



Quasi-Polynomial Time Approximation Schemes for Packing and Covering Problems in Planar Graphs

Michał Pilipczuk¹ · Erik Jan van Leeuwen² · Andreas Wiese³ 

Received: 7 November 2018 / Accepted: 30 December 2019 / Published online: 13 January 2020
© Springer Science+Business Media, LLC, part of Springer Nature 2020

Abstract

We consider two optimization problems in planar graphs. In MAXIMUM WEIGHT INDEPENDENT SET OF OBJECTS we are given a graph G and a family \mathcal{D} of *objects*, each being a connected subgraph of G with a prescribed weight, and the task is to find a maximum-weight subfamily of \mathcal{D} consisting of pairwise disjoint objects. In MINIMUM WEIGHT DISTANCE SET COVER we are given a graph G in which the edges might have different lengths, two sets \mathcal{D}, \mathcal{C} of vertices of G , where vertices of \mathcal{D} have prescribed weights, and a nonnegative radius r . The task is to find a minimum-weight subset of \mathcal{D} such that every vertex of \mathcal{C} is at distance at most r from some selected vertex. Via simple reductions, these two problems generalize a number of geometric optimization tasks, notably MAXIMUM WEIGHT INDEPENDENT SET for polygons in the plane and WEIGHTED GEOMETRIC SET COVER for unit disks and unit squares. We present *quasi-polynomial time approximation schemes* (QPTASs) for both of the above problems in planar graphs: given an accuracy parameter $\epsilon > 0$ we can compute a solution whose weight is within multiplicative factor of $(1 + \epsilon)$ from the optimum in time $2^{\text{poly}(1/\epsilon, \log |\mathcal{D}|)} \cdot n^{\mathcal{O}(1)}$, where n is the number of vertices of the input graph. We note that a QPTAS for MAXIMUM WEIGHT INDEPENDENT SET OF OBJECTS would follow from existing work. However, our main contribution is to provide a unified framework that works for both problems in both a planar and geometric setting and to transfer the techniques used for recursive approximation schemes for geometric problems due to Adamaszek and Wiese (in Proceedings of the FOCS 2013, IEEE, 2013; in Proceedings of the SODA 2014, SIAM, 2014) and Har-Peled and Sariel (in Proceedings of the SOCG 2014, SIAM, 2014) to the setting of planar graphs. In particular, this yields a purely combinatorial viewpoint on these methods as a phenomenon in planar graphs.

The work of Mi. Pilipczuk on these results was, at the initial stages, partially supported by Polish National Science Centre grant UMO-2013/11/D/ST6/03073, and, at the later stages, was a part of project TOTAL that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant Agreement No. 677651). The research of A. Wiese is supported by the Millennium Nucleus Information and Coordination in Networks ICM/FIC RC130003 and by the grant Fondecyt Regular 1170223.

Extended author information available on the last page of the article

Keywords Approximation schemes · Planar graphs · Independent set of objects · Geometric set cover

1 Introduction

INDEPENDENT SET and DOMINATING SET are fundamental optimization problems on graphs. Given a graph G where each vertex v has a weight $w(v)$, in INDEPENDENT SET one seeks to find a vertex subset $I \subseteq V(G)$ of maximum possible weight such that no two vertices in I are adjacent, whereas in DOMINATING SET one searches for a vertex subset D of minimum possible weight such that each vertex $v \in V(G)$ is contained in D or adjacent to a vertex in D . Even in the unit-weight setting, both problems are notoriously hard to approximate and they are also $W[1]$ -hard, i.e., we do not expect that they admit *fixed-parameter tractable (fpt)* algorithms running in time $f(k) \cdot n^{O(1)}$, where k is the expected solution size.

Therefore, special cases of the problems were investigated, for instance the case where the input graph is planar. On planar graphs, classic layering techniques can be applied to show that both problems admit EPTASs, i.e., $(1 + \epsilon)$ -approximation algorithms with a running time of $f(1/\epsilon)n^{O(1)}$ for some function f , and fpt algorithms for the parameterization by the solution size, i.e., for a parameter k , algorithms running in time $f(k)n^{O(1)}$ that find a best solution among those of size at most k . Given these results, it is natural to consider generalizations of the above problems on planar graphs.

In this paper we study the DISTANCE INDEPENDENT SET and the DISTANCE DOMINATING SET problems. Given additionally a value $r \in \mathbb{R}$ and lengths for the edges of G , in the DISTANCE INDEPENDENT SET problem we require that any two selected vertices in I are at distance larger than r from each other, and in the DISTANCE DOMINATING SET problem we require that each vertex $v \in V(G)$ is at distance at most r from some vertex of D . Let us stress that we assume that r is part of the input and in particular not assumed to be a constant; in fact, for constant r and unit edge lengths, it is well-known that the same layering techniques easily yield EPTASs and fpt algorithms on planar graphs. In the parameterized setting, both problems are $W[1]$ -hard even for unit weights and unit edge-lengths; however, the trivial $n^{O(k)}$ -time algorithms can be improved to $n^{O(\sqrt{k})}$ -time algorithms [5]. These parameterized algorithms extend a technique originally developed to design quasi-polynomial time approximation schemes (QPTASs) for INDEPENDENT SET and DOMINATING SET in the geometric (Euclidean) setting [1, 2, 4]. The idea is to guess a sparse separator that has only small intersection with the optimal solution and that splits the problem into two roughly equal-sized subproblems, and then to solve the subproblems recursively. The natural question arises whether one can transfer the insights obtained in the parameterized setting back to approximation algorithms, and obtain approximation schemes for DISTANCE INDEPENDENT SET and DISTANCE DOMINATING SET in planar graphs.

Our contribution. In this paper we show that this is indeed possible and we present the first quasi-polynomial time approximation schemes for DISTANCE INDEPENDENT SET and DISTANCE DOMINATING SET on planar graphs when r is part of the input.

In fact, we give QPTASs for two even more general problems, which we name MAXIMUM WEIGHT INDEPENDENT SET OF OBJECTS (MWISO) and MINIMUM WEIGHT DISTANCE SET COVER (MWDSC). In MWISO we are given a graph G and a family \mathcal{D} of *objects*, each being a connected subgraph of G with a prescribed weight, and the goal is to find a maximum-weight subfamily of \mathcal{D} consisting of pairwise disjoint objects. In MWDSC we are given an graph G and a length for each of its edges, subsets of vertices \mathcal{D} and \mathcal{C} where vertices of \mathcal{D} are weighted, and radius $r \in \mathbb{R}$. The goal is to find a minimum-weight subset $\mathcal{F} \subseteq \mathcal{D}$ that r -covers \mathcal{C} in the sense that each vertex of \mathcal{C} is at distance at most r from some vertex of \mathcal{F} . MWISO generalizes DISTANCE INDEPENDENT SET by taking \mathcal{D} to be the family $\{\{v : \text{dist}(u, v) \leq r/2\} : u \in V(G)\}$ of all balls of radius $r/2$ in the graph, after subdividing the edges appropriately¹. MWDSC generalizes DISTANCE DOMINATING SET by taking $\mathcal{C} = V(G)$. The following statements summarize our results.

Theorem 1 *The MAXIMUM WEIGHT INDEPENDENT SET OF OBJECTS problem in planar graphs admits a QPTAS with running time $2^{\text{poly}(1/\epsilon, \log N)} \cdot n^{O(1)}$, where n is the vertex count of the input graph and $N = |\mathcal{D}|$ is the number of objects in the input.*

Theorem 2 *The MINIMUM WEIGHT DISTANCE SET COVER problem in planar graphs admits a QPTAS with running time $2^{\text{poly}(1/\epsilon, \log N)} \cdot n^{O(1)}$, where n is the vertex count of the input graph and $N = |\mathcal{D}|$ is the number of vertices allowed to be selected to the solution.*

We remark that it is possible to derive Theorem 1 via a reduction to the MAXIMUM WEIGHT INDEPENDENT SET OF POLYGONS problem (see below) and then applying the known QPTAS [4] for this problem. However, our main focus in this paper is to transfer the machinery developed in [1, 2, 4] for optimization problems in geometric settings to problems in planar graphs. The heart of our technical contribution is to show that for any instance of the above problems there is a set of candidate separators of polynomial size such that one of them splits the given problem in a balanced way and intersects only a tiny fraction of the given solution. The latter is important since the intersected objects will be lost (in the case of MWISO) or might be paid twice (in the case of MWDSC) and hence we need to bound their total weight by ϵOPT . We state here an informal version of our separator lemma for the case of MWISO.

Lemma 1 (Informal) *In polynomial time we can compute a set $\mathbb{X} \subseteq 2^{\mathcal{D}}$ of separators such that for every solution $\mathcal{F} \subseteq \mathcal{D}$, say of weight W , there exists $\mathcal{X} \in \mathbb{X}$ such that $\mathbf{w}(\mathcal{F} \cap \mathcal{X}) \leq \epsilon W$ and in the intersection graph of $\mathcal{D} - \mathcal{X}$ each connected component \mathcal{C} satisfies $\mathbf{w}(\mathcal{C} \cap \mathcal{F}) \leq \frac{\epsilon}{10} W$.*

¹ For every pair of vertices u', v' with $\text{dist}(u', v') \leq r$, if $\{v : \text{dist}(u', v) \leq r/2\} \cap \{v : \text{dist}(v', v) \leq r/2\} = \emptyset$ then we subdivide the middle edge of the shortest path between u' and v' such that the sets $\{v : \text{dist}(u', v) \leq r/2\}, \{v : \text{dist}(v', v) \leq r/2\}$ intersect on the new vertex. This way we add at most $O(n^2)$ new vertices to the instance.

Using Lemma 1 as abstraction for finding separators, we can apply the same recursive scheme as [1, 2, 4]: we guess the correct separator $\mathcal{X} \in \mathbb{X}$, construct a subproblem for each connected component of the intersection graph of $\mathcal{D} - \mathcal{X}$, and recurse in each of them up to recursion depth $\mathcal{O}(\log |\mathcal{D}|)$. Thus, the only part of the reasoning that uses planarity is Lemma 1.

The proof of Lemma 1 follows the reasoning of Har-Peled [4]. The idea is to prove the following auxiliary result: for the optimal solution \mathcal{F} (and in fact for any feasible solution) there exists a separator of length roughly $s = \mathcal{O}(\frac{1}{\epsilon} \ln \frac{1}{\epsilon})$ that cuts through at most an ϵ -fraction of the weight of \mathcal{F} and splits the weight of \mathcal{F} in a balanced way. Lemma 1 then follows by enumerating all candidates for such separators. In [4], the separator was simply a polygon with roughly s vertices. We lift this concept to planar graphs using *Voronoi separators* as in the work of Marx and Pilipczuk [5]. Intuitively, a Voronoi separator of length r is an alternating cyclic sequence of r objects from \mathcal{D} and r faces of the graph, connected by shortest paths in order to form a closed curve; this curve splits the instance into two subinstances. Thus, shortest paths in the graph are the analogues of segments in the plane.

The auxiliary result is proved in [4] by showing that if \mathcal{S} is a sample of size roughly s^2 from \mathcal{F} , where each object is sampled independently with probability proportional to its weight, then a balanced separator of length s in the Voronoi diagram of \mathcal{S} satisfies all the required properties with high probability. We follow the same reasoning; however, again we need to properly understand how geometric concepts used in [4] — *spokes* and *corridors* — should be interpreted in planar graphs. Here, the technical toolbox for Voronoi diagrams and Voronoi separators developed in [5] becomes very useful. In particular, it turns out that a fine understanding of what faces are candidates for branching points of a Voronoi diagram, provided in [5], is essential to make the probabilistic argument work. Let us remark that we also somewhat simplify the original argument of Har-Peled by replacing the Exponential Decay Lemma with a direct probabilistic calculation.

To give the QPTAS for MWDSC we provide a variant of Lemma 1 suitable for this problem and then follow a similar recursive scheme as for Theorem 1. It is nice that we can reuse the above-mentioned auxiliary result introduced for Lemma 1 as a black-box, so the proof of the variant is relatively short. As in [7], the difference is that in the recursion instead of removing the guessed separator we preserve it in all the recursive subcalls, thus allowing double-buying objects from it.

Geometric problems. The above recursive machinery based on balanced separators was first introduced for obtaining a QPTAS for MAXIMUM WEIGHT INDEPENDENT SET OF RECTANGLES in the two-dimensional plane [1] and then extended for getting QPTASs for MAXIMUM WEIGHT INDEPENDENT SET OF POLYGONS [2, 4] and WEIGHTED GEOMETRIC SET COVER (WGSR) for pseudo-disks [7]. We prove that Theorems 1 and 2 generalize these results, with the exception that for WGSR we can treat only the cases of unit disks and axis-parallel unit squares, instead of general families of pseudo-disks. In “Appendix A” we explain how to derive the mentioned results from our theorems.

We would like to comment that it is possible to reduce MWISO to MAXIMUM WEIGHT INDEPENDENT SET OF POLYGONS and obtain a QPTAS for MWISO in this way: take the input graph and compute a straight line embedding for it. For each $p \in \mathcal{D}$

compute a spanning tree T_p and define a polygon P_p that consists of the edges of T_p . Now two polygons $P_p, P_{p'}$ overlap if and only if their corresponding objects p, p' overlap. Applying the QPTAS for MAXIMUM WEIGHT INDEPENDENT SET OF POLYGONS [4] to the resulting instance thus yields a QPTAS for MWISO. However, we believe that our QPTAS for MWISO is simpler in the sense that it works with the planar graph directly. Also note that such an approach does not work for MWDSC.

2 Proof of the Separator Lemma for MWISO

In this section we prove the Separator Lemma for MWISO, which was informally stated as Lemma 1 and is formally stated below. For a family \mathcal{D} of objects, $\text{IntGraph}(\mathcal{D})$ denotes the *intersection graph* of \mathcal{D} : graph with vertex set \mathcal{D} where two objects are adjacent iff they intersect. The reader may think of \mathcal{F} being the optimal solution and of W being its weight.

Lemma 2 (Separator Lemma for MWISO) *Let G be an n -vertex planar graph and \mathcal{D} be a weighted family of N objects in G . Let $0 < \epsilon < \frac{1}{10}$, and denote $s = 10^3 \cdot \frac{1}{\epsilon} \ln \frac{1}{\epsilon}$. Then there exists a family \mathbb{X} consisting of subsets of \mathcal{D} with the following properties:*

- (A1) $|\mathbb{X}| \leq 6^{3s} N^{15s}$ and \mathbb{X} can be computed in time $N^{O(s)} \cdot n^{O(1)}$; and
- (A2) for every real $W \geq 0$ and subfamily $\mathcal{F} \subseteq \mathcal{D}$ of pairwise disjoint objects such that $\mathbf{w}(\mathcal{F}) \leq W$ and $\mathbf{w}(p) \leq s^{-2}W$ for each $p \in \mathcal{F}$, there exist $\mathcal{X} \in \mathbb{X}$ such that $\mathbf{w}(\mathcal{F} \cap \mathcal{X}) \leq \epsilon W$ and for every connected component \mathcal{C} of $\text{IntGraph}(\mathcal{D}) - \mathcal{X}$ have $\mathbf{w}(\mathcal{C} \cap \mathcal{F}) \leq \frac{9}{10}W$.

The plan is as follows. We first recall the toolbox of *Voronoi separators*, introduced by Marx and Pilipczuk [5]. This allows us to state a stronger lemma, phrased as the existence of a short Voronoi separator appropriately breaking \mathcal{F} . We then show how Lemma 2 follows from this stronger result and subsequently prove the stronger result.

Before we proceed, let us set up the notation and basic assumptions about the input. Let G be the input graph. We assume the edges of G are assigned positive lengths² so that we have the shortest-path metric in G : $\text{dist}(u, v)$ denotes the shortest length of a path connecting u and v in G . We may assume that G is given with an embedding in a sphere Σ and that it is triangulated; that is, every face of G is a triangle. Indeed, adding edges of infinite length to triangulate the graph neither distorts the metric nor spoils the connectivity of the objects. Also, for every face f of G we fix any internal point c to be its *center*, and for each vertex u of f we fix some curve

² Obviously, edge lengths are immaterial for the MWISO problem. However, it is convenient to think of G as a graph whose edges have lengths so that we can define Voronoi diagrams. Furthermore, many results of this section will be reused for the MWDSC problem, where edge lengths do play a role.

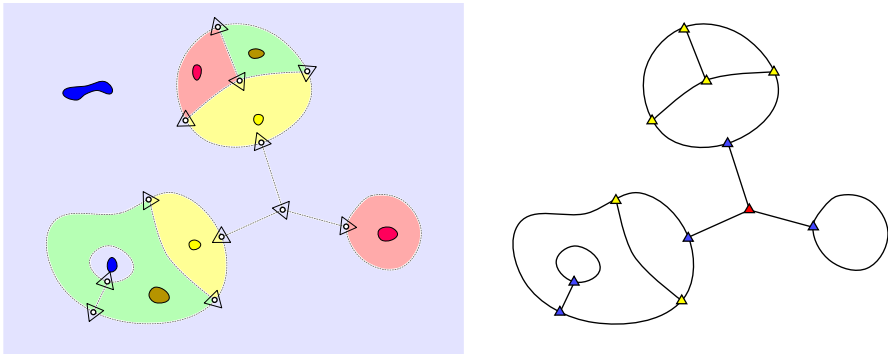


Fig. 1 Construction of the Voronoi diagram. The left panel depicts the Voronoi partition, branching points of the diagram are depicted as solid (primal) triangular faces, edges of the diagram are dashed. The right panel depicts the Voronoi diagram alone; note that it is a planar 3-regular multigraph. Branching points of types 1, 2, and 3 are depicted in yellow, blue, and red, respectively. The figures are taken from [5] with consent of the authors; the left one is slightly modified (Color figure online)

within f with endpoints u and c to be the *segment* connecting u and c so that segments connecting vertices of f with c pairwise do not cross.

