

Going Beyond Surfaces in Diameter Approximation

Michał Włodarczyk 

University of Warsaw, Poland

Abstract

Calculating the diameter of an undirected graph requires quadratic running time under the Strong Exponential Time Hypothesis and this barrier works even against any approximation better than $3/2$. For planar graphs with positive edge weights, there are known $(1 + \varepsilon)$ -approximation algorithms with running time $\text{poly}(1/\varepsilon, \log n) \cdot n$. However, these algorithms rely on shortest path separators and this technique falls short to yield efficient algorithms beyond graphs of bounded genus.

In this work we depart from embedding-based arguments and obtain diameter approximations relying on VC set systems and the local treewidth property. We present two orthogonal extensions of the planar case by giving $(1 + \varepsilon)$ -approximation algorithms with the following running times:

- $\mathcal{O}_h((1/\varepsilon)^{\mathcal{O}(h)} \cdot n \log^2 n)$ -time algorithm for graphs excluding an apex graph of size h as a minor,
- $\mathcal{O}_d((1/\varepsilon)^{\mathcal{O}(d)} \cdot n \log^2 n)$ -time algorithm for the class of d -apex graphs.

As a stepping stone, we obtain efficient $(1 + \varepsilon)$ -approximate distance oracles for graphs excluding an apex graph of size h as a minor. Our oracle has preprocessing time $\mathcal{O}_h((1/\varepsilon)^8 \cdot n \log n \log W)$ and query time $\mathcal{O}_h((1/\varepsilon)^2 \cdot \log n \log W)$, where W is the metric stretch. Such oracles have been so far only known for bounded genus graphs. All our algorithms are deterministic.

2012 ACM Subject Classification Theory of computation → Graph algorithms analysis

Keywords and phrases diameter, approximation, distance oracles, graph minors, treewidth

Digital Object Identifier 10.4230/LIPIcs.ESA.2025.39

Related Version Full Version: <https://arxiv.org/abs/2507.03447>

Funding Supported by Polish National Science Centre SONATA-19 grant 2023/51/D/ST6/00155.

1 Introduction

The diameter of a graph G , denoted $\text{diam}(G)$, is the largest distance between a pair of its vertices in the shortest path metric. We assume that each edge is assigned some real positive weight, and the length of a path is the sum of its weights. The obvious way to calculate $\text{diam}(G)$ is to check the distances for all vertex pairs, but this inevitably requires quadratic time. This trivial approach is essentially optimal because already unweighted graphs of diameter 2 or 3 cannot be distinguished in time $\mathcal{O}(m^{1.99})$ under the Strong Exponential Time Hypothesis [49] (see [2] for hardness in sparse graphs). On the other hand, truly subquadratic algorithms to compute diameter are available in planar graphs [10, 34] and graph classes excluding a minor [26, 42, 46]. Still, the best algorithm for planar graphs runs in time¹ $\tilde{\mathcal{O}}(n^{5/3})$ [34], which is far from linear.

The absence of fast algorithms calculating the diameter exactly led the researchers to seek fast approximation algorithms. Weimann and Yuster [55] gave the first almost-linear $(1 + \varepsilon)$ -approximation for planar diameter, running in time $\tilde{\mathcal{O}}(2^{\mathcal{O}(1/\varepsilon)} \cdot n)$. This was improved to randomized time $\tilde{\mathcal{O}}((1/\varepsilon)^5 \cdot n)$ by Chan and Skrepetos [12] (see also [18]). Li and Parter [47] gave a deterministic algorithm of this type, as well as a distributed algorithm with a small number of rounds. Observe that each of these can be used to recognize unweighted planar

¹ The variable n refers to the number of vertices in a graph but all the considered graph classes are sparse so the number of edges is $\mathcal{O}(n)$. The $\tilde{\mathcal{O}}$ -notation hides factors polylogarithmic in n . We refer to the dependence $\tilde{\mathcal{O}}(n)$ as *almost-linear*.



graphs of diameter D , simply by setting $\varepsilon = 1/(2D)$. There are also specialized algorithms for this task by Eppstein [28], with running time $2^{O(D \log D)} \cdot n$, and by Abboud, Mozes, and Weimann [1], with running time $\tilde{O}(D^2 \cdot n)$.

Shortest path separators. An area that is technique-wise related is the design of approximate distance oracles. After a preprocessing step, we would like to quickly estimate any (u, v) -distance within factor $(1 + \varepsilon)$. Thorup [53] exploited the fact that planar graphs admit balanced separators consisting of two shortest paths to give an oracle with preprocessing time $O((1/\varepsilon)^2 \cdot n \log^3 n)$, space $O((1/\varepsilon) \cdot n \log n)$, and query time $O(1/\varepsilon)$. The shortest path separators constitute a crucial ingredient of all the three aforementioned algorithms for estimating planar diameter as well.

Abraham and Gavaille [3] presented a generalization of shortest path separators to graph classes excluding a minor and extended Thorup’s oracle to such classes. However, their separators no longer enjoy such a nice structure. The main caveat is that already in a single balanced separator one might need to use paths P_1, P_2 such that P_2 is a shortest path only in $G - P_1$. Hence P_2 may be much longer than $\text{diam}(G)$, which is easily seen to be unavoidable already in the class of apex graphs², and it is unclear how to exploit such a separator for diameter computation. The second issue is that finding the separator requires highly superlinear time, which burdens the preprocessing time of the oracle. Almost-linear preprocessing time is currently available only for graphs of bounded genus [40] and for unit disk graphs [13]. The lack of well-structured balanced separators beyond surface-embeddable graphs seems to have limited the development of efficient metric-related algorithms so far.

Beyond surface embedding. There are two directions of generalizing the topological graph classes. First, one can pick $O(1)$ faces in the surface embedding of a graph and replace them with certain graphs of pathwidth $O(1)$, which are called the *vortices*. Second, one can insert a set of $O(1)$ vertices, called the *apices*, and connect them possibly to anyone. The structural theorem of Robertson and Seymour [48] states that for any graph H , all the graphs excluding H as a minor can be constructed from graphs embeddable in some fixed surface using these two operations and by taking a clique-sum of them.

When the excluded minor H is an apex graph, this description can be slightly restricted [25]. What is important to us, apex-minor-free graphs enjoy the *local treewidth property*: their treewidth is bounded in terms of the (unweighted) diameter and this relation is known to be linear [24]. Note that the diameter of an edge-weighted graph cannot be related to treewidth by any means because the diameter can be reduced by scaling the weights, while the topological structure of the graph remains the same. Nevertheless, we will see that the local treewidth property is still useful in the considered setting.

The main message of this work is that the local treewidth property can be used instead of shortest path separators for the sake of approximating diameter and radius, as well as for efficient construction of distance oracles. We first employ it to tackle apex-minor-free graphs but we also show that its usefulness is not limited to such graph classes.

Additive core-sets. Suppose we have a list of vertices $S = (s_0, s_1, \dots, s_{k-1})$ in G and we want to (approximately) learn all the distance patterns, between a vertex v and all of S . Chan and Skrepetos [12] employed Cabellos’s technique of abstract Voronoi diagrams [10] to

² G is called an *apex graph* if G can be made planar by removal of a single vertex.

build a data structure to browse such distance patterns in the case when G is planar and S lie on a single face. However, this technique relies heavily on the embedding of the graph.

A more versatile argument was given by Li and Parter [47] in order to approximate planar diameter in a distributed regime. Following their convention, a subset $U' \subseteq U$ is called a δ -additive core-set of U if for every $u \in U$ there exists $u' \in U'$ such that for each $i \in [0, k]$ we have $|d(u, s_i) - d(u', s_i)| \leq \delta$. When $\delta = \varepsilon \cdot \text{diam}(G)$, a trivial rounding argument yields a δ -additive core-set of size $(1/\varepsilon)^k$. Remarkably, the dependence on k can be in fact removed from the exponent [47]. This theorem relies on a bound on the VC dimension of a certain set system and this argument carries over to any minor-free graph class [46]. A direct consequence of these theorems reads as follows. (Here K_h stands for the h -vertex clique.)

► **Lemma 1** (Corollary of [46, Theorem 1.3]). *Let G be an edge-weighted K_h -minor-free graph G , $S = (s_0, \dots, s_{k-1}) \in V(G)^k$, $U \subseteq V(G)$, and $\varepsilon > 0$. Then U admits an $(\varepsilon \cdot \text{diam}(G))$ -additive core-set (with respect to S) of size $\mathcal{O}(k^{h-1} \cdot (1/\varepsilon)^h)$.*

For the sake of completeness, we provide a proof in Section 2. The local treewidth property will allow us to apply this lemma for $k = \text{poly}(1/\varepsilon)$ to efficiently compress information about distances in large subgraphs.

1.1 Our contribution

A crucial property of a separator S consisting of $\mathcal{O}(1)$ shortest paths is that S can be covered by $\mathcal{O}(1/\varepsilon)$ subgraphs of strong³ diameter $\varepsilon \cdot \text{diam}(G)$. Then the interaction between the two sides of S can be channeled through $\mathcal{O}(1/\varepsilon)$ portals on S when small distance distortion is allowed. We show that balanced separators of this kind exist and can be efficiently found in any class excluding an apex minor. Specifically, we construct a clustering of an apex-minor-free graph into pieces of small strong diameter such that the graph obtained by contracting each cluster has bounded treewidth. We use this construction to estimate the eccentricity $\text{ecc}_G(v)$ of every vertex v , that is, $\max_{u \in V(G)} d_G(v, u)$.

