompositions might change drastically after an edge insertion, and can therefore not be maintained incrementally in any obvious way. In particular, the result does not simply follow from the DYNFO-maintainability of regular tree languages shown in [GMS12]. We circumvent this problem by periodically recomputing a new tree decomposition (this can be done in logarithmic space [EJT10] and thus in AC^1) and by

showing that MSO-queries can be maintained for $\mathcal{O}(\log n)$ many change operations, even if they make the tree decomposition invalid.

Contributions. We briefly summarise the contributions described above. In this article, we introduce the notion of (\mathcal{C}, f) -maintainability and show that, amongst others, (almost domain independent) $(AC^1, \log n)$ -maintainable queries are in DYNFO. We show that MSO-definable decision problems and optimisation problems are $(AC^1, \log n)$ -maintainable and therefore in DYNFO, for structures of bounded treewidth. These proofs make use of a Feferman–Vaught-type composition theorem for MSO logic.

Related work. The simulation-based technique for proving that $(AC^1, \log n)$ -maintainable queries are in DYNFO is inspired by proof techniques from [DKM⁺15] and [SVZ18]. As mentioned above, in [GMS12] it has been shown that tree languages, i.e. MSO on trees, can be maintained in DYNFO. Independently, the maintenance of MSO definable queries on graphs of bounded treewidth is also studied in [BJM17], though in the restricted setting where the tree decomposition stays the same for all changes. A static but parallel version of Courcelle’s Theorem is given in [EJT10]: every MSO-definable problem for graphs of bounded treewidth can be solved with logarithmic space.

Organisation. Basic terminology is recalled in Section 2, followed by a short introduction into dynamic complexity in Section 3. In Section 4 we introduce the notion of (\mathcal{C}, f) -maintainability and show that $(AC^1, \log n)$ -maintainable queries are in DYNFO. A glimpse on the proof techniques for proving that MSO queries are in DYNFO for graphs of bounded treewidth is given in Section 5 via the example 3-COLOURABILITY. The proof of the general results is presented in Section 6. An extension to optimisation problems can be found in Section 7.

This article is the full version of [DMS⁺17].

2. PRELIMINARIES

We now introduce some notation and notions regarding logics, graph theory and complexity theory. We assume familiarity with first-order logic FO and other notions from finite model theory [Imm99, Lib04].

Relational structures. In this article we consider finite relational structures over relational signatures $\Sigma = \{R_1, \dots, R_\ell, c_1, \dots, c_m\}$, where each R_i is a relation symbol with a corresponding arity $\text{Ar}(R_i)$, and each c_j is a constant symbol. A Σ -structure \mathcal{A} consists of a finite *domain* A , also called the *universe* of \mathcal{A} , a relation $R_i^{\mathcal{A}} \subseteq A^{\text{Ar}(R_i)}$, and a constant $c_j^{\mathcal{A}} \in A$, for each $i \in \{1, \dots, \ell\}, j \in \{1, \dots, m\}$. The *active domain* $\text{adom}(\mathcal{A})$ of a structure \mathcal{A} contains all elements used in some tuple or as some constant of \mathcal{A} . For a set $B \subseteq A$ that contains all constants and a relation R , the restriction $R[B]$ of R to B is the relation $R \cap B^{\text{Ar}(R)}$. The structure $\mathcal{A}[B]$ induced by B is the structure obtained from \mathcal{A} by restricting the domain and all relations to B .

Sometimes, especially in Section 3, we consider relational structures as relational databases. This terminology is common in the context of dynamic complexity due to its original motivation from relational databases. Also for this reason, dynamic complexity

classes defined later will be defined as classes of queries of arbitrary arity, and not as a class of decision problems. However, we will mostly consider queries over structures with a single binary relation symbol E , that is, queries on graphs.

We will often use structures \mathcal{A} with a linear order \leq on the universe A , and compatible ternary relations encoding arithmetical operations $+$ and \times or, alternatively, a binary relation encoding the relation $\text{BIT} = \{(i, j) \mid \text{the } j\text{-th bit in the binary representation of } i \text{ is } 1\}$. The linear order in particular allows us to identify A with the first $|A|$ natural numbers. We write $\text{FO}(+, \times)$ or $\text{FO}(\text{BIT})$ to emphasise that we allow first-order formulas to use such additional relations.¹ We also use that $\text{FO}(+, \times) = \text{FO}(\text{BIT})$ [Imm99].

Tree decompositions and treewidth. A *tree decomposition* (T, B) of G consists of a (rooted, directed) tree $T = (I, F, r)$, with (tree) nodes I , (tree) edges F , a distinguished root node $r \in I$, and a function $B: I \rightarrow 2^V$ such that

- (1) the set $\{i \in I \mid v \in B(i)\}$ is non-empty for each node $v \in V$,
- (2) there is an $i \in I$ with $\{u, v\} \subseteq B(i)$ for each edge $(u, v) \in E$, and
- (3) the subgraph $T[\{i \in I \mid v \in B(i)\}]$ is connected for each node $v \in V$.

We refer to the number of children of a node i of T as its *degree*, and to the set $B(i)$ as its *bag*. We denote the parent node of i by $p(i)$. The *width* of a tree decomposition is defined as the maximal size of a bag minus 1. The *treewidth* of a graph G is the minimal width among all tree decompositions of G . A tree decomposition is *d-nice*, for some $d \in \mathbb{N}$, if

- (1) T has depth at most $d \log n$,
- (2) the degree of the nodes is at most 2, and
- (3) all bags are distinct.

Often we do not make the constant d explicit and just speak of *nice* tree decompositions.

Later we will use that tree decompositions can be transformed into nice tree decompositions with slightly increased width. This is formalized in the following lemma, whose proof is an adaption of [EJT10, Lemma 3.1].

Lemma 2.1. *For every $k \in \mathbb{N}$ there is a constant $d \in \mathbb{N}$ such that for every graph of treewidth k , a d -nice tree decomposition of width $4k + 5$ can be computed in logarithmic space.*

Proof. Let G be a graph of treewidth k . By [EJT10, Lemma 3.1] a tree decomposition (T, B) of width $4k + 3$ can be computed in logarithmic space, such that each non-leaf node has degree 2 and the depth is at most $d \log n$, for a constant d that only depends on k . To obtain a tree decomposition with distinct bags, we compose this algorithm with three further algorithms, each reading a tree decomposition (T, B) and transforming it into a tree decomposition (T', B') with a particular property. Since each of the four algorithms requires only logarithmic space, the same holds for their composition.

The first transformation algorithm produces a tree decomposition, in which for each leaf bag i it holds $B(i) \not\subseteq B(p(i))$. In particular, after this transformation, each bag of a leaf node i contains some graph node $u(i)$ that does not appear in any other bag. This transformation inspects each node i separately in a bottom-up fashion, and removes it if (1) $B(i) \subseteq B(p(i))$ and (2) every bag below i is a subset of $B(i)$. Clearly, logarithmic space suffices for this.

The second transformation (inductively) removes an inner node i of degree 1 with a child i' and inserts an edge between $B(p(i))$ and $B(i')$ whenever $B(i) \subseteq B(p(i))$ or $B(i) \subseteq B(i')$

¹The question of $<$ -invariance (c.f. [Lib04]) will not be relevant in the context of this article as a specific relation \leq will be available in the structure.




Figure 1: Illustration of (a) a proper triangle (i_0, i_1, i_2) , (b) a unary triangle (i_0, i_1, i_1) , and (c) an open triangle (i_0, i_0, i_0) . The blue shaded area is the part of the tree contained in the triangle.

holds. For this transformation, only one linear chain of nodes in T has to be considered at any time and therefore logarithmic space suffices again. Clearly, the connectivity property is not affected by these deletions.

The third transformation adds to every bag of an inner node i the nodes $u(i_1)$ and $u(i_2)$, guaranteed to exist by the first transformation, of the leftmost and rightmost leaf nodes i_1 and i_2 of the subtree rooted at i , respectively. Here, we assume the children of every node to be ordered by the representation of T as input to the algorithm. After this transformation, each node of degree 2 has a different bag than its two children thanks to the addition of $u(i_1)$ and $u(i_2)$. Each node of degree 1 has a different bag than its child, since this was already the case before (and to both of them the same two nodes might have been added). Altogether, all bags are pairwise distinct and the bag sizes have increased by at most 2.

We emphasise that, whenever a leftmost graph node $u(i_1)$ is added to $B(i)$, it is also added to all bags of nodes on the path from i to i_1 and therefore the connectivity property is not corrupted. It is easy to see that the third transformation can also be carried out in logarithmic space.

The three presented algorithms never increase the depth of a tree decomposition, so the final result is a d -nice tree decomposition for G of width $4k + 5$. \square

In this paper we only consider nice tree decompositions, and due to property (3) of these decompositions we can identify bags with nodes from I .

For two nodes i, i' of I , we write $i \preceq i'$ if i' is in the subtree of T rooted at i and $i \prec i'$ if, in addition, $i' \neq i$. A *triangle* δ of T is a triple (i_0, i_1, i_2) of nodes from I such that $i_0 \preceq i_1$, $i_0 \preceq i_2$, and (1) $i_1 = i_2$ or (2) neither $i_1 \preceq i_2$ nor $i_2 \preceq i_1$. In case of (2) we call the triangle *proper*, in case of (1) *unary*, unless $i_0 = i_1 = i_2$ in which we call it *open* (see Figure 1 for an illustration). The subtree $T(\delta)$ induced by a triangle consists of all nodes j of T for which the following holds:

- (i) $i_0 \preceq j$,
- (ii) if $i_0 \prec i_1$ then $i_1 \not\prec j$, and
- (iii) if $i_0 \prec i_2$ then $i_2 \not\prec j$.

That is, for a proper or unary triangle, $T(\delta)$ contains all nodes of the subtree rooted at i_0 which are not below i_1 or i_2 . For an open triangle $\delta = (i_0, i_0, i_0)$, $T(\delta)$ is just the subtree rooted at i_0 .

Each triangle δ induces a subgraph $G(\delta)$ of G as follows: $V(\delta)$ is the union of all bags of $T(\delta)$. By $B(\delta)$ we denote the set $B(i_0) \cup B(i_1) \cup B(i_2)$ of *interface nodes* of $V(\delta)$. All other nodes in $V(\delta)$ are called *inner nodes*. The edge set of $G(\delta)$ consists of all edges of G that involve at least one inner node of $V(\delta)$.

MSO-logic and MSO-types. MSO is the extension of first-order logic that allows existential and universal quantification over set variables X, X_1, \dots . The (*quantifier*) *depth* of an MSO formula is the maximum nesting depth of (second-order and first-order) quantifiers in the syntax tree of the formula.

For a signature Σ and a natural number $d \geq 0$, the *depth- d MSO type* of a Σ -structure \mathcal{A} is defined as the set of all MSO sentences φ over Σ of quantifier depth at most d , for which $\mathcal{A} \models \varphi$ holds. We also define the notion of types for structures with additional constants and MSO formulas with free variables. Let \mathcal{A} be a Σ -structure and $\bar{v} = (v_1, \dots, v_m)$ a tuple of elements from \mathcal{A} . We write (\mathcal{A}, \bar{v}) for the structure over $\Sigma \cup \{c_1, \dots, c_m\}$ which interprets c_i as v_i , for every $i \in \{1, \dots, m\}$. For a set \mathcal{Y} of first-order and second-order variables and an assignment α for the variables of \mathcal{Y} , the *depth- d MSO type* of $(\mathcal{A}, \bar{v}, \alpha)$ is the set of MSO formulas with free variables from \mathcal{Y} of depth d that hold in $(\mathcal{A}, \bar{v}, \alpha)$. For every depth- d MSO type τ , there is a depth- d MSO formula α_τ that is true in exactly the structures and for those assignments of type τ .