We assume that shortest paths are unique and the distances between pairs of vertices are pairwise different: for every pair of vertices u, v there is a unique shortest path connecting u and v and for $\{u, v\} \neq \{u', v'\}$ we have $\text{dist}(u, v) \neq \text{dist}(u', v')$. This can be ensured by using lexicographic tie-breaking rules and it increases the running time only by polynomial factors.

We are given a family \mathcal{D} of *objects* in G , where each object $p \in \mathcal{D}$ is a non-empty, connected subgraph of G . For any vertex u of G and any object $p \in \mathcal{D}$, let $\text{dist}(u, p)$ be the length of the shortest path connecting u with any vertex of p . For each object $p \in \mathcal{D}$, we fix any spanning tree $T(p)$ of p . We also assume that the family \mathcal{D} is weighted: every object $p \in \mathcal{D}$ is assigned a nonnegative real weight $\mathbf{w}(p)$. For $\mathcal{F} \subseteq \mathcal{D}$ we denote $\mathbf{w}(\mathcal{F}) = \sum_{p \in \mathcal{F}} \mathbf{w}(p)$.

2.1 Basic Toolbox

Voronoi partitions and diagrams. A subfamily $\mathcal{F} \subseteq \mathcal{D}$ is *independent* if the objects in \mathcal{F} are pairwise vertex-disjoint. Such an independent subfamily \mathcal{F} induces the *Voronoi partition* $\mathbb{M}_{\mathcal{F}}$, which is a partition of the vertex set of G into $|\mathcal{F}|$ parts according to the closest object from \mathcal{F} . Precisely, for $p \in \mathcal{F}$, we say that a vertex u belongs to the *Voronoi cell* $\mathbb{M}_{\mathcal{F}}(p)$ if $\text{dist}(u, p) < \text{dist}(u, p')$ for any $p' \in \mathcal{F}$, $p' \neq p$. Observe that ties do not happen due to distinctness of distances in G . We note that Marx and Pilipczuk consider in [5] a more general notion of a *normal subfamily*, but we do not need this generality here.

Assuming $|\mathcal{F}| \geq 4$, we define the *Voronoi diagram* induced by $\mathbb{M}_{\mathcal{F}}$ as follows; see Fig. 1 for a visualization. First, observe that every Voronoi cell $\mathbb{M}_{\mathcal{F}}(p)$, for $p \in \mathcal{F}$, induces a connected subgraph of G that contains p entirely (see Lemmas 4.1 and 4.2 in [5]). Extend $T(p)$ to a spanning tree $\hat{T}(p)$ of $G[\mathbb{M}_{\mathcal{F}}(p)]$ by adding, for each

vertex u of $\mathbb{M}_{\mathcal{F}}(p)$ that is not in p , the shortest path from u to p . Take the dual of G and remove all the edges dual to the edges of $\widehat{T}(p)$, for all $p \in \mathcal{F}$. Then exhaustively remove vertices of degree one, and finally replace each maximal path with internal vertices of degree 2 (so-called *2-path*) by a single edge; the embedding of this edge is defined as the union of embeddings of edges of G comprising the original 2-path. Thus, we obtain a connected, 3-regular plane multigraph H , called the *Voronoi diagram* of \mathcal{F} , whose faces bijectively correspond to the cells $\mathbb{M}_{\mathcal{F}}(p)$ for $p \in \mathcal{F}$. More precisely, every face of H is associated with a different object $p \in \mathcal{F}$ so that all the vertices of $\mathbb{M}_{\mathcal{F}}(p)$ are contained in this face. The 3-regularity of H follows from the assumption that G is triangulated. From Euler’s formula it follows that if $|\mathcal{F}| = k$, then H has k faces, $2k - 4$ vertices, and $3k - 6$ edges. See Lemmas 4.4 and 4.5 of [5] for a formal verification of these assertions, and Section 4.4 of [5] for a detailed description of the construction.

Branching points. If H is the Voronoi diagram of an independent subfamily $\mathcal{F} \subseteq \mathcal{D}$, then H is constructed from a subgraph of the dual of G by contracting maximal 2-paths. Hence, vertices of H correspond to faces of G . These primal faces, equivalently dual vertices, are called the *branching points* of the diagram H ; intuitively, these are faces where the boundaries of Voronoi cells meet non-trivially. A priori, every face of G could be a branching point of the Voronoi diagram of some independent subfamily $\mathcal{F} \subseteq \mathcal{D}$. However, in [5] it is proved that the number of candidates for branching points can be bounded polynomially in $|\mathcal{D}|$.

Theorem 3 (Theorem 4.7 of [5]) *There exists a family I of faces of G with $|I| \leq |\mathcal{D}|^4$ such that the following holds: for every independent subfamily of objects $\mathcal{F} \subseteq \mathcal{D}$, all branching points of the Voronoi diagram of \mathcal{F} are contained in I . Moreover, I can be computed in time polynomial in $|\mathcal{D}|$ and n .*

We fix the family I provided by Theorem 3 and call its members *\mathcal{D} -important faces* of G .

Voronoi separators. We now recall the concept of *Voronoi separators*; see Fig. 4 for a visualization. A *Voronoi separator* is a sequence of the form

$$S = \langle p_1, u_1, f_1, v_1, p_2, u_2, f_2, v_2, \dots, p_r, u_r, f_r, v_r \rangle,$$

where p_i are pairwise disjoint objects from \mathcal{D} , f_i are faces of G , and u_i, v_i are distinct vertices lying on the face f_i . For each $i \in \{1, \dots, r\}$, define P_i to be the shortest path from u_i to p_i and Q_i to be the shortest path from v_i to p_{i+1} , where indices behave cyclically. For a Voronoi separator S as above, its *length* is r and its *set of traversed objects* is $\mathcal{D}(S) = \{p_1, \dots, p_r\}$.

In the notation above, an object $q \in \mathcal{D}$ is *banned* by the separator S if either q intersects some object $p \in \mathcal{D}(S)$, or there is a vertex w on some path P_i such that $\text{dist}(w, q) < \text{dist}(w, p_i)$, or there is a vertex w on some path Q_i such that $\text{dist}(w, q) < \text{dist}(w, p_{i+1})$. In particular, $\mathcal{D}(S)$ is banned by S . Let $\text{Ban}(S)$ denote the set of objects in \mathcal{D} banned by S . Intuitively, the banned objects are those that are intersected by the separator and are lost when we recurse (in MWISO) or that

might be selected and paid twice (in MWDCS). Therefore, we will later ensure that their total weight is small.

The following result is the aforementioned key step toward the proof of Lemma 2. It may be regarded as a lift of Theorem 4.22 from [5] or of Lemma 4.1 from [4] to our setting.

Lemma 3 *Let W be a positive real, $0 < \epsilon < \frac{1}{10}$, and $s = 10^3 \cdot \frac{1}{\epsilon} \ln \frac{1}{\epsilon}$. Suppose $\mathcal{F} \subseteq \mathcal{D}$ is an independent subfamily of objects such that $|\mathcal{F}| \geq 4$, $\mathbf{w}(\mathcal{F}) \leq W$, and $\mathbf{w}(p) \leq s^{-2}W$ for all $p \in \mathcal{F}$. Then there exists a Voronoi separator S satisfying the following:*

- (B1) $\mathcal{D}(S) \subseteq \mathcal{F}$ and all faces traversed by S are \mathcal{D} -important;
- (B2) the length of S is at most $3s$;
- (B3) the total weight of objects of \mathcal{F} banned by S is at most ϵW ;
- (B4) for every connected component \mathcal{C} of $\text{IntGraph}(\mathcal{D}) - \text{Ban}(S)$, we have $\mathbf{w}(\mathcal{C} \cap \mathcal{F}) \leq \frac{9}{10}W$.

It is not hard to see that Lemma 2 follows from Lemma 3: we simply enumerate all candidates for a Voronoi separator S satisfying (B1) and (B2). A straightforward estimate using Theorem 3 shows that there are at most $6^{3s}N^{15s}$ candidates, and for each candidate S we add $\text{Ban}(S)$ to the constructed family \mathbb{X} .

Proof (of Lemma 2 assuming Lemma 3) Let us inspect in how many ways a Voronoi separator S satisfying properties (B1) and (B2) can be selected. Since S has length at most $3s$, there are at most N^{3s} ways to select the sequence of consecutive objects visited by S . Since every face traversed by S is \mathcal{D} -important and there are at most N^4 \mathcal{D} -important faces in total, there are at most N^{12s} ways to select the sequence of consecutive faces visited by S . On each of these faces we need to pick a pair of different vertices as the entry and leaving point, giving 6 ways per faces, so at most 6^{3s} in total. Thus, the separator S may be selected in at most $6^{3s}N^{15s}$ different ways. Furthermore, since the set of \mathcal{D} -important faces can be computed in polynomial time by Theorem 3, a family \mathcal{N} of at most $6^{3s}N^{15s}$ Voronoi separators satisfying (B1) and (B2) can be enumerated in time $N^{\mathcal{O}(s)} \cdot n^{\mathcal{O}(1)}$.

Now define the output family \mathbb{X} as follows:

$$\mathbb{X} = \{\text{Ban}(S) : S \in \mathcal{N}\}.$$

Clearly, \mathbb{X} satisfies property (A1), whereas that it also satisfies property (A2) follows directly from Lemma 3, properties (B3) and (B4). \square

Thus, we are left with proving Lemma 3. The idea, borrowed from Har-Peled [4], is that we construct a sufficiently large random sample from \mathcal{F} , where the probability of picking each object is proportional to its weight. Then we inspect the Voronoi diagram induced by the sample and we argue that with non-zero probability it has a short separator giving rise to the sought Voronoi separator S .

To implement this plan we need two ingredients: an appropriate lift of the sampling idea from [4] and the analysis of how separators in the Voronoi diagram give rise to Voronoi separators in the graph, which is essentially taken from [5] with some technical details added. These two ingredients are explained in the next two subsections.

2.2 Sampling

Spokes and diamonds. We first adjust technical notions used by Har-Peled [4] in the geometric context to our setting; see Fig. 5 for their visualization. Suppose we have an independent family of objects \mathcal{F} and its Voronoi diagram $H = H_{\mathcal{F}}$. Let f be any branching point of H and let u be any vertex of f . Let $p \in \mathcal{F}$ be the object of \mathcal{F} such that $u \in \mathbb{M}_{\mathcal{F}}(p)$. The *spoke* of u in H is the shortest path from u to p in G ; note all the vertices of this shortest path belong to the cell $\mathbb{M}_{\mathcal{F}}(p)$.

Consider any subfamily $\mathcal{S} \subseteq \mathcal{F}$ and let $H_{\mathcal{S}}$ be the Voronoi diagram induced by \mathcal{S} ; in the following, we consider spokes in the diagram $H_{\mathcal{S}}$. For any spoke P in $H_{\mathcal{S}}$, say connecting a vertex u with the object $p \in \mathcal{S}$ satisfying $u \in \mathbb{M}_{\mathcal{S}}(p)$, we say that P is *in conflict* with an object $p' \in \mathcal{F}$ if there is a vertex v on P such that $\text{dist}(v, p') < \text{dist}(v, p)$. Note that this implies $p' \notin \mathcal{S}$, because the spoke P has to be entirely contained in $\mathbb{M}_{\mathcal{S}}(p)$. Define the *weight* of P with respect to \mathcal{F} as the total weight of objects from \mathcal{F} that are in conflict with P .

Further, suppose e is an edge of $H_{\mathcal{S}}$, with endpoints f_1, f_2 (not necessarily different). Let p_1, p_2 be the objects of \mathcal{S} corresponding to the faces of $H_{\mathcal{S}}$ incident to e (possibly $p_1 = p_2$). Let $u_{1,1}, u_{1,2}$ be the vertices of f_1 such that $u_{1,1} \in \mathbb{M}_{\mathcal{S}}(p_1)$, $u_{1,2} \in \mathbb{M}_{\mathcal{S}}(p_2)$, and the edge $u_{1,1}u_{1,2}$ of G crosses the edge e of $H_{\mathcal{S}}$. Similarly pick vertices $u_{2,1}, u_{2,2}$ of f_2 . The *diamond* induced by e is the closed curve $\Delta_{\mathcal{S}}(e)$ on Σ formed by the union of: the segments connecting the center of f_1 with $u_{1,1}$ and $u_{1,2}$, the unique path between $u_{1,2}$ and $u_{2,2}$ in $\hat{T}(p_2)$, the segments connecting the center of f_2 with $u_{2,2}$ and $u_{2,1}$, and the unique path between $u_{2,1}$ and $u_{1,1}$ in $\hat{T}(p_1)$. The *interior* of $\Delta_{\mathcal{S}}(e)$ is the unique region of $\Sigma \setminus \Delta_{\mathcal{S}}(e)$ that contains e . The *weight* of $\Delta_{\mathcal{S}}(e)$ with respect to \mathcal{F} is the total weight of objects of \mathcal{F} that are entirely contained in the interior of $\Delta_{\mathcal{S}}(e)$; note that these objects do not belong to \mathcal{S} .

We remark that while spokes were used in [4] in the form roughly as above, diamonds correspond to the notion of a *corridor* from [4].

Sampling lemma: statement. We now state the crucial technical result: there is a bounded-size subfamily of the optimum solution that induces a Voronoi diagram where every spoke and every diamond has small weight. We will prove it using a probabilistic sampling argument.

Lemma 4 (Sampling lemma) *Suppose W is a positive real and \mathcal{F} is an independent, weighted family of objects in G such that $|\mathcal{F}| \geq 4$ and $\mathbf{w}(\mathcal{F}) \leq W$. Let $\ell \geq 10$ be an integer such that $\mathbf{w}(p) \leq \frac{W}{\ell}$ for each $p \in \mathcal{F}$. Then there exists a subfamily $\mathcal{S} \subseteq \mathcal{F}$ with $4 \leq |\mathcal{S}| \leq 2\ell$ such that in the Voronoi diagram $H_{\mathcal{S}}$, the weight with respect to \mathcal{F} of every spoke and of every diamond is at most $10 \ln \ell \cdot \frac{W}{\ell}$.*




Fig. 2 Singular faces of types 1, 2, and 3, respectively. Shortest paths from the vertices of the face f to respective objects are depicted in gray. The figure is taken verbatim from [5], with the consent of the authors (Color figure online)

Later, we will use Lemma 4 with $\ell = \mathcal{O}((\frac{1}{\epsilon} \ln \frac{1}{\epsilon})^2)$. The reader may imagine that we then apply a balanced planar graph separator on the Voronoi diagram H_S of size $\mathcal{O}(\sqrt{\ell})$ along which we partition \mathcal{F} into two parts, yielding the Voronoi separator claimed by Lemma 3. Since the weight with respect to \mathcal{F} of every spoke and every diamond is at most $10 \ln \ell \cdot \frac{W}{\ell}$, the total weight of the objects of \mathcal{F} banned by S will be bounded by ϵW .

Lemma 4 is a roughly an analogue of Lemma 3.3 from [4]. The main difference is that in the geometric setting, spokes and corridors have a simpler structure due to the fact that each branching point of the Voronoi diagram is defined by three objects from the solution — the three ones equidistant from it — so that the branching point is the meeting point of the three corresponding Voronoi regions. This is no longer the case in planar graphs, as observed in [5]. More precisely, out of the three regions around a branching point of the diagram, two or even three may be equal; this happens when there are bridges in the diagram, which is never the case in the geometric setting.

As part of their proof of Theorem 3, to understand these additional situations Marx and Pilipczuk define *singular faces*, which come in three types. The first one corresponds to “standard” branching points incident to three different regions, while the second and the third one correspond to branching points incident only to two, respectively one region.

Singular faces. For an independent triple of objects $\mathcal{F}_0 = \{p_1, p_2, p_3\} \subseteq \mathcal{D}$, a face f of G is called a *singular face of type 1* for (p_1, p_2, p_3) if in $\mathbb{M}_{\mathcal{F}_0}$, all the vertices of f belong to different cells (note that there are three cells in $\mathbb{M}_{\mathcal{F}_0}$). For an independent triple of objects $\mathcal{F}_0 = \{p_1, p_2, p_3\} \subseteq \mathcal{D}$, a face f is called a *singular face of type 2* for (p_1, p_2, p_3) if in $\mathbb{M}_{\mathcal{F}_0}$, one vertex v_1 of f belongs to $\mathbb{M}_{\mathcal{F}_0}(p_1)$, the other two vertices v_2, v_3 of f belong to $\mathbb{M}_{\mathcal{F}_0}(p_2)$, and the closed walk W obtained by taking the union of the unique path in $\hat{T}(p_2)$ between v_2 and v_3 and the edge $v_2 v_3$ on the boundary of f divides the plane into two regions, one containing p_1 and one containing p_3 . Finally, for an independent quadruple of objects $\mathcal{F}_0 = \{p_0, p_1, p_2, p_3\} \subseteq \mathcal{D}$, a face f is called a *singular face of type 3* for (p_0, p_1, p_2, p_3) if in $\mathbb{M}_{\mathcal{F}_0}$ all the vertices of f belong to $\mathbb{M}_{\mathcal{F}_0}(p_0)$, but the boundary of face f plus the minimal subtree of $\hat{T}(p_0)$ spanning the vertices of f divides the

plane into four regions: the face f itself, one region containing p_1 , one region containing p_2 , and one region containing p_3 . See Fig. 2 for a visualization.

It appears that for a fixed triple or quadruple of objects, there are only few singular faces.

