Line-Distortion, Bandwidth and Path-Length of a Graph

Feodor F. Dragan[1], Ekkehard Köhler[2], and Arne Leitert[1]

[1] Department of Computer Science, Kent State University, Kent, OH 44242, USA
{dragan,aleitert}@cs.kent.edu
[2] Mathematisches Institut, Brandenburgische Technische Universität Cottbus, D-03013 Cottbus, Germany
ekoehler@math.tu-cottbus.de

Abstract. We investigate the *minimum line-distortion* and the *minimum bandwidth* problems on unweighted graphs and their relations with the *minimum length* of a Robertson-Seymour's path-decomposition. The *length* of a path-decomposition of a graph is the largest diameter of a bag in the decomposition. The *path-length* of a graph is the minimum length over all its path-decompositions. In particular, we show: (i) if a graph G can be embedded into the line with distortion k, then G admits a Robertson-Seymour's path-decomposition with bags of diameter at most k in G; (ii) for every class of graphs with path-length bounded by a constant, there exist an efficient constant-factor approximation algorithm for the minimum line-distortion problem and an efficient constant-factor approximation algorithm for the minimum bandwidth problem; (iii) there is an efficient 2-approximation algorithm for computing the path-length of an arbitrary graph; (iv) AT-free graphs and some intersection families of graphs have path-length at most 2; (v) for AT-free graphs, there exist a linear time 8-approximation algorithm for the minimum line-distortion problem and a linear time 4-approximation algorithm for the minimum bandwidth problem.

1 Introduction and Previous Work

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 [21], as well as in computational chemistry and biology (see [15]). It 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 k such that $d_G(x, y) \leq |f(x) - f(y)| \leq k d_G(x, y)$ for every $x, y \in V$. The parameter k is called the *minimum line-distortion* of G and denoted by $\mathsf{ld}(G)$. The embedding f is called *non-contractive* since $d_G(x, y) \leq |f(x) - f(y)|$ for every $x, y \in V$.

In [3], 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

R Ravi and I.L. Gørtz (Eds.): SWAT 2014, LNCS 8503, pp. 158–169, 2014.

unweighted trees. In another paper [2] 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 k embedding of a weighted tree by a factor polynomial in k. Fast exponential-time exact algorithms for computing the line-distortion of a graph were proposed in [8,9]. Fomin et al. in [9] 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. in [8] gave an $\mathcal{O}(nk^4(2k+1)^{2k})$ time algorithm that for an unweighted graph G and integer k either constructs an embedding of G into the line with distortion at most k, or concludes that no such embedding exists. They extended their approach also to weighted graphs obtaining an $\mathcal{O}(nk^{4W}(2k+1)^{2kW})$ 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. Recently, Heggernes et al. in [13,14] 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 [14]. Furthermore, in [13], 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.

Minimum distortion embedding into the line may appear to be closely related to the widely known and extensively studied graph parameter *bandwidth*, denoted by $\mathsf{bw}(G)$. The only difference between the two parameters is that a minimum distortion embedding has to be *non-contractive*, whereas there is no such restriction for bandwidth. Formally, given an unweighted graph $G=(V,E)$ on n vertices, consider a 1-1 map f of the vertices V into integers in $[1,n]$; f is called a *layout* of G. The *bandwidth of layout* f is defined as the maximum stretch of any edge, i.e., $\mathsf{bw}(f)=\max_{uv\in E}|f(u)-f(v)|$. The *bandwidth* of a graph is defined as the minimum possible bandwidth achievable by any 1-1 map (layout) $V\rightarrow[1,n]$. That is, $\mathsf{bw}(G)=\min_{f:V\rightarrow[1,n]}\mathsf{bw}(f)$.

It is known that $\mathsf{bw}(G)\leq\mathsf{ld}(G)$ for every connected graph G (see, e.g., [14]). However, the bandwidth and the minimum line-distortion of a graph can be very different. For example, it is common knowledge that a cycle of length n has bandwidth 2, whereas its minimum line-distortion is exactly $n-1$ [14]. Bandwidth is known to be one of the hardest graph problems; it is NP-hard even for very simple graphs like caterpillars of hair-length at most 3 [18], and it is hard to approximate by a constant factor even for trees [1]. Polynomial-time algorithms for the exact computation of bandwidth are known for very few graph classes, including bipartite permutation graphs [12] and interval graphs (see, e.g., [17] and papers cited therein). A constant-factor approximation algorithm is known for AT-free graphs [16]. In [10] Golovach et al. showed also that the bandwidth minimization problem is Fixed Parameter Tractable on AT-free graphs by presenting an $n2^{\mathcal{O}(k \log k)}$ time algorithm. For general (unweighted) n-vertex graphs, the minimum bandwidth can be approximated within a factor of $\mathcal{O}(\log^{3.5} n)$ [7]. For n-vertex trees and chordal graphs, the minimum bandwidth can be approximated within a factor of $\mathcal{O}(\log^{2.5} n)$ [11].