The logic *guarded second-order logic* (GSO) extends MSO by *guarded* second-order quantification. Thus, it syntactically allows to quantify over non-unary relation variables. However, this quantification is semantically restricted: a tuple $\bar{t} = (a_1, \dots, a_m)$ can only occur in a quantified relation, if all elements from $\{a_1, \dots, a_m\}$ occur together in some tuple of the structure, in which the formula is evaluated.

For more background on MSO logic and types, readers might consult, e.g., [Lib04].

Complexity classes and descriptive characterisations. Our main result refers to the complexity class (uniform) AC^1 . It contains all queries that can be computed by (families of uniform) circuits of depth $\mathcal{O}(\log n)$, consisting of polynomially many “and”, “or” and “not” gates, where “and” and “or” gates may have unbounded fan-in. It contains the classes LOGSPACE and NL, and it can be characterised as the class $IND[\log n]$ of problems that can be expressed by applying a first-order formula $\mathcal{O}(\log n)$ times [Imm99, Theorem 5.22]. Here, n denotes the size of the universe and the formulas can use built-in relations $+$ and \times . More generally, this characterisation is also valid for the analogously defined classes $AC[f(n)]$ and $IND[f(n)]$, where the depth of the circuits and the number of applications of the first-order formula is $f(n)$, respectively, for some function $f: \mathbb{N} \rightarrow \mathbb{N}$. Technically, the function f needs to be *first-order constructible*, that is, there has to be a $FO(+, \times)$ formula $\psi_f(\bar{x})$ such that $\mathcal{A} \models \psi_f(\bar{a})$ if and only if \bar{a} is a base- n representation of $f(n)$, for any ordered structure \mathcal{A} with domain $\{0, \dots, n-1\}$.

Our proofs often assume that $\log n$ is a natural number, but they can be easily adapted to the general case.

3. DYNAMIC COMPLEXITY

We briefly repeat the essentials of dynamic complexity, closely following [SZ16, DKM⁺18].

The goal of a dynamic program is to answer a given query on an *input database* subjected to changes that insert or delete single tuples. The program may use an auxiliary data structure represented by an *auxiliary database* over the same domain. Initially, both input and auxiliary database are empty; and the domain is fixed during each run of the program.

A dynamic program has a set of update rules that specify how auxiliary relations are updated after a change of the input database. An *update rule* for updating an auxiliary relation T is basically a formula φ . As an example, if $\varphi(\bar{x}, \bar{y})$ is the update rule for auxiliary

relation T under insertions into input relation R , then the new version of T after insertion of a tuple \bar{a} to R is $T \stackrel{\text{def}}{=} \{\bar{b} \mid (\mathcal{I}, Aux) \models \varphi(\bar{a}, \bar{b})\}$ where \mathcal{I} and Aux are the current input and auxiliary databases. For a state $\mathcal{S} = (\mathcal{I}, Aux)$ of the dynamic program \mathcal{P} with input database \mathcal{I} and auxiliary database Aux we denote the state of the program after applying a sequence α of changes by $\mathcal{P}_\alpha(\mathcal{S})$. The dynamic program \mathcal{P} *maintains* a k -ary query q if, for each non-empty sequence α of changes and each empty input structure \mathcal{I}_\emptyset , a designated auxiliary relation Q in $\mathcal{P}_\alpha(\mathcal{S}_\emptyset)$ and $q(\alpha(\mathcal{I}_\emptyset))$ coincide. Here, $\mathcal{S}_\emptyset = (\mathcal{I}_\emptyset, Aux_\emptyset)$, where Aux_\emptyset denotes the empty auxiliary structure over the domain of \mathcal{I}_\emptyset , and $\alpha(\mathcal{I}_\emptyset)$ is the input database after applying α .

In this article, we are particularly interested in maintaining queries for structures of bounded treewidth. There are several ways to adjust the dynamic setting to restricted classes \mathcal{C} of structures. Sometimes it is possible that a dynamic program itself detects that a change operation would yield a structure outside the class \mathcal{C} . However, here we simply disallow change sequences that construct structures outside \mathcal{C} . That is, in the above definition, only change sequences α are considered, for which each prefix transforms an initially empty structure into a structure from \mathcal{C} . We say that a program maintains q for a class \mathcal{C} of structures, if Q contains its result after each change sequence α such that the application of each prefix of α to \mathcal{I}_\emptyset yields a structure from \mathcal{C} .

The class of queries that can be maintained by a dynamic program with first-order update formulas is called DYNFO. We say that a query q is in DYNFO for a class \mathcal{C} of structures, if there is such a dynamic program that maintains q for \mathcal{C} . Programs for queries in DYNFO(+, \times) have three particular auxiliary relations $\leq, +, \times$ that are initialised as a linear order and the corresponding addition and multiplication relations.

For a wide class of queries, membership in DYNFO(+, \times) implies membership in DYNFO [DKM⁺18]. Queries of this class have the property that the query result does not change considerably when elements are added to the domain but not to any relation. Informally, a query q is called *almost domain independent* if there is a constant c such that if a structure already has at least c “non-active” elements, adding more “non-active” elements does not change the query result with respect to the original elements. More formally, a query q is *almost domain independent* if there is a $c \in \mathbb{N}$ such that, for every structure \mathcal{A} and every set $B \subseteq A \setminus \text{adom}(\mathcal{A})$ with $|B| \geq c$ it holds $q(\mathcal{A})[(\text{adom}(\mathcal{A}) \cup B)] = q(\mathcal{A}[(\text{adom}(\mathcal{A}) \cup B)])$.

Examples 3.1.

- (a) The binary reachability query q_{Reach} , that maps a directed input graph $G = (V, E)$ to its transitive closure relation, is almost domain independent with $c = 0$: adding any set $B \subseteq V \setminus \text{adom}(G)$ of isolated nodes to a graph does not create or destroy paths in the remaining graph. Note that, for each node $v \in V$, the tuple (v, v) is part of the query result $q_{\text{Reach}}(G)$, so $q_{\text{Reach}}(G) \neq q_{\text{Reach}}(G[\text{adom}(G)])$ in general and therefore q_{Reach} is not *domain independent* in the sense of [DKM⁺15].
- (b) The FO definable Boolean query $q_{|\neg \text{adom}|=2}$, which is true if and only if exactly two elements are not in the active domain, is almost domain independent with $c = 3$.
- (c) The Boolean query q_{even} , which is true for domains of even size and false otherwise, is not almost domain independent.

Furthermore, all properties definable in monadic second-order logic are almost domain dependent.

Proposition 3.2. *All MSO-definable queries are almost domain independent.*

Proof sketch. This can be easily shown by an Ehrenfeucht game. Let φ be an MSO-formula of quantifier depth d with e free (node) variables and let $c = 2^d d + e$. Consider two graphs \mathcal{A}_1 and \mathcal{A}_2 that result from adding $c_1 \geq c$ and $c_2 \geq c$ isolated nodes to some graph \mathcal{A} , respectively. Since φ might have free variables, the game is played on $(\mathcal{A}_1, \bar{a}_1)$ and $(\mathcal{A}_2, \bar{a}_2)$, where \bar{a}_1 and \bar{a}_2 are tuples of elements of length e . Within d moves the spoiler can use at most d set moves and can therefore induce at most 2^d different “colours” on \mathcal{A}_1 and \mathcal{A}_2 . However, the duplicator can easily guarantee that, for each such “colour”, the number of isolated nodes (which do not occur in the initial tuples on both structures) of that colour is the same in \mathcal{A}_1 and \mathcal{A}_2 or it is larger than d , in both of them. Since the spoiler can have at most d node moves, he can not make use of this difference in the two structures. \square

The following proposition adapts Proposition 7 from [DKM⁺18].

Proposition 3.3. *Let q be an almost domain independent query. If $q \in \text{DYNFO}(+, \times)$ then also $q \in \text{DYNFO}$.*

The same statement is proved for weakly domain independent queries in [DKM⁺18, Proposition 7]. A query q is weakly domain independent, if $q(\mathcal{A})[\text{adom}(\mathcal{A})] = q(\mathcal{A}[\text{adom}(\mathcal{A})])$ for all structures \mathcal{A} , that is, if it is almost domain independent with $c = 0$. The proof of [DKM⁺18, Proposition 7] can easily be adapted for this more general statement, so we omit a full proof here. However, for readers who are familiar with the proof of [DKM⁺18, Proposition 7], we sketch the necessary changes.

Proof sketch. We assume familiarity with the proof of [DKM⁺18, Proposition 7] and only repeat its main outline. That proof explains, for a given $\text{DYNFO}(+, \times)$ -program \mathcal{P} , how to construct a DYNFO -program \mathcal{P}' that is equivalent to \mathcal{P} . This program \mathcal{P}' relies on the observation that a linear order, addition and multiplication can be maintained on the *activated* domain [Ete98], that is, on all elements that were part of the active domain at some time during the dynamic process. The linear order is determined by the order in which the elements are activated.

The program \mathcal{P}' can be regarded as the parallel composition of multiple copies, called *threads*, of the same dynamic program. Each thread simulates \mathcal{P} for a certain period of time: thread i starts when $f(i-1)$ elements are activated and does some initialisation, and it is *in charge* of answering the query whenever more than $f(i)$ but at most $f(i+1)$ elements are activated, for some function f . Thanks to weak domain independence, thread i only needs to simulate \mathcal{P} on a domain with $f(i+1)$ elements, for which it can define the arithmetic relations \leq , $+$ and \times based on the available arithmetic relations on the $f(i-1)$ activated elements.

For almost domain independent queries we extend this technique slightly. Let c be the constant from almost domain independence. The query result for a structure with domain size n coincides with the result of a simulation on a domain of size n' , if $n = n'$ or if there are at least c non-activated elements in both domains. To ensure that the simulation can in principle represent these c elements, phase i is generally in charge as long as the number of activated elements is at least $f(i) - c + 1$ but at most $f(i+1) - c$, as in phase i one can simulate \mathcal{P} on a domain of size $f(i+1)$.

However, if in the original structure the number of non-activated elements becomes smaller than c , the simulation has to switch to a domain of the same size as the original domain. Therefore in phase i there is not only one simulation with domain size $f(i+1)$, but one simulation (denoted by the pair (i, ℓ)) for each domain size $\ell \in \{f(i) - c + 1, \dots, f(i+1)\}$.

During phase i the simulation denoted by $(i, f(i+1))$ is in charge unless there are fewer than c non-activated elements in the original domain, which can be easily detected by a first-order formula. As soon as that happens, the simulation denoted by (i, n) takes over, where n is the size of the domain. \square

4. ALGORITHMIC TECHNIQUE

The definition of DYNFO requires that for the problem at hand each change can be handled by a first-order definable update operation. There are alternative definitions of DYNFO, where the initial structure is non-empty and the initial auxiliary relations can be computed within some complexity class [PI97, WS07]. However, in a practical scenario of dynamic query answering it is conceivable that the quality of the auxiliary relations decreases over time and that they are therefore recomputed from scratch at times. We formalise this notion by a relaxed definition of maintainability in which the initial structure is non-empty, the dynamic program is allowed to apply some preprocessing, and query answers need only be given for a certain number of change steps.

