

Efficient Approximation Schemes for Uniform-Cost Clustering Problems in Planar Graphs

Vincent Cohen-Addad

Sorbonne Université, CNRS, Laboratoire d'informatique de Paris 6, LIP6, F-75252 Paris, France
vincent.cohen.addad@ens-lyon.org

Marcin Pilipczuk

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

Michał Pilipczuk

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

Abstract

We consider the k -MEDIAN problem on planar graphs: given an edge-weighted planar graph G , a set of clients $C \subseteq V(G)$, a set of facilities $F \subseteq V(G)$, and an integer parameter k , the task is to find a set of at most k facilities whose opening minimizes the total connection cost of clients, where each client contributes to the cost with the distance to the closest open facility. We give two new approximation schemes for this problem:

- *FPT Approximation Scheme*: for any $\varepsilon > 0$, in time $2^{\mathcal{O}(k\varepsilon^{-3} \log(k\varepsilon^{-1}))} \cdot n^{\mathcal{O}(1)}$ we can compute a solution that has connection cost at most $(1 + \varepsilon)$ times the optimum, with high probability.
- *Efficient Bicriteria Approximation Scheme*: for any $\varepsilon > 0$, in time $2^{\mathcal{O}(\varepsilon^{-5} \log(\varepsilon^{-1}))} \cdot n^{\mathcal{O}(1)}$ we can compute a set of at most $(1 + \varepsilon)k$ facilities whose opening yields connection cost at most $(1 + \varepsilon)$ times the optimum connection cost for opening at most k facilities, with high probability.

As a direct corollary of the second result we obtain an EPTAS for UNIFORM FACILITY LOCATION on planar graphs, with same running time.

Our main technical tool is a new construction of a “coreset for facilities” for k -MEDIAN in planar graphs: we show that in polynomial time one can compute a subset of facilities $F_0 \subseteq F$ of size $k \cdot (\log n / \varepsilon)^{\mathcal{O}(\varepsilon^{-3})}$ with a guarantee that there is a $(1 + \varepsilon)$ -approximate solution contained in F_0 .

2012 ACM Subject Classification Theory of computation → Facility location and clustering; Theory of computation → Fixed parameter tractability

Keywords and phrases k -Median, Facility Location, Planar Graphs, Approximation Scheme

Digital Object Identifier 10.4230/LIPIcs.ESA.2019.33

Related Version Full version of the paper accessible at <http://arxiv.org/abs/1905.00656>.

Funding This work is a part of projects CUTACOMBS (Ma. Pilipczuk) and TOTAL (Mi. Pilipczuk) that have received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreements No. 714704 and No. 677651, respectively). Ce projet a bénéficié d’une aide de l’État gérée par l’Agence Nationale de la Recherche au titre du Programme Appel à projets générique JCJC 2018 portant la référence suivante: ANR-18-CE40-0004-01.



1 Introduction

We study approximation schemes for classic clustering objectives, formalized as follows. Given an edge-weighted graph G together with a set C of vertices called *clients*, a set F of vertices called *candidate facilities*, and an *opening cost* $\text{open} \in \mathbb{R}_{\geq 0}$, the UNIFORM FACILITY LOCATION problem asks for a subset of facilities (also called centers) $D \subseteq F$ that minimizes



© Vincent Cohen-Addad, Marcin Pilipczuk, and Michał Pilipczuk;
licensed under Creative Commons License CC-BY
27th Annual European Symposium on Algorithms (ESA 2019).

Editors: Michael A. Bender, Ola Svensson, and Grzegorz Herman; Article No. 33; pp. 33:1–33:14



Leibniz International Proceedings in Informatics

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

the cost defined as $|D| \cdot \text{open} + \sum_{c \in C} \min_{f \in D} \text{dist}(c, f)$. In the NON-UNIFORM FACILITY LOCATION variant, the opening costs may vary between facilities.

We also consider the related k -MEDIAN problem, where the tuple (G, F, C) comes with a hard budget k for the number of open facilities (as opposed to the opening cost `open`). That is, the problem asks for a set $D \subseteq F$ of size at most k that minimizes the connection cost $\sum_{c \in C} \text{dist}(c, D)$. Note that UNIFORM FACILITY LOCATION can be reduced to k -MEDIAN by guessing the number of open facilities in an optimal solution.

FACILITY LOCATION and k -MEDIAN model in an abstract way various clustering objectives appearing in applications. Therefore, designing approximation algorithms for them and their variants is a vibrant topic in the field of approximation algorithms. For NON-UNIFORM FACILITY LOCATION, a long line of work [1, 15, 22, 16] culminated with the 1.488-approximation algorithm by Li [18]. On the other hand, Guha and Khuller [13] showed that the problem cannot be approximated in polynomial time within factor better than 1.463 unless $\text{NP} \subseteq \text{DTIME}[n^{\mathcal{O}(\log \log n)}]$, which gives almost tight bounds on the best approximation factor achievable in polynomial time. For k -MEDIAN, the best known approximation ratio achievable in polynomial time is 2.67 due to Byrka et al. [3], while the lower bound of 1.463 due to Guha and Khuller [13] holds here as well.

Given the approximation hardness status presented above, it is natural to consider restricted metrics. In this work we consider *planar metrics*: we assume that the underlying edge-weighted graph G is planar.

