On the Minimum Eccentricity Shortest Path Problem

Feodor F. Dragan and Arne Leitert(✉)

Department of Computer Science, Kent State University, Kent, Ohio, USA
{dragan,aleitert}@cs.kent.edu

Abstract. In this paper, we introduce and investigate the *Minimum Eccentricity Shortest Path (MESP)* problem in unweighted graphs. It asks for a given graph to find a shortest path with minimum eccentricity. We demonstrate that:

– a minimum eccentricity shortest path plays a crucial role in obtaining the best to date approximation algorithm for a minimum distortion embedding of a graph into the line;
– the MESP-problem is NP-hard on general graphs;
– a 2-approximation, a 3-approximation, and an 8-approximation for the MESP-problem can be computed in $\mathcal{O}(n^3)$ time, in $\mathcal{O}(nm)$ time, and in linear time, respectively;
– a shortest path of minimum eccentricity k in general graphs can be computed in $\mathcal{O}(n^{2k+2}m)$ time;
– the MESP-problem can be solved in linear time for trees.

1 Introduction

All graphs occurring in this paper are connected, finite, unweighted, undirected, loopless and without multiple edges. For a graph $G = (V, E)$, we use $n = |V|$ and $m = |E|$ to denote the cardinality of the vertex set and the edge set of G. For a vertex v of G, $N_G(v) = \{u \in V \mid uv \in E\}$ is called the *open neighborhood*, and $N_G[v] = N_G(v) \cup \{v\}$ the *closed neighborhood* of v.

The *length* of a path from a vertex v to a vertex u is the number of edges in the path. The *distance* $d_G(u, v)$ of two vertices u and v is the length of a shortest path connecting u and v. The distance between a vertex v and a set $S \subseteq V$ is defined as $d_G(v, S) = \min_{u \in S} d_G(u, v)$. The *eccentricity* $\mathrm{ecc}_G(v)$ of a vertex v is $\max_{u \in V} d_G(u, v)$. For a set $S \subseteq V$, its eccentricity is $\mathrm{ecc}_G(S) = \max_{u \in V} d_G(u, S)$.

In this paper, we investigate the following problem.

Definition 1 (Minimum Eccentricity Shortest Path Problem). *For a given a graph G, find a shortest path P such that for each shortest path Q, $\mathrm{ecc}_G(P) \le \mathrm{ecc}_G(Q)$.*

F. Dehne et al. (Eds.): WADS 2015, LNCS 9214, pp. 276–288, 2015.
DOI: 10.1007/978-3-319-21840-3_23

Although this problem might be of an independent interest (it may arise in determining a "most accessible" speedy linear route in a network and can find applications in communication networks, transportation planning, water resource management and fluid transportation), our interest in this problem stems from the role it plays in obtaining the best to date approximation algorithm for a minimum distortion embedding of a graph into the line. In Section 2, we demonstrate that every graph G with a shortest path of eccentricity k admits an embedding f of G into the line with distortion at most $(8k+2)\,\mathrm{ld}(G)$, where $\mathrm{ld}(G)$ is the minimum line-distortion of G. Furthermore, if a shortest path of G of eccentricity k is given in advance, then such an embedding f can be found in linear time.

This fact augments the importance of investigating the *Minimum Eccentricity Shortest Path* problem (MESP-problem) in graphs. Fast algorithms for it will imply fast approximation algorithms for the minimum line distortion problem. Existence of low eccentricity shortest paths in special graph classes will imply low approximation bounds for those classes. For example, all AT-free graphs (and hence all interval, permutation, cocomparability graphs) enjoy a shortest path of eccentricity at most 1 [3], all convex bipartite graphs enjoy a shortest path of eccentricity at most 2 [5].

We prove also that for every graph G with $\mathrm{ld}(G) = \lambda$, the minimum eccentricity of a shortest path of G is at most $\lfloor \frac{\lambda}{2} \rfloor$. Hence, one gets an efficient embedding of G into the line with distortion at most $\mathcal{O}(\lambda^2)$.

In Section 3, we show that the MESP-problem is NP-hard on general graphs and that a shortest path of minimum eccentricity k in general graphs, can be computed in $\mathcal{O}(n^{2k+2}m)$ time. In Section 4, we design for the MESP-problem on general graphs a 2-approximation algorithm that runs in $\mathcal{O}(n^3)$ time, a 3-approximation algorithm that runs in $\mathcal{O}(nm)$ time and an 8-approximation algorithm that runs in linear time. In Section 5, we demonstrate that the MESP-problem can be solved in linear time for trees and distance-hereditary graphs, and in polynomial time for chordal graphs and dually chordal graphs.

Note that our Minimum Eccentricity Shortest Path problem is close but different from the *Central Path* problem in graphs introduced in [16]. It asks for a given graph G to find a path P (not necessarily shortest) such that any other path of G has eccentricity at least $\mathrm{ecc}_G(P)$. The Central Path problem generalizes the Hamiltonian Path problem and therefore is NP-hard even for chordal graphs [15]. Our problem is polynomial time solvable for chordal graphs.

In what follows we will need the following additional notions and notations.

