

Parameterized Algorithms for Modular-Width

Jakub Gajarský^{1*} Michael Lampis² Sebastian Ordyniak^{1†}

¹ Faculty of Informatics, Masaryk University, Brno, Czech Republic
 {gajarsky, ordyniak}@fi.muni.cz

² KTH Royal Institute of Technology, Stockholm, Sweden
 mlampis@kth.se

Abstract

It is known that a number of natural graph problems which are FPT parameterized by treewidth become W-hard when parameterized by clique-width. It is therefore desirable to find a different structural graph parameter which is as general as possible, covers dense graphs but does not incur such a heavy algorithmic penalty.

The main contribution of this paper is to consider a parameter called modular-width, defined using the well-known notion of modular decompositions. Using a combination of ILPs and dynamic programming we manage to design FPT algorithms for Coloring and Partitioning into paths (and hence Hamiltonian path and Hamiltonian cycle), which are W-hard for both clique-width and its recently introduced restriction, shrub-depth. We thus argue that modular-width occupies a sweet spot as a graph parameter, generalizing several simpler notions on dense graphs but still evading the “price of generality” paid by clique-width.

1 Introduction

The topic of this paper is the exploration of the algorithmic properties of some structural graph parameters. This area is typically dominated by an effort to achieve two competing goals: generality and algorithmic tractability. A good example of this tension is the contrast between treewidth and clique-width.

A large wealth of problems are known to be FPT when parameterized by treewidth [6, 5, 4]. One drawback of treewidth, however, is that this parameterization excludes a large number of interesting instances, since, in particular, graphs of small treewidth are necessarily sparse. The notion of clique-width (and its cousins rank-width [22] and boolean-width [7]) tries to ameliorate this problem by covering a significantly larger family of graphs, including many dense graphs. As it turns out though, the price one has to pay for this added generality is significant. Several natural problems which are known to be fixed-parameter tractable for treewidth become W-hard when parameterized by these measures [18, 17, 16].

It thus becomes an interesting problem to explore the trade-offs offered by these and other graph parameters. More specifically, one may ask: is there a natural graph parameter which covers dense graphs but still allows FPT algorithms for the problems lost to clique-width? This is the main question motivating this paper. We first attempt to use the recently introduced notion of shrub-depth for this role [20]. Shrubs are a restriction of clique-width which shows some hope, since it has been used to obtain improved algorithmic meta-theorems. Unfortunately, as we will establish, the hardness constructions for COLORING and HAMILTONIAN PATH used in [18] go through with small modifications for this restricted parameter as well.

*Research funded by the Czech Science Foundation under grant P202/11/0196

†Research funded by Employment of Newly Graduated Doctors of Science for Scientific Excellence (CZ.1.07/2.3.00/30.0009).

The main contribution of this paper is then the consideration of a parameter called modular-width which, we argue, nicely fills this niche. One way to define modular-width is by using the standard concept of modular decompositions (see e.g. [24]), as the maximum degree of the optimal modular decomposition tree. As a consequence, a graph's modular-width can be computed in polynomial time. Note that the concept of modular-width was already briefly considered in [8], but was then abandoned in that paper in favor of the more general clique-width. To the best of our knowledge, modular-width has not been considered as a parameter again, even though modular decompositions have found a large number of algorithmic applications, including in parameterized complexity (see [21] for a general survey and [26, 10, 1] for example applications in parameterized complexity).

We give here the first evidence indicating that modular-width is a structural parameter worthy of further study. On the algorithmic side, modular-width offers a significant advantage compared to clique-width, a fact we demonstrate by giving FPT algorithms for several variants of HAMILTONICITY and CHROMATIC NUMBER, all problems W-hard for clique-width. At the same time, we show that modular-width significantly generalizes several simpler parameters, such as neighborhood diversity [23] and twin-cover [19], which also allow FPT algorithms for these problems.

Our main algorithmic tool is a form of dynamic programming on the modular decomposition of the input graph. Unlike dynamic programming on the more standard tree decompositions, the main obstacle here is in combining the DP tables of the children of a node to compute the table for the node itself. This is in general a hard problem, but we show that it can sometimes be made tractable if every node of the decomposition has small degree, hence the parameterization by modular-width.

Even if the modular decomposition has small degree, combining the DP tables is still not necessarily a trivial problem. A second idea we rely on (in the case of HAMILTONICITY) is to use an Integer Linear Program, whose number of variables is bounded by the number of modules we are trying to combine. It is our hope that this technique, which seems quite general, will be applicable to other problems as well.

2 Preliminaries

We use standard notation from graph theory as can be found in, e.g., [9]. Let G be a graph. We denote the vertex set of G by $V(G)$ and the edge set of G by $E(G)$. Let $X \subseteq V(G)$ be a set of vertices of G . The *subgraph of G induced by X* , denoted $G[X]$, is the graph with vertex set X and edges $E(G) \cap [X]^2$. By $G \setminus X$ we denote the subgraph of G induced by $V(G) \setminus X$. Similarly, for $Y \subseteq E(G)$ we define $G \setminus Y$ to be the subgraph of G obtained by deleting all edges in Y from G . For a graph G and a vertex $v \in V(G)$, we denote by $N_G(v)$ and $N_G[v]$ the open and closed neighborhood of v in G , respectively.

2.1 Considered Problems

We consider the following problems on graphs. Let G be a graph. A *coloring* of G is a function $\lambda : V(G) \rightarrow \mathbb{N}$ such that for every edge $\{u, v\} \in E(G)$ it holds that $\lambda(u) \neq \lambda(v)$. We denote by $\lambda(G)$ the set of colors used by the coloring λ , i.e., $\lambda(G) = \{\lambda(v) : v \in V(G)\}$, and by $\Lambda(G)$ the set of all colorings of G that use at most $|V(G)|$ colors. The *chromatic number* of G , denoted by $\chi(G)$, is the smallest number c such that G has a coloring λ with $|\lambda(G)| \leq c$.

GRAPH COLORING

Input: A graph G .

Question: Compute $\chi(G)$.

Let G be a graph. A partition of G into paths is a set of disjoint paths of G whose union contains every vertex of G . We denote by $\text{ham}(G)$ the least integer p such that G has a partition into p paths.

PARTITIONING INTO PATHS

Input: A graph G .

Question: Compute $\text{ham}(G)$.

HAMILTONIAN PATH

Input: A graph G .

Question: Does G have a Hamiltonian Path?

HAMILTONIAN CYCLE

Input: A graph G .

Question: Does G have a Hamiltonian Cycle?

2.2 Parameterized Complexity

Here we introduce the relevant concepts of parameterized complexity theory. For more details, we refer to text books on the topic [12, 15, 25]. An instance of a parameterized problem is a pair (I, k) where I is the main part of the instance, and k is the parameter. A parameterized problem is *fixed-parameter tractable* if instances (I, k) can be solved in time $f(k)|I|^c$, where f is a computable function of k , and c is a constant. FPT denotes the class of all fixed-parameter tractable problems. Hardness for parameterized complexity classes is based on *fpt-reductions*. A parameterized problem L is fpt-reducible to another parameterized problem L' if there is a mapping R from instances of L to instances of L' such that (i) $(I, k) \in L$ if and only if $(I', k') = R(I, k) \in L'$, (ii) $k' \leq g(k)$ for a computable function g , and (iii) R can be computed in time $O(f(k)|I|^c)$ for a computable function f and a constant c . Central to the completeness theory of parameterized complexity is the hierarchy $\text{FPT} \subseteq \text{W}[1] \subseteq \text{W}[2] \subseteq \dots$. Each intractability class $\text{W}[t]$ contains all parameterized problems that can be reduced to a certain parameterized satisfiability problem under fpt-reductions.