► **Theorem 2.** *Let H be an apex graph of size h . There is an algorithm that, given an edge-weighted H -minor-free graph G and $\varepsilon > 0$, runs in time $\mathcal{O}_h((1/\varepsilon)^{3h+2} \cdot n \log^2 n)$ and estimates eccentricity of each vertex $v \in V(G)$ within additive error $\varepsilon \cdot \text{diam}(G)$. Moreover, for each $v \in V(G)$ the algorithm outputs vertex v' for which $d_G(v, v') \geq \text{ecc}_G(v) - \varepsilon \cdot \text{diam}(G)$.*

As a corollary, we get $(1 + \varepsilon)$ -approximation for both diameter and radius, which are respectively the largest and the smallest eccentricity.

To employ the additive core-sets in the proof of Theorem 2, we first need to precompute approximate distances for $\mathcal{O}((1/\varepsilon)^2 \cdot n \log n)$ pairs of vertices. To this end, we design a distance oracle with preprocessing and query times tailored to our needs. Similar oracles are available for minor-free graphs [3, 16, 40] but their preprocessing times are far from linear. An oracle with additive error $\varepsilon \cdot \text{diam}(G)$ suffices for the proof of Theorem 2 but we can improve it to multiplicative error by taking into account the metric stretch, i.e., the ratio between the largest and smallest positive distances.

► **Theorem 3.** *Let H be an apex graph. Given an edge-weighted H -minor-free graph G of metric stretch W and $\varepsilon > 0$, one can build an $(1 + \varepsilon)$ -approximate distance oracle with preprocessing time $\mathcal{O}_H((1/\varepsilon)^8 \cdot n \log n \log W)$ and query time $\mathcal{O}_H((1/\varepsilon)^2 \cdot \log n \log W)$.*

³ A subgraph H of G has *strong diameter* D if every pair of vertices $u, v \in V(H)$ can be connected by a path in H of length at most D . For comparison, *weak diameter* refers to distances measured in G .

The simplest case not covered by Theorem 2 is of course given by the class of apex graphs. This class no longer enjoys the local treewidth property. Moreover, some problems which are easy on planar graphs become much harder after insertion of a single apex [5, 11]. Note that removing just a single vertex can increase the diameter significantly and so our clustering technique cannot be directly applied to decompose such a graph.

We give a stronger theorem that in particular covers the apex graphs. As a consequence of Robertson and Seymour’s structural theorem, for every graph K there exists an apex graph H and a constant c , such that every K -minor-free graph can be constructed via a clique-sum from graphs $G_1, \dots, G_\ell$ with the following property: each G_i becomes H -minor-free after deletion of c vertices [23, Theorem VI.1]. We extend Theorem 2 to such graphs, assuming that the deletion set is known.

► **Theorem 4.** *Let $h, c \in \mathbb{N}$ and $\mathcal{H}$ be a graph class excluding an apex graph of size h . There is an algorithm that, given an edge-weighted graph G , $\varepsilon > 0$, and a set $A \subseteq V(G)$ of size c , for which $G - A \in \mathcal{H}$, runs in time $\mathcal{O}_{h+c}((1/\varepsilon)^\alpha \cdot n \log^2 n)$, where $\alpha = \mathcal{O}((h+1)(c+1))$, and estimates $\text{diam}(G)$ within factor $(1 + \varepsilon)$.*

Unlike the apex-minor-free case, this time we do not estimate eccentricities because Theorem 4 involves a certain “irrelevant vertex rule” that only preserves the diameter (approximately). The assumption that the set A is given in the input can be dropped for the class $\mathcal{H}$ of planar graphs, i.e., when the input is a d -apex graph. In this case, one can find the set A in time $2^{\mathcal{O}(d \log d)} \cdot n$ [37], which leads to the following unconditional statement.

► **Corollary 5.** *There is an $\mathcal{O}_d((1/\varepsilon)^{\mathcal{O}(d)} \cdot n \log^2 n)$ -time algorithm that estimates the diameter of an edge-weighted d -apex graph within factor $(1 + \varepsilon)$.*

This result can be regarded as a generalization of the planar case in a direction that is orthogonal to Theorem 2. Unfortunately, when the class $\mathcal{H}$ excludes an arbitrary apex graph H , we do not know whether one can drop the assumption about the set A being provided. Indeed, the fastest parameterized algorithm to find set A for which $G - A$ is H -minor-free (where H is some apex graph) requires quadratic time [50]. We discuss the perspective of handling more general graph classes in Section 6.

1.2 Related Work

Apart from the ones used in this work, there are more VC set systems considered in the context of minor-free graph metrics [7, 8, 21, 22, 26, 39]. Some other approaches to such metrics are based on a decomposition into low-diameter subgraphs, including the classic KPR theorem [41]. Our low-diameter clustering scheme bears resemblance to the *shortcut partitions* [14, 15, 16] which provide even stronger guarantees but it is not known how to construct them in almost-linear time. The related *buffered cop decompositions* [16] yield cluster graphs of bounded weak coloring numbers [6]. These in turn are connected to *sparse covers* [4, 9, 29], however there the low-diameter subgraphs need not to be disjoint.

An alternative way to represent a metric by a bounded treewidth graph is a *metric embedding*. For example, an edge-weighted planar graph G can be embedded into a metric given by a graph of treewidth $f(\varepsilon)$ so that the additive distortion is $\varepsilon \cdot \text{diam}(G)$ [32]. The function f can be estimated as $f(\varepsilon) = \mathcal{O}(\varepsilon^{-4})$ [14]. For apex-minor-free graphs, the current best treewidth bound is $f(\varepsilon) = 2^{\mathcal{O}(1/\varepsilon)}$ [16]. In the latter case, we again do not know if such an embedding can be computed in almost-linear time. See also [23, 30] for metric embeddings of minor-free graphs and a variant with multiplicative distortion.

For more advances in diameter computation, see [13, 17, 27], and for recent developments in planar distance oracles, see [18, 19, 33, 35, 44].

1.3 Organization of the Paper

We begin with Section 2 including preliminaries on VC set systems and additive core-sets. The following three sections contain the main technical lemmas. Sections 3 and 4 culminate with the proof of Theorem 2 while Section 5 contains the crucial lemma for Theorem 4. The formal proofs of Theorems 3 and 4 can be found in the full version of the article. The full version also contains the proofs of the statements marked with (★).

2 Preliminaries

All our statements concern undirected graphs with real positive edge weights. In an unweighted graph, each weight is set to one.

► **Definition 6** (Treewidth). *A tree decomposition of a graph G is a pair (T, χ) where T is a tree, and $\chi: V(T) \rightarrow 2^{V(G)}$ is a function, such that:*

1. *for each $v \in V(G)$ the nodes $\{t \mid v \in \chi(t)\}$ form a non-empty connected subtree of T ,*
 2. *for each edge $uv \in E(G)$ there is a node $t \in V(T)$ with $\{u, v\} \subseteq \chi(t)$.*
- The width of (T, χ) is defined as $\max_{t \in V(T)} |\chi(t)| - 1$. The treewidth of a graph G (denoted $\text{tw}(G)$) is the minimal width a tree decomposition of G .*

VC dimension. A set system is a collection $\mathcal{F}$ of subsets of some set U . We say that a subset $Y \subseteq A$ is *shattered* by $\mathcal{F}$ if $\{Y \cap S : S \in \mathcal{F}\} = 2^Y$, that is, every subset of Y is an intersection of Y with some $S \in \mathcal{F}$. The VC dimension of a set system is the size of the largest shattered subset. The notion of VC dimension was introduced by Vapnik and Chervonenkis [54]. What is important to us, a set system of bounded VC dimension cannot be too large.

► **Lemma 7** (Sauer–Shelah Lemma [51]). *Let $\mathcal{F}$ be a set system of VC dimension d over a set of size m . Then $|\mathcal{F}| = \mathcal{O}(m^d)$.*

2.1 Additive core-sets

For a vector $\mathbf{x} \in \mathbb{R}^k$ its ℓ_∞ -norm is defined as $\ell_\infty(\mathbf{x}) = \max_{i=1}^k |x_i|$.

► **Definition 8** ([47]). *For a sequence of vertices $S = (s_0, \dots, s_{k-1})$ from G we define the distance tuple of v with respect to S as $\Phi_S(v) = (d_G(v, s_0), \dots, d_G(v, s_{k-1}))$.*

For a vertex set V , a subset $V' \subseteq V$ is called a δ -additive core-set with respect to S if for every $v \in V$ there exists $v' \in V'$ such that $\ell_\infty(\Phi_S(v) - \Phi_S(v')) \leq \delta$.

We will take advantage of the bound on VC dimension of certain set systems related to a graph metric. Le and Wulff-Nilsen [46] considered a slight modification of the set systems introduced by Li and Parter [47], and so they denoted them by $\widehat{\mathcal{LP}}$.

► **Definition 9** ([46]). *Let G be an edge-weighted graph, S be a k -tuple of vertices from G , and $M \subseteq \mathbb{R}$. For $v \in V(G)$ we define:*

$$\hat{X}_v = \{(i, \Delta) : i \in [1, k-1], \Delta \in M, d_G(v, s_i) - d_G(v, s_0) \leq \Delta\}.$$

Let $\widehat{\mathcal{LP}}_{G,M}(S) = \{\hat{X}_v : v \in V(G)\}$ be a collection of subsets of the ground set $[k-1] \times M$.

► **Theorem 10** ([46, Theorem 1.3]). *Let G be an edge-weighted K_h -minor free graph and S, M be as in Definition 9. Then $\widehat{\mathcal{LP}}_{G,M}(S)$ has VC dimension at most $h-1$.*

We adapt the construction of additive core-sets from [47] to rely on the set system $\widehat{\mathcal{LP}}$.