Lemma 5 (Lemmas 4.8, 4.9, and 4.10 of [5]) *For each independent triple of objects (p_1, p_2, p_3) , there are at most 2 singular faces of type 1 for (p_1, p_2, p_3) , and at most 1 singular face of type 2 for (p_1, p_2, p_3) . For each independent quadruple of objects (p_0, p_1, p_2, p_3) , there is at most 1 singular face of type 3 for (p_0, p_1, p_2, p_3) .*

The next statement explains the connection between branching points and singular faces.

Lemma 6 (Lemma 4.12 of [5]) *Let $\mathcal{F} \subseteq \mathcal{D}$ be an independent subfamily of objects, and let H be the Voronoi diagram of \mathcal{F} . Then every branching point of H is either a type-1 singular face for some triple of objects from \mathcal{F} , or a type-2 singular face for some triple of objects from \mathcal{F} , or a type-3 singular face for some quadruple of objects from \mathcal{F} .*

Actually, the two results above may be combined into a proof of Theorem 3. Lemma 6 shows that every branching point of the Voronoi diagram of an independent subfamily of \mathcal{D} is among type-1, type-2, and type-3 singular faces for triples or quadruples of objects in \mathcal{D} , while using Lemma 5 we can bound their total number by $|\mathcal{D}|^4$.

Sampling lemma: proof We now have all the tools needed to prove Lemma 4. Contrary to Har-Peled [4] we do not use the Exponential Decay Lemma, but direct probability calculations; this makes the proof somewhat conceptually easier. The main complications are due to the need to handling different types of singular faces, instead of just one.

Proof (of Lemma 4) Denote $\eta = 10 \ln \ell \cdot \frac{w}{\ell}$. First observe that if $w(\mathcal{F}) \leq \eta$, then setting $\mathcal{S} = \mathcal{F}$ satisfies all the required properties, since no spoke may have larger weight than the whole of \mathcal{F} . Therefore, from now on we assume that $w(\mathcal{F}) > \eta$.

Construct \mathcal{S} by including every object $p \in \mathcal{F}$ independently with probability $q_p = w(p) \cdot \frac{\ell}{w}$; note that this value is at most 1 by the assumption of the lemma. Let X be the random variable equal to the cardinality of \mathcal{S} ; then $X = \sum_{p \in \mathcal{F}} X_p$, where X_p are indicator random variables, taking value 1 if p is included in \mathcal{S} and 0 otherwise. Note that $\mathbb{E}[X_p] = q_p$ and $\mathbb{E}[X] = \sum_{p \in \mathcal{F}} \mathbb{E}[X_p] = \ell \cdot \frac{w(\mathcal{F})}{w} \leq \ell$. Since X is a sum of independent indicator variables, standard concentration inequalities yield the following. □

Claim 1 *The probability that $|\mathcal{S}| > 2\ell$ or $|\mathcal{S}| \leq 4$ is at most $\frac{7}{12}$.*

Proof Since $\mathbb{E}[X] \leq \ell$, by Markov’s inequality we have that $|\mathcal{S}| > 2\ell$ happens with probability at most $\frac{1}{2}$. Hence, it suffices to prove that the probability that $|\mathcal{S}| \leq 4$ is that most $\frac{1}{12}$.

As $w(\mathcal{F}) > \eta = 10 \ln \ell \cdot \frac{w}{\ell}$, we have

$$\mathbb{E}[X] = \ell \cdot \frac{w(\mathcal{F})}{W} > 10 \ln \ell > 20.$$

Since X is a sum of independent $\{0, 1\}$ -random variables with mean larger than 200, by Chernoff’s inequality we infer that

$$\mathbb{P}(X \leq 4) \leq \exp\left(-\frac{1}{2}\left(\frac{4}{5}\right)^2 \cdot 20\right) \leq \frac{1}{12}.$$

This concludes the proof of the claim. □

Call a spoke in the Voronoi diagram H_S heavy if its weight with respect to \mathcal{F} is more than η . We now estimate the probability that there is a heavy spoke in H_S .

Claim 2 *The probability that there is a heavy spoke in H_S is at most $\frac{1}{6}$.*

Proof Take any triple p_1, p_2, p_3 of different objects from \mathcal{F} , and suppose f is a singular face of type 1 for (p_1, p_2, p_3) . Suppose further that P is one of the spokes incident to f in the Voronoi diagram of $\{p_1, p_2, p_3\} \subseteq \mathcal{F}$, and assume that P is heavy with respect to \mathcal{F} . Let us estimate the probability of the following event A : p_1, p_2, p_3 all belong to \mathcal{S} and moreover P remains a spoke in the Voronoi diagram H_S of \mathcal{S} (then it is obviously a heavy spoke in H_S as well). For A to happen, apart from the event that $p_1, p_2, p_3 \in \mathcal{S}$, we also need that the following event happens: for each $p' \in \mathcal{F} \setminus \{p_1, p_2, p_3\}$ that is in conflict with P , p' is not included in \mathcal{S} . Let $\mathcal{Z} \subseteq \mathcal{F}$ be the set of such objects p' . Using the standard inequality $1 - x \leq \exp(-x)$, we have

$$\begin{aligned} \mathbb{P}(A) &\leq q_{p_1} q_{p_2} q_{p_3} \cdot \prod_{p' \in \mathcal{Z}} (1 - q_{p'}) \leq q_{p_1} q_{p_2} q_{p_3} \cdot \exp\left(-\sum_{p' \in \mathcal{Z}} q_{p'}\right) \\ &\leq \frac{w(p_1)w(p_2)w(p_3)}{W^3} \cdot \ell^3 \cdot \exp\left(-\frac{\ell}{W} \sum_{p' \in \mathcal{Z}} w(p')\right) \\ &\leq \frac{w(p_1)w(p_2)w(p_3)}{W^3} \cdot \ell^3 \cdot \exp\left(-\frac{\ell \eta}{W}\right) \tag{1} \\ &\leq \frac{w(p_1)w(p_2)w(p_3)}{W^3} \cdot \ell^3 \cdot \exp(-10 \ln \ell) = \frac{w(p_1)w(p_2)w(p_3)}{W^3} \cdot \ell^{-7} \\ &\leq 10^{-7} \cdot \frac{w(p_1)w(p_2)w(p_3)}{W^3}. \end{aligned}$$

For a fixed triple $p_1, p_2, p_3 \in \mathcal{F}$, face f that is singular of type 1 for (p_1, p_2, p_3) , and spoke P incident to f , let us denote by $A_{P, (p_1, p_2, p_3)}^1$ the event A considered above. Let

B^1 denote the event that some event $A_{P,(p_1,p_2,p_3)}^1$ happens, i.e., B^1 is the union of all events $A_{P,(p_1,p_2,p_3)}^1$. By Lemma 5, for every triple $p_1, p_2, p_3 \in \mathcal{F}$ there are at most two faces f that are singular of type 1 for (p_1, p_2, p_3) , and for each of them there are at most 3 spokes P to consider. Thus, by applying the union bound to (1) we infer that

$$\mathbb{P}(B^1) \leq 6 \cdot 10^{-7} \cdot \sum_{p_1 \in \mathcal{F}} \sum_{p_2 \in \mathcal{F}} \sum_{p_3 \in \mathcal{F}} \frac{\mathbf{w}(p_1)\mathbf{w}(p_2)\mathbf{w}(p_3)}{W^3} = 6 \cdot 10^{-7} \cdot \frac{\mathbf{w}(\mathcal{F})^3}{W^3} \leq 10^{-6}.$$

Denote by B^2 the following event: there exists distinct $p_1, p_2, p_3 \in \mathcal{F}$, a singular face f of type 2 for (p_1, p_2, p_3) such that (i) $p_1, p_2, p_3 \in \mathcal{S}$ and (ii) at least one of the heavy (w.r.t. \mathcal{F}) spokes incident to f in the Voronoi diagram of $\{p_1, p_2, p_3\}$ remains a spoke in \mathcal{S} . Also, denote by B^3 the following event: there exists distinct $p_0, p_1, p_2, p_3 \in \mathcal{F}$, a singular face f of type 3 for (p_0, p_1, p_2, p_3) such that (i) $p_0, p_1, p_2, p_3 \in \mathcal{S}$ and (ii) at least one of the heavy (w.r.t. \mathcal{F}) spokes incident to f in the Voronoi diagram of $\{p_0, p_1, p_2, p_3\}$ remains a spoke in \mathcal{S} . Similar calculations as for B^1 yield the following:

$$\mathbb{P}(B^2) \leq 10^{-6} \quad \text{and} \quad \mathbb{P}(B^3) \leq 10^{-6}.$$

Here, when estimating $\mathbb{P}(B^3)$ we sum over quadruples of objects instead of triples. By Lemma 6, we conclude that the probability that any spoke in the diagram H_S is heavy is at most $3 \cdot 10^{-6} \leq \frac{1}{6}$. \square

We are left with diamonds. Call a diamond $\Delta_S(e)$ in H_S heavy if its weight with respect to \mathcal{F} is larger than η . The next check follows by essentially the same estimation as Claim 2.

Claim 3 *The probability that there is a heavy diamond in H_S is at most $\frac{1}{6}$.*

Proof Suppose e is an edge of H_S with endpoints f_1, f_2 , and let p_1, p_2 be the objects of \mathcal{S} corresponding to faces of H_S incident to e . By Lemma 6, each f_t ($t \in \{1, 2\}$) is either a type-1 singular face for a triple of objects from \mathcal{S} , or a type-2 singular face for a triple of objects from \mathcal{S} , or a type-3 singular face for a quadruple of objects from \mathcal{S} . Suppose for a moment that both f_1 and f_2 are type-1 singular faces for triples of objects from \mathcal{S} ; then it is easy to see that f_1 is a type-1 singular face for (p_1, p_2, q_1) for some object $q_1 \in \mathcal{S}$, while f_2 is a type-1 singular face for (p_1, p_2, q_2) for some object $q_2 \in \mathcal{S}$. We will further assume that $q_1 \neq q_2$ and discuss the other case, as well as the cases when f_1 or f_2 are not type-1 singular faces, at the end, as the reasoning for them is analogous.

All in all, we have a quadruple of pairwise different objects p_1, p_2, q_1, q_2 . For the diamond $\Delta = \Delta_S(e)$ to arise in H_S , the following two events need to happen simultaneously:

- objects p_1, p_2, q_1, q_2 need to simultaneously be included in \mathcal{S} ; and
- all objects entirely contained in the interior of Δ need to be not included in \mathcal{S} .

These two events are independent. The probability of the first is

$$\ell^4 \cdot \frac{\mathbf{w}(p_1)\mathbf{w}(p_2)\mathbf{w}(q_1)\mathbf{w}(q_2)}{W^4}.$$

Let \mathcal{G} be the family of objects entirely contained in the interior of Δ . By the assumption that Δ is heavy, we have that $\sum_{r \in \mathcal{G}} \mathbf{w}(r) \geq \eta$. The probability that no object of \mathcal{G} is included in \mathcal{S} is upper bounded by

$$\begin{aligned} \prod_{r \in \mathcal{G}} (1 - q_r) &\leq \prod_{r \in \mathcal{G}} \exp(-q_r) = \exp\left(-\sum_{r \in \mathcal{G}} q_r\right) \\ &= \exp\left(-\frac{\ell}{W} \cdot \sum_{r \in \mathcal{G}} \mathbf{w}(r)\right) \leq \exp(-10 \ln \ell) = \ell^{-10}. \end{aligned}$$

Therefore, the probability that Δ arises in H_S is upper bounded by

$$\frac{\mathbf{w}(p_1)\mathbf{w}(p_2)\mathbf{w}(q_1)\mathbf{w}(q_2)}{W^4} \cdot \ell^4 \cdot \ell^{-10} = \frac{\mathbf{w}(p_1)\mathbf{w}(p_2)\mathbf{w}(q_1)\mathbf{w}(q_2)}{W^4} \cdot \ell^{-6}.$$

Now for every quadruple (p_1, p_2, q_1, q_2) of objects from \mathcal{F} there are at most 4 diamonds induced by this quadruple as above, as there are at most 2 type-1 singular faces for (p_1, p_2, q_1) , and likewise for (p_1, p_2, q_2) . Therefore, the total probability that there exists a heavy diamond with both endpoints being singular faces of type-1 and $q_1 \neq q_2$ is bounded by

$$4 \cdot \ell^{-6} \cdot \sum_{(p_1, p_2, q_1, q_2) \in \mathcal{F}^4} \frac{\mathbf{w}(p_1)\mathbf{w}(p_2)\mathbf{w}(q_1)\mathbf{w}(q_2)}{W^4} = 4 \cdot \ell^{-6} \cdot \frac{\mathbf{w}(\mathcal{F})^4}{W^4} \leq 4 \cdot 10^{-6},$$

where the last inequality follows from the assumption that $\ell \geq 10$. The reasoning for the case when $q_1 = q_2$ is analogous (we sum over triples instead of quadruples), and similarly for the cases when we are dealing with singular faces of type different than 1. The number of different cases one needs to consider is bounded by 100, so summing all the probabilities we conclude that the probability that there is a heavy diamond in H_S is at most $\frac{1}{6}$. □

Concluding, the assertion $4 \leq |\mathcal{S}| \leq 2\ell$ does not hold with probability at most $\frac{7}{12}$, there is a heavy spoke in H_S with probability at most $\frac{1}{6}$, and there is a heavy diamond in H_S with probability at most $\frac{1}{6}$. Hence, with probability at least $\frac{1}{12}$ neither of the above holds, so there exists a subfamily \mathcal{S} satisfying all the postulated conditions. □

2.3 Balanced Nooses

We proceed with the proof of Lemma 3 by explaining the second ingredient: balanced separators in Voronoi diagrams. In general, short embedding-respecting separators in the Voronoi diagram—so-called *nooses*—correspond to the Voronoi separators we are looking for. We start by defining nooses and showing how the

existence of a *sphere-cut decomposition* of small width—a hierarchical decomposition of the diagram using nooses—implies the existence of a short noose that breaks the instance in a balanced way.

We remark that in [4], this part of the reasoning is essentially done by considering the *radial* graph of the Voronoi diagram and applying the weighted balanced separator theorem of Miller [6] to it. Such an approach would be also applicable in our case, but we find the approach via sphere-cut decompositions more explanatory regarding how separators in the (radial graph of the) diagram correspond to separators in the instance.

Sphere-cut decompositions. We now recall the framework of *sphere-cut decompositions*, which are embedding-respecting hierarchical decompositions of planar graphs.

A *branch decomposition* of a graph H is a pair (T, η) where T is a tree with all internal nodes having degree 3, and η is a bijection between the edge set of H and the leaf set of T (for clarity, we always use the term *node* for a vertex of a decomposition tree). Take any edge e of T and consider removing it from T ; then T breaks into two subtrees, say T_1, T_2 . Let A_1, A_2 be the preimages of the leaf sets of T_1, T_2 under η , respectively; then (A_1, A_2) is a partition of the edge set of H . The *width* of the edge e is the number of vertices of H incident to both an edge of A_1 and to an edge of A_2 , and the *width* of the branch decomposition (T, η) is the maximum among the widths of the edges of T . The *branchwidth* of H is the minimum possible width of a branch decomposition of H .

Let H be a connected plane graph embedded in a sphere Σ . A *noose* in H is a closed, directed curve γ on Σ without self-crossings that meets H only at its vertices and visits every face of H at most once. Note that removing γ from the sphere Σ breaks it into two open disks: for one of them γ is the clockwise traversal of the perimeter, and for the other it is the counter-clockwise traversal (fixing an orientation of Σ). The first disk shall be called **enc**(γ) while the second shall be called **exc**(γ) (for *enclosed* and *excluded*). Treat a curve γ on Σ as a continuous function $\gamma : [0, 1] \rightarrow \Sigma$. Two curves γ, γ' on Σ are called *homotopic* on Σ if there is a continuous function $h : [0, 1] \times [0, 1] \rightarrow \Sigma$ such that $h(0, y) = \gamma(0) = \gamma'(0)$ for all $y \in [0, 1]$, $h(1, y) = \gamma(1) = \gamma'(1)$ for all $y \in [0, 1]$, and $h(x, 0) = \gamma(x)$ and $h(x, 1) = \gamma'(x)$ for all $x \in [0, 1]$. Here h is the *homotopy*. Two nooses γ, γ' are *equivalent* if they are homotopic on Σ with a homotopy that fixes the embedding of H ; in other words, γ' can be obtained from γ by continuous transformations within the faces of H .

A *sphere-cut decomposition* of H is a triple (T, η, δ) where (T, η) is a branch decomposition of H and δ maps ordered pairs of adjacent nodes of T to nooses on Σ (w.r.t. H) such that the following conditions are satisfied for each pair of adjacent nodes of T :

- $\delta(x, y)$ is equal to $\delta(y, x)$ reversed (in particular **enc**($\delta(x, y)$) = **exc**($\delta(y, x)$));
- **enc**($\delta(x, y)$) contains all the edges of H mapped to the component of $T - xy$ containing y , while **exc**($\delta(y, x)$) contains all the edges of H mapped to the other component of $T - xy$.

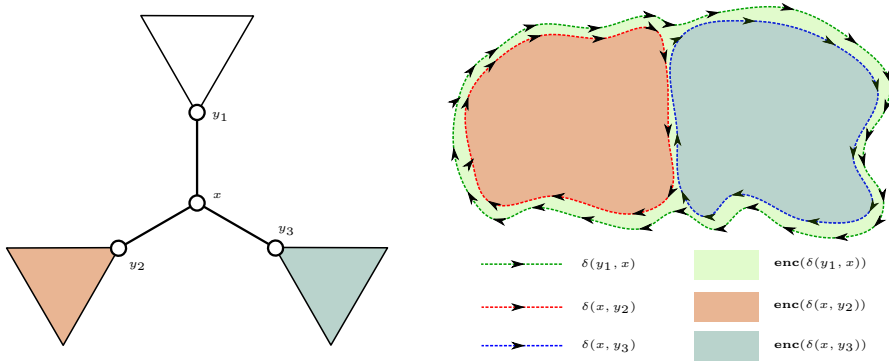


Fig. 3 The condition required from faithful sphere-cut decompositions

The following result follows from [3, 8] and was formulated in exactly this way in [5].

