Faster Diameter Computation in Graphs of Bounded Euler Genus

Kacper Kluk
Institute of Informatics, University of Warsaw, Poland

Marcin Pilipczuk
Institute of Informatics, University of Warsaw, Poland

Michał Pilipczuk
Institute of Informatics, University of Warsaw, Poland

Giannos Stamoulis
Université Paris Cité, CNRS, IRIF, F-75013, Paris, France

Abstract

We show that for any fixed integer $k \geqslant 0$, there exists an algorithm that computes the diameter and the eccentricities of all vertices of an input unweighted, undirected n-vertex graph of Euler genus at most k in time

$$\mathcal{O}_k(n^{2-\frac{1}{25}}).$$

Furthermore, for the more general class of graphs that can be constructed by clique-sums from graphs that are of Euler genus at most k after deletion of at most k vertices, we show an algorithm for the same task that achieves the running time bound

$$\mathcal{O}_k(n^{2-\frac{1}{356}} \log^{6k} n).$$

Up to today, the only known subquadratic algorithms for computing the diameter in those graph classes are that of [Ducoffe, Habib, Viennot; SICOMP 2022], [Le, Wulff-Nilsen; SODA 2024], and [Duraj, Konieczny, Potępa; ESA 2024]. These algorithms work in the more general setting of K_h-minor-free graphs, but the running time bound is $\mathcal{O}_h(n^{2-c_h})$ for some constant $c_h > 0$ depending on h. That is, our savings in the exponent of the polynomial function of n, as compared to the naive quadratic algorithm, are independent of the parameter k.

The main technical ingredient of our work is an improved bound on the number of distance profiles, as defined in [Le, Wulff-Nilsen; SODA 2024], in graphs of bounded Euler genus.

2012 ACM Subject Classification Mathematics of computing → Graph algorithms; Mathematics of computing → Graphs and surfaces; Mathematics of computing → Paths and connectivity problems; Theory of computation → Fixed parameter tractability; Theory of computation → Shortest paths

Keywords and phrases Diameter, eccentricity, subquadratic algorithms, surface-embeddable graphs

Digital Object Identifier 10.4230/LIPIcs.ICALP.2025.109

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: `https://arxiv.org/abs/2502.07501` [9]

Funding K.K. and Ma.P. are supported by Polish National Science Centre SONATA BIS-12 grant number 2022/46/E/ST6/00143. This work is a part of project BOBR (Mi.P., G.S.) that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 948057). In particular, a majority of work on this manuscript was done while G.S. was affiliated with University of Warsaw.

Acknowledgements Marcin thanks Jacob Holm, Eva Rotenberg, and Erik Jan van Leeuwen for many discussions on subquadratic algorithms for diameter while his stay on sabbatical in Copenhagen.

1 Introduction

Computing the diameter of an input (undirected, unweighted) graph G is a classic computational problem that can be trivially solved in $\mathcal{O}(nm)$ time[1]. In 2013, Roditty and Vassilevska-Williams showed that this running time bound cannot be significantly improved in general: any algorithm distinguishing graphs of diameter 2 and 3 running in time $\mathcal{O}(m^{2-\varepsilon})$, for any fixed $\varepsilon > 0$, would break the Strong Exponential Time Hypothesis [16]. This motivates the search for restrictions on G that would make the problem of computing the diameter more tractable.

As shown by Cabello and Knauer [3], sophisticated orthogonal range query data structures allow near-linear diameter computation in graphs of constant treewidth. A breakthrough result by Cabello [2] showed that the diameter of an n-vertex planar graph can be computed in $\widetilde{\mathcal{O}}(n^{11/6})$ time; this complexity has been later improved by Gawrychowski, Kaplan, Mozes, Sharir, and Weimann to $\widetilde{\mathcal{O}}(n^{5/3})$ [8][2]. A subsequent line of research [5, 6, 12] generalized this result to K_h-minor-free graphs: for every integer h, there exists a constant $c_h > 0$ such that the diameter problem in n-vertex K_h-minor-free graphs can be solved in time $\mathcal{O}_h(n^{2-c_h})$. In the works [6, 12], it holds that $c_h = \Omega\left(\frac{1}{h}\right)$; so the savings tend to zero as the size of the excluded clique minor increases.

However, known lower bounds, including the one of [16], does not exclude the possibility that c_h can be made a universal constant. That is, no known lower bound refutes the following conjecture:

▶ **Conjecture 1.** *There exists a constant $c > 0$ such that, for every integer $h > 1$, the diameter problem in (unweighted, undirected) n-vertex K_h-minor-free graphs can be solved in time $\mathcal{O}_h(n^{2-c})$.*

Graphs of bounded Euler genus. Our main result is the verification of Conjecture 1 for graphs of bounded Euler genus. Furthermore, our algorithm computes also the eccentricies of all the vertices of the input graph G. Recall here that the eccentricity of a vertex $v \in V(G)$ is defined as $\mathrm{ecc}(v) \coloneqq \max_{u \in V(G)} \mathrm{dist}_G(u, v)$, where $\mathrm{dist}_G(\cdot, \cdot)$ is the distance metric in G.

▶ **Theorem 2.** *For every integers $k \geqslant 1$, there exists an algorithm that, given an (unweighted, undirected) n-vertex graph G of Euler genus at most k, runs in time $\mathcal{O}_k(n^{2-\frac{1}{25}})$ and computes the diameter of G and the eccentricity of every vertex of G.*

We remark that in [2, Section 9], Cabello briefly speculated that his approach could be also generalized to graphs embeddable on surfaces of bounded genus. However, as noted in [2], this would require significant effort, as the technique works closely on the embedding and in surfaces of higher genus, additional topological hurdles arise. In contrast, in our proof of Theorem 2 the main ingredient is an improved combinatorial bound on the number of so-called *distance profiles* [12] in graphs of bounded Euler genus. This proof uses topology only very lightly, while the rest of the argument is rather standard and topology-free. All in all, we obtain a robust methodology of approaching the problem, which, as we will see, can be also used to attack Conjecture 1 to some extent.

---

[1] We follow the convention that the vertex and the edge count of the input graph are denoted by n and m, respectively.

[2] The $\widetilde{\mathcal{O}}(\cdot)$ notation hides factors polylogarithmic in n, and the $\mathcal{O}_k(\cdot)$ notation hides factors depending on a parameter k.

To explain our bound on distance profiles, we need to recall several relevant definitions.

Let G be a graph, $R \subseteq V(G)$ be a subset of vertices, and $s_R \in R$ be a vertex in R. The *distance profile* of a vertex $u \in V(G)$ to R (relative to s_R) is the function $\mathrm{prof}_{R,s_R}[u]\colon R \to \mathbb{Z}$ defined as follows:

$$\mathrm{prof}_{R,s_R}[u](s) = \mathrm{dist}_G(u,s) - \mathrm{dist}_G(u,s_R) \qquad \text{for all } s \in R.$$

Note that provided R is connected[3], we have $\mathrm{prof}_{R,s_R}[u](s) \in \{-|R|+1, -|R|+1, \ldots, |R|-2, |R|-1\}$. Li and Parter [13] were first to study (a bit differently defined) distance profiles in the context of distance problems in planar graphs and provide a nontrivial bound on the number of possible distance profiles. Inspired by their ideas, Le and Wulff-Nilsen [12] proved that if R is connected and G is K_h-minor-free, then the set system

$$\big\{\{(s,i) \in R \times \{-|R|+1,\ldots,|R|-1\} \mid i \leqslant \mathrm{prof}_{R,s_R}[u](s)\} \;:\; u \in V(G)\big\}$$

has VC dimension at most $h-1$. Hence, by applying the Sauer-Shelah Lemma we obtain that

▶ **Theorem 3** ([12])**.** *For every integer $h \geqslant 1$, K_h-minor-free graph G, connected set $R \subseteq V(G)$, and $s_R \in R$, there are at most $\mathcal{O}_h(|R|^{2h-2})$ different distance profiles to R relative to s_R.*

The VC dimension argument applied above inevitably leads to a bound with the exponent depending on h. We show that for graphs of bounded Euler genus, the bound of Theorem 3 can be improved to a polynomial of degree independent of the parameter.

▶ **Theorem 4.** *For every integer $k \geqslant 1$, (unweighted, undirected) graph G of Euler genus at most k, connected set $R \subseteq V(G)$, and $s_R \in R$, the number of distance profiles to R relative to s_R is at most $\mathcal{O}(k^4|R|^{12})$.*

The main idea behind the proof of Theorem 4 is the following simple observation: if P is a shortest path from some $u \in V(G)$ to s_R, then, as one walks along P from u to s_R, the distance profile of the current vertex to R can only (point-wise) increase. A slightly more technical modification of this argument works for shortest paths from $u \in V(G)$ to R. This allows us to reduce the case of bounded Euler genus graphs to the planar case by cutting along a constant number of shortest-to-R paths, and analysing how the distance profiles change during such a process.

One could ask whether an improvement similar to that of Theorem 4 would be possible even in the generality of K_h-minor-free graphs. Unfortunately, it seems that Theorem 4 is the limit of such improvements. More precisely, the following simple example shows that the linear dependency on h in the exponent of the bound on the number of profiles is inevitable even in graphs of treewidth h (which are K_{h+2}-minor-free).

Let $0 < k \ll \ell$ be positive integers. Let R be a path of length ℓ and $v_1, \ldots, v_k$ be k equidistant points on R (i.e., the distance between v_i and v_{i+1} is at least $p \coloneqq \lfloor \ell/(k-1) \rfloor$). For every vector $\mathbf{a} = (a_1, \ldots, a_k) \in \{\ell, \ldots, \ell+p\}^k$, construct a vertex $u(\mathbf{a})$ and, for every $i \in \{1,\ldots,k\}$, connect it with v_i using a path of length a_i. This finishes the construction of the graph G; see Figure 1 for an illustration. Note that G has treewidth at most $k+1$, because $G - \{v_1,\ldots,v_k\}$ is a forest. Furthermore, since the distance between consecutive vertices v_i is at least p, we have that $\mathrm{dist}_G(u(\mathbf{a}), v_i) = a_i$ for every vector $\mathbf{a}$ and $i \in \{1,\ldots,k\}$.