► **Lemma 1** (Corollary of [46, Theorem 1.3]). *Let G be an edge-weighted K_h -minor-free graph G , $S = (s_0, \dots, s_{k-1}) \in V(G)^k$, $U \subseteq V(G)$, and $\varepsilon > 0$. Then U admits an $(\varepsilon \cdot \text{diam}(G))$ -additive core-set (with respect to S) of size $\mathcal{O}(k^{h-1} \cdot (1/\varepsilon)^h)$.*

Proof. Let $D = \text{diam}(G)$, $\delta = \varepsilon \cdot D/2$, and $z = \lceil 2/\varepsilon \rceil$. Next, we choose $M = \{j \cdot \delta : j \in [-z, z]\}$. We will extend Definition 9 slightly. For vertices u, v we define $d'(u, v)$ as the largest integer j for which $j \cdot \delta \leq d_G(u, v)$. Note that $d'(u, v)$ is always bounded by z . For a vertex v let $\hat{X}_v^\circ$ be the pair $(d'(v, s_0), \hat{X}_v)$. By Theorem 10 and Lemma 7, the family $\{\hat{X}_v^\circ : v \in V(G)\}$ has size $\mathcal{O}((k-1) \cdot |M|^{h-1}) = \mathcal{O}((k/\varepsilon)^{h-1})$. It follows that the family $\{\hat{X}_v^\circ : v \in V(G)\}$ has $\mathcal{O}(k^{h-1}\varepsilon^{-h})$ elements.

We construct U' by picking for each $X \in \{\hat{X}_v^\circ : v \in V(G)\}$ an arbitrary vertex $v \in U$ with $\hat{X}_v^\circ = X$. Then $|U'| = \mathcal{O}(k^{h-1}\varepsilon^{-h})$. To show that U' is an (εD) -additive core-set, pick $u \in U$. We need to prove that there exists $u' \in U'$ such that $\ell_\infty(\Phi_S(u) - \Phi_S(u')) \leq \varepsilon D$. By construction, there is $u' \in U'$ for which $\hat{X}_u^\circ = \hat{X}_{u'}^\circ$. Now we need to argue that for each $i \in [0, k-1]$ it holds that $|d_G(u, s_i) - d_G(u', s_i)| \leq \varepsilon D$.

The simplest case is for $i = 0$. As $d'(u, s_0) = d'(u', s_0)$, there is an integer j such that both $d_G(u, s_0), d_G(u', s_0)$ belong to the interval $[j \cdot \delta, (j+1) \cdot \delta)$. Hence $|d_G(u, s_0) - d_G(u', s_0)| \leq \delta = \varepsilon D/2$.

Now consider $i \geq 1$. Let j be the largest integer (possibly negative) for which $d_G(u, s_i) - d_G(u, s_0) \leq j \cdot \delta$. We define j' analogously with respect to u' . Note that $j, j' \in [-z, z]$. Because $\hat{X}_u = \hat{X}_{u'}$ it must be $j = j'$. By the same argument as before, the numbers $x = d_G(u, s_i) - d_G(u, s_0)$, $y = d_G(u', s_i) - d_G(u', s_0)$ differ by at most δ . Finally, since $|d_G(u, s_0) - d_G(u', s_0)| \leq \delta$ we obtain that $|d_G(u, s_i) - d_G(u', s_i)| \leq 2\delta = \varepsilon D$. ◀

In our setting we will only know approximations of the distances from S . We show that this suffices to construct an additive core-set.


► **Lemma 11.** *Let G be an edge-weighted K_h -minor-free graph G , $S = (s_0, \dots, s_{k-1}) \in V(G)^k$, and $U \subseteq V(G)$. Let δ be given and denote $\varepsilon = \delta/\text{diam}(G)$. Next, suppose that for each $v \in U$ we are given $\mathbf{y}_v \in \mathbb{R}^k$ that satisfies $\ell_\infty(\mathbf{y}_v - \Phi_S(v)) \leq \delta$. Then in time $\mathcal{O}((k/\varepsilon)^h \cdot |U|)$ we can compute a (5δ) -additive core-set of U of size $\mathcal{O}(k^{h-1}\varepsilon^{-h})$.*

Proof. We first describe the algorithm, which is a simple greedy construction of a 3δ -cover in a metric space. Initialize $\hat{U} = \emptyset$ and iterate over $v \in U$ in any fixed order. When processing v , check if $\hat{U}$ contains u for which $\ell_\infty(\mathbf{y}_v - \mathbf{y}_u) \leq 3\delta$. If there is no such element, add v to $\hat{U}$.

First we show that every vector in $\{\mathbf{y}_v : v \in U\}$ is within distance 3δ in the ℓ_∞ metric from the set $\{\mathbf{y}_v : v \in \hat{U}\}$. Suppose otherwise and let $v \in U$ be a vertex whose vector $\mathbf{y}_v$ is further than 3δ from this set. In particular, this means that $\ell_\infty(\mathbf{y}_v - \mathbf{y}_u) > 3\delta$ for each u considered before v . So v must have been added to $\hat{U}$ at this point; a contradiction.

Now we argue that $\hat{U}$ is a (5δ) -additive core-set. For each $u \in U$ we know that there exists $v \in \hat{U}$ such that $\ell_\infty(\mathbf{y}_v - \mathbf{y}_u) \leq 3\delta$. Next, by assumption $\ell_\infty(\mathbf{y}_u - \Phi_S(u)) \leq \delta$ and likewise for v . Then by the triangle inequality we get $\ell_\infty(\Phi_S(v) - \Phi_S(u)) \leq 5\delta$, as intended.

Next, we bound the size $\hat{U}$. By construction, for each $u, v \in \hat{U}$ we have $\ell_\infty(\mathbf{y}_v - \mathbf{y}_u) > 3\delta$. We apply the triangle inequality again to get $\ell_\infty(\Phi_S(v) - \Phi_S(u)) > \delta$. By applying Lemma 1 for $\varepsilon' = \varepsilon/2$ we obtain $U' \subseteq U$ so that the $(\delta/2)$ -radius balls (in the ℓ_∞ metric) centered at $\{\Phi_S(v) : v \in U'\}$ cover the set $\{\Phi_S(v) : v \in U\}$. Each such ball has diameter δ so it can contain at most one element from the set $\{\Phi_S(v) : v \in \hat{U}\}$. This implies $|\hat{U}| \leq |U'|$ and so the size bound follows from Lemma 1.



■ **Figure 1** Construction of the clustering in Lemma 14. Left: The black edges form a shortest-path tree T rooted at r . The alternating background colors represented the layers. Each connected subgraph of T contained within a single layer forms a cluster. Right: An argument that G_C^I has bounded treewidth for $I = \{2, 3, 4\}$. For simplicity, the edges not originating from T are omitted. After contracting the layers 0 and 1 into a single vertex, we obtain a graph of (unweighted) diameter at most 6. Since this graph is apex-minor-free, we can use the local treewidth property.

Finally, we analyze the running time. When inspecting a vertex $u \in U$ we need to compute distances in $\mathbb{R}^k$ between $\mathbf{x}_u$ and the vectors of all the vertices already placed in the core-set. This requires at most $k \cdot |\hat{U}|$ operations and so the running time bound follows from the bound on $|\hat{U}|$. ◀

3 Low Radius Clustering

We begin with summoning the local treewidth property of apex-minor-free graphs.

► **Lemma 12** ([24]). *For every apex graph H , there exists a constant c such that for every unweighted H -minor-free graph G it holds that $\text{tw}(G) \leq c \cdot \text{diam}(G)$.*

Our first tool is a clustering of a graph into induced subgraphs of small strong diameter.

► **Definition 13.** *For an edge-weighted graph G , we define a δ -radius clustering of G as a collection of pairs $(C_i, c_i)_{i=1}^\ell$ with the following properties.*

1. $(C_i)_{i=1}^\ell$ forms a partition of $V(G)$.
2. For each $i \in [\ell]$ it holds that $G[C_i]$ is connected, $c_i \in C_i$, and each $v \in C_i$ is within distance δ from c_i in $G[C_i]$.

For $\mathcal{C} = (C_i, c_i)_{i=1}^\ell$ we denote by $G_{\mathcal{C}}$ the cluster graph obtained from G by contracting each C_i into a single vertex. The vertex c_i is called the center of the i -th cluster.

Note that each cluster in a δ -radius clustering has strong diameter at most 2δ . As our first goal, we aim to construct a δ -radius clustering, for $\delta = \varepsilon \cdot \text{diam}(G)$, whose cluster graph has treewidth $\mathcal{O}(1/\varepsilon)$. Secondly, even when δ is chosen to be much smaller $\text{diam}(G)$ we would like to put vertices u, v , satisfying $d_G(u, v) = \mathcal{O}(\delta)$, into “nearby” clusters.

For a graph G , a *layering* is simply a function $\text{layer}: V(G) \rightarrow \mathbb{N}$. When such a function is clear from the context, for $I \subseteq \mathbb{N}$ we denote by G^I the subgraph of G induced by the vertices v with $\text{layer}(v) \in I$. Given a clustering $\mathcal{C}$ of G and a layering of $G_{\mathcal{C}}$, we naturally extend it to a layering of G .