Theorem 4 *Every n -vertex sphere-embedded multigraph that is connected and bridgeless has a sphere-cut decomposition of width at most $\sqrt{4.5n}$.*

We note that in Theorem 4, the assumption that the multigraph is bridgeless is necessary, as multigraphs with bridges do not have sphere-cut decompositions at all.

Suppose (T, η, δ) is a sphere-cut decomposition of G . It is straightforward to see that we may adjust the nooses $\delta(x, y)$ for x, y ranging over adjacent nodes of T so that the following condition is satisfied: if node x has neighbors y_1, y_2, y_3 in T , then $\mathbf{enc}(\delta(y_1, x))$ is the disjoint union of $\mathbf{enc}(\delta(x, y_2))$, $\mathbf{enc}(\delta(x, y_3))$, and $(\delta(x, y_2) \cap \delta(x, y_3)) \setminus \delta(y_1, x)$, see Fig. 3. Sphere-cut decompositions satisfying this condition will be called *faithful*. It is easy to see that any sphere-cut decomposition can be made faithful by changing each noose to an equivalent one.

Separator theorem for nooses. We now state a separator theorem for nooses drawn from a sphere-cut decomposition of a given sphere-embedded multigraph. The theorem is weighted with respect to a measure defined as follows. Suppose \mathcal{R} is a finite family of pairwise disjoint *objects* on a sphere Σ , where each object $p \in \mathcal{R}$ is a nonempty arc-connected subset of Σ with associated nonnegative weight $\mathbf{w}(p)$. For an open disk $\Delta \subseteq \Sigma$, define its \mathcal{R} -measure $\mu_{\mathcal{R}}(\Delta)$ as the total weight of objects from \mathcal{R} that are entirely contained in Δ .

Lemma 7 *Let H be a connected, bridgeless multigraph embedded on a sphere Σ . Let \mathcal{R} be a weighted family of pairwise disjoint objects on Σ and let $W = \mathbf{w}(\mathcal{R})$. Suppose further (T, η, δ) is a faithful sphere-cut decomposition of H such that for every pair (x, y) of adjacent nodes in T such that x is a leaf, we have $\mu_{\mathcal{R}}(\mathbf{enc}(\delta(y, x))) < \frac{9}{20}W$. Then there exists a noose γ w.r.t H , which is one of the nooses in the sphere-cut decomposition (T, η, δ) , such that the following hold:*

$$\mu_{\mathcal{R}}(\mathbf{enc}(\gamma)) \leq \frac{9}{10}W \quad \text{and} \quad \mu_{\mathcal{R}}(\mathbf{exc}(\gamma)) \leq \frac{9}{10}W.$$

The proof of Lemma 7 is standard: we find a balanced edge (x, y) in the decomposition (T, η, δ) and $\delta(x, y)$ is the sought noose. The fact that nooses appearing in (T, η, δ) may intersect objects from \mathcal{R} requires some technical attention. We start by proving a general-use statement using which we can extract balanced separators from hierarchical decompositions, provided we want balancedness with respect to some measure $\mu(\cdot)$ that the decomposition roughly obeys.

Lemma 8 *Let T be a tree with all internal nodes of degree 3 and let W be a positive real. Suppose that there is a function μ which maps ordered pairs of adjacent nodes of T to nonnegative reals such that the following conditions are satisfied:*

(S1) *For every pair x, y of adjacent nodes in T , we have*

$$\frac{9}{10}W \leq \mu(x, y) + \mu(y, x) \leq W.$$

(S2) *If x is an internal node of T with neighbors y_1, y_2, y_3 , then*

$$\mu(x, y_2) + \mu(x, y_3) \leq \mu(y_1, x) \leq \mu(x, y_2) + \mu(x, y_3) + \frac{1}{10}W.$$

(S3) *Whenever x, y are adjacent nodes in T and x is a leaf, we have $\mu(y, x) < \frac{9}{20}W$.*

Then there exists an edge xy of T such that

$$\frac{1}{4}W < \mu(x, y) < \frac{3}{4}W \quad \text{and} \quad \frac{1}{4}W < \mu(y, x) < \frac{3}{4}W.$$

Proof Orient the edges of T as follows: if x, y are adjacent nodes, then orient xy toward x if $\mu(x, y) > \mu(y, x)$, toward y if $\mu(x, y) < \mu(y, x)$, and arbitrarily if $\mu(x, y) = \mu(y, x)$. Thus, T becomes an oriented tree with $|V(T)|$ nodes and $|V(T)| - 1$ edges, implying that there exists a node x of T that has outdegree 0.

We first check that x is not a leaf of T . Suppose otherwise and let y be the unique neighbor of x . By (S1) we have that $\mu(y, x) + \mu(x, y) \geq \frac{9}{10}W$. Since xy was oriented toward x , we infer that $\mu(y, x) \geq \mu(x, y)$, implying $\mu(y, x) \geq \frac{9}{20}W$. This is a contradiction with (S3).

Therefore x is an internal node of T . Let y_1, y_2, y_3 be the neighbors of x . Without loss of generality suppose $\mu(x, y_1)$ is the largest among $\mu(x, y_1), \mu(x, y_2), \mu(x, y_3)$. Using (S1) and (S2) we infer that

$$\mu(x, y_1) + \mu(x, y_2) + \mu(x, y_3) \geq \frac{9}{10}W - \mu(y_1, x) + \mu(x, y_2) + \mu(x, y_3) \geq \frac{4}{5}W,$$

which implies that $\mu(x, y_1) \geq \frac{4}{15}W > \frac{1}{4}W$. As xy_1 was oriented toward x , we also have $\mu(x, y_1) \leq \frac{1}{2}W$. By (S1) we infer that

$$\mu(y_1, x) \geq \frac{9}{10}W - \mu(x, y_1) \geq \frac{4}{10}W > \frac{1}{4}W \quad \text{and} \quad \mu(y_1, x) \leq W - \mu(x, y_1) \leq \frac{11}{15}W < \frac{3}{4}W.$$

Thus, the edge xy_1 satisfies the requirements from the statement. \square

With Lemma 8 in place, we now give the proof of Lemma 7.

Proof (of Lemma 7) We shall say that a closed non-self-crossing curve γ on Σ is \mathcal{R} -light if the total weight of objects in \mathcal{R} that intersect γ is at most $\frac{W}{10}$.

Suppose first that the sphere-cut decomposition (T, η, δ) contains at least one noose γ which is not \mathcal{R} -light. Then the total weight of objects of \mathcal{R} intersected by γ is more than $\frac{W}{10}$, hence both $\mu_{\mathcal{R}}(\mathbf{enc}(\gamma)) \leq \frac{9}{10}W$ and $\mu_{\mathcal{R}}(\mathbf{exc}(\gamma)) \leq \frac{9}{10}W$ hold, and γ satisfies the required properties. Hence, we may assume that all the nooses in (T, η, δ) are \mathcal{R} -light.

We now verify that for \mathcal{R} -light curves as separators, the requirements of Lemma 8 hold. \square

Claim 4 *Suppose γ is an \mathcal{R} -light curve on Σ . Then*

$$\frac{9}{10}W \leq \mu_{\mathcal{R}}(\mathbf{enc}(\gamma)) + \mu_{\mathcal{R}}(\mathbf{exc}(\gamma)) \leq W.$$

Proof The right inequality is obvious. For the left one, all objects in \mathcal{R} that do not intersect γ are entirely contained in either in $\mathbf{enc}(\gamma)$ or in $\mathbf{exc}(\gamma)$, thus they contribute their weight to the sum $\mu_{\mathcal{R}}(\mathbf{enc}(\gamma)) + \mu_{\mathcal{R}}(\mathbf{exc}(\gamma))$, while the total weight of objects intersecting γ is at most $\frac{W}{10}$ due to \mathcal{R} -lightness of γ . \square

Claim 5 *Suppose $\gamma, \gamma_1, \gamma_2$ are \mathcal{R} -light curves such that $\mathbf{enc}(\gamma)$ is equal to the disjoint union of $\mathbf{enc}(\gamma_1)$, $\mathbf{enc}(\gamma_2)$ and $(\gamma_1 \cap \gamma_2) \setminus \gamma$. Then*

$$\mu_{\mathcal{R}}(\mathbf{enc}(\gamma_1)) + \mu_{\mathcal{R}}(\mathbf{enc}(\gamma_2)) \leq \mu_{\mathcal{R}}(\mathbf{enc}(\gamma)) \leq \mu_{\mathcal{R}}(\mathbf{enc}(\gamma_1)) + \mu_{\mathcal{R}}(\mathbf{enc}(\gamma_2)) + \frac{W}{10}.$$

Proof The left inequality is obvious, because every subset from \mathcal{R} entirely contained either in $\mathbf{enc}(\gamma_1)$ or in $\mathbf{enc}(\gamma_2)$ is also entirely contained in $\mathbf{enc}(\gamma)$, while $\mathbf{enc}(\gamma_1)$ and $\mathbf{enc}(\gamma_2)$ are disjoint. For the right inequality, observe that every object from \mathcal{R} that is entirely contained in $\mathbf{enc}(\gamma)$ but not in $\mathbf{enc}(\gamma_1)$ or $\mathbf{enc}(\gamma_2)$ has to intersect $(\gamma_1 \cap \gamma_2) \setminus \gamma$. Since γ_1 is \mathcal{R} -light, the total weight of such objects is at most $\frac{W}{10}$. \square

We proceed with the proof of Lemma 7. For any ordered pair (x, y) of adjacent nodes in T , let $\mu(x, y) = \mu_{\mathcal{R}}(\mathbf{enc}(\delta(x, y)))$. Since all nooses in (T, η, δ) are \mathcal{R} -light and (T, η, δ) is faithful, Claims 4 and 5 verify that conditions (S1) and (S2) of Lemma 8 hold. Condition (S3), on the other hand, is satisfied by the assumptions of the lemma. Therefore, we may apply Lemma 8, yielding a pair (x, y) such that $\delta(x, y)$ satisfies the required assertions. \square

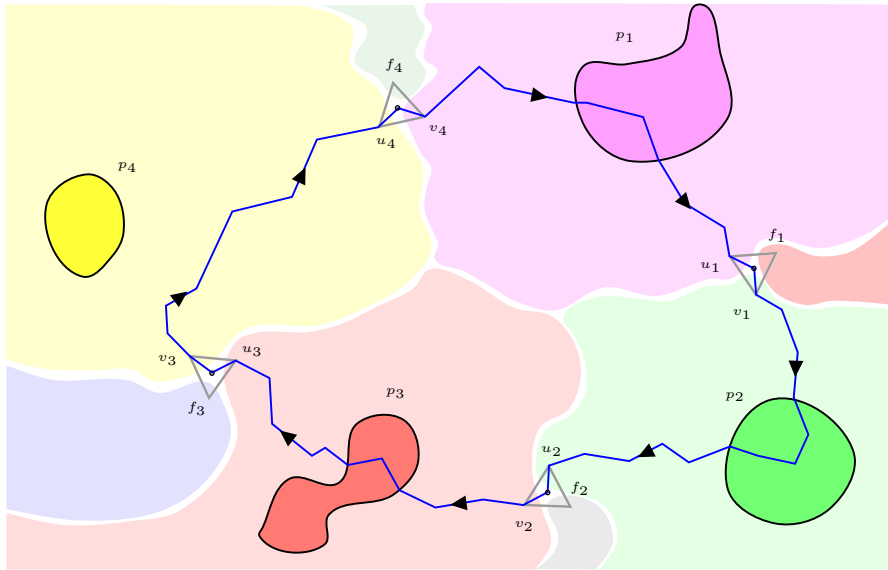


Fig. 4 A Voronoi separator S of length 4 and its perimeter $\pi(S)$, depicted in blue. The Voronoi regions are depicted in light colors, objects in respective solid colors. As in Fig. 5, the perimeter $\pi(S)$ may not cross some object p traversed by S . The figure is taken almost verbatim from [5], with slight modifications (Color figure online)

2.4 From Nooses to Separators: Proof of Lemma 3

For the proof of Lemma 3 we need to formalize the connection between nooses in the Voronoi diagram and Voronoi separators in the graph. This connection was largely explored in [5]. We now recall and adjust the basic observations; see Fig. 4 for a visualization.

Suppose G is an connected graph embedded in a sphere Σ , each of its edges has a given length, $\mathcal{S} \subseteq \mathcal{F}$ are independent families of objects in G with $|\mathcal{S}| \geq 4$, and $H = H_{\mathcal{S}}$ is the Voronoi diagram of \mathcal{S} in G . Recall that H is a connected 3-regular multigraph embedded in Σ . Suppose γ is a noose in H . Then γ naturally induces a Voronoi separator $S(\gamma)$ defined as the sequence:

$$\langle p_1, u_1, f_1, v_1, p_2, u_2, f_2, v_2, \dots, p_r, u_r, f_r, v_r \rangle \tag{2}$$

such that

- f_1, \dots, f_r are consecutive branching points of H visited by γ ;
- between f_{i-1} and f_i , γ travels through the face of H corresponding to the object p_i ;
- u_i is the vertex of f_i corresponding to the direction of γ entering f_i ; and
- v_i is the vertex of f_i corresponding to the direction of γ leaving f_i .




Fig. 5 Part of the Voronoi diagram with a spoke and a diamond highlighted. Depicted triangles are the branching points of the Voronoi diagram. Gray dashed curves are the edges of the Voronoi diagram. The spoke of u , which is the shortest path leading from u to the object p to whose region u belongs, is depicted in orange. The diamond induced by the edge $f_1 f_2$ is depicted in black (its interior is grayed). Note that the diamond does not need to cross objects p_1 and p_2 . For instance, as in the figure, it may happen that the path within $\hat{T}(p_2)$ connecting $u_{2,1}$ and $u_{2,2}$ does not intersect p_2 (Color figure online)

Here, by correspondence between vertices of f_i and directions entering/leaving f_i we mean the following. Recall that f_i is a triangular face of G and at the same time it is a degree-3 vertex of H . Edges incident to f_i in H correspond in the dual to edges of f_i in G , thus the vertices of f_i may be naturally associated with the three incidences between f_i and faces of H around it: the vertex of f_i in G incident to edges e_1, e_2 of f_i in G corresponds to the incidence with the face lying between the duals of e_1, e_2 in H . A noose in H may enter/leave f_i in three directions corresponding to these three incidences. Hence, we have a correspondence between the vertices of f_i and the directions.

For a Voronoi separator S as in (2), define its *perimeter* $\pi(S)$ to be the noose constructed by concatenating the following $2r$ curves in order, two for each $i = 1, 2, \dots, r$:

- the unique path in $\hat{T}(p)$ from v_{i-1} to u_i , and
- a curve within f_i from u_i to v_i , obtained by concatenating the segment from u_i to the center of f_i and the segment from the center of f_i to v_i .

Given a noose γ with respect to H , the *canonical version* of γ is the noose $\pi(S(\gamma))$. Note that γ and its canonical version are equivalent as nooses with respect to H . A noose that is equal to its own canonical version shall be called *canonical*. It can be easily seen that if γ is a canonical noose with respect to H , then $\mathbf{enc}(\gamma)$ is the

union of interiors of diamonds induced by edges of H enclosed by γ ; symmetrically for $\text{exc}(\gamma)$. Then the following is immediate.

Lemma 9 *Every sphere-cut decomposition of H for which all nooses are canonical is faithful.*

We are ready to give a proof of Lemma 3.

Proof (of Lemma 3) Apply Lemma 4 to \mathcal{F} and $\ell = s^2 = 10^6 \cdot \left(\frac{1}{\epsilon} \ln \frac{1}{\epsilon}\right)^2$. Note that thus $\ell \geq 10$, as required by Lemma 4. This yields a subfamily $\mathcal{S} \subseteq \mathcal{F}$ and the corresponding Voronoi diagram $H = H_{\mathcal{S}}$ satisfying the following properties:

- (N1) $4 \leq |\mathcal{S}| \leq 2\ell$;
- (N2) every spoke in H has weight with respect to \mathcal{F} upper bounded by $10 \ln \ell \cdot \frac{W}{\ell}$; and
- (N3) every diamond in H has weight with respect to \mathcal{F} upper bounded by $10 \ln \ell \cdot \frac{W}{\ell}$.

Recall that H is a 3-regular connected multigraph embedded in Σ . Let us continue the proof under the assumption that H is bridgeless. The general case when H may have bridges requires attention to a few more technical details but is conceptually no different; we present it later.