[3] A subset of vertices R of a graph G is *connected* if the induced subgraph $G[R]$ is connected.

Consequently, if we restrict to vectors $\mathbf{a}$ with $a_1 = \ell$, every vertex $u(\mathbf{a})$ has a different distance profile to R relative to v_1. Finally, note that there are $(p+1)^{k-1} \geqslant (\ell/(k-1))^{k-1} = \Omega_k(\ell^k)$ different vectors $\mathbf{a}$ with $a_1 = \ell$, giving that many different profiles.

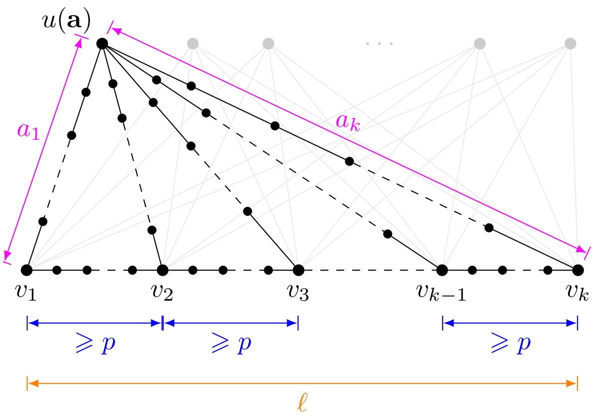

Figure 1 Illustration of a construction that shows that linear dependency on h in the exponent of the bound on the number of profiles is inevitable, even in graphs of treewidth h.

Our algorithm for Theorem 2 follows closely the approach of Le and Wulff-Nilsen [12] augmented by the bound provided by Theorem 4. Namely, we first compute an r-division of the input graph G into regions of size $r = n^{\delta}$, for some small $\delta > 0$. Then we use Theorem 4 for individual regions R to speed up the computation of distances between R and $V(G) - R$, by grouping vertices outside R according to their distance profiles to R. Each group is batch-processed in a single step.

Generalizations. Further, we show that our techniques combine well with the techniques for bounded treewidth graphs of Cabello and Knauer [3]. First, we show that Conjecture 1 holds for classes of graphs of bounded Euler genus with a constant number of *apices*, i.e., vertices that are arbitrarily connected to the rest of the graph.

▶ **Theorem 5.** *For every integers $g, k \geqslant 1$, there exists an algorithm that, given an (unweighted, undirected) n-vertex graph G and a set $A \subseteq V(G)$ such that $|A| \leqslant k$ and $G - A$ is of Euler genus at most g, runs in time $\mathcal{O}_{g,k}(n^{2-\frac{1}{25}} \log^{k-1} n)$ and computes the diameter of G and the eccentricity of every vertex of G.*

Second, we show that Conjecture 1 holds for classes of graphs constructed by clique-sums of graphs as in Theorem 5. To state this result formally, we need some definitions. For a graph G, a *tree decompostion* of G is a pair (T, β) where T is a tree and β is a function that assigns to every $t \in V(T)$ a *bag* $\beta(t) \subseteq V(G)$ such that (1) for every $v \in V(G)$, the set $\{t \in V(T) \mid v \in \beta(t)\}$ is nonempty and connected in T, and (2) for every $uv \in E(G)$ there exists $t \in V(T)$ with $u, v \in \beta(t)$. The *torso* of the bag $\beta(t)$ is constructed from $G[\beta(t)]$ by adding, for every neighbor s of t in T, all edges between the vertices of $\beta(s) \cap \beta(t)$.

▶ **Theorem 6.** *For every integer $k \geqslant 1$, there exists an algorithm with the following specification. The input consists of an (unweighted, undirected) n-vertex graph G together with a tree decomposition (T, β) of G and a set $A(t) \subseteq \beta(t)$ for every $t \in V(T)$ satisfying the following properties:*

- *For every node $t \in V(T)$, we have that $|A(t)| \leqslant k$ and the torso of $\beta(t)$ with the vertices of $A(t)$ deleted is a graph of Euler genus at most k.*

- *For every edge $st \in E(T)$, we have $|\beta(s) \cap \beta(t)| \leqslant k$.*

The algorithm runs in time $\mathcal{O}_k(n^{2-\frac{1}{356}} \log^{6k} n)$ and computes the diameter of G and the eccentricity of every vertex of G.

Note that the statements of Theorems 5 and 6 require the set A and the decomposition (T, β), respectively, to be provided explicitly on input; this should be compared with more general statements where the algorithm is given only G with a promise that such set A or decomposition (T, β) exist. At this point, we are not aware of any existing algorithm that would find in subquadratic time a set A as in Theorem 5, or the decomposition (T, β) with the sets A as in Theorem 6, even in the approximate sense. However, we were informed by Korhonen, Pilipczuk, Stamoulis, and Thilikos [11] that it seems likely that the techniques introduced in the recent almost linear-time algorithm for minor-testing [10] could be used to construct such an algorithm, with almost linear time complexity. With this result in place, the assumption about the decomposition and/or apex sets being provided on input could be lifted in Theorems 5 and 6; this is, however, left to future work.

Discussion. As one of the main outcomes of their theory of graph minors, Robertson and Seymour proved the following Structure Theorem [15]: every K_h-minor-free graph G admits a tree decomposition (T, β) such that

- for every pair s, t of adjacent nodes of T, the set $\beta(t) \cap \beta(s)$ has size $\mathcal{O}_h(1)$; and
- the torso of every bag $\beta(t)$ is "nearly embedable" into a surface of bounded (in terms of h) Euler genus.

The notion of being "nearly embeddable" encompasses adding a constant number of apices (which can be handled by Theorem 6) and a constant number of so-called vortices (which are not handled by Theorem 6). Thus, our methods fall short of verifying Conjecture 1 in full generality due to vortices.

We remark that recently, Thilikos and Wiederrecht [19] proved a variant of the Structure Theorem, where under the stronger assumption of excluding a minor of a *shallow vortex grid*, instead of a clique minor, they gave a decomposition as above, but with torsos devoid of vortices. Thus, the decomposition for shallow-vortex-grid-minor-free graphs provided by [19] can be directly plugged into Theorem 6, with the caveat that [19] does not provide a subquadratic algorithm to compute the decomposition.

Coming back to Conjecture 1, the simplest case that we are currently unable to solve is the setting when the input is a planar graph plus a single vortex. More formally, for a fixed integer k, let $\mathcal{G}_k$ be the class of graphs defined as follows. We have $G \in \mathcal{G}_k$ if there exist two subgraphs G_0, G_1 of G and a sequence of vertices $v_1, \ldots, v_b$ in $V(G_0) \cap V(G_1)$ such that:

- $V(G) = V(G_0) \cup V(G_1)$,
- $E(G) = E(G_0) \cup E(G_1)$,
- G_0 admits a planar embedding where the vertices $v_1, \ldots, v_b$ lie on one face in this order, and
- G_1 admits a tree decomposition (T_1, β_1), where T_1 is a path on nodes $t_1, \ldots, t_b$ and for every $i \in \{1, \ldots, b\}$, the bag $\beta_1(t_i)$ contains v_i and is of size at most k.

It is easy to see that graphs from $\mathcal{G}_k$ are $K_{k+\mathcal{O}(1)}$-minor-free. Do they satisfy Conjecture 1? That is, is there a constant $c > 0$ such that the diameter problem in $\mathcal{G}_k$ can be solved in time $\mathcal{O}_k(n^{2-c})$?

Organization. We prove Theorem 4 in Section 3. Theorem 5 is proven in Section 4; note that Theorem 2 follows from Theorem 5 for $k = 1$. The proof of Theorem 6 can be found in the full version of the paper [9].

2 Preliminaries

Set systems and VC-dimension. A *set system* is a collection $\mathcal{F}$ of subsets of a given set A, which we call *ground set* of $\mathcal{F}$. 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, the intersections of Y and the sets in $\mathcal{F}$ contain every subset of Y. The *VC-dimension* of a set system $\mathcal{F}$ with ground set A is the size of the largest subset $Y \subseteq A$ shattered by $\mathcal{F}$. The notion of VC-dimension was introduced by Vapnik and Chervonenkis [20].

We will use the following well-known Sauer-Shelah Lemma [17, 18], which gives a polynomial upper bound on the size of a set system of bounded VC-dimension.

▶ **Lemma 7** (Sauer-Shelah Lemma). *Let $\mathcal{F}$ be a set system with ground set A. If the VC-dimension of $\mathcal{F}$ is at most k, then $|\mathcal{F}| = \mathcal{O}(|A|^k)$.*

Basic graph notation. All our graphs are undirected. For a graph G, the neighborhood of a vertex u is defined as $N_G(u) = \{v : uv \in E(G)\}$ and for $X \subseteq V(G)$ we have $N_G(X) = \bigcup_{u \in X} N_G(u) - X$.

The *length* of a path P, denoted $|P|$, is the number of edges of P. For two vertices u, v of a graph G, the *distance* between u and v, denoted $\mathrm{dist}_G(u, v)$, is defined as the minimum length of a path in G with endpoints u and v. For every $v \in V(G)$ and set $R \subseteq V(G)$, we set $\mathrm{dist}_G(v, R) \coloneqq \min\{\mathrm{dist}_G(v, y) \colon y \in R\}$. For vertices x, y appearing on a path P, by $P[x, y]$ we denote the subpath of P with endpoints x and y. The set of vertices traversed by a path P is denoted by $V(P)$. In all above notation, we sometimes drop the subscript if the graph is clear from the context.

For a nonnegative integer q, we use the shorthand $[q] \coloneqq \{1, \ldots, q\}$. For a vertex $v \in V(G)$ and a set $X \subseteq V(G)$, we define the *X-eccentricity* of v as $\mathrm{ecc}_X(v) \coloneqq \max_{x \in X} \mathrm{dist}(v, x)$. Thus, the eccentricity of v in G is the same as its $V(G)$-eccentricity.

The *Euler genus* of a graph G is the minimum Euler characteristic of a surface, where G is embeddable. We refer to the textbook of Mohar and Thomassen for more on surfaces and embedded graphs [14].