2.3 Treewidth

The treewidth of a graph is defined using the following notion of a tree decomposition (see, e.g., [3]). A *tree decomposition* of an (undirected) graph $G = (V, E)$ is a pair (T, χ) where T is a tree and χ is a labeling function that assigns each tree node t a set $\chi(t)$ of vertices of the graph G such that the following conditions hold: (1) Every vertex of G occurs in $\chi(t)$ for some tree node t , (2) For every edge $\{u, v\}$ of G there is a tree node t such that $u, v \in \chi(t)$, and (3) For every vertex v of G , the tree nodes t with $v \in \chi(t)$ form a connected subtree of T . The *width* of a tree decomposition (T, χ) is the size of a largest bag $\chi(t)$ minus 1 among all nodes t of T . A tree decomposition of smallest width is *optimal*. The *treewidth* of a graph G is the width of an optimal tree decomposition of G .

2.4 Shrub-depth

The recently introduced notion of *shrub-depth* [20] is the “low-depth” variant of clique-width, similar to the role that tree-depth plays with respect to treewidth.

Definition 1. We say that a graph G has a tree-model of m colors and depth $d \geq 1$ if there exists a rooted tree T (of height d) such that

1. the set of leaves of T is exactly $V(G)$,
2. the length of each root-to-leaf path in T is exactly d ,
3. each leaf of T is assigned one of m colors (this is not a graph coloring, though),
4. and the existence of a G -edge between $u, v \in V(G)$ depends solely on the colors of u, v and the distance between u, v in T .

The class of all graphs having a tree-model of m colors and depth d is denoted by $\mathcal{TM}_m(d)$.

Definition 2. A class of graphs $\mathcal{G}$ has shrub-depth d if there exists m such that $\mathcal{G} \subseteq \mathcal{TM}_m(d)$, while for all natural m it is $\mathcal{G} \not\subseteq \mathcal{TM}_m(d-1)$.

Note that Definition 2 is asymptotic as it makes sense only for infinite graph classes. Particularly, classes of shrub-depth 1 are known as the graphs of bounded *neighborhood diversity* in [23], i.e., those graph classes on which the twin relation on pairs of vertices (for a pair to share the same set of neighbors besides this pair) has a finite index.

2.5 Properties of Shrub-depth

In this section we will show some technical results that will later help us for our hardness proofs. For two graphs G and H , let $G \bowtie H$ denote the set of graphs such that they consist of a copy of G , a copy of H (disjoint from G) and arbitrary edges between vertices of G and H .

Proposition 1. *Let G be a graph with tree-model TM_G of height d which uses m colors and let H be a graph. Then there exists, for every $G' \in G \bowtie H$, a tree-model of height d which uses at most $m \cdot 2^{|V(H)|} + |V(H)|$ colors.*

Proof. For simplicity, we refer to copies of G and H in G' by G and H . We fix arbitrary ordering of $V(H)$.

Since each vertex of G is adjacent to some subset of $V(H)$ in G' , there are at most $2^{V(H)}$ “types” of vertices in G according to their neighborhoods in H . We extend the colors used in the leaves of TM_G by $|V(H)|$ bits and for each leaf l of TM_G , we set these bits as follows: we set the i -th bit to 1 iff the vertex v of G which corresponds to l in TM_G is adjacent to the i -th vertex of H in G' . We denote this new tree-model by TM_G^+ .

Next, we construct a tree-model TM_H of height one for H – it consists of a root vertex and $V(|H|)$ leaves, each of different color (we use new colors, different from those of TM_G^+).

Now we construct a tree-model $\text{TM}_{G'}$ of G' – we connect the root of TM_H to the root of TM_G^+ by a path of length $d - 1$ (to ensure that the leaf-to-root distance is the same in the whole tree-model; if the height of TM_G^+ is one, we identify the root of TM_H with the root of TM_G^+ .) It is not hard to see that $\text{TM}_{G'}$ is a tree-model of G' :

- the edges in the “ H ”-part of G' depend only on the colors of the subtree TM_H of $\text{TM}_{G'}$,
- the edges in the “ G ”-part of G' depend only on the colors of the subtree TM_G^+ of $\text{TM}_{G'}$ and the original distances in TM_G (simply by looking at the colors of leaves before adding new bits, i.e. ignoring additional bits),
- the edges between the vertices of “ H ”- and “ G ”-parts depend only on the colors used in the “ H ”-part and the newly added bits in the “ G ”-part of G' (all are at distance $2d$ in $\text{TM}_{G'}$),
- none of these three edge-dependencies affect the other two.

□

Proposition 2. *Let G be a graph, S a subset of its vertices, and TM_G its tree-model of height d_G and m_G colors. Let H be a graph, R a subset of its vertices, and TM_H its tree-model of height d_H and m_H colors. Assume that $d_H \leq d_G$. Then the graph G' obtained by taking a copy of G , a copy of H , and making each vertex from S adjacent to each vertex of R has a tree-model $\text{TM}_{G'}$ of height d and at most $2m_G + 2m_H$ colors.*

Proof. We assume that the sets of colors in TM_G and TM_H are disjoint. The proof uses the same idea as the proof of Proposition 1. First, we extend the colors in both TM_H and TM_G by one additional bit. For a leaf of TM_G , we set this bit to 1 iff the corresponding vertex of G is in S . Similarly, for a leaf of TM_H , we set this bit to 1 iff the corresponding vertex of H is in R . We denote the tree-models obtained in this way by TM_G^+ and TM_H^+ .

Now we construct $\text{TM}_{G'}$ by connecting the root of TM_H^+ to the root of TM_G^+ by a path of length $d_G - d_H$ (or identify these two roots if $d_G = d_H$). The distances between the pairs of leaves in both “ G ”- and “ H ”-parts of $\text{TM}_{G'}$ are the same as they were in TM_G and TM_H ; their original colors can be “recovered” by simply ignoring the newly added bit. The existence of an edge between a vertex from G

and a vertex from H in G' is determined by $\text{TM}_{G'}$ by the value of the additional bit and the fact that their colors come from different tree-models (the sets of colors were disjoint); the distance between them is always $2d$. Thus, $\text{TM}_{G'}$ is indeed a tree-model of G' . Since the numbers d_G and d_H were doubled by adding a new bit to each color and since $\text{TM}_{G'}$ does not use any other color, the Proposition is now proved. $\square$

Proposition 3. *Let G be a graph with a tree-model TM_G of height d and m colors and let R be a subset of $V(G)$. Let H be a graph and S be a subset of $V(H)$. Let G' be the graph obtained from G by creating $|R|$ copies of H indexed by vertices of R and making each $v \in R$ adjacent to the vertices S_v of H_v . Then G' has a tree-model $\text{TM}_{G'}$ of height $d + 1$ and $m + |V(H)|$ colors.*

Proof. We extend the set of colors used in TM_G by $|V(H)|$ new colors. The construction of $\text{TM}_{G'}$ from TM_G is simple – for each $v \in R$, replace its corresponding leaf l in TM_G by a tree of height one which has:

- leaf l' with the same color as l
- $|V(H)|$ leaves, each with different color from the newly created colors

For each $v \notin R$, subdivide once the edge to which its corresponding leaf was adjacent in TM_G (to have the same leaf-to-root distance in the whole $\text{TM}_{G'}$).

