Minimum Eccentricity Shortest Path Problem with Respect to Structural Parameters

Martin Kučera[✉] and Ondřej Suchý[✉]

Department of Theoretical Computer Science, Faculty of Information Technology,
Czech Technical University in Prague, Prague, Czech Republic
`martin@mkucera.cz`, `ondrej.suchy@fit.cvut.cz`

Abstract. The MINIMUM ECCENTRICITY SHORTEST PATH PROBLEM consists in finding a shortest path with minimum eccentricity in a given undirected graph. The problem is known to be NP-complete and W[2]-hard with respect to the desired eccentricity. We present fpt algorithms for the problem parameterized by the modular width, distance to cluster graph, the combination of distance to disjoint paths with the desired eccentricity, and maximum leaf number.

Keywords: Graph theory · Minimum eccentricity shortest path · Parameterized complexity · Fixed-parameter tractable

1 Introduction

The MINIMUM ECCENTRICITY SHORTEST PATH (MESP) problem asks, given an undirected graph and an integer k , to find a shortest path with eccentricity at most k —a shortest path (between its endpoints) whose distance to all other vertices in the graph is at most k . The shortest path achieving the minimum k may be viewed as the “most accessible”, and as such, may find applications in communication networks, transportation planning, water resource management, and fluid transportation [7]. Some large graphs constructed from reads similarity networks of genomic data appear to have very long shortest paths with low eccentricity [17]. Furthermore, MESP can be used to obtain the best to date approximation for a minimum distortion embedding of a graph into the line [7] which has applications in computer vision [16], computational biology and chemistry [11, 12]. The eccentricity of MESP is closely tied to the notion of laminarity (minimum eccentricity of the graph’s diameter) [1].

MESP was introduced by Dragan and Leitert [7] who showed that it is NP-hard on general graphs and constructed a slice-wise polynomial (XP) algorithm,

This paper is eligible for the best student paper award.

M. Kučera—Work of this author was supported by the Student Summer Research Program 2020 of FIT CTU in Prague, and by grant SGS20/212/OHK3/3T/18.

O. Suchý—The author acknowledges the support of the OP VVV MEYS funded project CZ.02.1.01/0.0/0.0/16.019/0000765 “Research Center for Informatics”.

© The Author(s) 2021

P. Flocchini and L. Moura (Eds.): IWOCA 2021, LNCS 12757, pp. 442–455, 2021.

https://doi.org/10.1007/978-3-030-79987-8_31

Fig. 1. Hasse diagram of the boundedness relation between structural parameters explored in this paper and related ones. An edge between a parameter **A** above and a parameter **B** below means that whenever **A** is bounded for some graph class, then so is **B**. The parameters for which MESP is FPT are in green (dark if the result is described in this paper, light if implied by those described). Yellow represents a parameter for which MESP is FPT in combination with the desired eccentricity. The complexity for parameters in gray is open. (Color figure online)

which finds a shortest path with eccentricity at most k in a graph with n vertices and m edges in $\mathcal{O}(n^{2k+2}m)$ time. They also presented a linear-time algorithm for trees. Additionally, they developed a 2-approximation, a 3-approximation, and an 8-approximation algorithm that runs in $\mathcal{O}(n^3)$ time, $\mathcal{O}(nm)$ time, and $\mathcal{O}(m)$ time, respectively. Birmelé et al. [1] further improved the 8-approximation to a 3-approximation, which still runs in linear time. Dragan and Leitert [8] showed that MESP can be solved in linear time for distance-hereditary graphs (generalizing the previous result for trees) and in polynomial time for chordal and dually chordal graphs. Later, they proved [9] that the problem is NP-hard even for bipartite subcubic planar graphs, and W[2]-hard with respect to the desired eccentricity for general graphs. Furthermore, they showed that in a graph with a shortest path of eccentricity k , a minimum k -dominating set can be found in $n^{\mathcal{O}(k)}$ time. A related problem of finding shortest isometric cycle was studied by Birmelé et al. [3]. Birmelé et al. [2] studied a generalization of MESP, where the task is to decompose a graph into subgraphs with bounded shortest-path eccentricity, the hub-laminar decomposition.

Our Contribution. We continue the research direction of MESP in structured graphs [8], focusing on parameters which can measure the amount of structure present in the graph. As treewidth, the most famous of such measures, remains out of reach, we focus on parameters providing even more structure (see [14] for an overview of some graph measures). We provide fpt algorithms for the problem with respect to the modular width, distance to cluster graph, distance to disjoint paths combined with the desired eccentricity, and maximum leaf number (see Fig. 1 for an overview of our results).

Outline. In Sect. 2, we provide necessary notations and formal definitions. In Sect. 3, we describe our parameterized algorithms. In Sect. 4, we discuss possible future work.

Proofs omitted in this extended abstract due to space restrictions can be found in the full version of the paper, preprint is available on ArXiv [13].

2 Preliminaries

We consider finite connected unweighted undirected simple loopless graphs.

We refer to Diestel [6] for graph notions.

For a graph $G = (V, E)$ we denote $n = |V|$ and $m = |E|$. We denote $G[S]$ the induced subgraph of G on vertices $S \subseteq V$ and $G \setminus S = G[V \setminus S]$.