Our main tool in this paper is Robertson-Seymour's path-decomposition and its length. A *path-decomposition* ([20]) of a graph $G = (V, E)$ is a sequence of subsets $\{X_i : i \in I\}$ ($I := \{1, 2, \dots, q\}$) of vertices of G, called *bags*, with three properties: (1) $\bigcup_{i \in I} X_i = V$; (2) For each edge $uv \in E$, there is a bag X_i such that $u, v \in X_i$; (3) For every three indices $i \leq j \leq k$, $X_i \cap X_k \subseteq X_j$ (equivalently, the subsets containing any particular vertex form a contiguous subsequence of the whole sequence). We denote a path-decomposition $\{X_i : i \in I\}$ of a graph G by $\mathcal{P}(G)$. The *width* of a path-decomposition $\mathcal{P}(G) = \{X_i : i \in I\}$ is $\max_{i \in I} |X_i| - 1$. The *path-width* of a graph G, denoted by $\mathsf{pw}(G)$, is the minimum width over all path-decompositions $\mathcal{P}(G)$ of G [20]. The caterpillars are exactly the graphs with path-width 1. Following [5] (where the notion of tree-length of a graph was introduced), we define the *length* of a path-decomposition $\mathcal{P}(G)$ of a graph G to be $\lambda := \max_{i \in I} \max_{u,v \in X_i} d_G(u, v)$ (i.e., each bag X_i has diameter at most λ in G). The *path-length* of G, denoted by $\mathsf{pl}(G)$, is the minimum length over all path-decompositions of G. Interval graphs are exactly the graphs with path-length 1; it is known that G is an interval graph if and only if G has a path-decomposition with each bag being a maximal clique of G. Following [6] (where the notion of tree-breadth of a graph was introduced), we define the *breadth* of a path-decomposition $\mathcal{P}(G)$ of a graph G to be the minimum integer r such that for every $i \in I$ there is a vertex $v_i \in V$ with $X_i \subseteq D_G(v_i, r)$ (i.e., each bag X_i can be covered by a disk $D_G(v_i, r)$ of radius at most r in G). Note that vertex v_i does not need to belong to X_i. The *path-breadth* of G, denoted by $\mathsf{pb}(G)$, is the minimum breadth over all path-decompositions of G. Evidently, for any graph G with at least one edge, $1 \leq \mathsf{pb}(G) \leq \mathsf{pl}(G) \leq 2\mathsf{pb}(G)$ holds. Hence, if one parameter is bounded by a constant for a graph G then the other parameter is bounded for G as well.

Recently, Robertson-Seymour's *tree-decompositions* with bags of bounded radius proved to be very useful in designing an efficient approximation algorithm for the problem of minimum stretch embedding of an unweighted graph into its spanning tree [6]. The decision version of the problem is the *tree t-spanner problem* which asks, for a given graph $G = (V, E)$ and an integer t, if a spanning tree T exists such that $d_T(x, y) \leq t\ d_G(x, y)$ for every $x, y \in V$. It was shown in [6] that: (a) if a graph G can be embedded to a spanning tree with stretch t, then G admits a Robertson-Seymour's tree-decomposition with bags of radius at most $\lceil t/2 \rceil$ and diameter at most t in G (i.e., the tree-breadth $\mathsf{tb}(G)$ of G is at most $\lceil t/2 \rceil$ and the tree-length $\mathsf{tl}(G)$ of G is at most t); (b) there is an efficient algorithm which constructs for an n-vertex unweighted graph G with $\mathsf{tb}(G) \leq \rho$ a spanning tree with stretch at most $2\rho \log_2 n$. As a consequence, an efficient $(\log_2 n)$-approximation algorithm for the problem of minimum stretch embedding of an unweighted graph into its spanning tree was obtained [6].

Contribution of This Paper: Motivated by [6], in this paper, we investigate possible connections between the line-distortion and the path-length (path-breadth) of a graph. We show that, for every graph G, $\mathsf{pl}(G) \leq \mathsf{ld}(G)$ and $\mathsf{pb}(G) \leq \lceil \mathsf{ld}(G)/2 \rceil$ hold. Furthermore, we demonstrate that for every class of graphs with path-length bounded by a constant, there is an efficient constant-factor approximation algorithm for the minimum line-distortion problem. As a consequence, every graph G with $\mathsf{ld}(G) = c$ can be embedded in polynomial time into the line with distortion at most $\mathcal{O}(c^2)$ (reproducing a result from [3]). Additionally, using the same technique, we show that, for every class of graphs with path-length bounded by a constant, there is an efficient constant-factor approximation algorithm for the minimum bandwidth problem. We also investigate (i) what particular graph classes have constant bounds on path-length and (ii) how fast the path-length of an arbitrary graph can be computed or sharply estimated. We present an efficient 2-approximation (3-approximation) algorithm for computing the path-length (resp., the path-breadth) of a graph. We show that AT-free graphs and some intersection families of graphs have small path-length and path-breadth. In particular, the path-length of every AT-free graph is at most 2. Using this and some additional structural properties, we give a linear time 8-approximation algorithm for the minimum line-distortion problem and a linear time 4-approximation algorithm for the minimum bandwidth problem for AT-free graphs.

2 Preliminaries

All graphs occurring in this paper are connected, finite, unweighted, undirected, loopless and without multiple edges. We call $G = (V, E)$ an *n-vertex m-edge graph* if $|V| = n$ and $|E| = m$. In this paper we consider only graphs with $n > 1$. A *clique* is a set of pairwise adjacent vertices of G. By $G[S]$ we denote a subgraph of G induced by vertices of $S \subseteq V$. For a vertex v of G, the sets $N_G(v) = \{w \in V : vw \in E\}$ and $N_G[v] = N_G(v) \cup \{v\}$ are called the *open neighborhood* and the *closed neighborhood* of v, respectively.

