Inf. Process. Lett. 185 (2024) 106471

Contents lists available at ScienceDirect

Information Processing Letters

journal homepage: www.elsevier.com/locate/ipl

FPT approximation and subexponential algorithms for covering few or many edges

Fedor V. Fomin [a,*,1], Petr A. Golovach [a,1], Tanmay Inamdar [b,2], Tomohiro Koana [c,3]

[a] Department of Informatics, University of Bergen, Bergen, 5004, Norway
[b] Indian Institute of Technology, Jodhpur, Jodhpur, 342030, India
[c] Algorithmics and Computational Complexity, Technische Universitat Berlin, Berlin, Germany

ARTICLE INFO

Keywords:
Parameterized complexity
Approximation algorithms
Partial vertex cover

ABSTRACT

We study the α-Fixed Cardinality Graph Partitioning (α-FCGP) problem, the generic local graph partitioning problem introduced by Bonnet et al. [Algorithmica 2015]. In this problem, we are given a graph G, two numbers k, p and $0 \le \alpha \le 1$, the question is whether there is a set $S \subseteq V$ of size k with a specified coverage function $\mathrm{cov}_\alpha(S)$ at least p (or at most p for the minimization version). The coverage function $\mathrm{cov}_\alpha(\cdot)$ counts edges with exactly one endpoint in S with weight α and edges with both endpoints in S with weight $1-\alpha$. α-FCGP generalizes a number of fundamental graph problems such as Densest k-Subgraph, Max k-Vertex Cover, and Max $(k, n-k)$-Cut.

A natural question in the study of α-FCGP is whether the algorithmic results known for its special cases, like Max k-Vertex Cover, could be extended to more general settings. One of the simple but powerful methods for obtaining parameterized approximation [Manurangsi, SOSA 2019] and subexponential algorithms [Fomin et al. IPL 2011] for Max k-Vertex Cover is based on the greedy vertex degree orderings. The main insight of our work is that the idea of greedy vertex degree ordering could be used to design fixed-parameter approximation schemes (FPT-AS) for $\alpha > 0$ and subexponential-time algorithms for the problem on apex-minor free graphs for maximization with $\alpha > 1/3$ and minimization with $\alpha < 1/3$.[4]

1. Introduction

In this work, we study a broad class of problems called α-Fixed Cardinality Graph Partitioning (α-FCGP), originally introduced by Bonnet et al. [2].[5] The input is a graph $G = (V, E)$, two non-negative integers k, p, and a real number $0 \le \alpha \le 1$. The question is whether there is a set $S \subseteq V$ of size exactly k with $\mathrm{cov}_\alpha(S) \ge p$ ($\mathrm{cov}_\alpha(S) \le p$ for the minimization variant), where

$$\mathrm{cov}_\alpha(S) := (1-\alpha) \cdot m(S) + \alpha \cdot m(S, V \setminus S).$$

Here, $m(S)$ is the number of edges with both endpoints in S, and $m(S, V \setminus S)$ is the number of edges with one endpoint in S and other in $V \setminus S$. We will call the maximization and minimization problems Max α-FCGP and Min α-FCGP, respectively. This problem generalizes many problems, namely, Densest k-Subgraph (for $\alpha = 0$), Max k-Vertex Cover[6] (for $\alpha = 1/2$), Max $(k, n-k)$-Cut (for $\alpha = 1$), and their minimization counterparts.

Although there are plethora of publications that study these special cases, the general α-FCGP has not received much attention, except for the work of Bonnet et al. [2], Koana et al. [20], and Schachnai and

* Corresponding author.
E-mail addresses: Fedor.Fomin@uib.no (F.V. Fomin), Petr.Golovach@uib.no (P.A. Golovach), taninamdar@gmail.com (T. Inamdar), tomohiro.koana@tu-berlin.de (T. Koana).

1 Supported by the Research Council of Norway via the project BWCA (grant no. 314528).
2 Supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (LOPRE grant no. 819416).
3 Supported by the DFG project DiPa (NI 369/21).
4 A preliminary version of the paper was published in MFCS 2023 [12].
5 Bonnet et al. [2] called the problem 'local graph partitioning problem', however we adopt the nomenclature from Koana et al. [20].
6 This is problem is also referred to as Partial Vertex Cover.

https://doi.org/10.1016/j.ipl.2024.106471
Received 31 August 2023; Received in revised form 22 December 2023; Accepted 3 January 2024
Available online 9 January 2024

Zehavi [24]. In this paper, we aim to demonstrate the wider potential of the existing algorithms designed for specific cases, such as Max k-Vertex Cover, by presenting an algorithm that can handle the more general problem of α-FCGP. Algorithms for these specific cases often rely on greedy vertex degree orderings. For instance, Manurangsi [21], showing that a $(1-\epsilon)$-approximate solution can be found in the set of $\mathcal{O}(k/\epsilon)$ vertices with the largest degrees, gave a $(1-\epsilon)$-approximation algorithm for Max k-Vertex Cover that runs in time $(1/\epsilon)^{\mathcal{O}(k)} \cdot n^{\mathcal{O}(1)}$. Fomin et al. [15] gave a $2^{\mathcal{O}(\sqrt{k})} \cdot n^{\mathcal{O}(1)}$-time algorithm for Max k-Vertex Cover on apex-minor graphs via bidimensionality arguments, by showing that an optimal solution S is adjacent to every vertex of degree at least $d+1$, where d is the minimum degree over vertices in S. In this work, we will give approximation algorithms as well as subexponential-time algorithms for apex-minor free graphs exploiting the greedy vertex ordering.