We will use the following result of Le and Wulff-Nilsen [12, Theorem 1.3] for planar graphs. Note that the set R is not necessarily connected.

▶ **Theorem 8.** *Let $h \geqslant 1$ be an integer, G be a K_h-minor-free (unweighted, undirected) graph, R be a subset of $V(G)$, and $s_R \in R$. Then the set system*

$$\{\{(s, i) \in R \times \mathbb{Z} \mid i \leqslant \mathrm{dist}_G(u, s) - \mathrm{dist}_G(u, s_R)\} \ : \ u \in V(G)\}$$

has VC-dimension at most $h - 1$.

Algorithmic tools. All our algorithms assume the word RAM model.

To cope with apices, we will need the following classic data structure due to Willard [21].

▶ **Theorem 9** ([21]). *Let V be a set of n points in $\mathbb{R}^d$ and let $w \colon V \to \mathbb{R}$ be a weight function. By a* suffix range*, we mean any set of the form*

$$\mathsf{Range}(r_1, \ldots, r_d) \coloneqq \{(x_1, \ldots, x_d) \in \mathbb{R}^d \ | \ x_i \geqslant r_i \text{ for all } i \in [d]\}$$

for some range parameters $r_1, \ldots, r_d \in \mathbb{R}$.

There is a data structure that uses $\mathcal{O}\left(n \log^{d-1} n\right)$ preprocessing time, $\mathcal{O}\left(n \log^{d-1} n\right)$ memory and answers the following suffix range queries in time $\mathcal{O}\left(\log^{d-1} n\right)$: given a tuple $(r_i)_{i \in [d]}$, find the maximum value of $w(v)$ over all $v \in V \cap \mathsf{Range}(r_1, \ldots, r_d)$.

We will also need the following standard statement about r-divisions.

▶ **Theorem 10** ([22]). *Let G be a K_t-minor-free graph on n vertices. For any fixed constant $\varepsilon > 0$, and for any parameter r with $Ct^2 \log n \leqslant r \leqslant n$, where C is some absolute constant, we can construct in time $\mathcal{O}\left(n^{1+\varepsilon}\sqrt{r}\right)$ an r-division of G, that is, a collection $\mathcal{R}$ of connected subsets of vertices of G such that:*
- $\bigcup \mathcal{R} = V(G)$,
- $|R| \leqslant r$ *for every* $R \in \mathcal{R}$, *and*
- $\sum_{R\in\mathcal{R}} |\partial R| \leqslant \mathcal{O}(nt/\sqrt{r})$, *where* $\partial R = R \cap N_G(V(G) - R)$.

3 Distance profiles in graphs of bounded Euler genus

In this section we prove Theorem 4. Our argument consists of a reduction to the planar case, where we can use the constant bound on the VC-dimension of the set system given by the distance profiles due to Le and Wulff-Nilsen [12]. The main idea behind the reduction is to consider certain notions of "extended" profiles, where the extension is built along a collection of shortest paths. These shortest paths can be chosen in such a way that by cutting the graph along these paths we obtain a plane graph. Then a bound on the number of the extended profiles in the obtained plane graph translates to a bound on the number of (standard) distance profiles in the original graph.

Preliminary definitions and results needed for defining profiles with respect to shortest paths are given in Section 3.1. These extended profiles are then defined in Section 3.2. There, we also prove that a fundamental lemma that equality of extended profiles entails equality of (standard) distance profiles. The main reduction providing the proof of Theorem 4 is given at the end of this section.

3.1 Milestones

Let G be a graph, R be a subset of $V(G)$, v_0 be a vertex in $V(G)$, and P be a shortest path from v_0 to R. Let x be the endpoint of P in R. Further, let $\leqslant_P$ be the linear ordering of the vertices traversed by P: for two vertices $v, u \in V(P)$, we have $v \leqslant_P u$ if u belongs to $P[v, x]$. We say that a vertex $v \in V(P)$ is a *milestone of* P if either $v = x$ or we have $\text{prof}_{R,x}[v] \neq \text{prof}_{R,x}[u]$, where u is the successor of v in $\leqslant_P$. We denote by $M_R(P)$ the set of all milestones of P. Given a milestone $v \in M_R(P)$, the *neutral prefix of v in P* is defined as the vertex set of the maximal subpath Q of $P[v_0, v]$ satisfying the following: v is the only milestone of P that belongs to Q.

The next lemma shows that minimum-length paths towards R that contain a vertex in the neutral prefix of a milestone can be assumed to pass through that milestone vertex.

▶ **Lemma 11.** *Let G be a graph, R be a subset of $V(G)$, v_0 be a vertex in $V(G)$ and P be a shortest path from v_0 to R. Then for every $v \in M_R(P)$, every u in the neutral prefix of v, and every $y \in R$, it holds that* $\text{dist}(u, y) = |P[u, v]| + \text{dist}(v, y)$.

Proof. Let x be the endpoint of P in R. Note that, by definition, $\text{prof}_{R,x}[v] = \text{prof}_{R,x}[u]$. Also, $\text{dist}(u, x) = \text{dist}(u, v) + \text{dist}(v, x)$ and $\text{dist}(u, v) = |P[u, v]|$. Therefore, $\text{dist}(u, y) = |P[u, v]| + \text{dist}(v, y)$ for every $y \in R$. ◀

We also give an upper bound on the number of milestones.

▶ **Lemma 12.** *Let G be a graph, R be a connected subset of $V(G)$, v_0 be a vertex of G, and P be a shortest path from v_0 to R. Then the number of milestones of P is at most* $|R|^2 - |R| + 1$.

Proof. Let x be the endpoint of P in R. First observe that since P is a shortest path from v_0 to R, we have $\mathrm{dist}(v, y) \geqslant \mathrm{dist}(v, x)$ for every $v \in V(P)$ and every $y \in R$; hence $\mathrm{prof}_{R,x}[v](y) \geqslant 0$. Also, since R is connected, for every $y \in R$ we have $\mathrm{prof}_{R,x}[x](y) \leqslant |R|-1$. To conclude the proof, it suffices to prove that for all $v_1, v_2 \in V(P)$ with $v_1 \leqslant_P v_2$, we have

$$\mathrm{prof}_{R,x}[v_1](y) \leqslant \mathrm{prof}_{R,x}[v_2](y) \qquad \text{for all } y \in R. \tag{1}$$

Indeed, (1) together with the previous observations shows that all the distinct distance profiles of the form $\mathrm{prof}_{R,x}[v]$ for $v \in V(P)$ can be treated as vectors of length $|R|$ with entries in $\{0, \ldots, |R|-1\}$, and they all have distinct sums $\sum_{y \in R} \mathrm{prof}_{R,x}[v](y)$. Since these sums range between 0 and $|R|^2 - |R|$, the total number of distinct profiles is at most $|R|^2 - |R| + 1$, implying the same bound on the number of milestones.

To see why (1) holds, note that $\mathrm{dist}(v_1, y) \leqslant \mathrm{dist}(v_1, v_2) + \mathrm{dist}(v_2, y)$ implies that

$$\mathrm{dist}(v_1, y) \leqslant \mathrm{dist}(v_1, v_2) + \mathrm{prof}_{R,x}[v_2](y) + \mathrm{dist}(v_2, x) = \mathrm{prof}_{R,x}[v_2](y) + \mathrm{dist}(v_1, x);$$

the last equality follows from P being a shortest path containing v_1, v_2, and x (in this order). This in turn implies that $\mathrm{prof}_{R,x}[v_1](y) = \mathrm{dist}(v_1, y) - \mathrm{dist}(v_1, x) \leqslant \mathrm{prof}_{R,x}[v_2](y)$, as claimed. ◀

3.2 Anchor-distance profiles

Shortest path collections. Let G be a graph and R be a subset of vertices of G. We say that a collection $\mathcal{P}$ of paths in G is an *R-shortest path collection* if

- every $P \in \mathcal{P}$ is a shortest path from some $v^P \in V(G)$ to R, i.e., $|P| = \mathrm{dist}(v^P, R)$; and
- $R \subseteq \bigcup_{P \in \mathcal{P}} V(P)$.

For each $P \in \mathcal{P}$, we denote by x^P the endpoint of P in R. Note that the collection $\mathcal{P}$ obtained by taking, for every $y \in R$, the zero-length path from y to y, is an R-shortest path collection.

We say that an R-shortest path collection is *consistent* if, for every $P_1, P_2 \in \mathcal{P}$ and $v \in V(P_1) \cap V(P_2)$ it holds that $P_1[v, x^{P_1}] = P_2[v, x^{P_2}]$. That is, once two paths intersect, they continue together towards R.

The following statement is a direct consequence of the definition of an R-shortest path collection.

▶ **Observation 13.** *Let G be a graph, R be a subset of vertices of G, and $\mathcal{P}$ be an R-shortest path collection. Then for every two paths $P_1, P_2 \in \mathcal{P}$ and every $v \in V(P_1) \cap V(P_2)$, we have $|P_1[v, x^{P_1}]| = |P_2[v, x^{P_2}]|$.*

Anchor vertices and their prefixes. Let G be a graph, R be a subset of $V(G)$, and $\mathcal{P}$ be an R-shortest path collection. We denote by $H_{\mathcal{P}}$ the union of the paths in $\mathcal{P}$, i.e., the graph $(\bigcup_{P \in \mathcal{P}} V(P), \bigcup_{P \in \mathcal{P}} E(P))$. We say that a vertex is an *anchor vertex* if either it has degree more than two in $H_{\mathcal{P}}$ or it is a milestone of a path $P \in \mathcal{P}$. We denote by $A_R(P)$ the set of all anchor vertices lying on a path $P \in \mathcal{P}$ and by $A_R(\mathcal{P})$ the set of all anchor vertices for $\mathcal{P}$. Given a path $P \in \mathcal{P}$ with endpoints v_0 and $y \in R$, and an anchor vertex $w \in A_R(P)$, the *prefix of w in P* is the vertex set of the maximal subpath Q of $P[v_0, w]$ satisfying the following: w is the only anchor vertex of P that belongs to Q. Note that for every anchor $w \in V(P)$ there is a milestone w' of P (possibly $w = w'$) such that the prefix of w in P is a subset of the neutral prefix of w' in P. Finally, for an anchor vertex w, the *tail* of w, denoted $\mathrm{tail}(w)$, is the subgraph of G consisting of the union of all prefixes of w in P over all paths $P \in \mathcal{P}$ that contain w.