We denote an ordered sequence of elements $\mathbf{s} = (s_1, \dots, s_{|\mathbf{s}|})$. For two sequences $\mathbf{s} = (s_1, \dots, s_{|\mathbf{s}|})$, $\mathbf{t} = (t_1, \dots, t_{|\mathbf{t}|})$ we denote their concatenation

$$\mathbf{s} \frown \mathbf{t} = (s_1, \dots, s_{|\mathbf{s}|}, t_1, \dots, t_{|\mathbf{t}|}).$$

A *path* is a sequence of vertices where every two consecutive vertices are adjacent. The first and last vertices of the path are called its *endpoints*. A *path between u and v* or *u - v -path* is a path with endpoints u and v . The length of a path P is the number of edges in it, i.e., $|P| - 1$. A u - v -path is *shortest* if it has the least length among all u - v -paths. The distance $d_G(u, v)$ between two vertices $u, v \in V$ is the length of the shortest u - v -path.

The distance between a vertex $u \in V$ and a set of vertices $S \subseteq V$ is $d_G(u, S) = \min_{s \in S} d_G(u, s)$. The eccentricity of a set $S \subseteq V$ is $\text{ecc}_G(S) = \max_{u \in V} d_G(u, S)$. For a path P , we use P instead of $V(P)$ for its set of vertices, if there is no risk of confusion, e.g., $d_G(u, P) = d_G(u, V(P))$ and $\text{ecc}_G(P) = \text{ecc}_G(V(P))$.

For vertex $u \in V$ we denote $N_G(u) = \{v \mid \{u, v\} \in E\}$ the open neighborhood, $N_G[u] = N_G(u) \cup \{u\}$ the closed neighborhood, and $N_G^k[u] = \{v \in V \mid d_G(u, v) \leq k\}$ the closed k -neighborhood of u .

In this paper, we focus on the following problem.

MINIMUM ECCENTRICITY SHORTEST PATH PROBLEM (MESP)

Input: An undirected graph G , desired eccentricity $k \in \mathbb{N}$.

Question: Is there a path P in G which is a shortest path between its endpoints with $\text{ecc}_G(P) \leq k$?

A parameterized problem Π is *fixed parameter tractable (FPT)* with respect to a parameter k if there is an algorithm solving any instance of Π with size n in $f(k) \cdot n^{O(1)}$ time for some computable function f . Such an algorithm is called a *parameterized* or an *fpt algorithm*. See Cygan et al. [5] for more information on parameterized algorithms.

In this paper, we present fpt algorithms for MESP with respect to the following structural parameters.

Definition 1 (Modular width, [10]). Consider graphs that can be obtained from an algebraic expression that uses the following operations:

- (O1) create an isolated vertex;
- (O2) the disjoint union of 2 disjoint graphs (the disjoint union of graphs G_1 and G_2 is the graph $(V(G_1) \cup V(G_2), E(G_1) \cup E(G_2))$);
- (O3) the complete join of 2 disjoint graphs (the complete join of graphs G_1 and G_2 is the graph $(V(G_1) \cup V(G_2), E(G_1) \cup E(G_2) \cup \{\{v, w\} \mid v \in V(G_1), w \in V(G_2)\})$);

(O4) the substitution with respect to some pattern graph T (for a graph T with vertices $t_1, \dots, t_n$ and disjoint graphs $G_1, \dots, G_n$ the substitution of the vertices of T by the graphs $G_1, \dots, G_n$ is the graph with vertex set $\bigcup_{i=1}^n V(G_i)$ and edge set $\bigcup_{i=1}^n E(G_i) \cup \{\{u, v\} \mid u \in V(G_i), v \in V(G_j), \text{ and } \{t_i, t_j\} \in E(T)\}$).

We define the width of an algebraic expression A as the maximum number of operands used by any occurrence of the operation (O4) in A . The modular-width of a graph G , denoted $\text{mw}(G)$, can be defined as the least integer m such that G can be obtained from such an algebraic expression of width at most m .

Given a graph $G = (V, E)$, an algebraic expression of width $\text{mw}(G)$ describing G can be constructed in $O(|V| + |E|)$ time [15].

Definition 2 (Distance to cluster graph). For a graph $G = (V, E)$, a modulator to cluster graph is a vertex subset $X \subseteq V$ such that $G \setminus X$ is a vertex-disjoint union of cliques. The distance to cluster graph is the size of the smallest modulator to cluster graph.

A modulator to cluster graph of a graph with distance to cluster graph p can be found in $\mathcal{O}(1.9102^p \cdot (n + m))$ time [4].

Definition 3 (Distance to disjoint paths). For a graph $G = (V, E)$ a modulator to disjoint paths is a vertex subset $X \subseteq V$, such that $G \setminus X$ is a vertex-disjoint union of paths. The distance to disjoint paths is the size of the smallest modulator to disjoint paths.

For completeness, we include the following result which is rather folklore.

Lemma 1 ($\star^1$). The modulator to disjoint paths C of a graph G with distance to disjoint paths c can be found in $\mathcal{O}(4^c(n + m))$ time.

Definition 4 (Maximum leaf number). The maximum leaf number of a graph G is the maximum number of leaves in a spanning tree of G .

The presented algorithms rely on the following lemma.

Lemma 2. For any graph $G = (V, E)$, any set $S \subseteq V$, and any vertex $s \in V$, at most one permutation $\pi = (\pi_1, \dots, \pi_{|S|})$ of the vertices in S exists, such that there is a shortest path P with the following properties:

1. The first vertex on P is s ,
2. P contains all vertices from S , and
3. the vertices from S appear on P in exactly the order given by π .

Moreover, given a precomputed distance matrix for G , the permutation π can be found in $\mathcal{O}(|S| \log |S|)$ time.