For approximation algorithms, we will show that both Max α-FCGP and Min α-FCGP admit *FPT Approximation Schemes* (FPT-AS) for $\alpha > 0$, i.e., there is an algorithm running in time $(\frac{k}{\epsilon\alpha})^{\mathcal{O}(k)} \cdot n^{\mathcal{O}(1)}$ that finds a set S of size k with $\mathsf{cov}_\alpha(S) \geq (1-\epsilon) \cdot \mathsf{OPT}$ (or $\mathsf{cov}_\alpha(S) \leq (1+\epsilon) \cdot \mathsf{OPT}$ for the minimization variant), where OPT denotes the optimal value of p. Previously, the special cases were known to admit FPT approximation schemes; see [23,17,18,21] for $\alpha = 1/2$ and [2] for $\alpha = 1$. In particular, the state-of-the-art running time for Max α-FCGP with $\alpha = 1/2$ is the aforementioned algorithm of Manurangsi that runs in time $(1/\epsilon)^{\mathcal{O}(k)} \cdot n^{\mathcal{O}(1)}$ for maximization (also for the minimization variant). We generalize this argument for $\alpha \geq 1/3$, leading to a faster FPT-AS for Max α-FCGP in this range. For $\alpha = 0$, the situation is more negative; Max α-FCGP (namely, Densest k-Subgraph) does not admit any $o(k)$-approximation algorithm with running time $f(k) \cdot n^{\mathcal{O}(1)}$ under the Strongish Planted Clique Hypothesis [22]. Min α-FCGP is also hard to approximate when $\alpha = 0$ since it encompasses Independent Set as a special case for $p = 0$.

Next, we discuss the regime of subexponential-time algorithms. Amini et al. [1] showed that Max k-Vertex Cover is FPT on graphs of bounded degeneracy, including planar graphs, giving a $k^{\mathcal{O}(k)} \cdot n^{\mathcal{O}(1)}$-time algorithm. They left it open whether it can be solved in time $2^{o(k)} \cdot n^{\mathcal{O}(1)}$. This was answered in the affirmative by Fomin et al. [15], who showed that Max k-Vertex Cover on apex-minor free graphs can be solved in time $2^{\mathcal{O}(\sqrt{k})} \cdot n^{\mathcal{O}(1)}$ time. Generalizing this result, we give a $2^{\mathcal{O}(\sqrt{k})} \cdot n^{\mathcal{O}(1)}$-time algorithm for Max α-FCGP with $\alpha > 1/3$ and Min α-FCGP with $\alpha < 1/3$. The complexity landscape of Max α-FCGP with $\alpha < 1/3$ (and Min α-FCGP with $\alpha > 1/3$) is not well understood. It is a long-standing open question whether Densest k-Subgraph on planar graphs is NP-hard [4]. Note that the special case Clique is trivially polynomial-time solvable on planar graphs because a clique on 5 vertices does not admit a planar embedding.

Further related work. As mentioned, special cases of α-FCGP when $\alpha \in \{0, 1/2, 1\}$ have been extensively studied. For instance, the W[1]-hardness for the parameter k has been long known for these special cases [3,11,16]. Both Max α-FCGP and Min α-FCGP are actually W[1]-hard for every $\alpha \in [0,1]$ with the exception $\alpha \neq 1/3$, as can be seen from a parameterized reduction from Clique and Independent Set on regular graphs. Note that α-Fixed Cardinality Graph Partitioning becomes trivial when $\alpha = 1/3$ because $\mathsf{cov}_\alpha(S) = \frac{1}{3} \cdot \sum_{v \in S} d(v)$ for any $S \subseteq V$ where $d(v)$ is the degree of v.

Bonnet et al. [2] gave a $(\Delta k)^{2k} \cdot n^{\mathcal{O}(1)}$-time algorithm for α-FCGP where Δ is the maximum degree. They also gave an algorithm with running time $\Delta^k \cdot n^{\mathcal{O}(1)}$ for Max α-FCGP with $\alpha > 1/3$ and Min α-FCGP with $\alpha < 1/3$. This result was strengthened by Schachnai and Zehavi [24]; they gave a $4^{k+o(k)}\Delta^k \cdot n^{\mathcal{O}(1)}$-time algorithm for any value of α. Koana et al. [20] showed that Max α-FCGP admits polynomial kernels on sparse families of graphs when $\alpha > 1/3$. For instance, Max α-FCGP admits a $k^{\mathcal{O}(d)}$-sized kernel where d is the degeneracy of the input graph. They also showed analogous results for Min α-FCGP with $\alpha < 1/3$.

Preliminaries. For an integer n, let $[n]$ denote the set $\{1, \cdots, n\}$.

We use the standard graph-theoretic notation and refer to the textbook of Diestel [10] for undefined notions. In this work, we assume that all graphs are simple and undirected. For a graph G and a vertex set S, let $G[S]$ be the subgraph of G induced by S. For a vertex v in G, let $d(v)$ be its *degree*, i.e., the number of its neighbors. For vertex sets X, Y, let $m(X) := |\{uv \in E \mid u, v \in X\}|$ and $m(X,Y) := |\{uv \in E \mid u \in X, v \in Y\}|$. In this work, an optimal solution for Max α-FCGP (and Min α-FCGP) is a vertex set S of size k such that $\mathsf{cov}_\alpha(S) \geq \mathsf{cov}_\alpha(S')$ (resp., $\mathsf{cov}_\alpha(S) \leq \mathsf{cov}_\alpha(S')$) for every vertex set of size k. A graph H is a *minor* of G if a graph isomorphic to H can be obtained from G by vertex and edge removals and edge contractions. Given a graph H, a family of graph $\mathcal{H}$ is said to be *H-minor free* if there is no $G \in \mathcal{H}$ having H as a minor. A graph H is an *apex graph* if H can be made planar by the removal of a single vertex.

We refer to the textbook of Cygan et al. [5] for an introduction to Parameterized Complexity and we refer to the paper of Marx [23] for an introduction to the area of parameterized approximation.

2. FPT approximation algorithms

In this section, we design an FPT Approximation Schemes for Max α-FCGP as well as Min α-FCGP parameterized by k and α, assuming $\alpha > 0$.

