

# Fast Parallel Fixed-Parameter Algorithms via Color Coding

Max Bannach<sup>1,2</sup>, Christoph Stockhusen<sup>1</sup>, and Till Tantau<sup>1</sup>

- 1 Institute for Theoretical Computer Science  
Universität zu Lübeck  
Lübeck, Germany  
{bannach,stockhus,tantau}@tcs.uni-luebeck.de
- 2 Graduate School for Computing in Medicine and Life Sciences  
Universität zu Lübeck

---

## Abstract

Fixed-parameter algorithms have been successfully applied to solve numerous difficult problems within acceptable time bounds on large inputs. However, most fixed-parameter algorithms are inherently *sequential* and, thus, make no use of the parallel hardware present in modern computers. We show that parallel fixed-parameter algorithms do not only exist for numerous parameterized problems from the literature – including vertex cover, packing problems, cluster editing, cutting vertices, finding embeddings, or finding matchings – but that there are parallel algorithms working in *constant* time or at least in time *depending only on the parameter* (and not on the size of the input) for these problems. Phrased in terms of complexity classes, we place numerous natural parameterized problems in parameterized versions of  $AC^0$ . On a more technical level, we show how the *color coding* method can be implemented in constant time and apply it to embedding problems for graphs of bounded tree-width or tree-depth and to model checking first-order formulas in graphs of bounded degree.

**1998 ACM Subject Classification** F.1.3 Complexity Measures and Classes

**Keywords and phrases** color coding, parallel computation, fixed-parameter tractability, graph packing, cutting  $\ell$  vertices, cluster editing, tree-width, tree-depth, model checking

## 1 Introduction

The classical objective of parameterized complexity theory is to determine for a parameterized problem whether it can be solved by an algorithm running in time  $f(k) \cdot n^c$ , where  $f$  is some function,  $k$  is a parameter,  $n$  is the input length, and  $c$  is some constant. Such algorithms are nowadays routinely used to solve large instances for NP- or even PSPACE-hard problems within acceptable amounts of time. Nevertheless, “acceptable” is not the same as “small” and one would like to further reduce the runtime by using multiple cores to speed up the computation. For this, one needs *parallel* fixed-parameter algorithms, but most fixed-parameter algorithms have been devised with a sequential computation model in mind. Indeed, the most important tool of parameterized complexity theory, namely kernelization, is inherently sequential: It asks us to repeatedly apply rules to an input, each time modifying the input slightly and making it a little smaller, until the input’s size only depends on the parameter. There is no straightforward way of parallelizing such algorithms since later modifications strongly depend on what happened earlier, forcing us to apply the typically very large number of kernelization steps in a sequential manner.



© M. Bannach, C. Stockhusen, and T. Tantau;  
licensed under Creative Commons License CC-BY

Conference title on which this volume is based on.

Editors: Billy Editor and Bill Editors; pp. 1–19



Leibniz International Proceedings in Informatics

LIPIC Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

**Our Contributions.** The purpose of the present paper is to show that not only do parallel fixed-parameter algorithms exist for many natural, well-studied problems from the literature; for certain problems there are even parallel algorithms that require only *constant* time in a concurrent-read, concurrent-write PRAM model (so the runtime is totally independent of the input) or at least time *depending only on the parameter* (so the length of the input is irrelevant). In all cases, the *work* done by the algorithms is still  $f(k) \cdot n^c$ , that is, the same as the time bound for sequential fixed-parameter algorithms.<sup>1</sup> Phrased more formally, our objective is to identify parameterized problems that lie in the complexity classes  $\text{para-AC}_{O(1)}$  and  $\text{para-AC}_{f(k)}$  (formal definitions will be given later).

In order to tackle the parallel parameterized complexity of natural problems like the vertex cover problem, we introduce three technical tools. The first and foremost is *color coding*: all of our proofs employ this technique at least indirectly and we show that the *universal coloring families* that lie at the heart of the technique can be computed in constant time. Second, numerous natural “packing problems” are special cases of the following embedding problem: Given graphs  $H$  and  $G$ , find a (not necessarily induced) subgraph of  $G$  that is isomorphic to  $H$ . We give new bounds on the complexity of this problem when  $H$  has bounded tree-width or bounded tree-depth; and these bounds later translate directly to bounds on different packing problems. Third, we translate an algorithmic meta-theorem of Flum and Grohe [16] to the parallel world: We show that model checking first-order properties of graphs can be done in parallel in time depending only on the parameters (actually, only on the locality rank of the formula), where the parameters are the to-be-checked formula and the degree of the graph.

We then apply the tools to a wide variety of natural graph problems, namely *packing problems*, *covering problems*, *clustering problems*, and *separation problems*. For *packing problems* the objective is to determine whether a given graph  $G$  contains  $k$  vertex-disjoint copies of some fixed graph  $H$  like, say, a triangle. Even for triangles, this problem is already NP-complete, but when  $k$  is considered as a parameter, the triangle packing problem lies in FPT [15]. We show that there is a constant-time, FPT-work algorithm for triangle packing – and indeed for packing any graph of fixed size. The *covering problems* we study include the vertex cover problem and its partial version. We present a constant parallel time algorithm for the first problem and an algorithm for the second needing time depending only on the parameter. These results nicely reflect on a theoretical basis the “empirical” observation that  $p$ -VERTEX-COVER is one of the “easiest” parameterized problems and that the partial version is a bit harder to solve. For *clustering problems*, also known as *cluster editing problems*, the objective is to transform a graph by adding or deleting few edges into a collection of “clusters” – which are just cliques in the simplest case. We present a constant time, FPT-work algorithm for cluster editing. For *graph separation problems* the objective is to “cut away” a special part of a graph using few vertices. We show that certain versions of these problems can be solved by a parallel fixed-parameter algorithm in time depending only on the parameter and FPT work (while other versions are known to be W[1]-hard).

---

<sup>1</sup> The *work* done by a parallel algorithm is the total number of computational steps made by all computational units during a computation. Since “all work needs to be done,” in practice the runtime of a parallel algorithm is its work divided by the number of available cores. In particular, the *work* done by a *parallel algorithm* should not exceed the *runtime* of a *sequential algorithm* for the same problem. In our case, this means that in order to compete with sequential algorithms running in “FPT time,” our parallel algorithm must not only be fast, but may only do “FPT work.”

**Related Work.** There is a growing body of literature reporting on the practicalities of implementing fixed-parameter algorithms in parallel [1]. In contrast, there are only few results addressing parallel fixed-parameter tractability on a theoretical level (as we do in the present paper), see for instance Cesati and Di Ianni [9]. Since it is well-known from classical complexity theory that problems solvable in logarithmic space can be parallelized well, previous research on *parameterized logarithmic space* contributes to our understanding of which parameterized problems can be parallelized in principle. This research was started by Cai, Chen, Downey, and Fellows [8]. First (quite technical) complete problems for parameterized logarithmic space were later introduced by Chen, Flum, and Grohe [11], and by Flum and Grohe [16]. A more structural study of parameterized space and circuit classes (which addresses parallelization more directly) was later made by Elberfeld and the last two authors [14]. Parameterized Circuit Complexity was also studied by Downey et al. with respect to the Weft Hierarchy [13]. Recently, Chen and Müller [10] connected color coding and parameterized space in an algorithm for finding embedding of bounded tree-depth graphs in parameterized logarithmic space (a result which we strengthen considerably in Corollary 3.7).

The first use of the color coding technique can be traced back to Alon, Yuster, and Zwick [2]. They used the technique to provide an FPT-algorithm that decides whether there is an embedding of a graph  $H$  of bounded tree-width into another graph  $G$ , where  $H$  is the parameter.

**Organization of This Paper.** In Section 2 we give formal definitions of the classes of problems solvable by parallel fixed-parameter algorithms. While most of our definitions and classes are standard, the class of problems solvable in “time depending on the parameter and FPT work” seems to be new. In Section 3 we introduce our three technical tools – color coding, embeddings, and model checking – and prove the results mentioned earlier. In Section 4 we study the complexity of the natural parameterized graph problems and establish new upper bounds on their complexities. Due to lack of space, proofs have been moved to the appendix; we give proof sketches for some of them in the main text.

## 2 Classes of Fixed-Parameter Parallelism

For our definition of parallel fixed-parameter tractability, we mostly use the standard terminology of parameterized complexity theory, see for instance [17]: A *parameterized problem* is a tuple  $(Q, \kappa)$  of a language  $Q \subseteq \Sigma^*$  over an alphabet  $\Sigma$  and a parameterization  $\kappa: \Sigma^* \rightarrow \mathbb{N}$  that maps instances to parameter values. In the classical definition, Downey and Fellows [12] require the parameterization to be computable, while Flum and Grohe [17] require it to be computable in polynomial time. Elberfeld and the last two authors require it to be computable in logarithmic space [14] and mention that it would be better if the parameterization is first-order computable (FO-computable) or, equivalently, to be computable by logarithmic-time-uniform constant depth circuits [24]. Since we will only deal with parameterized circuit classes that lie within parameterized logarithmic space, we will require all parameterizations to be FO-computable. We denote parameterized problems with a leading “ $p$ –” as in  $p$ -VERTEX-COVER and, when the parameter may be unclear, add it as an index as in  $p_{|H|}$ -EMB.