In a graph G the *length* of a path from a vertex v to a vertex u is the number of edges in the path. The *distance* $d_G(u, v)$ between vertices u and v is the length of a shortest path connecting u and v in G. The *diameter* in G of a set $S \subseteq V$ is $\max_{x,y \in S} d_G(x, y)$ and its *radius* in G is $\min_{x \in V} \max_{y \in S} d_G(x, y)$ (in some papers they are called the *weak diameter* and the *weak radius* to indicate that the distances are measured in G not in $G[S]$). The distance between a vertex v and a set S of G is measured as $d_G(v, S) = \min_{u \in S} d_G(v, u)$. The *disk* of G of radius k centered at vertex v is the set of all vertices at distance at most k to v: $D_G(v, k) = \{w \in V : d_G(v, w) \leq k\}$.

The following result generalizes a characteristic property of the famous class of AT-free graphs (see [4]). An independent set of three vertices such that each pair is joined by a path that avoids the neighborhood of the third is called an *asteroidal triple*. A graph G is an *AT-free graph* if it does not contain any asteroidal triples [4]. Proofs of statements in this section are omitted.

Proposition 1. *Let G be a graph with $\mathsf{pl}(G) \leq \lambda$. Then, for every three vertices u, v, w of G there is one vertex, say v, such that the disk of radius λ centered at v intercepts every path connecting u and w, i.e., the removal of disk $D_G(v, \lambda)$ from G disconnects u and w.*

We will also need the following property of graphs with $\mathsf{pl}(G) \leq \lambda$. A path P of a graph G is called *k-dominating path* of G if every vertex v of G is at distance at most k from a vertex of P, i.e., $d_G(v, P) \leq k$. A pair of vertices x, y of G is called a *k-dominating pair* if every path between x and y is a k-dominating path of G. It is known that every AT-free graph has a 1-dominating pair [4].

Corollary 1. *Every graph G with $\mathsf{pl}(G) \leq \lambda$ has a λ-dominating pair.*

The following proposition further strengthens the connections between graphs with small path-length and AT-free graphs. Recall that the k-power of a graph $G = (V, E)$ is a graph $G^k = (V, E')$ such that for every $x, y \in V$ $(x \neq y)$, $xy \in E'$ if and only if $d_G(x, y) \leq k$.

Proposition 2. *For a graph G with $\mathsf{pl}(G) \leq \lambda$, $G^{2\lambda}$ is an AT-free graph.*

A subset of vertices of a graph is called *connected* if the subgraph induced by those vertices is connected. We say that two connected sets S_1, S_2 of a graph G *see each other* if they have a common vertex or there is an edge in G with one end in S_1 and the other end in S_2. A family of connected subsets of G is called a *bramble* if every two sets of the family see each other. We say that a bramble $\mathcal{F} = \{S_1, \ldots, S_h\}$ of G is k-dominated by a vertex v of G if in every set S_i of $\mathcal{F}$ there is a vertex $u_i \in S_i$ with $d_G(v, u_i) \leq k$.

Proposition 3. *For a graph G with $\mathsf{pb}(G) \leq \rho$, every bramble of G is ρ-dominated by a vertex.*

Corollary 2. *Let G be a graph with $\mathsf{pb}(G) \leq \rho$, S be a subset of vertices of G and $r\colon S \to \mathbb{N}$ be a radius function defined on S such that the disks of the family $\mathcal{F} = \{D_G(x, r(x)) : x \in S\}$ pairwise intersect. Then the disks $\{D_G(x, r(x) + \rho) : x \in S\}$ have a nonempty common intersection.*

3 Bandwidth of Graphs with Bounded Path-Length

In this section we show that there is an efficient algorithm that for any graph G with $\mathsf{pl}(G) = \lambda$ produces a layout f with bandwidth at most $(4\lambda + 2)\mathsf{bw}(G)$. Moreover, this statement is true even for all graphs with λ-dominating shortest paths. Recall that a shortest path P of a graph G is a *k-dominating shortest path* of G if every vertex v of G is at distance at most k from a vertex of P, i.e., $d_G(v, P) \leq k$. We will need the following auxiliary lemma.

Lemma 1 ([19]). *For each vertex $v \in V$ of an arbitrary graph G and each positive integer r, $\frac{|D_G(v,r)|-1}{2r} \leq \mathsf{bw}(G)$.*

The main result of this section is the following.

Proposition 4. *Every graph G with a k-dominating shortest path has a layout f with bandwidth at most $(4k+2)\mathsf{bw}(G)$. If a k-dominating shortest path of G is given in advance, then such a layout 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 k-dominating shortest path of G. Consider a Breadth-First-Search-tree T_P of G started from path P, i.e., BFS(P)-tree of G. For each vertex x_i of P, let X_i be the set of vertices of G that are located in the branch of T_P that is rooted at x_i. We have $x_i \in X_i$. Since P k-dominates G, we have $d_G(v, x_i) \leq k$ for every $i \in \{1, \ldots, q\}$ and every $v \in X_i$. Now create a layout f of G by placing vertices of X_i before all vertices of X_j, if $i < j$, and by placing vertices within each X_i in an arbitrary order.