2.1. FPT-AS for Max/Min α-FCGP for any α > 0

Theorem 1. *For any $0 < \alpha \leq 1$ and $0 < \epsilon \leq 1$,* Max α-FCGP *and* Min α-FCGP *each admits an FPT-AS parameterized by k, ϵ and α. More specifically, given a graph $G = (V, E)$ and an integer k, there exists an algorithm that runs in time $\left(\frac{k}{\epsilon\alpha}\right)^{\mathcal{O}(k)} \cdot n^{\mathcal{O}(1)}$, and finds a set $S \subseteq V$ such that $\mathsf{cov}_\alpha(S) \geq (1-\epsilon) \cdot \mathsf{cov}_\alpha(O)$ for* Max α-FCGP *and $\mathsf{cov}_\alpha(S) \leq (1+\epsilon) \cdot \mathsf{cov}_\alpha(O)$ for* Min α-FCGP*, where $O \subseteq V$ is an optimal solution.*

For the case that $\mathsf{OPT} := \mathsf{cov}_\alpha(O)$ is large, the following greedy argument will be helpful.

Lemma 1. *For* Max α-FCGP*, let S be the set of k vertices with the largest degrees. Then, $\mathsf{cov}_\alpha(S) \geq \mathsf{OPT} - 2k^2$. For* Min α-FCGP*, let S be the set of k vertices with the smallest degrees. Then, $\mathsf{cov}_\alpha(S) \leq \mathsf{OPT} + 2k^2$.*

Proof. Without loss of generality, we assume that $O \neq S$. Let $S \setminus O = \{y_1, y_2, \ldots, y_t\}$, and $O \setminus S = \{w_1, w_2, \ldots, w_t\}$, where $1 \leq t \leq k$. Here, we index the vertices so that $d(y_i) \geq d(y_j)$ and $d(w_i) \geq d(w_j)$ (for Min α-FCGP, $d(y_i) \leq d(y_j)$ and $d(w_i) \leq d(w_j)$) for $i < j$. Note that due to the choice of S, it holds that $d(y_i) \geq d(w_i)$ ($d(y_i) \leq d(w_i)$ for Min α-FCGP) for each $1 \leq i \leq t$.

Now we define a sequence of solutions $O_0, O_1, \ldots, O_t$, where $O_0 = O$, and for each $1 \leq i \leq t$, $O_i := (O_{i-1} \setminus \{w_i\}) \cup \{y_i\}$. Note that $O_t = S$. We claim that for each $1 \leq i \leq t$, $\mathsf{cov}_\alpha(O_i) \geq \mathsf{cov}_\alpha(O_{i-1}) - 2k$ for Max α-FCGP and $\mathsf{cov}_\alpha(O_i) \leq \mathsf{cov}_\alpha(O_{i-1}) + 2k$ for Min α-FCGP. To this end, we note that O_i is obtained from O_{i-1} by removing w_i and adding y_i. Thus, $\mathsf{cov}_\alpha(O_i) = \mathsf{cov}_\alpha(O_{i-1}) - (\alpha m_1 + ((1-\alpha) - \alpha) \cdot m_2) + \alpha m_3 + ((1-\alpha) - \alpha) \cdot m_4$, where

$$m_1 := m(\{w_i\}, V \setminus O_{i-1}), \qquad m_2 := m(\{w_i\}, O_{i-1} \setminus \{w_i\}),$$
$$m_3 := m(\{y_i\}, V \setminus O_i), \qquad m_4 := m(\{y_i\}, O_i \setminus \{w_i\}).$$

Observe that $d(w_i) - k \leq m_1 \leq d(w_i)$, $d(y_i) - k \leq m_3 \leq d(y_i)$, and $0 \leq m_2, m_4 \leq k$. We consider Max α-FCGP first. We have that

$$\begin{aligned}\mathsf{cov}_\alpha(O_i) &= \mathsf{cov}_\alpha(O_{i-1}) + \alpha(m_3 - m_1) + (1-2\alpha)(m_4 - m_2) \\ &\geq \mathsf{cov}_\alpha(O_{i-1}) + \alpha(m_3 - m_1) - |(1-2\alpha)(m_4 - m_2)|.\end{aligned}$$

Since $m_3 - m_1 \geq d(y_i) - d(w_i) - k \geq -k$ and $|(1-2\alpha)(m_4 - m_2)| \leq k$, we obtain $\mathrm{cov}_\alpha(O_i) \geq \mathrm{cov}_\alpha(O_{i-1}) - 2k$, regardless of the value of α. We consider MIN α-FCGP next. It holds that

$$\begin{aligned}\mathrm{cov}_\alpha(O_i) &= \mathrm{cov}_\alpha(O_{i-1}) + \alpha(m_3 - m_1) + (1-2\alpha)(m_4 - m_2)\\ &\leq \mathrm{cov}_\alpha(O_{i-1}) + \alpha(m_3 - m_1) + |(1-2\alpha)(m_4 - m_2)|.\end{aligned}$$

Since $m_3 - m_1 \leq d(y_i) - d(w_i) + k \leq k$ and $|(1-2\alpha)(m_4 - m_2)| \leq k$, we obtain $\mathrm{cov}_\alpha(O_i) \leq \mathrm{cov}_\alpha(O_{i-1}) + 2k$, regardless of the value of α.

Therefore, $\mathrm{cov}_\alpha(O_t) \geq \mathrm{cov}_\alpha(O_0) - 2kt \geq \mathrm{OPT} - 2k^2$ for MAX α-FCGP and $\mathrm{cov}_\alpha(O_t) \leq \mathrm{cov}_\alpha(O_0) + 2kt \leq \mathrm{OPT} + 2k^2$ for MIN α-FCGP. □

Lemma 1 allows us to find an approximate solution when OPT is sufficiently large. The case that OPT is small remains. We use different approaches for MAX α-FCGP and MIN α-FCGP.

Algorithm for MAX α-FCGP. Let v_1 be a vertex with the largest degree. Our algorithm considers two cases depending on whether $d(v_1) > \Delta := \frac{2k^2}{\varepsilon\alpha} + k$. If $d(v_1) > \Delta$, we can argue that the set S from Lemma 1 a $(1-\varepsilon)$-approximate solution. To that end, we make the following observation.