Hat-distances. Let G be a graph, R be a subset of vertices of G, and $\mathcal{P}$ be an R-shortest path collection. We denote by

$$U_{\mathcal{P}} \coloneqq V(G) - \bigcup_{P\in\mathcal{P}} V(P).$$

For every $u \in U_{\mathcal{P}}$, and every anchor vertex $w \in A_R(\mathcal{P})$, we set

$$\widehat{\mathrm{dist}}(u,w) \coloneqq \min\{|Q_{u,z}| + |P[z,w]| \colon P \in \mathcal{P} \wedge w \in V(P) \wedge z \text{ is in the prefix of } w \text{ in } P\},$$

where $Q_{u,z}$ is a shortest path from u to z with all its internal vertices in $U_{\mathcal{P}}$. If such $Q_{u,z}$ does not exist for any $z \in V(\mathrm{tail}(w))$, we set $\widehat{\mathrm{dist}}(u,w) \coloneqq \infty$.

The following statement is a direct consequence of the definition of $\widehat{\mathrm{dist}}(\cdot,\cdot)$.

▶ **Observation 14.** *Let G be a graph, R be a subset of vertices of G, and $\mathcal{P}$ be an R-shortest path collection. Then for every $u \in U_{\mathcal{P}}$, we have that*

$$\mathrm{dist}(u,R) = \min\left\{\widehat{\mathrm{dist}}(u,w) + \mathrm{dist}(w,R) \colon w \in A_R(\mathcal{P})\right\}.$$

Anchor-distance profiles. Let G be a graph, R be a subset of vertices of G, and $\mathcal{P}$ be an R-shortest path collection. The *anchor-distance profile* of a vertex $u \in U_{\mathcal{P}}$ to R with respect to $\mathcal{P}$ is a function $\mathrm{prof}^{\star}_{R,\mathcal{P}}[u]$ mapping each $w \in A_R(\mathcal{P})$ to

$$\mathrm{prof}^{\star}_{R,\mathcal{P}}[u](w) \coloneqq \widehat{\mathrm{dist}}(u,w) + \mathrm{dist}(w,R) - \mathrm{dist}(u,R).$$

Observation 14 implies that we always have $\mathrm{prof}^{\star}_{R,\mathcal{P}}[u](w) \geqslant 0$. We set

$$\widehat{\mathrm{prof}}_{R,\mathcal{P}}[u](w) \coloneqq \min\{\mathrm{prof}^{\star}_{R,\mathcal{P}}[u](w), |R|+1\}.$$

▶ **Lemma 15.** *Let G be a graph, let R be a connected subset of vertices of G, and $s_R \in R$. Also, let $\mathcal{P}$ be an R-shortest path collection. Then for all $u_1, u_2 \in U_{\mathcal{P}}$,*

$$\widehat{\mathrm{prof}}_{R,\mathcal{P}}[u_1] = \widehat{\mathrm{prof}}_{R,\mathcal{P}}[u_2] \qquad \textit{implies} \qquad \mathrm{prof}_{R,s_R}[u_1] = \mathrm{prof}_{R,s_R}[u_2].$$

Proof. Fix $u_1, u_2 \in U_{\mathcal{P}}$ with $\widehat{\mathrm{prof}}_{R,\mathcal{P}}[u_1] = \widehat{\mathrm{prof}}_{R,\mathcal{P}}[u_2]$. We start by proving the following.

▷ Claim 16. Let $u \in U_{\mathcal{P}}$ and $y \in R$. There is an anchor $w \in A_R(\mathcal{P})$ such that

- $\widehat{\mathrm{dist}}(u,w) + \mathrm{dist}(w,y) = \mathrm{dist}(u,y)$ and
- $\widehat{\mathrm{prof}}_{R,\mathcal{P}}[u](w) \leqslant |R|$.

Proof. Let Q be a shortest path from u to y and let $P \in \mathcal{P}$ be the path which Q first intersects (if the first vertex of Q in $\bigcup_{P\in\mathcal{P}} V(P)$ belongs to more than one paths in $\mathcal{P}$, we choose P arbitrarily among these paths). Also, let u' be the first vertex of Q (when ordering from u to y) in $V(P)$ and w be the anchor of P that contains u' in its prefix (in P). Note that $u' \in V(\mathrm{tail}(w))$.

We first show that

$$\widehat{\mathrm{dist}}(u,w) + \mathrm{dist}(w,y) = \mathrm{dist}(u,y). \tag{2}$$

By Lemma 11 and the fact that $|Q[u',y]| = \mathrm{dist}(u',y)$, we have

$$\mathrm{dist}(w,y) = |Q[u',y]| - |P[u',w]|. \tag{3}$$

Also, by definition, we have

$$\widehat{\text{dist}}(u, w) \leqslant |Q[u, u']| + |P[u', w]|. \tag{4}$$

By (3) and (4), we get that $\widehat{\text{dist}}(u, w) + \text{dist}(w, y) \leqslant |Q|$. Moreover, since Q is a shortest path from u to y and $\widehat{dist}(u, w) \geqslant \text{dist}(u, w)$, we have

$$|Q| = \text{dist}(u, y) \leqslant \text{dist}(u, w) + \text{dist}(w, y) \leqslant \widehat{dist}(u, w) + dist(w, y).$$

This proves (2).

Next, we show that $\widehat{\text{prof}}_{R,\mathcal{P}}[u](w) \leqslant |R|$. Note that

$$\begin{aligned}\text{prof}^{\star}_{R,\mathcal{P}}[u](w) + \text{dist}(u, R) &= \widehat{\text{dist}}(u, w) + \text{dist}(w, R) \\ &\leqslant \widehat{\text{dist}}(u, w) + \text{dist}(w, y) = \text{dist}(u, y).\end{aligned}$$

The connectivity of R implies that $\text{dist}(u, y) \leqslant \text{dist}(u, R) + |R|$, which gives $\text{prof}^{\star}_{R,\mathcal{P}}[u](w) \leqslant |R|$, and the claim follows. ◁

We next show that there is an integer c such that for every $y \in R$, we have

$$\text{dist}(u_1, y) = \text{dist}(u_2, y) + c.$$

Note that this will immediately imply that $\text{prof}_{R,s_R}[u_1] = \text{prof}_{R,s_R}[u_2]$.

By Observation 14, for every $h \in \{1, 2\}$, there is an anchor $w_h \in A_R(\mathcal{P})$ such that $\text{dist}(u_h, R) = \widehat{\text{dist}}(u_h, w_h) + \text{dist}(w_h, R)$, which is equivalent to $\text{prof}^{\star}_{R,\mathcal{P}}[u_h](w_h) = 0$. If w_h lies on $P_h \in \mathcal{P}$, then $\text{dist}(u_h, R) = \text{dist}(u_h, x^{P_h})$. Therefore, as $\text{prof}^{\star}_{R,\mathcal{P}}[u_1] = \text{prof}^{\star}_{R,\mathcal{P}}[u_2]$, we can choose $w_1 = w_2$ and $P_1 = P_2$, hence $x^{P_1} = x^{P_2}$. In other words, there exists $x \in R$ such that $\text{dist}(u_1, R) = \text{dist}(u_1, x)$ and $\text{dist}(u_2, R) = \text{dist}(u_2, x)$. We set $c \coloneqq \text{dist}(u_1, x) - \text{dist}(u_2, x) = \text{dist}(u_1, R) - \text{dist}(u_2, R)$.

Now, fix $y \in R$. Let $w_1 \in A_R(P)$ be the anchor from Claim 16 (applied for u_1 and y). As $\widehat{\text{prof}}_{R,\mathcal{P}}[u_1] = \widehat{\text{prof}}_{R,\mathcal{P}}[u_2]$ and $\text{prof}^{\star}_{R,\mathcal{P}}[u_1](w_1) \leqslant |R|$, we get that $\text{prof}^{\star}_{R,\mathcal{P}}[u_1](w_1) = \text{prof}^{\star}_{R,\mathcal{P}}[u_2](w_1)$, i.e.,

$$\widehat{\text{dist}}(u_1, w_1) + \text{dist}(w_1, R) - \text{dist}(u_1, R) = \widehat{\text{dist}}(u_2, w_1) + \text{dist}(w_1, R) - \text{dist}(u_2, R).$$

Therefore,

$$\begin{aligned}\text{dist}(u_1, y) &= \widehat{\text{dist}}(u_1, w_1) + \text{dist}(w_1, y) \\ &= \widehat{\text{dist}}(u_2, w_1) + \text{dist}(w_1, y) + c \geqslant \text{dist}(u_2, y) + c;\end{aligned}$$

the first equality follows from Claim 16. Thus $\text{dist}(u_2, y) + c \leqslant \text{dist}(u_1, y)$. A symmetric reasoning shows that also $\text{dist}(u_1, y) - c \leqslant \text{dist}(u_2, y)$. Therefore we get $\text{dist}(u_1, y) = \text{dist}(u_2, y) + c$, as required. ◀

3.3 Reduction from bounded genus graphs to planar graphs

We next recall several definitions related to embeddings of graphs on surfaces. Our basic terminology follows [14]. We say that a graph H embedded in a surface Σ is a *simple cut-graph* of Σ if H has a single face that is also homeomorphic to an open disk; equivalently, H has a single facial walk. Given a surface Σ and a simple cut-graph H on Σ, we denote by $\Sigma \rtimes H$ the surface obtained by cutting Σ along H. Note that, provided H is a simple cut-graph, $\Sigma \rtimes H$ is always a disk.