A parameterized problem is *fixed-parameter tractable* if it can be decided in time  $f(\kappa(x)) \cdot |x|^c$  for any input  $x$ , where  $f$  is some computable function and  $c$  a constant. An equivalent definition is that there exists a set  $R \in P$ , where  $P$  denotes the class of languages decidable

in polynomial time, such that  $x \in Q$  iff  $(x, 1^{f(\kappa(x))}) \in R$ . The first definition of fixed-parameter tractability gave rise to the class name FPT in the literature, while the second definition gives rise to the name para-P for the same class. The advantage of the second definition is that we can replace the class P in the definition by arbitrary complexity classes and arrive at classes like *parameterized logarithmic space*, para-L, or *parameterized constant depth circuits*, para-AC<sup>0</sup>. These parameterized classes inherit their inclusion structure from the classical classes, so we have

$$\text{para-AC}^0 \subsetneq \text{para-TC}^0 \subseteq \text{para-NC}^1 \subseteq \text{para-L} \subseteq \text{para-NL} \subseteq \text{para-AC}^1 \subseteq \text{para-P}.$$

It is not quite obvious, but the class para-AC<sup>0</sup> already captures one of the types of algorithms mentioned in the introduction, namely “constant time, FPT-work,” while none of the above classes seems to capture “parameter time, FPT-work.” For this reason and in order to explicitly spell out what para-AC<sup>0</sup> contains, we provide a new definition:<sup>2</sup>

► **Definition 2.1** (Classes of Parallel Fixed-Parameter Tractability). Let  $d: \mathbb{N}^2 \rightarrow \mathbb{N}$  be a *depth bounding function* and  $w: \mathbb{N}^2 \rightarrow \mathbb{N}$  be a *width bounding function* which both map each pair of an input length and a parameter to a number. We define para-AC $[d, w]$  as the class of parameterized problems  $(Q, \kappa)$  for which there exists a DLOGTIME-uniform<sup>3</sup> family  $(C_{n,k})_{n,k \in \mathbb{N}}$  of AC-circuits (only NOT-, AND-, and OR-gates are allowed, AND- and OR-gates may have unbounded fan-in) such that:

1. For all  $x \in \Sigma^*$ , the circuit  $C_{|x|, \kappa(x)}$  evaluates to 1 on input  $x$  if, and only if,  $x \in Q$ .
2. The depth of each  $C_{n,k}$  is at most  $d(n, k)$ .
3. The size of each  $C_{n,k}$  is at most  $w(n, k)$ .

In the present paper we exclusively study parallel algorithms with “FPT-work” and are therefore only interested in the case where  $w$  is member of the family  $W$  of functions of the form  $f(k) \cdot n^c$  for a computable function  $f$  and a constant  $c$ . We introduce for arbitrary families  $D$  of functions  $d: \mathbb{N}^2 \rightarrow \mathbb{N}$  the abbreviation para-AC $_D$  for  $\bigcup_{d \in D, w \in W} \text{para-AC}[d, w]$ . For constant depth bounding functions the resulting class para-AC $_{O(1)}$  is the same as the class para-AC<sup>0</sup>.<sup>4</sup> For arbitrary  $i > 0$  we obtain  $\text{para-AC}_{\{f(k)+c \cdot \log^i n \mid f: \mathbb{N} \rightarrow \mathbb{N} \wedge c \in \mathbb{N}\}} = \text{para-AC}^i$  (in slight abuse of notation we will write such classes simply as para-AC $_{f(k)+O(\log^i n)}$ ).

When the depth bounding function just depends on the parameter, so  $d(n, k) = f(k)$ , we get a new class para-AC $_{f(k)}$  that we abbreviate with para-AC<sup>0 $\uparrow$</sup> . This class does not seem to arise from substituting some classical class for P in the definition of para-P. In particular, this class seems to be incomparable with all classes between para-TC<sup>0</sup> and para-NL. It is, however, clearly contained in para-AC<sup>1</sup>, and is strictly more powerful than para-AC<sup>0</sup> as we will see later. This class captures the problems solvable in “parameter time, FPT-work” and we have

$$\text{para-AC}^0 \subsetneq \text{para-AC}^{0\uparrow} \subseteq \text{para-AC}^1.$$

Let us define for arbitrary  $i \geq 0$  the class para-AC <sup>$i\uparrow$</sup>  as para-AC $_{f(k) \cdot O(\log^i n)}$ . Notice that we have by definition the inclusion structure  $\text{para-AC}^i \subseteq \text{para-AC}^{i\uparrow} \subseteq \text{para-AC}^{i+\epsilon}$ .

<sup>2</sup> The definition can trivially be adjusted to use TC-circuits or NC-circuits, but we will not need them.

<sup>3</sup> We use DLOGTIME-uniform families since they are equivalent to first-order definable families and constitute one of the strongest forms of uniformity [4].

<sup>4</sup> Since the designation para-AC<sup>0</sup> has been used in previous publications and is a bit shorter, we will use it in the following.

### 3 Technical Tools

#### 3.1 Color Coding in Constant Parallel Time

The idea of color coding is best understood by a concrete application, for instance to the well-known matching problem: Given an undirected graph  $G$  and a number  $k$ , does  $G$  contain  $k$  edges such that no two of them share any endpoints? Directly solving this problem is not easy since the known polynomial-time algorithms for it are rather involved. Consider, however, what happens when we randomly color the graph with  $k$  colors and then check whether *the vertices of each color class contain at least one edge*. Clearly, if this is the case, there is a matching of size  $k$  – and if there is no such matching, then no coloring will pass the test.

We now formalize the idea behind color coding and then show how the colorings can be computed in constant time. It turns out that one can derandomize the computation of a coloring: instead of random colorings we use sets of colorings such that for every set of  $k$  vertices and “desired” colors for them, at least one coloring colors the vertices as desired:

► **Definition 3.1** (Universal Coloring Families). For natural numbers  $n$ ,  $k$ , and  $c$ , an  $(n, k, c)$ -universal coloring family is a set  $\Lambda$  of functions  $\lambda: \{1, \dots, n\} \rightarrow \{1, \dots, c\}$  such that for every subset  $S \subseteq \{1, \dots, n\}$  of size  $|S| = k$  and for every mapping  $\mu: S \rightarrow \{1, \dots, c\}$  there is at least one function  $\lambda \in \Lambda$  with  $\forall s \in S: \mu(s) = \lambda(s)$ .

The matching problem can be solved easily when we have access to a  $(n, 2k, k)$ -universal coloring family: If there is a matching of size  $k$ , the family will contain some coloring that colors the two endpoints of the first edge with color 1, the endpoints of the second edge with color 2, and so on. Thus there is, indeed, a matching of size  $k$  in the graph if for at least one coloring every color class contains an edge. Since we can easily check in parallel for all colorings whether this is the case for one of them, the complexity of  $p_k$ -MATCHING hinges critically on the complexity of computing the universal coloring family and the size of this family. The next theorem shows that  $(n, k, c)$ -universal coloring families of reasonable size can be computed “in constant time and work  $f(k, c) \cdot n^{O(1)}$ ,” which implies that  $p_k$ -MATCHING  $\in$  para-AC<sup>0</sup> holds:

► **Theorem 3.2.** *There is a DLOGTIME-uniform family  $(C_{n,k,c})_{n,k,c \in \mathbb{N}}$  of AC-circuits without inputs such that each  $C_{n,k,c}$*

1. *outputs an  $(n, k, c)$ -universal coloring family (coded as a sequence of function tables),*
2. *has constant depth (independent of  $n$ ,  $k$ , or  $c$ ), and*
3. *has size at most  $O(n \log c \cdot c^{k^2} \cdot k^4 \log^2 n)$ .*

**Sketch of Proof.** The family of universal coloring functions we construct is based on the concept of  $k$ -perfect hash functions [17], that, after slight modifications, provide us with the desired coloring properties. The crucial part is to implement them using circuits that are DLOGTIME-uniform. However, we can achieve this, since the numbers  $n$ ,  $k$ , and  $c$  are encoded in unary and the operations required to compute the functions are only additions, multiplications, and modulo operations. ◀

Investigating a parameterized version of matching may seem a bit strange at first sight – matching is even known to be solvable in randomized polylogarithmic parallel time. However, the exact parallel time complexity is still open in the classical setting while from a parameterized perspective, we just saw that the matching can be solved *very* quickly in parallel. Another problem that one would maybe not expect to be studied in the parameterized setting, but which will be useful in a number of situations, is  $p$ -THRESHOLD. The

inputs are a bitstring  $b \in \{0, 1\}^n$  and a parameter  $t$ . The question is whether there are at least  $t$  many 1's in  $b$ . Clearly, the unparameterized version is complete for  $\text{TC}^0$ , and using the fact that the problem lies in  $\text{AC}^0$  for polylogarithmic thresholds [23] yields the fact that its parameterized version lies in  $\text{para-AC}^0$ . However, this result requires profound result of circuit complexity and is rather involved, but using color coding we can give a very simple proof of this fact:

► **Lemma 3.3.**  $p\text{-THRESHOLD} \in \text{para-AC}^0$ .