The *diameter* of a graph G is $\mathrm{diam}(G) = \max_{u,v\in V} d_G(u,v)$. The diameter $\mathrm{diam}_G(S)$ of a set $S \subseteq V$ is defined as $\max_{u,v\in S} d_G(u,v)$. A pair of vertices x, y of G is called a *diametral pair* if $d_G(u,v) = \mathrm{diam}(G)$. In this case, every shortest path connecting x and y is called a *diametral path*.

A path P of a graph G is called a *k-dominating path* of G if $\mathrm{ecc}_G(P) \leq k$. In this case, we say also that P *k-dominates* each vertex of G. A pair of vertices x, y of G is called a *k-dominating pair* if every path connecting x and y has eccentricity at most k.

For a vertex s, let $L_i^{(s)} = \{v \mid d_G(s,v) = i\}$ denote the vertices with distance i from s. We will also refer to $L_i^{(s)}$ as the i-th layer.

2 Motivation Through the Line-Distortion of a Graph

Computing a minimum distortion embedding of a given n-vertex graph G into the line ℓ was recently identified as a fundamental algorithmic problem with important applications in various areas of computer science, like computer vision [17], as well as in computational chemistry and biology (see [12,13]). The *minimum line distortion* problem asks, for a given graph $G = (V, E)$, to find a mapping f of vertices V of G into points of ℓ with minimum number λ such that $d_G(x,y) \leq |f(x) - f(y)| \leq \lambda\, d_G(x,y)$ for every $x, y \in V$. The parameter λ is called the *minimum line-distortion* of G and denoted by $\mathrm{ld}(G)$. The embedding f is called *non-contractive* since $d_G(x,y) \leq |f(x) - f(y)|$ for every $x, y \in V$.

In [2], Bădoiu et al. showed that this problem is hard to approximate within a constant factor. They gave an exponential-time exact algorithm and a polynomial-time $\mathcal{O}(n^{1/2})$-approximation algorithm for arbitrary unweighted input graphs, along with a polynomial-time $\mathcal{O}(n^{1/3})$-approximation algorithm for unweighted trees. In another paper [1], Bădoiu et al. showed that the problem is hard to approximate by a factor $\mathcal{O}(n^{1/12})$, even for weighted trees. They also gave a better polynomial-time approximation algorithm for general weighted graphs, along with a polynomial-time algorithm that approximates the minimum line-distortion λ embedding of a weighted tree by a factor that is polynomial in λ.

Fast exponential-time exact algorithms for computing the line-distortion of a graph were proposed in [7,8]. Fomin et al. [8] showed that a minimum distortion embedding of an unweighted graph into the line can be found in time $5^{n+o(n)}$. Fellows et al. [7] gave an $\mathcal{O}(n\lambda^4(2\lambda+1)^{2\lambda})$ time algorithm that for an unweighted graph G and integer λ either constructs an embedding of G into the line with distortion at most λ, or concludes that no such embedding exists. They extended their approach also to weighted graphs obtaining an $\mathcal{O}(n\lambda^{4W}(2\lambda+1)^{2\lambda W})$ time algorithm, where W is the largest edge weight. Thus, the problem of minimum distortion embedding of a given n-vertex graph G into the line ℓ is Fixed Parameter Tractable.

Heggernes et al. [10,11] initiated the study of minimum distortion embeddings into the line of specific graph classes. In particular, they gave polynomial-time algorithms for the problem on bipartite permutation graphs and on threshold graphs [11]. Furthermore, in [10], Heggernes et al. showed that the problem of computing a minimum distortion embedding of a given graph into the line remains NP-hard even when the input graph is restricted to a bipartite, cobipartite, or split graph, implying that it is NP-hard also on chordal, cocomparability, and AT-free graphs. They also gave polynomial-time constant-factor approximation algorithms for split and cocomparability graphs.

Recently, in [5], a more general result for unweighted graphs was proven: for every class of graphs with path-length bounded by a constant, there exists an efficient constant-factor approximation algorithm for the minimum line-distortion problem. As a byproduct, an efficient algorithm was obtained which for each unweighted graph G with $\mathrm{ld}(G) = \lambda$ constructs an embedding with distortion at most $\mathcal{O}(\lambda^2)$. Furthermore, for AT-free graphs, a linear time 8-approximation algorithm for the minimum line-distortion problem was obtained. Note that AT-free graphs contain all cocomparability graphs and hence all interval, permutation and trapezoid graphs.

In this section we simplify and improve the result of [5]. We show that a minimum eccentricity shortest path plays a crucial role in obtaining the best to date approximation algorithm for the minimum line-distortion problem.

We will need the following simple "local density" lemma.