Suppose now that a graph G embedded on Σ and H is a subgraph of G that is a simple cut-graph of H. We define $G \text{✂} H$ to be the graph embedded on $\Sigma \text{✂} H$ obtained from G as follows. First, let σ be the (unique) facial walk of H and note that each edge e of H is contained exactly twice in σ and each vertex v of H is contained in σ as many times as the degree of v in H. To obtain $G \text{✂} H$, we replace H with a simple cycle C_σ whose vertex set is the set of copies of vertices of H and its edge set is the set of copies of edges of H in the obvious way. Notice that σ also prescribes for every edge uv of G between a vertex $u \in V(G) - V(H)$ and a vertex $v \in V(H)$, to which copy of v in $G \text{✂} H$ the vertex u should be adjacent to (in $G \text{✂} H$). See Figure 2 for an illustration.

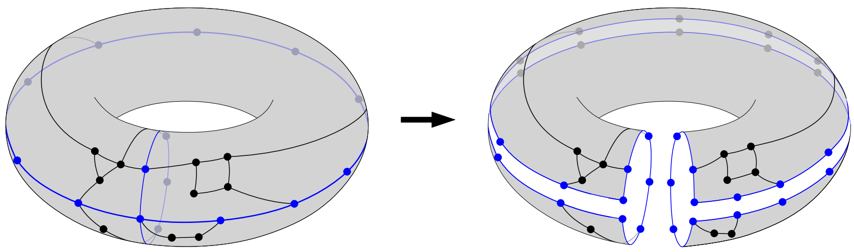

Figure 2 Left: A graph G embedded on a surface Σ and a subgraph H of G (in blue) that is a simple cut-graph of Σ. Right: The graph $G \text{✂} H$ embedded on the surface $\Sigma \text{✂} H$ (which is homeomorphic to a disk); the blue vertices/edges are copies of the vertices/edges of H.

We will use the following well-known result which appears in the literature under different formulations; see e.g. [1, 4, 7].

▶ **Lemma 17.** *For every integer $k \geqslant 1$ and for every edge-weighted connected graph G embedded on a surface Σ of Euler genus at most k and every vertex $u \in V(G)$, there is a subgraph H of G with the following properties:*

- *H is a simple cut-graph of Σ, and*
- *$V(H)$ is the union of the vertex sets of $\mathcal{O}(k)$ shortest paths in G that have u as a common endpoint.*

We are now ready to proceed to the proof of Theorem 4. Employing Lemma 15, we aim at bouding the VC-dimension of the set system defined by the anchor-distance profiles. This is can be done by a suitable reduction to the planar setting using Lemma 17.

Proof of Theorem 4. We assume that G is connected – the distance profiles of all vertices that are not connected to R are equal. Let T_R be a spanning tree of $G[R]$ and let G_0 be the graph obtained from G after contracting T_R into a single vertex v_R. Consider an embedding of G_0 on a surface Σ of Euler genus at most k. By Lemma 17, there is a subgraph H_0 of G_0 that is a simple cut-graph of Σ and a family $\mathcal{P}_0$ of $\mathcal{O}(k)$ shortest paths in G_0, each with v_R as an endpoint, such that $V(H_0) = \bigcup_{P \in \mathcal{P}_0} V(P)$. Furthermore, as Lemma 17 handles edge weights, we can slightly perturb the weights so that shortest paths in G_0 are unique and, in particular, all shortest paths with one endpoint in v_R form a tree. Since H_0 is a simple cut-graph of Σ, $G_0 \text{✂} H_0$ is disk-embedded. Uncontracting T_R, we get a subgraph H of G such that $G \text{✂} H$ is disk-embedded. Let $\mathcal{P}$ be the family of projections of the paths of $\mathcal{P}_0$ onto G plus, for every $y \in R$, a zero-length path from y to y. Hence, $\mathcal{P}$ is an R-shortest paths collection of size $\mathcal{O}(k)$ with $V(\mathcal{P}) = V(H)$. Furthermore, since in G_0 the paths of $\mathcal{P}_0$ formed a tree rooted at v_R, $\mathcal{P}$ is consistent.

Note that due to Lemma 12 we have that $\sum_{P\in\mathcal{P}} |M_R(P)| \leqslant \mathcal{O}(k|R|^2)$. Also, since $\mathcal{P}$ is consistent, if B are the vertices that are not in R (recall that vertices in R are milestones) and have degree more than two in the graph obtained by the union of the paths in $\mathcal{P}$, then $|B| \leqslant |\mathcal{P}| - 1$. Hence,

$$\sum_{P\in\mathcal{P}} |A_R(P)| \leqslant \mathcal{O}(k|R|^2). \tag{5}$$

We set $\mathcal{T}$ be the set of all vertices of $G \mathbin{\text{✂}} H$ that are copies of the anchor vertices $A_R(\mathcal{P})$. Since $|\mathcal{P}| = \mathcal{O}(k)$, we can strengthen (5) to

$$|\mathcal{T}| = \mathcal{O}(k|R|^2). \tag{6}$$

For $s \in \mathcal{T}$, let $w(s) \in A_R(\mathcal{R})$ be the anchor vertex whose copy (in $G \mathbin{\text{✂}} H$) is s. In the other direction, for $w \in A_R(\mathcal{R})$, let $S(w)$ be the set of copies of w in $G \mathbin{\text{✂}} H$.

Let U be the set of vertices of $G \mathbin{\text{✂}} H$ that are *not* copies of vertices from H (i.e., $U = V(G) - V(H)$). We set E_{out} be the set of all edges uv of $G \mathbin{\text{✂}} H$ where $u \in U$ and v is a copy of a vertex from H, i.e., $v \in V(G \mathbin{\text{✂}} H) - U$. We also set E_{next} be the set of all edges uv of $G \mathbin{\text{✂}} H$ where u is a copy of an anchor vertex $w \in A_R(P)$ for some $P \in \mathcal{P}$ and v is a copy of the neighbor of w in P that *is not* in the prefix of w in P.

Let now $\widehat{G}$ be the graph obtained from $G \mathbin{\text{✂}} H$ after the following modifications:

- we subdivide $|V(G)|$-many times each edge in $E_{\mathsf{out}} \cup E_{\mathsf{next}}$,
- we introduce a new vertex t and add, for every $s \in \mathcal{T}$, a path between t and s of length

$$d_{w(s),t} \coloneqq |V(G)| + \mathrm{dist}_G(w(s), R).$$

See Figure 3. Observe that since $G \mathbin{\text{✂}} H$ is disk-embedded, $\widehat{G}$ is planar, because we may embed t together with all the added paths outside of the disk containing $G \mathbin{\text{✂}} H$.

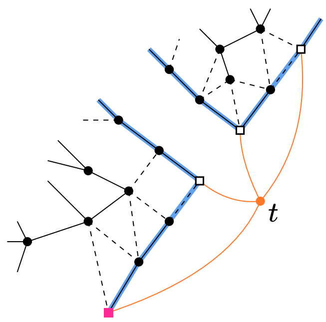

Figure 3 An illustration of (a part of) the construction of the graph $\widehat{G}$. The squared vertices are copies of anchor vertices. The marked squared vertex is also a copy of a vertex in R. The highlighted edges are copies of edges of H in $G \mathbin{\text{✂}} H$, while the paths obtained by subdividing the edges of $E_{\mathsf{out}} \cup E_{\mathsf{next}}$ are depicted with dashed edges. Edges adjacent to t correspond to paths of appropriate length.

For every $u \in U$, we define a function $\pi[u]$, mapping every $w \in A_R(\mathcal{P})$ to

$$\pi[u](w) \coloneqq \min\{\mathrm{dist}_{\widehat{G}}(u, s) : s \in S(w)\} + d_{w,t} - \mathrm{dist}_{\widehat{G}}(u, t).$$

Also, we set $\widehat{\mathcal{X}} \coloneqq \{\widehat{X}_u \mid u \in U\}$, where for $u \in U$,

$$\widehat{X}_u \coloneqq \{(w, i) \in A_R(\mathcal{P}) \times \{0, \ldots, |R| + 1\} \mid i \leqslant \pi[u](w)\}.$$

▷ Claim 18. The set system $\widehat{\mathcal{X}}$ has size $\mathcal{O}(k^4|R|^{12})$.

Proof. We set $\mathcal{T}^+ \coloneqq \mathcal{T} \cup \{t\}$. We start with the set system $\mathcal{C}^1 \coloneqq \{C_u^1 \ : \ u \in U\}$, where

$$C_u^1 \coloneqq \left\{(s,i) \in \mathcal{T}^+ \times \mathbb{Z} \mid i \leqslant \mathrm{dist}_{\widehat{G}}(u,s) - \mathrm{dist}_{\widehat{G}}(u,t)\right\}.$$

As $\widehat{G}$ is planar, by Theorem 8 we infer that $\mathcal{C}$ has VC-dimension at most 4.

We now "shift columns" of $\mathcal{C}^1$. That is, define $\mathcal{C}^2 \coloneqq \{C_u^2 \ : \ u \in U\}$, where

$$C_u^2 \coloneqq \left\{(s,i) \in \mathcal{T}^+ \times \mathbb{Z} \mid i \leqslant \mathrm{dist}_{\widehat{G}}(u,s) + d_{w(s),t} - \mathrm{dist}_{\widehat{G}}(u,t)\right\}.$$

Clearly, the VC-dimension of $\mathcal{C}^1$ and $\mathcal{C}^2$ are equal: a set $Z \subseteq \mathcal{T}^+ \times \mathbb{Z}$ shatters $\mathcal{C}^1$ if and only if the set $\{(s, d_{w(s),t} + i) \ : \ (s,i) \in Z\}$ shatters $\mathcal{C}^2$.

Now, let $\mathcal{C}^3$ be "cropped" $\mathcal{C}^2$: $\mathcal{C}^3 \coloneqq \{C_u^3 \ : \ u \in U\}$, where

$$C_u^3 \coloneqq C_u^2 \cap \left(\mathcal{T}^+ \times \{0,\ldots,|R|+1\}\right).$$

Since restricting to a smaller universe cannot increase VC-dimension, $\mathcal{C}^3$ has VC-dimension at most 4. Since $|\mathcal{T}^+| = \mathcal{O}(k|R|^2)$, by Sauer-Shelah lemma (Lemma 7) we have $|\mathcal{C}^3| = \mathcal{O}(k^4|R|^{12})$.