¹ Proofs of lemmas marked with $(\star)$ can be found in the full version of the paper, preprint is available on ArXiv [13].

Proof. For the sake of deriving a contradiction, suppose that there are two different permutations π and π' satisfying the conditions, and let P, P' be the respective shortest paths. Let $i \in \{1, \dots, |S| - 1\}$ be the first position such that $\pi_i \neq \pi'_i$ and let $j \in \{2, \dots, |S|\}$ be the position of π_i in π' (clearly, $j > i$). Let P_1 be the subpath of P from s to π_i , and P'_2 the subpath of P' from π'_j to $\pi'_{|S|}$ (excluding the first vertex π'_j). Then, $P'' = P_1 \frown P'_2$ is a path which is strictly shorter than P' and has the same endpoints. This contradicts P' being a shortest path.

Sorting all vertices in S by increasing distance from the starting endpoint s yields our permutation π . It corresponds to some shortest path P if and only if

$$d_G(s, \pi_1) + \sum_{i=1}^{|S|-1} d_G(\pi_i, \pi_{i+1}) = d_G(s, \pi_{|S|}).$$

□

3 Parameterized Algorithms

In this section, we present several fpt algorithms for MESP. In Subsect. 3.1, we present an algorithm parameterized by the modular width. In Subsect. 3.2 we define the CONSTRAINED SET COVER (CSC) problem. In Subsect. 3.3 we show an fpt algorithm for MESP parameterized by the distance to cluster graph which reduces MESP to CSC. In Subsect. 3.4, we present an fpt algorithm parameterized by the distance to disjoint paths and the desired eccentricity, combined. This algorithm also depends on the solution of the CSC problem. Subsect. 3.5 presents an algorithm parameterized by the maximum leaf number.

3.1 Modular Width

We present an fpt algorithm for MESP parameterized by the modular width.

Let $G = (V, E)$ be a graph with modular width w and A be the corresponding algebraic expression describing the graph. We take a look at the last operation applied in A . Operation (O1) is trivial and (O2) yields a disconnected graph, therefore we suppose the last operation is either (O3) or (O4).

If it is (O3) and G is a path (of length at most 3), then the whole path is trivially a shortest path with eccentricity 0. If G is not a path, then the minimum eccentricity shortest path is any single edge connecting the two original graphs with eccentricity 1.

If it is (O4), the pattern graph is $T = (V_T, E_T)$ with $V_T = \{v_1, \dots, v_w\}$ and the substituted graphs are $G_1, \dots, G_w$, then we suppose that $w \geq 3$ and T is not a clique (otherwise (O3) could be used as the last operation). We continue by showing that the structure of the pattern graph restricts the structure of any shortest path in the resulting graph significantly.

Lemma 3 (★). *If the last operation in A is (O4), then there is a minimum eccentricity shortest path in G which contains at most one vertex from each G_i for $i \in \{1, \dots, w\}$.*

Now, we show that with respect to eccentricity, all vertices in the same graph G_i are equivalent. I.e., a minimum eccentricity shortest path in G can be found by trying all shortest paths in T .

Lemma 4 (*). *Let P be a shortest path in G and $p \in P \cap G_i$. We create a path P' by substituting p in P by any $p' \in G_i$. Then, $\text{ecc}_G(P') = \text{ecc}_G(P)$.*

Based on what we have shown, we can construct an algorithm to solve MESP. We handle separately the graphs created using (O1) or (O3) as the last operation. For (O4), we iterate through all possible shortest paths π in T . For each of them and each $i \in \{1, \dots, |\pi|\}$ we let $p_i \in G_{\pi_i}$ arbitrarily, and let $P := (p_1, \dots, p_{|\pi|})$. Then we check whether P is a shortest path with eccentricity at most k in G . By the above arguments, if there is a shortest path of eccentricity at most k , we will find one.

All shortest paths in a graph may be found by simply performing n DFS traversals (one starting in each vertex). Each forward step of the DFS represents a new shortest path; we skip edges that would break the shortest path property (this can easily be checked with a precomputed distance matrix).

By assuming a trivial upper bound 2^w on the number of shortest paths in T , we arrive at the following theorem.

Theorem 1. *There is an algorithm that solves MESP in $\mathcal{O}(2^w \cdot n^3)$ time, where w is the modular width of the input graph.*

3.2 Constrained Set Cover

In this subsection we define the CONSTRAINED SET COVER (CSC) problem. In the following two subsections it will be used as a subroutine to solve MESP.

CONSTRAINED SET COVER

Input: A set $\mathcal{C} = C_1 \cup \dots \cup C_m$ of candidates, a set $\mathcal{R} = \{r_1, \dots, r_n\}$ of requirements to be satisfied, and a function $\Psi : \mathcal{C} \rightarrow 2^{\mathcal{R}}$ that determines for each candidate which requirements it satisfies.

Question: Is there a *constrained set cover*, that is, a set of candidates, exactly one from each set $s_1 \in C_1, \dots, s_m \in C_m$ such that together they satisfy all the requirements, i.e., $\Psi(s_1) \cup \dots \cup \Psi(s_m) = \mathcal{R}$?

Each candidate can be thought of as a set of (satisfied) requirements. Hence, if we drop the constraints $s_i \in C_i$, we get the ordinary SET COVER. In our definition several candidates can satisfy the same set of requirements.

In the next two subsections we use the following lemma.