We claim that this layout f has bandwidth at most $(4k+2)\mathsf{bw}(G)$. Consider any edge uv of G and assume $u \in X_i$ and $v \in X_j$ $(i \leq j)$. For this edge uv we have $f(v) - f(u) \leq |\bigcup_{l=i}^{j} X_l| - 1$. We know also that $d_P(x_i, x_j) = j - i \leq 2k+1$, since P is a shortest path of G and $d_P(x_i, x_j) = d_G(x_i, x_j) \leq d_G(x_i, u) + 1 + d_G(x_j, v) \leq 2k+1$. Consider vertex x_c of P with $c = i + \lfloor (j-i)/2 \rfloor$, i.e., a middle vertex of subpath of P between x_i and x_j. Consider an arbitrary vertex w in X_l, $i \leq l \leq j$. Since $d_G(x_c, w) \leq d_G(x_c, x_l) + d_G(x_l, w)$, $d_G(x_c, x_l) \leq \lceil 2k+1 \rceil / 2$ and $d_G(x_l, w) \leq k$, we get $d_G(x_c, w) \leq 2k+1$. In other words, disk $D_G(x_c, 2k+1)$ contains all vertices of $\bigcup_{l=i}^{j} X_l$. Applying Lemma 1 to $|D_G(x_c, 2k+1)| \geq |\bigcup_{l=i}^{j} X_l|$, we conclude $f(v) - f(u) \leq |\bigcup_{l=i}^{j} X_l| - 1 \leq |D_G(x_c, 2k+1)| - 1 \leq 2(2k+1)\mathsf{bw}(G) = (4k+2)\mathsf{bw}(G)$. □

Corollary 3. *For every n-vertex m-edge graph G, a layout with bandwidth at most $(4\mathsf{pl}(G)+2)\mathsf{bw}(G)$ can be found in $\mathcal{O}(n^2 m)$ time.*

Proof. For an n-vertex m-edge graph G, a k-dominating shortest path with $k \leq \mathsf{pl}(G)$ can be found in $\mathcal{O}(n^2 m)$ time as follows. Iterate over all vertex pairs of G. For each pair pick a shortest path P connecting them and run BFS(P) to find most distant vertex v_P from P. Finally, report that path P for which $d_G(v_P, P)$ is minimum. By Corollary 1, this minimum is at most $\mathsf{pl}(G)$. □

Thus, we have the following interesting conclusion.

Theorem 1. *For every class of graphs with path-length bounded by a constant, there is an efficient constant-factor approximation algorithm for the minimum bandwidth problem.*

In Section 6 we show that the path-length of every AT-free graph is at most 2. Using additional structural properties of AT-free graphs, we give for them a linear time 4-approximation algorithm for the minimum bandwidth problem. This result reproduces an approximation result from [16] with a better run-time.

4 Path-Length and Line-Distortion

In this section, we first show that the line-distortion of a graph gives an upper bound on its path-length and then demonstrate that if the path-length of a graph G is bounded by a constant then there is an efficient constant-factor approximation algorithm for the minimum line-distortion problem on G.

Proposition 5. *For an arbitrary graph G, $\mathsf{pl}(G) \leq \mathsf{ld}(G)$, $\mathsf{pw}(G) \leq \mathsf{ld}(G)$ and $\mathsf{pb}(G) \leq \lceil \mathsf{ld}(G)/2 \rceil$.*

Proof. It is known (see, e.g., [14]) that every connected graph $G = (V, E)$ has a minimum distortion embedding f into the line ℓ (called a *canonic* embedding) such that $|f(x) - f(y)| = d_G(x, y)$ for every two vertices of G that are placed next to each other in ℓ by f. Assume, in what follows, that f is such a canonic embedding and let $k := \mathsf{ld}(G)$.

Consider the following path-decomposition of G created from f. For each vertex v, form a bag B_v consisting of all vertices of G which are placed by f in the interval $[f(v), f(v)+k]$ of the line ℓ. Order these bags with respect to the left ends of the corresponding intervals. Evidently, for every vertex $v \in V$, $v \in B_v$, i.e., each vertex belongs to a bag. More generally, a vertex u belongs to a bag B_v if and only if $f(v) \leq f(u) \leq f(v) + k$. Since $\mathsf{ld}(G) = k$, for every edge uv of G, $|f(u) - f(v)| \leq k$ holds. Hence, both ends of edge uv belong either to bag B_u (if $f(u) < f(v)$) or to bag B_v (if $f(v) < f(u)$). Consider now three bags B_a, B_b, and B_c with $f(a) < f(b) < f(c)$ and a vertex v of G that belongs to B_a and B_c. We have $f(a) < f(b) < f(c) \leq f(v) \leq f(a) + k < f(b) + k$. Hence, necessarily, v belongs to B_b as well.

It remains to show that each bag B_v, $v \in V$, has in G diameter at most k, radius at most $\lceil k/2 \rceil$ and cardinality at most $k + 1$. Indeed, for any two vertices $x, y \in B_v$, we have $|f(x) - f(y)| \leq k$, i.e., $d_G(x, y) \leq |f(x) - f(y)| \leq k$. Furthermore, any interval $[f(v), f(v) + k]$ (of length k) can have at most $k + 1$ vertices of G as the distance between any two vertices placed by f to this interval is at least 1 ($|f(x) - f(y)| \geq d_G(x, y) \geq 1$). Thus, $|B_v| \leq k + 1$ for every $v \in V$.