A query q is called (\mathcal{C}, f) -*maintainable*, for some complexity class² \mathcal{C} and some function $f : \mathbb{N} \rightarrow \mathbb{R}$, if there is a dynamic program \mathcal{P} and a \mathcal{C} -algorithm \mathcal{A} such that for each input database \mathcal{I} over a domain of size n , each linear order \leq on the domain, and each change sequence α of length $|\alpha| \leq f(n)$, the relation Q in $\mathcal{P}_\alpha(\mathcal{S})$ and $q(\alpha(\mathcal{I}))$ coincide where $\mathcal{S} = (\mathcal{I}, \mathcal{A}(\mathcal{I}, \leq))$.

Although we feel that (\mathcal{C}, f) -maintainability deserves further investigation, in this paper we exclusively use it as a tool to prove that queries are actually maintainable in DYNFO. To this end, we show next that every $(AC^1, \log n)$ -maintainable query is actually in DYNFO and prove later that the queries in which we are interested are $(AC^1, \log n)$ -maintainable.

Theorem 4.1. *Every $(AC^1, \log n)$ -maintainable, almost domain independent query is in DYNFO.*

We do not prove this theorem directly, but instead give a more general result, strengthening the correspondence between depth of the initialising circuit families and number of change steps the query has to be maintained.

Theorem 4.2. *Let $f : \mathbb{N} \rightarrow \mathbb{N}$ be a first-order constructible function with $f \in \mathcal{O}(n)$. Every $(AC[f(n)], f(n))$ -maintainable, almost domain independent query is in DYNFO.*

Proof. Let $f : \mathbb{N} \rightarrow \mathbb{N}$ be first-order constructible with $f \in \mathcal{O}(n)$ and assume that an $AC[f(n)]$ algorithm \mathcal{A} and a dynamic program \mathcal{P} witness that an almost domain independent query q is $(AC[f(n)], f(n))$ -maintainable. Thanks to Proposition 3.3 it suffices to construct a dynamic program \mathcal{P}' that witnesses $q \in \text{DYNFO}(+, \times)$. We restrict ourselves to graphs, for simplicity.

The overall idea is to use a simulation technique similar to the ones used in Proposition 3.3 and in [SVZ18, Theorem 8.1]. We first present the computations performed by \mathcal{P}' intuitively and later explain how they can be expressed in first-order logic.

We consider each application of one change as a *time step* and refer to the graph after time step t as $G_t = (V, E_t)$. After each time step t , the program \mathcal{P}' starts a thread that is in

²Strictly speaking \mathcal{C} should be a complexity class of functions. In this paper, the implied class of functions will always be clear from the stated class of decision problems.

charge of answering the query at time point $t + f(n)$. Each thread works in two phases, each lasting $\frac{f(n)}{2}$ time steps. Roughly speaking, the first phase is in charge of simulating \mathcal{A} and in the second phase \mathcal{P} is used to apply all changes that occur from time step $t + 1$ to time step $t + f(n)$. Using $f(n)$ many threads, \mathcal{P}' is able to answer the query from time point $f(n)$ onwards.

We now give more details on the two phases and describe afterwards how to deal with time points earlier than $f(n)$. For the first phase, we make use of the equality $\text{AC}[f(n)] = \text{IND}[f(n)]$, see [Imm99]. Let ψ be an inductive formula that is applied $d f(n)$ times, for some d , to get the auxiliary relations $\mathcal{A}(G, \leq)$ for a given graph G and the given order \leq . The program \mathcal{P}' applies ψ to G_t , $2d$ times during each time step, and thus the result of \mathcal{A} on (G_t, \leq) is obtained after $\frac{f(n)}{2}$ steps. The change operations that occur during these steps are not applied to G_t directly but rather stored in some additional relation. If some edge e is changed multiple times, the stored change for e is adjusted accordingly.

During the second phase the $\frac{f(n)}{2}$ stored change operations and the $\frac{f(n)}{2}$ change operations that happen during the next $\frac{f(n)}{2}$ steps are applied to the state after phase 1. To this end, it suffices for \mathcal{P}' to apply two changes during each time step by simulating two update steps of \mathcal{P} . Observe that \mathcal{P}' processes the changes in a different order than they actually occur. However, both change sequences result in the same graph. Since \mathcal{P} can maintain q for $f(n)$ changes, the program \mathcal{P}' can give the correct query answer for q about $G_{t+f(n)}$ at the end of phase 2, that is, at time point $t + f(n)$.

The following auxiliary relations are used by thread i :

- a binary relation \hat{E}_i that contains the edges currently considered by the thread,
- binary relations Δ_i^+ and Δ_i^- that cache edges inserted and deleted during the first phase, respectively,
- a relation \hat{R}_i for each auxiliary relation R of \mathcal{P} with the same arity,
- and a relation C_i that is used as a counter: it contains exactly one tuple which is interpreted as the counter value, according to its position in the lexicographic order induced by \leq .

When thread i starts its first phase at time point t , it sets \hat{E}_i to E_t and the counter C_i to 0; its other auxiliary relations are empty in the beginning. Whenever an edge (u, v) is inserted (or deleted), \hat{E}_i is not changed, (u, v) is inserted into Δ_i^+ (or Δ_i^-)³, and the counter C_i is incremented by one. The relations \hat{R}_i are replaced by the result of applying their defining first-order formulas $2d$ times, as explained above.

When the counter value is at least $\frac{f(n)}{2}$, the thread enters its second phase and proceeds as follows. When an edge (u, v) is inserted (or deleted), it applies this change and the change implied by the lexicographically smallest tuple in Δ_i^+ and Δ_i^- , if these relations are not empty: it simulates \mathcal{P} for these changes using the edge set \hat{E}_i and auxiliary relations \hat{R}_i , replaces the auxiliary relations accordingly and adjusts \hat{E}_i , Δ_i^+ and Δ_i^- . Again, the counter C_i is incremented. If the counter value is $f(n)$, the thread's query result is used as the query result of \mathcal{P}' , and the thread stops. All steps are easily seen to be first-order expressible.

So far we have seen how \mathcal{P}' can give the query answer from time step $f(n)$ onwards. For time steps earlier than $f(n)$ the approach needs to be slightly adapted as the program does not have enough time to simulate \mathcal{A} . The idea is that for time steps $t < f(n)$ the active domain is small and, exploiting the almost domain independence of q , it suffices to compute

³If an edge (u, v) with $(u, v) \in \Delta_i^-$ is inserted, it is instead deleted from Δ_i^- , and accordingly for deletions of edges $(u, v) \in \Delta_i^+$.

the query result with respect to this small domain extended by c isolated elements, where c is the constant from almost domain independence. The result on this restricted domain can afterwards be used to define the result for the whole domain.

Towards making this idea more precise, let n_0, b be such that $bn \geq f(n)$ for all $n \geq n_0$. We focus on explaining how \mathcal{P}' handles structures with $n \geq n_0$, as small graphs with less than n_0 nodes can be dealt with separately.

The program \mathcal{P}' starts a new thread at time $\frac{t}{2}$ for the graph $G_{\frac{t}{2}}$ with at most $\frac{t}{2}$ edges. Such a thread is responsible for providing the query result after t time steps, and works in two phases that are similar to the phases described above. It computes relative to a domain D_t of size $\min\{2t + c, n\}$, where c is the constant from (almost) domain independence. The size of D_t is large enough to account for possible new nodes used in edge insertions in the following $\frac{t}{2}$ change steps. The domain D_t is chosen as the first $|D_t|$ elements of the full domain (with respect to \leq). The program \mathcal{P}' maintains a bijection π between the active domain D_G of the current graph G and the first $|D_G|$ elements of the domain to allow a translation between D_G and D_t .

The first phase of the thread for t starts at time point $\frac{t}{2} + 1$ and applies ψ for $8bcd$ times during each of the next time steps. This simulation of \mathcal{A} is finished after at most $\frac{(2t+c)b}{8bc} \leq \frac{t}{4}$ time steps, and therefore the auxiliary relations are properly initialised at time point $\frac{3t}{4}$.

In the second phase, starting at time step $\frac{3t}{4} + 1$ and ending at time step t , the changes that occurred in the first phase are applied, two at a time. The thread is then ready to answer q at time point t . Since at time t at most $2t$ elements are used by edges, the almost domain independence of q guarantees that the result computed by the thread relative to D_t coincides with the D_t -restriction of the query result for $\pi(G_t)$. The query result for G_t is obtained by translating the obtained result according to π^{-1} , and extending it to the full domain. More precisely, a tuple \bar{t} is included in the query result, if it can be generated from a tuple \bar{t}' of the restricted query result by replacing elements from $\pi^{-1}(D_t) \setminus \text{adom}(G_t)$ by elements from $V \setminus \text{adom}(G_t)$ (under consideration of equality constraints among these elements). Again, all steps are easily seen to be first-order definable using the auxiliary relations from above.

The above presentation assumes a separate thread for each time point and each thread uses its own relations. These threads can be combined into one dynamic program as follows. Since at each time point at most $f(n)$ threads are active, we can number them in a round robin fashion with numbers $1, \dots, f(n)$ that we can encode by tuples of constant arity. The arity of all auxiliary relations is incremented accordingly and the additional dimensions are used to indicate the number of the thread to which a tuple belongs. \square

5. WARM-UP: 3-COLOURABILITY

In this section, we show that the 3-colourability problem 3COL for graphs of bounded treewidth can be maintained in DYNFO. Given an undirected graph, 3COL asks whether its vertices can be coloured with three colours such that adjacent vertices have different colours.

Theorem 5.1. *For every k , 3COL is in DYNFO for graphs with treewidth at most k .*

The remainder of this section is dedicated to a proof for this theorem. Thanks to Theorem 4.1 and the fact that 3COL is almost domain independent, it suffices to show that 3COL is $(AC^1, \log n)$ -maintainable for graphs with treewidth at most k . In a nutshell, our approach can be summarised as follows.

The AC¹ initialisation computes a nice tree decomposition (T, B) of width at most $4k + 5$ and maximum bag size $\ell \stackrel{\text{def}}{=} 4k + 6$, as well as information about the 3-colourability of induced subgraphs of G . More precisely, it computes, for each triangle δ of T and each 3-colouring C of the nodes of $B(\delta)$, whether there exists a colouring C' of the inner vertices of $G(\delta)$, such that all edges involving at least one inner vertex are consistent with $C \cup C'$.

During the following $\log n$ change operations, the dynamic program does not need to do much. It only maintains a set S of *special* bags: for each *affected* graph node v that participates in any changed (i.e. deleted or inserted) edge, S contains one bag in which v occurs. Also, if two bags are special, their least common ancestor is considered special and is included in S . It will be guaranteed that there are at most $4 \log n$ special bags. With the auxiliary information, a first-order formula φ can test whether G is 3-colourable as follows. By existentially quantifying 8ℓ variables, the formula can choose two bits of information for each of the at most $4\ell \log n$ nodes in special bags. For each such node, these two bits are interpreted as encoding of one of three colours and the formula φ checks that this colouring of the special bags can be extended to a colouring of G . This can be done with the help of the auxiliary relations computed during the initialisation which provide all necessary information about colourability of subgraphs induced by triangles consisting of special bags.