Lemma 5 (*). *CONSTRAINED SET COVER can be solved in $\mathcal{O}(2^{2^{|\mathcal{R}|}} |\mathcal{R}| \cdot |\mathcal{C}|)$ time.*

3.3 Distance to Cluster Graph

In this subsection, we present an fpt algorithm for MESP parametrized by the distance to cluster graph. The trivial case where distance to cluster graph is 0 is omitted. Note that if G is a graph with a modulator to cluster graph U , then, for any edge $\{u, v\}$ in $G \setminus U$, u and v have the same neighborhood in $G \setminus U$.

The high-level idea of the algorithm is that we iteratively guess (by trying all possible combinations), for each vertex in the modulator to cluster U , whether it lies on the desired shortest path (we say it belongs to the set L), or it is at distance 1 or 2 from the shortest path (it belongs to the set $\mathcal{R}_1$ or $\mathcal{R}_2$, respectively), or at an even further distance. Then, we try to find a shortest path such that all vertices from L lie on it, and all vertices in $\mathcal{R}_1, \mathcal{R}_2$ have the respective distance from the path. Finding such a path is reduced to solving the CSC problem presented in Subsect. 3.2. Once we guess the correct combination of these sets, we actually construct the MESP.

First, we discuss some properties of graphs having the desired path.

Lemma 6 ($\star$). *Let G be a graph with modulator to cluster graph U and let P be a shortest path with $\text{ecc}_G(P) = k$. Then, there exists a shortest path P' such that it contains at least one vertex from U and $\text{ecc}_G(P') \leq k$.*

Definition 5. *Let G be a graph with a modulator to cluster graph U . Let P be a shortest path in G with $\text{ecc}_G(P) \leq k$ and $U \cap P \neq \emptyset$. We denote $L^P = P \cap U$ and $\pi^P = (\pi_1^P, \dots, \pi_{|L^P|}^P)$ the permutation/order of vertices from L^P in which they appear on the path P . We denote $\mathcal{R}_i^P = \{u \in U \mid d_G(u, P) = i\}$ the set of vertices in U that are at distance i from P , for $i \in \{1, 2\}$.*

Let $V = V(G) \setminus U$. Since $G[V]$ is a disjoint union of cliques, and for every $i \in \{1, \dots, |L^P| - 1\}$, all vertices that are between π_i^P and π_{i+1}^P on P are from V , we have $d_G(\pi_i^P, \pi_{i+1}^P) \leq 3$, as otherwise P would not be a shortest path.

Let $\pi = (\pi_1, \dots, \pi_{|\pi|})$ be a candidate (guess) on the value of π^P . Intuitively, if we had the correct values of $\pi = \pi^P$, we would only need to select the (at most two) vertices between each π_i, π_{i+1} .

To help us refer to those pairs π_i, π_{i+1} between which we still need to choose some vertices we denote $\mathbf{h}_\pi = (h_1, \dots, h_\ell)$ the increasing sequence of all indices i such that $\{\pi_i, \pi_{i+1}\} \notin E$. For every $i \notin \mathbf{h}_\pi$, we have $\{\pi_i, \pi_{i+1}\} \in E(G)$ and, thus, there is no vertex between π_i and π_{i+1} on P . For every $h_i \in \mathbf{h}_\pi$: If $d_G(\pi_{h_i}, \pi_{h_i+1}) = 2$, then there is one vertex on P between π_{h_i} and π_{h_i+1} , and it is from V . If $d_G(\pi_{h_i}, \pi_{h_i+1}) = 3$, then there are two vertices from V on P between π_{h_i} and π_{h_i+1} .

Definition 6. *We define the set C_{h_i} of candidate vertices between π_{h_i} and π_{h_i+1} for each $h_i \in \mathbf{h}_\pi$.*

$$C_{h_i} = \begin{cases} \left\{ (u, v) \in V^2 \mid \{ \{ \pi_{h_i}, u \}, \{ u, \pi_{h_i+1} \} \} \subseteq E \right\} & \text{if } d_G(\pi_{h_i}, \pi_{h_i+1}) = 2 \\ \left\{ (u, v) \in V^2 \mid \{ \{ \pi_{h_i}, u \}, \{ u, v \}, \{ v, \pi_{h_i+1} \} \} \subseteq E \right\} & \text{if } d_G(\pi_{h_i}, \pi_{h_i+1}) = 3 \\ \emptyset & \text{otherwise} \end{cases}$$

For $h_i \in \mathbf{h}_\pi$ with $d_G(\pi_{h_i}, \pi_{h_{i+1}}) = 2$, the set C_{h_i} contains pairs of the same vertices (u, u) . To avoid adding some vertex into a path twice, we define a function μ which maps a pair of two elements to a sequence of length 1 or 2:

$$\mu(u, v) = \begin{cases} (u) & \text{if } u = v, \\ (u, v) & \text{if } u \neq v. \end{cases}$$

To solve MESP, we need to choose exactly one pair from each of $C_{h_1}, \dots, C_{h_\ell}$. Later, we show that the problem of choosing these pairs is an instance of CSC.

First, we define a function $\delta^P : U \cup V \rightarrow \mathbb{N}$ that will help us prove that the path constructed from the CSC solution will have a small eccentricity:

$$\delta^P(u) = \min \{d_G(u, L^P), d_G(u, \mathcal{R}_1^P) + 1, d_G(u, \mathcal{R}_2^P) + 2\}.$$

Lemma 7 ($\star$). *Function δ^P is a good estimate of the distance from P :*

1. $\delta^P(u) = d_G(u, P)$ for every $u \in U$, and
2. $\delta^P(u) = d_G(u, P \setminus (N_G[u] \cap V))$ for every $u \in V$.

