

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

Michał Pilipczuk¹

Institute of Informatics, University of Warsaw, Poland
michal.pilipczuk@mimuw.edu.pl

Erik Jan van Leeuwen

Department of Information and Computing Sciences, Utrecht University, The Netherlands
e.j.vanleeuwen@uu.nl

Andreas Wiese²

Department of Industrial Engineering and Center for Mathematical Modeling,
Universidad de Chile, Chile
awiese@dii.uchile.cl

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 an edge-weighted graph G , 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. Our main technical contribution is to transfer the techniques used for recursive approximation schemes for geometric problems due to Adamaszek, Har-Peled, and Wiese [1, 2, 4] to the setting of planar graphs. In particular, this yields a purely combinatorial viewpoint on these methods.

2012 ACM Subject Classification Theory of computation → Packing and covering problems

Keywords and phrases QPTAS, planar graphs, Voronoi diagram

Digital Object Identifier 10.4230/LIPIcs.ESA.2018.65

Related Version A full version of the paper is available at <http://arxiv.org/abs/1807.07626>.

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

¹ The research of Mi. Pilipczuk was partially supported by Polish National Science Centre grant UMO-2013/11/D/ST6/03073, and is 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).

² 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.

© Michał Pilipczuk, Erik Jan van Leeuwen, and Andreas Wiese;
licensed under Creative Commons License CC-BY
26th Annual European Symposium on Algorithms (ESA 2018).

Editors: Yossi Azar, Hannah Bast, and Grzegorz Herman; Article No. 65; pp. 65:1–65:13
Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



European Research Council
Established by the European Commission

1 Introduction

INDEPENDENT SET and DOMINATING SET are fundamental optimization problems on graphs. Given a graph G where each vertex v has a weight $\mathbf{w}(v)$, 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 $f(k) \cdot n^{\mathcal{O}(1)}$, 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 $f(1/\epsilon)n^{\mathcal{O}(1)}$ for some function f , and fpt algorithms for the parameterization by the solution size, i.e., for a parameter k , algorithms running in time $f(k)n^{\mathcal{O}(1)}$ 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 weights on 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 weights, 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; however, the trivial $n^{\mathcal{O}(k)}$ -time algorithms can be improved to $n^{\mathcal{O}(\sqrt{k})}$ -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 edge-weighted graph G , 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, while 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 $2^{\text{poly}(1/\epsilon, \log N)} \cdot n^{\mathcal{O}(1)}$, 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 $2^{\text{poly}(1/\epsilon, \log N)} \cdot n^{\mathcal{O}(1)}$, 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.*

To obtain our QPTASs for MWISO, we extend 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 3 (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 $\mathbf{w}(\mathcal{C} \cap \mathcal{F}) \leq \frac{9}{10} W$.*

Using Lemma 3 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 3.

The proof of Lemma 3 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 3 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 concept 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 3 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 3 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. Due to space constraints, the proof of Theorem 2 is entirely deferred to the full version, while in this extended abstract we focus on proving Theorem 1.

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 the full version we explain how to derive the mentioned results from our theorems.

2 Proof of the Separator Lemma for MWISO

In this section we prove the Separator Lemma for MWISO, which was informally stated as Lemma 3 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 4** (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 $N^{\mathcal{O}(s)} \cdot n^{\mathcal{O}(1)}$; 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 $\text{IntGraph}(\mathcal{D}) \setminus \mathcal{X}$ we 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 4 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 weights³ 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 weight 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 its internal point c to be its *center*, and for each vertex u of f we fix some curve 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.

³ Obviously, edge weights are immaterial for the MWISO problem. However, it is convenient to think of G as edge-weighted so that we can define Voronoi diagrams. Furthermore, many results of this section will be reused for the MWDSC problem, where edge weights play a role in the problem.

We are given a family \mathcal{D} of *objects* in G , where each object $p \in \mathcal{D}$ is a nonempty, 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 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. 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 $\widehat{T}(p)$ 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, 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 the Voronoi diagram of \mathcal{F} 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 the diagram 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 nontrivially. 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 5** (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 5 and call its members *\mathcal{D} -important* faces of G .

Voronoi separators. We now recall the concept of *Voronoi separators*. 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 MWDSC). Therefore, we will later ensure that their total weight is small.

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

► **Lemma 6.** *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 exist 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 $\mathbf{w}(\mathcal{C}) \leq \frac{9}{10}W$.

It is not hard to see that Lemma 4 follows from Lemma 6: we simply enumerate all candidates for a Voronoi separator S satisfying 1 and 2, a straightforward estimate using Theorem 5 shows that there are at most $6^{3s}N^{15s}$ of them, and for each candidate S we add $\text{Ban}(S)$ to the constructed family \mathbb{X} . Details can be found in the full version.