Since H is bridgeless, by Theorem 4 we may find a sphere-cut decomposition (T, η, δ) of H of width at most $\sqrt{4.5|V(H)|} \leq \sqrt{18\ell} \leq 5\sqrt{\ell}$. We may further assume that each noose used by (T, η, δ) is canonical, hence (T, η, δ) is faithful by Lemma 9. The following claims now follow easily from the properties of \mathcal{S} provided by Lemma 4. □

Claim 6 *For every pair (x, y) of adjacent nodes of T , the total weight of objects of \mathcal{F} banned by the Voronoi separator $S(\delta(x, y))$ is at most ϵW .*

Proof Since the width of (T, η, δ) is at most $5\sqrt{\ell}$, the noose $\delta(x, y)$ traverses at most $5\sqrt{\ell} = 5s$ faces of H . Each traversal, say of the face corresponding to an object $p \in \mathcal{S}$, is a path contained in the union of two spokes and possibly a path within p . Observe that the objects banned by $S(\delta(x, y))$ are exactly the objects that are in conflict with any of these at most $10\sqrt{\ell}$ spokes. By (N2), each of these spokes has weight with respect to \mathcal{F} upper bounded by $10 \ln \ell \cdot \frac{W}{\ell}$, which means that the total weight of objects banned by $S(\delta(x, y))$ is at most

$$10\sqrt{\ell} \cdot 10 \ln \ell \cdot \frac{W}{\ell} = 100 \ln \ell \cdot \frac{W}{\sqrt{\ell}} = 200 \ln s \cdot \frac{W}{s} \leq \epsilon W,$$

where the last inequality holds due to $\ln s \leq 5 \cdot \ln \frac{1}{\epsilon}$ being satisfied for $\epsilon < \frac{1}{10}$. □

Claim 7 *If (x, y) is a pair of adjacent nodes in T such that x is a leaf, then the total weight of objects of \mathcal{F} contained in $\mathbf{enc}(\delta(y, x))$ is at most ϵW .*

Proof Let $e = \eta^{-1}(x)$. Then $\delta(y, x)$ is the diamond induced by e and $\mathbf{enc}(\delta(y, x))$ is its interior. By (N3), every diamond in H has weight with respect to \mathcal{F} upper bounded by

$$10 \ln \ell \cdot \frac{W}{\ell} \leq \epsilon W,$$

Similarly as in Claim 6. The claim follows. \square

The next claim is the key insight from [5]: Voronoi separators induced by nooses in (T, η, δ) break the intersection graph $\text{IntGraph}(\mathcal{D})$ in the expected way.

Claim 8 *Suppose (x, y) is a pair of adjacent nodes of T and $p, p' \in \mathcal{D}$ are two objects not banned by $S(\delta(x, y))$ such that p and p' are contained in different regions of $\Sigma - \delta(x, y)$. Then p and p' are non-adjacent in $\text{IntGraph}(\mathcal{D})$.*

Proof Since p and p' are contained in different regions of $\Sigma - \delta(x, y)$, they are in particular disjoint. Hence, they are non-adjacent in the intersection graph $\text{IntGraph}(\mathcal{D})$. \square

For an open disk Δ on Σ , define the measure $\mu_{\mathcal{F}}(\Delta)$ as before: $\mu_{\mathcal{F}}(\Delta)$ is the sum of weights of objects $p \in \mathcal{F}$ that are entirely contained in Δ . Since $\epsilon \leq \frac{9}{20}$, Claim 7 verifies that the prerequisites of Lemma 7 are satisfied. Thus, we may apply Lemma 7 to (T, η, δ) , yielding a pair of adjacent nodes (x, y) of T such that within each of the disks $\mathbf{enc}(\delta(x, y))$ and $\mathbf{enc}(\delta(y, x))$ the total weight of objects from \mathcal{F} is at most $\frac{9}{10}W$.

We claim that the Voronoi separator S satisfies all the required conditions. For condition (B1), we have that $\mathcal{D}(S) \subseteq S \subseteq \mathcal{F}$ and all faces traversed by S are branching points of the Voronoi diagram of $S \subseteq \mathcal{D}$; hence, they are \mathcal{D} -important. For condition (B2), the length of $S(\delta(x, y))$ is equal to the length of $\delta(x, y)$, which is at most $3s$. Condition (B3) follows directly from Claim 6. Finally, for condition (B4), Claim 8 together with the construction of S imply that every connected component of $\text{IntGraph}(\mathcal{D}) - \text{Ban}(S)$ contains objects from \mathcal{F} of total weight at most $\frac{9}{10}W$. Hence, S has all the prescribed properties.

We are left with discussing the case when the Voronoi diagram H has bridges. Note that diamonds in H are well-defined also for bridges; hence, as in Claim 7, every diamond in H has weight at most $10 \ln \ell \cdot \frac{W}{\ell} \leq \epsilon W$.

Consider any bridge b in H , say with endpoints f_1 and f_2 ; recall that these are branching points of H . For $t = 1, 2$, let us construct the canonical noose (w.r.t. H) of length 1 traveling through f_t and the unique face of H incident to the bridge b . Call this noose $\beta(f_t, b)$ and orient it so that $\mathbf{enc}(\beta(f_t, b))$ contains the bridge b ; nooses constructed in this manner shall be called *bridge nooses*.

The same calculations as in Claim 6 show that for any bridge noose β , the total weight of objects of \mathcal{F} banned by the Voronoi separator $S(\beta)$ is at most ϵW ; this is

because the length of a bridge noose is 1, which is never larger than $3s$. This means that if for any bridge noose β we have that for both $\mathbf{enc}(\beta)$ and $\mathbf{exc}(\beta)$, the total weight of objects from \mathcal{F} contained in each of these disks is at most $\frac{9}{10}W$, then $S(\beta)$ satisfies all the required conditions. Hence, from now on we assume that for every bridge noose β , either $\mu_{\mathcal{F}}(\mathbf{enc}(\beta))$ or $\mu_{\mathcal{F}}(\mathbf{exc}(\beta))$ exceeds $\frac{9}{10}W$.

Let us inspect any bridge b in H , say with endpoints f_1 and f_2 . Denote $\beta_1 = \beta(f_1, b)$ and $\beta_2 = \beta(f_2, b)$. As $\mathbf{exc}(\beta_1)$ and $\mathbf{exc}(\beta_2)$ are disjoint, it cannot happen that $\mu_{\mathcal{F}}(\mathbf{exc}(\beta_1)) > \frac{9}{10}W$ and $\mu_{\mathcal{F}}(\mathbf{exc}(\beta_2)) > \frac{9}{10}W$. Also, if it happened that $\mu_{\mathcal{F}}(\mathbf{enc}(\beta_1)) > \frac{9}{10}W$ and $\mu_{\mathcal{F}}(\mathbf{enc}(\beta_2)) > \frac{9}{10}W$, then the objects from \mathcal{F} contained in the interior of the diamond induced by b in H (which is equal to $\mathbf{enc}(\beta_1) \cap \mathbf{enc}(\beta_2)$) would have weight more than $\frac{8}{10}W$, a contradiction with the fact that every diamond in H has weight at most ϵW . Therefore, either $\mu_{\mathcal{F}}(\mathbf{enc}(\beta_1)) > \frac{9}{10}W$ and $\mu_{\mathcal{F}}(\mathbf{exc}(\beta_2)) > \frac{9}{10}W$, or $\mu_{\mathcal{F}}(\mathbf{exc}(\beta_1)) > \frac{9}{10}W$ and $\mu_{\mathcal{F}}(\mathbf{enc}(\beta_2)) > \frac{9}{10}W$. Orient b toward f_2 in the former case and toward f_1 in the latter case.

Thus, every bridge of H becomes oriented. Since H has one fewer bridge than bridgeless components, there exists a bridgeless component B of H such that every bridge incident to B is oriented toward its endpoint residing in B .

We first resolve the corner case when B consists of one vertex, say u . Since H is 3-regular, u has three different neighbors v_1, v_2, v_3 such that uv_1, uv_2, uv_3 are the three bridges incident to u . These bridges are oriented toward u , which means that each of the disks

$$\mathbf{exc}(\beta(u, uv_1)), \mathbf{exc}(\beta(u, uv_2)), \mathbf{exc}(\beta(u, uv_3))$$

contains objects from \mathcal{F} of total weight more than $\frac{9}{10}W$. However, each object from \mathcal{F} can be contained in at most 2 of these disks, and a simple counting argument yields a contradiction.

Hence, from now on suppose B is a bridgeless component of H consisting of more than one vertex. Since B is bridgeless and nontrivial and H is 3-regular, every vertex of u of B is incident to at most one bridge of H . Thus, it is easy to see that there is an injective mapping ϕ from bridges incident to B to edges of B so that each bridge b shares an endpoint with $\phi(b)$.

For a bridge b incident to B , let H_b be the component of $H - b$ that is disjoint with B . Now, for every noose γ with respect to B we may define a noose $\hat{\gamma}$ with respect to H so that γ and $\hat{\gamma}$ are equivalent as nooses w.r.t. B , and for every bridge b incident to B , either all of H_b, b , and $\phi(b)$ are contained in $\mathbf{enc}(\gamma)$, or all of them are contained in $\mathbf{exc}(\gamma)$. See Fig. 6 for an example. We may further require that $\hat{\gamma}$ is canonical (as a noose with respect to H).

By Theorem 4, let (T, η, δ) be a sphere-cut decomposition of B of width at most $\sqrt{4.5|V(B)|} \leq 3\sqrt{\ell}$. Replace each noose γ appearing in (T, η, δ) by $\hat{\gamma}$; it is easy to see that thus (T, η, δ) is still a sphere-cut decomposition of B and becomes faithful. We would like now to apply the whole reasoning to (T, η, δ) and B . Claim 6 works exactly as before, because the only condition we needed is an upper bound on the weights of spokes in H .

However, Claim 7 may fail, because disks of the form $\mathbf{enc}(\delta(y, x))$, where x is a leaf of T and y is the node of T adjacent to it, are no longer diamonds in H . More




Fig. 6 Construction of $\hat{\gamma}$ (a noose w.r.t. H) out of γ (a noose w.r.t. B). The left panel depicts the bridgeless component B alone together with the noose γ . The right panel depicts the whole graph H , with components $H_{b_1}, H_{b_2}, H_{b_3}$ and bridges b_1, b_2, b_3 reintroduced, and the adjusted noose $\hat{\gamma}$. The mapping ϕ is depicted with gray, dashed arrows

precisely, if $e = \eta^{-1}(x)$ is the edge of H associated with x and $\Delta(e)$ is the interior of the diamond of H induced by it, then the following two cases may happen:

- if e is not in the image of ϕ , then $\mathbf{enc}(\delta(y, x)) = \Delta(e)$;
- if $e = \phi(b)$ for some bridge b incident to B , say with u being the endpoint of b in B , then $\mathbf{enc}(\delta(y, x))$ is the union of $\Delta(e)$, the disk $\mathbf{enc}(\beta(u, b))$, and the intersection of their boundaries.

In the first case, we already know that $\mu_{\mathcal{F}}(\mathbf{enc}(\delta(y, x))) \leq \epsilon W < \frac{9}{20} W$. In the second case, observe that every object of \mathcal{F} contained in $\mathbf{enc}(\delta(y, x))$ is either contained $\Delta(e)$ or is not contained in $\mathbf{enc}(\beta(u, b))$. The total weight of objects satisfying the first condition is at most ϵW , whereas the total weight of objects satisfying the second condition is less than $\frac{W}{10}$, because $\mu_{\mathcal{F}}(\mathbf{enc}(\beta(u, b))) > \frac{9}{10} W$. Hence,

$$\mu_{\mathcal{F}}(\mathbf{enc}(\delta(y, x))) \leq \epsilon W + \frac{W}{10} < \frac{9}{20} W.$$

Summing up, in both cases we have verified that $\mu_{\mathcal{F}}(\mathbf{enc}(\delta(y, x))) < \frac{9}{20} W$, so we may apply Lemma 7 and proceed as in the case without bridges. \square

3 A QPTAS for MAXIMUM WEIGHT INDEPENDENT SET OF OBJECTS

In this section we use Lemma 2 to design a QPTAS for MWISO, that is, we prove Theorem 1. Recall that the setting is as follows. The input is (G, \mathcal{D}) , where G is a graph embedded in a sphere Σ together with a family of objects \mathcal{D} , each being a connected subgraph of G with a prescribed positive weight. Moreover, we are given an accuracy parameter $\epsilon > 0$ and we may assume w.l.o.g. that $\epsilon < \frac{1}{10}$. The goal is to find an independent subfamily $\mathcal{F} \subseteq \mathcal{D}$ with the largest possible weight; more precisely, the algorithm shall compute a solution of weight at least $(1 - \epsilon)$ times the optimum. Let $n = |V(G)|$ and $N = |\mathcal{D}|$; w.l.o.g. we assume $N \geq 2$.

First, we assume that all objects in \mathcal{D} have weights between 1 and $M = 2\epsilon^{-1}N$. For the given instance (G, \mathcal{D}) let us define

$$\text{span}(\mathcal{D}) = \frac{\max_{p \in \mathcal{D}} \mathbf{w}(p)}{\min_{p \in \mathcal{D}} \mathbf{w}(p)}.$$

The following standard claim shows that we may assume that the span in the input instance is bounded polynomially in N and $1/\epsilon$; more precisely, by $M = 2\epsilon^{-1}N$.

Claim 9 *Suppose there is a QPTAS as described in Theorem 1, but working on instances (G, \mathcal{D}) that additionally satisfy $\text{span}(\mathcal{D}) < 2\epsilon^{-1}|\mathcal{D}|$. Then Theorem 1 holds.*

Proof Let $M = 2\epsilon^{-1}|\mathcal{D}|$. For every object $p \in \mathcal{D}$ construct a subfamily $\mathcal{D}_p \subseteq \mathcal{D}$ consisting of all objects whose weight is not larger than $\mathbf{w}(p)$ but larger than $\mathbf{w}(p)/M$. Clearly, $\text{span}(\mathcal{D}_p) < M$. We claim that there exists $p \in \mathcal{D}$ such that the optimum solution for (G, \mathcal{D}_p) has weight at most $(1 - \epsilon/2)$ smaller than the optimum solution for (G, \mathcal{D}) . If this is the case, then it suffices to run the assumed QPTAS on each of the instances (G, \mathcal{D}_p) , for $p \in \mathcal{D}$, using accuracy parameter $\epsilon/2$, and output the heaviest of the solutions found.

Let $\mathcal{F}_{\text{OPT}} \subseteq \mathcal{D}$ be the optimum solution for (G, \mathcal{D}) . Let p be the heaviest object in \mathcal{F}_{OPT} . Consider $\mathcal{F} = \{q \in \mathcal{F}_{\text{OPT}} : \mathbf{w}(q) > \mathbf{w}(p)/M\}$ and observe that $\mathcal{F} \subseteq \mathcal{D}_p$, thus \mathcal{F} is a solution for (G, \mathcal{D}_p) . On the other hand, we have

$$\mathbf{w}(\mathcal{F}_{\text{OPT}} \setminus \mathcal{F}) \leq \mathbf{w}(p)/M \cdot |\mathcal{F}_{\text{OPT}} \setminus \mathcal{F}| \leq \mathbf{w}(p)/M \cdot |\mathcal{D}| = \epsilon/2 \cdot \mathbf{w}(p) \leq \epsilon/2 \cdot \mathbf{w}(\mathcal{F}_{\text{OPT}}),$$

where the last inequality follows from $p \in \mathcal{F}_{\text{OPT}}$. Hence $\mathbf{w}(\mathcal{F}) \geq (1 - \epsilon/2)\mathbf{w}(\mathcal{F}_{\text{OPT}})$ and we are done. □

By Claim 9 in order to prove Theorem 1 it suffices to work under the assumption that $\text{span}(\mathcal{D}) < M$. Hence, by rescaling the weights we may assume that all the weights of objects in \mathcal{D} are at least 1 and smaller than M .

Before we proceed to the algorithm, we fix the following parameters:

$$d_{\max} = 10 \ln(MN), \quad \hat{\epsilon} = \frac{\epsilon}{d_{\max}}, \quad s = 10^3 \cdot \frac{1}{\hat{\epsilon}} \ln \frac{1}{\hat{\epsilon}}. \quad (3)$$

Parameter d_{\max} is the maximum recursion depth of the algorithm; note that $d_{\max} = \mathcal{O}(\log(N/\epsilon))$. Next, $\hat{\epsilon} = \mathcal{O}(\epsilon/\log(N))$ is the refined accuracy parameter which will be used throughout the recursion instead of ϵ . Similarly as in [1, 2, 4], intuitively we lose a factor of $1 - \hat{\epsilon}$ in each recursion level which yields an overall approximation ratio of $(1 - \hat{\epsilon})^{d_{\max}} = (1 - \epsilon/d_{\max})^{d_{\max}} = 1 - \mathcal{O}(\epsilon)$. Let us stress that although the algorithm uses recursion and the number of objects changes in subsequent recursive calls, the values of $d_{\max}, \hat{\epsilon}, s$ are fixed as above and their definitions always refer to the initial number of objects.

We now explain the algorithm; it is also summarized using pseudocode as Algorithm 1. Let us fix an optimum solution \mathcal{F}_{OPT} and denote $W = \mathbf{w}(\mathcal{F}_{\text{OPT}})$. We

shall analyze \mathcal{F}_{OPT} , which will lead to the formulation of the algorithm as a recursive search for \mathcal{F}_{OPT} .

We would like to use Lemma 2 to guess a Voronoi separator that breaks \mathcal{F}_{OPT} in a balanced way. However, we first need make sure that every object in question constitutes only a small fraction of W . This is done by a standard method of guessing exactly “heavy” objects in the solution, whose number is small, and proceeding only with the “light” ones.

More precisely, call an object $p \in \mathcal{F}_{\text{OPT}}$ *heavy* if $w(p) > s^{-2}W$. Observe that the number of heavy objects in \mathcal{F}_{OPT} is at most s^2 , hence there are at most N^{s^2} possible ways to select those heavy objects from \mathcal{D} . The algorithm branches into all possible such ways, in each branch fixing a different candidate for the set of heavy objects. Hence, by increasing the number of subproblems by a multiplicative factor N^{s^2} we may assume that the algorithm fixes the set \mathcal{F}_{hv} consisting of all heavy objects in \mathcal{F}_{OPT} . Let \mathcal{D}' be obtained from \mathcal{D} by removing \mathcal{F}_{hv} and all objects intersecting any object from \mathcal{F}_{hv} , and let $\mathcal{F}'_{\text{OPT}} = \mathcal{F}_{\text{OPT}} - \mathcal{F}_{\text{hv}}$. Note that $w(\mathcal{F}'_{\text{OPT}}) \leq W$ and $w(p) \leq s^{-2}W$ for all $p \in \mathcal{F}'_{\text{OPT}}$.