Before we give a detailed proof, we need some more notation. Let $G = (V, E)$ be a graph and (T, B) with $T = (I, F, r)$ a nice tree decomposition with bags of size at most ℓ . A *colouring* of a set U of vertices is just a mapping from U to $\{1, 2, 3\}$. An edge (u, v) is *properly coloured* if u and v are mapped to different colours. For a triangle δ , we say that a colouring C of $B(\delta)$ is *consistent*, if there exists a colouring C' of the inner vertices of $G(\delta)$ such that all edges of $G(\delta)$ are properly coloured by $C \cup C'$. Recall that $G(\delta)$ only contains edges that involve at least one inner vertex.

We say that a tuple $\bar{v}(i) = (v_1, \dots, v_\ell)$ *represents a tree node* $i \in I$ (or, the bag $B(i)$) if $B(i) = \{v_1, \dots, v_\ell\}$. A tuple $\bar{v}(\delta) = (\bar{v}(i_0), \bar{v}(i_1), \bar{v}(i_2))$ *represents the triangle* $\delta = (i_0, i_1, i_2)$. If $\bar{v}(\delta) = (v_1, \dots, v_{3\ell})$ represents the triangle δ and \bar{c} is a tuple from $\{1, 2, 3\}^{3\ell}$ such that $c_j = c_{j'}$ whenever $v_j = v_{j'}$ for $j, j' \in \{1, \dots, 3\ell\}$, we write $C_{\bar{c}, \bar{v}}$ for the colouring of $B(\delta)$ defined by $C_{\bar{c}, \bar{v}}(v_j) = c_j$, for every $j \in \{1, \dots, 3\ell\}$.

Proof (of Theorem 5.1). Let $G = (V, E)$ be a graph of treewidth at most k .

The AC¹ initialisation first computes a d -nice tree decomposition (T, B) with bags of size at most $\ell = 4k + 6$, for the constant d guaranteed to exist by Lemma 2.1, and the predecessor relation \preceq of T . Also, it initialises the relations \leq and BIT. Next, it computes the following auxiliary relations in a bottom-up fashion with respect to $T = (I, F, r)$. For each tuple $\bar{c} \in \{1, 2, 3\}^{3\ell}$ the auxiliary relation $R_{\bar{c}}$ contains all tuples $\bar{v}(\delta)$ from $V^{3\ell}$ that represent some triangle δ such that $C_{\bar{c}, \bar{v}}$ is a consistent colouring of $B(\delta)$.

The auxiliary relations are computed inductively and bottom-up, that is, the auxiliary information for a tuple representing a triangle (i_0, i_1, i_2) is computed by using the information for the triangles rooted at the two children of i_0 . It will be easy to see that each inductive step can be defined by a first-order formula and, since T has depth $d \log n$, the induction reaches a fixpoint after $d \log n$ iterations. Therefore the initialisation is in $\text{IND}[\log n] = \text{AC}[\log n]$. We recall that triangles can be open, unary or proper, depending on whether they are induced by a single bag, by two, or by three bags.

For the base case, a tuple \bar{v} representing an open triangle corresponding to a leaf of T is in $R_{\bar{c}}$ if and only if, for each $j \in \{1, \dots, \ell\}$, $c_j = c_{\ell+j} = c_{2\ell+j}$, since there are no inner vertices to worry about.




Figure 2: Illustration of the inductive step in the computation of colourability information for triangles in the proof of Theorem 5.1.

The inductive cases are straightforward. We only describe in detail the case of a proper triangle $\delta = (i_0, i_1, i_2)$ where i_1 and i_2 are in different subtrees of i_0 ; the other cases are similar. Let i'_1 and i'_2 be the two children of i_0 such that $i'_1 \preceq i_1$ and $i'_2 \preceq i_2$. Figure 2 illustrates this situation. By the induction hypothesis, the auxiliary information for all triangles rooted at i'_1 and i'_2 has already been computed. A tuple $\bar{v} = \bar{v}(\delta)$ is in some $R_{\bar{c}}$, if there are tuples $\bar{d}, \bar{e} \in \{1, 2, 3\}^{3\ell}$ such that $\bar{u} = \bar{v}((i'_1, i_1, i_1)) \in R_{\bar{d}}$, $\bar{w} = \bar{v}((i'_2, i_2, i_2)) \in R_{\bar{e}}$ and it holds that

- $C_{\bar{c}, \bar{v}}$ and $C_{\bar{d}, \bar{u}}$ coincide on $B(i_0) \cap B(i'_1)$ and on $B(i_1)$,
- $C_{\bar{c}, \bar{v}}$ and $C_{\bar{e}, \bar{w}}$ coincide on $B(i_0) \cap B(i'_2)$ and on $B(i_2)$, and
- all edges over $B(i_0) \cup B(i'_1) \cup B(i'_2) \cup B(i_1) \cup B(i_2)$ of which at least one node is not in $B(i_0) \cup B(i_1) \cup B(i_2)$ are properly coloured by $C_{\bar{c}, \bar{v}} \cup C_{\bar{d}, \bar{u}} \cup C_{\bar{e}, \bar{w}}$.

We next describe how a dynamic program can maintain 3-colourability for $\log n$ change steps starting from the above initial auxiliary relations with the help of an additional ℓ -ary relation S and another binary relation N . Whenever an edge (u, v) is inserted into or deleted from E , we consider both u and v as *affected*. With each affected graph node v we associate a tree node $i(v) \in I$ such that $v \in B(i(v))$. Tree nodes of the form $i(v)$ for affected nodes v are called *special*. Furthermore, if node i is the least common ancestor of two special nodes i_1, i_2 it becomes special, as well. The dynamic program keeps track of all special nodes using the relation S which contains all tuples \bar{v} that represent some special node. Furthermore, using the relation N it maintains a bijection between the first $\ell|S|$ nodes of V with respect to the linear order \leq and the graph nodes in S . We call a triangle $\delta = (i_0, i_1, i_2)$ of T *clean* if there are no special nodes in $T(\delta)$ apart from i_0, i_1, i_2 .

It only remains to describe how a first-order formula can check 3-colourability of G given the relation S and the relations $R_{\bar{c}}$.

We note first that within $\log n$ steps at most $2 \log n$ graph nodes can be affected resulting in at most $4 \log n$ special tree nodes altogether (since each new special node can contribute at most one new least common ancestor of special nodes). That is, the set Z of graph nodes occurring in some tuple of S contains at most $4\ell \log n$ nodes. A colouring of Z can be represented by $8\ell \log n$ bits and can thus be guessed by a first-order formula by quantifying over 8ℓ first-order variables $x_1, \dots, x_{4\ell}, y_1, \dots, y_{4\ell}$. More precisely, the j -th bits of x_r and y_r together represent the colour of the special node at position $(r-1) \log n + j$ with respect to the linear order represented by N . The first-order formula can easily check that the colouring C of S represented by $x_1, \dots, x_{4\ell}, y_1, \dots, y_{4\ell}$ is consistent for edges between special nodes and

that for each clean triangle of T induced by special nodes it can be extended to a consistent colouring of the inner nodes. The latter information is available in the relations $R_{\bar{c}}$. \square

6. MSO QUERIES

In this section we prove a dynamic version of Courcelle’s Theorem: all MSO properties can be maintained in DYNFO for graphs with bounded treewidth. More precisely, for a given MSO sentence φ we consider the model checking problem MC_φ that asks whether a given graph G satisfies φ , that is, whether $G \models \varphi$ holds.

Theorem 6.1. *For every MSO sentence φ and every k , MC_φ is in DYNFO for graphs with treewidth at most k .*

Since, for every k , guarded second-order logic (GSO) has the same expressive power as MSO on graphs with treewidth at most k [Cou94, Theorem 2.2], we can immediately conclude the following corollary.

Corollary 6.2. *For every GSO sentence φ and every k , MC_φ is in DYNFO for graphs with treewidth at most k .*

We first give a rough sketch of the proof. Let φ be a fixed MSO formula of quantifier depth d and k a treewidth. We show that MC_φ is $(\text{AC}^1, \log n)$ -maintainable for graphs with treewidth at most k . The construction of a dynamic program for MC_φ is similar to the one in the proof for 3-COLOURABILITY (Theorem 5.1). At each point, the program needs to evaluate φ on a graph $G' = (V, (E \setminus E^-) \cup E^+)$ with n nodes, where $|E^- \cup E^+| = \mathcal{O}(\log n)$, using a nice tree decomposition (T, B) of width $4k + 5$ for the initial graph $G = (V, E)$ and auxiliary information on the MSO type of depth d for each triangle of T (defined as in Section 2).

The graph G' can be viewed as having a *center* $C \subseteq V$ of logarithmic size, that contains the nodes with edges in $E^- \cup E^+$ and, additionally, for each of these nodes v all nodes of one bag that contains v .

Furthermore, there are node sets D_1, \dots, D_ℓ that, together with C , contain all nodes from V , such that the sets $D_i - C$ are pairwise disjoint and disconnected, and each set $D_i \cap C$ has size $\mathcal{O}(1)$ (cf. Figure 3). From the type information for the triangles, the program can infer the depth- d MSO types of each $G'[D_i]$.

The situation is similar as in the *Composition Theorem* of Elberfeld, Grohe and Tantau [EGT16, Theorem 3.8]. However, in their setting, the size of C is bounded by a constant, and they show, very roughly, that there is a first-order formula that can be evaluated on a suitable extension of $G[C]$ by information on the types of the $G[D_i]$ to yield the same result as φ on G .

We show that in our setting one can construct an MSO formula ψ such that $G' \models \varphi$ if and only if $\mathcal{B} \models \psi$, where \mathcal{B} is a structure, which extends $G'[C]$ by type information about the $G'[D_i]$. This construction is detailed in Section 6.1 below. It uses well-known techniques and, in particular, a composition theorem by Shelah (cf. Theorem 6.7).

The dynamic program then uses a first-order formula (to be evaluated in a suitable extension of G' with auxiliary relations) that is obtained from ψ by replacing the second-order quantification over C by first-order quantification over V . This is possible, since sets of size $\mathcal{O}(\log n)$ over C can be encoded by $\mathcal{O}(1)$ elements of V .




Figure 3: Sketch of a graph with center C of connection width 2, highlighted in blue, and petals D_1, \dots, D_6 .

In the remainder of this section we make these ideas more precise. In the next subsection we state and prove a composition theorem for graphs with a center of the form described above. Then, in Section 6.2, we show how this theorem is applied to dynamically evaluate an MSO formula φ .

6.1. A Feferman–Vaught-type composition theorem. In the following, we give an adaptation of the Feferman–Vaught-type composition theorem from [EGT16] that will be useful for maintaining MSO properties. Intuitively, the idea is very easy, but the formal presentation will come with some technicalities. For ease of presentation, we explain the basic idea for graphs first.

We consider graphs $G = (V, E)$ with a *center* $C \subseteq V$, such that there are sets D_1, \dots, D_ℓ such that, for some $w > 0$, the following conditions hold.

- $C \cup \bigcup_{i=1}^{\ell} D_i = V$.
- For all $i \neq j$, $D_i \cap D_j \subseteq C$.
- All edges in E have both end nodes in C or in some D_i .
- For every i , $|D_i \cap C| \leq w$.
- For each i there is some element $v_i \in D_i \cap C$ that is not contained in any D_j , for $j \neq i$.

In this case, we say that C has *connection width* w in G . See Figure 3 for an illustration.

We refer to the sets D_1, \dots, D_ℓ as *petals* and the nodes v_1, \dots, v_ℓ as *identifiers* of their respective petals. We emphasise that ℓ is bounded by $|C|$, but not assumed to be bounded by a constant. Readers who have read the proof of Theorem 5.1 can roughly think of C as the set of vertices from special bags.