Lemma 1. *For every vertex set $S \subseteq V$ of an arbitrary graph $G = (V, E)$,*

$$|S| - 1 \leq \mathrm{diam}_G(S)\,\mathrm{ld}(G).$$

Proof. Consider an embedding f^* of G into the line ℓ with distortion $\mathrm{ld}(G)$. Let a and b be the leftmost and the rightmost, respectively, in ℓ vertices of S, i. e., $f^*(a) = \min\{f^*(v) \mid v \in S\}$ and $f^*(b) = \max\{f^*(v) \mid v \in S\}$. Consider a shortest path P in G between a and b. Since for each edge xy of G (and hence of P) $|f^*(x) - f^*(y)| \leq \mathrm{ld}(G)$ holds, we get $f^*(b) - f^*(a) \leq d_G(a, b)\,\mathrm{ld}(G) \leq \mathrm{diam}_G(S)\,\mathrm{ld}(G)$. On the other hand, since all vertices of S are mapped to points of ℓ between $f^*(a)$ and $f^*(b)$, we have $f^*(b) - f^*(a) \geq |S| - 1$. □

The main result of this section is the following.

Theorem 1. *Every graph G with a shortest path of eccentricity k admits an embedding f of G into the line with distortion at most $(8k+2)\,\mathrm{ld}(G)$. If a shortest path of G of eccentricity k is given in advance, then such an embedding f can be found in linear time.*

Proof. Let $P = (x_0, x_1, \ldots, x_i, \ldots, x_j, \ldots, x_q)$ be a shortest path of G of eccentricity k. Build a $BFS(P, G)$-tree T of G (i. e., a Breadth-First-Search tree of G started at path P). Denote by $\{X_0, X_1, \ldots, X_q\}$ the decomposition of the vertex set V of G obtained from T by removing the edges of P. That is, X_i is the vertex set of a subtree (branch) of T growing from vertex x_i of P. See Fig. 1(a) for an illustration. Since eccentricity of P is k, we have $d_G(v, x_i) \leq k$ for every $i \in \{1, \ldots, q\}$ and every $v \in X_i$.

We define an embedding f of G into the line ℓ by performing a preorder traversal of the vertices of T starting at vertex x_0 and visiting first vertices of X_i and then vertices of X_{i+1}, $i = 0, \ldots, q-1$. We place vertices of G on the line in that order, and also, for each $i \in \{0, \ldots, q-1\}$, we leave a space of length $d_T(v_i, v_{i+1})$ between any two vertices v_i and v_{i+1} placed next to each other (this can be done during the preorder traversal). Alternatively, f can be defined by creating a twice around tour of the tree T, which visits vertices of X_i prior to vertices of X_{i+1}, $i = 0, \ldots, q-1$, and then returns to x_0 from x_q along edges of P. Following vertices of T from x_0 to x_q as shown in Fig. 1(b) (i. e., using upper part of the twice around tour), $f(v)$ can be defined as the first appearance of vertex v in that subtour (see Fig. 1(c)).

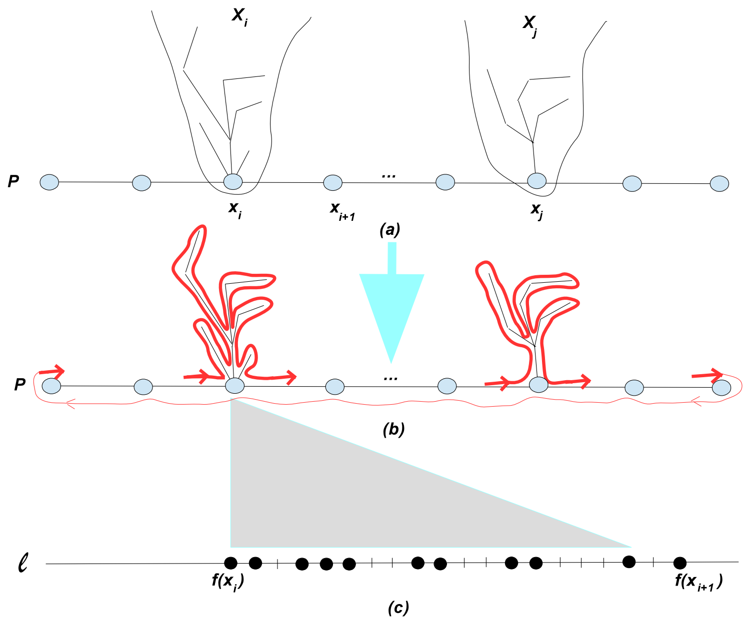

Fig. 1. Illustration to the proof of Theorem 1. (a) The decomposition $\{X_0, X_1, \ldots, X_q\}$ of the vertex set V of G. (b) The upper part of the twice around tour. (c) An embedding f obtained from following the upper part of the twice around tour.

We claim that f is a (non-contractive) embedding with distortion at most $(8k+2)\,\mathrm{ld}(G)$. It is sufficient to show that $d_G(x,y) \leq |f(x)-f(y)|$ for every two vertices of G that are placed by f next to each other in ℓ and that $|f(v)-f(u)| \leq (8k+2)\,\mathrm{ld}(G)$ for every edge uv of G (see, e. g., [2,11]).

Let x, y be arbitrary two vertices of G that are placed by f next to each other in ℓ. By construction, we know that $|f(x)-f(y)| = d_T(x,y)$. Since $d_G(x,y) \leq d_T(x,y)$, we get also $d_G(x,y) \leq |f(x)-f(y)|$, i. e., f is non-contractive.