### 3.2 Finding Embeddings of Graphs of Bounded Tree-Width and Depth

A different way of looking at the matching problem is to see it as an embedding problem: Instead of trying to find  $k$  edges in a graph  $G$  that have no endpoints in common, we can try to “embed” the graph  $H = kK_2$ , consisting of  $k$  isolated edges, into  $G$ . The advantage of this different point of view is, of course, that it generalizes nicely:

► **Problem 3.4** ( $p\text{-EMB}(\mathcal{H})$  for some class  $\mathcal{H}$  of undirected graphs).

*Instance:* Two undirected graphs  $H = (V_H, E_H) \in \mathcal{H}$  and  $G = (V_G, E_G)$ .

*Parameter:*  $H$

*Question:* Is there a injective homomorphism  $\phi: V_H \rightarrow V_G$ , that is, is  $H$  isomorphic to a (not necessarily induced) subgraph of  $G$ ?

For arbitrary  $\mathcal{H}$ , the problem is easily be seen to be  $\text{W}[1]$ -hard by a reduction from  $p\text{-CLIQUE}$ . However, for restricted  $\mathcal{H}$ , the problem becomes fixed-parameter tractable. The best results so far are by Chen and Müller [10] who show that when  $\mathcal{H}$  has bounded tree-depth,  $p\text{-EMB}(\mathcal{H}) \in \text{para-L}$ ; when  $\mathcal{H}$  has bounded path-width,  $p\text{-EMB}(\mathcal{H})$  is the para-L-reduction closure of the distance problem in graphs, parameterized by the distance; and when  $\mathcal{H}$  has bounded tree-width,  $p\text{-EMB}(\mathcal{H})$  is the para-L-reduction closure of the embedding problem for trees, parameterized by the tree-size. In contrast to these results, Amano showed for the unparameterized setting, in which we consider the size of  $H$  to be a constant, that the problem can be solved in  $\text{AC}^0$  with similar techniques [3]. We improve considerably on the first result of Chen and Müller by proving that embeddings of graphs of bounded tree-depth can actually be computed in  $\text{para-AC}^0$ . We complement their other results, without improving them, by showing that for graphs of bounded tree-width (and, thereby, also for bounded path-width) the embedding problem lies in  $\text{para-AC}^{0\uparrow}$ .

In order to formulate our results, we first need to review the definition of a tree-decomposition, see [17] for a more detailed introduction. A *tree-decomposition* of a graph  $H = (V, E)$  is a tree  $T$  together with a mapping  $\iota$  from the nodes of  $T$  to subsets (called *bags*) of  $V$  such that (1) for every edge  $\{u, v\} \in E$  there is some bag containing  $u$  and  $v$ , that is, there is some  $x \in V$  with  $\{u, v\} \subseteq \iota(x)$  and (2) for every vertex  $x \in V$  the set of nodes of  $T$  whose bags contain  $x$  forms a connected subset of  $T$ . The *width* of tree-decomposition is the size of its largest bag minus 1, its depth is the maximum of the width and the depth of  $T$ . Define  $\text{tw}(H)$  as the minimum width any tree-decomposition of  $H$  must have; define  $\text{td}(H)$  similarly for the tree-depth.

► **Theorem 3.5.** Given two graphs  $H = (V_H, E_H)$  and  $G = (V_G, E_G)$  together with a tree-decomposition  $(T, \iota)$  of  $H$ . An embedding of  $H$  into  $G$  can be computed by an  $\text{AC}$ -circuit of depth  $O(\text{depth}(T))$  and size  $f(|V_H|) \cdot O(|V_G|^{\text{width}(T)})$ , if such an embedding exists.

**Sketch of Proof.** Color the vertices of  $H$  uniquely and compute a  $(|V_G|, |V_H|, |V_H|)$ -universal coloring family. Starting from the leaves of the tree-decomposition, merge compatible partial



homomorphisms for the vertices of the bags until we reach the root of the decomposition, and, thus, obtain a homomorphism for  $H$ . The number of iterative steps required for this equals the depth of the tree-decomposition. ◀

If  $H$  is a parameter, we can compute a width- or depth-bounded tree-decomposition  $(T, \iota)$  of  $H$  in a preprocessing step. This implies the following corollaries:

► **Corollary 3.6.** *Let  $\mathcal{H}$  be the class of all graphs of tree-width at most  $d$  for some constant  $d$ . Then  $p\text{-EMB}(\mathcal{H}) \in \text{para-AC}_{f(|H|)} \subseteq \text{para-AC}^{0\uparrow}$ .*

► **Corollary 3.7.** *Let  $\mathcal{H}$  be the class of all graphs of tree-depth at most  $d$  for some constant  $d$ . Then  $p\text{-EMB}(\mathcal{H}) \in \text{para-AC}^0$ .*

We make two remarks at this point: First, one cannot generalize Theorem 3.5 to clique-width since the embedding problem for cliques, which have clique-width 1, is already hard for  $W[1]$ . Second, the theorem and the corollary also hold for relational structures  $\mathcal{H}$  and  $\mathcal{G}$  and if we bound the tree-width of  $\mathcal{H}$ 's Gaifman graph. Since paths have tree-width 1, the complexity of one of the canonical problems for color-coding – the  $p_k\text{-PATH}$  problem – can be determined:  $p_k\text{-PATH} \in \text{para-AC}^{0\uparrow}$ . This allows us to give a short proof of the following lemma on the complexity of the distance problem for directed graphs where the distance is the parameter (one can also prove this lemma directly quite easily):

► **Lemma 3.8.**  $p_d\text{-DISTANCE} \in \text{para-AC}_{f(d)} \subseteq \text{para-AC}^{0\uparrow}$ .

A known fact from circuit complexity states that a polynomial-sized AC-circuit that decides whether a given graph  $G$  contains a path of length at most  $d$  between two vertices  $s$  and  $t$  requires depth  $\Omega(\log \log d)$  [5]. This implies  $p_d\text{-DISTANCE} \notin \text{para-AC}^0$ .

► **Corollary 3.9.**  $\text{para-AC}^0 \subsetneq \text{para-AC}^{0\uparrow}$ .

### 3.3 First-Order Model Checking

Our last result in this section on tools is an algorithmic meta-theorem: We show that the model checking problem for first-order logic on graphs of bounded degree lies in  $\text{para-AC}^{0\uparrow}$ . We build strongly on a previous result by Flum and Grohe [16], who showed that this model checking problem lies in  $\text{para-L}$ , but differ in three regards: First, we use color coding in our proof, which simplifies the argument somewhat, second, we identify the parameterized distance problem on bounded degree graphs as the only part of the computation that is presumably not in  $\text{para-AC}^0$ , and, third, we observe that the degree of the graphs can be made a parameter and need not be considered constant.

► **Problem 3.10** ( $p_{\phi, \delta}\text{-MC}(\text{FO})$ ).

*Instance:* A logical structure  $\mathcal{A}$  and a first-order formula  $\phi$ .

*Parameter:* The (size of) the formula  $\phi$  and the maximum degree  $\delta$  of  $\mathcal{A}$ 's Gaifman graph.

*Question:*  $\mathcal{A} \models \phi$ ?

► **Theorem 3.11.**  $p_{\phi, \delta}\text{-MC}(\text{FO}) \in \text{para-AC}_{f(\phi + \delta)} \subseteq \text{para-AC}^{0\uparrow}$ .

**Sketch of Proof.** By Gaifman's Theorem [20] we can rewrite the given formula as a formula  $\phi'$  in Gaifman normal form. Thus, what essentially remains is to check whether the structure (which we can interpret as a graph) contains  $k$  disjoint “balls” of size bounded in the parameter (due to the maximum degree of the underlying Gaifman graph) that satisfy the subformulas in  $\phi'$ . To find these substructures, we make use of color coding and apply Lemma 3.8 to compute the corresponding connecting components. Finally, we only have to model check the resulting parameter-sized substructures. ◀

We conclude with the remark that the depth of the circuits constructed in the above theorem just depends on the degree  $\delta$  of the graph and on the radius  $r$  of the balls, which measure how “local” the formula  $\phi$  is. The smallest  $r$  for which  $\phi$  can be rewritten as in the proof is known as the *locality rank*  $\text{lr}(\phi)$  and the proof actually shows that  $p_{\phi,\delta}\text{-MC(FO)} \in \text{para-AC}_{O(\delta \text{lr}(\phi))}$ .