Our goal is to show the following. If a graph G has a center C of connection width w , for some constant w , then $G[C]$ can be extended by the information about the MSO types of its petals in a suitable way, resulting in a structure \mathcal{B} with universe C , such that MSO formulas over G have equivalent MSO formulas over \mathcal{B} .

In the following, we work out the above plan in more detail. Although, for Theorem 6.1, we need the composition theorem only for coloured graphs with some constants, we deal in the following with arbitrary signatures. We fix some relational signature Σ and assume that it contains a unary relation symbol C .

The definition of the connection width of sets C easily carries over to Σ -structures. In particular, tuples need to be entirely in C or in some petal D_i , and all constants of the structure need to be included in C . For every i , we call the set $I_i \stackrel{\text{def}}{=} D_i \cap C$ the *interface* of D_i and the nodes of $D_i - C$ *inner elements* of D_i .

Let \mathcal{A} be a Σ -structure, C a center of connection width w with petals D_1, \dots, D_ℓ . For every i , let $\bar{u}^i = (u_1^i, \dots, u_w^i)$ be a tuple of elements from the interface I_i of D_i such that u_1^i is an identifier of its petal D_i and every node from I_i occurs in \bar{u}^i . By $(\mathcal{A}_i, \bar{u}^i)$ we denote the substructure of \mathcal{A} induced by D_i with u_1^i, \dots, u_w^i as constants but *without* all tuples over C , i.e., $(\mathcal{A}_i, \bar{u}^i)$ only contains tuples with at least one inner element of D_i .

Let $d > 0$. The *depth d , width w MSO indicator structure* of \mathcal{A} relative to C and tuples \bar{u}^i is the unique structure \mathcal{B} which expands $\mathcal{A}[C]$ by the following relations:

- a w -ary relation J that contains all tuples \bar{u}^i , and
- for every depth- d MSO-type τ over $\Sigma \cup \{c_1, \dots, c_w\}$, a unary relation R_τ containing those identifier nodes u_1^i for which the depth- d MSO-type of $(\mathcal{A}_i, \bar{u}^i)$ is τ .

We note that different choices of the tuples \bar{u}^i result in different indicator structures and we denote the set of all indicator structures of \mathcal{A} relative to C by $\mathcal{S}(\mathcal{A}, C, w, d)$.

We are now ready to formulate the desired composition theorem.

Theorem 6.3. *For each $d > 0$, every MSO sentence φ with depth d , and each w , there is a number d' and a MSO sentence ψ such that for every Σ -structure \mathcal{A} , every center C of \mathcal{A} with connection width w and every $\mathcal{B} \in \mathcal{S}(\mathcal{A}, C, w, d')$ it holds $\mathcal{A} \models \varphi$ if and only if $\mathcal{B} \models \psi$.*

The proof of Theorem 6.3 uses Shelah's *generalised sums* [She75]. We follow the exposition from Blumensath et al. [BCL08]. In a nutshell, a generalised sum is a composition of several disjoint *component structures* along an *index structure*. Shelah's composition theorem states that MSO sentences on a generalised sum can be translated to MSO sentences on the index structure enriched by MSO type information on the components.

We apply the composition theorem on the basis of the following ideas, illustrated in Figure 4. From the center C and the petals D_i we define an index structure \mathcal{I} and component structures \mathcal{D}_i , respectively. In the generalised sum, these disjoint structures are again composed into a structure that is very similar to \mathcal{A} . More precisely, \mathcal{A} can be defined in the generalised sum by a first-order interpretation, and thus, thanks to Lemma 6.5 below, we can translate the formula φ for \mathcal{A} into a formula φ' on the generalised sum. Shelah's composition theorem then provides a translation of φ' into an MSO formula ψ' on the structure \mathcal{I} enriched with type information on the structures \mathcal{D}_i . This enriched index structure is again very similar to an MSO indicator structure \mathcal{B} : there is a first-order interpretation that defines the enriched index structure in \mathcal{B} . As a consequence, by Lemma 6.5 again, the formula ψ' can be translated into a formula φ on \mathcal{B} .

Before we proceed to the proof of Theorem 6.3, we formally introduce the notions of a generalised sum and a first-order interpretation, and state the corresponding results on translations of MSO formulas. We start with first-order interpretations.

Definition 6.4. Let Σ, Γ be relational signatures. A *first-order interpretation* Υ from Σ to Γ consists of a first-order formula $\varphi_U(x)$ and first-order formulas $\varphi_R(x_1, \dots, x_r)$ for each r -ary relation symbol $R \in \Gamma$, each over signature Σ .

The first-order interpretation Υ *interprets*, in a Σ -structure \mathcal{A} , the Γ -structure $\Upsilon(\mathcal{A})$ with universe $U^{\Upsilon(\mathcal{A})} \stackrel{\text{def}}{=} \{a \mid \mathcal{A} \models \varphi_U(a)\}$ and relations

$$R^{\Upsilon(\mathcal{A})} \stackrel{\text{def}}{=} \{(a_1, \dots, a_r) \mid \mathcal{A} \models \varphi_R(a_1, \dots, a_r), a_1, \dots, a_r \in U^{\Upsilon(\mathcal{A})}\}$$




Figure 4: Overview of the proof strategy of Theorem 6.3.

for each $R \in \Gamma$.

A first-order interpretation from Σ to Γ not only interprets a Γ -structure in a Σ -structure, it also translates Γ -formulas to Σ -formulas.

Lemma 6.5 (see e.g. [BCL08, Proposition 3.2]). *Let Υ be a first-order interpretation from Σ to Γ . For every FO (MSO) formula $\varphi(x_1, \dots, x_\ell)$ over Γ there is an FO (MSO) formula $\varphi^\Upsilon(x_1, \dots, x_\ell)$ over Σ such that $\mathcal{A} \models \varphi^\Upsilon(a_1, \dots, a_\ell) \Leftrightarrow \Upsilon(\mathcal{A}) \models \varphi(a_1, \dots, a_\ell)$ for all Σ -structures \mathcal{A} and all elements $a_i \in U^{\Upsilon(\mathcal{A})}$.*

We now turn to the definition of generalised sums.

Definition 6.6 ([She75], formulation following [BCL08]). Let $\mathcal{I} = (I, S_1, \dots, S_r)$ be a structure and $(\mathcal{D}_i)_{i \in I}$ a sequence of structures $\mathcal{D}_i = (D_i, R_1^i, \dots, R_\ell^i)$ indexed by elements i of \mathcal{I} . The *generalised sum* of $(\mathcal{D}_i)_{i \in I}$ is the structure

$$\sum_{i \in I} \mathcal{D}_i \stackrel{\text{def}}{=} (U, \sim, R'_1, \dots, R'_\ell, S'_1, \dots, S'_r)$$

with universe $U \stackrel{\text{def}}{=} \{\langle i, a \rangle \mid i \in I, a \in D_i\}$ and relations

- $\langle i, a \rangle \sim \langle i', a' \rangle$ if and only if $i = i'$
- $R'_j \stackrel{\text{def}}{=} \{\langle \langle i, a_1 \rangle, \dots, \langle i, a_\ell \rangle \rangle \mid i \in I, (a_1, \dots, a_\ell) \in R_j^i\}$
- $S'_j \stackrel{\text{def}}{=} \{\langle \langle i_1, a_1 \rangle, \dots, \langle i_\ell, a_\ell \rangle \rangle \mid (i_1, \dots, i_\ell) \in S_j, a_k \in D_{i_k} \text{ for all } k \in \{1, \dots, \ell\}\}$

The structures \mathcal{I} and \mathcal{D}_i in this definition are also referred to as *index structure* and *component structures*, respectively.

Theorem 6.7 ([She75], formulation from [BCL08, Theorem 3.16]). *From every MSO sentence φ , a finite sequence χ_1, \dots, χ_s of MSO formulas and an MSO formula ψ can be constructed such that*

$$\sum_{i \in I} \mathcal{D}_i \models \varphi \Leftrightarrow (\mathcal{I}, \llbracket \chi_1 \rrbracket, \dots, \llbracket \chi_s \rrbracket) \models \psi$$

for all index structures \mathcal{I} and component structures \mathcal{D}_i , where $\llbracket \chi \rrbracket \stackrel{\text{def}}{=} \{i \in I \mid \mathcal{D}_i \models \chi\}$.

Intuitively, the formulas χ_i encode the MSO type information on the component structures of the generalized sum.

With the necessary notions in place, we can now prove Theorem 6.3.

Proof (of Theorem 6.3). Suppose that \mathcal{A} is a Σ -structure with center C of connection width w , tuples \bar{u}^i collecting the interface nodes for each petal D_i as described above, and let φ be an MSO formula over signature Σ . The proof is in three steps, depicted in Figure 4:

- (A) We present an index structure \mathcal{I} and component structures \mathcal{D}_i , and show that there is an FO-interpretation Υ that interprets the structure \mathcal{A} in the generalised sum $\sum_{i \in I} \mathcal{D}_i$. Thus there is an MSO formula φ' such that $\mathcal{A} \models \varphi$ if and only if $\sum_{i \in I} \mathcal{D}_i \models \varphi'$ by Lemma 6.5.
- (B) From Theorem 6.7 we obtain formulas ψ' and χ_1, \dots, χ_s such that $\sum_{i \in I} \mathcal{D}_i \models \varphi'$ if and only if $(\mathcal{I}, \llbracket \chi_1 \rrbracket, \dots, \llbracket \chi_s \rrbracket) \models \psi'$.
- (C) Then we show that there is an FO-interpretation Υ' that interprets $(\mathcal{I}, \llbracket \chi_1 \rrbracket, \dots, \llbracket \chi_s \rrbracket)$ in each MSO indicator structure $\mathcal{B} \in \mathcal{S}(\mathcal{A}, C, w, d')$ with appropriate d' . Thus there is an MSO formula ψ that satisfies $\mathcal{B} \models \psi$ if and only if $(\mathcal{I}, \llbracket \chi_1 \rrbracket, \dots, \llbracket \chi_s \rrbracket) \models \psi'$ by Lemma 6.5.

Combining these three steps allows us to conclude

$$\mathcal{A} \models \varphi \stackrel{(A)}{\iff} \sum_{i \in I} \mathcal{D}_i \models \varphi' \stackrel{(B)}{\iff} (\mathcal{I}, \llbracket \chi_1 \rrbracket, \dots, \llbracket \chi_s \rrbracket) \models \psi' \stackrel{(C)}{\iff} \mathcal{B} \models \psi$$

for any MSO indicator structure $\mathcal{B} \in \mathcal{S}(\mathcal{A}, C, w, d')$.

It remains to prove (A) and (C), since (B) is a direct application of Theorem 6.7.

Towards proving (A) we construct structures \mathcal{I} and \mathcal{D}_i which are closely related to the substructure $\mathcal{A}[C]$ of \mathcal{A} and the structures $(\mathcal{A}_i, \bar{u}^i)$, respectively. Let R_1, \dots, R_q be the relation symbols of Σ and let $\Gamma = \{S_1, \dots, S_q\}$. The structure \mathcal{I} has universe C , Γ -relations S_i defined as the restriction of the respective Σ -relation R_i of \mathcal{A} to C , and the relation J as described above. Each structure \mathcal{D}_v for $v \in C$ is over signature $\Sigma \cup \{U_1, \dots, U_w\}$ and defined as follows. If v is an identifier u_1^i of a petal, then $\mathcal{D}_v = (\mathcal{A}_i, \{u_1^i\}, \dots, \{u_w^i\})$, that is, the restriction of \mathcal{A} to the elements from D_i , without any tuples consisting only of elements from $C \cap D_i$, and with additional unary, singleton relations U_1, \dots, U_w such that U_j includes only the j -th interface node u_j^i . If v is no identifier node then \mathcal{D}_v is the structure with universe $\{v\}$ and empty relations.