It is easy to see that $\text{TM}_{G'}$ is indeed a tree-model of G' – for each $v \in R$ the leaves corresponding to v and $V(H_v)$ in $\text{TM}_{G'}$ are at distance two from each other and the adjacency depends only on the colors. The color of each leaf corresponding to $v \in V(G)$ remained the same, and the distance between each pair of such leaves was increased by 2. $\square$

Lemma 1. *Let $\mathcal{G}$ be a class of graphs which admit a tree-decomposition of width w and height d . Then $\mathcal{G}$ has shrub-depth $d + 1$.*

Proof. Let $G \in \mathcal{G}$. We proceed by induction on the height d of a tree-decomposition (T, χ) of G of width at most w .

If the height is 0, then the tree-decomposition consists of 1 bag which contains at most w vertices. A tree-model of height 1 for the graph induced by this bag can be easily constructed – create a tree of height 1, with 1 leaf for every vertex in the bag and assign each leaf a different color. For graphs of treewidth w , these tree-models always use at most w colors.

If the height of T is d , we take a root bag B_r (by root bag we mean any bag whose distance to any other bag is at most d) of the tree-decomposition and create a tree-model TM_r of height one for the graph induced by the bag; this requires at most w colors. After deleting this bag from the decomposition, we are left with a collection of tree-decompositions of height $d - 1$. We delete all vertices of B_r from all bags of these decompositions. By the induction hypothesis, the graphs induced by these decompositions have tree-models of height d with at most m colors (for some natural m , which depends only on d and w). Let us denote these tree-models by TM_i .

Now we construct a tree-model TM^- of height $d + 1$ by introducing a root vertex r and connecting the root of each TM_i to r . This is a tree-model of $G \setminus B_r$, which uses m colors. Since $G \in G[B_r] \bowtie G[G \setminus B_r]$, we can use Proposition 1 to conclude that there exists a tree-model TM_G of G of height $d + 1$ which uses at most $m \cdot 2^w + w$ colors. Since neither m or w depend on particular choice of G , this concludes the proof. $\square$

2.6 Modular-Width

For our algorithms we consider graphs that can be obtained from an algebraic expression that uses the following operations:

- (O1) create an isolated vertex;

(O2) the disjoint union of 2 graphs, i.e., the *disjoint union* of 2 graph G_1 and G_2 , denoted by $G_1 \oplus G_2$, is the graph with vertex set $V(G_1) \cup V(G_2)$ and edge set $E(G_1) \cup E(G_2)$;

(O3) the complete join of 2 graphs, i.e., the *complete join* of 2 graphs G_1 and G_2 , denoted by $G_1 \otimes G_2$, is the graph with vertex set $V(G_1) \cup V(G_2)$ and edge set $E(G_1) \cup E(G_2) \cup \{ \{v, w\} : v \in V(G_1) \text{ and } w \in V(G_2) \}$;

(O4) the substitution operation with respect to some graph G with vertices $v_1, \dots, v_n$, i.e., for graphs $G_1, \dots, G_n$ the *substitution* of the vertices of G by the graphs $G_1, \dots, G_n$, denoted by $G(G_1, \dots, G_n)$, is the graph with vertex set $\bigcup_{1 \leq i \leq n} V(G_i)$ and edge set $\bigcup_{1 \leq i \leq n} E(G_i) \cup \{ \{u, v\} : u \in V(G_i) \text{ and } v \in V(G_j) \text{ and } i \neq j \}$. Hence, $G(G_1, \dots, G_n)$ is obtained from G by substituting every vertex $v_i \in V(G)$ with the graph G_i and adding all edges between the vertices of a graph G_i and the vertices of a graph G_j whenever $\{v_i, v_j\} \in E(G)$.

Let A be an algebraic expression that uses only the operations (O1)–(O4). We define the *width* of A as the maximum number of operands used by any occurrence of the operation (O4) in A . It is well-known that the *modular-width* of a graph G , denoted $\text{mw}(G)$, is the least integer m such that G can be obtained from such an algebraic expression of width at most m . Furthermore, an algebraic expression of width $\text{mw}(G)$ can be constructed in linear time [27].

2.7 Integer Linear Programming

For our algorithms, we use the well-known result that INTEGER LINEAR PROGRAMMING is fixed-parameter tractable parameterized by the number of variables.

INTEGER LINEAR PROGRAMMING FEASIBILITY

Parameter: p

Input: A matrix $A \in \mathbb{Z}^{m \times p}$ and a vector $b \in \mathbb{Z}^m$.

Question: Is there a vector $x \in \mathbb{Z}^p$ such that $Ax \leq b$?

Proposition 4 ([14]). INTEGER LINEAR PROGRAMMING FEASIBILITY is *fixed-parameter tractable* and can be solved in time $O(p^{2.5p+o(p)} \cdot L)$ where L is the number of bits in the input.

3 Hardness for Problems on Shrub-depth

In this section we give evidence that the recently introduced parameter shrub-depth is not restrictive enough to obtain fixed-parameter algorithms for problems that are W[1]-hard on graphs of bounded cliquewidth. In particular, we show that GRAPH COLORING and HAMILTONIAN PATH are W[1]-hard parameterized by the number of colors (used in a tree-model of the input graph) on classes of graphs of shrub-depth 5. Note that restricting the shrub-depth means restricting the height of the tree-model that can be employed and for every restriction on the height of the tree-model the number of colors needed to model the graph gives a different parameter. In particular, if we restrict the shrub-depth to 1 the number of colors of the tree-model corresponds to the neighborhood diversity of a graph. This implies that GRAPH COLORING and HAMILTONIAN PATH become fixed-parameter tractable when parameterized by the number of colors (used in a tree-model of the input graph) on classes of graphs of shrub-depth 1 [23]. It is an interesting open question what is the least possible shrub-depth that allows for fixed-parameter algorithms for the problems GRAPH COLORING and HAMILTONIAN PATH.

Theorem 1. GRAPH COLORING parameterized by the number of colors (used in a tree-model of the input graph) is W[1]-hard on classes of graphs of shrub-depth 5.

We will prove the above theorem after showing the following lemma.

Lemma 2. EQUITABLE COLORING is W[1]-hard parameterized by the variant of treewidth, where the height of the tree (of the tree-decomposition) is at most 3.

Proof. Our proof uses the same reduction as the proof of [13, Theorem 5.]. We only show that the tree-decompositions of the graphs constructed there have height at most 3.

The graph G' constructed in [13] consists of a disjoint union of trees of height 2. Hence, the tree-decomposition of G' has height at most 3. The next step in their construction is to add an appropriate number of isolated vertices to G' . This step does not increase the height of the tree-decomposition any further.

The final step of their reduction consists of adding a clique and connecting some vertices of the clique to some vertices of the graph constructed above. Hence, a tree-decomposition of height 3 for the whole graph can be obtained by adding all vertices of the clique to every bag of the already constructed tree-decomposition. $\square$

Proof of Theorem 1. Our proof uses the same reduction from EQUITABLE COLORING on graphs of bounded treewidth as the proof of [16, Theorem 3.2]. We only show that the graphs constructed there have tree-models of height at most 5 that use at most $f(w)$ colors, where the input graph G has a tree-decomposition of width w and height 3.