Now we show how to choose optimal vertices from each C_i by solving CSC.

Lemma 8 ($\star$). *Suppose that P is a shortest path in G with $\text{ecc}_G(P) \leq k$, both endpoints of P are in U , and we have the corresponding values of $L^P, \pi^P, \mathcal{R}_1^P, \mathcal{R}_2^P$ as described in Definition 5. Let $\mathbf{h}_{\pi^P} = (h_1, \dots, h_\ell)$ and $(s_{h_1}, \dots, s_{h_\ell})$ be a solution of the CSC instance with requirements $\mathcal{R}^P = \mathcal{R}_1^P \cup \mathcal{R}_2^P$, sets of candidates $\mathcal{C} = C_{h_1} \cup \dots \cup C_{h_\ell}$, and function $\Psi(u, v) = N_G(u) \cup N_G(v) \cup \left((N_G^2[u] \cup N_G^2[v]) \cap \mathcal{R}_2^P \right)$. Then*

$$\begin{aligned} P' &= (\pi_1^P, \dots, \pi_{h_1}^P) \frown \mu(s_{h_1}) \frown (\pi_{h_1+1}^P, \dots, \pi_{h_2}^P) \\ &\cdots \frown \mu(s_{h_i}) \frown (\pi_{h_i+1}^P, \dots, \pi_{h_{i+1}}^P) \frown \mu(s_{h_{i+1}}) \\ &\cdots \frown (\pi_{h_{\ell-1}+1}^P, \dots, \pi_{h_\ell}^P) \frown \mu(s_{h_\ell}) \frown (\pi_{h_\ell+1}^P, \dots, \pi_{|L^P|}^P) \end{aligned}$$

is a shortest path and $\text{ecc}_G(P') \leq \max\{2, k\}$.

Clearly, if $k \geq 2$, then we can use Lemma 8 to construct a shortest path with eccentricity at most k . Now, we discuss the case when $k = 1$.

Observation 1 ($\star$). *If $\text{ecc}_G(P) = 1$, then for every $u \in U \cup V$ we have $\delta^P(u) \leq 2$.*

Corollary 1 ($\star$). *If $\text{ecc}_G(P) = 1$, then a path P' with $\text{ecc}_G(P') \leq 1$ can be constructed similarly as in Lemma 8 with the following modification. For each candidate set C_i which contains some pair $(x, y) \in V^2$ such that there is a neighbor $z \in V$ of x with $\delta^P(z) = 2$, remove every $(u, v) \in V^2$ such that z is not a neighbor of u from C_i .*

We have shown how to construct a shortest path with eccentricity at most k by solving the CSC problem, even if $k = 1$. Finally, we observe that such a path can be constructed even if one or both of its endpoints are in V .

Lemma 9 $(\star)$. *If P has an endpoint $s \in V$, its neighbor $t \in P$ might also be in V . Let $P = (s, t, \dots)$. We may obtain a path P' with $\text{ecc}_G(P') \leq k$ by removing $\Psi(s, s)$ (and $\Psi(t, t)$ if $t \in V$) from $\mathcal{R}^P$, finding $s_{h_1}, \dots, s_{h_\ell}$ by solving the CSC, and prepending s (and t if $t \in V$) to P' .*

MESP can be solved by trying all possible combinations of $(L, s, \mathcal{R}) : L \subseteq U, s \in L, \mathcal{R} = \mathcal{R}_1 \cup \mathcal{R}_2 \subseteq (U \setminus L)$. For each combination, do:

1. Find a permutation π of L , such that $\pi_1 = s$ and $\sum_{i=1}^{|L|-1} d_G(\pi_i, \pi_{i+1}) = d_G(\pi_1, \pi_{|L|})$. If it does not exist, continue with the next combination.
2. For each $h_i \in \mathbf{h}_\pi$: create set C_{h_i} according to Definition 6 and Corollary 1.
3. Solve the CSC instance as described in Lemma 8.
4. If the CSC instance has a solution, construct path P' as in Lemma 8.
5. Check if $\text{ecc}_G(P') \leq k$. If yes, return P' . If not, try the same after prepending and/or appending all combinations of single vertices and of pairs of vertices to P' (see Lemma 9).

Note that by Lemma 2, there is at most one such permutation π in step 1.

Theorem 2 $(\star)$. *In a graph with distance to cluster graph p , MESP can be solved in $\mathcal{O}(2^{4p} \cdot n^6)$ time.*

3.4 Distance to Disjoint Paths

In this subsection, we present an fpt algorithm for MESP parameterized by the distance to disjoint paths and the desired eccentricity, combined.

The high-level idea of the algorithm is similar to that in Subsect. 3.3. We iteratively guess (by trying all possible combinations), for each vertex in the modulator to disjoint paths C , what is the distance to the desired shortest path. Then, we try to find a shortest path which satisfies all the guessed distance requirements by solving an instance of the CSC problem. We argue that if these requirements are guessed correctly, the resulting path will indeed be the desired MESP.

We start by discussing some properties of graphs in which a shortest path $P = (p_1, \dots, p_{|P|})$ with $\text{ecc}_G(P) \leq k$ does exist. Assume that P is such a path, fixed for the next few lemmas and definitions.

Definition 7. Let $\widehat{C}^P = C \cup \{p_1, p_{|P|}\}$. Let $L^P = P \cap \widehat{C}^P$. We denote $\pi^P = (\pi_1^P, \dots, \pi_{|L^P|}^P)$ the permutation/order of vertices from L^P on the path P . We define function $\delta^P(v) = d_G(v, P)$ for every $v \in V$.