In the generalised sum $\mathcal{S} \stackrel{\text{def}}{=} \sum_{v \in I} \mathcal{D}_v = (U, \sim, R'_1, \dots, R'_q, U'_1, \dots, U'_w, S'_1, \dots, S'_q, J')$, the universe U consists of elements of the form $\langle u_1^i, w \rangle$, where $w \in D_i$, and of the form $\langle u, u \rangle$ where $u \in C$ is not an identifier of any petal. We emphasise that the formulas of the interpretation that defines (a copy of) \mathcal{A} in \mathcal{S} can *not* access the components u and v of an element $\langle u, v \rangle \in U$. However, U can be partitioned into four kinds of elements, each of which can easily be distinguished from the others in a first-order fashion:

- (i) Elements of the form $\langle v, v \rangle$, for which D_v is not a petal. They can be identified since they constitute an equivalence class of size 1 with respect to \sim ;
- (ii) Elements of the form $\langle v, v \rangle$, for which D_v is a petal. These are precisely the elements in U'_1 ;
- (iii) Elements of the form $\langle v, u \rangle$, where $u \in C$ is in the petal D_v . These elements occur in some U'_i , for $i > 1$ (but not in U'_1);
- (iv) Elements of the form $\langle v, u \rangle$, where $u \notin C$ is in the petal D_v . They do not occur in any U'_i and are not of type (i).

Elements of type (iv) are in one-to-one correspondence with the inner elements of petals D_i . Elements from C might have several copies in U , but only one of the types (i) or (ii). Thus,

the formula that defines the universe for the first-order interpretation Υ of \mathcal{A} in $\sum_{i \in I} \mathcal{D}_i$ simply drops all elements of type (iii). Tuples of \mathcal{A} of a relation R_i that entirely consist of nodes from C (that is, elements of type (i) or (ii) in \mathcal{S}) are directly induced by the corresponding relation S'_i .

In order to define tuples with at least one node of type (iv) in \mathcal{A} , we first observe that it can be expressed in a first order fashion, whether for a type (iii) element $\langle v_1, u_1 \rangle$ and a type (i) element $\langle v_2, v_2 \rangle$ it holds $u_1 = v_2$, i.e., that, intuitively, $\langle v_2, v_2 \rangle$ is the copy representing u_1 in \mathcal{S} that survives in the universe of the interpretation. We claim that this condition holds, if and only if $\langle v_2, v_2 \rangle$ occurs as the i -th entry in some tuple of J' with first entry $\langle v_1, u_1 \rangle$, where i is the unique number such that $\langle v_1, u_1 \rangle \in U'_i$. From this claim, first-order expressibility follows instantly. The “only if”-part of the claim is straightforward. For the “if”-part, it follows from the latter condition that there is a tuple with first entry v_1 and i -th entry v_2 in J , by the definition of J' . Since $\langle v_1, u_1 \rangle \in U'_i$, there is also a tuple in J with v_1 as first entry and u_1 as i -th entry. However, since J has at most one tuple with any given value as first entry, $u_1 = v_2$ follows, as claimed.

A tuple with some element $\langle v, u \rangle$ of type (iv) is now in a relation R_i of the interpretation, if it can be transformed into a tuple of R'_i by replacing some elements $\langle w, w \rangle$ of type (i) with $\langle v, w \rangle$.

It follows from the construction that $\Upsilon(\mathcal{S})$ is isomorphic to \mathcal{A} and therefore $\mathcal{A} \models \varphi$ if and only if $\Upsilon(\mathcal{S}) \models \varphi$. We obtain the formulas φ' , χ_1, \dots, χ_s and ψ' as explained above. Let d' be the maximal quantifier depth of any formula χ_j . This concludes step (A) of the proof.

Towards proving (C), recall that we need to show that there is a first-order interpretation Υ' which interprets $(\mathcal{I}, \llbracket \chi_1 \rrbracket, \dots, \llbracket \chi_s \rrbracket)$ in \mathcal{B} , for any $\mathcal{B} \in \mathcal{S}(\mathcal{A}, C, w, d')$. Let \mathcal{B} be such a structure. The universe of \mathcal{I} is C , that is, the same as the universe of \mathcal{B} . Thus the formula of Υ' that defines the universe is trivial.

For the definitions of the relations $\llbracket \chi_j \rrbracket$, the idea is as follows. If v is *not* an identifier u_1^i of a petal, then $v \in \llbracket \chi_j \rrbracket$ if and only if χ_j holds in the structure consisting of only one element and with empty relations, which can be hard-coded in the defining formula. Otherwise, if $v = u_1^i$ for some i , we need to determine whether χ_j holds in the structure $\mathcal{D}_v = (\mathcal{A}_i, \{u_1^i\}, \dots, \{u_w^i\})$, which is a structure over signature $\Sigma \cup \{U_1, \dots, U_w\}$. The structure \mathcal{B} contains information about the MSO types of the structures $(\mathcal{A}_i, \bar{u}^i)$, but $(\mathcal{A}_i, \bar{u}^i)$ is a structure over signature $\Sigma \cup \{c_1, \dots, c_w\}$. Yet it is easy to see that for the formula χ'_j that is obtained from χ_j by replacing every atom $U_k(x)$ by $x = c_k$ it holds $(\mathcal{A}_i, \bar{u}^i) \models \chi'_j$ if and only if $\mathcal{D}_v \models \chi_j$. So, in this case $v \in \llbracket \chi_j \rrbracket$ if and only if $v \in R_\tau^{\mathcal{B}}$ for a depth- d' MSO type τ with $\chi'_j \in \tau$. All these conditions can even be expressed by quantifier-free first-order formulas for fixed formulas χ_j . As a result, by Lemma 6.5 we obtain from Υ' and ψ' a formula ψ with $\mathcal{B} \models \psi \Leftrightarrow \mathcal{A} \models \varphi$. \square

6.2. The dynamic program. We proceed to show that every MSO-definable property can be maintained in DYNFO, and thus prove Theorem 6.1. Thanks to Theorem 4.1 and Proposition 3.2 it suffices to show that MC_φ is $(\text{AC}^1, \log n)$ -maintainable for graphs G with treewidth at most k .

The idea for our dynamic program is similar to the idea for maintaining 3-colourability: during its initialisation the program constructs a tree decomposition and appropriate MSO types for all triangles (instead of partial colourings as in the proof of Theorem 5.1). During the change sequence, a set C of nodes is defined that contains, for each affected graph node

v , all nodes of at least one special bag containing v . The set C has connection width w for some constant w and the dynamic program basically maintains an MSO indicator structure \mathcal{B} for G relative to C . As there are only $\log n$ many change steps, the size of C is bounded by $\mathcal{O}(\log n)$.

By Theorem 6.3 there is an MSO formula ψ with the property that $G \models \varphi$ if and only if $\mathcal{B} \models \psi$. Although the dynamic program maintains \mathcal{B} , it cannot directly evaluate ψ , as it is restricted to use first-order formulas. For this reason we first show that second-order quantification over sets of size $\mathcal{O}(\log n)$ can be simulated in first-order logic, if a particular relation is present. Afterwards we present the details of the dynamic program.

We call an MSO-formula C -restricted, if all its quantified subformulas are of one of the following forms.

- $\exists x (C(x) \wedge \varphi)$ or $\forall x (C(x) \rightarrow \varphi)$,
- $\exists X (\forall x (X(x) \rightarrow C(x)) \wedge \varphi)$ or $\forall X (\forall x (X(x) \rightarrow C(x)) \rightarrow \varphi)$.

Let \mathcal{A} be a structure with a unary relation C and a $(k+1)$ -ary relation Sub , for some k . We say that Sub *encodes subsets of C* if, for each subset $C' \subseteq C$, there is a k -tuple \bar{t} such that, for every element $c \in C$ it holds $c \in C'$ if and only if $(\bar{t}, c) \in \text{Sub}$. Clearly, such an encoding of subsets only exists if $|V|^k \geq 2^{|C|}$ and thus if $|C| \leq k \log |V|$.

Proposition 6.8. *For each C -restricted MSO-sentence ψ over a signature Σ (containing C) and every k there is a first-order sentence χ over $\Sigma \cup \{S\}$ where S is a $(k+1)$ -ary relation symbol such that, for every Σ -structure \mathcal{A} and $(k+1)$ -ary relation Sub that encodes subsets of C (of \mathcal{A}), it holds $\mathcal{A} \models \psi$ if and only if $(\mathcal{A}, \text{Sub}) \models \chi$.*

Proof. The proof is straightforward. Formulas $\exists X (\forall x (X(x) \rightarrow C(x)) \wedge \varphi)$ are translated into formulas $\exists \bar{x} \varphi'$, where \bar{x} is a tuple of k variables and φ' results from φ by simply replacing every atomic formula $X(y)$ by $\text{Sub}(\bar{x}, y)$. Universal set quantification is translated analogously. \square

Proof (of Theorem 6.1). Thanks to Theorem 4.1 and Proposition 3.2 it suffices to show that MC_φ is $(\text{AC}^1, \log n)$ -maintainable in DYNFO for graphs with treewidth at most k . Let d be the quantifier depth of φ and let d' and ψ be the number and the MSO sentence guaranteed to exist by Theorem 6.3.

Given a graph $G = (V, E)$, the AC^1 initialisation first ensures that relations $\leq, +, \times$ and BIT are available. Then it computes a d_{tree} -nice tree decomposition (T, B) with $T = (I, F, r)$ with bags of size at most $\ell \stackrel{\text{def}}{=} 4k+6$, for the constant d_{tree} guaranteed to exist by Lemma 2.1, together with the predecessor order \preceq on I . With each node $i \in I$, we associate a tuple $\bar{v}(i) = (v_1, \dots, v_m, v_1, \dots, v_1)$ of length ℓ , where $B(i) = \{v_1, \dots, v_m\}$ and $v_1 < \dots < v_m$. That is, if the bag size of i is ℓ , this tuple just contains all graph nodes of the bag in increasing order. If the bag size is smaller, the smallest graph node is repeated. Similarly, with each triangle $\delta = (i_0, i_1, i_2)$ such that the subgraph $G(\delta)$ has at least one inner node, we associate a tuple $\bar{v}(\delta) = (v(\delta), \bar{v}(i_0), \bar{v}(i_1), \bar{v}(i_2))$, where $v(\delta)$ denotes the smallest inner node of $G(\delta)$ with respect to \preceq .

The dynamic program further uses auxiliary relations S, C, N , and D_τ , for each depth- d' MSO type τ over the signature that consists of the binary relation symbol E and $3\ell+1$ constant symbols $c_1, \dots, c_{3\ell+1}$. The intended meaning is that C is a center of G with connection width $3\ell+1$ and that from these relations an MSO indicator structure \mathcal{B} relative to C can be defined in first-order.




Figure 5: Tree of a tree decomposition. Tree nodes representing special bags are coloured blue, the corresponding maximal clean triangles are indicated as coloured areas. The union of all special bags are a center of the graph, with the graphs induced by the maximal clean triangles as petals.