Observation 1. *If $d(v_1) > \Delta$, then $2k^2 \leq \varepsilon \cdot \mathrm{cov}_\alpha(S)$.*

Proof. Note that $m(S, V \setminus S) = \sum_{u \in S} m(\{u\}, V \setminus S) \geq m(\{v_1\}, V \setminus S) \geq d(v_1) - k$, where the inequality follows from the fact that at most k edges incident to v_1 can have the other endpoint in S. This implies that

$$\mathrm{cov}_\alpha(S) \geq \alpha \cdot m(S, V \setminus S) \geq \alpha \cdot (d(v_1) - k) \geq \frac{2k^2}{\varepsilon},$$

where we use the assumptions that $0 < \alpha \leq 1$ and $d(v_1) \geq \Delta$. □

Thus, for $d(v_1) > \Delta$, we have $\mathrm{OPT} \leq \mathrm{cov}_\alpha(S) + 2k^2 \leq (1+\varepsilon) \cdot \mathrm{cov}_\alpha(S)$, and thus $\mathrm{cov}_\alpha(S) \geq (1-\varepsilon) \cdot \mathrm{OPT}$.

So assume that $d(v_1) < \Delta$. In this case, the maximum degree of the graph is bounded by $\Delta = \frac{2k^2}{\varepsilon\alpha} + k = \mathcal{O}(\frac{k^2}{\varepsilon\alpha})$. In this case, we solve the problem optimally using the algorithm of Shachnai and Zehavi [24] for MAX α-FCGP, that runs in time $4^{k+o(k)} \cdot \Delta^k \cdot n^{\mathcal{O}(1)}$, which is at most $\left(\frac{k^2}{\varepsilon\alpha}\right)^{\mathcal{O}(k)} \cdot n^{\mathcal{O}(1)}$. Combining both cases, we conclude the proof of Theorem 1 for MAX α-FCGP.

Algorithm for MIN α-FCGP. For MIN α-FCGP, our algorithm considers two cases depending on the value of OPT. If $\mathrm{OPT} \geq \frac{2k^2}{\varepsilon}$, then our algorithm returns the set S from Lemma 1. Note that $\mathrm{cov}_\alpha(S) \leq \mathrm{OPT} + 2k^2 \leq (1+\varepsilon) \cdot \mathrm{OPT}$.

Now suppose that $\mathrm{OPT} < \frac{2k^2}{\varepsilon}$. In this case, we know that O cannot contain a vertex of degree larger than $\Delta := \frac{2k^2}{\alpha\varepsilon} + k$, for otherwise, $\mathrm{cov}_\alpha(O) > \alpha(\Delta - k) \geq \mathrm{OPT}$, which is a contradiction. Thus, in this case the maximum degree of the graph is bounded by Δ, and again we can solve the problem optimally in time $\left(\frac{k^2}{\varepsilon\alpha}\right)^{\mathcal{O}(k)} \cdot n^{\mathcal{O}(1)}$, using the algorithm of Shachnai and Zehavi [24] for MIN α-FCGP.

Since the value of OPT is unknown to us, we cannot directly conclude which case is applicable. So we find a solution for each case and return a better one. This completes the proof of Theorem 1 for MIN α-FCGP.

2.2. *Faster FPT-AS for* MAX α-FCGP *when* $\alpha \geq 1/3$

In this section, we show that a simpler idea of Manurangsi [21] gives a faster FPT-AS for MAX α-FCGP when $\alpha \geq 1/3$, i.e., $\alpha \geq 1 - 2\alpha$, leading to the following theorem.

Theorem 2. *For any $1/3 \leq \alpha \leq 1$,* MAX α-FCGP *admits an FPT-AS running in time $\left(\frac{1}{\varepsilon}\right)^{\mathcal{O}(k)} \cdot n^{\mathcal{O}(1)}$.*

Proof. Let $0 < \varepsilon < 1$ be fixed and let us sort the vertices of $V(G)$ by their degrees (breaking ties arbitrarily). Let $V' \subseteq V(G)$ denote the $k + \lceil \frac{4k}{\varepsilon^2} \rceil$ vertices of the largest degrees. We show that V' contains a $(1-\varepsilon)$-approximate solution. Let O denote an optimal solution for MAX α-FCGP. Further define $O_i := O \cap V'$, $O_o := O \setminus V'$.

Let $U := V' \setminus O_i$ and let $U^* \subseteq U$ be a subset of size $|O_o|$ chosen uniformly at random from U. Let $\rho := \frac{|O_o|}{|U|} \leq \frac{k}{|U|} \leq \varepsilon^2/4$. In Lemma 2, we show that $\mathbb{E}[\mathrm{cov}_\alpha(O_i \cup U^*)] \geq (1-\varepsilon) \cdot \mathrm{cov}_\alpha(O)$, which implies that V' contains a $(1-\varepsilon)$-approximate solution. The algorithm simply enumerates all subsets of size k from V' and returns the best solution found. It follows that the running time of the algorithm is $\binom{|V'|}{k} \cdot n^{\mathcal{O}(1)} = \left(\frac{1}{\varepsilon}\right)^{\mathcal{O}(k)} \cdot n^{\mathcal{O}(1)}$. All that remains is the proof of the following lemma.

Lemma 2. $\mathbb{E}[\mathrm{cov}_\alpha(O_i \cup U^*)] \geq (1-\varepsilon) \cdot \mathrm{cov}_\alpha(O)$.

Proof. We fix some notation. For a vertex $u \in V$ and a subset $R \subseteq V$, we use $d_R(u)$ to denote the number of neighbors of u in R. When $R = V$, we use $d(u)$ instead of $d_V(u)$. Let $S = O_i \cup U^* = (O \setminus O_o) \cup U^*$. We want to analyze the expected value of $\mathrm{cov}_\alpha(S)$. To this end, we write $\mathrm{cov}_\alpha(S) = \mathrm{cov}_\alpha(O) - A + B$, where A is the "loss" in the objective due to removal of O_o and B is the "gain" in the objective due to addition of U^*, defined as follows.

$$A = \alpha \cdot m(O_o, V \setminus O_o) + (1-\alpha) \cdot m(O_o)$$

$$B = Q_1 + Q_2 - \alpha \cdot m(O_i, U^*), \text{where,}$$

$$Q_1 = \alpha \cdot m(U^*, V \setminus (U^* \cup O_i)) + (1-\alpha) \cdot m(U^*)$$

$$Q_2 = (1-\alpha) \cdot m(O_i, U^*)$$

Q_1 is the total contribution of the edges with at least one endpoint in U^* and other outside S, and Q_2 is the total contribution of edges with one endpoint in U^* and other in O_i. Note that the lemma is equivalent to showing that $\mathbb{E}[B - A] \geq -\varepsilon \cdot \mathrm{cov}_\alpha(O)$, where the expectation is over the choice of U^*.