We may now apply Lemma 2 to $\mathcal{F}'_{\text{OPT}} \subseteq \mathcal{D}'$ with W being the upper bound on its weight. Thus, in time $N^{O(s)} \cdot n^{O(1)}$ we may compute a family \mathbb{X} consisting of subsets of \mathcal{D}' with $|\mathbb{X}| \leq 6^{3s}N^{15s}$ and satisfying the following property: there exists $\mathcal{X} \in \mathbb{X}$ such that $w(\mathcal{F}'_{\text{OPT}} \setminus \mathcal{X}) \leq \epsilon W$ and within each connected component of the graph $\text{IntGraph}(\mathcal{D}') - \mathcal{X}$ the total weight of objects from $\mathcal{F}'_{\text{OPT}}$ does not exceed $\frac{9}{10}W$. By branching into all the members of \mathbb{X} , via increasing the number of subproblems by a multiplicative factor $|\mathbb{X}|$ we may henceforth assume that the algorithm fixes \mathcal{X} with properties as above.

For a fixed choice of \mathcal{F}_{hv} and \mathcal{X} as above, let us inspect the connected components of $\text{IntGraph}(\mathcal{D}') - \mathcal{X}$; let their vertex sets be $\mathcal{D}_1, \dots, \mathcal{D}_k$. We apply the algorithm recursively to the instances (G, \mathcal{D}_i) for $i = 1, \dots, k$, yielding independent families $\mathcal{F}_1, \dots, \mathcal{F}_k$ with $\mathcal{F}_i \subseteq \mathcal{D}_i$. We record the family $\mathcal{F} = \mathcal{F}_{\text{hv}} \cup \bigcup_{i=1}^k \mathcal{F}_i$ as a candidate for the solution; it is straightforward to see from the construction that this family is independent. Finally, as the final solution we output the heaviest among the recorded candidates; that is, the heaviest solution found for all choices of \mathcal{F}_{hv} and \mathcal{X} . We remark that if for some choice of \mathcal{F}_{hv} it turned out that $\mathcal{D}' = \emptyset$, i.e., every object intersects some objects from \mathcal{F}_{hv} , then \mathbb{X} contains only one choice of \mathcal{X} being \emptyset , hence we include $\mathcal{F} = \mathcal{F}_{\text{hv}}$ among the candidates without invoking any recursive calls.

The base case of the recursion is provided by trimming it at level d_{max} . More precisely, all subcalls at depth larger than d_{max} return empty solutions. This concludes the description of the algorithm; as mentioned, it is summarized using pseudocode as Algorithm 1.

Algorithm 1: Algorithm AlgMWISO

Input: An instance (G, \mathcal{D}) , recursion depth d
Output: An independent family $\mathcal{F} \subseteq \mathcal{D}$

```

if  $d > d_{\max}$  then
  return  $\emptyset$ 
 $\mathcal{F} \leftarrow \emptyset$ 
forall  $\mathcal{F}_{\text{hv}}$ : independent subfamily of  $\mathcal{D}$  with  $|\mathcal{F}_{\text{hv}}| \leq s^2$  do
   $\mathcal{D}' \leftarrow \mathcal{D}$  – (all objects that intersect any object of  $\mathcal{F}_{\text{hv}}$ )
   $\mathbb{X} \leftarrow$  family computed for  $\mathcal{D}'$  using Lemma 2
  forall  $\mathcal{X} \in \mathbb{X}$  do
     $\mathcal{D}_1, \dots, \mathcal{D}_k \leftarrow$  vertex sets of the connected components of  $\text{IntGraph}(\mathcal{D}') - \mathcal{X}$ 
    for  $i = 1$  to  $k$  do
       $\mathcal{F}_i \leftarrow \text{AlgMWISO}(G, \mathcal{D}_i, d + 1)$ 
       $\mathcal{F}_{\text{cand}} \leftarrow \mathcal{F}_{\text{hv}} \cup \bigcup_{i=1}^k \mathcal{F}_i$ 
      if  $w(\mathcal{F}_{\text{cand}}) > w(\mathcal{F})$  then
         $\mathcal{F} \leftarrow \mathcal{F}_{\text{cand}}$ 
  return  $\mathcal{F}$ 

```

Running time analysis. On a high level, the above algorithm runs in time $2^{\text{poly}(1/\epsilon, \log N)} \cdot n^{\mathcal{O}(1)}$ because at each node of the recursion tree the algorithm uses polynomial time and calls itself on $N^{\mathcal{O}(s^2)}$ subinstances, where $s = \text{poly}(1/\epsilon, \log N)$. Together with the bound of $d_{\max} = \mathcal{O}(\log(N/\epsilon))$ on the recursion depth this yields the promised running time.

Formally, we argue as follows. Throughout this analysis N denotes the initial number of objects. It is straightforward to see that the computation at each recursive call, excluding time spent on executing the invoked subcalls, takes time $N^{\mathcal{O}(s)} \cdot n^{\mathcal{O}(1)}$. Therefore, to bound the running time it suffices to bound the total number of nodes in the recursion tree. As the recursion is trimmed at depth d_{\max} , its height is at most d_{\max} . Further, in each call we investigate:

- at most N^{s^2} ways to select heavy objects \mathcal{F}_{hv} from \mathcal{D} ;
- at most $6^{3s} N^{15s}$ ways to select the family \mathcal{X} ; and
- at most N connected components of $\text{IntGraph}(\mathcal{D}') - \mathcal{X}$.

This yields at most $6^{3s} \cdot N^{s^2+18s+1} \leq N^{2s^2}$ instances on which a recursive subcall is invoked. This means that each node of the recursion tree has at most N^{2s^2} children, which together with the height of at most $d_{\max} = 10 \ln(NM)$ yields an upper bound of $N^{20s^2 \ln(NM)}$ on the number of nodes in the recursion tree. Since the computation at each node takes time $N^{\mathcal{O}(s)} \cdot n^{\mathcal{O}(1)}$, we arrive at the total running time of

$$N^{\mathcal{O}(s^2 \log(NM))} \cdot n^{\mathcal{O}(1)} = 2^{\text{poly}(1/\epsilon, \log N)} \cdot n^{\mathcal{O}(1)},$$

as promised.

Approximation factor analysis. As mentioned before, the claimed approximation ratio follows as intuitively in each of the d_{\max} recursion levels we lose a factor of $1 - \hat{\epsilon}$, accumulating to $(1 - \hat{\epsilon})^{d_{\max}} = 1 - \mathcal{O}(\epsilon)$ overall. Formally, we argue as follows. It is clear that the family output by the algorithm is indeed independent.

In order to prove that its weight is at least $(1 - \epsilon)W$, it suffices to show the following claim and apply it for $d = 0$; i.e., for the initial instance (G, \mathcal{D}) . Recall here that \mathcal{F}_{OPT} is a fixed optimum solution for the whole instance.

Claim 10 *Let $(G, \tilde{\mathcal{D}})$ be an instance on which the algorithm is called in the recursion tree, say at depth d . Suppose further that $\mathbf{w}(\mathcal{F}_{\text{OPT}} \cap \tilde{\mathcal{D}}) \leq (9/10)^d \cdot W$. Then the call of the algorithm to $(G, \tilde{\mathcal{D}})$ returns an independent subfamily $\tilde{\mathcal{F}} \subseteq \tilde{\mathcal{D}}$ with $\mathbf{w}(\tilde{\mathcal{F}}) \geq (1 - d'\epsilon)\mathbf{w}(\mathcal{F}_{\text{OPT}} \cap \tilde{\mathcal{D}})$, where $d' = d_{\max} - d$.*

Proof Denote $\tilde{\mathcal{F}}_{\text{OPT}} = \mathcal{F}_{\text{OPT}} \cap \tilde{\mathcal{D}}$. We prove the claim by induction on the depth d , starting with the case $d \geq d_{\max}$ and then decreasing d . For the base of the induction, for $d \geq d_{\max}$ we have

$$(9/10)^d \leq (1 - 1/10)^{10 \ln(MN)} < e^{-\ln(MN)} = \frac{1}{MN}.$$

On the other hand, \mathcal{F}_{OPT} contains at most N objects of weight less than M each, hence $W < MN$. Combining these two observations we infer that $(9/10)^d \cdot W < 1$. Since each object of \mathcal{D} has weight at least 1, we conclude that $\tilde{\mathcal{F}}_{\text{OPT}} = \emptyset$, so the claim holds vacuously.

Suppose now that $d < d_{\max}$. Let $\tilde{W} = \mathbf{w}(\tilde{\mathcal{F}}_{\text{OPT}})$; by assumption, $\tilde{W} \leq (9/10)^d \cdot W$. Call an object $p \in \tilde{\mathcal{F}}_{\text{OPT}}$ *heavy* if $\mathbf{w}(p) \geq s^{-2}\tilde{W}$ and let $\tilde{\mathcal{F}}_{\text{hv}} \subseteq \tilde{\mathcal{F}}_{\text{OPT}}$ be the family of heavy objects in $\tilde{\mathcal{F}}_{\text{OPT}}$. Observe that $|\tilde{\mathcal{F}}_{\text{hv}}| \leq s^2$, so in the first loop of the algorithm there is an iteration in which we correctly fix the set $\tilde{\mathcal{F}}_{\text{hv}}$. We continue the reasoning under this assumption. As in the description, let $\tilde{\mathcal{D}}'$ be \mathcal{D} with all objects intersecting any object from \mathcal{F}_{hv} removed.

Now let $\tilde{\mathcal{F}}'_{\text{OPT}} = \tilde{\mathcal{F}}_{\text{OPT}} \setminus \tilde{\mathcal{F}}_{\text{hv}}$. Since $\mathbf{w}(p) < s^{-2}\tilde{W}$ for all $p \in \tilde{\mathcal{F}}'_{\text{OPT}}$, we may apply Lemma 2 to $\tilde{\mathcal{F}}'_{\text{OPT}}$ and \tilde{W} as the upper bound on its weight, and conclude that the enumerated family \times contains at least one member \mathcal{X} satisfying the properties as in the description of the algorithm. In particular, in one iteration of the second loop the algorithm correctly fixes such \mathcal{X} and proceeds with it. From now on we continue under this assumption.

Having fixed $\tilde{\mathcal{F}}_{\text{hv}}$ and \mathcal{X} , the algorithm investigates the connected components $\tilde{\mathcal{D}}_1, \dots, \tilde{\mathcal{D}}_k$ of $\text{IntGraph}(\tilde{\mathcal{D}}') - \mathcal{X}$ and applies itself recursively to each instance $(G, \tilde{\mathcal{D}}_i)$, for $i = 1, \dots, k$, yielding families $\tilde{\mathcal{F}}_i \subseteq \tilde{\mathcal{D}}_i$. By the properties of \mathcal{X} , we have that

$$\mathbf{w}(\mathcal{F}_{\text{OPT}} \cap \tilde{\mathcal{D}}_i) \leq (9/10) \cdot \mathbf{w}(\mathcal{F}_{\text{OPT}} \cap \tilde{\mathcal{D}}) \leq (9/10)^{d+1} \cdot W,$$

for each $i \in \{1, \dots, k\}$. Since recursive subcalls to instances $(G, \tilde{\mathcal{D}}_i)$ are at level $d + 1$ in the recursion tree, by the induction assumption we infer that

$$\mathbf{w}(\mathcal{F}_i) \geq (1 - (d' - 1)\epsilon) \cdot \mathbf{w}(\mathcal{F}_{\text{OPT}} \cap \tilde{\mathcal{D}}_i) \tag{4}$$

for each $i \in \{1, \dots, k\}$. Since the total weight of objects of $\tilde{\mathcal{F}}_{\text{OPT}} \cap \mathcal{X}$ is at most $\epsilon\tilde{W}$, we have that

$$\mathbf{w}(\tilde{\mathcal{F}}_{\text{hv}}) + \sum_{i=1}^k \mathbf{w}(\tilde{\mathcal{F}}_{\text{OPT}} \cap \tilde{\mathcal{D}}_i) \geq (1 - \hat{\epsilon})\tilde{W}. \tag{5}$$

Recall that the algorithm returns a family of weight not smaller than the weight of $\tilde{\mathcal{F}} = \tilde{\mathcal{F}}_{\text{hv}} \cup \bigcup_{i=1}^k \tilde{\mathcal{F}}_i$. By combining (4) with (5) we have

$$\begin{aligned} \mathbf{w}(\tilde{\mathcal{F}}) &= \mathbf{w}(\tilde{\mathcal{F}}_{\text{hv}}) + \sum_{i=1}^k \mathbf{w}(\tilde{\mathcal{F}}_i) \\ &\geq \mathbf{w}(\tilde{\mathcal{F}}_{\text{hv}}) + (1 - (d' - 1)\hat{\epsilon}) \cdot \sum_{i=1}^k \mathbf{w}(\tilde{\mathcal{F}}_{\text{OPT}} \cap \tilde{\mathcal{D}}_i) \\ &\geq (1 - (d' - 1)\hat{\epsilon}) \cdot \left(\mathbf{w}(\tilde{\mathcal{F}}_{\text{hv}}) + \sum_{i=1}^k \mathbf{w}(\tilde{\mathcal{F}}_{\text{OPT}} \cap \tilde{\mathcal{D}}_i) \right) \\ &\geq (1 - (d' - 1)\hat{\epsilon})(1 - \hat{\epsilon})\tilde{W} \geq (1 - d'\hat{\epsilon})\tilde{W}. \end{aligned}$$

This concludes the proof of Theorem 1. □

4 A QPTAS for MINIMUM WEIGHT DISTANCE SET COVER

In this section we prove Theorem 2. We first state and prove a suitable variant of Lemma 2; its proof will follow from Lemma 3 in a similar way.

Lemma 10 (Separator Lemma for MWDCS) *Let G be an n -vertex planar graph, \mathcal{D} be a weighted set of vertices of G with $|\mathcal{D}| = N$, \mathcal{C} be a set of vertices of G , and r be a nonnegative real. Let $0 < \epsilon < \frac{1}{20}$ and denote $s = 10^3 \cdot \frac{1}{\epsilon} \ln \frac{1}{\epsilon}$. Then there is a family \mathbb{X} consisting of quadruples of the form $(\mathcal{D}_1, \mathcal{D}_2, \mathcal{C}_1, \mathcal{C}_2)$ such that:*

- (D1) *For each $(\mathcal{D}_1, \mathcal{D}_2, \mathcal{C}_1, \mathcal{C}_2) \in \mathbb{X}$, we have that $\mathcal{D}_1, \mathcal{D}_2 \subseteq \mathcal{D}$, $\mathcal{C}_1, \mathcal{C}_2 \subseteq \mathcal{C}$, and $\mathcal{C}_1 \cup \mathcal{C}_2 = \mathcal{C}$;*
- (D2) *$|\mathbb{X}| \leq 6^{3s}N^{15s}$ and \mathbb{X} can be computed in time $N^{\mathcal{O}(s)} \cdot n^{\mathcal{O}(1)}$; and*
- (D3) *For every real $W \geq 0$ and subset $\mathcal{F} \subseteq \mathcal{D}$ such that \mathcal{F} r -covers \mathcal{C} , $\mathbf{w}(\mathcal{F}) \leq W$, and $\mathbf{w}(p) \leq s^{-2}W$ for each $p \in \mathcal{F}$, there exists a quadruple $(\mathcal{D}_1, \mathcal{D}_2, \mathcal{C}_1, \mathcal{C}_2) \in \mathbb{X}$ such that:*
 - $\mathbf{w}(\mathcal{F} \cap \mathcal{D}_1 \cap \mathcal{D}_2) \leq \epsilon W$;
 - $\mathbf{w}(\mathcal{F} \cap \mathcal{D}_1) \leq \frac{19}{20}W$ and $\mathbf{w}(\mathcal{F} \cap \mathcal{D}_2) \leq \frac{19}{20}W$; and
 - $\mathcal{F} \cap \mathcal{D}_1$ r -covers \mathcal{C}_1 and $\mathcal{F} \cap \mathcal{D}_2$ r -covers \mathcal{C}_2 .

Proof As before, we may assume we are given an embedding of G in a sphere Σ , G is triangulated using edges of infinite lengths, and distances in G are unique.

Consider applying Lemma 3 to \mathcal{D} (where we interpret vertices in \mathcal{D} as single-vertex objects) and a subset $\mathcal{F} \subseteq \mathcal{D}$ as in (D3). This yields a Voronoi separator S of

length at most $3s$ such that all faces traversed by S are \mathcal{D} -important and S breaks \mathcal{F} in a balanced way as described in Lemma 3. As explained in the proof of Lemma 2, given only \mathcal{D} there are at most $6^{3s}N^{15s}$ ways to choose such a Voronoi separator S and in time $N^{\mathcal{O}(s)} \cdot n^{\mathcal{O}(1)}$ we can enumerate a family \mathcal{N} of at most $6^{3s}N^{15s}$ candidates for S .