Now observe that for every $u_1, u_2 \in U$

$$C_{u_1}^3 = C_{u_2}^3 \qquad \text{implies} \qquad \widehat{X}_{u_1} = \widehat{X}_{u_2}. \tag{7}$$

Indeed, the assumption $C_{u_1}^3 = C_{u_2}^3$ implies that for every $w \in A_R(\mathcal{P})$ and $s \in S(W)$ we have

$$\begin{aligned}&\max(0, \min(|R|+1, \mathrm{dist}_{\widehat{G}}(u_1,s) + d_{w,t} - \mathrm{dist}_{\widehat{G}}(u_1,t)))\\ &= \max(0, \min(|R|+1, \mathrm{dist}_{\widehat{G}}(u_2,s) + d_{w,t} - \mathrm{dist}_{\widehat{G}}(u_2,t))).\end{aligned}$$

For fixed $w \in A_R(\mathcal{P})$, we take a minimum of the above expression over all $s \in S(w)$, obtaining:

$$\begin{aligned}&\max(0, \min(|R|+1, \min\{\mathrm{dist}_{\widehat{G}}(u_1,s) \ : \ s \in S(w)\} + d_{w,t} - \mathrm{dist}_{\widehat{G}}(u_1,t)))\\ &= \max(0, \min(|R|+1, \min\{\mathrm{dist}_{\widehat{G}}(u_2,s) \ : \ s \in S(w)\} + d_{w,t} - \mathrm{dist}_{\widehat{G}}(u_2,t))).\end{aligned}$$

This proves (7). From (7), we infer $|\widehat{\mathcal{X}}| \leqslant |\mathcal{C}^3| = \mathcal{O}(k^4|R|^{12})$, as desired. ◁

We next relate the distance from a vertex $u \in U$ to R (in G) and to t (in $\widehat{G}$).

▷ Claim 19. For every $u \in U$, $\mathrm{dist}_G(u,R) = \mathrm{dist}_{\widehat{G}}(u,t) - 2|V(G)|$.

Proof. Fix $u \in U$. We first show that $\mathrm{dist}_G(u,R) \leqslant \mathrm{dist}_{\widehat{G}}(u,t) - 2|V(G)|$. For this, consider a shortest path $\widehat{Q}$ in $\widehat{G}$ from u to t. Observe that there is a vertex $s \in \mathcal{T}$ that is a copy of an anchor vertex w, such that $\widehat{Q}[s,t]$ is the path from s to t of length $d_{w,t}$ added in the construction of $\widehat{G}$ from $G \between H$. Recall that $d_{w,t} = \mathrm{dist}_G(w,R) + |V(G)|$. Also, observe that $\widehat{Q}[u,s]$ contains at least one subdivided edge of E_{out}, as it starts in U and ends outside U, and otherwise corresponds to a walk from u to w in G. Therefore, we have

$$\begin{aligned}\mathrm{dist}_{\widehat{G}}(u,t) = |\widehat{Q}| &= |\widehat{Q}[u,s]| + |\widehat{Q}[s,t]|\\ &= |\widehat{Q}[u,s]| + \mathrm{dist}_G(w,R) + |V(G)|\\ &\geqslant |V(G)| + \mathrm{dist}_G(u,w) + \mathrm{dist}_G(w,R) + |V(G)|\\ &\geqslant \mathrm{dist}_G(u,R) + 2|V(G)|.\end{aligned}$$

We next show that $\mathrm{dist}_G(u,R) \geqslant \mathrm{dist}_{\widehat{G}}(u,t) - 2|V(G)|$. For this, consider a shortest path Q in G from u to R. Let $y \in R$ be the endpoint of Q in R. Also, let z be the first vertex of Q (when ordering from u to y) in $\bigcup_{P\in\mathcal{P}} V(P)$ and let $P \in \mathcal{P}$ be the path that z is contained (if z is contained to more than one paths, pick one of them arbitrarily). Also, let w be the first vertex of $P[z, x^P]$ (when ordering from z to x^P) that is an anchor vertex. Observe that $Q[u,z]$ corresponds to a path in $\widehat{G}$ from u to a copy s' of z that contains exactly one subdivided edge of E_{out} (and no edge of E_{next}) and there is a copy of $P[z,w]$ in $\widehat{G}$ from s' to a copy s of w that contains no edge of $E_{\mathsf{out}} \cup E_{\mathsf{next}}$. Therefore,

$$\begin{aligned}
|Q| &= |Q[u,z]| + |Q[z,y]| \\
&= |Q[u,z]| + |P[z,x^P]| \\
&\qquad\qquad\qquad (Q[z,y] \text{ and } P[z,x^P] \text{ being shortest paths from } z \text{ to } R) \\
&= |Q[u,z]| + |P[z,w]| + |P[w,x^P]| \\
&= |Q[u,z]| + |P[z,w]| + \mathrm{dist}_G(w,R) \\
&\qquad\qquad\qquad (P \text{ being shortest path from a vertex } v^P \text{ to } R) \\
&\geqslant \mathrm{dist}_{\widehat{G}}(u,s) - |V(G)| + d_{w,t} - |V(G)| \\
&\geqslant \mathrm{dist}_{\widehat{G}}(u,t) - 2|V(G)|.
\end{aligned}$$

Thus, we have $\mathrm{dist}_G(u,R) = |Q| \geqslant \mathrm{dist}_{\widehat{G}}(u,t) - 2|V(G)|$, as desired. ◁

▷ Claim 20. For every $u \in U$ and $w \in A_R(\mathcal{P})$, it holds that

$$\begin{aligned}
\widehat{\mathrm{dist}}(u,w) < \infty \quad &\text{if and only if} \quad \widehat{\mathrm{dist}}(u,w) = \min\left\{\mathrm{dist}_{\widehat{G}}(u,s)\ :\ s \in S(w)\right\} - |V(G)|, \\
\widehat{\mathrm{dist}}(u,w) = \infty \quad &\text{if and only if} \quad \min\left\{\mathrm{dist}_{\widehat{G}}(u,s)\ :\ s \in S(w)\right\} > 2|V(G)|.
\end{aligned}$$

Proof. We first show that if $\widehat{\mathrm{dist}}(u,w) < \infty$, then there exists $s \in S(w)$ with $\mathrm{dist}_{\widehat{G}}(u,s) \leqslant |V(G)| + \widehat{\mathrm{dist}}(u,w)$. To this end, let Q be a path from u to w in G of length $\widehat{\mathrm{dist}}(u,w)$, as in the definition of $\widehat{\mathrm{dist}}(u,w)$. There exists $P \in \mathcal{P}$ with $w \in A_R(P)$ and a vertex $z \in V(P) \cap V(Q)$ such that Q decomposes into $Q[u,z]$ and $Q[z,w] = P[z,w]$, with all internal vertices of $Q[u,z]$ in U. Then, $\widehat{G}$ contains a copy s' of z such that $Q[u,z]$ projects to a path from u to s' with one subdivided edge of E_{out} (and no edge of E_{next}) and also a copy of $P[z,w]$ from s' to a copy s of w with no subdivided edge of $E_{\mathsf{out}} \cup E_{\mathsf{next}}$. The concatenation of these two paths witness that $\mathrm{dist}_{\widehat{G}}(u,s) \leqslant |V(G)| + \widehat{\mathrm{dist}}(u,w)$, as desired.

To finish the proof of the claim, it suffices to show that if there exists $s \in S(w)$ with $\mathrm{dist}_{\widehat{G}}(u,s) \leqslant 2|V(G)|$, then $\widehat{\mathrm{dist}}(u,w) \leqslant \mathrm{dist}_{\widehat{G}}(u,s) - |V(G)|$ (in particular, $\widehat{\mathrm{dist}}(u,w) \neq \infty$). To this end, let $\widehat{Q}$ be a path in $\widehat{G}$ from u to s of minimum length. Since $u \in U$ but $s \notin U$, $\widehat{Q}$ necessarily contains at least one subdivided edge of E_{out}. Since $|\widehat{Q}| \leqslant 2|V(G)|$, $\widehat{Q}$ contains exactly one edge of E_{out}, no edge of E_{next}, and no edge incident with t. Consequently, there exists a vertex s' on $\widehat{Q}$ which is a copy of a vertex z that lies in the prefix of w on some path $P \in \mathcal{P}$ such that $\widehat{Q}$ decomposes as $\widehat{Q}[u,s']$, which has all internal vertices in U, and $\widehat{Q}[s',s]$ going along a copy of $P[z,w]$ to $s \in S(w)$. Hence, $\widehat{Q}$ corresponds to a path Q in G from u to w that satisfies the requirements of the definition of $\widehat{\mathrm{dist}}(u,w)$ and witnesses $\widehat{\mathrm{dist}}(u,w) \leqslant |\widehat{Q}| - |V(G)|$, as desired.

This finishes the proof of the claim. ◁

Using the two previous claims, we infer that for every $u \in U$ and $w \in A_R(\mathcal{P})$ it holds that

$$\widehat{\mathrm{prof}}_{R,\mathcal{P}}[u](w) = \min(|R| + 1, \pi[u](w)). \tag{8}$$

Indeed,

$$\begin{aligned}
&\min\left(|R|+1, \pi[u](w)\right) \\
&\quad = \min\left(|R|+1, \min\left\{\mathrm{dist}_{\widehat{G}}(u,s) \ :\ s \in S(w)\right\} + d_{w,t} - \mathrm{dist}_{\widehat{G}}(u,t)\right) \\
&\quad = \min\left(|R|+1, \min\left\{\mathrm{dist}_{\widehat{G}}(u,s) \ :\ s \in S(w)\right\} - |V(G)| \right. \\
&\qquad\qquad\qquad \left. + \mathrm{dist}_G(w,R) - \mathrm{dist}_G(u,R)\right) \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{by Claim 19} \\
&\quad = \min\left(|R|+1, \widehat{\mathrm{dist}}(u,w) + \mathrm{dist}_G(w,R) - \mathrm{dist}_G(u,R)\right) \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{by Claim 20} \\
&\quad = \widehat{\mathrm{prof}}_{R,\mathcal{P}}[u](w).
\end{aligned}$$

Here, in the third step we used the estimate $\mathrm{dist}_G(u,R) - \mathrm{dist}_G(w,R) \leqslant |U| \leqslant |V(G)| - |R|$, so if $\min\left\{\mathrm{dist}_{\widehat{G}}(u,s) \ :\ s \in S(w)\right\} > 2|V(G)|$ (which is equivalent to $\widehat{\mathrm{dist}}(u,w) = \infty$ by Claim 20), then the minimum is attained at value $|R|+1$.