Consider now the point $p_v := f(v) + \lfloor k/2 \rfloor$ in the interval $[f(v), f(v) + k]$ of ℓ. Assume, without loss of generality, that p_v is between $f(x)$ and $f(y)$, the images of two vertices x and y of G placed next to each other in ℓ by f. Let $f(x) \leq p_v < f(y)$. Since f is a canonic embedding, there must exist in G a vertex c on a shortest path between x and y such that $d_G(x, c) = p_v - f(x)$ and $d_G(c, y) = f(y) - p_v = d_G(x, y) - d_G(x, c)$. We claim that for every vertex $w \in B_v$, $d_G(c, w) \leq \lceil k/2 \rceil$ holds. Assume $f(w) \geq f(y)$ (the case when $f(w) \leq f(x)$ is similar). Then, we have $d_G(c, w) \leq d_G(c, y) + d_G(y, w) \leq (f(y) - p_v) + (f(w) - f(y)) = f(w) - p_v \leq f(w) - f(v) - \lfloor k/2 \rfloor \leq k - \lfloor k/2 \rfloor \leq \lceil k/2 \rceil$. □

It should be noted that the difference between the path-length and the line-distortion of a graph can be very large. A complete graph K_n on n vertices has path-length 1, whereas the line-distortion of K_n is $n - 1$. Note also that the bandwidth and the path-length of a graph do not bound each other. The bandwidth of K_n is $n - 1$ while its path-length is 1. On the other hand, the path-length of cycle C_{2n} is n while its bandwidth is 2.

Now we show that there is an efficient algorithm that for any graph G with $\mathsf{pl}(G) = \lambda$ produces an embedding f of G into the line ℓ with distortion at most $(12\lambda + 7)\mathsf{ld}(G)$. Again, this statement is true even for all graphs with λ-dominating shortest paths. We will need the following auxiliary lemma from [3]. We reformulate it slightly. Recall that a subset of vertices of a graph is called *connected* if the subgraph induced by those vertices is connected.

Lemma 2 ([3]). *Any connected subset $S \subseteq V$ of a graph $G = (V, E)$ can be embedded into the line with distortion at most $2|S| - 1$ in time $\mathcal{O}(|V| + |E|)$. In particular, there is a mapping f, computable in $\mathcal{O}(|V| + |E|)$ time, of vertices from S into points of the line such that $d_G(x, y) \leq |f(x) - f(y)| \leq 2|S| - 1$ for every $x, y \in S$.*

The main result of this section is the following.

Proposition 6. *Every graph G with a k-dominating shortest path admits an embedding f of G into the line with distortion at most $(8k + 4)\mathsf{ld}(G) + (2k)^2 + 2k + 1$. If a k-dominating shortest path of G is given in advance, then such an embedding f can be found in linear time.*

Proof. Like in the proof of Proposition 4, consider a k-dominating shortest path $P = (x_0, x_1, \ldots, x_i, \ldots, x_j, \ldots, x_q)$ of G and identify by BFS(P) the sets X_i, $i \in \{1, \ldots, q\}$. We had $d_G(v, x_i) \leq k$ for every $i \in \{1, \ldots, q\}$ and every $v \in X_i$. It is clear also that each X_i is a connected subset of G. Similar to [3], we define an embedding f of G into the line ℓ by placing vertices of X_i before all vertices of X_j, if $i < j$, and by placing vertices within each X_i in accordance with the embedding mentioned in Lemma 2. Also, for each $i \in \{1, \ldots, q - 1\}$, leave a space of length $2k + 1$ between the interval of ℓ spanning the vertices of X_i and the interval spanning the vertices of X_{i+1}.

We claim that f is a (non-contractive) embedding with distortion at most $(8k+4)\mathsf{ld}(G)+(2k)^2+2k+1$. It is sufficient to show that $d_G(x, y) \leq |f(x)-f(y)|$ for every two vertices of G that are placed next to each other in ℓ by f and that $|f(v) - f(u)| \leq (8k + 4)\mathsf{ld}(G) + (2k)^2 + 2k + 1$ for every edge uv of G (see, e.g., [3,14]).

From Lemma 2, we know that $d_G(x, y) \leq |f(x) - f(y)| \leq 2|X_l| - 1$ for every $x, y \in X_l$ and $l \in \{1, 2, \ldots, q\}$. Additionally, for every $x \in X_i$ and $y \in X_{i+1}$ ($i \in \{1, 2, \ldots, q-1\}$), we have $d_G(x, y) \leq d_G(x, x_i) + 1 + d_G(y, x_{i+1}) \leq 2k + 1 \leq |f(y) - f(x)|$ (as a space of length $2k+1$ is left between the interval of ℓ spanning the vertices of X_i and the interval spanning the vertices of X_{i+1}). Hence, f is non-contractive.

Consider now an arbitrary edge uv of G and assume $u \in X_i$ and $v \in X_j$ ($i \leq j$). For this edge uv we have $f(v) - f(u) \leq \sum_{l=i}^{j}(2|X_l| - 1 + 2k + 1) - 2k - 1 \leq 2|\bigcup_{l=i}^{j} X_l| + 2k(j - i + 1) - 2k - 1 = 2|\bigcup_{l=i}^{j} X_l| + 2k(j - i) - 1$. Recall that $d_P(x_i, x_j) = j - i \leq 2k + 1$, since P is a shortest path of G and $d_P(x_i, x_j) = d_G(x_i, x_j) \leq d_G(x_i, u)+1+d_G(x_j, v) \leq 2k+1$. Hence, $f(v)-f(u) \leq 2|\bigcup_{l=i}^{j} X_l| + 2k(2k + 1) - 1$.