Let $\widehat{C}, L$ be candidates for $\widehat{C}^P, L^P$, respectively. Similarly as in Subsect. 3.3, the permutation $\pi = (\pi_1, \dots, \pi_{|L|})$ of the vertices in L is unique (if it exists), and

can be found in polynomial time. For each consecutive pair of vertices $\pi_i, \pi_{i+1} \in L$, there may be multiple shortest paths connecting them, such that they do not contain any other vertices from $\widehat{C}$. Exactly one of these shortest paths is contained in P for each pair. We say $\bar{\sigma}$ is a *candidate segment* if it is a sequence of vertices on some shortest path from π_i to π_{i+1} excluding the endpoints π_i, π_{i+1} and $\bar{\sigma} \cap \widehat{C} = \emptyset$. We define $\mathcal{S}(\pi_i, \pi_{i+1})$ as a set of all candidate segments $\bar{\sigma}$ between π_i and π_{i+1} . We denote $\widetilde{\mathcal{S}} = \bigcup_{i=1}^{|L|-1} \mathcal{S}(\pi_i, \pi_{i+1})$ the set of all candidate segments in G . We say that a candidate segment $\sigma \in \widetilde{\mathcal{S}}$ is a *necessary segment* if it must be part of any shortest path P' such that $\widehat{C} = \widehat{C}^{P'}$, $L = L^{P'}$, $\pi = \pi^{P'}$, and $\text{ecc}_G(P') \leq k$.

Intuitively, if we had the correct values of π , we would only need to select one segment out of each $\mathcal{S}(\pi_i, \pi_{i+1})$ for $i \in \{1, \dots, |L| - 1\}$, in order to construct the path P . To do so, we need the following function, which estimates the distance from a vertex to the path P .

Definition 8 (estimate distance to P). For a graph $G = (V, E)$, a set of vertices $\widehat{C} \subseteq V$ and a function $\delta : \widehat{C} \rightarrow \mathbb{N}$ we define $d_G^\delta : V \times 2^{\widehat{C}} \rightarrow \mathbb{N}$ as

$$d_G^\delta(v, S) = \min_{s \in S} d_G(v, s) + \delta(s).$$

Observation 2 ($\star$). If $\widehat{C} = \widehat{C}^P$ and $\delta = \delta^P|_{\widehat{C}}$ (that is, the restriction of δ^P to $\widehat{C}$) for some shortest path P in G , then for every $v \in V$ we have $d_G(v, P) \leq d_G^\delta(v, \widehat{C})$. In particular, if $d_G^\delta(v, \widehat{C}) \leq k$, then $d_G(v, P) \leq k$.

If we had the correct values for the permutation π of vertices from $\widehat{C}$ that are on P , we would still have to take care of those vertices $v \in V$ with $d_G^\delta(v, \widehat{C}) > k$, in order to solve MESP. In particular, we would have to choose a segment from each $\mathcal{S}(\pi_i, \pi_{i+1})$ in a way that for every vertex v with $d_G^\delta(v, \widehat{C}) > k$, there would be some chosen segment at distance at most k from v . We say that a candidate segment $\bar{\sigma} \in \widetilde{\mathcal{S}}$ satisfies $v \in V \setminus \widehat{C}$ if $d_G(v, \bar{\sigma}) \leq k < d_G^\delta(v, \widehat{C})$.

We continue by showing that the number of vertices v with $d_G^\delta(v, \widehat{C}) > k$ which do not lie on P is bounded by the size of L .

Lemma 10 ($\star$). Let $\bar{\sigma} \in \widetilde{\mathcal{S}}$ be a candidate segment and $D = \{v \in V \setminus P \mid \bar{\sigma} \text{ satisfies } v\}$. Then $|D| \leq 2$.

Corollary 2. Let P be a shortest path in G with $\text{ecc}_G(P) \leq k$. Let $U = \{v \in V \setminus P \mid d_G^{\delta^P}(v, \widehat{C}^P) > k\}$. There are $|L^P| - 1$ segments on P , therefore $|U| \leq 2(|L^P| - 1)$.

We have shown that there are not many vertices $v \notin P$ with $d_G^{\delta^P}(v, \widehat{C}^P) > k$. Now, we show that all such vertices actually have $d_G^{\delta^P}(v, \widehat{C}^P) = k + 1$.

Lemma 11 ($\star$). Let P be a shortest path in G with $\text{ecc}_G(P) \leq k$. Let $v \in V$ be such that $d_G^{\delta^P}(v, \widehat{C}^P) \geq k + 1$. Let $u \in P$ be the nearest vertex to v on P . Then either $u = v$, or $d_G(u, v) = k$ and $d_G(u, \widehat{C}^P) = 1$.

Fig. 2. Example of a situation from Lemma 11. Path P is red, candidate segments are blue. The segment containing vertex u satisfies v . If all the dashed parts are present in G , then the segment containing vertex w also satisfies v . (Color figure online)

Let $v \in V \setminus P$ be such that $d_G^{\delta^P}(v, \widehat{C}^P) = k+1$ and $S \subseteq \widetilde{\mathcal{S}}$ be a set of candidate segments that satisfy v . As shown in Fig. 2, there may be multiple such segments in S . However, if v itself lies on some segment, then only segments in the same $\mathcal{S}(\pi_i, \pi_{i+1})$ may satisfy v .