Now, construct \mathbb{X} by including the following quadruple $(\mathcal{D}_1, \mathcal{D}_2, \mathcal{C}_1, \mathcal{C}_2)$ for each $S \in \mathcal{N}$. Let $\gamma = \pi(S)$ be the perimeter of S (recall that the perimeter of S is the noose “following” S ; we defined it in Sect. 2.4). Then define $(\mathcal{D}_1, \mathcal{D}_2, \mathcal{C}_1, \mathcal{C}_2)$ as follows:

- \mathcal{D}_1 is the union of $\text{Ban}(S)$ and all vertices of \mathcal{D} contained in $\mathbf{enc}(\gamma)$;
- \mathcal{D}_2 is the union of $\text{Ban}(S)$ and all vertices of \mathcal{D} contained in $\mathbf{exc}(\gamma)$;
- \mathcal{C}_1 comprises all vertices of \mathcal{C} contained in $\mathbf{enc}(\gamma) \cup \gamma$; and
- \mathcal{C}_2 comprises all vertices of \mathcal{C} contained in $\mathbf{exc}(\gamma) \cup \gamma$.

This concludes the construction of \mathbb{X} ; we are left with verifying its properties. Conditions (D1) and (D2) are clear from the construction, so we are left with condition (D3). Suppose W and \mathcal{F} are as in the statement of condition (D3).

For the first assertion, observe that $\mathcal{D}_1 \cap \mathcal{D}_2 = \text{Ban}(S)$ and by Lemma 3, condition (B3), we have that $\mathbf{w}(\mathcal{F} \cap \text{Ban}(S)) \leq \epsilon W$. The claim follows.

For the second assertion, observe that $\mathcal{F} \cap \mathcal{D}_1$ comprises the vertices of \mathcal{F} contained in $\mathbf{enc}(\gamma)$ and the vertices of \mathcal{F} banned by S . By Lemma 3, conditions (B3) and (B4), the weights of these sets of vertices are at most $\frac{9}{10}W$ and at most $\epsilon W \leq \frac{W}{20}$, respectively. It follows that $\mathbf{w}(\mathcal{F} \cap \mathcal{D}_1) \leq \frac{19}{20}W$; a symmetric reasoning shows that $\mathbf{w}(\mathcal{F} \cap \mathcal{D}_2) \leq \frac{19}{20}W$ as well.

We are left with the third assertion. Take any vertex $c \in \mathcal{C}_1$. Let p be the vertex from \mathcal{F} that is closest to c and let P be the shortest path connecting c with p . Since \mathcal{F} r -covers \mathcal{C} , we have that $\text{length}(P) = \text{dist}(p, c) \leq r$. If $p \in \mathcal{F} \cap \mathcal{D}_1$ then we are done, hence assume that $p \in \mathcal{F} \setminus \mathcal{D}_1$. In particular, p is contained in $\mathbf{exc}(\gamma)$ and p is not banned by S . Since $c \in \mathcal{C}_1$, c is contained in $\mathbf{enc}(\gamma) \cup \gamma$, so P has to cross γ at some vertex, say w . Let $q \in \mathcal{D}(S)$ be the vertex on the Voronoi separator S that is the closest to w ; recall that $\mathcal{D}(S) \subseteq \mathcal{F}$, hence $q \in \mathcal{F}$. Since p is not banned by S , we have that $\text{dist}(w, p) > \text{dist}(w, q)$. As P is the shortest path connecting c and p , we infer that

$$r \geq \text{dist}(c, p) = \text{dist}(c, w) + \text{dist}(w, p) > \text{dist}(c, w) + \text{dist}(w, q) \geq \text{dist}(c, q).$$

This means that c is r -covered by q . But q is banned by S due to $q \in \mathcal{D}(S) \subseteq \text{Ban}(S)$, which implies that c is r -covered by \mathcal{F}_1 . Since c was chosen arbitrarily, we infer that \mathcal{C}_1 is r -covered by $\mathcal{F} \cap \mathcal{D}_1$, and a symmetric reasoning shows that \mathcal{C}_2 is r -covered by $\mathcal{F} \cap \mathcal{D}_2$. □

Reduction of the weight span. Similarly as in Sect. 3, we first bound the weight span of the input instance, defined as

$$\text{span}(\mathcal{D}) = \frac{\max_{p \in \mathcal{D}} \mathbf{w}(p)}{\min_{p \in \mathcal{D}} \mathbf{w}(p)}.$$

by a polynomial of $|\mathcal{D}|$ and $1/\epsilon$.

Claim 11 *Suppose there is a QPTAS as described in Theorem 2, but working on instances $(G, \mathcal{D}, \mathcal{C}, r)$ that additionally satisfy $\text{span}(\mathcal{D}) \leq 2\epsilon^{-1}|\mathcal{D}|$. Then Theorem 2 holds.*

Proof Let $M = 2\epsilon^{-1}|\mathcal{D}|$. For every vertex $p \in \mathcal{D}$ construct a weighted set of vertices \mathcal{D}_p by modifying \mathcal{D} as follows: remove all vertices of weight larger than $w(p)$ and for every vertex of weight smaller than $w(p)/M$, increase its weight to $w(p)/M$. For clarity, the weight function in \mathcal{D}_p will be denoted by $w_p(\cdot)$. That is, for each $q \in \mathcal{D}$ we have:

- if $w(p) < w(q)$, then $q \notin \mathcal{D}_p$;
- if $w(p)/M < w(q) \leq w(p)$, then $q \in \mathcal{D}_p$ and $w_p(q) = w(q)$; and
- if $w(q) \leq w(p)/M$, then $q \in \mathcal{D}_p$ and $w_p(q) = w(p)/M$.

Clearly $\text{span}(\mathcal{D}_p) \leq M$. We claim that there exists $p \in \mathcal{D}$ such that the optimum solution for $(G, \mathcal{D}_p, \mathcal{C}, r)$ has weight at most $(1 + \epsilon/2)$ larger than the optimum solution for $(G, \mathcal{D}, \mathcal{C}, r)$. If this is the case, then it suffices to run the assumed QPTAS on each of the instances $(G, \mathcal{D}_p, \mathcal{C}, r)$, for $p \in \mathcal{D}$, using accuracy parameter $\epsilon/2$, and output the lightest of the solutions found. Indeed, instances $(G, \mathcal{D}_p, \mathcal{C}, r)$ are derived from $(G, \mathcal{D}, \mathcal{C}, r)$ only by dropping some vertices and increasing the weights of some others, so every solution for $(G, \mathcal{D}_p, \mathcal{C}, r)$ naturally projects to a solution for $(G, \mathcal{D}, \mathcal{C}, r)$ of not larger weight.

Let $\mathcal{F}_{\text{OPT}} \subseteq \mathcal{D}$ be the optimum solution for $(G, \mathcal{D}, \mathcal{C}, r)$. Let p be the heaviest vertex in \mathcal{F}_{OPT} . Observe that \mathcal{F}_{OPT} is still a solution in $(G, \mathcal{D}_p, \mathcal{C}, r)$, yet its weight with respect to $w_p(\cdot)$ may be larger than with respect to $w(\cdot)$. However, for each vertex $q \in \mathcal{F}_{\text{OPT}}$ we have

$$w_p(q) \leq w(q) + w(p)/M.$$

Thus,

$$\begin{aligned} w_p(\mathcal{F}_{\text{OPT}}) &\leq w(\mathcal{F}_{\text{OPT}}) + |\mathcal{F}_{\text{OPT}}| \cdot w(p)/M \\ &= w(\mathcal{F}_{\text{OPT}}) + \epsilon/2 \cdot w(p) \cdot \frac{|\mathcal{F}_{\text{OPT}}|}{|\mathcal{D}|} \leq (1 + \epsilon/2) \cdot w(\mathcal{F}_{\text{OPT}}), \end{aligned}$$

where the last inequality follows from $p \in \mathcal{F}_{\text{OPT}}$. □

Denoting $M = 2\epsilon^{-1}N = 2\epsilon^{-1}|\mathcal{D}|$, by Claim 11 in order to prove Theorem 2 it suffices to work under the assumption that $\text{span}(\mathcal{D}) < M$. By rescaling the weights from now on we assume that all the weights of vertices in \mathcal{D} are at least 1 and smaller than M .

Parameters. Similarly as in Sect. 3, we set parameters as follows:

$$d_{\max} = 20 \ln(MN), \quad \hat{\epsilon} = \frac{\epsilon}{2d_{\max}}, \quad s = 10^3 \cdot \frac{1}{\hat{\epsilon}} \ln \frac{1}{\hat{\epsilon}}. \quad (6)$$

Note that compared to the proof of Theorem 1, the recursion depth d_{\max} is twice larger and the rescaled accuracy parameter $\hat{\epsilon}$ is four times smaller. This is for technical reasons.

Description of the algorithm. We perform a similar recursive scheme as in Sect. 3, but using Lemma 10 instead of Lemma 2. The algorithm is summarized using pseudocode as Algorithm 2. We now explain it viewing it as a recursive search for the optimum solution.

Recall that we are given an instance $(G, \mathcal{D}, \mathcal{C}, r)$ and an accuracy parameter ϵ , and we have fixed parameters $d_{\max}, \hat{\epsilon}, s$ as in (6). Fix an optimum solution \mathcal{F}_{OPT} and denote $W = \mathbf{w}(\mathcal{F}_{\text{OPT}})$.

We first guess heavy vertices in the solution as in the proof of Theorem 1. Call a vertex $p \in \mathcal{F}_{\text{OPT}}$ heavy if $\mathbf{w}(p) > s^{-2}W$. Observe that \mathcal{F}_{OPT} contains at most s^2 heavy vertices; denote them by \mathcal{F}_{hv} . The algorithm iterates through all subsets of \mathcal{D} of size at most s^2 , in each case fixing the considered subset as \mathcal{F}_{hv} . In one of the cases \mathcal{F}_{hv} will be fixed correctly, hence we may proceed with the assumption that the algorithm knows \mathcal{F}_{hv} .

Let $\mathcal{F}'_{\text{OPT}} = \mathcal{F}_{\text{OPT}} \setminus \mathcal{F}_{\text{hv}}$ and $\mathcal{D}' = \mathcal{D} \setminus \mathcal{F}_{\text{hv}}$. Further, let \mathcal{C}' be constructed from \mathcal{C} by removing all vertices that are r -covered by \mathcal{F}_{hv} . Note that \mathcal{C}' is r -covered by $\mathcal{F}'_{\text{OPT}}$. Moreover, observe that $\mathbf{w}(p) \leq s^{-2}W$ for each $p \in \mathcal{F}'_{\text{OPT}}$, hence we may apply Lemma 10 to \mathcal{D}' and \mathcal{C}' . Thus, in time $N^{\mathcal{O}(s)} \cdot n^{\mathcal{O}(1)}$ we construct a family \mathbb{X} consisting of at most $6^{3s}N^{15s}$ quadruples with the following guarantee: there exists a quadruple $(\mathcal{D}_1, \mathcal{D}_2, \mathcal{C}_1, \mathcal{C}_2) \in \mathbb{X}$ such that

- $\mathcal{D}_1, \mathcal{D}_2 \subseteq \mathcal{D}', \mathcal{C}_1, \mathcal{C}_2 \subseteq \mathcal{C}'$, and $\mathcal{C}_1 \cup \mathcal{C}_2 = \mathcal{C}'$;
- $\mathbf{w}(\mathcal{F}'_{\text{OPT}} \cap \mathcal{D}_1 \cap \mathcal{D}_2) \leq \hat{\epsilon}W$;
- $\mathbf{w}(\mathcal{F}'_{\text{OPT}} \cap \mathcal{D}_1) \leq \frac{19}{20}W$ and $\mathbf{w}(\mathcal{F}'_{\text{OPT}} \cap \mathcal{D}_2) \leq \frac{19}{20}W$; and
- $\mathcal{F}'_{\text{OPT}} \cap \mathcal{D}_1$ r -covers \mathcal{C}_1 and $\mathcal{F}'_{\text{OPT}} \cap \mathcal{D}_2$ r -covers \mathcal{C}_2 .

By iterating through all members of \mathbb{X} , we may henceforth assume that the algorithm has fixed a quadruple $(\mathcal{D}_1, \mathcal{D}_2, \mathcal{C}_1, \mathcal{C}_2)$ with the properties as above.

The algorithm now recurses on instances $(G, \mathcal{D}_1, \mathcal{C}_1, r)$ and $(G, \mathcal{D}_2, \mathcal{C}_2, r)$, thus computing two solutions $\mathcal{F}_1 \subseteq \mathcal{D}_1$ and $\mathcal{F}_2 \subseteq \mathcal{D}_2$ such that \mathcal{C}_1 is r -covered by \mathcal{F}_1 and \mathcal{C}_2 is r -covered by \mathcal{F}_2 . Since $\mathcal{C}_1 \cup \mathcal{C}_2 = \mathcal{C}'$ and all members of $\mathcal{C} \setminus \mathcal{C}'$ are r -covered by \mathcal{F}_{hv} , the set $\mathcal{F} = \mathcal{F}_{\text{hv}} \cup \mathcal{F}_1 \cup \mathcal{F}_2$ r -covers \mathcal{C} . The algorithm outputs the lightest of the sets \mathcal{F} computed as above for all choices of \mathcal{F}_{hv} and $(\mathcal{D}_1, \mathcal{D}_2, \mathcal{C}_1, \mathcal{C}_2) \in \mathbb{X}$.

The base case of the recursion is given by trimming it at depth d_{\max} . More precisely, all recursive calls at depth larger than d_{\max} return $\mathcal{F} = \emptyset$ in the case when \mathcal{C} is empty (then it is an optimal solution), or a special marker \perp in the case when \mathcal{C} is not empty. The marker \perp should be interpreted as “error”, that is, it marks that a recursive call has failed to find a solution, and it propagates in the recursion using the following convention: $\mathbf{w}(\perp) = +\infty$ and the union of \perp with any other set is equal to \perp .

Algorithm 2: Algorithm AlgMWDSC

Input: An instance $(G, \mathcal{D}, \mathcal{C}, r)$, recursion depth d
Output: A subset $\mathcal{F} \subseteq \mathcal{D}$ that r -covers \mathcal{C}

```

    if  $\mathcal{C} = \emptyset$  then
    |   return  $\emptyset$ 
    if  $d > d_{\max}$  then
    |   return  $\perp$ 
    |    $\mathcal{F} \leftarrow \perp$ 
    |   forall  $\mathcal{F}_{\text{hv}}$ : subset of  $\mathcal{D}$  with  $|\mathcal{F}_{\text{hv}}| \leq s^2$  do
    |   |    $\mathcal{D}' \leftarrow \mathcal{D} \setminus \mathcal{F}_{\text{hv}}$ 
    |   |    $\mathcal{C}' \leftarrow \mathcal{C} \setminus$  (vertices of  $\mathcal{C}$  that are  $r$ -covered by  $\mathcal{F}_{\text{hv}}$ )
    |   |    $\mathbb{X} \leftarrow$  family computed for  $\mathcal{D}'$  and  $\mathcal{C}'$  using Lemma 10
    |   |   forall  $(\mathcal{D}_1, \mathcal{D}_2, \mathcal{C}_1, \mathcal{C}_2) \in \mathbb{X}$  do
    |   |   |    $\mathcal{F}_1 \leftarrow \text{AlgMWDSC}(G, \mathcal{D}_1, \mathcal{C}_1, r, d + 1)$ 
    |   |   |    $\mathcal{F}_2 \leftarrow \text{AlgMWDSC}(G, \mathcal{D}_2, \mathcal{C}_2, r, d + 1)$ 
    |   |   |    $\mathcal{F}_{\text{cand}} \leftarrow \mathcal{F}_{\text{hv}} \cup \mathcal{F}_1 \cup \mathcal{F}_2$ 
    |   |   |   if  $\mathbf{w}(\mathcal{F}_{\text{cand}}) < \mathbf{w}(\mathcal{F})$  then
    |   |   |   |    $\mathcal{F} \leftarrow \mathcal{F}_{\text{cand}}$ 
    |   |   return  $\mathcal{F}$ 
    return  $\mathcal{F}$ 

```

Again, we need to argue that the running time of the algorithm is as promised and that the output solution has weight at most $(1 + \epsilon)W$. The running time analysis is exactly the same as in the proof of Theorem 1, so we skip it and proceed directly to arguing the approximation factor.

Approximation factor analysis. Similarly as in the case of Theorem 1, the fact that the algorithm outputs a solution of weight at most $(1 + \epsilon)W$ follows directly from the following claim by applying it to the original instance $(G, \mathcal{D}, \mathcal{C}, r)$ and $d = 0$. Recall here that \mathcal{F}_{OPT} is a fixed optimum solution to the original instance and its weight is W .

Claim 12 *Let $(G, \tilde{\mathcal{D}}, \tilde{\mathcal{C}}, r)$ be an instance on which the algorithm is called in the recursion tree, say at depth d . Suppose further that $\mathbf{w}(\mathcal{F}_{\text{OPT}} \cap \tilde{\mathcal{D}}) \leq (19/20)^d \cdot W$ and that $\mathcal{F}_{\text{OPT}} \cap \tilde{\mathcal{D}}$ r -covers $\tilde{\mathcal{C}}$. Then the call of the algorithm to $(G, \tilde{\mathcal{D}}, \tilde{\mathcal{C}}, r)$ returns a subset $\tilde{\mathcal{F}} \subseteq \tilde{\mathcal{D}}$ that r -covers $\tilde{\mathcal{C}}$ and satisfies $\mathbf{w}(\tilde{\mathcal{F}}) \leq (1 + 2d'\hat{\epsilon})\mathbf{w}(\mathcal{F}_{\text{OPT}} \cap \tilde{\mathcal{D}})$, where $d' = d_{\max} - d$.*

Proof We prove the claim by induction on d , starting with the case $d \geq d_{\max}$ and then proceeding with decreasing d . When $d \geq d_{\max}$, we have

$$(19/20)^d \leq (1 - 1/20)^{20 \ln(MN)} < e^{-\ln(MN)} = \frac{1}{MN}.$$

Since \mathcal{F}_{OPT} consists of at most N vertices of weight less than M each, we have $W < MN$. Therefore

$$\mathbf{w}(\mathcal{F}_{\text{OPT}} \cap \tilde{\mathcal{D}}) \leq (19/20)^d \cdot W < \frac{1}{MN} \cdot MN = 1.$$

As each vertex of \mathcal{D} has weight at least 1, we infer that $\mathcal{F}_{\text{OPT}} \cap \tilde{\mathcal{D}} = \emptyset$. This means that $\tilde{\mathcal{C}}$ is empty as well and the algorithm correctly outputs \emptyset as a solution.