As in the proof of Proposition 4, $|\bigcup_{l=i}^{j} X_l| - 1 \leq (4k + 2)\mathsf{bw}(G)$. As $\mathsf{bw}(G) \leq \mathsf{ld}(G)$ for every graph G (see, e.g., [14]), we get $f(v)-f(u) \leq 2|\bigcup_{l=i}^{j} X_l|+2k(2k+1) - 1 \leq 2(4k + 2)\mathsf{bw}(G) + 2k(2k + 1) + 1 \leq (8k + 4)\mathsf{ld}(G) + 2k(2k + 1) + 1$. □

Corollary 4. *For every n-vertex m-edge graph G, an embedding into the line with distortion at most* $(12\mathsf{pl}(G)+7)\mathsf{ld}(G)$ *can be found in* $\mathcal{O}(n^2m)$ *time.*

Proof. See the proof of Corollary 3 and note that, by Proposition 5, $\mathsf{pl}(G) \leq \mathsf{ld}(G)$. Hence, the distortion established in Proposition 6 becomes $\leq (8\mathsf{pl}(G)+4)\mathsf{ld}(G) + 2(2\mathsf{pl}(G)+1)\mathsf{ld}(G) + 1 \leq (12\mathsf{pl}(G)+7)\mathsf{ld}(G)$. □

Thus, we have the following interesting conclusion.

Theorem 2. *For every class of graphs with path-length bounded by a constant, there is an efficient constant-factor approximation algorithm for the minimum line-distortion problem.*

Using inequality $\mathsf{pl}(G) \leq \mathsf{ld}(G)$ in Corollary 4 once more, we reproduce a result of [3].

Corollary 5 ([3]). *For every graph G with* $\mathsf{ld}(G) = c$, *an embedding into the line with distortion at most* $\mathcal{O}(c^2)$ *can be found in polynomial time.*

It should be noted that, since the difference between the path-length and the line-distortion of a graph can be very large (close to n), the result in Corollary 4 seems to be stronger.

Theorem 1 and Theorem 2 stress the importance of investigations of (i) what particular graph classes have constant bounds on path-length and of (ii) how fast the path-length of an arbitrary graph can be computed or sharply estimated.

5 Constant-Factor Approximation of Path-Length

Let $G = (V, E)$ be an arbitrary graph and s be its arbitrary vertex. A *layering* $\mathcal{L}(s, G)$ of G with respect to a start vertex s is the decomposition of V into the layers $L_i = \{u \in V : d_G(s, u) = i\}, i = 0, 1, \ldots, q$. We can get a path-decomposition of G by adding to each layer L_i $(i > 0)$ all vertices from layer L_{i-1} that have a neighbor in L_i. Let $L_i^+ := L_i \cup (\bigcup_{v \in L_i}(N_G(v) \cap L_{i-1}))$. Clearly, the sequence $\{L_1^+, \ldots, L_q^+\}$ is a path-decomposition of G and can be constructed in $\mathcal{O}(|E|)$ total time. We call this path-decomposition an *extended layering* of G and denote it by $\mathcal{L}^+(s, G)$. It turns out that this type of path-decomposition has length at most twice as large as the path-length of the graph.

Theorem 3. *For every graph G with* $\mathsf{pl}(G) = \lambda$ *there is a vertex s such that the length of the extended layering* $\mathcal{L}^+(s, G)$ *of G is at most* 2λ. *In particular, a factor 2 approximation of the path-length of an arbitrary n-vertex graph can be computed in* $\mathcal{O}(n^3)$ *total time.*

Proof. Consider a path-decomposition $\mathcal{P}(G) = \{X_1, \ldots, X_p\}$ of length $\mathsf{pl}(G) = \lambda$ of G. Let s be an arbitrary vertex from X_1. Consider the layering $\mathcal{L}(s, G)$ of G with respect to s where $L_i = \{u \in V : d_G(s, u) = i\}$ $(i = 0, 1, \ldots, q)$. Let x and y be two arbitrary vertices from L_i $(i \in \{1, \ldots, q\})$ and x' and y' be arbitrary vertices from L_{i-1} with $xx', yy' \in E$. We will show that $\max\{d_G(x,y), d_G(x,y'), d_G(x',y)\} \leq 2\lambda$. By induction on i, we may assume that $d_G(y',x') \leq 2\lambda$ as $x', y' \in L_{i-1}$.

If there is a bag in $\mathcal{P}(G)$ containing both vertices x and y, then $d_G(x,y) \leq \lambda$ and therefore $d_G(x,y') \leq \lambda + 1 \leq 2\lambda$, $d_G(y,x') \leq \lambda + 1 \leq 2\lambda$. Assume now that all bags containing x are earlier in $\mathcal{P}(G) = \{X_1, X_2, \ldots, X_p\}$ than the bags containing y. Let B be a bag of $\mathcal{P}(G)$ containing both ends of edge xx' (such a bag necessarily exists by properties of path-decompositions). By the position of this bag B in $\mathcal{P}(G)$ and the fact that $s \in X_1$, any shortest path connecting s with y must have a vertex in B. Let w be a vertex of B that is on a shortest path of G connecting vertices s and y and containing edge yy'. Such a shortest path must exist because of the structure of the layering $\mathcal{L}(s,G)$ that starts at s and puts y' and y in consecutive layers. We have $\max\{d_G(x,w), d_G(x',w)\} \leq \lambda$. If $w = y'$ then we are done; $\max\{d_G(x,y), d_G(x,y'), d_G(x',y)\} \leq \lambda + 1 \leq 2\lambda$. So, assume that $w \neq y'$. Since $d_G(x,s) = d_G(s,y) = i$ (by the layering) and $d_G(x,w) \leq \lambda$, we must have $d_G(w,y') + 1 = d_G(w,y) = d_G(s,y) - d_G(s,w) = d_G(s,x) - d_G(s,w) \leq d_G(w,x) \leq \lambda$. Hence, $d_G(y,x) \leq d_G(y,w) + d_G(w,x) \leq 2\lambda$, $d_G(y,x') \leq d_G(y,w) + d_G(w,x') \leq 2\lambda$ and $d_G(y',x) \leq d_G(y',w) + d_G(w,x) \leq 2\lambda - 1$.