The relation S stores tuples $\bar{v}(i)$ representing special bags, as in the proof of Theorem 5.1. The relations D_τ provide MSO type information for all triangles. More precisely, for each triangle $\delta = (i_0, i_1, i_2)$ for which the subgraph $G(\delta)$ has at least one inner node, D_τ contains the tuple $\bar{v}(\delta)$ if and only if the MSO depth- d' type of $(G(\delta), \bar{v}(\delta))$ is τ .

The set C always contains all graph nodes that occur in special bags (and thus in S), plus one inner node $v(\delta)$, for each maximal clean triangle⁴ δ with at least one inner node. The relation N defines a bijection between C and an initial segment of \leq .

We observe that C is a center of G and that the petals induced by C correspond to the maximal clean triangles, with respect to the special nodes stored in S , with at least two inner nodes. The interface $I(\delta)$ of a petal corresponding to a maximal clean triangle $\delta = (i_0, i_1, i_2)$ contains the nodes from $B(i_0), B(i_1)$, and $B(i_2)$ as well as the node $v(\delta)$, so C has connection width $w \stackrel{\text{def}}{=} 3\ell + 1$. Figure 5 gives an illustration.

Now, an indicator structure $\mathcal{B} \in \mathcal{S}(G, C, w, d')$ can be first-order defined as follows. Clearly, maximal clean triangles can be easily first-order defined from the relation S . For each maximal clean triangle $\delta = (i_0, i_1, i_2)$ with at least two inner nodes, the relation J contains the tuple $\bar{v}(\delta)$, and the relation R_τ contains $v(\delta)$ if and only if $\bar{v}(\delta) \in D_\tau$. We translate the MSO formula ψ to a C -restricted MSO formula χ' such that $\mathcal{B} \models \psi \Leftrightarrow (G, Aux) \models \chi'$, where Aux is the auxiliary database stored by the dynamic program. This translation is basically as described by Lemma 6.5. The formula χ' results from ψ by

- C -restricting every quantified subformula, so, for example, replacing every quantified subformula $\exists X \theta$ by $\exists X (\forall x (X(x) \rightarrow C(x)) \wedge \theta)$ and every quantified subformula $\forall X \theta$ by $\forall X (\forall x (X(x) \rightarrow C(x)) \rightarrow \theta)$, and
- replacing every atom $A(\bar{x})$ by $\theta_A(\bar{x})$, where θ_A is the first-order formula that defines A in (G, Aux) .

It clearly holds that $(G, Aux) \models \chi' \Leftrightarrow \mathcal{B} \models \psi$, and by Theorem 6.3 also $(G, Aux) \models \chi' \Leftrightarrow G \models \varphi$.

⁴For such a triangle $\delta = (i_0, i_1, i_2)$, the nodes i_0, i_1, i_2 are exactly the special nodes in $T(\delta)$.

We now define a relation Sub that encodes subsets of C . We observe that C is of size at most $b \log n$ for some $b \in \mathbb{N}$. Thus a subset C' of C can be represented by a tuple (a_1, \dots, a_b) of nodes, where an element $c \in C$ is in C' if and only if c is the m -th element of C with respect to the mapping defined by N , $m = (\ell - 1) \log n + j$ and the j -th bit of a_ℓ is one. By Proposition 6.8 we finally obtain a first-order formula χ such that $(G, \text{Aux}, \text{Sub}) \models \chi \Leftrightarrow (G, \text{Aux}) \models \chi' \Leftrightarrow G \models \varphi$. That means that a dynamic program that maintains the auxiliary relations as intended can maintain the query MC_φ .

It thus remains to describe how the auxiliary relations can be initialised and updated. The set C is initially the bag $B(r)$ of the root of T plus one inner node $v(r, r, r)$, and S contains the tuple $\bar{v}(r)$.

The relations D_τ are computed in $d_{\text{tree}} \log n$ inductive steps, each of which can be defined in first-order logic, and therefore this computation can be carried out in AC^1 , thanks to $\text{IND}[\log n] = \text{AC}^1$. More precisely, the computation of the relations D_τ proceeds inductively in a bottom-up fashion. It starts with triangles $\delta = (i_0, i_1, i_2)$ for which $T(\delta)$ has exactly one or two inner tree nodes (i.e., nodes different from i_0, i_1, i_2). Since such graphs $G(\delta)$ have at most 5ℓ nodes, their type can be determined by a first-order formula.⁵

For larger triangles, several cases need to be distinguished. Here we explain the case of a triangle $\delta = (i_0, i_1, i_2)$, for which i_0 has child nodes i'_1 and i'_2 such that $i'_1 \preceq i_1$ and $i'_2 \preceq i_2$ (cf., Figure 2). In this case, the type τ of $(G(\delta), \bar{v}(\delta))$ can be determined from the types τ_1 of $(G(\delta_1), \bar{v}(\delta_1))$ and τ_2 of $(G(\delta_2), \bar{v}(\delta_2))$, where $\delta_1 = (i'_1, i_1, i_1)$ and $\delta_2 = (i'_2, i_2, i_2)$, and the type τ_0 of the graph G_0 that includes all edges of $G(\delta)$ that are not already in $G(\delta_1)$ or $G(\delta_2)$. More precisely, τ_0 is the type of $(G_0, \bar{v}(i_0), \bar{v}(i_1), \bar{v}(i_2), \bar{v}(i'_1), \bar{v}(i'_2))$ and G_0 is the subgraph of G with node set $V_0 = \bigcup \{B(i_0), B(i_1), B(i_2), B(i'_1), B(i'_2))\}$ and all edges from $G[V_0]$ that have at least one endpoint in $B(i'_1) \cup B(i'_2)$. These types are either already computed or the graphs are of size at most 5ℓ and their type can therefore be determined by a first-order formula as before.

We make this more precise. We observe that $(G(\delta), \bar{v}(\delta))$ can be composed from the graphs $(G_0, \bar{v}(i_0), \bar{v}(i_1), \bar{v}(i_2), \bar{v}(i'_1), \bar{v}(i'_2))$, $(G(\delta_1), \bar{v}(\delta_1))$ and $(G(\delta_2), \bar{v}(\delta_2))$ by first taking the disjoint union of these graphs and afterwards fusing nodes according to the identities induced by the additional constants. For both operations, the depth- d MSO type of the resulting structure only depends on the depth- d MSO type(s) of the original structure(s) [Mak04, Theorem 1.5 (ii) and Proposition 3.6]. The type τ of $(G(\delta), \bar{v}(\delta))$ is therefore determined by a finite function f as $\tau = f(\tau_0, \tau_1, \tau_2)$, which can be directly encoded into first-order formulas.

Finally, we describe how a dynamic program can maintain S, C , and N for $\log n$ many changes. The relation D_τ is not adapted during the changes. Whenever an edge (u, v) is inserted to or deleted from G , the nodes u and v are viewed as *affected*. For every *affected* node u that is not yet in a bag stored in S , a *special* tree node is selected (in some canonical way, e.g. always the smallest node with respect to \preceq is selected) such that $u \in B(j)$. Furthermore, if node i is the least common ancestor of j and another special node it becomes special, as well. It is easy to see that when selecting j as a special node, at most one further node becomes special. The tuples $\bar{v}(j)$ and $\bar{v}(i)$ (if i exists) are added to S , their elements are added to C , the identifier nodes in C for maximal clean triangles are corrected, and N is updated accordingly. \square

⁵Basically, all isomorphism types of such graphs and their respective MSO types can be directly encoded into first-order formulas.

7. MSO OPTIMISATION PROBLEMS

With the techniques presented in the previous section, also MSO definable optimisation problems can be maintained in DYNFO for graphs with bounded treewidth. An MSO definable optimisation problem OPT_φ is induced by an MSO formula $\varphi(X_1, \dots, X_m)$ with free set variables X_1, \dots, X_m . Given a graph G with vertex set V , it asks for sets $A_1, \dots, A_m \subseteq V$ of minimal⁶ size $\sum_{i=1}^m |A_i|$ such that $G \models \varphi(A_1, \dots, A_m)$. Examples (with $m = 1$) for such problems are the vertex cover problem and the dominating set problem.

We require from a dynamic program for such a problem that it maintains unary query relations Q_1, \dots, Q_m that store, at any time, an optimal solution for the current graph.

Theorem 7.1. *For every MSO formula $\varphi(X_1, \dots, X_m)$ and every k , OPT_φ is in DYNFO for graphs with treewidth at most k .*

As already mentioned in the previous section, for every k and every GSO formula φ there is an MSO formula ψ that is equivalent on graphs with treewidth k [Cou94, Theorem 2.2]. Moreover, if $\varphi = \exists X_1 \dots \exists X_m \varphi'$, then ψ is of the form $\exists X_1^1 \dots \exists X_1^\ell \dots \exists X_m^1 \dots \exists X_m^\ell \psi'$, for some natural number ℓ . So, we can conclude the following corollary.

Corollary 7.2. *For every GSO formula $\varphi(X_1, \dots, X_m)$ and every k , OPT_φ is in DYNFO for graphs with treewidth at most k .*

Given the machinery from the previous section, the plan for a dynamic program for an MSO-definable optimisation problem is relatively straightforward. Again, it suffices to show $(\text{AC}^1, \log n)$ -maintainability. The affected nodes of the graph after $\log n$ changes are again collected in a center C of the graph (with $\mathcal{O}(\log n)$ additional nodes as before). For each petal D_i and each relevant MSO-type τ we basically maintain a collection (B_1, \dots, B_m) of subsets of $D_i - C$ that yields type τ in D_i and is minimal with respect to $\sum_{i=1}^m |B_i|$. Then, it is easy to compute in a first-order fashion, for every possible colouring of C , the minimum achievable overall sum for extensions of the colouring that make φ true.

Proof (of Theorem 7.1). We only prove the special case of $m = 1$, the extension to the general case is straightforward. Let $\varphi(X)$ be an MSO formula of quantifier depth d . The proof of Theorem 6.1 shows how one can obtain a dynamic program that $(\text{AC}^1, \log n)$ -maintains the model checking problem MC_ψ for $\psi \stackrel{\text{def}}{=} \exists X \varphi$. We adapt this proof, and reuse its notation, in order to obtain such a dynamic program for OPT_φ , using almost the same auxiliary relations. Together with Theorem 4.1 and Proposition 3.2, the result follows.

In the following, we sketch the proof idea. We consider φ to be an MSO sentence over the signature $\{E, X\}$. Let $G^{+X} = (V, E, X)$ be an arbitrary expansion of a graph G with a center C of connection width w , for some constant w . By Theorem 6.3 there is a number d' and an MSO sentence ψ such that for every $\mathcal{B}^{+X} \in \mathcal{S}(G^{+X}, C, w, d')$ it holds that $G^{+X} \models \varphi$ if and only if $\mathcal{B}^{+X} \models \psi$. So, the formula ψ uses the type information on the petals provided by \mathcal{B}^{+X} as well as $G^{+X}[C]$ directly to check whether the relation X represents a *feasible* solution of the problem OPT_φ . In the proof of Theorem 6.1 we explained how to obtain a first-order formula χ from ψ such that $(G, \text{Aux}, \text{Sub}) \models \chi \Leftrightarrow \mathcal{B} \models \psi$, for the auxiliary database Aux maintained by the dynamic program constructed in the proof of Theorem 6.1 and a relation Sub encoding subsets of C .