For the input graph G , we consider its tree-model TM_G of height 4 and with at most m colors (for some m which does not depend on G) whose existence is guaranteed by Lemma 1 and Lemma 2. We replace each leaf of TM_G by a tree of height one which has the following leaves:

- vertex u which was replaced from the original tree-model, with the same color as u had in TM_G
- r vertices colored by S_1 to S_r
- r vertices colored by P_1 to P_r

We create a tree-model for the clique P_M of size r , where each vertex has a different color p_i , $1 \leq i \leq r$. Finally, we create tree-models of height 1 for cliques C_i ($1 \leq i \leq r$) of size $n \frac{r-1}{r}$, each of them fully colored by color c_i .

For each $i \in \{1, \dots, r\}$, the union of all vertices colored by S_i (P_i) corresponds to the set S_i (P_i) in [16, Theorem3.2], respectively. Similarly graphs $P^M, C_1, \dots, C_r$ correspond to their counterparts of the same name in [16, Theorem3.2].

We connect the roots of tree-models of $P^M, C_1, \dots, C_r$ to the root of TM_G (by a path of appropriate length, so that the leaf-to-root distance is the same in the whole tree) to obtain TM'_G .

The height of TM'_G is 5, because the height of TM_G was 4 and during the construction the height increased by one (by replacing leaves by trees of height one).

Now we describe the edges and their dependencies in TM'_G :

- all edges from the original graph remain the same, TM' clearly encodes this information, since the only thing that changed is that the distance between every pair u, v of leaves from the original tree-model increased by 2;
- we make each u adjacent to each vertex with label P_i (for all $i \in \{1, \dots, r\}$, in the whole tree-model), i.e. for all distances between leaves, all the vertices colored by colors of the original tree-model are adjacent to all the vertices colored by P_i ;
- we make each u adjacent to “its” vertices colored by S_i , (for all $i \in \{1, \dots, r\}$) i.e. for distance 2, all vertices which have colors from the original tree-model are adjacent to all the vertices colored by S_i , $i \in \{1, \dots, r\}$;
- we make each vertex colored by S_i adjacent to the whole clique P_M except for a vertex colored by p_i , i.e. for distance 10, each vertex colored by S_i is adjacent to each vertex colored by p_j , except when $i = j$;
- for distance 2, vertices colored P_i and S_j are not adjacent if $i = j$, and are adjacent if $i \neq j$. For distance higher than 2, vertices colored by P_i and S_j are adjacent, whether $i = j$ or not;
- we make, for each i , every vertex of C_i adjacent to all vertices of colored by S_i , i.e. for distance 10, vertices colored by c_i and S_i are adjacent, for $i \in \{1, \dots, r\}$

- we make, every vertex of C_i adjacent to all vertices colored by P_j , $i \neq j$, i.e. for distance 10, vertices colored by c_i and P_j , for all $i, j \in \{1, \dots, r\}$ are adjacent, except when $i = j$.

The reduction from [16, Theorem 3.2] is now completed, and TM'_G represents the resulting graph of this reduction. We used at most m colors of the tree-model of original graph and $4r$ new colors S_i, P_i, c_i, p_i . The number of colors in TM'_G therefore bounded by $m + 4r$, which concludes the proof. $\square$

Theorem 2. HAMILTONIAN PATH parameterized by the number of colors (used in a tree-model of the input graph) is $W[1]$ -hard on class of graphs of shrub-depth 5.

We observe that the reduction showing $W[1]$ -hardness of CAPACITATED DOMINATING SET on bounded treewidth given in [11] actually shows that this problem is $W[1]$ -hard with respect to the variant of treewidth, where the height of the tree (of the tree-decomposition) is at most 2.

We then proceed by showing that the reduction given in [18] from CAPACITATED DOMINATING SET parameterized by treewidth to HAMILTONIAN PATH parameterized by cliquewidth produces graphs of shrub-depth at most 5.

Proposition 5. CAPACITATED DOMINATING SET is $W[1]$ -hard when parametrized by the variant of treewidth, where the height of the tree (of the tree-decomposition) is at most 2.

Proof. Our proof uses the same reduction as the proof of [11, Theorem 1.]. We only show that the tree-decompositions of the graphs constructed there have height at most 2. Indeed, in the graph H constructed in [11, Section 3.1.], delete the sets S_i and $S_{i,j}$ and all the vertices $x_i, y_i, z_i, p_{i,j}, q_{i,j}$. After this, we are left with a forest of height 1, which has a tree-decomposition of height 2. After adding all deleted vertices to each bag, we obtain a tree-decomposition of the whole graph of height 2 and of width $O(k^4)$. $\square$

Proof of Theorem 2. We use the reduction from [18, Theorem 5.1]. We show that when the instances of CAPACITATED DOMINATING SET have a tree-decomposition of height 2, then the graphs we obtain by their reduction have shrub-depth at most 5.

First we argue that, when we begin with the graph G which has a tree-decomposition of height 2, the graph $G'(c')$ has a tree-decomposition of height 3. The graph $G'(c')$ is the graph obtained after the first step of the reduction with an assumption that the capacity of each vertex is exactly one. This graph is basically obtained by replacing each vertex of G by a gadget and replacing each edge of G by two gadgets. Since these gadgets have constant size, one can just replace each occurrence of vertex v in the tree-decomposition of G by a gadget, and “hang” the bag containing the pair of gadgets corresponding to an edge u, v in G to any bag which contained these two vertices in the original tree-decomposition. This increases the height of the tree-decomposition by at most one.

By Lemma 1, the resulting graph has a tree-model TM_1 of height 4 and at most m colors, for some natural m (which depends only on the height of the newly created tree-decomposition (3) and its width w' , which is bounded by a constant multiple of w).

Graph $G'(c)$ is obtained from $G'(c')$ by adding vertices, each of which is a twin of some vertex from $G'(c')$. Even though this can in general affect the number of colors in the tree-model of a graph, in this particular instance it is not the case – each vertex to which we add twin vertices has two neighbors and is the only vertex which has this neighborhood. We can therefore construct the tree-model TM_2 of $G'(c)$ by simply adding siblings (with the same color) to the corresponding leafs of TM_1 .

The second part of the reduction consists of introducing a graph which consists of 2 new vertices connected by an appropriate number of paths of length two. Let Y denote the set of middle vertices of these paths. All vertices from Y are connected to some subset of $V(G'(c))$. It is easy to construct a tree-model of height 1 with 2 colors for this graph (root, one color for vertices of Y and one color for remaining two vertices) and therefore by Proposition 3 there exists a tree-model TM_3 of the resulting graph of height 4 and with $2m + 4$ colors.

Next, to some subset of vertices of the graph induced by TM_3 , a copy of a gadget of fixed size c is connected (in the same fashion for all vertices). The tree-model of this gadget of height 1 with c colors is easy to construct and therefore by Proposition 2 the tree-model TM_4 of the graph obtained in this way has height 5 and at most $2m + 4 + c$ colors.




Figure 1: Relationship between vertex cover (vc), neighborhood diversity (nd), twin-cover (tw), modular-width (mw), shrub-depth (sd), treewidth (tw), and clique-width (cw). Arrows indicate generalizations, e.g., modular-width generalizes both neighborhood diversity and twin-cover.