Observation 3 ($\star$). *Let P be a shortest path in G with $\text{ecc}_G(P) \leq k$. Let $\bar{\sigma} \in \mathcal{S}(\pi_i^P, \pi_{i+1}^P)$ be a candidate segment that contains some vertex u such that $d_G^{\delta^P}(u, \widehat{C}^P) = k+1$. Let $u' \in P$ be the nearest vertex to u on P . Then u' lies on a segment $\sigma \in \mathcal{S}(\pi_i^P, \pi_{i+1}^P)$.*

We already know from Lemma 11 that if a candidate segment contains some vertex v with $d_G^{\delta}(v, \widehat{C}) > k+1$, then it is a necessary segment. Now, we show another sufficient condition for a candidate segment to be a necessary segment.

Lemma 12 ($\star$). *Let $\bar{\sigma} \in \mathcal{S}(\pi_i, \pi_{i+1})$ be a candidate segment that contains some vertices $u, v \in \bar{\sigma}$ such that $u \neq v$ and $d_G^{\delta}(u, \widehat{C}) = d_G^{\delta}(v, \widehat{C}) = k+1$. Then, $\bar{\sigma}$ is a necessary segment.*

Let us summarize what we have shown so far. If we had the correct values for the permutation π of vertices from $\widehat{C}$ that are on P , and of δ , we would only need to select one segment out of each $\mathcal{S}(\pi_i, \pi_{i+1})$ to find a shortest path with eccentricity at most k . There are some vertices $u \in V$ such that $d_G^{\delta}(u, \widehat{C}) \leq k$ and for these vertices, the distance to the resulting path will be at most k , no matter which segments we choose.

A segment which contains some vertex v with $d_G^{\delta}(v, \widehat{C}) > k+1$ is a necessary segment. A segment which contains two vertices $u \neq v$ with $d_G^{\delta}(u, \widehat{C}) = d_G^{\delta}(v, \widehat{C}) = k+1$ is a necessary segment as well. For the remaining segments, we know that for every $u \in V$ with $d_G^{\delta}(u, \widehat{C}) > k$, the shortest path with eccentricity at most k needs to contain some $\sigma_u \in \widetilde{\mathcal{S}}$ such that $d_G(u, \sigma_u) \leq k$. Furthermore,

for every $v \in \widehat{C}$ with $d_G(v, L) > \delta(v)$, the path needs to contain some $\sigma_v \in \widetilde{\mathcal{S}}$ such that $d_G(v, \sigma_v) \leq \delta(v)$.

Clearly, the problem of selecting one segment out of each set of candidate segments is an instance of CSC: the sets of candidates are $\mathcal{C} = \mathcal{S}(\pi_1, \pi_2) \cup \dots \cup \mathcal{S}(\pi_{|L|-1}, \pi_{|L|})$, the requirements are $\mathcal{R} = \{v \in V \setminus \widehat{C} \mid d_G^\delta(v, \widehat{C}) > k\} \cup \{v \in \widehat{C} \setminus L \mid d_G(v, L) > \delta(v)\}$, and the function $\Psi(\bar{\sigma}) = \{v \in V \setminus \widehat{C} \mid \bar{\sigma} \text{ satisfies } v\} \cup \{v \in \widehat{C} \setminus L \mid d_G(v, \sigma) \leq \delta(v)\}$.

We know that the number of vertices outside of P that the segments can satisfy is bounded by the size of L . Furthermore, we know that if a segment contains at least two vertices that need to be satisfied, then it is a necessary segment. Lastly, we know that if a segment from some $\mathcal{S}(\pi_i, \pi_{i+1})$ contains one vertex v with $d_G^\delta(v, \widehat{C}) = k + 1$, then only segments from the same $\mathcal{S}(\pi_i, \pi_{i+1})$ may satisfy v . Thus, all segments in $\mathcal{S}(\pi_i, \pi_{i+1})$ that do not satisfy v may be disregarded. By this, we ensure that v will be satisfied no matter which segment is chosen, and v does not need to be added to the requirements $\mathcal{R}$. Hence, the requirements $\mathcal{R}$ do not need to contain any vertices from P , and the size of $\mathcal{R}$ is bounded by the size of C .

In the following lemma, we show that we do not need to explicitly check whether a segment contains some vertex v with $d_G^\delta(v, \widehat{C}) > k + 1$ to decide that it is a necessary segment. This will simplify our algorithm a bit.

Lemma 13 ($\star$). *If a segment $\sigma \in \widetilde{\mathcal{S}}$ contains a vertex u such that $d_G^\delta(u, \widehat{C}) > k + 1$, then it must also contain two vertices v, v' with $d_G^\delta(v, \widehat{C}) = d_G^\delta(v', \widehat{C}) = k + 1$.*

Finally, we propose an algorithm that solves MESP. It finds the correct values for $p_1, \widehat{C}, L$, and $\delta: \widehat{C} \rightarrow \{0, \dots, k\}$ by trying all possible combinations. For each combination, it performs the following steps.