Consider now an arbitrary edge uv of G and assume $u \in X_i$ and $v \in X_j$ $(i \leq j)$. Note that $d_P(x_i,x_j) = j-i \leq 2k+1$, since P is a shortest path of G and $d_P(x_i,x_j) = d_G(x_i,x_j) \leq d_G(x_i,u)+1+d_G(x_j,v) \leq 2k+1$. Set $S = \bigcup_{h=i}^{j} X_h$. For any two vertices $x, y \in S$, $d_G(x,y) \leq d_G(x,P)+2k+1+d_G(y,P) \leq k+2k+1+k = 4k+1$ holds. Hence, $\mathrm{diam}_G(S) \leq 4k+1$. Consider subtree T_S of T induced by S. Clearly, T_S is connected and has $|S|-1$ edges. Therefore, $f(v)-f(u) \leq 2(|S|-1)$ since each edge of T_S contributes to $f(v)-f(u)$ at most 2 units. Now, by Lemma 1, $f(v)-f(u) \leq 2(|S|-1) \leq 2\,\mathrm{diam}_G(S)\,\mathrm{ld}(G) \leq (8k+2)\,\mathrm{ld}(G)$. □

Recall that a pair x, y of vertices of a graph G forms a k-dominating pair if every path connecting x and y in G has eccentricity at most k. It turns out that the following result is true.

Proposition 1. *If the minimum line-distortion of a graph G is λ, then G has a $\lfloor \frac{\lambda}{2} \rfloor$-dominating pair.*

Proof. Let f be an optimal line embedding for G. This embedding has a first vertex v_1 and a last vertex v_n. Let u be an arbitrary vertex and P an arbitrary path from v_1 to v_n. If u is not on this path, there is an edge $v_i v_j$ of P with $f(v_i) < f(u) < f(v_j)$. Without loss of generality, we can say that $f(u) - f(v_i) \leq \lfloor (f(v_j) - f(v_i))/2 \rfloor \leq \lfloor \frac{\lambda}{2} \rfloor$. Thus, each vertex is $\lfloor \frac{\lambda}{2} \rfloor$-dominated by each path from v_1 to v_n, i. e., v_1, v_n is a $\lfloor \frac{\lambda}{2} \rfloor$-dominating pair. □

Corollary 1. *For every graph G with $\mathrm{ld}(G) = \lambda$, the minimum eccentricity of a shortest path of G is at most $\lfloor \frac{\lambda}{2} \rfloor$.*

Theorem 1 and Corollary 1 stress the importance of investigating the *Minimum Eccentricity Shortest Path* problem (MESP-problem) in graphs. As we will show later, although the MESP-problem is NP-hard on general graphs, there are much better (than for the minimum line distortion problem) approximation algorithms for it. We design for the MESP-problem on general graphs a 2-approximation algorithm that runs in $\mathcal{O}(n^3)$ time, a 3-approximation algorithm that runs in $\mathcal{O}(nm)$ time and an 8-approximation algorithm that runs in linear time.

Combining Theorem 1 and Corollary 1 with those approximation results, we reproduce a result of [2] and [5].

Corollary 2 ([2,5]). *For every graph G with $\mathrm{ld}(G) = \lambda$, an embedding into the line with distortion at most $\mathcal{O}(\lambda^2)$ can be found in polynomial time.*

It should be noted that, since the difference between the minimum eccentricity of a shortest path and the line-distortion of a graph can be very large (close to n), the result in Theorem 1 seems to be stronger. Furthermore, one version of our algorithm (that uses an 8-approximation algorithm for the MESP-problem) runs in total linear time.

3 NP-Completeness Result

In this section, we will show that in general it is NP-complete to find a minimum eccentricity shortest path. For this, we define the decision version of this problem (k-ESP) as follows: Given a graph G and an integer k, does G contain a shortest path P with eccentricity at most k?

Theorem 2. *The decision version of the minimum eccentricity shortest path problem is NP-complete.*

Proof. We will proof this by reducing SAT to k-ESP.

Let $\mathcal{I}$ be an instance of SAT with the variables $\mathcal{P} = \{p_1, \ldots, p_n\}$ and the clauses $\mathcal{C} = \{c_1, \ldots, c_m\}$. We assume $\mathcal{I}$ is a formula given in CNF. Also, let $k = \max\{n, m\}$. We create a graph G as shown in Figure 2. For each variable p_i

create two vertices, one representing p_i and one representing $\neg p_i$. Create one vertex c_i for every clause c_i. Additionally, create two vertices u_0, u_n and, for each i with $0 \leq i \leq n$, a vertex v_i.

Connect each variable vertex p_i and $\neg p_i$ with v_{i-1} and v_i directly with an edge. Connect each clause with the variables containing it with a path of length k. Also connect v_0 with u_0 and v_n with u_n with a path of length k.

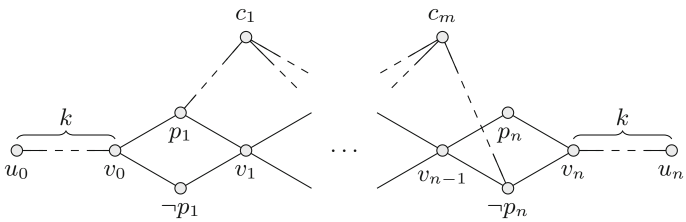

Fig. 2. Reduction from SAT to k-DSP. Illustration to the proof of Theorem 2.