After this, two specific vertices x and y of each of these newly added gadgets are connected to the whole Y . By the construction from the proof of Proposition 2 each vertex from Y now has one of two colors in TM_4 and no vertex outside Y has any of these colors. By the construction from the proof of Proposition 3 the copies of vertices x and y have the same colors in the whole TM_4 , and no other vertex has this color. It follows that these newly added edges depend only on the colors of the leaves of the tree-model.

The final step of reduction consists of adding an independent set, which is connected to some subset of G' . An independent set has tree-model of height 1 and 1 color and therefore by Proposition 2 the resulting graph has tree-model TM_5 of height 5 and at most $2(2m + 4 + c) + 2$ colors.

Since neither m nor c depend on the choice of G , the theorem is now proved. $\square$

4 Modular-Width and Other Parameters

In this section we study the relationships of modular-width, shrub-depth and other important width parameters. Of particular importance is the observation that modular-width generalizes the recently introduced parameters neighborhood diversity [23] and twin-cover [19]. Both of these parameters have been introduced to obtain FPT algorithms on dense graphs for problems that are hard for clique-width. Figure 1 summarizes these relationships. Most of these relationships are well-known or have recently been shown in [23, 19, 20, 8]. Consequently, we only show the relationships whose proofs cannot be found anywhere else.

Theorem 3. *Let G be a graph. Then $\text{mw}(G) \leq \text{nd}(G)$ and $\text{mw}(G) \leq 2^{\text{tc}(G)} + \text{tc}(G)$. Furthermore, both inequalities are strict, i.e., there are graphs with bounded modular-width and unbounded neighborhood diversity (or unbounded twin-cover number).*

Proof. Let G be a graph. Using the definition of neighborhood diversity from [23] it follows that G has a partition $\{V_1, \dots, V_{\text{nd}(G)}\}$ of its vertex set such that for every $1 \leq i \leq \text{nd}(G)$ it holds that the graph $G[V_i]$ is either a clique or an independent set and for every $1 \leq i < j \leq \text{nd}(G)$, either all vertices in V_i are adjacent to all vertices of V_j or G contains no edges between vertices in V_i and vertices in V_j . Let G' be the graph with vertex set $v_1, \dots, v_{\text{nd}(G)}$ and an edge between v_i and v_j if and only if the graph G contains all edges between vertices in V_i and vertices in V_j . Then $G = G'(G[V_1], \dots, G[V_{\text{nd}(G)}])$. Furthermore, because for every $1 \leq i \leq \text{nd}(G)$, V_i is either a clique or an independent set, we can obtain the graph $G[V_i]$ from an algebraic expression A_i that uses only the operations (O1)–(O3). Substituting the algebraic expressions A_i for every $G[V_i]$ into $G'(G[V_1], \dots, G[V_{\text{nd}(G)}])$ gives us the desired algebraic expression for G of width $\text{nd}(G)$.

Let G be a graph. Using the definition of twin-cover from [19] it follows that there is a set C of at most $\text{tc}(G)$ vertices of G such that every component C' of $G \setminus C$ is a clique and every vertex in C' is connected to the same vertices in C . Let $C_1, \dots, C_l$ be sets of components of $G \setminus C$ such that 2

components of $G \setminus C$ are contained in the same set C_i if and only if their vertices have the same neighbors in C . Because there are at most $2^{|C|}$ possible such neighborhoods, we obtain that $l \leq 2^{|C|}$. Let G' be the graph with vertices $C \cup \{c_1, \dots, c_l\}$ and edges $E(G[C]) \cup \{\{c_i, v_i\} : v_i \in N_G(C_i)\}$. Then $G = G'(v_1, \dots, v_{|C|}, G[C_1], \dots, G[C_l])$. Furthermore, because the graphs $G[C_i]$ are disjoint unions of cliques, we can obtain each of these graphs from an algebraic expression A_i that uses only the operations (O1)–(O3). Substituting the algebraic expressions A_i for every $G[C_i]$ into $G'(v_1, \dots, v_{|C|}, G[C_1], \dots, G[C_l])$ gives us the desired algebraic expression for G of width $2^{|C|} + |C| \leq 2^{\text{tc}(G)} + \text{tc}(G)$.

To see that both inequalities are strict we refer the reader to [20, Example 5.4 a)]. The example exhibits a family of co-graphs, i.e., graphs of modular-width 0, with unbounded neighborhood diversity and unbounded twin-cover number. $\square$

The following theorem shows that modular-width and shrub-depth are orthogonal to each other.

Theorem 4. *There are classes of graphs with unbounded modular-width and bounded shrub-depth and vice versa.*

Proof. Let S_n be the graph obtained from a star (with n leaves) after subdividing every edge exactly once and let $\mathcal{S}$ be the class of all these graphs. Because S_n is a prime graph, its modular-width equals its number of vertices, i.e., $2n + 1$. Consequently, $\mathcal{S}$ has unbounded modular-width. We next show that $\mathcal{S}$ has shrub-depth at most 2 by showing that S_n has a tree-model of height 2 that uses at most 3 colors. The tree-model consists of a root r that has $n + 1$ children, such that n of these children (which correspond to the n edges around the center of the star) have 2 children themselves colored by color 1 and 2, respectively, and the remaining child has only 1 child (which corresponds to the center of the star) which is colored by color 2. Then S_n can be defined in this tree-model.

To show the “vice versa” part of the theorem, we refer the reader to [20, Example 5.4 a)]. The example gives a family of co-graphs, i.e., graphs with modular-width 0, that has unbounded shrub-depth. $\square$

The next theorem shows that also shrub-depth generalizes neighborhood diversity and twin-cover.

Theorem 5. *Let $\mathcal{G}$ be a class of graphs. If $\mathcal{G}$ has bounded neighborhood diversity or bounded twin-cover number then $\mathcal{G}$ has shrub-depth at most 2. Furthermore, there are classes of graphs that have unbounded neighborhood diversity and twin-cover number but shrub-depth 2.*

Proof. Suppose that $\mathcal{G}$ has bounded neighborhood diversity, i.e., there is some natural number c such that $\text{nd}(G) \leq c$ for every $G \in \mathcal{G}$. We show that every graph $G \in \mathcal{G}$ has a tree-model of height at most 1 that uses at most $\text{nd}(G) \leq c$ colors. Using the definition of neighborhood diversity from [23] it follows that G has a partition $\{V_1, \dots, V_{\text{nd}(G)}\}$ of its vertex set such that for every $1 \leq i \leq \text{nd}(G)$ it holds that the graph $G[V_i]$ is either a clique or an independent set and for every $1 \leq i < j \leq \text{nd}(G)$, either all vertices in V_i are adjacent to all vertices of V_j or G contains no edges between vertices in V_i and vertices in V_j . Then the tree-model for G consists of a root r and 1 leaf for every vertex v in G with color i if $v \in V_i$.