## 4 Fast Parallel Fixed-Parameter Algorithms for Natural Problems

The tools we have developed are now applied to a number of natural parameterized problems found in the literature.

**Packing Problems.** We have already pointed out that the parameterized matching problem can be seen as an embedding problem, where the objective is to embed the graph  $H = kK_2$ , consisting of  $k$  disjoint copies of a single edge, into a graph  $G$ . Embedding multiple disjoint copies of the same graph into another graph is also known as “packing”. Clearly, instead of edges we can also pack other things as long as taking any number of copies of these “other things” still has bounded tree-depth. For instance, we can try to “pack”  $k$  different triangles into  $G$ , that is, we can check whether there are  $k$  vertex-disjoint triangles in  $G$ . Unlike the matching problem, triangle packing is known to be NP-complete.

► **Theorem 4.1.**  $p\text{-TRIANGLE-PACKING} \in \text{para-AC}^0$ .

**Proof.** Just observe that a graph  $H$  consisting of any number of disjoint copies of a triangle has tree-depth 3. The claim follows from Corollary 3.7. ◀

Indeed, for any fixed graph  $H_0$  the packing problem  $p\text{-}H_0\text{-PACKING}$  lies in  $\text{para-AC}^0$ , where the question is whether we can find  $k$  disjoint copies of  $H_0$  in  $G$  and  $k$  is the parameter:

► **Theorem 4.2.**  $p\text{-}H_0\text{-PACKING} \in \text{para-AC}^0$  for every fixed graph  $H_0$ .

Further variants arise when, instead of a single graph  $H_0$ , we are given a whole multiset of graphs as inputs and we must find disjoint copies of all of them in  $G$ . Again, as long as there is a fixed bound on the size of the graphs, the tree-depth of their disjoint union is bounded and, hence, the packing problem lies in  $\text{para-AC}^0$ .

The complexity of packing problem changes when the to-be-packed graphs no longer have constant size as in the following problem:

► **Problem 4.3** ( $p_{k,l}\text{-CYCLE-PACKING}$ ).

*Instance:* An undirected graph  $G$  and two numbers  $k$  and  $l$ .

*Parameter:*  $k$  and  $l$

*Question:* Are there  $k$  vertex-disjoint cycles in  $G$ , each having length  $l$ ?

The graph  $H = kC_l$  consisting of  $k$  copies of a cycle of length  $l$  no longer has bounded tree-depth; it does have tree-width 2, however. Thus, by Theorem 3.5 we get:

► **Theorem 4.4.**  $p_{k,l}\text{-CYCLE-PACKING} \in \text{para-AC}_{f(k+l)} \subseteq \text{para-AC}^{0\uparrow}$ .

The same result obviously also holds for  $p_{k,l}\text{-PATH-PACKING}$  and it also holds for  $p\text{-FOREST-PACKING}$ , where we are given a forest as input and the parameter is the total numbers of vertices in it. We conclude with the remark that these ideas cannot be extended to packing graphs whose tree-width is not bounded: Already embedding cliques, let alone packing them, is  $W[1]$ -hard.



**Covering Problems.** In *covering problems* we must choose vertices in a graph (or sometimes hypergraph) such that all ( $p$ -VERTEX-COVER) or some ( $p$ -PARTIAL-VERTEX-COVER) of the edges are “covered,” that is, they intersect with the set of chosen vertices. The best-known covering problem is undoubtedly  $p$ -VERTEX-COVER, whose complexity has been scrutinized extensively in parameterized complexity theory. We now prove  $p$ -VERTEX-COVER  $\in$  para-AC<sup>0</sup>; a fact that nicely reflects on a theoretical basis the “empirical” observation that  $p$ -VERTEX-COVER is one of the “easiest” parameterized problems. The problem was one of the first shown to lie in para-P, was then shown to lie in para-L by Cai et al. [8], then in para-TC<sup>0</sup> by Elberfeld and the last two authors [14].

► **Theorem 4.5.**  $p$ -VERTEX-COVER  $\in$  para-AC<sup>0</sup>.

*Partial* covering problems ask us not to cover all edges, but only  $t$  of them:

► **Problem 4.6** ( $p_{k,t}$ -PARTIAL-VERTEX-COVER).

*Instance:* An undirected graph  $G = (V, E)$  and two numbers  $k$  and  $t$ .

*Parameter:*  $k, t$

*Question:* Is there a set  $S \subseteq V$  of cardinality  $|S|$  at most  $k$  such that the cardinality of  $\{\{u, v\} \in E \mid u \in S \vee v \in S\}$  is at least  $t$ ?

Another version is  $p_t$ -EXACT-PARTIAL-VERTEX-COVER, where the size of  $S$  is no longer restricted, but the cardinality of  $\{\{u, v\} \in E \mid u \in S \vee v \in S\}$  must be *exactly*  $t$ .

These problems, which are generally considered to be harder than the plain vertex cover problem, lie in the class para-AC<sup>0</sup><sup>†</sup>. Our proofs make an interesting use of Theorem 3.11. Recall that this “meta-theorem” states that all first-order properties of graphs, parameterized by the first-order property and the maximum degree of the graph, can be decided in para-AC<sup>0</sup><sup>†</sup>. Covering properties can be expressed using first-order formulas – but we make *no* requirement concerning the degree of the input graph. The trick is to first reduce the inputs to graphs of bounded degree and then apply the meta theorem. Such a two-step approach is typically in advanced applications of algorithmic meta-theorems.

► **Theorem 4.7.**  $p_{k,t}$ -PARTIAL-VERTEX-COVER  $\in$  para-AC <sub>$f(k+t)$</sub>   $\subseteq$  para-AC<sup>0</sup><sup>†</sup>.

► **Theorem 4.8.**  $p_t$ -EXACT-PARTIAL-VERTEX-COVER  $\in$  para-AC <sub>$f(t)$</sub>   $\subseteq$  para-AC<sup>0</sup><sup>†</sup>.

We conclude with the remark that the above results on finding vertex coverings for graphs cannot easily be extended to hypergraphs since for hypergraphs covering problems are typically hard for at least W[1].

**Clustering Problems.** Clustering algorithms have a wide variety of applications, for example in computational biology where we want to cluster genes and proteins or process transcription data [7]. A basic clustering problem for graphs is the following:

► **Problem 4.9** ( $p_{k,\ell}$ -CLUSTER-EDITING).

*Instance:* An undirected graph  $G = (V, E)$  and a numbers  $\ell$  and  $k$ .

*Parameter:*  $\ell, k$

*Question:* Can we add and/or delete up to  $k$  edges to or from  $G$  such that the resulting graph consists of  $\ell$  connected components, each of which is a clique?

A variant is  $p_k$ -MANY-CLUSTER-EDITING, where we just require that the edited graph consists of cliques and do not prescribe the number of clusters beforehand. This variant has been extensively studied, most notably by Gramm et al. [21] and Böcker [6] who showed its

fixed-parameter tractability. For the first version, algorithms based on color coding result in reasonable running times, but were recently outperformed by other approaches [19]. However, using a color coding approach is useful when we consider parallel algorithms:

► **Theorem 4.10.**  $p_{k,\ell}$ -CLUSTER-EDITING  $\in$  para-AC<sup>0</sup>.

► **Corollary 4.11.**  $p_k$ -MANY-CLUSTER-EDITING  $\in$  para-AC<sup>0</sup>.

We remark that if  $\ell$  is not no longer considered a parameter in cluster editing, the problem complexity increases only moderately:

► **Corollary 4.12.**  $p_k$ -CLUSTER-EDITING  $\in$  para-TC<sup>0</sup>.

Theorem 4.10 has another interesting corollary: Let  $p_{k,p}$ -COMPLETE- $p$ -PARTITE-EDITING be the problem of determining whether in a graph  $G$  we can add and/or remove up to  $k$  edges such that the resulting graph is complete  $p$ -partite, that is, its vertex set can be partitioned into exactly  $p$  non-empty sets such that there is an edge between two vertices if, and only if, they belong to two different sets. Since the complement of a complete  $p$ -partite graph is exactly a collection of  $p$  cliques, we get the following corollary:

► **Corollary 4.13.**

1.  $p_{k,p}$ -COMPLETE- $p$ -PARTITE-EDITING  $\in$  para-AC<sup>0</sup>.
2.  $p_k$ -COMPLETE- $p$ -PARTITE-EDITING  $\in$  para-TC<sup>0</sup>.

Finally, instead of looking for just one complete  $p$ -partite graph, we can look for several at the same time:

► **Problem 4.14** ( $p_{k,p}$ -MULTIPARTITE-CLUSTER-EDITING).

*Instance:* An undirected graph  $G = (V, E)$ , a natural number  $k$ , and a sequence of natural numbers  $p_1, p_2, \dots, p_\ell$ .

*Parameter:*  $k, p = p_1 + \dots + p_\ell$

*Question:* Can we add or delete  $k$  edges of  $G$  such that the resulting graph consist of  $d$  connected components  $C_1$  to  $C_\ell$  such that each  $C_i$  is a complete  $p_i$ -partite graph?