Since A does not depend on the choice of U^*, we have

$$\begin{aligned}\mathbb{E}[A] = A &= \alpha \cdot m(O_o, V \setminus O_o) + (1-\alpha) \cdot m(O_o)\\ &\leq \alpha \cdot m(O_o, V \setminus O_o) + 2\alpha \cdot m(O_o) = \alpha \sum_{v \in O_o} d(v) \qquad (1)\end{aligned}$$

Here the inequality follows from $\alpha \geq 1/3$. Now let us consider $\mathbb{E}[B] = \mathbb{E}[Q_1 + Q_2 - \alpha \cdot m(U^*, O_i)]$. For any pair of distinct vertices u, v, let $X_{uv} = 1$ if $\{u, v\}$ is an edge and $X_{uv} = 0$ otherwise. Then, consider

$$\begin{aligned}\mathbb{E}[m(U^*, O_i)] &= \sum_{u \in U} \sum_{v \in O_i} X_{uv} \cdot \Pr(v \in U^*)\\ &= \rho \sum_{u \in U} \sum_{v \in O_i} X_{uv} \leq \frac{\varepsilon^2}{4} \cdot m(O_i, U) \qquad (2)\end{aligned}$$

Now we analyze $\mathbb{E}[Q_1]$. For every edge with one endpoint in U and the other in $V \setminus (U \cup O_i)$, there is a contribution α to Q_1 with probability ρ. Moreover, for every edge with both endpoints in U, the contribution to Q_1 is α with probability $2\rho(1-\rho)$ and $1-\alpha$ with probability ρ^2. Thus, we obtain

$$\begin{aligned}\mathbb{E}[Q_1] &= \alpha\rho \cdot m(U, V \setminus (U \cup O_i)) + (2\alpha\rho(1-\rho) + (1-\alpha)\rho^2) \cdot m(U)\\ &= \alpha\rho \cdot m(U, V \setminus (U \cup O_i)) + (2\alpha\rho + (1-3\alpha)\rho^2) \cdot m(U)\\ &\geq \alpha\rho \cdot (m(U, V \setminus (U \cup O_i)) + 2m(U)) = \alpha\rho \sum_{u \in U} d_{V \setminus O_i}(u). \qquad (3)\end{aligned}$$

Here the inequality is due to $\alpha \geq 1/3$.

Note that for any $u \in U$, and $v \in O_o$, $d(u) \geq d(v)$, which implies that for any $u \in U$, $d(u) \geq \frac{\sum_{v \in O_o} d(v)}{|O_o|}$. Therefore,

$$\sum_{u\in U} d(u) \geq \frac{|U|}{|O_o|}\sum_{v\in O_o} d(v) = \frac{1}{\rho}\cdot\sum_{v\in O_o} d(v) \tag{4}$$

Now we consider two cases.

Case 1: $\sum_{u\in U} d(u) \leq \frac{4}{\varepsilon}\cdot m(O_i, U)$. Then,

$$\frac{4}{\varepsilon}\cdot m(O_i,U) \geq \sum_{u\in U} d(u) \geq \frac{1}{\rho}\cdot\sum_{v\in O_o} d(v) \qquad \text{(Using (4))}$$

$$\implies \frac{4}{\varepsilon}\cdot m(O_i,U) \geq \frac{8}{\varepsilon^2}\cdot\sum_{v\in O_o} d(v) \qquad \text{(Since } \rho \leq \varepsilon^2/4\text{)}$$

$$\implies \varepsilon/2\cdot\alpha\cdot m(O_i,U) \geq \alpha\sum_{v\in O_o} d(v) \geq \mathbb{E}[A] \tag{5}$$

Note that we use (1) in the last inequality. Then consider,

$$\begin{aligned}\mathbb{E}[B-A] &\geq -\alpha\cdot\mathbb{E}[m(U^*,O_i)] - \mathbb{E}[A]\\ &\geq \frac{\varepsilon^2}{8}\alpha\cdot m(O_i,U) - \frac{\varepsilon}{2}\cdot\alpha\cdot m(O_i,U) \geq -\varepsilon\alpha\cdot m(O_i,U)\\ &\qquad\qquad \text{(Using (2) and (5))}\\ &\geq -\varepsilon\cdot \mathsf{cov}_\alpha(O) \end{aligned}\tag{6}$$

This finishes the first case.

Case 2: $\sum_{u\in U} d(u) > \frac{4}{\varepsilon}\cdot m(O_i, U)$. This implies that,

$$\frac{\varepsilon}{4}\cdot\sum_{u\in U} d(u) > \sum_{u\in U} d_{O_i}(u)$$

$$\implies \sum_{u\in U} d_{V\setminus O_i}(u) \geq \left(1-\frac{\varepsilon}{4}\right)\cdot\sum_{u\in U} d(u) \tag{7}$$

Then, plugging back in (3), we obtain,

$$\begin{aligned}\mathbb{E}[Q_1] &\geq \alpha\rho(1-\rho/2)\cdot(1-\varepsilon/4)\cdot\sum_{u\in U} d(u)\\ &\geq \alpha\rho(1-\varepsilon/2)\cdot\sum_{u\in U} d(u)\\ &\geq \alpha\rho(1-\varepsilon/2)\cdot\frac{|U|}{|O_o|}\sum_{v\in O_i} d(v) \qquad \text{(From (4))}\\ &\geq \alpha(1-\varepsilon/2)\cdot\sum_{v\in O_i} d(v)\\ &\geq A\cdot(1-\varepsilon/2)\end{aligned}\tag{8}$$

Note that we use (1) in the last inequality. Then, by (2) and (8), we obtain that,

$$\mathbb{E}[B-A] = \mathbb{E}[B] - \mathbb{E}[A] \geq -\varepsilon/2\cdot\mathbb{E}[A] - \alpha\varepsilon\cdot m(O_i, V\setminus O) \tag{9}$$