► **Lemma 14 (★).** *Let H be an apex graph. Given a connected edge-weighted n -vertex H -minor-free graph G and $\delta > 0$, we can in time $\mathcal{O}(n \log n)$ find a δ -radius clustering $\mathcal{C}$ of G and a layering function $\text{layer}: V(G_{\mathcal{C}}) \rightarrow \mathbb{N}$ such that the following hold.*

1. *For an interval $I \subseteq \mathbb{N}$ of size p , the graph $G_{\mathcal{C}}^I$ has treewidth $\mathcal{O}_H(p)$.*
2. *The treewidth of $G_{\mathcal{C}}$ is bounded by $\mathcal{O}_H(\text{diam}(G)/\delta)$.*
3. *For each $u, v \in V(G)$ and $p \in \mathbb{N}$ satisfying $d_G(u, v) \leq \delta \cdot p$, it holds that $|\text{layer}(u) - \text{layer}(v)| \leq p$. In this case, there is an interval $I \subseteq \mathbb{N}$ depending only on $(\text{layer}(u), \text{layer}(v), p)$, such that $|I| \leq 2p + 1$ and the (u, v) -distance is the same in G and G^I .*

The proof constructs the layering by analyzing a shortest-path tree T rooted at an arbitrary vertex r . Then the j -th layer corresponds to vertices whose distance from r belongs to $[j \cdot \delta, (j+1) \cdot \delta)$. The clustering $\mathcal{C}$ is obtained by contracting all the edges in T connecting vertices sharing the same layer.

Clearly, each root-to-leaf path in T traverses at most $\text{diam}(G)/\delta$ layers, which bounds the (unweighted) diameter of $G_{\mathcal{C}}$. Since $G_{\mathcal{C}}$ is a contraction of G , it is H -minor-free as well and we can take advantage of the local treewidth property. A similar argument yields Item 1 while Item 3 follows from the triangle inequality.

Lemma 12 only guarantees that $\text{tw}(G_{\mathcal{C}})$ can be bounded in terms of ε , without providing an efficient way to construct a tree decomposition. Moreover, the known algorithms for computing optimal (or $\mathcal{O}(1)$ -approximate) tree decomposition require time exponential in treewidth, whereas we aim at polynomial dependence on ε . Luckily, there is one algorithm available that suits our needs, with a bit worse width guarantee.

While several algorithms rely on dynamic programming over a tree decomposition to compute the diameter exactly [2, 28, 36], we cannot afford such an strategy due to prohibitive error accumulation. Instead, we will turn the decomposition into a balanced one, again by increasing its width slightly, and use it to build our data structures.

► **Lemma 15.** *There is an algorithm that, given a graph G of treewidth at most k , runs in time $\mathcal{O}(k^7 \cdot n \log n)$ and outputs a binary rooted tree decomposition of G of width $\mathcal{O}(k^2)$, depth $\mathcal{O}(\log n)$, and size $\mathcal{O}(n)$.*

Proof. First, in time $\mathcal{O}(k^7 \cdot n \log n)$ one can compute a tree decomposition (T, χ) of G of width $k' = \mathcal{O}(k^2)$ [31]. By a standard argument, we can modify (T, χ) to have size $\mathcal{O}(n)$. Next, we turn (T, χ) into a binary rooted tree decomposition of depth $\mathcal{O}(\log n)$ and width at most $4k' + 3$ in time $\mathcal{O}(n)$ [20]. This also preserves the bound $|V(T)| = \mathcal{O}(n)$. ◀

4 Processing a Tree Decomposition

We will need an argument to reroute paths to go through a center c_i of some cluster C_i .

► **Lemma 16.** *Let G be an edge-weighted graph, $u, v, w \in V(G)$, and let C be a vertex set contained in the δ -radius ball around w . Suppose there is a shortest (u, v) -path P in G that intersects C . Then $d_G(u, w) + d_G(w, v) \leq d_G(u, v) + 2\delta$.*

Proof. Orient the path P from u to v and define u', v' as the first and the last vertices from C appearing on P . Since u', v' lie on a shortest (u, v) -path, we have $d_G(u, u') + d_G(v', v) \leq d_G(u, v)$. Next, it holds that each of $d_G(w, u')$, $d_G(w, v')$ is at most δ . Hence $d_G(u, w) \leq d_G(u, u') + \delta$ and $d_G(w, v) \leq d_G(v', v) + \delta$, which yields the desired inequality. ◀

We will now show how, using a tree decomposition of the cluster graph of width $\mathcal{O}((1/\varepsilon)^2)$, to construct a distance oracle and estimate eccentricities. We give a unified proof for both statements because the constructed oracle is directly applied to obtain the second result. The presented algorithm is arguably simpler than the recursive schemes based on shortest path separators since in the analysis we only need to manage $\mathcal{O}(1)$ rounding errors, rather than $\mathcal{O}(\log n)$. We no longer need to assume that the excluded minor is an apex graph and this will allow us later to apply Theorem 17 in a more general setting.

► **Theorem 17.** *Suppose we are given an edge-weighted n -vertex K_h -minor-free graph G , $\delta > 0$, a δ -radius clustering $\mathcal{C}$ of G , and a binary rooted tree decomposition (T, χ) of $G_{\mathcal{C}}$ of width w , depth $\mathcal{O}(\log n)$, and size $\mathcal{O}(n)$.*

1. *In time $\mathcal{O}_h(w \cdot n \log n)$ we can build a data structure that, when queried for $u, v \in V(G)$, outputs $x \in [d_G(u, v), d_G(u, v) + 2 \cdot \delta]$ in time $\mathcal{O}(w \cdot \log n)$.*
2. *Assume additionally that G is connected and let $\tau = \text{diam}(G)/\delta$. In time $\mathcal{O}_h(w^{h+1} \cdot \tau^h \cdot n \log^2 n)$ we can compute eccentricity of each vertex $v \in V(G)$ within additive error $16 \cdot \delta$. Moreover, within the same time we can also output a vertex $u \in V(G)$ for which $d_G(v, u) \geq \text{ecc}_G(v) - 16 \cdot \delta$.*

Let us clarify that the parameter τ does not have to be given and it only appears in the analysis. However, we can always assume that the diameter is known within factor 2 after computing any eccentricity.

Proof. We begin with some notation that will be used throughout the proof. We will use variables u, v to denote vertices of G and variables i, j to index the set of clusters. For a cluster $i \in \mathcal{C}$, let $\text{loc}(i)$ denote the node $t \in V(T)$ that is closest to the root among those satisfying $i \in \chi(t)$. We naturally extend this notation to vertices of G : all vertices $v \in C_i$ receive $\text{loc}(v) = \text{loc}(i)$. This mapping can be computed for all clusters and vertices in linear time by scanning T top-down.

For $t \in V(T)$ let $U_t \subseteq \mathcal{C}$ denote the set of clusters $i \in \mathcal{C}$ for which $\text{loc}(i)$ lies in the subtree of T rooted at t , inclusively. Next, let $V_t = \bigcup_{i \in U_t} C_i$.

► **Claim 18.** It holds that $\sum_{t \in V(T)} |V_t| = \mathcal{O}(n \log n)$, and all the sets V_t can be computed in total time $\mathcal{O}(n \log n)$.


Proof. For each $i \in \mathcal{C}$ the nodes t satisfying $i \in U_t$ must lie on the $\text{loc}(i)$ -to-root path in T and so there are $\mathcal{O}(\log n)$ of them. Hence each $i \in \mathcal{C}$ contributes $\mathcal{O}(|C_i| \cdot \log n)$ to the sum $\sum_{t \in V(T)} |V_t|$, which yields the upper bound as $\mathcal{C}$ is a partition of $V(G)$. This argument leads directly to an algorithm computing the sets V_t . ◁

Distance oracle. Consider $t \in V(T)$, $i \in \mathcal{C}$, and $v \in V(G)$. We call (t, i, v) a *valid triple* when $t = \text{loc}(i)$ and $v \in V_t$. For fixed $t \in V(T)$, the number of valid triples in which t appears is at most $(w + 1) \cdot |V_t|$, and so by Claim 18 the number of valid triples is $\mathcal{O}(w \cdot n \log n)$.

Recall that for $i \in \mathcal{C}$ the vertex $c_i \in C_i$ is the center of cluster i . We will populate table **inner-dist** indexed by valid triples, so that **inner-dist** $[t, i, v]$ will store the exact (v, c_i) -distance within the graph $G[V_t]$. If there is no such path, we set **inner-dist** $[t, i, v] = \infty$.

► **Claim 19.** We can compute **inner-dist** $[t, i, v]$ for all valid triples in total time $\mathcal{O}_h(w \cdot n \log n)$.

Proof. Consider $t \in V(T)$. For each $i \in \mathcal{C}$ with $\text{loc}(i) = t$, we compute all distances from c_i in $G[V_t]$ using the linear single-source shortest path algorithm for K_h -minor-free graphs [52]. The total running time, over all $t \in V(T)$, is proportional to the number of valid triples. ◁



■ **Figure 2** Illustration to Claim 20. By inspecting the nodes $t \in V(T)$ that are ancestral to both $\text{loc}(u), \text{loc}(v)$, we obtain a family of subgraphs $H_1 \subset H_2 \subset H_3 \subset H_4 \subset H_5 = G$, induced by the sets V_t . Here, H_4 is the smallest among them that accommodates the shortest (u, v) -path Q . This path must intersect some cluster C_i for which $\text{loc}(i) = t$, where $G[V_t] = H_4$, as otherwise it would be contained in H_3 . Hence the (u, v) -distance in H_4 is the same as in G and we can stretch Q within H_4 to go through c_i using Lemma 16.

We now argue that information gathered in table `inner-dist` is sufficient to retrieve all distances in G , up to small additive error. This is the only place where we rely on the fact that each cluster has small *strong* diameter.

▷ **Claim 20.** For each $u, v \in V(G)$ there exist $t \in V(T)$ and $i \in \mathcal{C}$ such that:

1. (t, i, u) and (t, i, v) are valid triples,
2. $\text{inner-dist}[t, i, u] + \text{inner-dist}[t, i, v] \leq d_G(u, v) + 2\delta$.