Thus, we are left with proving Lemma 6. 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. 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: 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 $\widehat{T}(p_2)$, 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 $\widehat{T}(p_1)$. 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 7 (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 $10 \ln \ell \cdot \frac{W}{\ell}$.*

Later, we will use Lemma 7 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 $H_{\mathcal{S}}$ of size $\mathcal{O}(\sqrt{\ell})$ along which we partition \mathcal{F} into two parts, yielding the Voronoi separator claimed by Lemma 6. 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 7 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 5, 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 $\widehat{T}(p_2)$ between v_2 and v_3 and the edge v_2v_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 $\widehat{T}(p_0)$ 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 Figure 8 in [5] for a visualization.

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

► **Lemma 8** (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 9** (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 5. Lemma 9 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 8 we can bound their total number by $|\mathcal{D}|^4$.

Sampling lemma: proof. We now have all the tools needed to prove Lemma 7. 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 7. Denote $\eta = 10 \ln \ell \cdot \frac{W}{\ell}$. First observe that if $\mathbf{w}(\mathcal{F}) \leq \eta$, 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 $\mathbf{w}(\mathcal{F}) > \eta$.

Construct \mathcal{S} by including every object $p \in \mathcal{F}$ independently with probability $q_p = \mathbf{w}(p) \cdot \frac{\ell}{W}$; 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 $\mathbb{E}[X] = \sum_{p \in \mathcal{F}} \mathbb{E}[X_p] = \ell \cdot \frac{\mathbf{w}(\mathcal{F})}{W} \leq \ell$. Since X is a sum of independent indicator variables, standard concentration inequalities yield the following.

► **Claim 10** (\spadesuit^4). *The probability that $|\mathcal{S}| > 2\ell$ or $|\mathcal{S}| \leq 100$ is at most $\frac{7}{12}$.*

Call a spoke in the Voronoi diagram $H_{\mathcal{S}}$ *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 $H_{\mathcal{S}}$.

► **Claim 11** (\spadesuit). *The probability that there is a heavy spoke in $H_{\mathcal{S}}$ is at most $\frac{1}{6}$.*

Proof sketch. Fix any triple of objects $p_1, p_2, p_3 \in \mathcal{F}$ and face f and consider the following event $A^{p_1, p_2, p_3, f}$: p_1, p_2, p_3 are all included in \mathcal{S} , f is a branching point of $H_{\mathcal{S}}$ with vertices u_1, u_2, u_3 belonging to the cells of p_1, p_2, p_3 , respectively, and moreover the spoke P of u_1 (which is the shortest path from u_1 to p_1) is heavy. Let $\mathcal{Z} \subseteq \mathcal{F}$ be the family of those objects from \mathcal{F} that are in conflict with P ; then $\mathbf{w}(\mathcal{Z}) > \eta$. In order for $A^{p_1, p_2, p_3, f}$ to happen, all of

⁴ Proofs of claims or lemmas marked with a \spadesuit are deferred to the full version.

$\{p_1, p_2, p_3\}$ have to be included in \mathcal{S} and none of \mathcal{Z} may be included in \mathcal{S} . Since objects are included in \mathcal{S} independently, we have

$$\mathbb{P}(A^{p_1, p_2, p_3, f}) \leq \mathbf{w}(p_1)\mathbf{w}(p_2)\mathbf{w}(p_3) \cdot \frac{\ell^3}{W^3} \cdot \prod_{q \in \mathcal{Z}} (1 - \mathbf{w}(q)) \leq \frac{\mathbf{w}(p_1)\mathbf{w}(p_2)\mathbf{w}(p_3)}{W^3} \cdot \ell^{-7};$$

this follows from simple calculations using the lower bound on $\mathbf{w}(\mathcal{Z})$. By Lemma 8, for each triple $p_1, p_2, p_3 \in \mathcal{F}$ there are at most two singular faces of type 1, hence at most two faces f for which $A^{p_1, p_2, p_3, f}$ has a non-zero probability, each with three vertices. Summing through all the triples $p_1, p_2, p_3 \in \mathcal{F}$, we see that the probability that there is a heavy spoke in $H_{\mathcal{S}}$ incident to a type-1 branching point is at most $6 \cdot \ell^{-7} \leq 10^{-6}$. Applying a similar reasoning to the other two types of branching points yields the claim. \square

We are left with diamonds. Call a diamond $\Delta_{\mathcal{S}}(e)$ in $H_{\mathcal{S}}$ *heavy* if its weight with respect to \mathcal{F} is larger than η . The next check follows by essentially the same estimation as Claim 11.