Now we argue that $\alpha\cdot m(O_i, V\setminus O) + \mathbb{E}[A] = \alpha\cdot m(O_i, V\setminus O) + A \leq \mathsf{cov}_\alpha(O)$. All edges with one endpoint in O_i and other outside O contribute α to the objective, which corresponds to the first term. Note that A is exactly the contribution of edges with at least one endpoint in O_o to the objective. Further, note that no such edge has one endpoint in O_i and other outside O, and thus is not counted in the first term. Thus, the sum of two terms is upper bounded by the objective, $\mathsf{cov}_\alpha(O)$. Plugging it back in (9), we obtain that $\mathbb{E}[B-A] \geq -\varepsilon\cdot\mathsf{cov}_\alpha(O)$ in the second case as well. ☐

With the proof of Lemma 2, the proof of Theorem 2 is complete. ☐

Example showing a gap for $\alpha < 1/3$. We now describe examples showing that for each fixed $\alpha < 1/3$, the above strategy of focusing on a bounded number of vertices of the largest degree does not lead to a $(1-\varepsilon)$-approximation, for large enough k. Let $f(k,\varepsilon)$ be an arbitrary function. Consider any $\alpha = 1/3 - \mu$, where $0 < \mu \leq 1/3$, and let $N \geq f(k,\varepsilon)$ be a large positive integer. The graph $G = (V, E)$ showing a gap is defined as follows. $V = H \uplus L \uplus O$, where $|H| = N, |L| = kN$ and $|O| = k$, thus $|V| = N(k+1)+k$. For each vertex $v\in H$, we attach k distinct vertices from L as pendants. Finally, we add all $\binom{k}{2}$ edges among the vertices of O, making in into a complete graph.

Each vertex of H has degree exactly k, each vertex of O has degree exactly $k-1$, and each vertex of L has degree exactly 1. This gives the sorted order of vertices by non-increasing degrees. It follows that the first $f(k,\varepsilon)$ vertices in the sorted order, say T, all belong to H. Furthermore, for any subset $S\subseteq T$ of size k, $\mathsf{cov}_\alpha(S) = \alpha\cdot k^2 = (\frac{1}{3}-\mu)\cdot k^2$. On the other hand, $\mathsf{cov}_\alpha(O) = (1-\alpha)\cdot\binom{k}{2} = (\frac{2}{3}+\mu)\cdot\frac{k^2-k}{2} \approx (\frac{1}{3}+\frac{\mu}{2})\cdot k^2$, assuming k is large enough. Hence, any k-sized subset $S\subseteq T$, $\frac{\mathsf{cov}_\alpha(S)}{\mathsf{cov}_\alpha(O)} < \frac{1/3-\mu}{1/3+\mu/2} \leq 1-3\mu$. Thus, T does not contain a $(1-\varepsilon)$-approximate solution for any $\varepsilon < 3\mu$. This shows that our analysis of Theorem 2 is tight for the range of $\alpha \geq 1/3$.

3. Subexponential FPT algorithm for Max α-FCGP on apex-minor free graphs

Fomin et al. [15] showed that Partial Vertex Cover on apex-minor free graphs can be solved in time $2^{\mathcal{O}(\sqrt{k})}\cdot n^{\mathcal{O}(1)}$. In this section, we will prove its generalization to Max α-FCGP as well as Min α-FCGP:

Theorem 3. *For an apex graph H, let $\mathcal{H}$ be a family of H-minor free graphs.*

- *For any $\alpha \geq 1/3$, Max α-FCGP for $\mathcal{H}$ can be solved in $2^{\mathcal{O}(\sqrt{k})}\cdot n^{\mathcal{O}(1)}$ time.*
- *For any $\alpha \leq 1/3$, Min α-FCGP for $\mathcal{H}$ can be solved in $2^{\mathcal{O}(\sqrt{k})}\cdot n^{\mathcal{O}(1)}$ time.*

We will give a proof for the maximization variant. The minimization variant follows analogously. Let $\sigma = v_1, v_2, \ldots, v_n$ be an ordering of vertices of V in the non-increasing order of degrees, with ties broken arbitrarily. That is, $d(v_1) \geq d(v_2) \geq \ldots \geq d(v_{n-1}) \geq d(v_n)$. We will denote the graph by $G = (V_\sigma, E)$ to emphasize the fact that the vertex set is ordered w.r.t. σ. We also let $V_\sigma^j = \{v_1,\ldots,v_j\}$. We first prove the following lemma.

Lemma 3. *Let $G = (V_\sigma, E)$ be a yes-instance for Max α-FCGP, where $1/3 \leq \alpha \leq 1$. Let $C = \left\{u_{i_1}, u_{i_2}, \ldots, u_{i_k}\right\}$ be the lexicographically smallest solution for Max α-FCGP and $u_{i_k} = v_j$ for some j. Then C is a dominating set of size k for $G[V_\sigma^j]$.*

Proof. Suppose for the contradiction that C is not a dominating set for $G[V_\sigma^j]$. Then, there exists a vertex v_i with $1 \leq i < j$ such that $N[v_i]\cap C = \emptyset$. Set $C' = (C\setminus\{v_j\})\cup\{v_i\}$. Note that $d(v_i)\geq d(v_j)$. Define the following:

$$m_1 = m(\{v_j\}, V\setminus C),$$

$$m_2 = m(\{v_j\}, C\setminus\{v_j\}),$$

$$m_3 = m(\{v_i\}, (V\setminus C)\cup\{v_j\}) = d(v_i),$$

$$m_4 = m(\{v_i\}, C\setminus\{v_j\}) = 0.$$

We will show that C' is another solution for the Max α-FCGP instance. Since $C'\setminus\{v_i\} = C\setminus\{v_j\}$, it suffices to show that

$$\mathsf{cov}_\alpha(C') - \mathsf{cov}_\alpha(C) = (\mathsf{cov}_\alpha(C') - \mathsf{cov}_\alpha(C'\setminus\{v_i\})) - (\mathsf{cov}_\alpha(C) - \mathsf{cov}_\alpha(C\setminus\{v_j\}))$$