For every $u \in U$, we set

$$B_u \coloneqq \left\{(w,i) \in A_R(\mathcal{P}) \times \mathbb{Z} \mid i \leqslant \widehat{\mathrm{prof}}_{G,R}[u](w)\right\}.$$

Claim 18 and (8) imply that the set system $\{B_u \ :\ u \in U\}$ has size $\mathcal{O}(k^4|R|^{12})$.

Now, for every $v \in V(G)$, we set

$$S_v \coloneqq \left\{(s,i) \in R \times \{-|R|, \ldots, |R|\} \mid i \leqslant \mathrm{prof}_{R,s_R}[v](s)\right\}.$$

The bound on the size of the set system $\{B_u \ :\ u \in U\}$ and Lemma 15 imply that the size of $\{S_u \ :\ u \in U\}$ is bounded by $\mathcal{O}(k^4|R|^{12})$. We conclude the proof of the lemma by bounding the size of $\{S_u \ :\ u \in V(G) - U\}$. For this, note that every vertex $v \in V(G) - U$ is either a milestone for some path $P \in \mathcal{P}$ or a vertex in the neutral prefix of a milestone. In the latter case, there is a path $P \in \mathcal{P}$ and a milestone $w \in M_R(P)$ such that $S_v = S_w$. Therefore, we have

$$|\{S_u \ :\ u \in V(G) - U\}| \leqslant \sum_{P \in \mathcal{P}} |M_R(P)| \leqslant \mathcal{O}(k|R|^2),$$

where the second inequality follows from (5). Hence, the size of $\{S_v \ :\ v \in V(G)\}$ is at most $|\{S_u \ :\ u \in U\}| + |\{S_u \ :\ u \in V(G) - U\}| = \mathcal{O}(k^4|R|^{12})$. This finishes the proof of Theorem 4. ◀

4 Bounded Euler genus graphs with apices: proof of Theorem 5

In this section we prove Theorem 5. (Note that Theorem 2 is a special case of Theorem 5 for $k = 1$.) We start by deriving the following corollary from Theorem 9.

▶ **Corollary 21.** *Let V be a set of n points in $\mathbb{R}^d$. There is a data structure that uses $\mathcal{O}\left(dn\log^{d-2} n\right)$ preprocessing time, $\mathcal{O}\left(dn\log^{d-2} n\right)$ memory and answers the following queries in time $\mathcal{O}\left(d\log^{d-2} n\right)$: given $r_1, \ldots, r_d \in \mathbb{R}$, find $\max_{v \in V} \min_{i \in [d]}(v_i + r_i)$, where v_i denotes the ith coordinate of v.*

Proof. Fix query parameters $r_1, \ldots, r_d \in \mathbb{R}$. Let $\lambda \coloneqq \max_{v \in V} \min_{i \in [d]}(v_i + r_i)$ denote the answer we want to find.

We say that a pair $(v, i) \in V \times [d]$ is *good* if for every $j \in [d]$, it holds that $v_j - v_i \geqslant r_i - r_j$. Let

$$\lambda' = \max\left\{v_i + r_i : i \in [d], v \in V, \text{ and } (v, i) \text{ is good}\right\}.$$

We claim that

$$\lambda = \lambda'. \tag{9}$$

Let $v' = \text{argmax}_{v \in V}(\min_{i \in [d]} v_i + r_i)$ and let $i' = \text{argmin}_{i \in [d]} v'_i + r_i$. By the choice of i', for each j we have $v'_j + r_j \geqslant v'_{i'} + r_{i'}$, implying $v'_j - v'_{i'} \geqslant r_{i'} - r_j$. Hence (v', i') is good, so $\lambda' \geqslant v'_{i'} + r_{i'} = \lambda$.

On the other hand, consider a good pair (v', i') maximizing $v'_{i'} + r_{i'}$. The goodness of (v', i') implies that $i' = \text{argmin}_{i \in [d]} v'_i + r_i$, hence $\lambda \geqslant \min_{i \in [d]} v'_i + r_i = v'_{i'} + r_{i'} = \lambda'$. This proves (9).

For every $i \in [d]$, we set V_i to be the set

$$\{(v_1 - v_i, v_2 - v_i, \ldots, v_{i-1} - v_i, v_{i+1} - v_i, \ldots, v_d - v_i) : v \in V\} \subseteq \mathbb{R}^{d-1},$$

and set $w_i(v) \coloneqq v_i$. Let $\mathbb{D}_i$ be the data structure obtained by applying Theorem 9 to V_i and w_i. Consider the suffix range

$$R_i \coloneqq \mathsf{Range}(r_i - r_1, r_i - r_2, \ldots, r_i - r_{i-1}, r_i - r_{i+1}, \ldots, r_i - r_d) \subseteq \mathbb{R}^{d-1}.$$

Now, by (9) we have that

$$\lambda = \max\left\{r_i + \max\{w_i(v) : v \in V_i \cap R_i\} : i \in [d]\right\}.$$

This value can be computed by asking d queries to the data structures $\mathbb{D}_i$, for $i \in [d]$. This gives us a data structure satisfying the conditions given in the lemma statement. ◀

The main work in the proof of Theorem 5 will be done in the following lemma, which provides a fast computation of eccentricities once a suitable division is provided on input. We adopt the notation for divisions introduced in the statement of Theorem 10.

▶ **Lemma 22.** *Fix constants $0 < \alpha, \gamma, \rho < 1$ and $k \in \mathbb{N}$. Assume we are given a connected graph G on n vertices with $O(n)$ edges with positive integer weights, a subset of vertices X, a subset of apices $A \subseteq V(G)$ of size at most k, and a family $\mathcal{R}$ with $V(G) - A = \bigcup \mathcal{R}$ such that the following conditions are satisfied:*

- $\sum_{R \in \mathcal{R}} |\partial R| \leqslant \mathcal{O}(n^\gamma)$;
- *for every $R \in \mathcal{R}$, $|R| \leqslant \mathcal{O}(n^\rho)$ and $G[R]$ is connected and contains $O(|R|)$ edges; and*
- *for every $R \in \mathcal{R}$, the number of distance profiles in $G - A$ on ∂R is of $\mathcal{O}(n^\alpha)$.*

Then, we can compute X-eccentricity of every vertex of G in time $\mathcal{O}(n^{\gamma+2\rho} \log n + (n^{1+\gamma} + n^{1+\alpha}) \log^{k-1} n)$.

Proof. Let $G' \coloneqq G - A$ and $X' \coloneqq X \cap V(G')$. Denote $A \coloneqq \{a_1, a_2, \ldots, a_k\}$. We first describe the procedure, and then discuss its time complexity.

For every $a \in A$ and $u \in V(G)$, we compute distance between a and u denoted $d_A(a, u)$.

Step 1. We start by precomputing the following information for every region $R \in \mathcal{R}$. For all $u, v \in R$, we compute the distance between u and v in $G'[R]$, denoted $d_R(u, v)$. For all $s \in \partial R, u \in V(G')$, we compute the distance between s and u in G', denoted $d_{\partial R}(u, s)$. We arbitrarily pick a pivot vertex $s_R \in \partial R$, and for brevity denote $p_R[u] := \mathrm{prof}_{\partial R, s_R}[u]$, where the profile is considered in G'. That is, $p_R[u]$ is the (∂R)-profile of u with respect to s_R:

$$p_R[u](s) = d_{\partial R}(u, s) - d_{\partial R}(u, s_R), \quad \text{for all } u \in V(G') \text{ and } s \in \partial R.$$

We define $P_R := \{p_R[u] : u \in V(G')\}$. By our assumption, we have $|P_R| \leqslant \mathcal{O}(n^\alpha)$. Finally, for every profile $p \in P_R$, we list all vertices $v \in X' - R$ such that $p_R[v] = p$ and set up the data structure of Corollary 21 for the points $(d_A(a_1, v), \ldots, d_A(a_k, v), d_{\partial R_u}(s_R, v))$; denote it by $\mathbb{D}_{R,p}$.

Step 2. For every $u \in V(G)$, we compute $\mathrm{ecc}_X(u)$ as follows. If $u \in A$, the answer is $\max_{v \in X} d_A(u, v)$. Hence, we may assume $u \notin A$. Let R_u denote any region of $\mathcal{R}$ containing u. For every $v \in R_u$, the shortest path from u to v in G either:

- goes through an apex, in which case its length is $\min_{a \in A} d_A(a, u) + d_A(a, v)$; or
- is disjoint from A and intersects ∂R_u, in which case its length is $\min_{s \in \partial R_u} d_{\partial R_u}(s, u) + d_{\partial R_u}(s, v)$; or
- is contained entirely in R_u, in which case its length is $d_{R_u}(u, v)$.

The length of this path is therefore the minimum among the three quantities. Using the above observation, we compute $\mathrm{dist}_G(u, v)$ explicitly for each $v \in R_u$, and define $\Delta_u^{R_u} := \max_{v \in R_u \cap X} \mathrm{dist}_G(u, v)$.