Now, suppose that $\mathcal{G}$ has bounded twin-cover, i.e., there is some natural number c such that $\text{tc}(G) \leq c$ for every $G \in \mathcal{G}$. We show that every graph $G \in \mathcal{G}$ has a tree-model of height at most 2 that uses at most $2^{\text{tc}(G)} + \text{tc}(G) \leq 2^c + c$ colors. Using the definition of twin-cover from [19] it follows that there is a set $W = \{w_1, \dots, w_{\text{tc}(G)}\}$ of $\text{tc}(G)$ vertices of G such that every component C' of $G \setminus W$ is a clique and every vertex in C' is connected to the same vertices in W . Let $C_1^1, \dots, C_{p_1}^1, \dots, C_1^l, \dots, C_{p_l}^l$ be all the components of $G \setminus W$ such that 2 components $C_{i_1}^{j_1}$ and $C_{i_2}^{j_2}$ have the same neighborhood in W if and only if $j_1 = j_2$. Because there are at most $2^{|W|}$ possible such neighborhoods, we obtain that $l \leq 2^{|W|}$. We construct a tree-model of G as follows. We start with the root node r , which has 1 child, say C_i^j , for every component of $G \setminus W$, and 1 child (which is also a leaf of the tree-model), say w_i , with color $l + i$ for every $1 \leq i \leq |W|$. Finally, every node C_i^j has $|V(C_i^j)|$ children (which are also leaves of the tree-model) with color j . This finishes the construction of the tree-model for G .

To see that there are classes of graphs that have unbounded neighborhood diversity and twin-cover number but shrub-depth 2, consider the class $\mathcal{S}$ from the proof of Theorem 4. As shown in Theorem 4

$\mathcal{S}$ has unbounded modular-width but shrub-depth 2. It now follows from Theorem 3 that $\mathcal{S}$ has also unbounded neighborhood diversity and unbounded twin-cover number. $\square$

5 Algorithms on Modular-width

In this section we show that PARTITIONING INTO PATHS, HAMILTONIAN PATH, HAMILTONIAN CYCLE, and COLORING are fixed-parameter tractable parameterized by the modular-width of the input graph. Our algorithms use a bottom-up dynamic programming approach along the parse-tree of an algebraic expression as defined in Section 2.6. That is for every node of such a parse-tree we compute a solution (or a record representing a solution) given solutions (or records) for the children of the node in the parse-tree. The running time of our algorithms is then the number of nodes in the parse-tree times the maximum time spend at any node of the parse-tree. Because the number of nodes of a parse-tree is linear in the number of vertices of the created graph, it suffices to bound the maximum time spend at any node of the parse-tree. Furthermore, because the operations (O1)–(O3) can be replaced by one operation of type (O4) that uses at most 2 operands, we only need to bound the time spend to compute a record for the graph obtained by operation (O4). To avoid cumbersome run-time bounds we use the notation O^+ to suppress poly-logarithmic factors, i.e., we write $O^+(f)$ when we have $O(f \log^d f)$ for some constant d .

5.1 Coloring

This section is devoted to a proof of the following theorem. Recall the definition of GRAPH COLORING and related notions from Section 2.1.

Theorem 6. GRAPH COLORING parameterized by the modular-width of the input graph is fixed-parameter tractable.

As outlined above we only need to bound the time spend to compute a record for a node of type (O4) of the parse-tree. In the case of GRAPH COLORING a record is simply the chromatic number of the graph. Hence, we will have shown the theorem after showing the following lemma.

Lemma 3. Let G be a graph with vertices $v_1, \dots, v_n$, $G_1, \dots, G_n$ be graphs, and $H := G(G_1, \dots, G_n)$. Then $\chi(H)$ can be computed from $\chi(G_1), \dots, \chi(G_n)$ in time $O^+(2^n n^2 \max_{1 \leq i \leq n} \chi(G_i))$.

We will prove the lemma by reducing the coloring problem to the following problem.

MAX WEIGHTED PARTITION

Input: An n -element set N and functions $f_1, \dots, f_k$ from the subsets of N to integers from the range $[-M, M]$.

Question: A k -partition $(S_1, \dots, S_k)$ of N that maximizes $f_1(S_1) + \dots + f_k(S_k)$.

Proposition 6. ([2, Theorem 4.]) MAX WEIGHTED PARTITION can be solved in time $O^+(2^n k^2 M)$.

To simplify the reduction to MAX WEIGHTED PARTITION we need the following Proposition and Lemma.

Proposition 7 ([23]). Let G be a graph with vertices $v_1, \dots, v_n$ and $s_1, \dots, s_n$ be natural numbers. Then $\chi(G(K_{s_1}, \dots, K_{s_n})) = \min_{\lambda \in \Lambda(G)} (\sum_{c \in \lambda(G)} \max\{s_i : v_i \in \lambda^{-1}(c)\})$.

Lemma 4. Let G be a graph with vertices $v_1, \dots, v_n$, $G_1, \dots, G_n$ be graphs, $H_K := G(K_{\chi(G_1)}, \dots, K_{\chi(G_n)})$, and $H := G(G_1, \dots, G_n)$. Then $\chi(H_K) = \chi(H)$.

Proof. We will show that for every coloring λ of H there is a coloring λ_K of H_K that uses no more colors, i.e., $|\lambda_K(H_K)| \leq |\lambda(H)|$, and vice versa. Let λ be a coloring for H . Then for every $1 \leq i \leq n$ the number of colors used to color the copy of G_i in H is at least $\chi(G_i)$. Hence, we can use a subset of these colors to color the copy of $K_{\chi(G_i)}$ in H_K .

For the reverse direction, let λ_K be a coloring for H_K . Then for every $1 \leq i \leq n$ the number of colors used to color the copy of $K_{\chi(G_i)}$ in H_K is $\chi(G_i)$. Hence, we can use the same colors to color the copy of G_i in H . $\square$

We can now proceed with a proof of Lemma 3.

of Lemma 3. We reduce the coloring problem to the MAX WEIGHTED PARTITION problem as follows: We set $N := V(G)$ and $f_1(S) = \dots = f_k(S) = -\max\{\chi(G_i) : v_i \in S\}$ for every subset S of N . It follows from Proposition 7 and Lemma 4 that the maximum weight of a partition of this instance corresponds to the chromatic number $\chi(H)$. Hence, the lemma follows from Proposition 6. $\square$

5.2 Partitioning into Paths

This section is devoted to a proof of the following theorem. Recall the definition of PARTITIONING INTO PATHS and related notions from Section 2.1.

Theorem 7. PARTITIONING INTO PATHS (and hence also HAMILTONIAN PATH and HAMILTONIAN CYCLE) parameterized by the modular width of the input graph is fixed-parameter tractable.

As outlined above we only need to bound the time spend to compute a record for a node of type (O4) of the parse-tree. In the case of PARTITIONING INTO PATH a record of a graph G is the pair $(\text{ham}(G), |V(G)|)$. Hence, we will have shown the theorem after showing the following lemma. From now on we will assume that G is a graph with vertices $v_1, \dots, v_n$, $G_1, \dots, G_n$ are graphs, $H = G(G_1, \dots, G_n)$, and $m = |E(G)|$.

Lemma 5. Given the graph G and the pairs $(\text{ham}(G_1), |V(G_1)|), \dots, (\text{ham}(G_n), |V(G_n)|)$ the pair $(\text{ham}(H), |V(H)|)$ can be computed in time $O^+\left(\text{ham}(H) \left((m+n)2^n + n(2(m+n))^{5(m+n)+o(m+n)}\right)\right)$

The remainder of this section is devoted to a proof of this lemma.

For a graph G and an integer i we define the graph $G \oplus i$ as the graph with vertex set $V(G) \cup \{1, \dots, i\}$ and edge set $E(G) \cup \{\{v, j\} : v \in V(G) \text{ and } 1 \leq j \leq i\}$, i.e., the graph $G \oplus i$ is obtained from G by adding i vertices and connect them to every vertex in G .