Proof. Let Q be a shortest (u, v) -path in G . Next, let $\mathcal{C}_Q \subseteq \mathcal{C}$ denote the set of clusters intersected by Q and let $T_Q \subseteq V(T)$ be the set of nodes in which $\mathcal{C}_Q$ appear, i.e., $T_Q = \chi^{-1}(\mathcal{C}_Q)$. Since G_C is a contraction of G , the subgraph $G_C[\mathcal{C}_Q]$ is connected. Consequently, T_Q induces a connected subtree of T . Let $t \in T_Q$ be the node that is closest to the root among T_Q . Note that $t = \text{loc}(i)$ for some $i \in \mathcal{C}_Q$.


We claim that the pair (t, i) satisfies the requested conditions. Note that for each $j \in \mathcal{C}_Q$ the node $\text{loc}(j)$ lies in the subtree of T rooted at t . Hence $\mathcal{C}_Q \subseteq U_t$ and so $V(Q) \subseteq V_t$. In particular, we obtain that $u, v \in V_t$ so (t, i, u) , (t, i, v) are valid triples (Item 1).

Since Q is a shortest (u, v) -path and $V(Q) \subseteq V_t$, we infer that the (u, v) -distance in $G[V_t]$ equals their distance in G . Now we apply Lemma 16 to the graph $H = G[V_t]$ and $C = C_i$, $w = c_i$. Because $i \in \mathcal{C}_Q$, we know that Q intersects C_i . It is crucial that all C_i -vertices are within distance δ from c_i when considering distances within $G[C_i]$ (Definition 13(2)), not just in G . As $C_i \subseteq V_t$, this property carries over to the graph H . Lemma 16 implies that $d_H(u, c_i) + d_H(c_i, v) \leq d_H(u, v) + 2\delta$. We have $\text{inner-dist}(t, i, u) = d_H(u, c_i)$, $\text{inner-dist}(t, i, v) = d_H(c_i, v)$, and $d_H(u, v) = d_G(u, v)$, which yields Item 2. ◁

▷ **Claim 21 (Oracle).** Let $u, v \in V(G)$ be given. Using the table `inner-dist` we can output $x \in [d_G(u, v), d_G(u, v) + 2\delta]$ in time $\mathcal{O}(w \cdot \log n)$.

Proof. We iterate over all pairs $t \in V(T)$, $i \in \mathcal{C}$ for which (t, i, u) and (t, i, v) are valid triples. There are only $\mathcal{O}(\log n)$ candidates for t just because (t, i, u) is a valid triple and so t must lie on the $\text{loc}(u)$ -to-root path. For each such t there are clearly at most $w + 1$ candidates for i .

We return x equal to the minimum value of $\text{inner-dist}[t, i, u] + \text{inner-dist}[t, i, v]$ taken over all such pairs (t, i) . For each (t, i) the considered sum is either ∞ or it equals the length of some (u, v) -walk in $G[V_t]$. This implies that $x \geq d_G(u, v)$. The inequality $x \leq d_G(u, v) + 2\delta$ follows from Claim 20. ◁



■ **Figure 3** Illustration to Claim 23. The blue area corresponds to the union of clusters C_i over $i \in \chi(t)$. For any $v \in V_{t_1}$, $u \in V_{t_2}$ the shortest (v, u) -path (red) must traverse some C_i so by Lemma 16 there is a slightly longer path that goes through c_i (purple). The hollow discs represent the additive core-set of V_{t_2} with respect to the centers from $\chi(t)$. This core-set contains a vertex u' which is comparable to u with respect to distances from $\{c_i : i \in \chi(t)\}$ (e.g., the orange path). Hence u' can simulate u for the task of estimating the eccentricity of v .

Claims 19 and 21 yield the additive distance oracle (Item 1). Now we move on to a proof sketch of Item 2, whose complete proof can be found in the full version of the paper.

Eccentricities. For a non-root $t \in V(T)$ we refer to its parent in T as $\text{parent}(t)$. For $u, v \in V(G)$ let $\text{apx-dist}(u, v)$ denote the value returned by the oracle from Claim 21 when queried for the pair u, v . This is a well-defined function because the oracle is deterministic. Recall the concept of an additive core-set from Lemma 1 and that $\tau = \text{diam}(G)/\delta$.

▷ **Claim 22 (★).** In total time $\mathcal{O}((w \cdot \tau)^h \cdot n + w^2 \cdot n \log^2 n)$ we can compute for each non-root $t \in V(T)$ a subset $V'_t \subseteq V_t$ of size $\mathcal{O}(w^{h-1} \tau^h)$ that is a $(10 \cdot \delta)$ -additive core-set of V_t with respect to $\chi(\text{parent}(t))$.

The total number of distances we care about is $(w + 1) \cdot \sum_{t \in V(T)} |V_t| = \mathcal{O}(w \cdot n \log n)$. Each of them can be estimated using the oracle in time $\mathcal{O}(w \log n)$. We would like to use those numbers to compute the additive core-sets, however a priori the bound from Lemma 1 might not hold for approximate distances. Nevertheless, we prove that a greedy procedure still constructs a small core-set, even after perturbation of distances. A similar issue is handled in [47] via randomization but we can afford a simpler argument because our strategy allows us to first query the oracle for all the distances between V_t and $\chi(\text{parent}(t))$.

In the next claim, we argue that the core-sets suffice to estimate the maximal distance between v and u , where v is fixed and u is quantified over all vertices from some subtree. There are three sources of accuracy loss: relying on the additive core-set, on the approximate distance oracle, and stretching the paths along a center from $\chi(t)$ (Lemma 16).

▷ **Claim 23 (★).** Consider $t \in V(T)$ with children t_1, t_2 . Let $v \in V_{t_1}$ and $V' \subseteq V_{t_2}$ be a ρ -additive core-set with respect to $\chi(t)$. Let $y = \max d_G(v, u)$ over all $u \in V_{t_2}$, and x be defined as follows:

$$x = \max_{u \in V'} \min_{i \in \chi(t)} (\text{apx-dist}(v, c_i) + \text{apx-dist}(c_i, u)).$$

Then $|x - y| \leq \rho + 6 \cdot \delta$. Furthermore, the vertex $u^* \in V'$ realizing x satisfies $d_G(v, u^*) \geq y - \rho - 6 \cdot \delta$.

39:12 Going Beyond Surfaces in Diameter Approximation

We now use the additive core-sets to estimate the eccentricity of a vertex. Essentially, we apply Claim 23 for each choice of the lowest common ancestor of v, u ; hence $\mathcal{O}(\log n)$ times.

▷ **Claim 24 (★).** For a given vertex $v \in V(G)$ we can compute its eccentricity within additive error $16 \cdot \delta$ in time $\mathcal{O}(w^{h+1} \cdot \tau^h \cdot \log^2 n)$. Within the same time, we can output a vertex u for which $d_G(v, u) \geq \text{ecc}_G(v) - 16 \cdot \delta$.

Theorem 17(2) follows by applying Claim 24 to each vertex. ◀

► **Theorem 2.** Let H be an apex graph of size h . There is an algorithm that, given an edge-weighted H -minor-free graph G and $\varepsilon > 0$, runs in time $\mathcal{O}_h((1/\varepsilon)^{3h+2} \cdot n \log^2 n)$ and estimates eccentricity of each vertex $v \in V(G)$ within additive error $\varepsilon \cdot \text{diam}(G)$. Moreover, for each $v \in V(G)$ the algorithm outputs vertex v' for which $d_G(v, v') \geq \text{ecc}_G(v) - \varepsilon \cdot \text{diam}(G)$.

Proof. One can assume that G is connected as otherwise the diameter is ∞ . We compute the eccentricity D of an arbitrary vertex v : this can be done in linear time [52]. The diameter of G must be contained in $[D, 2D]$. We apply Lemma 14 for $\delta = \varepsilon D/16$ in time $\mathcal{O}(n \log n)$; let $\mathcal{C}$ denote the obtained δ -radius clustering. Next, let $\tau = \text{diam}(G)/\delta$; note that $\tau = \mathcal{O}(1/\varepsilon)$. By Lemma 14(2) we know that $\text{tw}(G_{\mathcal{C}}) = \mathcal{O}_h(1/\varepsilon)$. We use Lemma 15 to build a binary tree decomposition of $\text{tw}(G_{\mathcal{C}})$ of width $w = \mathcal{O}_h(\varepsilon^{-2})$, depth $\mathcal{O}(\log n)$ and size $\mathcal{O}(n)$; this takes time $\mathcal{O}_h(\varepsilon^{-7} \cdot n \log n)$. Now we can take advantage of Theorem 17(2) to compute the eccentricities of all vertices within additive error $16 \cdot \delta \leq \varepsilon \cdot \text{diam}(G)$, together with vertices that realize these distance approximately. This takes time $\mathcal{O}_h(w^{h+1} \cdot \tau^h \cdot n \log^2 n) = \mathcal{O}_h((1/\varepsilon)^{3h+2} \cdot n \log^2 n)$. Finally, observe that h is greater than 1 in any non-trivial instance so $7 < 3h + 2$ and so the time needed to compute the tree decomposition is negligible. ◀

In the full version of the article we show how to replace additive error in the distance oracle with multiplicative one, which yields Theorem 3.

5 Handling Apices

We move on to the proof of Theorem 4, extending the diameter approximation to graphs G with an apex set A , for which $G - A$ is apex-minor-free. We do not formulate an analogous theorem generalizing the distance oracle because it is trivial. To see this, first precompute all the distances from each $w \in A$ and then, given u, v , return the minimum of their distance in $G - A$ and $\min_{w \in A} d(u, w) + d(w, v)$. The first scenario can be readily handled with the same data structure as in Theorem 3.

The reason why we cannot apply the same reduction for diameter computation is that $\text{diam}(G - A)$ might be huge compared to $\text{diam}(G)$. Hence we lose the relation between the diameter and radii of clusters (governing the additive error), crucial to bound the treewidth of the cluster graph. As a remedy, we will first preprocess the graph $G - A$ to chop it into connected components of bounded diameter.