► **Theorem 4.15.**

1.  $p_{k,p}$ -MULTIPARTITE-CLUSTER-EDITING  $\in$  para-AC<sup>0</sup>
2.  $p_{k,\ell}$ -MULTIPARTITE-CLUSTER-EDITING  $\in$  para-TC<sup>0</sup>.

**Graph Separation Problems.** Graph separation problems are problems where we ask to separate a set of  $\ell$  vertices from the remaining graph by deleting at most  $k$  other vertices. They play a key role in many real-world network applications like finding communities or isolating dangerous vertices. While this problem is well-known to be NP-complete in the unparameterized setting and W[1]-hard in the parameterized setting for parameters  $k, \ell$ , and  $k + \ell$ , the complexity of the problem changes dramatically if we require the separated set of vertices to be connected:

► **Problem 4.16** ( $p_{k,\ell}$ -CUTTING- $\ell$ -CONNECTED-VERTICES).

*Instance:* An undirected graph  $G = (V, E)$  and two natural numbers  $k$  and  $\ell$ .

*Parameter:*  $k, \ell$

*Question:* Is there a partitioning of  $V$  into three sets  $X, S$ , and  $Y$  with  $|X| = \ell$  and  $|S| \leq k$  such that  $X$  is connected and for all  $\{x, y\} \in E$  with  $x \in X$  we have  $y \notin Y$ ?

Marx [22] showed that this problem is fixed-parameter tractable; Fomin, Golovach, and Korhonen [18] studied a similar version, namely  $p_{k,\ell}$ -CUTTING-AT-MOST- $\ell$ -VERTICES, in which the set  $X$  is not required to be connected and may be of size *at most*  $\ell$ , i. e.,  $1 < |X| \leq \ell$ , and for which Fomin et al. gave an FPT-algorithm based on color coding. The main idea is to colorize the given graph with two colors such that the vertices of the set  $X$  get colored with the first color and the vertices in  $S$  get the second color. Thus, we only have to find the solution within the vertices of the first color. This algorithm can be implemented in  $\text{para-AC}^{0\uparrow}$  and, moreover, works for  $p_{k,\ell}$ -CUTTING- $\ell$ -CONNECTED-VERTICES as well.

► **Theorem 4.17.**

1.  $p_{k,\ell}$ -CUTTING- $\ell$ -CONNECTED-VERTICES  $\in \text{para-AC}_{f(\ell)} \subseteq \text{para-AC}^{0\uparrow}$ .
2.  $p_{k,\ell}$ -CUTTING-AT-MOST- $\ell$ -VERTICES  $\in \text{para-AC}_{f(\ell)} \subseteq \text{para-AC}^{0\uparrow}$ .

We conclude with the remark that both problems can also be solved with algorithms similar to the ones presented above if we consider the *terminal versions* of these problems [18], i. e., there is a special terminal vertex  $t$  which has to be part of  $X$ . For this, we have to modify the above algorithms to consider only blue components that contain  $t$ .

## 5 Conclusion

We have seen that many natural parameterized problems can be solved in constant parallel time or in parallel time depending only on the parameters while doing only “FPT work.” We stress that our results are of a theoretical nature and do not directly give practical parallel implementations for the problems presented; but they show that such implementations are possible *in principle* for them. The core technique used in all proofs (at least indirectly) was *color coding*, which can be done in constant time and which is already used in practice.

This paper did not address lower bounds. While for  $\text{para-AC}^0$  this is not problematic since this class lies at the bottom of almost any hierarchy of parameterized classes, some problems in  $\text{para-AC}^{0\uparrow}$  might well “fall down” to  $\text{para-AC}^0$ . Here we only know an explicit lower bound for the distance problem, which does not lie in  $\text{para-AC}^0$ . Establishing lower bounds for other problems in  $\text{para-AC}^{0\uparrow}$  is therefore a reasonable research goal.

---

## References

- 1 F. N. Abu-Khzam, M. A. Langston, P. Shanbhag, and C. T. Symons. Scalable Parallel Algorithms for FPT Problems. *Algorithmica*, 45(3):269–284, 2006.
- 2 N. Alon, R. Yuster, and U. Zwick. Color-Coding. *Journal of the ACM*, 42(4):844–856, 1995.
- 3 Kazuyuki Amano.  $k$ -Subgraph Isomorphism on  $\text{AC}^0$  Circuits. *Computational Complexity*, 19(2):1016–3328, 2010.
- 4 David A. Mix Barrington and Neil Immerman. On uniformity within  $\text{NC}^1$ . *Journal of Computer and System Sciences*, 41(3):274–306, 1990.
- 5 Paul Beame, Russell Impagliazzo, and Toniann Pitassi. Improved depth lower bounds for small distance connectivity. *Computational Complexity*, 7(4):325–345, 1998.
- 6 S. Böcker. A Golden Ratio Parameterized Algorithm for Cluster Editing. In *Proceedings of the Twenty-Second International Workshop on Combinatorial Algorithms*, volume 7056 of *IWOCA '11*, pages 85–95. Springer Berlin Heidelberg, 2011.
- 7 S. Böcker and J. Baumbach. Cluster Editing. In *Proceedings of the Ninth Conference on Computability in Europe*, volume 7921 of *CiE '13*, pages 33–44. Springer Berlin Heidelberg, 2013.

- 8 L. Cai, J. Chen, R. G. Downey, and M. R. Fellows. Advice Classes of Parameterized Tractability. *Annals of Pure and Applied Logic*, 84(1):119–138, 1997.
- 9 Marco Cesati and Miriam Di Ianni. Parameterized parallel complexity. In *Proceedings of the Fourth International Euro-Par Conference*, volume 1470 of *Euro-Par '89*, pages 892–896. Springer Berlin Heidelberg, 1998.
- 10 H. Chen and M. Müller. The Fine Classification of Conjunctive Queries and Parameterized Logarithmic Space. *ACM Transactions on Computation Theory*, 7(2):7:1–7:27, 2014.
- 11 Y. Chen, J. Flum, and M. Grohe. Bounded Nondeterminism and Alternation in Parameterized Complexity Theory. In *Proceedings of the Eighteenth IEEE Conference on Computational Complexity, CCC '03*, pages 13–29. IEEE Computer Society, Los Alamitos, California, 2003.
- 12 R.G. Downey and M.R. Fellows. *Parameterized Complexity*. Springer, 1999.
- 13 Rodney G. Downey, Michael R. Fellows, and Kenneth W. Regan. Parameterized Circuit Complexity and the  $W$  Hierarchy. *Theoretical Computer Science*, 191(1–2):97–115, 1998.
- 14 M. Elberfeld, C. Stockhusen, and T. Tantau. On the Space Complexity of Parameterized Problems: Classes and Completeness. *Algorithmica*, 71(3):661–701, 2014.
- 15 M. Fellows, P. Heggernes, F. Rosamond, C. Sloper, and J. A. Telle. Finding  $k$  Disjoint Triangles in an Arbitrary Graph. In *Proceedings of the Thirtieth International Workshop on Graph-Theoretic Concepts in Computer Science*, volume 3353 of *WG '04*, pages 235–244. Springer Berlin Heidelberg, 2004.
- 16 J. Flum and M. Grohe. Describing Parameterized Complexity Classes. In *Proceedings of the Nineteenth Annual Symposium on Theoretical Aspects of Computer Science*, volume 2285 of *STACS '02*, pages 359–371. Springer Berlin Heidelberg, 2002.
- 17 J. Flum and M. Grohe. *Parameterized Complexity Theory*. Springer, Heidelberg, Germany, 2006.
- 18 F. V. Fomin, P. A. Golovach, and J. H. Korhonen. On the Parameterized Complexity of Cutting a Few Vertices from a Graph. In *Proceedings of the Thirty-Eight International Symposium of Mathematical Foundations of Computer Science*, volume 8087 of *MFCS '13*, pages 421–432. Springer, Heidelberg, Germany, 2013.
- 19 F. V. Fomin, S. Kratsch, M. Pilipczuk, M. Pilipczuk, and Y. Villanger. Tight Bounds for Parameterized Complexity of Cluster Editing. In *Proceedings of the Thirtieth International Symposium on Theoretical Aspects of Computer Science*, volume 20 of *STACS '13*, pages 30–43. International Symposium on Theoretical Aspects of Computer Science, 2013.
- 20 H. Gaifman. On Local and Non-Local Properties. In *Proceedings of the Herbrand Symposium, Logic Colloquium '81*, pages 105–135. North Holland, 1982.
- 21 J. Gramm, J. Guo, F. Hüffner, and R. Niedermeier. Graph-Modeled Data Clustering: Fixed-Parameter Algorithms for Clique Generation. In *Proceedings of the Fifth Italian Conference of Algorithms and Complexity*, volume 2653 of *CIAC '03*, pages 108–119. Springer Berlin Heidelberg, 2003.
- 22 D. Marx. Parameterized Graph Separation Problems. *Theoretical Computer Science*, 351(3):394–406, 2006.
- 23 I. Newman, P. Ragde, and A. Wigderson. Perfect Hashing, Graph Entropy, and Circuit Complexity. In *Proceedings of the Fifth Annual Structure in Complexity Theory Conference*, pages 91–99. IEEE Computer Society, Los Alamitos, California, 1990.
- 24 H. Vollmer. *Introduction to Circuit Complexity*. Springer, 1999.