Note that every shortest path in G not containing v_0 and v_n has an eccentricity larger than k. Also, a shortest path from v_0 to v_n has length $2n$ ($d_G(v_{i-1}, v_i) = 2$, passing p_i or $\neg p_i$). Since $k \geq n$, no shortest path from v_0 to v_n is passing a vertex c_i; in this case the minimal length would be $2k + 2$. Additionally, note that for all vertices in G except the vertices which represent clauses, the distance to a vertex v_i with $0 \leq i \leq n$ is at most k.

We will now show that $\mathcal{I}$ is satisfiable if and only if G has a shortest path with eccentricity k.

First assume $\mathcal{I}$ is satisfiable. Let $f \colon \mathcal{P} \to \{T, F\}$ be a satisfying assignment for the variables. As shortest path P we choose a shortest path from v_0 to v_n. Thus, we have to chose between p_i and $\neg p_i$. We will chose p_i if and only if $f(p_i) = T$. Because $\mathcal{I}$ is satisfiable, there is a p_i for each c_j such that either $f(p_i) = T$ and $d_G(c_j, p_i) = k$, or $f(p_i) = F$ and $d_G(c_j, \neg p_i) = k$. Thus, P has eccentricity k.

Next consider a shortest path P in G of eccentricity k. As mentioned above, P contains either p_i or $\neg p_i$. Now we define $f \colon \mathcal{P} \to \{T, F\}$ as follows:

$$f(p_i) = \begin{cases} T & \text{if } p_i \in P, \\ F & \text{else, i.e. } \neg p_i \in P. \end{cases}$$

Because P has eccentricity k and only vertices representing a variable in the clause c_j are at distance k to vertex c_j, f is a satisfying assignment for $\mathcal{I}$. □

V. B. Le[1] pointed out that, by slightly modifying the created graph, it can be shown that the problem remains NP-complete even if the graph has a bounded vertex-degree of 3.

[1] University of Rostock, Germany.

Note that the factor k in this reduction depends on the input size. In [14] it was already mentioned that for $k = 1$ the problem can be solved in $\mathcal{O}(n^3 m)$ time by modifying an algorithm given in [4]. There, the problem was called *Dominating Shortest Path* problem. In the full version of this paper, we show that the k-ESP problem can be solved in $\mathcal{O}(n^{2k+2}m)$ time for every fixed $k \geq 0$.

We can slightly modify the MESP problem such that a start vertex s and an end vertex t of the path are given. This is, for a given a graph G and two vertices s and t, find a shortest (s,t)-path P such that for each shortest (s,t)-path Q, $\mathrm{ecc}_G(P) \leq \mathrm{ecc}_G(Q)$. We call this the (s,t)-MESP problem. From the reduction above, it follows that the decision version of this problem is NP-complete, too.

Corollary 3. *The decision version of the (s,t)-MESP problem is NP-complete.*

4 Approximation Algorithms

In this section we will present different approximation algorithms. The algorithms differ in their approximation factor and runtime. Base for them are the following two lemmas.

Lemma 2. *In a graph G, let P be a shortest path from s to t of eccentricity at most k. For each layer $L_i^{(s)}$ there is a vertex $p_i \in P$ such that the distance from p_i to each vertex $v \in L_i^{(s)}$ is at most $2k$. Additionally, $p_i \in L_i^{(s)}$ if $i \leq d_G(s,t)$, and $p_i = t$ if $i \geq d_G(s,t)$.*

Proof. For each vertex v, let $p(v) \in P$ be a vertex with $d_G(p(v), v) \leq k$.

For each $i \leq d_G(s,t)$, let $p_i \in P \cap L_i^{(s)}$ be the vertex in P with distance i to s. For an arbitrary vertex $v \in L_i^{(s)}$, let $j = d_G(s, p(v))$. Because $\mathrm{ecc}_G(P) \leq k$ and P is a shortest path, $|i - j| \leq k$. Thus, $d_G(p_i, v) \leq d_G(p_i, p(v)) + d_G(p(v), v) \leq 2k$.

Let $L' = \{v \mid d_G(s,v) \geq d_G(s,t)\}$. Because P has eccentricity at most k, $d_G(p,t) \leq k$ for all $p \in \{p(v) \mid v \in L'\}$. Therefore, $d_G(t,v) \leq 2k$ for all $v \in L'$. □

Lemma 3. *If G has a shortest path of eccentricity at most k from s to t, then every path Q with $s \in Q$ and $d_G(s,t) \leq \max_{v \in Q} d_G(s,v)$ has eccentricity at most $3k$.*

Proof. Let P be a shortest path from s to t with $\mathrm{ecc}_G(P) \leq k$ and Q an arbitrary path with $s \in Q$ and $d_G(s,t) \leq \max_{v \in Q} d_G(s,v)$. Without loss of generality, we can assume that Q starts at s. Also let u be an arbitrary vertex. Since $\mathrm{ecc}_G(P) \leq k$, there is a vertex $p \in P$ with $d_G(u,p) \leq k$. Because $d_G(s,t) \leq \max_{v \in Q} d_G(s,v)$, there is a vertex $q \in Q$ with $d_G(s,p) = d_G(s,q)$. By Lemma 2, the distance between p and q is at most $2k$. Thus, the distance from q to u is at most $3k$. □

Corollary 4. *For a given graph G and two vertices s and t, each shortest (s,t)-path is a 3-approximation for the (s,t)-MESP problem.*

Theorem 3. *Algorithm 1 calculates a 3-approximation for the MESP problem in $\mathcal{O}(nm)$ time.*

Algorithm 1. A 3-approximation for the MESP problem.