► **Lemma 25.** Let H be a fixed graph and $c \in \mathbb{N}$. Let G be a connected edge-weighted n -vertex connected graph, $A \subseteq V(G)$ be such that $G - A$ is H -minor-free and $|A| = c$, and $D \in \mathbb{R}$ satisfy $2D \geq \text{diam}(G)$. There is an algorithm that, given the above and $\varepsilon > 0$, runs in time $\mathcal{O}_{H,c}((1/\varepsilon)^c \cdot n)$ and either correctly reports that $\text{diam}(G) > D$ or outputs a connected edge-weighted n -vertex graph G' and $A' \subseteq V(G')$ of size c , with the following guarantees.

1. $G' - A'$ is H -minor-free.
2. $\text{diam}(G') \geq \text{diam}(G)$.
3. If $\text{diam}(G) \leq D$ then $\text{diam}(G') \leq (1 + \varepsilon) \cdot D$.
4. Each connected component of $G' - A'$ has strong diameter at most $8 \cdot (2/\varepsilon)^c \cdot D$.

Proof. First of all, we remove all edges with weight larger than $2D$ as they cannot contribute to any shortest path. Observe that G is K_h -minor-free for $h = |V(H)| + c$. We compute shortest paths from each vertex $w \in A$ using the linear SSSP algorithm for minor-free graphs [52]. If any computed distance is greater than D , we can terminate the algorithm. Otherwise for each $w \in A$, $v \in V(G)$, we insert the wv -edge with weight $d_G(w, v)$ or update the weight if the edge is already present. This preserves all the pairwise distances in G and does not affect the graph $G - A$. We will keep the variable G to refer to the modified graph. The performed modification ensures the following property.

▷ **Claim 26.** Let $u, v \in V(G)$ and suppose that there is a shortest (u, v) -path in G that intersects A . Then there is a shortest (u, v) -path with one internal vertex that belongs to A .

For $u, v \in V(G)$ let $d'(u, v)$ denote the largest integer j for which $j \cdot \varepsilon D/2 \leq d_G(u, v)$. For $v \in V(G) \setminus A$ we define its *signature* as $S_v = (d'(v, a_1), \dots, d'(v, a_c))$ where $(a_1, \dots, a_c)$ is some fixed ordering of A . Let $\mathcal{S} = \{S_v : v \in V(G) \setminus A\}$. Since we have not terminated the algorithm in the first step, each $d'(v, a_i)$ is most $\lfloor 2/\varepsilon \rfloor$ and so $|\mathcal{S}| \leq (2/\varepsilon)^c$.

We say that a pair $S_1, S_2 \in \mathcal{S}$ is *relevant* if for each $i \in [c]$ we have $S_1[i] + S_2[i] > 2/\varepsilon$. We also call $S \in \mathcal{S}$ relevant if it appears in at least one relevant pair.

▷ **Claim 27.** Let $u, v \in V(G) \setminus A$ be such that (S_u, S_v) is a relevant pair. Then there is no (u, v) -path of length $\leq D$ that goes through A .

Proof. If there was such a path then $d_G(u, w) + d_G(w, v) \leq D$ for some $w \in A$. But for each $w \in A$ we have $d_G(u, w) + d_G(w, v) \geq (\varepsilon D/2) \cdot (d'(u, w) + d'(w, v)) > D$. ◁

▷ **Claim 28.** Let $u, v \in V(G) \setminus A$ be such that $S_u = S_v$ and S_u is relevant. If $d_{G-A}(u, v) > 2D$ then $\text{diam}(G) > D$.


Proof. Suppose that $\text{diam}(G) \leq D$. Since S_u is relevant, there exists $z \in V(G) \setminus A$ such that (S_u, S_z) is a relevant pair. Due to Claim 27, the shortest (u, z) -path avoids A so $d_{G-A}(u, z) \leq D$. By the same argument $d_{G-A}(v, z) \leq D$ and the claim follows from the triangle inequality. ◁

For each relevant $S \in \mathcal{S}$ let us choose an arbitrary v with $S_v = S$; we denote the set of chosen vertices as B . Clearly $|B| \leq |\mathcal{S}| \leq (2/\varepsilon)^c$. For each $b \in B$ we compute all the distances from b in $G - A$ in linear time and check whether there is any $v \in V(G) \setminus A$ with $S_v = S_b$ and $d_{G-A}(b, v) > 2D$. If so, then by Claim 28 we can again terminate the algorithm.

▷ **Claim 29.** Suppose that the algorithm has not been terminated and let $Z \subseteq V(G) \setminus A$ be the set of vertices z satisfying $d_{G-A}(z, b) > 3D$ for all $b \in B$. Next, let $F \subseteq E(G - A)$ be the set of edges from $G - A$ with at least one endpoint in Z . Suppose that $\text{diam}(G) \leq D$. Then $\text{diam}(G \setminus F) \leq (1 + \varepsilon) \cdot D$.

Proof. Consider $u, v \in V(G)$ and let P be a shortest (u, v) -path in G . By assumption, P is not longer than D . We need to show that there is a (u, v) -path of length $\leq (1 + \varepsilon) \cdot D$ that avoids F . If P intersects A then by Claim 26 we can assume that P consists only of edges incident to A (at most two). Then clearly P avoids F .

Suppose now that $V(P) \cap A = \emptyset$ and that P uses some edge from F . Then there is $z \in V(P) \cap Z$ whose distance in $G - A$ to each of u, v is at most D . We claim that S_u, S_v are irrelevant. Assume w.l.o.g. that S_u is relevant and let $b \in B$ be the chosen vertex with $S_u = S_b$. We have $d_{G-A}(u, b) \leq 2D$ as otherwise we would have terminated the algorithm due to Claim 28. But $d_{G-A}(u, z) \leq D$ so $d_{G-A}(z, b) \leq 3D$. This contradicts the property used to define Z .



■ **Figure 4** Illustration to Claim 29. The set B of representatives for relevant signatures is given by the squares. By Claim 28, every vertex sharing the same signature as $b \in B$ (represented by the same color) must lie in the $2D$ -radius ball around b in $G - A$. If z is outside all $3D$ -radius balls, then every path in $G - A$ of length $\leq D$ traversing z must connect vertices u, v for which (S_u, S_v) is an irrelevant pair. For such pairs, there must be an almost shortest path traversing A .

We established that S_u, S_v are irrelevant, so the pair (S_u, S_v) is also irrelevant. By definition, this means that there is $w \in A$ such that $d'(u, w) + d'(w, v) \leq 2/\varepsilon$. We have $d_G(u, w) \leq (\varepsilon D/2) \cdot (d'(u, w) + 1)$ and likewise for the pair (v, w) . We estimate

$$d_G(u, w) + d_G(w, v) \leq (\varepsilon D/2) \cdot (d'(u, w) + d'(w, v) + 2) \leq D + (\varepsilon D/2) \cdot 2 = (1 + \varepsilon) \cdot D.$$

Hence the path (u, w, v) is sufficiently short and it avoids F . The claim follows. $\triangleleft$

We showed that removing all edges from F is safe for approximating the diameter of G . It remains to prove that this operation splits $G - A$ into pieces of bounded strong diameter.

▷ **Claim 30.** Let $C \subseteq V(G) \setminus A$ be a set of vertices inducing a connected component of $(G - A) \setminus F$. Then $\text{diam}(G[C]) \leq 8 \cdot (2/\varepsilon)^c \cdot D$.

Proof. For $b \in B$ let C_b be the set of vertices within distance $3D$ from b in $G - A$. Observe that C is either a single vertex from Z or a union of C_b for $b \in B'$ for some $B' \subseteq B$. In the first case the diameter is 0 so let us focus on the second case.

Let $u, v \in C$ and P be a shortest (u, v) -path in $G[C]$. Then $u \in C_{b_1}$ for some $b_1 \in B'$. Let v_1 be the farthest vertex on P (counting from u) that belongs to C_{b_1} . The (u, v_1) -distance in $G[C]$ is at most their distance in $G[C_{b_1}]$ so $d_P(u, v_1) \leq 6D$ as P is a shortest path. Next, take u_2 to be the next vertex on P after v_1 and pick $b_2 \in B$ satisfying $u_2 \in C_{b_2}$. Analogously as before, we find the vertex $v_2 \in C_{b_2}$ that is farthest on P and continue this process. By construction, every vertex $b \in B'$ appears at most once as b_i so this splits P into at most $|B'|$ subpaths of length $\leq 6D$. Also for each i the edge between v_i and u_{i+1} has weight at most $2D$ because we have removed all the heavier edges at the very beginning. The claim follows from $|B'| \leq |B| \leq (2/\varepsilon)^c$. $\triangleleft$

We check that $G' = G \setminus F$ and $A' = A$ satisfy all the points of the lemma. First, $G' - A$ is a subgraph of $G - A$ so it remains H -minor-free. Next, removing edges can only increase the diameter so $\text{diam}(G') \geq \text{diam}(G)$ while the second inequality follows from Claim 29. The diameter bound for components of $G' - A'$ has been given in Claim 30. Finally, the sets Z and F can be identified in time $\mathcal{O}_H(|B| \cdot n) = \mathcal{O}_{H,c}((1/\varepsilon)^c \cdot n)$. $\blacktriangleleft$

We briefly explain how Lemma 25 leads to Theorem 4. Given $D \in [\text{diam}(G), 2 \cdot \text{diam}(G)]$ we want to learn that $\text{diam}(G) > D$ or that $\text{diam}(G) \leq (1 + \varepsilon) \cdot D$. Let (G', A') be the output given by Lemma 25 and $\delta = \Theta(\varepsilon D)$. Each connected component G^i of $G' - A'$ satisfies $\text{diam}(G^i) = \mathcal{O}((1/\varepsilon)^c \cdot D)$. We apply Lemma 14 to get a δ -radius clustering $\mathcal{C}^i$ of G^i so that the treewidth of the cluster graph $G_{\mathcal{C}^i}^i$ is $\mathcal{O}(\text{diam}(G^i)/\delta) = \mathcal{O}((1/\varepsilon)^{c+1})$. We combine those to get a δ -radius clustering $\mathcal{C}'$ of G' by assigning each $w \in A'$ a singleton cluster.