## A Technical Appendix: Proofs

For the readers convenience, the claims of the proofs given in this appendix are repeated before the proofs.

**Claim of Theorem 3.2.** *There is a DLOGTIME-uniform family  $(C_{n,k,c})_{n,k,c \in \mathbb{N}}$  of AC-circuits without inputs such that each  $C_{n,k,c}$*

1. *outputs an  $(n, k, c)$ -universal coloring family (coded as a sequence of function tables),*
2. *has constant depth (independent of  $n$ ,  $k$ , or  $c$ ), and*
3. *has size at most  $O(n \log c \cdot c^{k^2} \cdot k^4 \log^2 n)$ .*

**Proof.** Define

$$\begin{aligned} \lambda_{p,a}(x) &= (a \cdot x \bmod p) \bmod k^2, \\ \Lambda'_{n,k} &= \{ \lambda_{p,a} \mid p \text{ is a prime with } p < k^2 \log n \text{ and } a \in \{0, \dots, p-1\} \}, \\ \Lambda_{n,k,c} &= \{ \omega \circ \lambda_{p,a} \mid \omega: \{0, \dots, k^2-1\} \rightarrow \{1, \dots, c\}, p < k^2 \log n, a \in \{0, \dots, p-1\} \}. \end{aligned}$$

It is well-known that  $\Lambda'_{n,k}$  is a family of  $k$ -perfect hash functions, i. e., for every subset  $S \subseteq \{1, \dots, n\}$  with  $|S| = k$  it contains a function that is injective on  $S$ , see [17]. Therefore, given a subset  $S$  and a function  $\mu: S \rightarrow \{1, \dots, c\}$ , some member of  $\lambda_{p,a} \in \Lambda'_{n,k}$  will map the members of  $S$  injectively to a subset  $S'$  of  $\{0, \dots, k^2-1\}$  and, then, some function  $\omega: \{0, \dots, k^2-1\} \rightarrow \{1, \dots, c\}$  will map  $S'$  in such a way that  $\omega \circ \lambda_{p,a}$  equals  $\mu$  on  $S$ . Consequently the set  $\Lambda_{n,k,c}$  is an  $(n, k, c)$ -universal coloring family. Notice that we use all  $p < k^2 \log n$  in the definition of  $\Lambda_{n,k,c}$  and, thus, including the prime numbers only indirectly. The sizes of the two sets can be bounded by  $|\Lambda'_{n,k}| \leq (k^2 \cdot \log n)^2$  and  $|\Lambda_{n,k,c}| \leq c^{k^2} \cdot (k^2 \cdot \log n)^2 = c^{k^2} k^4 \log^2 n$ . Each function in  $\Lambda_{n,k,c}$  can clearly be encoded in  $n \log_2 c$  bits.

For the construction of circuits  $C_{n,k,c}$  observe that they have no inputs and must just output  $\Lambda_{n,k,c}$  in a fixed encoding. Thus, we can, in principle, hardwire the complete output of  $C_{n,k,c}$  into a depth-0 circuit. The tricky part is, of course, arguing that the circuit family is DLOGTIME-uniform. However, this is surprisingly simple: Having a look at the definition of  $\Lambda_{n,k,c}$ , we see that computing the  $i$ th bit of its encoding only involves simple computations consisting of additions, multiplications, and modulo operations on  $i$  and the numbers  $n$ ,  $k$ , and  $c$ . Now, DLOGTIME-uniformity means that we have time *logarithmic* in the *unary* encodings of these numbers and hence *polynomial* time in their *binary* encodings. Since addition, multiplication, and modulo are clearly polynomial-time computable, we get the claim. ◀

**Claim of Lemma 3.3.**  $p$ -THRESHOLD  $\in$  para-AC<sup>0</sup>.

**Proof.** On input of a bitstring  $b$  of length  $n$  and a number  $t$ , use Theorem 3.2 to compute an  $(n, t, t)$ -universal coloring family. Now, if  $b$  contains at least  $t$  many 1's, then there is a coloring of the positions of  $b$  such that each color class contains at least one 1. Thus, it suffice to test in parallel for all colorings whether this is the case. ◀

**Claim of Theorem 3.5.** *Given two graphs  $H = (V_H, E_H)$  and  $G = (V_G, E_G)$  together with a tree-decomposition  $(T, \iota)$  of  $H$ . An embedding of  $H$  into  $G$  can be computed by an AC-circuit of depth  $O(\text{depth}(T))$  and size  $f(|V_H|) \cdot O(|V_G|^{\text{width}(T)})$ , if such an embedding exists.*

**Proof.** Let  $N$  denote the set of nodes of  $T$ . For a node  $n$  of  $T$ , let  $T_n$  be the subtree of  $T$  rooted at  $n$  and let  $N_n$  be the set of its vertices. Color  $G$  by a  $(|V_G|, |V_H|, |V_H|)$ -universal coloring family using Theorem 3.2 and test all members of this family in parallel. To simplify the notation, let us identify the  $|V_H|$  colors in each coloring with the vertices of  $V_H$ .

Let us call a subset  $X \subseteq V_G$  *colorful* if all vertices of  $X$  have a different color; let  $Y \subseteq V_H$  be the set of these colors and let  $\mu_X: Y \rightarrow X$  map each color  $y$  in  $Y$  to the vertex  $x$  in  $X$  having this color. Note that if there is an embedding  $\phi$  of  $H$  into  $G$ , for at least one coloring there is a  $\mu_X = \phi$  such that  $X = \phi(V_H) \subseteq V_G$  is the image of  $V_H$ .

Let us call a colorful subset  $I \subseteq V_G$  of size at most  $\text{width}(T) + 1$  *good for a node  $n$  of  $T$*  if there is a colorful superset  $J \supseteq I$  of vertices of  $V_G$  such that  $\mu_J$  is an injective homomorphism  $\mu_J: \bigcup_{m \in N_n} \iota(m) \rightarrow J$ . In other words,  $I$  can be extended to a solution of the embedding problem for the tree rooted at  $n$ .

Clearly, since the size of the  $\iota(n)$ 's and the  $I$ 's are restricted by the  $\text{width}(T) + 1$ , for a leaf  $n$  of  $T$  we can decide whether a subset is good for some node of  $T$  using only a constant number of AC-layers of width bounded by a function in  $\text{width}(T)$ . Also observe that the number of subsets of  $V_G$  of size at most  $\text{width}(T) + 1$  is bounded by  $|V_G|^{\text{width}(T)+1}$  and, thus, we can consider all of them in parallel in each layer of an AC-circuit.

We must now show that we can decide, using only as many layers as the depth of  $T$ , whether  $H$  has an embedding in  $G$ . In a first layer, we first determine for each leaf  $n$  of  $T$  the set of all good sets  $I$  for  $n$ . In the inductive step, consider a node  $n$  such that for all its children we have already determined which sets are good for them. Let  $I$  be a colorful set for which we must determine whether it is good for  $n$ . We claim that this is the case when two conditions are met:

1. The set  $I$  is “a correct embedding itself,” meaning that  $\mu_I$  is an injective homomorphism  $\mu_I: \iota(n) \rightarrow I$ .
2. The children of  $n$  “can be made consistent with  $I$ ,” meaning that for each child  $c$  of  $n$  in  $T$ , there is a colorful set  $I_c \subseteq V_G$  that is good for  $c$  and  $\mu_I$  and  $\mu_{I_c}$  are identical on  $\iota(n) \cap \iota(c)$ .

To see that these tests suffice in order to test whether  $I$  is good for  $n$ , just observe that by the definition of a tree-decomposition, all bags that contain a particular vertex  $h$  of  $V_H$  must form a connected subset of  $T$ . Our second condition ensures that when a given vertex  $g \in V_G$  has been picked as the image of  $h$  in some  $I$ , the same vertex must have been picked in all children and, thus, a partial homomorphism on  $I$  can be extended to a homomorphism on the vertices in the whole tree rooted at  $n$ .

To conclude the proof, we just observe that the two tests can clearly be implemented with a constant number of AC-layers.  $\blacktriangleleft$

**Claim of Lemma 3.8.**  $p_d\text{-DISTANCE} \in \text{para-AC}_{f(d)} \subseteq \text{para-AC}^{0\uparrow}$ .

**Proof.** Determining whether there is a path of length  $d$  from  $s$  to  $t$  in  $G$  is the same as asking whether the path  $P_d$  of length  $d$  can be embedded into a directed graph  $G$  with the start and end of the path marked appropriately so that they must be mapped to  $s$  and  $t$ , respectively. Since paths have tree-width 1, Theorem 3.5 and the second remark following it give the claim.  $\blacktriangleleft$

**Claim of Theorem 3.11.**  $p_{\phi, \delta}\text{-MC(FO)} \in \text{para-AC}_{f(\phi+\delta)} \subseteq \text{para-AC}^{0\uparrow}$ .