Input: A graph $G = (V, E)$.
Output: A shortest path with eccentricity at most $3k$, where k is the minimum eccentricity of all paths in G.

```
1 foreach s ∈ V do
2     Find a vertex v for which the distance to s is maximal. Also find a shortest
      path P(s) from s to v.
3     Calculate k(s) = ecc_G(P(s)).
4 Among all computed paths P(s), select one for which k(s) is minimal.
```

Proof. Assume a given graph G has a shortest path P from s to t with $\text{ecc}_G(P) = k$ and s is the vertex selected by the loop in line 1. Let v be a vertex such that $d_G(s, v)$ is maximal (line 2). Because $d_G(s, v)$ is maximal, $d_G(s, t) \leq d_G(s, v)$. Thus, by Lemma 3, each path from s to v has eccentricity at most $3k$, i. e. $k(s) \leq 3k$ (line 3). Therefore, the eccentricity of the path selected in line 4 is also at most $3k$.

It is easy to see that line 2 and line 3 run in $\mathcal{O}(m)$ time for a given s. Therefore, the overall runtime for the algorithm is $\mathcal{O}(nm)$. □

Theorem 4. *Algorithm 2 calculates a 2-approximation for the MESP problem in $\mathcal{O}(n^3)$ time.*

Proof (Correctness). Assume a given graph G has a shortest path P from s to t with $\text{ecc}_G(P) = k$ and s is the vertex selected by the loop in line 2. Let Q be a shortest path from s to v. We say the *layer-wise eccentricity* of Q is ϕ if for each layer $L_i^{(s)}$ $(i \leq d_G(s, v))$ there is a vertex $q_i \in Q \cap L_i^{(s)}$ with $\max\{d_G(q_i, u) \mid u \in L_i^{(s)}\} \leq \phi$.

We will now show that lines 4 to 8 calculate for each v the minimal $\phi(v)$ such that there is a shortest path Q from s to v with a layer-wise eccentricity $\phi(v)$.

By induction assume this is true for all vertices $u \in L_j^{(s)}$ with $j \leq i - 1$. Now let v be an arbitrary vertex in $L_i^{(s)}$. Line 6 calculates the maximal distance $\phi(v)$ from v to all other vertices in $L_i^{(s)}$. Since v is the only vertex in $Q \cap L_i^{(s)}$ for every shortest path Q from s to v, the layer-wise eccentricity of each Q is at least $\phi(v)$. Let u be a neighbour of v in the previous layer. By induction $\phi(u)$ is optimal. Therefore, $\phi(v) := \max\{\min_{u \in N_G^-[v]} \phi(u), \phi(v)\}$ (line 7) is optimal for v.

Since line 9 selects the vertex u with the smallest $\phi(u)$ as parent for v, each path Q from s to v in $T(s)$ has an optimal layer-wise eccentricity of $\phi(v)$. Line 8 calculates the maximal distance from v to all vertices in $\{u \mid d_G(s, u) \geq d_G(s, v)\}$. Thus, $\text{ecc}_G(Q) \leq \phi'(v)$ and line 10 and 11 select a shortest path which has an eccentricity at most $\phi'(v)$.

By Lemma 2, we know that P has a layer-wise eccentricity of at most $2k$. Thus, the path Q from s to t in $T(s)$ has a layer-wise eccentricity of at most $2k$. Additionally, Lemma 2 says that t $2k$-dominates all vertices in $\{v \mid d_G(s, v) \geq$

$d_G(s,t)\}$. Therefore, $\text{ecc}_G(Q) \leq 2k$. Thus, the path selected in line 11 is a shortest path with eccentricity at most $2k$. □

Proof (Complexity). Line 1 runs in $\mathcal{O}(nm)$ time. If the distances are stored in an array, they can be later accessed in constant time. Therefore, line 6 and line 8 run in $\mathcal{O}(n)$ time for a given s and v or in $\mathcal{O}(n^3)$ time overall. For a given s, line 7 runs in $\mathcal{O}(m)$ time and therefore has an overall runtime of $\mathcal{O}(nm)$. Line 9 has an overall runtime of $\mathcal{O}(nm)$, line 11 takes $\mathcal{O}(n^2)$ time, and line 10 runs in $\mathcal{O}(n)$ time. Adding all together, the total runtime is $\mathcal{O}(n^3)$. □

Algorithm 2. A 2-approximation for the MESP problem.