is nonnegative. By definition,

$$\mathsf{cov}_\alpha(C') - \mathsf{cov}_\alpha(C'\setminus\{v_i\}) = \alpha\cdot m_3 + ((1-\alpha)-\alpha)\cdot m_4 = \alpha\cdot d(v_i) \text{ and}$$

$$\begin{aligned}\mathsf{cov}_\alpha(C) - \mathsf{cov}_\alpha(C\setminus\{v_j\}) &= \alpha\cdot m_1 + ((1-\alpha)-\alpha)\cdot m_2 \leq \alpha\cdot(m_1+m_2)\\ &= \alpha\cdot d(v_j),\end{aligned}\tag{10}$$

where the inequality is due to the assumption that $\alpha \geq 1/3$. Therefore,

$$\mathsf{cov}_\alpha(C') - \mathsf{cov}_\alpha(C) = \alpha \cdot (d(v_i) - d(v_j)) \geq 0,$$

which is a contradiction to the assumption that C is the lexicographically smallest solution for MAX α-FCGP. □

In view of Lemma 3, we can use the following approach to search for the lexicographically smallest solution C. First, we guess the last vertex v_j of C in the ordering σ, i.e., we search for a solution C such that $v_j \in C$ and $C \subseteq V_\sigma^j$. If $G[V_\sigma^j]$ has no dominating set of size at most, say $2k$, then we reject. This can be done in polynomial time, since DOMINATING SET admits a PTAS on apex-minor free graphs [7]. We thus may assume that there is a dominating set of size $2k$ in $G[V_\sigma^j]$. It is known that an apex-minor free graph with a dominating set of size κ has treewidth $\mathcal{O}(\sqrt{\kappa})$, where $\mathcal{O}$ hides a factor depending on the apex graph whose minors are excluded [6,9,13]. We can use a constant-factor approximation algorithm of Demaine [8] to find a tree decomposition $\mathcal{T}$ of width $w \in \mathcal{O}(\sqrt{k})$. Finally, we solve the problem via dynamic programming over the tree decomposition. Bonnet et al. [2] gave a $\mathcal{O}^*(2^w)$-time algorithm that solves MAX α-FCGP with a tree decomposition of width w given. We need to solve a slightly more general problem because $\mathcal{T}$ is the tree decomposition is over V_σ^j. To remove $V \setminus V_\sigma^j$, we introduce a weight $\omega : V_\sigma^j \to \mathbb{N}$ defined by $\omega(v) = |N(v) \cap (V \setminus V_\sigma^j)|$. The objective is then to maximize $\mathsf{cov}_\alpha(C) + \alpha \sum_{v \in C} \omega(C)$. The dynamic programming algorithm of Bonnet et al. can be adapted to solve this weighted variant in the same running time. Thus, we obtain a $2^{\mathcal{O}(\sqrt{k})} \cdot n^{\mathcal{O}(1)}$-time algorithm for MAX α-FCGP.

For MIN α-FCGP, we can show the following lemma whose proof is omitted because it is almost analogous to the previous one. The only change is that, V_σ refers to the vertices in the non-decreasing order of degrees. Also, we consider the regime where $0 \leq \alpha \leq 1/3$, which implies $\alpha \leq 1 - 2\alpha$, which would give the reverse inequality in (10).

Lemma 4. *Let $G = (V_\sigma, E)$ be a yes-instance for* MAX *α-FCGP, where $0 \leq \alpha \leq 1/3$. Let $C = \left\{u_{i_1}, u_{i_2}, \ldots, u_{i_k}\right\}$ be the lexicographically smallest solution for* MAX *α-FCGP and $u_{i_k} = v_j$ for some j. Then C is a dominating set of size k for $G[V_\sigma^j]$.*

With this lemma at hand, an analogous algorithm solves MIN α-FCGP in $2^{\mathcal{O}(\sqrt{k})} \cdot n^{\mathcal{O}(1)}$ time, thereby proving Theorem 3.

4. Conclusion

In this paper, we demonstrated that the algorithms exploiting the "degree-sequence" that have been successful for designing algorithms for MAX k-VERTEX COVER naturally generalize to MAX/MIN α-FCGP. Specifically, we designed FPT approximations for MAX/MIN α-FCGP parameterized by k, α, and ε, for any $\alpha \in (0, 1]$. For MAX α-FCGP, this result is tight since, when $\alpha = 0$, the problem is equivalent to DENSEST k-SUBGRAPH, which is hard to approximate in FPT time [22]. We also designed subexponential FPT algorithms for MAX α-FCGP (resp. MIN α-FCGP) for the range $\alpha \geq 1/3$ (resp. $\alpha \leq 1/3$) on any apex-minor closed family of graphs. It is a natural open question whether one can obtain subexponential FPT algorithms for MAX/MIN α-FCGP for the entire range $\alpha \in [0, 1]$. A notable special case is that of DENSEST k-SUBGRAPH on planar graphs. In this case, the problem is not even known to be NP-hard, if the subgraph is allowed to be disconnected. For the DENSEST CONNECTED k-SUBGRAPH problem, it was shown by Keil and Brecht [19] that the problem is NP-complete on planar graphs. From the other side, it can be shown that DENSEST CONNECTED k-SUBGRAPH admits a subexponential in k randomized algorithm on apex-minor free graphs using the general results of Fomin et al. [14]. Thus, dealing with disconnected dense subgraphs is difficult for both algorithms and lower bounds.

CRediT authorship contribution statement

Fedor V. Fomin: Conceptualization, Writing – original draft, Writing – review & editing. **Petr A. Golovach:** Conceptualization, Writing – original draft, Writing – review & editing. **Tanmay Inamdar:** Writing – original draft, Writing – review & editing, Conceptualization. **Tomohiro Koana:** Conceptualization, Writing – original draft, Writing – review & editing.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Tanmay Inamdar reports financial support was provided by European Research Council. Fedor V. Fomin, Petr A. Golovach reports financial support was provided by Research Council of Norway. Tomohiro Koana reports financial support was provided by German Research Foundation.

Data availability

No data was used for the research described in the article.