**Proof.** The first part of our proof is identical to the one given by Flum and Grohe in [16]: Let  $\phi$  be a formula given as input. For simplicity of presentation, we assume that the structure



$\mathcal{A}$  is actually an undirected graph  $G = (V, E)$  of maximum degree  $\delta$ . Let  $d(a, b)$  denote the distance of two vertices in  $G$  and let  $N_r(a) = \{b \in V \mid d(a, b) \leq r\}$  be the ball around  $a$  of radius  $r$  in  $G$ . Let  $G[N_r(a)]$  denote the subgraph of  $G$  induced on  $N_r(a)$ . By Gaifman's Theorem [20] we can rewrite  $\phi$  as a Boolean combination of formulas of the following form:

$$\exists x_1 \cdots \exists x_k \left( \bigwedge_{i \neq j} \psi_{\text{dist} > 2r}(x_i, x_j) \wedge \bigwedge_i \psi(x_i) \right)$$

where  $\psi_{\text{dist} > 2r}(x_i, x_j)$  is a standard formula expressing that  $d(x_i, x_j) > 2r$  and  $\psi$  is  $r$ -local, meaning that for all  $a \in V$  we have  $G \models \psi(a) \iff G[N_r(a)] \models \psi(a)$ . What remains to be done is to determine whether there are  $k$  vertices  $a_1$  to  $a_k$  in  $G$  such that the balls  $N_r(a_i)$  do not intersect and  $G[N_r(a_i)] \models \psi(a_i)$  holds for them.

At this point, we digress from the line of argument of Flum and Grohe, who now give a slightly involved space-efficient algorithm for determining the existence of such  $a_i$  without having to write them down (which is not possible in parameterized logarithmic space). Instead, we use color coding at this point: Introduce colors 1 to  $k$ . Since the maximum degree  $\delta$  is a parameter and  $r$  depends only on a parameter, the maximum size  $M$  of any  $N_r(a)$  is bounded by the parameter. This means that there is an  $(|V|, Mk, k+1)$ -universal coloring family such that for the vertices  $a_i$  from above at least one coloring has the following property: All vertices in  $N_r(a_i)$  have the same color  $i$ . This means that each  $N_r(a_i)$  is contained in a monochromatic connected component of  $G$  having color  $i$ .

It remains to test whether for each color  $i$  there is a vertex  $a_i$  such that  $N_r(a_i)$  has color  $i$  and  $G[N_r(a_i)] \models \psi(a_i)$  holds. For this, let some candidate  $a_i$  be given. We need to determine for a given vertex  $b$  whether  $d(a_i, b) \leq r$  where the distance is computed in the subgraph of  $G$  induced by the vertices of color  $i$ . In other words, we need to solve the problem  $p_{d, \delta}$ -UNDIRECTED-DISTANCE, which is parameterized over the distance  $d$  and the maximum degree  $\delta$  and which can be solved in  $\text{para-AC}_{f(r)}$  by Lemma 3.8.<sup>5</sup> Once the set  $N_r(a)$  of vertices reachable from a vertex  $a$  in at most  $r$  steps has been determined, we can create an isomorphic copy of  $G[N_r(a)]$  consisting just of an  $|N_r(a)| \times |N_r(a)|$  adjacency matrix in  $\text{para-AC}^0$ : Number the vertices of  $G$  in some manner (for instance, in the order they appear in the input), which also induces an ordering on the vertices of  $N_r(a)$ . The entry in row  $i$  and column  $j$  of the matrix is a 1 if the  $i$ th and the  $j$ th vertex in  $N_r(a)$  are connected by an edge in  $E_G$ . Determining which vertex is the  $i$ th vertex of  $N_r(a)$  can be done by a  $\text{para-AC}^0$  circuit by Lemma 3.3.

Given the adjacency matrix of  $G[N_r(a)]$ , we can clearly decide in  $\text{para-AC}^0$  whether  $G[N_r(a)] \models \phi$ , since the size of  $G[N_r(a)]$  depends only on the original input parameters. ◀

**Claim of Theorem 4.5.**  $p$ -VERTEX-COVER  $\in \text{para-AC}^0$ .

**Proof.** Let us reiterate the steps of the well-known Buss kernelization: Let  $G$  be an input graph and let  $k$  be the size of the sought vertex cover. First, we can determine, in parallel, all vertices  $v$  that have degree at least  $k+1$  and, as observed by Buss, all of these vertices must be part of any vertex cover of size at most  $k$ . Remove these vertices from  $G$  in parallel and then remove all isolated vertices. Buss' second observation is that if the remaining graph has more than  $k(k+1)$  vertices, no vertex cover of size  $k$  exists. Thus, we get a quadratic problem kernel.

<sup>5</sup> Flum and Grohe argue that the undirected distance problem can be solved in space  $f(r, \delta) + O(\log n)$ , since we just need  $r \log_2 \delta$  bits to describe a path of length  $r$  starting at a vertex  $a$  and can iterate over all possible paths with that many bits.

Elberfeld et al. [14] observe that the essential parallel steps of the algorithm are the following: (1) Check whether the degree of a vertex is at least  $k$ . (2) Checking whether there are at most  $k$  such vertices. (3) Checking whether there are at most  $k(k+1)$  vertices that are not only connected to the high-degree vertices. (4) Computing the subgraph induced by these  $k(k+1)$  vertices.

While Elberfeld et al. conclude at this point that the computation can be implemented by para-TC<sup>0</sup> circuits (“we just have to count”), Lemma 3.3 shows that the computation can be implemented using a para-AC<sup>0</sup> circuit: All counting involves thresholds depending only of the parameter. ◀

**Claim of Theorem 4.7.**  $p_{k,t}$ -PARTIAL-VERTEX-COVER  $\in$  para-AC <sub>$f(k+t)$</sub>   $\subseteq$  para-AC<sup>0†</sup>.

**Proof.** On input of a graph  $G$ , first test whether there is a vertex of degree at least  $t$ . If so, we can accept since this vertex alone already constitutes the desired cover. Otherwise, we know that the graph has a maximum degree bounded by the parameter and we can apply Theorem 3.11 to the following first-order formula, which depends only on  $k$  and  $t$ :

$$\underbrace{\exists x_1 \cdots \exists x_k}_{\text{the size-}k \text{ cover}} \underbrace{\exists a_1 \exists b_1 \cdots \exists a_t \exists b_t}_{\text{the } t \text{ covered edges}} \left( \phi_{\text{dist}}(a_1, b_1, \dots, a_t, b_t) \wedge \bigwedge_{i=1}^t (E(a_i, b_i) \wedge \bigvee_{j=1}^k a_i = x_j) \right).$$

Here,  $\phi_{\text{dist}}$  is a standard formula expressing that  $\{a_1, b_1\}, \dots, \{a_t, b_t\}$  are distinct sets. ◀

**Claim of Theorem 4.8.**  $p_t$ -EXACT-PARTIAL-VERTEX-COVER  $\in$  para-AC <sub>$f(t)$</sub>   $\subseteq$  para-AC<sup>0†</sup>.

**Proof.** We again wish to apply Theorem 3.11, but now the preprocessing step is a bit more complicated: Vertices of degree higher than  $t$  no longer constitute a solution, indeed, these vertices *cannot* be selected as part of a solution. However, we also cannot simply remove them as we did in the proof of Theorem 4.5 since parts of their neighbors might be chosen and the edges attached to them are then part of the  $t$  covered edges. The trick is to replace all vertices  $v$  of degree  $d > t$  by  $d$  new vertices and to add an edge from each of the former neighbors of  $v$  to exactly one of these  $d$  new vertices. (It is not difficult to implement these replacement steps in constant depth by also adding some unnecessary isolated vertices.) Let us color all “new” vertices red.

Once the graph has been preprocessed, it will once more have degree bounded in the parameter and we can apply Theorem 3.11 to a formula stating “there exist  $k$  vertices and  $t$  distinct edges such that the  $k$  vertices are not red, the  $t$  edges always have one endpoint among the  $k$  vertices and all edges of the graph having an endpoint among the  $k$  vertices are among the  $t$  edges.” ◀

**Claim of Theorem 4.10.**  $p_{k,\ell}$ -CLUSTER-EDITING  $\in$  para-AC<sup>0</sup>.

**Proof.** Let  $G = (V, E)$  and  $k$  be given as input. For the moment, assume that  $G$  can be clustered after  $k$  edge modifications and let  $C = \{C_1, \dots, C_\ell\}$  be a solution, that is, a partitioning of  $V$  such that  $R = \{\{u, v\} \mid u \in C_i, v \in C_j, i \neq j\} \cap E$  (these edges need to be removed) and  $A = \{\{u, v\} \mid u, v \in C_i\} \setminus E$  (these edges need to be added) together have size at most  $k$ . Define  $M = \bigcup R \cup \bigcup A$  as the set of all vertices attached to edges that need to be modified. Let us call a cluster  $C_i$  *partly modified* if  $C_i \not\subseteq M$ , and *completely modified* if  $C_i \subseteq M$ .