Proposition 8. Let G be a graph and $h(G) = \min\{i : G \oplus i \text{ has a Hamiltonian cycle}\}$. Then $\text{ham}(G) = h(G)$.

Proof. We first show that $h(G) \leq \text{ham}(G)$. Let $\{P_1, \dots, P_{\text{ham}(G)}\}$ be a partition of G into paths. Then $G \oplus \text{ham}(G)$ contains the Hamiltonian cycle $1, P_1, 2, P_2, \dots, \text{ham}(G), P_{\text{ham}(G)}, 1$, as required. It remains to show that $\text{ham}(G) \leq h(G)$. Let C be a Hamiltonian cycle in $G \oplus h(G)$. Then $C \setminus \{1, \dots, h(G)\}$ is a partition of G into at most $h(G)$ disjoint path, as required. $\square$

A slightly less general version of the following lemma has already been proven in [23].

Lemma 6. Let HAMILTONIAN CYCLE be the ILP with variables $\{e_{ij}, e_{ji} : \{v_i, v_j\} \in E(G)\}$ and constraints:

For every $1 \leq i \leq n$:

- (1) $\sum_{j \in \{l : v_l \in N_G(v_i)\}} e_{ij} = \sum_{j \in \{l : v_l \in N_G(v_i)\}} e_{ji}$ (“incoming = outgoing”)
- (2) $\sum_{j \in \{l : v_l \in N_G(v_i)\}} e_{ij} \leq |V(G_i)|$ (at most $|V(G_i)|$)
- (3) $\text{ham}(G_i) \leq \sum_{j \in \{l : v_l \in N_G(v_i)\}} e_{ij}$ (at least $\text{ham}(G_i)$)

For every partition of $V(G)$ into vertex sets A and B :

- (4) $\sum_{1 \leq i < j \leq n : \{v_i, v_j\} \in E(G) \wedge |e \cap A| = 1} e_{ij} + e_{ji} \geq 1$ (“connectivity”)

For every variable e_{ij} :

- (5) $e_{ij} \geq 0$.

Then H has a Hamiltonian cycle if and only if the ILP HAMILTONIAN CYCLE is feasible. Furthermore, the size of the ILP is at most $O^+(m2^n)$ and it has $2m$ variables.

Proof. The size bound on the ILP HAMILTONIAN CYCLE is obvious. Suppose that H has a Hamiltonian cycle C . W.l.o.g. we can assume that C is directed. For every $\{v_i, v_j\} \in E(G)$ we set e_{ij} to be the number of arcs (x, y) in C such that $x \in V(G_i)$ and $y \in V(G_j)$ and similarly we set e_{ji} to be the number of arcs (x, y) in C such that $x \in V(G_j)$ and $y \in V(G_i)$. Then, because C is a Hamiltonian cycle of H , this assignment of e_{ij} and e_{ji} satisfies the constraints (1)–(5), as required.

For the reverse direction, suppose that the ILP HAMILTONIAN CYCLE is feasible and let β be an assignment of the variables e_{ij} and e_{ji} witnessing this. Let G' be the directed multigraph obtained from G by replacing every edge $\{v_i, v_j\}$ with $\beta(e_{ij})$ parallel arcs from v_i to v_j and $\beta(e_{ji})$ parallel arcs from v_j to v_i . Because of the constraints (1), (4), and (5), it follows that G' contains a directed eulerian tour T , i.e., a closed directed walk that visits all the arcs of G' exactly once. Clearly, when fixing any vertex of G' , the tour T defines an ordering of the arcs of G' . Let π be any such ordering of the arcs of G' . For every $1 \leq i \leq n$, let $P_i = (P_1^i, \dots, P_{p_i}^i)$ be a partition of G_i into p_i disjoint paths, where $p_i = \sum_{j \in \{l : v_l \in N_G(v_i)\}} e_{ij}$. Because of the constraints (2) and (3) we know that such a partition exists for every $1 \leq i \leq n$. For every arc $a = (v_i, v_j)$ in T where a is the l -th arc leaving v_i in T and a is the l' -th arc entering v_j in T (according to the ordering π), we denote by $e(a)$ the edge of H from the second endpoint of P_l^i to the first endpoint of $P_{l'}^j$. Then the edges in $\{e(a) : a \in T\}$ together with the edges of all the path $P_1^1, \dots, P_{p_n}^n$ form a Hamiltonian cycle in H , as required. $\square$

Lemma 7. *Given the graph G and the pairs $(\text{ham}(G_1), |V(G_1)|), \dots, (\text{ham}(G_n), |V(G_n)|)$ it can be decided whether the graph H has a Hamiltonian cycle in time $O^+(m2^n + (2m)^{5m+o(m)}n)$.*

Proof. To decide whether the graph H has a Hamiltonian cycle we construct and solve the ILP HAMILTONIAN CYCLE from Lemma 6. The running time of this algorithm is then the time it takes to construct the ILP, i.e., $O^+(m2^n)$, plus the time needed to solve the ILP, i.e., $O^+((2m)^{5m+o(m)} \log(m2^n)) \in O^+((2m)^{5m+o(m)}n)$ using Proposition 4. This concludes the proof of the Lemma. $\square$

We are now ready to show Lemma 5.

Proof of Lemma 5. Clearly, $|V(H)| = \sum_{1 \leq i \leq n} |V(G_i)|$ so it remains to show how to compute $\text{ham}(H)$. Because of Proposition 8 $\text{ham}(H)$ is equal to the minimum positive integer $1 \leq l \leq |V(H)|$ such that the graph $H \oplus l$ has a Hamiltonian cycle. For every $1 \leq l \leq |V(H)|$ the graph $H \oplus l$ is equal to the graph $G'(G_1, \dots, G_n, I_l)$ where G' is the graph obtained from G by adding one vertex v_{n+1} and making it adjacent to all vertices of G , and the graph I_l is the independent set on l vertices. Because $\text{ham}(I_l) = |V(I_l)| = l$ we can use Lemma 7 to decide whether the graph $H \oplus l$ has an Hamiltonian cycle in time $O(m'2^{n'} + (2m')^{5m'+o(m')}n')$ where $n' = n + 1$ and $m' = m + n$. This concludes the proof of the lemma. $\square$

6 Conclusion

We examined some of the algorithmic properties of modular-width, a natural structural parameter. Our results indicate that this is a notion which may be worthy of further study independently of its more famous cousin, clique-width, since its decreased generality does offer some algorithmic pay-off.

As a direction for further research, it would be interesting to see if more problems which are hard for clique-width (or even for treewidth) become tractable for modular-width. Two prime suspects in this category are EDGE DOMINATING SET and PARTITION INTO TRIANGLES.

Beyond that, it would be interesting to investigate if the techniques of this paper can be further generalized, perhaps eventually leading to meta-theorem-like results. In particular, our ILP-based solution for HAMILTONICITY may be applicable (with some modifications) to other problems. One may ask: what properties must a problem possess for us to be able to give a straightforward DP algorithm that uses ILPs to combine the tables?

The main property that a problem should satisfy for these ideas to apply is that the sets of partial solutions arising in the dynamic programming formulation should be *convex*. Convexity is important here, since we would like to be able to express the information contained in the DP tables using linear constraints, in order to use an ILP. Convexity was easy to establish in the case of PARTITION INTO PATHS and similar problems, since the set of feasible partial solutions is the set of integers k such that there exists a partition of a subgraph into k paths. If one knows the minimum feasible k , all larger integers are also feasible and this is trivially a convex set. The question then becomes, are there any other natural problems where convexity can be established (perhaps non-trivially) and used in this way?