1. Find a permutation π of L , such that $\pi_1 = p_1$ and $\sum_{i=1}^{|L|-1} d_G(\pi_i, \pi_{i+1}) = d_G(\pi_1, \pi_{|L|})$. If it does not exist, continue with the next combination.
2. For each π_i, π_{i+1} , check all candidate segments in $\mathcal{S}(\pi_i, \pi_{i+1})$.
 - (a) If there are any segments containing a vertex u with $d_G^\delta(u, \widehat{C}) = k + 1$, then we may disregard all candidate segments which do not satisfy u .
 - (b) After disregarding these segments, if there is only one candidate segment left, it is a necessary segment. If there is no candidate segment left, then no solution exists.
3. If there is a vertex v such that $d_G^\delta(v, \widehat{C}) > k + 1$, and it does not lie on a segment that we have marked as a necessary segment, then no solution exists.
4. Construct the set U of vertices v that are not contained in any segment and have $d_G^\delta(v, \widehat{C}) = k + 1$. If $|U| > 2(|L| - 1)$, then no solution exists.
5. Choose the rest of the segments from all candidate segments (except those disregarded in step 1) by solving the CSC instance, with requirements $u \in \widehat{C} \setminus L$ whose distance to the parts of P selected so far is greater than $\delta(u)$, and all of U .

6. If the CSC instance has a solution, construct a path from π and from the chosen candidate segments. If the resulting path has eccentricity at most k , return it.

Note that by Lemma 2, there is at most one such permutation π in step 1. We arrive at the following theorem.

Theorem 3 ($\star$). *For a graph with distance to disjoint paths c , MESP can be solved in $\mathcal{O}(2^{5c} k^c \cdot n^4)$.*

3.5 Maximum Leaf Number

Theorem 4 ($\star$). *There is an algorithm that solves MESP in $\mathcal{O}(2^\ell \cdot n^3)$ time, where ℓ is the maximum leaf number of the input graph.*

4 Future Directions

We have shown that MESP is fixed-parameter tractable with respect to several structural parameters. This partially answers an open question of Dragan and Leitert [8] on classes where the problem is polynomial time solvable, as this is the case whenever we limit ourselves to a class where one of the studied parameters is a constant.

The natural next steps in the research of parameterized complexity of MESP would be to investigate the existence of fpt algorithms with respect to the distance to disjoint paths alone, and with respect to other structural parameters, such as tree depth or feedback vertex set number (see Fig. 1).

References

1. Birmelé, É., de Montgolfier, F., Planche, L.: Minimum eccentricity shortest path problem: an approximation algorithm and relation with the k -laminarity problem. In: Chan, T.-H.H., Li, M., Wang, L. (eds.) COCOA 2016. LNCS, vol. 10043, pp. 216–229. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-48749-6_16
2. Birmelé, E., de Montgolfier, F., Planche, L., Viennot, L.: Decomposing a graph into shortest paths with bounded eccentricity. In: ISAAC 2017. LIPIcs, vol. 92, pp. 15:1–15:13. Dagstuhl (2017)
3. Birmelé, E., de Montgolfier, F., Planche, L., Viennot, L.: Decomposing a graph into shortest paths with bounded eccentricity. Discrete Appl. Math. **284**, 353–374 (2020)
4. Boral, A., Cygan, M., Kociumaka, T., Pilipczuk, M.: A fast branching algorithm for cluster vertex deletion. Theory Comput. Syst. **58**(2), 357–376 (2016)
5. Cygan, M., et al.: Parameterized Algorithms. Springer, Cham (2015)
6. Diestel, R.: Graph Theory, 5th Edition, Graduate texts in mathematics, vol. 173. Springer, Cham (2016). <https://doi.org/10.1007/978-3-319-31940-7>
7. Dragan, F.F., Leitert, A.: On the minimum eccentricity shortest path problem. In: Dehne, F., Sack, J.-R., Stege, U. (eds.) WADS 2015. LNCS, vol. 9214, pp. 276–288. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-21840-3_23

8. Dragan, F.F., Leitert, A.: Minimum eccentricity shortest paths in some structured graph classes. *J. Graph Algorithms Appl.* **20**(2), 299–322 (2016)
9. Dragan, F.F., Leitert, A.: On the minimum eccentricity shortest path problem. *Theor. Comput. Sci.* **694**, 66–78 (2017)
10. Gajarský, J., Lampis, M., Ordyniak, S.: Parameterized algorithms for modular-width. In: Gutin, G., Szeider, S. (eds.) *IPEC 2013. LNCS*, vol. 8246, pp. 163–176. Springer, Cham (2013). https://doi.org/10.1007/978-3-319-03898-8_15
11. Indyk, P.: Algorithmic applications of low-distortion geometric embeddings. In: *FOCS 2001*, pp. 10–33. IEEE Computer Society (2001)
12. Indyk, P., Matousek, J.: Low-distortion embeddings of finite metric spaces. In: *Handbook of Discrete and Computational Geometry*, pp. 177–196. CRC Press, Boca Raton (2004)
13. Kučera, M., Suchý, O.: Minimum eccentricity shortest path problem with respect to structural parameters (2021). <https://arxiv.org/abs/2008.07898>
14. Sorge, M., Weller, M.: The graph parameter hierarchy (2016). <https://manyu.pro/assets/parameter-hierarchy.pdf>
15. Tedder, M., Corneil, D., Habib, M., Paul, C.: Simpler linear-time modular decomposition via recursive factorizing permutations. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfssdóttir, A., Walukiewicz, I. (eds.) *ICALP 2008. LNCS*, vol. 5125, pp. 634–645. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-70575-8_52
16. Tenenbaum, J.B., Silva, V.D., Langford, J.C.: A global geometric framework for nonlinear dimensionality reduction. *Science* **290**(5500), 2319–2323 (2000)
17. Völk, F., Baptiste, E., Habib, M., Lopez, P., Vigliotti, C.: Read networks and k-laminar graphs. *CoRR* abs/1603.01179 (2016)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.