We conclude that the distance between any two vertices from L_i^+ is at most 2λ, that is, the length of tree decomposition $\mathcal{L}^+(s,G)$ of G is at most 2λ. □

Theorem 4. *For every graph G with $\mathsf{pb}(G) = \rho$ there is a vertex s such that the breadth of the extended layering $\mathcal{L}^+(s,G)$ of G is at most 3ρ. In particular, a factor 3 approximation of the path-breadth of an arbitrary n-vertex graph can be computed in $\mathcal{O}(n^3)$ total time.*

Proof. Since $\mathsf{pl}(G) \leq 2\mathsf{pb}(G)$, by Theorem 3, there is a vertex s in G such that the length of extended layering $\mathcal{L}^+(s,G) = \{L_1^+, \ldots, L_q^+\}$ of G is at most 4ρ. Consider a bag L_i^+ of $\mathcal{L}^+(s,G)$ and a family $\mathcal{F} = \{D_G(x, 2\rho) : x \in L_i^+\}$ of disks of G. Since $d_G(u,v) \leq 4\rho$ for every pair $u, v \in L_i^+$, the disks of $\mathcal{F}$ pairwise intersect. Hence, by Corollary 2, the disks $\{D_G(x, 3\rho) : x \in L_i^+\}$ have a nonempty common intersection. A vertex w from that common intersection has all vertices of L_i^+ within distance at most 3ρ. That is, for each $i \in \{1, \ldots, q\}$ there is a vertex w_i with $L_i^+ \subseteq D_G(w_i, 3\rho)$. □

6 Approximation of Line-Distortions of AT-Free Graphs

The path-length of every AT-free graph is bounded by 2 (proof is omitted).

Proposition 7. *If G is an AT-free graph, then $\mathsf{pb}(G) \leq \mathsf{pl}(G) \leq 2$.*

The class of AT-free graphs contains a number of intersection families of graphs, among them the permutation graphs, the trapezoid graphs and the cocomparability graphs. Theorem 2 implies already that there is an efficient constant-factor approximation algorithm for the minimum line-distortion problem on permutation graphs, trapezoid graphs, cocomparability graphs as well as AT-free graphs. Recall that for arbitrary (unweighted) graphs the minimum line-distortion problem is hard to approximate within a constant factor [3]. Furthermore, the problem remains NP-hard even when the input graph is restricted to a chordal, cocomparability, or AT-free graph [13]. Polynomial-time constant-factor approximation algorithms were known only for split and cocomparability graphs [13]. As far as we know, for AT-free graphs (the class which contains all cocomparability graphs), no prior efficient approximation algorithm was known.

In this section, using additional structural properties of AT-free graphs we give a better approximation algorithm for all AT-free graphs. It is an 8-approximation algorithm and runs in linear time. The following nice structural result from [16] will be very useful.

Lemma 3 ([16]). *Let $G = (V, E)$ be an AT-free graph. Then, there is a dominating path $\pi = (v_0, \dots, v_k)$ and a layering $\mathcal{L} = \{L_0, \dots, L_k\}$ with $L_i = \{u \in V : d_G(u, v_0) = i\}$ such that for all $u \in L_i$ ($i \geq 1$), $uv_i \in E$ or $uv_{i-1} \in E$. Computing π and $\mathcal{L}$ can be done in linear time.*

Theorem 5. *There is a linear time algorithm to compute an 8-approximation of the line-distortion of an AT-free graph.*

Proof. Let G be an AT-free graph. We first compute a path $\pi = (v_0, \dots, v_k)$ and a layering $\mathcal{L} = \{L_0, \dots, L_k\}$ as defined in Lemma 3. To define an embedding f of G into the line, we partition every layer L_i in three sets: $\{v_i\}$, $X_i = \{x : x \in L_i, v_i x \in E\}$, and $\overline{X}_i = L_i \setminus (\{v_i\} \cup X_i)$. Note that if $x \in \overline{X}_i$, then $v_{i-1}x \in E$. Since each vertex in X_i is adjacent to v_i and each vertex in $\overline{X}_i$ is adjacent to v_{i-1}, for all $x, y \in X_i$, $d_G(x, y) \leq 2$, and for all $x, y \in \overline{X}_i$, $d_G(x, y) \leq 2$. Also, for all $x \in X_i$ and $y \in \overline{X}_i$, $d_G(x, y) \leq 3$. The embedding f places vertices of G into the line in the following order: $(v_0, \dots, v_{i-1}, \overline{X}_i, X_i, v_i, \overline{X}_{i+1}, X_{i+1}, v_{i+1}, \dots, v_k)$. Between every two vertices x, y placed next to each other in the line, to guarantee non-contractiveness, f leaves a space of length $d_G(x, y)$ (which is either 1 or 2 or 3, where 3 occurs only when $x \in \overline{X}_i$ and $y \in X_i$ for some i).