Suppose now that $d < d_{\text{max}}$. Denote $\tilde{\mathcal{F}}_{\text{OPT}} = \mathcal{F}_{\text{OPT}} \cap \tilde{\mathcal{D}}$ and let $\tilde{W} = \mathbf{w}(\tilde{\mathcal{F}}_{\text{OPT}})$; by assumption, $\tilde{W} \leq (19/20)^d \cdot W$. Call a vertex $p \in \tilde{\mathcal{F}}_{\text{OPT}}$ heavy if $\mathbf{w}(p) \geq s^{-2}\tilde{W}$ and let $\tilde{\mathcal{F}}_{\text{hv}} \subseteq \tilde{\mathcal{F}}_{\text{OPT}}$ be the set of heavy vertices in $\tilde{\mathcal{F}}_{\text{OPT}}$. Observe that $|\tilde{\mathcal{F}}_{\text{hv}}| \leq s^2$, so in the first loop of the algorithm there is an iteration in which we correctly fix the set $\tilde{\mathcal{F}}_{\text{hv}}$. We continue the reasoning under this assumption. As in the description, let $\tilde{\mathcal{D}}'$ be $\tilde{\mathcal{D}}$ with all vertices from $\tilde{\mathcal{F}}_{\text{hv}}$ removed, and let $\tilde{\mathcal{C}}$ be $\tilde{\mathcal{C}}$ with all vertices r -covered by $\tilde{\mathcal{F}}_{\text{hv}}$ removed.

Now let $\tilde{\mathcal{F}}'_{\text{OPT}} = \tilde{\mathcal{F}}_{\text{OPT}} \setminus \tilde{\mathcal{F}}_{\text{hv}}$ and observe that $\tilde{\mathcal{F}}'_{\text{OPT}}$ r -covers \mathcal{C}' . Since $\mathbf{w}(p) < s^{-2}\tilde{W}$ for all $p \in \tilde{\mathcal{F}}'_{\text{OPT}}$, we may apply Lemma 10 to $\tilde{\mathcal{F}}'_{\text{OPT}}$, \mathcal{C}' , and \tilde{W} as the upper bound on the weight of \mathcal{F}_{OPT} . Thus we conclude that the enumerated family \times contains at least one quadruple $(\tilde{\mathcal{D}}_1, \tilde{\mathcal{D}}_2, \tilde{\mathcal{C}}_1, \tilde{\mathcal{C}}_2)$ satisfying the following:

- (i) $\mathbf{w}(\tilde{\mathcal{F}}_{\text{OPT}} \cap \tilde{\mathcal{D}}_1 \cap \tilde{\mathcal{D}}_2) \leq \hat{\epsilon}\tilde{W}$;
- (ii) $\mathbf{w}(\tilde{\mathcal{F}}_{\text{OPT}} \cap \tilde{\mathcal{D}}_1) \leq \frac{19}{20}\tilde{W}$ and $\mathbf{w}(\tilde{\mathcal{F}}_{\text{OPT}} \cap \tilde{\mathcal{D}}_2) \leq \frac{19}{20}\tilde{W}$; and
- (iii) $\tilde{\mathcal{F}}_{\text{OPT}} \cap \tilde{\mathcal{D}}_1$ r -covers \mathcal{C}_1 and $\tilde{\mathcal{F}}_{\text{OPT}} \cap \tilde{\mathcal{D}}_2$ r -covers \mathcal{C}_2 .

In particular, in one iteration of the second loop the algorithm correctly fixes such a quadruple and proceeds with it. From now on we continue under this assumption.

Having fixed $\tilde{\mathcal{F}}_{\text{hv}}$ and $(\tilde{\mathcal{D}}_1, \tilde{\mathcal{D}}_2, \tilde{\mathcal{C}}_1, \tilde{\mathcal{C}}_2)$, the algorithm recursively calls itself on instances $(G, \tilde{\mathcal{D}}_1, \tilde{\mathcal{C}}_1, r)$ and $(G, \tilde{\mathcal{D}}_2, \tilde{\mathcal{C}}_2, r)$, yielding families $\tilde{\mathcal{F}}_1 \subseteq \tilde{\mathcal{D}}_1$ and $\tilde{\mathcal{F}}_2 \subseteq \tilde{\mathcal{D}}_2$. By property (ii), we have that

$$\mathbf{w}(\mathcal{F}_{\text{OPT}} \cap \tilde{\mathcal{D}}_i) \leq (19/20) \cdot \mathbf{w}(\mathcal{F}_{\text{OPT}} \cap \tilde{\mathcal{D}}) \leq (19/20)^{d+1} \cdot W,$$

for $i \in \{1, 2\}$. As the recursive subcalls to instances $(G, \tilde{\mathcal{D}}_i, \tilde{\mathcal{C}}_i, r)$ are at level $d + 1$ in the recursion tree, and due to property (iii), by the induction assumption we infer that $\tilde{\mathcal{F}}_i \neq \perp$ and

$$\mathbf{w}(\tilde{\mathcal{F}}_i) \leq (1 + 2(d' - 1)\hat{\epsilon}) \cdot \mathbf{w}(\mathcal{F}_{\text{OPT}} \cap \tilde{\mathcal{D}}_i), \tag{7}$$

for $i \in \{1, 2\}$. By property (i), we have that

$$\begin{aligned} & \mathbf{w}(\tilde{\mathcal{F}}_{\text{hv}}) + \mathbf{w}(\tilde{\mathcal{F}}_{\text{OPT}} \cap \tilde{\mathcal{D}}_1) + \mathbf{w}(\tilde{\mathcal{F}}_{\text{OPT}} \cap \tilde{\mathcal{D}}_2) \\ &= \mathbf{w}(\tilde{\mathcal{F}}_{\text{hv}}) + \mathbf{w}(\tilde{\mathcal{F}}_{\text{OPT}} \cap \tilde{\mathcal{D}}') + \mathbf{w}(\tilde{\mathcal{F}}_{\text{OPT}} \cap \tilde{\mathcal{D}}_1 \cap \tilde{\mathcal{D}}_2) \\ &= \mathbf{w}(\tilde{\mathcal{F}}_{\text{OPT}}) + \mathbf{w}(\tilde{\mathcal{F}}_{\text{OPT}} \cap \tilde{\mathcal{D}}_1 \cap \tilde{\mathcal{D}}_2) \\ &\leq \tilde{W} + \hat{\epsilon}\tilde{W} = (1 + \hat{\epsilon})\tilde{W}. \end{aligned} \tag{8}$$

Recall that the algorithm returns a solution of weight not larger than the weight of $\tilde{\mathcal{F}} = \tilde{\mathcal{F}}_{\text{hv}} \cup \tilde{\mathcal{F}}_1 \cup \tilde{\mathcal{F}}_2$. By combining (7) with (8) we have

$$\begin{aligned}
 \mathbf{w}(\tilde{\mathcal{F}}) &\leq \mathbf{w}(\tilde{\mathcal{F}}_{\text{hv}}) + \mathbf{w}(\tilde{\mathcal{F}}_1) + \mathbf{w}(\tilde{\mathcal{F}}_2) \\
 &\leq \mathbf{w}(\tilde{\mathcal{F}}_{\text{hv}}) + (1 + 2(d' - 1)\hat{\epsilon}) \cdot \left(\mathbf{w}(\mathcal{F}_{\text{OPT}} \cap \tilde{\mathcal{D}}_1) + \mathbf{w}(\mathcal{F}_{\text{OPT}} \cap \tilde{\mathcal{D}}_2) \right) \\
 &\leq (1 + 2(d' - 1)\hat{\epsilon}) \cdot \left(\mathbf{w}(\tilde{\mathcal{F}}_{\text{hv}}) + \mathbf{w}(\tilde{\mathcal{F}}_{\text{OPT}} \cap \tilde{\mathcal{D}}_1) + \mathbf{w}(\tilde{\mathcal{F}}_{\text{OPT}} \cap \tilde{\mathcal{D}}_2) \right) \\
 &\leq (1 + 2(d' - 1)\hat{\epsilon})(1 + \hat{\epsilon})\tilde{W} \leq (1 + 2d'\hat{\epsilon})\tilde{W},
 \end{aligned}$$

where the last inequality follows from $2(d' - 1)\hat{\epsilon}^2 \leq \hat{\epsilon}$, which is true for $d' \leq d_{\text{max}} \leq \frac{1}{40\hat{\epsilon}}$. This concludes the proof. □

This finishes the proof of Theorem 2.

Acknowledgements We thank Dániel Marx for insightful discussions on the approach to optimization problems in geometric and planar settings via Voronoi diagrams.

Appendix: Geometric Problems

In the MAXIMUM WEIGHT INDEPENDENT SET OF POLYGONS (MWISP) problems we are given a family \mathcal{P} of polygons in the plane, each with a prescribed weight, and the task is to find a maximum-weight subset of pairwise disjoint polygons. In the WEIGHTED GEOMETRIC SET COVER problem we are given a family \mathcal{P} of subsets of the plane, each with a prescribed weight, and a set \mathcal{U} of points in the plane. The goal is to find a minimum-weight subfamily of \mathcal{P} whose union covers \mathcal{U} . A QPTAS for MWISP with running time $2^{\text{poly}(1/\epsilon, \log N)} \cdot n^{O(1)}$, where $N = |\mathcal{P}|$ and n is the total number of vertices of the input polygons, was given by Har-Peled [4]. A QPTAS for WGSR under the assumption that \mathcal{P} is a family of *pseudo-disks* (i.e. compact, simply connected subsets of the plane such that the boundaries of each two meet in at most two points) with running time $2^{\text{poly}(1/\epsilon, n+N)}$, where $n = |\mathcal{U}|$ and $N = |\mathcal{P}|$, was given by Mustafa et al. [7].

We now prove that the abovementioned results follow from Theorems 1 and 2 via simple reductions from the geometric to the planar setting; the exception is that for WGSR we can tackle only unit disks and unit squares instead of general pseudo-disks. These reductions were already observed by Marx and Pilipczuk in [5], who used them for the parameterized variants of the problems. We give them for completeness.

Corollary 1 *The MAXIMUM WEIGHT INDEPENDENT SET OF POLYGONS problem admits a QPTAS with running time $2^{\text{poly}(1/\epsilon, \log N)} \cdot n^{O(1)}$, where N is the number of polygons and n is the total number of vertices of those polygons.*

Corollary 2 *The WEIGHTED GEOMETRIC SET COVER problems for unit disks and axis-parallel unit squares admits a QPTAS with running time $2^{\text{poly}(1/\epsilon, \log N)} \cdot n^{O(1)}$, where N is the number of disks/squares on input and n is the number of points to be covered.*

Before we proceed to proving the statements above, let us first give a general-use definition of the *crossing graph* of a point set in the plane; see Fig. 7. Consider a finite set of points X in the plane and a metric $d(\cdot, \cdot)$ on the plane induced by some norm (i.e. $d(x, y) = \|x - y\|$ for some norm $\|\cdot\|$). By $G(X, d)$ we denote a graph defined as follows. For every pair of distinct points $a, b \in X$, draw the segment with endpoints a and b in the plane. Let $Y \supseteq X$ be the set consisting of X and all intersections of all the segments drawn; then set Y to be the vertex set of $G(X, d)$. Two points $x, y \in Y$ are connected by an edge in $G(X, d)$ if they are two consecutive points from Y on any of the drawn segments; that is, the segment connecting x and y is contained in one of the segments connecting pairs of vertices from X , and there is no other point from Y inside this segment. The length of this edge is set to $d(x, y)$. Note that $G(X, d)$ has at most $|X|^4$ vertices, is planar with a straight-line embedding described above, and for every two vertices $a, b \in X$ we have $\text{dist}_{G(X, d)}(a, b) = d(a, b)$.

For further reference, let $d_2(\cdot, \cdot)$ and $d_\infty(\cdot, \cdot)$ denote the metrics on the plane induced by the ℓ_2 - and ℓ_∞ -norm, respectively. We now prove Corollaries 1 and 2.

Proof (of Corollary 1) Let \mathcal{P} be the given weighted set of polygons and let X be the set of all their vertices; then $|X| = n$. Consider the graph $G = G(X, d_\infty)$; note that G has at most n^4 vertices. For every polygon $P \in \mathcal{P}$ construct a corresponding object p in G defined as follows: p is the subgraph of G induced by all the vertices of G whose embeddings are contained in P . It is easy to see that p defined in this manner is connected. The weight of p is equal to the weight of P . Let \mathcal{D} be the family comprising all the constructed objects. It readily follows that sets of pairwise disjoint polygons from \mathcal{P} are in one-to-one correspondence with independent subfamilies of \mathcal{D} , hence it suffices to apply the algorithm of Theorem 1 to the instance (G, \mathcal{D}) . \square

Proof (of Corollary 2) Let us first concentrate on the case of unit disks. Let the input be $(\mathcal{P}, \mathcal{U})$, where \mathcal{P} is a weighted set of unit disks and \mathcal{U} is a set of points to be covered. Let X be the set consisting of all the centers of disks from \mathcal{P} and all the points from \mathcal{U} ; then $|X| \leq N + n$. Consider the graph $G = G(X, d_2)$; note that G has at most

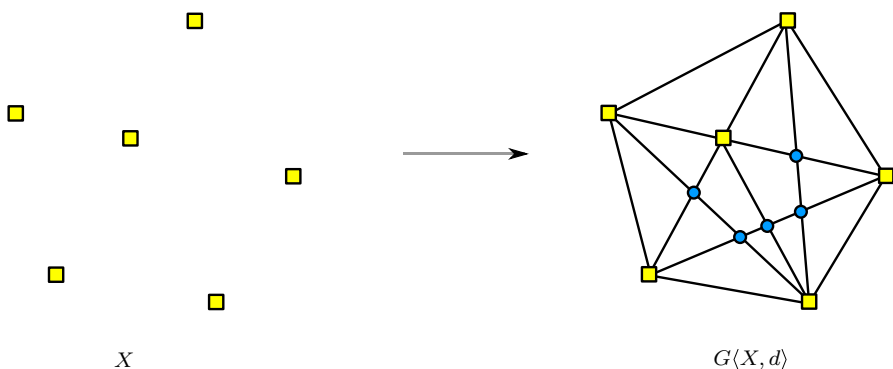


Fig. 7 A point set X and the corresponding crossing graph $G(X, d)$, for an unspecified metric d . Note that two triples of points in X are collinear

$(N + n)^4$ vertices. Construct a weighted set of points \mathcal{D} as follows: for every disk $P \in \mathcal{P}$, add the center of P to \mathcal{D} and assign it weight equal to the weight of P . It can be easily seen that a subset of disks $\mathcal{Q} \subseteq \mathcal{P}$ covers all the points from \mathcal{U} if and only if the centers of disks from \mathcal{Q} cover \mathcal{U} in G , where we consider domination radius $\frac{1}{2}$. Hence it suffices to apply the algorithm of Theorem 2 to the instance $(G, \mathcal{D}, \mathcal{U}, \frac{1}{2})$.

For the case of unit squares we may follow the same reasoning except we use metric $d_\infty(\cdot, \cdot)$ instead of $d_2(\cdot, \cdot)$. \square

References

1. Adamaszek, A., Wiese, A.: Approximation schemes for maximum weight independent set of rectangles. In: Proceedings of the FOCS 2013, pp. 400–409. IEEE (2013)
2. Adamaszek, A., Wiese, A.: A QPTAS for maximum weight independent set of polygons with polylogarithmically many vertices. In: Proceedings of the SODA 2014, pp. 645–656. SIAM (2014)
3. Gu, Q.-P., Tamaki, H.: Improved bounds on the planar branchwidth with respect to the largest grid minor size. *Algorithmica* **64**(3), 416–453 (2012)
4. Har-Peled, S.: Quasi-polynomial time approximation scheme for sparse subsets of polygons. In: Proceedings of the SOCG 2014, pp. 120–129. SIAM (2014)
5. Marx, D., Pilipczuk, M.: Optimal parameterized algorithms for planar facility location problems using Voronoi diagrams. In: Proceedings of the ESA 2015, Volume 9294 of LNCS, pp. 865–877. Springer (2015)
6. Miller, G.L.: Finding small simple cycle separators for 2-connected planar graphs. *J. Comput. Syst. Sci.* **32**(3), 265–279 (1986)
7. Mustafa, N.H., Raman, R., Ray, S.: Quasi-polynomial time approximation scheme for weighted geometric set cover on pseudodisks and halfspaces. *SIAM J. Comput.* **44**(6), 1650–1669 (2015)
8. Seymour, P.D., Thomas, R.: Call routing and the ratcatcher. *Combinatorica* **14**(2), 217–241 (1994)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

Michał Pilipczuk¹ · Erik Jan van Leeuwen² · Andreas Wiese³ 

✉ Andreas Wiese
awiese@dii.uchile.cl

Michał Pilipczuk
michal.pilipczuk@mimuw.edu.pl

Erik Jan van Leeuwen
e.j.vanleeuwen@uu.nl

¹ Institute of Informatics, University of Warsaw, Warsaw, Poland

² Department of Information and Computing Sciences, Utrecht University, Utrecht, The Netherlands

³ Department of Industrial Engineering and Center for Mathematical Modeling, Universidad de Chile, Santiago de Chile, Chile