Input: A graph $G = (V, E)$.
Output: A shortest path with eccentricity at most $2k$, where k is the minimum eccentricity of all paths in G.

```
1  Calculate the distances d_G(u,v) for all vertex pairs u and v, including
   L_i^(u) = {v ∈ V | d_G(u,v) = i} with 0 ≤ i ≤ ecc_G(u) for each u.
2  foreach s ∈ V do
3      Set φ(s) := 0.
4      for i := 1 to ecc_G(s) do
5          foreach v ∈ L_i^(s) do
6              Set φ(v) := max_{u ∈ L_i^(s)} d_G(u,v).
7              Let N_G^-(v) = L_{i-1}^(s) ∩ N_G(v) denote the neighbours of v in the
               previous layer. Set φ(v) := max{min_{u ∈ N_G^-(v)} φ(u), φ(v)}.
8              Set φ^+(v) := max{d_G(u,v) | d_G(s,u) ≥ i}.
9      Calculate a BFS-tree T(s) starting from s. If multiple vertices u are possible
       as parent for a vertex v, select one with the smallest φ(u).
10     Let t be the vertex for which φ'(t) := max{φ(t), φ^+(t)} is minimal. Set
       k(s) := φ'(t).
11 Among all computed pairs s and t, select a pair (and corresponding path in
   T(s)) for which k(s) is minimal.
```

Algorithm 1 and 2 both iterate over all vertices of the graph to find the best start vertex. Lemma 4 will show that a constant factor approximation can be found with a simple algorithm which starts at an arbitrary vertex. However, the approximation factor will be much higher.

Lemma 4. *Let G be a graph having a shortest path of eccentricity k. Let x be a vertex most distant from some arbitrary vertex, and y be a vertex most distant from x. Then, x, y is a $8k$-dominating pair of G.*

Proof. Let p be an end vertex of a shortest path of eccentricity k in a given graph G. By Lemma 2, the diameter in G of each layer $L_i^{(p)}$ is at most $4k$. Assume, x is most distant from an arbitrary vertex s.

If there is a layer containing both s and x, then $d_G(s,x) \leq 4k$. By the choice of x, each vertex of G is within distance at most $4k$ from s, hence, within distance at most $8k$ from x. Evidently, in this case, x, y is a $8k$-dominating pair of G.

Assume now, without loss of generality, that $x \in L_i^{(p)}$ and $s \in L_l^{(p)}$ with $i < l$. Consider an arbitrary vertex v of G which belongs to a layer with an index smaller than i. We show that $d_G(x,v) \leq 8k$. As $L_i^{(p)}$ separates v from s, a shortest path $P(s,v)$ of G between s and v must have a vertex u in $L_i^{(p)}$. We have $d_G(s,x) \geq d_G(s,v) = d_G(s,u) + d_G(u,v)$ and, by the triangle inequality, $d_G(s,x) \leq d_G(s,u) + d_G(u,x)$. Hence, $d_G(u,v) \leq d_G(u,x)$ and, since both u and x belong to same layer $L_i^{(p)}$, $d_G(u,x) \leq 4k$. That is, $d_G(x,v) \leq d_G(x,u) + d_G(u,v) \leq 2d_G(u,x) \leq 8k$.

If $d_G(x,y) \leq 8k$ then, by the choice of y, each vertex of G is within distance at most $8k$ from x. Hence, x, y is a $8k$-dominating pair of G. So, assume that $d_G(x,y) > 8k$, i. e., the layer $L_j^{(p)}$ with $i < j$ contains y. Repeating the arguments of the previous paragraph, we can show that $d_G(y,v) \leq 8k$ for every vertex v that belongs to a layer with an index greater than j.

Consider now an arbitrary path P of G connecting vertices x and y. P has a vertex in every layer $L_h^{(p)}$ with $i \leq h \leq j$. Hence, for each vertex v of G that belongs to layer $L_h^{(p)}$ $(i \leq h \leq j)$, there is a vertex $u \in P \cap L_h^{(p)}$ such that $d_G(v,u) \leq 4k$. As $d_G(v,x) \leq 8k$ for each vertex v from $L_{i'}^{(p)}$ with $i' < i$ and $d_G(v,y) \leq 8k$ for each vertex v from $L_{j'}^{(p)}$ with $j' > j$, we conclude that $\mathrm{ecc}_G(P) \leq 8k$. □

Corollary 5. *An* 8*-approximation for the MESP problem can be calculated in linear time.*

5 MESP for Certain Graph Classes

So far, we investigated the MESP problem in general graphs. Next, we will show that the problem is solvable in linear or polynomial time for certain graph classes.

Lemma 5. *If a tree has a shortest path of eccentricity k, then any diametral path has eccentricity at most k.*

Proof. In a tree T, let P be a shortest path from s to t with $\mathrm{ecc}_G(P) = k$ and D be a diametral path from x to y. Assume P and D do not intersect. Then there is a vertex $u \in P$ with minimal distance to D and a vertex $z \in D$ with minimal distance to P. Thus, the paths from u to x and from u to y contain z. Because $d_T(x,P) \leq k$, $d_T(y,P) \leq k$, and $d_T(u,z) > 0$, we have $d_T(z,x) < k$ and $d_T(z,y) < k$. Therefore, $d_T(x,y) < 2k$. Each diametral path of length l in a tree contains a vertex c with $\mathrm{ecc}_T(c) = \lceil l/2 \rceil$ [9]. Thus, $\mathrm{ecc}_G(D) \leq k$.