Let $\mathcal{B} \in \mathcal{S}(G, C, w, d' + 1)$ be an MSO indicator structure for G . Our goal is to maintain some relations that augment the type information provided by \mathcal{B} such that a formula χ'

⁶The adaptation to maximisation problems is straightforward.

similar to χ can “guess” a relation X , check that it is a feasible solution, compute its size and verify that no feasible solution of smaller size exists. Of course, a relation X of unrestricted size cannot be quantified in first-order logic, even in the presence of Sub, but we will see that the restriction of X to C and the type information on the petals can be quantified, which is sufficient for our purpose.

We now give the details of the construction. The structure \mathcal{B} contains relations R_τ such that $u_j^i \in R_\tau$ if and only if the depth- $(d' + 1)$ MSO type of $(\mathcal{A}_i, \bar{u}^i)$ over signature $\Sigma = \{E, c_1, \dots, c_{3\ell+1}\}$ is τ , where the subgraph \mathcal{A}_i over universe D_i and the tuple \bar{u}^i are defined as in Subsection 6.1. We say that a depth- d' type τ' over signature $\Sigma^{+X} \stackrel{\text{def}}{=} \Sigma \cup \{X\}$ can be realised in $(\mathcal{A}_i, \bar{u}^i)$ by a set $A_i \subseteq D_i$, if τ' is the depth- d' MSO type of $(\mathcal{A}_i, \bar{u}^i, A_i)$. If $(\mathcal{A}_i, \bar{u}^i)$ has depth- $(d' + 1)$ MSO type τ , the existence of such a set is equivalent to the statement $\exists X \alpha_{\tau'} \in \tau$. We note that τ' already determines whether $u_j^i \in A_i$ shall hold, for each constant u_j^i from the tuple \bar{u}^i .

The dynamic program maintains relations $\#R_{\tau'}$ and $Q_{\tau'}$, for each depth- d' MSO type over Σ^{+X} . The relations $\#R_{\tau'}$ give the minimal size of a set that realises the type τ' . So, if τ' can be realised in $(\mathcal{A}_i, \bar{u}^i)$ by some set A , and s is the minimal size of such a set, then $\#R_{\tau'}$ shall contain the tuple (u_1^i, v_s) , where v_s is the $(s + 1)$ -th element⁷ with respect to \leq . Furthermore, for the lexicographically minimal set A of this kind and size s , $Q_{\tau'}$ shall contain all tuples (u_1^i, a) , where $a \in A$.

We construct a first-order formula χ' from χ that is able to define an *optimal* solution X for OPT_φ from $(G, \text{Aux}, \text{Sub})$ expanded by the relations $\#R_{\tau'}$ and $Q_{\tau'}$. First, this formula quantifies for each depth- d' MSO type τ' the set of identifiers u_1^i such that X realises τ' in $(\mathcal{A}_i, \bar{u}^i)$ and checks consistency: as for each node $v \in C$ that appears in \bar{u}^i the type τ' already determines whether $v \in X$ shall hold, the respective types need to agree for nodes that appear in multiple petals. For each u_1^i the assigned type also needs to be realisable in the respective substructure $(\mathcal{A}_i, \bar{u}^i)$, which can be checked using the relations R_τ of \mathcal{B} . Using this information, χ' can apply χ to check that the implied set X is a feasible solution. With the help of $\#R_{\tau'}$ it can compute the size of X , as $\text{FO}(+, \times)$ is able to add up logarithmically many numbers [Vol99, Theorem 1.21] and C is only of logarithmic size in $|V|$. Also χ' checks that no other assignment of types τ' to identifier nodes results in feasible solutions of smaller size. Finally, χ' uses the relations $Q_{\tau'}$ to actually return an optimal solution X .

Building on the proof of Theorem 6.1, it remains to show that the additional auxiliary relations $\#R_{\tau'}$ and $Q_{\tau'}$ can be initialised and maintained. Actually, we maintain similarly defined relations $\#D_{\tau'}$ and $F_{\tau'}$, the relations $\#R_{\tau'}$ and $Q_{\tau'}$ are then first-order definable by the dynamic program using these relations.

Let δ be a triangle such that $G(\delta)$ has at least one inner node. Similar to the relations D_τ used in the proof of Theorem 6.1, here a relation $\#D_{\tau'}$ contains the tuple $(\bar{v}(\delta), u)$ if and only if (1) the depth- d' MSO type τ' is realisable in $(G(\delta), \bar{v}(\delta))$, and (2) u is the $(s + 1)$ -th element with respect to \leq , where s is the minimal size of a set that realises this type. Furthermore, for the lexicographically minimal set A of this kind and size s , $F_{\tau'}$ contains all tuples $(\bar{v}(\delta), a)$, where $a \in A$. It is clear that these relations suffice to define the relations $\#R_{\tau'}$ and $Q_{\tau'}$, given the other relations of the proof of Theorem 6.1.

The proof of Theorem 6.1 can be extended to show that the initial versions of these auxiliary relations can be computed in AC^1 . For the inductive step of this computation,

⁷We ignore the case that the size could be as large as $|V|$, which can be handled by some additional encoding.

a type τ' realisable in a structure $(G(\delta), \bar{v}(\delta))$ might be achievable by a finite number of combinations of types of its substructures. Here, the overall size of the realising set for X needs to be computed and the minimal solution needs to be picked. This is possible by a $\text{FO}(+, \times)$ -formula since the number of possible combinations is bounded by a constant depending only on d and k .

The updates of the auxiliary relations are exactly as in the proof of Theorem 6.1. Since D_τ needs no updates there, neither $\#D_{\tau'}$ nor $F_{\tau'}$ do, here. \square

From the proof it is easy to see that a dynamic program can also maintain the *size* s of an optimal solution, either implicitly as $\sum_{j=1}^m |Q_j|$ for distinguished relations Q_j , or as $\{v_s\}$. Additionally, it can easily be adapted for optimisation problems on *weighted* graphs, where nodes and edges have polynomial weights in n .

8. CONCLUSION

In this paper, we introduced a strategy for maintaining queries by periodically restarting its computation from scratch and limiting the number of change steps that have to be taken into account. This has been captured in the notion of (\mathcal{C}, f) -maintainable queries, and we proved in particular that all $(\text{AC}^1, \log n)$ -maintainable, almost domain independent queries are actually in DYNFO. As a consequence, decision and optimisation queries definable in MSO- and GSO-logic are in DYNFO for graphs of bounded treewidth. For this, we stated a Feferman-Vaught-type composition theorem for these logics, which might be interesting in its own right. Though we phrase our results for MSO and GSO for graphs only, their proofs translate swiftly to general relational structures.

Apart from this paper, this strategy is already used in [DMVZ18] to prove that the reachability query can be maintained dynamically under insertions and deletions of a non-constant number of edges per change step.

We believe that our strategy will find further applications. For instance, it is conceivable that interesting queries on planar graphs, such as the shortest-path query, can be maintained for a bounded number of changes using auxiliary data computed by an AC^1 algorithm (in particular since many important data structures for planar graphs can be constructed in logarithmic space and therefore in AC^1).

ACKNOWLEDGMENT

We thank an anonymous referee for valuable suggestions that greatly simplified the proof of Theorem 6.3.

REFERENCES

- [BCL08] Achim Blumensath, Thomas Colcombet, and Christof Löding. Logical theories and compatible operations. In Jörg Flum, Erich Grädel, and Thomas Wilke, editors, *Logic and Automata: History and Perspectives [in Honor of Wolfgang Thomas]*, volume 2 of *Texts in Logic and Games*, pages 73–106. Amsterdam University Press, 2008.
- [BJM17] Patricia Bouyer-Decitre, Vincent Jugé, and Nicolas Markey. Courcelle’s theorem made dynamic. *CoRR*, abs/1702.05183, 2017.
- [Cou90] Bruno Courcelle. The monadic second-order logic of graphs. I. recognizable sets of finite graphs. *Inf. Comput.*, 85(1):12–75, 1990.

- [Cou94] Bruno Courcelle. The monadic second order logic of graphs VI: on several representations of graphs by relational structures. *Discrete Applied Mathematics*, 54(2-3):117–149, 1994.
- [DKM⁺15] Samir Datta, Raghav Kulkarni, Anish Mukherjee, Thomas Schwentick, and Thomas Zeume. Reachability is in DynFO. In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part II*, volume 9135 of *Lecture Notes in Computer Science*, pages 159–170. Springer, 2015.
- [DKM⁺18] Samir Datta, Raghav Kulkarni, Anish Mukherjee, Thomas Schwentick, and Thomas Zeume. Reachability is in DynFO. *J. ACM*, 65(5):33:1–33:24, August 2018.
- [DMS⁺17] Samir Datta, Anish Mukherjee, Thomas Schwentick, Nils Vortmeier, and Thomas Zeume. A strategy for dynamic programs: Start over and muddle through. In Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl, editors, *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland*, volume 80 of *LIPICs*, pages 98:1–98:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017.
- [DMVZ18] Samir Datta, Anish Mukherjee, Nils Vortmeier, and Thomas Zeume. Reachability and distances under multiple changes. In Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella, editors, *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic*, volume 107 of *LIPICs*, pages 120:1–120:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018.
- [DST95] Guozhu Dong, Jianwen Su, and Rodney W. Topor. Nonrecursive incremental evaluation of datalog queries. *Ann. Math. Artif. Intell.*, 14(2-4):187–223, 1995.
- [EGT16] Michael Elberfeld, Martin Grohe, and Till Tantau. Where first-order and monadic second-order logic coincide. *ACM Trans. Comput. Log.*, 17(4):25:1–25:18, 2016.
- [EJT10] Michael Elberfeld, Andreas Jakoby, and Till Tantau. Logspace versions of the theorems of Bodlaender and Courcelle. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 143–152. IEEE Computer Society, 2010.
- [Ete98] Kousha Etesami. Dynamic tree isomorphism via first-order updates. In *Proceedings of the Seventeenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS)*, pages 235–243, 1998.
- [GMS12] Wouter Gelade, Marcel Marquardt, and Thomas Schwentick. The dynamic complexity of formal languages. *ACM Trans. Comput. Log.*, 13(3):19, 2012.
- [Imm99] Neil Immerman. *Descriptive complexity*. Graduate texts in computer science. Springer, 1999.
- [Lib04] Leonid Libkin. *Elements of Finite Model Theory*. Springer, 2004.
- [Mak04] Johann A. Makowsky. Algorithmic uses of the Feferman-Vaught theorem. *Ann. Pure Appl. Logic*, 126(1-3):159–213, 2004.
- [PI97] Sushant Patnaik and Neil Immerman. Dyn-FO: A parallel, dynamic complexity class. *J. Comput. Syst. Sci.*, 55(2):199–209, 1997.
- [She75] Saharon Shelah. The monadic theory of order. *Annals of Mathematics*, pages 379–419, 1975.
- [SVZ18] Thomas Schwentick, Nils Vortmeier, and Thomas Zeume. Dynamic complexity under definable changes. *ACM Trans. Database Syst.*, 43(3):12:1–12:38, October 2018.
- [SZ16] Thomas Schwentick and Thomas Zeume. Dynamic complexity: recent updates. *SIGLOG News*, 3(2):30–52, 2016.
- [Vol99] Heribert Vollmer. *Introduction to Circuit Complexity - A Uniform Approach*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 1999.
- [WS07] Volker Weber and Thomas Schwentick. Dynamic complexity theory revisited. *Theory Comput. Syst.*, 40(4):355–377, 2007.