Using Lemma 15 we can compute a balanced tree decomposition of $G_{\mathcal{C}'}^i$ of width $w = \mathcal{O}((1/\varepsilon)^{2(c+1)})$. It can be now processed with Theorem 17(2): in time $\mathcal{O}(w^{h+1}(1/\varepsilon)^h \cdot n \log^2 n)$ we can estimate $\text{diam}(G')$ within additive error $16 \cdot \delta = \Theta(\varepsilon D)$. If this value is too large, we learn that $\text{diam}(G') > (1 + \varepsilon) \cdot D$, implying $\text{diam}(G) > D$. Otherwise we infer that $\text{diam}(G) \leq \text{diam}(G') \leq (1 + \varepsilon) \cdot D$. Repeating this procedure for different D allows us to estimate $\text{diam}(G)$ within multiplicative error $(1 + \varepsilon)$. The details can be found in the full version of the article.

6 Conclusion and Open Problems

We have presented almost-linear algorithms approximating diameter that bypass topological arguments and work way beyond surface-embeddable graphs. A natural next step would be to generalize these results further, to all graph classes excluding a minor. The reason why we could not extend Corollary 5 to all cases covered in Theorem 4 is that we need to know the deletion set to the apex-minor-free class. The best known algorithm finding such a set requires quadratic time [50]. To apply the same strategy for the general case of H -minor-free graphs, one would need to compute some form of the Robertson and Seymour's decomposition, which seems an even harder task. However, recent advances in minor testing [43] suggest that this may be possible in time $\mathcal{O}_H(n^{1+o(1)})$ (see the section "Outlook" there). We have not attempted to generalize Theorem 4 to work on such decompositions because (1) this would require a tedious analysis of a balanced tree decomposition of the clique-sum decomposition, and (2) this strategy currently seems hopeless to achieve running time $\mathcal{O}_H(\text{poly}(1/\varepsilon, \log n) \cdot n)$. Is there an alternative approach that would work without the decomposition at hand?

Another question is whether one can improve the dependence on n to $\mathcal{O}(n \log n)$ or even $\mathcal{O}(n)$? The current running time bound follows from the fact that we need to precompute $\Omega(n \log n)$ distances and each oracle call requires time $\Omega(\log n)$. Can this be improved? For example, for planar metrics there exist distance oracles with constant query time (for constant ε) albeit with preprocessing time $\Omega(n \log^3 n)$ [12, 45].

When it comes to the dependence on ε , our running time is burdened by the fact that for k -treewidth graphs we compute a tree decomposition of width $\mathcal{O}(k^2)$. Unfortunately, all the known almost-linear $\mathcal{O}(1)$ -approximation algorithms for treewidth require time exponential in k . Since we rely on a construction from [31] that works for general graphs, it is tempting to improve the width guarantee therein when the input graph is H -minor-free. Again, this is doable in planar graphs [38].

References

- 1 Amir Abboud, Shay Mozes, and Oren Weimann. What Else Can Voronoi Diagrams Do for Diameter in Planar Graphs? In Inge Li Gørtz, Martin Farach-Colton, Simon J. Puglisi, and Grzegorz Herman, editors, *31st Annual European Symposium on Algorithms (ESA 2023)*, volume 274 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 4:1–4:20, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ESA.2023.4.

- 2 Amir Abboud, Virginia Vassilevska Williams, and Joshua R. Wang. Approximation and fixed parameter subquadratic algorithms for radius and diameter in sparse graphs. In Robert Krauthgamer, editor, *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 377–391. SIAM, 2016. doi:10.1137/1.9781611974331.CH28.
- 3 Ittai Abraham and Cyril Gavoille. Object location using path separators. In *Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*, pages 188–197, 2006. doi:10.1145/1146381.1146411.
- 4 Ittai Abraham, Cyril Gavoille, Dahlia Malkhi, and Udi Wieder. Strong-diameter decompositions of minor free graphs. In *Proceedings of the nineteenth annual ACM symposium on Parallel algorithms and architectures*, pages 16–24, 2007. doi:10.1145/1248377.1248381.
- 5 Francisco Barahona. The max-cut problem on graphs not contractible to K_5 . *Operations Research Letters*, 2(3):107–111, 1983. doi:10.1016/0167-6377(83)90016-0.
- 6 Romain Bourneuf and Marcin Pilipczuk. Bounding ϵ -scatter dimension via metric sparsity. In Yossi Azar and Debmalya Panigrahi, editors, *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025, New Orleans, LA, USA, January 12-15, 2025*, pages 3155–3171. SIAM, 2025. doi:10.1137/1.9781611978322.101.
- 7 Nicolas Bousquet and Stéphan Thomassé. VC-dimension and Erdős–Pósa property. *Discrete Mathematics*, 338(12):2302–2317, 2015. doi:10.1016/j.disc.2015.05.026.
- 8 Vladimir Braverman, Shaofeng H-C Jiang, Robert Krauthgamer, and Xuan Wu. Coresets for clustering in excluded-minor graphs and beyond. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2679–2696. SIAM, 2021. doi:10.1137/1.9781611976465.159.
- 9 Costas Busch, Ryan LaFortune, and Srikanta Tiruthapura. Sparse covers for planar graphs and graphs that exclude a fixed minor. *Algorithmica*, 69:658–684, 2014. doi:10.1007/s00453-013-9757-4.
- 10 Sergio Cabello. Subquadratic algorithms for the diameter and the sum of pairwise distances in planar graphs. *ACM Trans. Algorithms*, 15(2):21:1–21:38, 2019. doi:10.1145/3218821.
- 11 Dibyayan Chakraborty and Kshitij Gajjar. Finding geometric representations of apex graphs is NP-hard. *Theor. Comput. Sci.*, 971:114064, 2023. doi:10.1016/J.TCS.2023.114064.
- 12 Timothy M. Chan and Dimitrios Skrepetos. Faster Approximate Diameter and Distance Oracles in Planar Graphs. In Kirk Pruhs and Christian Sohler, editors, *25th Annual European Symposium on Algorithms (ESA 2017)*, volume 87 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 25:1–25:13, Dagstuhl, Germany, 2017. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ESA.2017.25.
- 13 Timothy M. Chan and Dimitrios Skrepetos. Approximate shortest paths and distance oracles in weighted unit-disk graphs. *Journal of Computational Geometry*, 10(2):3–20, February 2019. doi:10.20382/jocg.v10i2a2.
- 14 Hsien-Chih Chang, Jonathan Conroy, Hung Le, Lazar Milenkovic, Shay Solomon, and Cuong Than. Covering planar metrics (and beyond): $O(1)$ trees suffice. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 2231–2261. IEEE, 2023. doi:10.1109/FOCS57990.2023.00139.
- 15 Hsien-Chih Chang, Jonathan Conroy, Hung Le, Lazar Milenkovic, Shay Solomon, and Cuong Than. Resolving the steiner point removal problem in planar graphs via shortcut partitions. *arXiv preprint arXiv:2306.06235*, 2023. doi:10.48550/arXiv.2306.06235.
- 16 Hsien-Chih Chang, Jonathan Conroy, Hung Le, Lazar Milenković, Shay Solomon, and Cuong Than. Shortcut partitions in minor-free graphs: Steiner point removal, distance oracles, tree covers, and more. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 5300–5331. SIAM, 2024. doi:10.1137/1.9781611977912.191.
- 17 Hsien-Chih Chang, Jie Gao, and Hung Le. Computing Diameter+2 in Truly-Subquadratic Time for Unit-Disk Graphs. In Wolfgang Mulzer and Jeff M. Phillips, editors, *40th International Symposium on Computational Geometry (SoCG 2024)*, volume 293 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 38:1–38:14, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.SoCG.2024.38.