Next, assume P and D intersect. Then there is a vertex $x' \in P \cap D$ with $d_T(x,x') = d_T(x,P) \leq k$ and $y' \in P \cap D$ with $d_T(y,y') = d_T(y,P) \leq k$. Assume there is a vertex v with $d_T(v,D) > k$. Thus, there is a vertex $v' \in P \setminus D$ with $d_T(v,v') \leq k$ and, without loss of generality, $d_T(s,v') < d_T(s,x')$. Therefore,

x' is the vertex in D with minimal distance to v. It follows that $d_T(y,v) = d_T(y,x') + d_T(x',v) > d_T(y,x') + d_T(x',x) = d_T(y,x)$. This contradicts with D being a diametral path. □

Recall that a diametral path in a tree can be found as follows: Select an arbitrary vertex v. Find a most distant vertex x from v and then a most distant vertex y from x. The path from x to y is a diametral path. Thus, it follows from Lemma 5:

Theorem 5. *The MESP problem can be solved for trees in linear time.*

In [6] we show that the MESP problem can be solved in linear time for distance-hereditary graphs and in polynomial time for chordal graphs and dually chordal graphs.

References

1. Bădoiu, M., Chuzhoy, J., Indyk, P., Sidiropoulos, A.: Low-distortion embeddings of general metrics into the line. In: Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC 2005), pp. 225–233. ACM (2005), Baltimore
2. Bădoiu, M., Dhamdhere, K., Gupta, A., Rabinovich, Y., Raecke, H., Ravi, R., Sidiropoulos, A.: Approximation algorithms for low-distortion embeddings into low-dimensional spaces. In: Proceedings of the ACM/SIAM Symposium on Discrete Algorithms (2005)
3. Corneil, D.G., Olariu, S., Stewart, L.: Linear Time Algorithms for Dominating Pairs in Asteroidal Triple-free Graphs. SIAM J. Computing **28**, 292–302 (1997)
4. Deogun, J.S., Kratsch, D.: Diametral path graphs. In: Nagl, M. (ed.) WG 1995. LNCS, vol. 1017, pp. 344–357. Springer, Heidelberg (1995)
5. Leitert, A., Dragan, F.F., Köhler, E.: Line-distortion, bandwidth and path-length of a graph. In: Ravi, R., Gørtz, I.L. (eds.) SWAT 2014. LNCS, vol. 8503, pp. 158–169. Springer, Heidelberg (2014)
6. Dragan, F.F., Leitert, A.: Minimum eccentricity shortest paths in some structured graph classes. In: WG 2015: 41st International Workshop on Graph-Theoretic Concepts in Computer Science, June 17–19, 2015, Munich, Germany, Lecture Notes in Computer Science (2015) (to appear)
7. Fellows, M.R., Fomin, F.V., Lokshtanov, D., Losievskaja, E., Rosamond, F.A., Saurabh, S.: Distortion is fixed parameter tractable. In: Albers, S., Marchetti-Spaccamela, A., Matias, Y., Nikoletseas, S., Thomas, W. (eds.) ICALP 2009, Part I. LNCS, vol. 5555, pp. 463–474. Springer, Heidelberg (2009)
8. Fomin, F.V., Lokshtanov, D., Saurabh, S.: An exact algorithm for minimum distortion embedding. Theor. Comput. Sci. **412**, 3530–3536 (2011)
9. Handler, G.Y.: Minimax location of a facility in an undirected tree graph. Transportation Science **7**, 287–293 (1973)
10. Heggernes, P., Meister, D.: Hardness and approximation of minimum distortion embeddings. Information Processing Letters **110**, 312–316 (2010)
11. Heggernes, P., Meister, D., Proskurowski, A.: Computing minimum distortion embeddings into a path of bipartite permutation graphs and threshold graphs. Theoretical Computer Science **412**, 1275–1297 (2011)

12. Indyk, P.: Algorithmic applications of low-distortion geometric embeddings. In: Proceedings of FOCS 2001, pp. 10–35. IEEE (2005)
13. Indyk, P., Matousek, J.: Low-distortion embeddings of finite metric spaces, Handbook of Discrete and Computational Geometry, 2nd edn., pp. 177–196. CRC Press (2004)
14. Kratsch, D.: Domination and total domination on asteroidal triple-free graphs. Discrete Applied Mathematics **99**, 111–123 (2000)
15. Müller, H.: Hamiltonian circuits in chordal bipartite graphs. Discrete Mathematics **156**, 291–298 (1996)
16. Slater, P.J.: Locating central paths in a graph. Transportation Science **16**, 1–18 (1982)
17. Tenenbaum, J.B., de Silva, V., Langford, J.C.: A global geometric framework for nonlinear dimensionality reduction. Science **290**, 2319–2323 (2000)