For every $v \in V(G) - (A \cup R_u)$, the shortest path between u and v either crosses A or ∂R_u. The length of such path avoiding A is

$$\min_{s \in \partial R_u} d_{\partial R_u}(s, u) + d_{\partial R_u}(s, v) = d_{\partial R_u}(s_R, v) + \min_{s \in \partial R_u} \left(d_{\partial R_u}(s, u) + p_{R_u}[v](s)\right).$$

We partition the vertices v by their profile $p_{R_u}[v]$ and for every $p \in P_{R_u}$, we compute the maximum distance to a vertex with profile p separately. Let $V_p = \{v \in X' - R_u \mid p_{R_u}[v] = p\}$. For every $v \in V_p$, we have

$$\mathrm{dist}_G(u, v) = \min\left(\min_{a \in A} d_A(a, u) + d_A(a, v), d_{\partial R_u}(s_R, v) + \min_{s \in \partial R_u} \left(d_{\partial R_u}(s, u) + p(s)\right)\right).$$

We set $r_i := d_A(a, u)$ for $i \in [k]$, and $r_{k+1} := \min_{s \in \partial R_u} \left(d_{\partial R_u}(s, u) + p(s)\right)$. Now,

$$\max_{v \in V_p} \mathrm{dist}_G(u, v) = \max_{v \in V_p} \min(r_1 + d_A(a_1, v), \ldots, r_k + d_A(a_k, v), r_{k+1} + d_{\partial R_u}(s_R, v)).$$

This value can be computed by querying $r_1, \ldots, r_{k+1}$ to the data structure $\mathbb{D}_{R_u,p}$. We define $\Delta_u^{V(G)-(A \cup R_u)}$ as the maximum of such values over all $p \in P_{R_u}$.

Finally, we set $\Delta_u^A := \max_{a \in A \cap X} d_A(a, u)$, and report

$$\mathrm{ecc}_X(u) = \max\left(\Delta_u^A, \Delta_u^{R_u}, \Delta_u^{V(G)-(A \cup R_u)}\right).$$

It remains to argue that this algorithm can be implemented in the desired running time. For any source $u \in V(G)$, distance from u to all vertices of G can be calculated in time $\mathcal{O}((|V(G)| + |E(G)|) \log |V(G)|)$ using Dijkstra's algorithm. Therefore:

- computing $d_A(a, \cdot)$ for all a can be done in time $\mathcal{O}(n \log n)$,
- computing $d_{\partial R}(\cdot, \cdot)$ for all R can be done in time $\mathcal{O}(n \log n \cdot \sum_{R \in \mathcal{R}} |\partial R|) \leqslant \mathcal{O}(n^{1+\gamma} \log n)$,

- computing $d_R(\cdot,\cdot)$ for all $R \in \mathcal{R}$ can be done in time $\mathcal{O}(|\mathcal{R}|n^{2\rho}\log n) \leqslant \mathcal{O}(n^{\gamma+2\rho}\log n)$; constructing $G[R]$ takes $\mathcal{O}(|R|^2\log n) = \mathcal{O}(n^{2\rho}\log n)$ time and calculating all pairs shortest paths can be done in time $\mathcal{O}(|R||E(G[R])|\log n) = \mathcal{O}(n^{2\rho}\log n)$.

Finally, the total size of the data structures $\mathbb{D}_{R,p}$ over all R, p is $\mathcal{O}(|\mathcal{R}|n) = \mathcal{O}(n^{1+\gamma})$, hence we can construct them in time $\mathcal{O}(n^{1+\gamma}\log^{k-1} n)$.

Consider $u \in V(G) - A$ fixed in step 2. Computing $\Delta_u^{R_u}$ takes $\mathcal{O}(|R| \cdot |\partial R_u|)$ time. Computing Δ_u^A can be done in constant time. Computing $\Delta_u^{V(G)-(A\cup R_u)}$ requires asking $|P_{R_u}|$ queries to some $\mathbb{D}_{R,p}$, which takes $\mathcal{O}(n^\alpha \log^{k-1} n)$ time in total. In total, step 2 for all vertices u can be done in time $\mathcal{O}(n^{1+\alpha}\log^{k-1} n + n^\rho \cdot \sum_{u\in V(G)-A} |\partial R_u|) = \mathcal{O}(n^{1+\alpha}\log^{k-1} n + n^{\gamma+2\rho})$.

We conclude that the total running time is $\mathcal{O}(n^{\gamma+2\rho}\log n + (n^{1+\gamma} + n^{1+\alpha})\log^{k-1} n)$. ◀

The next statement is a reformulation of Theorem 5.

▶ **Theorem 23.** *Fix constants $k, g \in \mathbb{N}$. Let $\mathcal{C}$ denote the class of all graphs that can be obtained by taking a graph G of Euler genus bounded by g, and adding k apices adjacent arbitrarily to the rest of G and to each other. Then there is an algorithm that given an unweighted graph G belonging to $\mathcal{C}$, together with its set of apices A, computes the eccentricity of every vertex in time $\mathcal{O}_{k,g}\left(n^{1+\frac{24}{25}}\log^{k-1} n\right)$.*

Proof. Let $A = \{a_1, \ldots, a_k\}$ denote the set of apices and let $G' = G - A$. Fix $\rho \coloneqq \frac{2}{25}$. Since graphs of bounded genus exclude some fixed clique as a minor, by Theorem 10 (with $\varepsilon = \rho/2$) we can find an $\mathcal{O}(n^\rho)$-division $\mathcal{R}$ of G' satisfying $\sum_{R\in\mathcal{R}} |\partial R| = \mathcal{O}(n^{1-\frac{\rho}{2}})$ in time $\mathcal{O}(n^{1+\rho})$. By Theorem 4, the graph G' has a degree 12 polynomial bound on the number of distance profiles. In particular, the number of profiles on every ∂R is of $\mathcal{O}(|R|^{12}) = \mathcal{O}(n^{12\rho})$. Let $X \coloneqq V(G)$, $\gamma \coloneqq 1 - \frac{\rho}{2} = \frac{24}{25}$ and $\alpha \coloneqq 12\rho = \frac{24}{25}$. Now, applying Lemma 22 gives us an algorithm computing all eccentricities in time $\mathcal{O}(n^{1+\frac{24}{25}}\log^{k-1} n)$. ◀

References

1 Glencora Borradaile, Erik D. Demaine, and Siamak Tazari. Polynomial-time approximation schemes for subset-connectivity problems in bounded-genus graphs. *Algorithmica*, 68(2):287–311, 2014. doi:10.1007/S00453-012-9662-2.

2 Sergio Cabello. Subquadratic algorithms for the diameter and the sum of pairwise distances in planar graphs. *ACM Transactions on Algorithms*, 15(2):21:1–21:38, 2019. doi:10.1145/3218821.

3 Sergio Cabello and Christian Knauer. Algorithms for graphs of bounded treewidth via orthogonal range searching. *Computational Geometry*, 42(9):815–824, 2009. doi:10.1016/J.COMGEO.2009.02.001.

4 Sergio Cabello, Éric Colin de Verdière, and Francis Lazarus. Algorithms for the edge-width of an embedded graph. *Computational Geometry*, 45(5):215–224, 2012. Special issue: 26th Annual Symposium on Computation Geometry at Snowbird, Utah, USA. doi:10.1016/j.comgeo.2011.12.002.

5 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.

6 Lech Duraj, Filip Konieczny, and Krzysztof Potępa. Better diameter algorithms for bounded VC-dimension graphs and geometric intersection graphs. In *32nd Annual European Symposium on Algorithms, ESA 2024*, volume 308 of *LIPIcs*, pages 51:1–51:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICS.ESA.2024.51.

7 Jeff Erickson and Kim Whittlesey. Greedy optimal homotopy and homology generators. In *Proc. of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2005*, pages 1038–1046. SIAM, 2005. URL: http://dl.acm.org/citation.cfm?id=1070432.1070581.

8 Paweł Gawrychowski, Haim Kaplan, Shay Mozes, Micha Sharir, and Oren Weimann. Voronoi diagrams on planar graphs, and computing the diameter in deterministic $\widetilde{O}(n^{5/3})$ time. *SIAM Journal on Computing*, 50(2):509–554, 2021. doi:10.1137/18M1193402.

9 Kacper Kluk, Marcin Pilipczuk, Michał Pilipczuk, and Giannos Stamoulis. Faster diameter computation in graphs of bounded euler genus. *CoRR*, abs/2502.07501, 2025. doi:10.48550/arXiv.2502.07501.

10 Tuukka Korhonen, Michał Pilipczuk, and Giannos Stamoulis. Minor Containment and Disjoint Paths in almost-linear time. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 53–61, 2024. doi:10.1109/FOCS61266.2024.00014.

11 Tuukka Korhonen, Michał Pilipczuk, Giannos Stamoulis, and Dimitrios Thilikos, 2024. Private communication.

12 Hung Le and Christian Wulff-Nilsen. VC set systems in minor-free (di)graphs and applications. In *2024 ACM-SIAM Symposium on Discrete Algorithms, SODA 2024*, pages 5332–5360. SIAM, 2024. doi:10.1137/1.9781611977912.192.

13 Jason Li and Merav Parter. Planar diameter via metric compression. In Moses Charikar and Edith Cohen, editors, *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 152–163. ACM, 2019. doi:10.1145/3313276.3316358.

14 Bojan Mohar and Carsten Thomassen. *Graphs on Surfaces*. Johns Hopkins series in the mathematical sciences. Johns Hopkins University Press, 2001. URL: https://jhupbooks.press.jhu.edu/title/graphs-surfaces.

15 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.

16 Liam Roditty and Virginia Vassilevska Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In *45th Symposium on Theory of Computing Conference, STOC 2013*, pages 515–524. ACM, 2013. doi:10.1145/2488608.2488673.

17 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.

18 Saharon Shelah. A combinatorial problem; stability and order for models and theories in infinitary languages. *Pacific Journal of Mathematics*, 41(1):247–261, 1972. doi:10.2140/pjm.1972.41.247.

19 Dimitrios M. Thilikos and Sebastian Wiederrecht. Killing a vortex. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022*, pages 1069–1080. IEEE, 2022. doi:10.1109/FOCS54457.2022.00104.

20 V. N. Vapnik and A. Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications*, 16(2):264–280, 1971. doi:10.1137/1116025.

21 Dan E. Willard. New data structures for orthogonal range queries. *SIAM Journal on Computing*, 14(1):232–253, 1985. doi:10.1137/0214019.

22 Christian Wulff-Nilsen. Separator theorems for minor-free and shallow minor-free graphs with applications. *CoRR*, abs/1107.1292, 2011. arXiv:1107.1292.