► **Claim 12 (♠).** *The probability that there is a heavy diamond in $H_{\mathcal{S}}$ is at most $\frac{1}{6}$.*

Concluding, assertion $4 \leq |\mathcal{S}| \leq 2\ell$ does not hold with probability at most $\frac{7}{12}$, there is a heavy spoke in $H_{\mathcal{S}}$ with probability at most $\frac{1}{6}$, and there is a heavy diamond in $H_{\mathcal{S}}$ 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. \blacktriangleleft

2.3 Balanced nooses

We proceed with the proof of Lemma 6 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 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 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 $\mathbf{enc}(\gamma)$ while the second shall be called $\mathbf{exc}(\gamma)$ (for *enclosed* and *excluded*). 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 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 $\mathbf{enc}(\delta(x, y)) = \mathbf{exc}(\delta(y, x))$);
- $\mathbf{enc}(\delta(x, y))$ contains all the edges of H mapped to the component of $T - xy$ containing y , while $\mathbf{exc}(\delta(y, x))$ contains all the edges of H mapped to the other component of $T - xy$.

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

► **Theorem 13.** *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 13, 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 $(\delta(x, y_2) \cap \delta(x, y_3)) \setminus \delta(y_1, x)$. 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 14 (♠).** *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 14 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. Details are in the full version.

Proof sketch of Lemma 6. The proof of Lemma 6 now essentially follows by combining Theorem 13, Lemma 7, and Lemma 14 as follows. Applying Lemma 7 to \mathcal{F} with $\ell = s^2$ yields a suitable subfamily $\mathcal{S} \subseteq \mathcal{F}$. We investigate the Voronoi diagram $H_{\mathcal{S}}$ induced by \mathcal{S} . Applying Theorem 13 to $H_{\mathcal{S}}$ yields a sphere-cut decomposition of width at most $3s$, which we can feed to Lemma 14 to obtain a balanced noose γ . This noose naturally corresponds to

a Voronoi separator $S(\gamma)$, obtained by essentially tracing γ and marking the parts of the Voronoi diagram of \mathcal{S} that it visits. The connection between nooses in the Voronoi diagram and Voronoi separators was largely explored in [5]. Then properties asserted by Lemma 7, in particular the fact that every spoke in $H_{\mathcal{S}}$ has small weight, imply that $S(\gamma)$ satisfies all the requested conditions. There is one technical caveat in that the Voronoi diagram $H_{\mathcal{S}}$ may have bridges, so a priori we cannot apply Theorem 13 to it; this requires special technical treatment. The detailed proof of Lemma 6 can be found in the full version.

3 A QPTAS for Maximum Weight Independent Set of Objects

In this section we use Lemma 4 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 $N \geq 2$.

We will also assume that all objects in \mathcal{D} have weights between 1 and $M = 2\epsilon^{-1}N$. This assumption is easy to achieve as follows: guess the heaviest object p from the optimum solution, remove all objects of weight less than $\mathbf{w}(p)/M$ from \mathcal{D} , rescale the weights to the interval $[1, M]$, and then look for an $(1 - \epsilon/2)$ -approximate solution. To see that this is correct, observe that in the optimum solution, objects of weight less than $\mathbf{w}(p)/M$ in total constitute at most an $\epsilon/2$ -fraction of the weight of p alone, so by removing them we lose at most an $\epsilon/2$ fraction of the optimum. A formal reasoning is presented in the full version.

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

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

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 $(1 - \hat{\epsilon})^{d_{\max}} = (1 - \epsilon/\log(N))^{\mathcal{O}(\log(N))} = 1 - \mathcal{O}(\epsilon)$. 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 4 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 $\mathbf{w}(p) > s^{-2}W$. 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

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
┌    $\mathcal{D}' \leftarrow \mathcal{D} - (\text{all objects that intersect any object of } \mathcal{F}_{\text{hv}})$ 
    $\mathbb{X} \leftarrow \text{family computed for } \mathcal{D}' \text{ using Lemma 4}$ 
   forall  $\mathcal{X} \in \mathbb{X}$  do
┌    $\mathcal{D}_1, \dots, \mathcal{D}_k \leftarrow \text{vertex sets of the connected components of}$ 
    $\text{IntGraph}(\mathcal{D}') - \mathcal{X}$ 
   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  $\mathbf{w}(\mathcal{F}_{\text{cand}}) > \mathbf{w}(\mathcal{F})$  then
┌    $\mathcal{F} \leftarrow \mathcal{F}_{\text{cand}}$ 
└   return  $\mathcal{F}$ 

```

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 $\mathbf{w}(\mathcal{F}'_{\text{OPT}}) \leq W$ and $\mathbf{w}(p) \leq s^{-2}W$ for all $p \in \mathcal{F}'_{\text{OPT}}$.

We may now apply Lemma 4 to $\mathcal{F}'_{\text{OPT}} \subseteq \mathcal{D}'$ with W being the upper bound on its weight. Thus, in time $N^{\mathcal{O}(s)} \cdot n^{\mathcal{O}(1)}$ 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 $\mathbf{w}(\mathcal{F}'_{\text{OPT}} \setminus \mathcal{X}) \leq \epsilon W$ 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.

The above algorithm runs in time $2^{\text{poly}(1/\epsilon, N)} \cdot n^{\mathcal{O}(1)}$, 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. As mentioned, 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. Formal proofs are in the full version.

References

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