- 18 Hsien-Chih Chang, Robert Krauthgamer, and Zihan Tan. Almost-linear ϵ -emulators for planar graphs. In Stefano Leonardi and Anupam Gupta, editors, *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, pages 1311–1324. ACM, 2022. doi:10.1145/3519935.3519998.
- 19 Panagiotis Charalampopoulos, Paweł Gawrychowski, Yaowei Long, Shay Mozes, Seth Pettie, Oren Weimann, and Christian Wulff-Nilsen. Almost optimal exact distance oracles for planar graphs. *Journal of the ACM*, 70(2):1–50, 2023. doi:10.1145/3580474.
- 20 Krishnendu Chatterjee, Rasmus Ibsen-Jensen, and Andreas Pavlogiannis. Optimal tree-decomposition balancing and reachability on low treewidth graphs. *IST Austria*, 2014. doi:10.15479/AT:IST-2014-314-v1-1.
- 21 Victor Chepoi, Bertrand Estellon, and Yann Vaxès. Covering planar graphs with a fixed number of balls. *Discret. Comput. Geom.*, 37(2):237–244, 2007. doi:10.1007/S00454-006-1260-0.
- 22 Vincent Cohen-Addad, Andrew Draganov, Matteo Russo, David Saulpic, and Chris Schwegelshohn. A tight VC-dimension analysis of clustering coresets with applications. In Yossi Azar and Debmalya Panigrahi, editors, *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025, New Orleans, LA, USA, January 12-15, 2025*, pages 4783–4808. SIAM, 2025. doi:10.1137/1.9781611978322.162.
- 23 Vincent Cohen-Addad, Hung Le, Marcin Pilipczuk, and Michał Pilipczuk. Planar and minor-free metrics embed into metrics of polylogarithmic treewidth with expected multiplicative distortion arbitrarily close to 1. In *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023, Santa Cruz, CA, USA, November 6-9, 2023*, pages 2262–2277. IEEE, 2023. doi:10.1109/FOCS57990.2023.00140.
- 24 Erik D. Demaine and Mohammad Taghi Hajiaghayi. Equivalence of local treewidth and linear local treewidth and its algorithmic applications. In J. Ian Munro, editor, *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2004, New Orleans, Louisiana, USA, January 11-14, 2004*, pages 840–849. SIAM, 2004. URL: <https://dl.acm.org/doi/10.5555/982792.982919>.
- 25 Erik D. Demaine, Mohammad Taghi Hajiaghayi, and Ken-ichi Kawarabayashi. Approximation algorithms via structural results for apex-minor-free graphs. In Susanne Albers, Alberto Marchetti-Spaccamela, Yossi Matias, Sotiris E. Nikolettseas, and Wolfgang Thomas, editors, *Automata, Languages and Programming, 36th International Colloquium, ICALP 2009, Rhodes, Greece, July 5-12, 2009, Proceedings, Part I*, volume 5555 of *Lecture Notes in Computer Science*, pages 316–327. Springer, 2009. doi:10.1007/978-3-642-02927-1_27.
- 26 Guillaume Ducoffe, Michel Habib, and Laurent Viennot. Diameter, eccentricities and distance oracle computations on H-minor free graphs and graphs of bounded (distance) Vapnik–Chervonenkis dimension. *SIAM Journal on Computing*, 51(5):1506–1534, 2022. doi:10.1137/20M136551X.
- 27 Lech Duraj, Filip Konieczny, and Krzysztof Potepa. Better diameter algorithms for bounded vc-dimension graphs and geometric intersection graphs. In Timothy M. Chan, Johannes Fischer, John Iacono, and Grzegorz Herman, editors, *32nd Annual European Symposium on Algorithms, ESA 2024, September 2-4, 2024, Royal Holloway, London, United Kingdom*, volume 308 of *LIPICs*, pages 51:1–51:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICs.ESA.2024.51.
- 28 David Eppstein. Diameter and treewidth in minor-closed graph families. *Algorithmica*, 27(3):275–291, 2000. doi:10.1007/S004530010020.
- 29 Arnold Filtser. On sparse covers of minor free graphs, low dimensional metric embeddings, and other applications. *CoRR*, abs/2401.14060, 2024. doi:10.48550/arXiv.2401.14060.
- 30 Arnold Filtser and Hung Le. Low treewidth embeddings of planar and minor-free metrics. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1081–1092, 2022. doi:10.1109/FOCS54457.2022.00105.

- 31 Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, Michal Pilipczuk, and Marcin Wrochna. Fully polynomial-time parameterized computations for graphs and matrices of low treewidth. *ACM Trans. Algorithms*, 14(3):34:1–34:45, 2018. doi:10.1145/3186898.
- 32 Eli Fox-Epstein, Philip N. Klein, and Aaron Schild. Embedding planar graphs into low-treewidth graphs with applications to efficient approximation schemes for metric problems. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1069–1088. SIAM, 2019. doi:10.1137/1.9781611975482.66.
- 33 Viktor Fredslund-Hansen, Shay Mozes, and Christian Wulff-Nilsen. Truly Subquadratic Exact Distance Oracles with Constant Query Time for Planar Graphs. In Hee-Kap Ahn and Kunihiko Sadakane, editors, *32nd International Symposium on Algorithms and Computation (ISAAC 2021)*, volume 212 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 25:1–25:12, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ISAAC.2021.25.
- 34 Pawel Gawrychowski, Haim Kaplan, Shay Mozes, Micha Sharir, and Oren Weimann. Voronoi diagrams on planar graphs, and computing the diameter in deterministic $\tilde{O}(n^{5/3})$ time. *SIAM J. Comput.*, 50(2):509–554, 2021. doi:10.1137/18M1193402.
- 35 Pawel Gawrychowski, Shay Mozes, Oren Weimann, and Christian Wulff-Nilsen. Better tradeoffs for exact distance oracles in planar graphs. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 515–529. SIAM, 2018. doi:10.1137/1.9781611975031.34.
- 36 Thore Husfeldt. Computing Graph Distances Parameterized by Treewidth and Diameter. In Jiong Guo and Danny Hermelin, editors, *11th International Symposium on Parameterized and Exact Computation (IPEC 2016)*, volume 63 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 16:1–16:11, Dagstuhl, Germany, 2017. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.IPEC.2016.16.
- 37 Bart M. P. Jansen, Daniel Lokshtanov, and Saket Saurabh. A near-optimal planarization algorithm. In Chandra Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 1802–1811. SIAM, 2014. doi:10.1137/1.9781611973402.130.
- 38 Frank Kammer and Torsten Tholey. Approximate tree decompositions of planar graphs in linear time. *Theor. Comput. Sci.*, 645:60–90, 2016. doi:10.1016/J.TCS.2016.06.040.
- 39 Adam Karczmarz and Da Wei Zheng. Subquadratic algorithms in minor-free digraphs: (weighted) distance oracles, decremental reachability, and more. In Yossi Azar and Debmalya Panigrahi, editors, *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025, New Orleans, LA, USA, January 12-15, 2025*, pages 4338–4351. SIAM, 2025. doi:10.1137/1.9781611978322.147.
- 40 Ken-ichi Kawarabayashi, Philip N Klein, and Christian Sommer. Linear-space approximate distance oracles for planar, bounded-genus and minor-free graphs. In *International Colloquium on Automata, Languages, and Programming*, pages 135–146. Springer, 2011. doi:10.1007/978-3-642-22006-7_12.
- 41 Philip Klein, Serge A. Plotkin, and Satish Rao. Excluded minors, network decomposition, and multicommodity flow. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, STOC '93*, pages 682–690, New York, NY, USA, 1993. Association for Computing Machinery. doi:10.1145/167088.167261.
- 42 Kacper Kluk, Marcin Pilipczuk, Michal Pilipczuk, and Giannos Stamoulis. Faster diameter computation in graphs of bounded Euler genus. In Keren Censor-Hillel, Fabrizio Grandoni, Joël Ouaknine, and Gabriele Puppis, editors, *52nd International Colloquium on Automata, Languages, and Programming, ICALP 2025, July 8-11, 2025, Aarhus, Denmark*, volume 334 of *LIPIcs*, pages 109:1–109:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPICS.ICALP.2025.109.

- 43 Tuukka Korhonen, Michal Pilipczuk, and Giannos Stamoulis. Minor containment and disjoint paths in almost-linear time. In *65th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2024, Chicago, IL, USA, October 27-30, 2024*, pages 53–61. IEEE, 2024. doi:10.1109/FOCS61266.2024.00014.
- 44 Hung Le. Approximate distance oracles for planar graphs with subpolynomial error dependency. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 1877–1904. SIAM, 2023. doi:10.1137/1.9781611977554.CH72.
- 45 Hung Le and Christian Wulff-Nilsen. Optimal approximate distance oracle for planar graphs. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 363–374. IEEE, 2021. doi:10.1109/FOCS52979.2021.00044.
- 46 Hung Le and Christian Wulff-Nilsen. VC set systems in minor-free (di)graphs and applications. In David P. Woodruff, editor, *Proceedings of the 2024 ACM-SIAM Symposium on Discrete Algorithms, SODA 2024, Alexandria, VA, USA, January 7-10, 2024*, pages 5332–5360. SIAM, 2024. doi:10.1137/1.9781611977912.192.
- 47 Jason Li and Merav Parter. Planar diameter via metric compression. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 152–163, 2019. doi:10.1145/3313276.3316358.
- 48 Neil Robertson and Paul D Seymour. Graph minors. XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory, Series B*, 89(1):43–76, 2003. doi:10.1016/S0095-8956(03)00042-X.
- 49 Liam Roditty and Virginia Vassilevska Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing, STOC '13*, pages 515–524, New York, NY, USA, 2013. Association for Computing Machinery. doi:10.1145/2488608.2488673.
- 50 Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. k -apices of Minor-closed Graph Classes. II. Parameterized Algorithms. *ACM Trans. Algorithms*, 18(3):21:1–21:30, 2022. doi:10.1145/3519028.
- 51 Norbert Sauer. On the density of families of sets. *Journal of Combinatorial Theory, Series A*, 13(1):145–147, 1972. doi:10.1016/0097-3165(72)90019-2.
- 52 Siamak Tazari and Matthias Müller-Hannemann. Shortest paths in linear time on minor-closed graph classes, with an application to steiner tree approximation. *Discret. Appl. Math.*, 157(4):673–684, 2009. doi:10.1016/J.DAM.2008.08.002.
- 53 Mikkel Thorup. Compact oracles for reachability and approximate distances in planar digraphs. *J. ACM*, 51(6):993–1024, November 2004. doi:10.1145/1039488.1039493.
- 54 Vladimir N Vapnik and A Ya Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. In *Measures of complexity: festschrift for alexey chervonenkis*, pages 11–30. Springer, 2015. doi:10.1007/978-3-319-21852-6_3.
- 55 Oren Weimann and Raphael Yuster. Approximating the diameter of planar graphs in near linear time. *ACM Transactions on Algorithms (TALG)*, 12(1):1–13, 2015. doi:10.1145/2764910.