References

[1] Omid Amini, Fedor V. Fomin, Saket Saurabh, Implicit branching and parameterized partial cover problems, J. Comput. Syst. Sci. 77 (6) (2011) 1159–1171.

[2] Édouard Bonnet, Bruno Escoffier, Vangelis Th. Paschos, Emeric Tourniaire, Multi-parameter analysis for local graph partitioning problems: using greediness for parameterization, Algorithmica 71 (3) (2015) 566–580.

[3] Leizhen Cai, Parameterized complexity of cardinality constrained optimization problems, Comput. J. 51 (1) (2008) 102–121.

[4] Derek G. Corneil, Yehoshua Perl, Clustering and domination in perfect graphs, Discrete Appl. Math. 9 (1) (1984) 27–39, https://doi.org/10.1016/0166-218X(84)90088-X.

[5] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, Saket Saurabh, Parameterized Algorithms, Springer, 2015.

[6] Erik D. Demaine, Fedor V. Fomin, Mohammad Taghi Hajiaghayi, Dimitrios M. Thilikos, Bidimensional parameters and local treewidth, SIAM J. Discrete Math. 18 (3) (2004) 501–511, https://doi.org/10.1137/S0895480103433410.

[7] Erik D. Demaine, Mohammad Taghi Hajiaghayi, Bidimensionality: new connections between FPT algorithms and PTASs, in: Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2005, Vancouver, British Columbia, Canada, January 23–25, 2005, SIAM, 2005, pp. 590–601, http://dl.acm.org/citation.cfm?id=1070432.1070514.

[8] Erik D. Demaine, Mohammad Taghi Hajiaghayi, Ken-ichi Kawarabayashi, Algorithmic graph minor theory: decomposition, approximation, and coloring, in: 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005), Proceedings, 23–25 October 2005, Pittsburgh, PA, USA, IEEE Computer Society, 2005, pp. 637–646.

[9] Erik D. Demaine, Mohammad Taghi Hajiaghayi, Linearity of grid minors in treewidth with applications through bidimensionality, Combinatorica 28 (1) (2008) 19–36, https://doi.org/10.1007/s00493-008-2140-4.

[10] Reinhard Diestel, Graph Theory, 4th edition, Graduate Texts in Mathematics, vol. 173, Springer, 2012.

[11] Rodney G. Downey, Vladimir Estivill-Castro, Michael R. Fellows, Elena Prieto-Rodriguez, Frances A. Rosamond, Cutting up is hard to do: the parameterized complexity of k-cut and related problems, in: Computing: the Australasian Theory Symposiumm, CATS 2003, Adelaide, SA, Australia, February 4–7, 2003, in: Electronic Notes in Theoretical Computer Science, vol. 78, Elsevier, 2003, pp. 209–222.

[12] Fedor V. Fomin, Petr A. Golovach, Tanmay Inamdar, Tomohiro Koana, FPT approximation and subexponential algorithms for covering few or many edges, in: Jérôme Leroux, Sylvain Lombardy, David Peleg (Eds.), 48th International Symposium on Mathematical Foundations of Computer Science, MFCS 2023, August 28 to September 1, 2023, Bordeaux, France, in: LIPIcs, vol. 272, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023, 46.

[13] Fedor V. Fomin, Petr A. Golovach, Dimitrios M. Thilikos, Contraction bidimensionality: the accurate picture, in: Algorithms - ESA 2009, 17th Annual European Symposium, Proceedings, Copenhagen, Denmark, September 7–9, 2009, in: Lecture Notes in Computer Science, vol. 5757, Springer, 2009, pp. 706–717.

[14] Fedor V. Fomin, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, Saket Saurabh, Subexponential parameterized algorithms for planar and apex-minor-free graphs via low treewidth pattern covering, SIAM J. Comput. 51 (6) (2022) 1866–1930, https://doi.org/10.1137/19m1262504.

[15] Fedor V. Fomin, Daniel Lokshtanov, Venkatesh Raman, Saket Saurabh, Subexponential algorithms for partial cover problems, Inf. Process. Lett. 111 (16) (2011) 814–818, https://doi.org/10.1016/j.ipl.2011.05.016.

[16] Jiong Guo, Rolf Niedermeier, Sebastian Wernicke, Parameterized complexity of vertex cover variants, Theory Comput. Syst. 41 (3) (2007) 501–520.

[17] Anupam Gupta, Euiwoong Lee, Jason Li, Faster exact and approximate algorithms for k-cut, in: 59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7–9, 2018, IEEE Computer Society, 2018, pp. 113–123.

[18] Anupam Gupta, Euiwoong Lee, Jason Li, An FPT algorithm beating 2-approximation for k-cut, in: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7–10, 2018, SIAM, 2018, pp. 2821–2837.

[19] J. Mark Keil, Timothy B. Brecht, The complexity of clustering in planar graphs, J. Comb. Math. Comb. Comput. 9 (1991) 155–159.

[20] Tomohiro Koana, Christian Komusiewicz, André Nichterlein, Frank Sommer, Covering many (or few) edges with k vertices in sparse graphs, in: Proceedings of the 39th International Symposium on Theoretical Aspects of Computer Science (STACS '22), in: LIPIcs, vol. 219, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, 42.

[21] Pasin Manurangsi, A note on max k-vertex cover: faster FPT-AS, smaller approximate kernel and improved approximation, in: 2nd Symposium on Simplicity in Algorithms, SOSA 2019, January 8–9, 2019, San Diego, CA, USA, in: OASIcs, vol. 69, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019, 15.

[22] Pasin Manurangsi, Aviad Rubinstein, Tselil Schramm, The strongish planted clique hypothesis and its consequences, in: 12th Innovations in Theoretical Computer Science Conference, ITCS 2021, Virtual Conference, January 6–8, 2021, in: LIPIcs, vol. 185, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021, 10.

[23] Dániel Marx, Parameterized complexity and approximation algorithms, Comput. J. 51 (1) (2008) 60–78, https://doi.org/10.1093/comjnl/bxm048.

[24] Hadas Shachnai, Meirav Zehavi, Parameterized algorithms for graph partitioning problems, Theory Comput. Syst. 61 (3) (2017) 721–738, https://doi.org/10.1007/s00224-016-9706-0.