References

- [1] Stéphane Bessy, Christophe Paul, and Anthony Perez. Polynomial kernels for 3-leaf power graph modification problems. *Discrete Applied Mathematics*, 158(16):1732–1744, 2010.
- [2] Andreas Björklund, Thore Husfeldt, and Mikko Koivisto. Set partitioning via inclusion-exclusion. *SIAM J. Comput.*, 39(2):546–563, 2009.
- [3] Hans L. Bodlaender. A tourist guide through treewidth. *Acta Cybernetica*, 11(1–2):1–22, 1993.
- [4] Hans L. Bodlaender. The algorithmic theory of treewidth. *Electronic Notes in Discrete Mathematics*, 5:27–30, 2000.
- [5] Hans L. Bodlaender. Treewidth: Characterizations, applications, and computations. In Fedor V. Fomin, editor, *Graph-Theoretic Concepts in Computer Science, 32nd International Workshop, WG 2006, Bergen, Norway, June 22–24, 2006, Revised Papers*, volume 4271 of *Lecture Notes in Computer Science*, pages 1–14. Springer, 2006.
- [6] Hans L. Bodlaender and Arie M. C. A. Koster. Combinatorial optimization on graphs of bounded treewidth. *Comput. J.*, 51(3):255–269, 2008.
- [7] Binh-Minh Bui-Xuan, Jan Arne Telle, and Martin Vatshelle. Boolean-width of graphs. *Theor. Comput. Sci.*, 412(39):5187–5204, 2011.
- [8] Bruno Courcelle, Johann A. Makowsky, and Udi Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory of Computer Systems*, 33(2):125–150, 2000.
- [9] Reinhard Diestel. *Graph Theory*, volume 173 of *Graduate Texts in Mathematics*. Springer Verlag, New York, 2nd edition, 2000.
- [10] Michael Dom, Jiong Guo, Falk Hüffner, Rolf Niedermeier, and Anke Truß. Fixed-parameter tractability results for feedback set problems in tournaments. In Tiziana Calamoneri, Irene Finocchi, and Giuseppe F. Italiano, editors, *Algorithms and Complexity, 6th Italian Conference, CIAC 2006, Rome, Italy, May 29–31, 2006, Proceedings*, volume 3998 of *Lecture Notes in Computer Science*, pages 320–331. Springer, 2006.
- [11] Michael Dom, Daniel Lokshtanov, Saket Saurabh, and Yngve Villanger. Capacitated domination and covering: A parameterized perspective. In Martin Grohe and Rolf Niedermeier, editors, *Parameterized and Exact Computation, Third International Workshop, IWPEC 2008, Victoria, Canada, May 14–16, 2008. Proceedings*, volume 5018 of *Lecture Notes in Computer Science*, pages 78–90. Springer, 2008.
- [12] R. G. Downey and M. R. Fellows. *Parameterized Complexity*. Monographs in Computer Science. Springer Verlag, New York, 1999.
- [13] Michael R. Fellows, Fedor V. Fomin, Daniel Lokshtanov, Frances A. Rosamond, Saket Saurabh, Stefan Szeider, and Carsten Thomassen. On the complexity of some colorful problems parameterized by treewidth. *Inf. Comput.*, 209(2):143–153, 2011.

- [14] Michael R. Fellows, Daniel Lokshtanov, Neeldhara Misra, Frances A. Rosamond, and Saket Saurabh. Graph layout problems parameterized by vertex cover. In Seok-Hee Hong, Hiroshi Nagamochi, and Takuro Fukunaga, editors, *Algorithms and Computation, 19th International Symposium, ISAAC 2008, Gold Coast, Australia, December 15-17, 2008. Proceedings*, volume 5369 of *Lecture Notes in Computer Science*, pages 294–305. Springer, 2008.
- [15] Jörg Flum and Martin Grohe. *Parameterized Complexity Theory*, volume XIV of *Texts in Theoretical Computer Science. An EATCS Series*. Springer Verlag, Berlin, 2006.
- [16] Fedor V. Fomin, Petr A. Golovach, Daniel Lokshtanov, and Saket Saurabh. Clique-width: on the price of generality. In Claire Mathieu, editor, *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2009, New York, NY, USA, January 4-6, 2009*, pages 825–834. SIAM, 2009.
- [17] Fedor V. Fomin, Petr A. Golovach, Daniel Lokshtanov, and Saket Saurabh. Algorithmic lower bounds for problems parameterized with clique-width. In Moses Charikar, editor, *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*, pages 493–502. SIAM, 2010.
- [18] Fedor V. Fomin, Petr A. Golovach, Daniel Lokshtanov, and Saket Saurabh. Intractability of clique-width parameterizations. *SIAM J. Comput.*, 39(5):1941–1956, 2010.
- [19] Robert Ganian. Twin-cover: Beyond vertex cover in parameterized algorithmics. In Dániel Marx and Peter Rossmanith, editors, *Parameterized and Exact Computation - 6th International Symposium, IPEC 2011, Saarbrücken, Germany, September 6-8, 2011. Revised Selected Papers*, volume 7112 of *Lecture Notes in Computer Science*, pages 259–271. Springer, 2011.
- [20] Robert Ganian, Petr Hlinený, Jaroslav Nesetril, Jan Obdržálek, Patrice Ossona de Mendez, and Reshma Ramadurai. When trees grow low: Shrubs and fast mso1. In Branislav Rován, Vladimiro Sassone, and Peter Widmayer, editors, *Mathematical Foundations of Computer Science 2012 - 37th International Symposium, MFCS 2012, Bratislava, Slovakia, August 27-31, 2012. Proceedings*, volume 7464 of *Lecture Notes in Computer Science*, pages 419–430. Springer, 2012.
- [21] Michel Habib and Christophe Paul. A survey of the algorithmic aspects of modular decomposition. *Computer Science Review*, 4(1):41–59, 2010.
- [22] Sang il Oum. Rank-width and vertex-minors. *J. Comb. Theory, Ser. B*, 95(1):79–100, 2005.
- [23] Michael Lampis. Algorithmic meta-theorems for restrictions of treewidth. *Algorithmica*, 64(1):19–37, 2012.
- [24] Ross M. McConnell and Jeremy Spinrad. Modular decomposition and transitive orientation. *Discrete Mathematics*, 201(1–3):189–241, 1999.
- [25] Rolf Niedermeier. *Invitation to Fixed-Parameter Algorithms*. Oxford Lecture Series in Mathematics and its Applications. Oxford University Press, Oxford, 2006.
- [26] Fábio Protti, Maise Dantas da Silva, and Jayme Luiz Szwarcfiter. Applying modular decomposition to parameterized cluster editing problems. *Theory of Computing Systems*, 44(1):91–104, 2009.
- [27] Marc Tedder, Derek G. Corneil, Michel Habib, and Christophe Paul. Simpler linear-time modular decomposition via recursive factorizing permutations. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfssdóttir, and Igor Walukiewicz, editors, *Automata, Languages and Programming, 35th International Colloquium, ICALP 2008, Reykjavik, Iceland, July 7-11, 2008, Proceedings, Part I: Track A: Algorithms, Automata, Complexity, and Games*, volume 5125 of *Lecture Notes in Computer Science*, pages 634–645. Springer, 2008.