Our objective is to determine the clusters  $C_i$  without knowing them. Towards this aim, we apply color coding for an  $(n, 2k + \ell, 2)$ -universal coloring family with the colors blue and orange. If a clustering  $C$  exists, at least one coloring has the following two properties:

1. All vertices in  $M$  are colored blue. (Hence, all completely modified clusters will be completely blue.)
2. In each partly modified cluster  $D$ , at least one vertex  $D \setminus M$  is colored orange. Let  $d$  be the smallest such vertex with respect to the ordering of the vertices in the input.

To identify the partly modified clusters, we consider only orange vertices. Since all vertices in  $M$  are colored blue, edges incident to orange vertices will not change. Adjacent orange vertices will therefore belong to the same cluster and non-adjacent ones will belong to different clusters. Hence, we can identify the vertex  $d$  in each partly modified cluster: It is an orange vertex that is not adjacent to a smaller orange vertex. Observe that in  $G$  the vertex  $d$  has the following property:

(\*) All vertices connected to  $d$  form a partly modified cluster  $D$ .

Thus, the orange vertices that are not adjacent to smaller orange vertices induce a set of partly modified clusters. We count their number and count the total number  $m_1$  of modifications needed to form them. Since  $m_1 \leq k$  must hold, we can compute  $m_1$  in constant parallel time by Lemma 3.3.

It remains to consider the number of modifications needed to form the completely modified clusters. However, the total number of vertices in these clusters is at most  $2k$ ; so after conceptually removing all vertices that are part of partly modified clusters, at most  $2k$  vertices may remain. We can compute the subgraph induced by these vertices in constant time (using the same argument as in the proof of Theorem 3.11 for the construction of the subgraph induced on  $N_r(a)$ ) and then solve this kernel in constant time, yielding a minimum number  $m_2$  of modifications needed to create the completely covered clusters. We accept when  $m_1 + m_2 \leq k$  and the number of partly modified clusters and completely modified clusters is  $\ell$ . ◀

**Claim of Corollary 4.11.**  $p_k$ -MANY-CLUSTER-EDITING  $\in$  para-AC<sup>0</sup>.

**Proof.** The argument is similar to the one from Theorem 4.10, but we first apply a preprocessing to find components of  $G$  that are already cliques. Since the number of sought clusters (cliques) is not limited, no optimal solution will ever modify edges adjacent to vertices in such a clique and, thus, we can conceptually remove them from the input. To identify these vertices, let us call a vertex *cliquish* if all its neighbors are pairwise connected in  $G$ ; a property that we can easily test in constant time. The set  $X$  of cliquish vertices contains exactly all vertices of clusters already present in  $G$ .

The argument now continues as in Theorem 4.10, only we (1) completely ignore the vertices in  $X$  and (2) look for *up to*  $\ell = 2k$  partly or completely modified clusters rather than exactly  $\ell$  such clusters. ◀

**Claim of Corollary 4.12.**  $p_k$ -CLUSTER-EDITING  $\in$  para-TC<sup>0</sup>.

**Proof.** Our argument starts as in the proof of Corollary 4.11: We identify components that are already cliques. However, this time, we only add such a component to  $X$  if (1) its size is larger than  $k$  because, then, it cannot be part of any editing or if (2) there are  $2k$  such components of the same size  $s \leq k$  earlier in the input<sup>6</sup>, because we can apply any necessary

<sup>6</sup> A component  $A$  is “earlier in the input than a component  $B$ ” if there is a vertex in  $A$  whose position in the input is before all vertices in  $B$ .

modifications to these  $2k$  earlier components. Now, we count the number  $x$  of clusters in  $X$  (this is the only place where we need a  $\text{TC}^0$  circuit). If  $\ell' := \ell - x > 2k + k(k+1)/2$ , we know the coloring cannot lead to a solution. Otherwise, we ask whether the graph  $G$  without  $X$  together with the numbers  $\ell'$  and  $k$  is an instance of  $p_{k,\ell}$ -CLUSTER-EDITING. ◀

**Claim of Theorem 4.15.**

1.  $p_{k,p}$ -MULTIPARTITE-CLUSTER-EDITING  $\in$  para- $\text{AC}^0$
2.  $p_{k,\ell}$ -MULTIPARTITE-CLUSTER-EDITING  $\in$  para- $\text{TC}^0$ .

**Proof.** The start of our proof is identical to the one of Theorem 4.10 and we use the same terminology. The first difference concerns the property (\*) from the proof: When all clusters are cliques, an orange vertex in such a clique immediately identifies all vertices in it, namely as the set of its neighbors. For  $p$ -partite graphs, this is more difficult and we use a new definition of equivalence: Let us call two orange vertices *equivalent* if they are adjacent in  $G$  or if they have the same neighborhood in  $G$ . On orange vertices, this is, indeed, an equivalence relation and some vertices  $d$  will have the property that they are minimal with respect to this relation. For such vertices, we make a new observation:

(\*\*) The partly modified cluster containing  $d$  consists of  $d$ , all orange vertices equivalent to  $d$ , all blue vertices connected to any of these vertices, and possibly some additional vertices from  $M$ .

The “additional vertices from  $M$ ” arise for instance in a bipartite cluster when one shore contains only blue vertices and the other contains some blue and some orange vertices. Then the blue vertices of the second shore do not have any orange neighbors. Let  $Y$  be the set of all vertices that are “definitely identified” by rule (\*\*), that is, the set of all orange vertices and together with its neighborhood. If  $|V| - |Y| > 2k$  we can stop, since only vertices from  $M$  may be missing from  $Y$ . We can also stop when the number of identified partly modified clusters is more than  $\ell$ .

We can now identify all partly modified clusters  $D$  – except that some of the vertices in  $M$  may still be lacking –, but we do not yet know which one is  $C_1$ , which one is  $C_2$ , and so on. We try out, in parallel, all possible injective mappings from the identified clusters to the set of indices  $\{1, \dots, \ell\}$  together with all possible ways of mapping the at most  $2k$  vertices in  $V \setminus Y$  to  $\{1, \dots, \ell\}$ . Each pair of mappings determines a possible clustering  $\{C_1, \dots, C_\ell\}$  and we can now (1) compute the number of edge removals that are necessary to remove all edges between clusters and (2) use Corollary 4.13 to determine the minimal number of editing operations necessary to make the  $i$ th cluster  $C_i$  a complete  $p_i$ -partite graph in para- $\text{AC}^0$  or para- $\text{TC}^0$ , depending on whether the values of the  $p_i$  are parameters or not. We accept when the total number of modifications is at most  $k$ . ◀

**Claim of Theorem 4.17.**

1.  $p_{k,\ell}$ -CUTTING- $\ell$ -CONNECTED-VERTICES  $\in$  para- $\text{AC}_{f(\ell)} \subseteq$  para- $\text{AC}^{0\uparrow}$ .
2.  $p_{k,\ell}$ -CUTTING-AT-MOST- $\ell$ -VERTICES  $\in$  para- $\text{AC}_{f(\ell)} \subseteq$  para- $\text{AC}^{0\uparrow}$ .

**Proof.** We begin with the first item. For this, we make use of a family of  $(n, k + \ell, 2)$ -universal coloring functions. If  $G$  contains a set  $X$  of  $\ell$  vertices that can be separated from the remaining vertices by removing a set  $S$  of at most  $k$  vertices, then the family of coloring functions contains a coloring such that the vertices of  $X$  are colored with the first color, say blue, and the vertices of the set  $S$  are colored with the second color, say orange. Hence, we iterate over these colorings, and for each coloring we try to find  $X$  and  $S$  by searching

connected components of blue vertices in the graph. For this, we iterate over all vertices  $x$  of the graph and each time check whether it is part of a set  $X$  with the desired properties:

We can find out whether a vertex  $y$  has distance at most  $d$  from  $x$  in the blue subgraph in  $\text{para-AC}_{f(d)}$  by Lemma 3.8. If there is some  $y$  at distance  $\ell + 1$ , we know that the component containing  $x$  is too large and we can stop. Otherwise, we can identify all vertices  $y$  reachable from  $x$  inside the blue component in  $\text{para-AC}_{f(\ell)}$ . If the number of such vertices in  $\ell$  (we can test this even in constant depth using Lemma 3.3), test whether the number of orange vertices connected to any such  $y$  is at most  $k$  (again, this test can be done in constant time).

To prove the second item, we proceed in a similar way, but instead of searching for connected components of size  $\ell$ , we search for a blue component of size at most  $\ell$  that has at most  $k$  orange neighbors. ◀