We will now show that f approximates the minimum line-distortion of G. Since $\mathcal{L}$ is a BFS layering from v_0, i.e., it represents the distances of vertices from v_0, there is no edge xy with $x \in L_{i-1}$ and $y \in L_{i+1}$. Also note that $D_G(v_i, 2) \supseteq L_i \cup L_{i+1} \cup \{v_{i-1}\}$. By the definition of f, for all $xy \in E$ with $x, y \in L_i \cup L_{i+1}$, $|f(x) - f(y)| < |f(v_{i-1}) - f(v_{i+1})|$. Therefore, counting how many vertices are placed by f between $f(v_{i-1})$ and $f(v_{i+1})$ and the distance in G between vertices placed next to each other, we get $|f(x) - f(y)| \leq 2(|D_G(v_i, 2)| - 2) + 2 = 2(|D_G(v_i, 2)| - 1)$. Using Lemma 1 and the fact that $\mathsf{bw}(G) \leq \mathsf{ld}(G)$, we get $|f(x) - f(y)| \leq 8\ \mathsf{ld}(G)$ for all $xy \in E$. □

It is easy to see that the order in which vertices of G placed by f into the line gives also a layout of G with bandwidth at most $4\mathsf{bw}(G)$. This reproduces an approximation result from [16] (in fact, their algorithm had complexity $\mathcal{O}(m + n \log n)$ for an n-vertex m-edge graph, since it involved a known $\mathcal{O}(n \log n)$ time algorithm to find an optimal layout for a caterpillar tree).

Corollary 6 ([16]). *There is a linear time algorithm to compute a 4-approximation of the minimum bandwidth of an AT-free graph.*

References

1. Blache, G., Karpinski, M., Wirtgen, J.: On approximation intractability of the bandwidth problem, Technical report TR98-014, University of Bonn (1997)
2. Bădoiu, M., Chuzhoy, J., Indyk, P., Sidiropoulos, A.: Low-distortion embeddings of general metrics into the. In: STOC 2005, pp. 225–233. ACM (2005)
3. 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: SODA 2005, pp. 119–128. ACM/SIAM (2005)
4. Corneil, D.G., Olariu, S., Stewart, L.: Asteroidal Triple-Free Graphs. SIAM Journal on Discrete Mathematics 10, 399–430 (1997)
5. Dourisboure, Y., Gavoille, C.: Tree-decompositions with bags of small diameter. Discr. Math. 307, 208–229 (2007)
6. Dragan, F.F., Köhler, E.: An Approximation Algorithm for the Tree t-Spanner Problem on Unweighted Graphs via Generalized Chordal Graphs. In: Goldberg, L.A., Jansen, K., Ravi, R., Rolim, J.D.P. (eds.) RANDOM 2011 and APPROX 2011. LNCS, vol. 6845, pp. 171–183. Springer, Heidelberg (2011)
7. Feige, U.: Approximating the bandwidth via volume respecting embedding. J. of Computer and System Science 60, 510–539 (2000)
8. 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)
9. Fomin, F.V., Lokshtanov, D., Saurabh, S.: An exact algorithm for minimum distortion embedding. Theor. Comput. Sci. 412, 3530–3536 (2011)
10. Golovach, P.A., Heggernes, P., Kratsch, D., Lokshtanov, D., Meister, D., Saurabh, S.: Bandwidth on AT-free graphs. Theor. Comput. Sci. 412, 7001–7008 (2011)
11. Gupta, A.: Improved Bandwidth Approximation for Trees and Chordal Graphs. J. Algorithms 40, 24–36 (2001)
12. Heggernes, P., Kratsch, D., Meister, D.: Bandwidth of bipartite permutation graphs in polynomial time. Journal of Discrete Algorithms 7, 533–544 (2009)
13. Heggernes, P., Meister, D.: Hardness and approximation of minimum distortion embeddings. Information Processing Letters 110, 312–316 (2010)
14. 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)
15. Indyk, P., Matousek, J.: Low-distortion embeddings of finite metric spaces. In: Handbook of Discrete and Computational Geometry, pp. 177–196. CRC (2004)
16. Kloks, T., Kratsch, D., Müller, H.: Approximating the Bandwidth for Asteroidal Triple-Free Graphs. J. Algorithms 32, 41–57 (1999)
17. Kratsch, D., Stewart, L.: Approximating Bandwidth by Mixing Layouts of Interval Graphs. SIAM J. Discrete Math. 15, 435–449 (2002)
18. Monien, B.: The Bandwidth-Minimization Problem for Caterpillars with Hair Length 3 is NP-Complete. SIAM J. Alg. Disc. Meth. 7, 505–512 (1986)
19. Räcke, H.: `http://ttic.uchicago.edu/~harry/teaching/pdf/lecture15.pdf`
20. Robertson, N., Seymour, P.: Graph minors. I. Excluding a forest. Journal of Combinatorial Theory, Series B 35, 39–61 (1983)
21. Tenenbaum, J.B., de Silva, V., Langford, J.C.: A global geometric framework for nonlinear dimensionality reduction. Science 290, 2319–2323 (2000)