It was a long-standing open problem whether FACILITY LOCATION admits a polynomial-time approximation scheme (PTAS) in planar metrics. For the uniform case, this question has been resolved in affirmative by Cohen-Addad et al. [8] in an elegant way: they showed that local search of radius $\mathcal{O}(1/\varepsilon^2)$ actually yields a $(1 + \varepsilon)$ -approximation, giving a PTAS with running time $n^{\mathcal{O}(1/\varepsilon^2)}$. This approach also gives a PTAS for k -MEDIAN with a similar running time, and works even in metrics induced by graphs from any fixed proper minor-closed class.

Very recently, Cohen-Addad et al. [9] also gave a PTAS for NON-UNIFORM FACILITY LOCATION in planar metrics using a different approach. Roughly, the idea is to first apply Baker layering scheme to reduce the problem to the case when in all clusters (sets of clients connected to the same facility) in the solution, all clients are within distance between 1 and r from the center, for some constant r depending only on ε . This case is then resolved by another application of Baker layering scheme, followed by a dynamic programming on a hierarchical decomposition of the graph using shortest paths as balanced separators.

Both the schemes of [8] and of [9] are PTASes: they run in time $n^{g(\varepsilon)}$ for some function g . It is therefore natural to ask for an *efficient PTAS* (EPTAS): an approximation scheme with running time $f(\varepsilon) \cdot n^{\mathcal{O}(1)}$ for some function f . Such an EPTAS was given by Cohen-Addad [5] for k -MEANS in low-dimensional Euclidean spaces; this is a variant of k -MEDIAN where every client contributes to the connection cost with the *square* of its distance from the closest open facility. Here, the idea is to apply local search as in [8], but to use the properties of the metric to explore the local neighborhood faster. Unfortunately, this technique mainly relies on the Euclidean structure (or on the bounded doubling dimension of the input) and seems hard to lift to the general planar case. Also the techniques of [9] are far from yielding an EPTAS: essentially, one needs to use a logarithmic number of portals at every step of the final dynamic programming in order to tame the accumulation of error through $\log n$ levels of the decomposition. Recently, a near-linear time EPTAS was given by Cohen-Addad et al. [6] for k -median, k -means and non-uniform facility location where the input points lie in a metric of bounded doubling dimension. This result is obtained using a randomized dissection of the metric, a technique that is not available for planar inputs.

The goal of this work is to circumvent these difficulties and give an EPTAS for UNIFORM FACILITY LOCATION in planar metrics.

Our results. Our main technical contribution is the following theorem. In essence, it states that when solving k -MEDIAN on a planar graph one can restrict the facility set to a subset of size $k \cdot (\varepsilon^{-1} \log n)^{\mathcal{O}(\varepsilon^{-3})}$, at the cost of losing a multiplicative factor of $(1 + \varepsilon)$ on the optimum connection cost. This can be seen as the planar version of the classic result by Matoušek [20] who showed that for Euclidean metrics of dimension d , it is possible to reduce the number of candidate centers to $\text{poly}(k)\varepsilon^{-\mathcal{O}(d)}$ at the cost of losing a multiplicative factor of $(1 + \varepsilon)$ on the optimum connection cost (through the use of coresets as well). For general metrics, obtaining such a result seems challenging, since this would imply a $(1 + \varepsilon)$ -approximation algorithm with running time $f(k, \varepsilon)n^{\mathcal{O}(1)}$, which would contradict Gap-ETH [7].

From now on, by *with constant probability* we mean with probability at least $1/2$; this can be boosted by independent repetition.

► **Theorem 1.** *Given a k -Median instance (G, F, C, k) , where G is a planar graph, and an accuracy parameter $\varepsilon > 0$, one can in randomized polynomial time compute a set $F_0 \subseteq F$ of size $k \cdot (\varepsilon^{-1} \log n)^{\mathcal{O}(\varepsilon^{-3})}$ satisfying the following condition with constant probability: there exists a set $D_0 \subseteq F_0$ of size at most k such that for every set $D \subseteq F$ of size at most k it holds that $\sum_{c \in C} \min_{u \in D_0} \text{dist}(u, c) \leq (1 + \varepsilon) \sum_{c \in C} \min_{u \in D} \text{dist}(u, c)$.*

A direct corollary of Theorem 1 is a fixed-parameter approximation scheme for the k -MEDIAN problem in planar graphs. This continues the line of work on fixed-parameter approximation schemes for k -median and k -means in Euclidean spaces [12, 17], where the goal is to design an algorithm running in time $f(k, \varepsilon) \cdot n^{\mathcal{O}(1)}$ for a computable function f .

► **Corollary 2.** *Given a k -Median instance (G, F, C, k) , where G is a planar graph, and an accuracy parameter $\varepsilon > 0$, one can in randomized time $2^{\mathcal{O}(k\varepsilon^{-3} \log(k\varepsilon^{-1}))} \cdot n^{\mathcal{O}(1)}$ compute a solution $D \subseteq F$ that has connection cost at most $(1 + \varepsilon)$ times the minimum possible connection cost with constant probability.*

Proof. Apply the algorithm of Theorem 1 and let $F_0 \subseteq F$ be the obtained subset of facilities. Then run a brute-force search through all subsets of F_0 of size at most k and output one with the smallest connection cost. Thus, the running time is

$$\left(k \cdot (\varepsilon^{-1} \log n)^{\mathcal{O}(\varepsilon^{-3})}\right)^k \cdot n^{\mathcal{O}(1)} \leq 2^{\mathcal{O}(k\varepsilon^{-3} \log(k\varepsilon^{-1}))} \cdot (\log n)^{\mathcal{O}(k\varepsilon^{-3})} \cdot n^{\mathcal{O}(1)}$$

which is at most $2^{\mathcal{O}(k\varepsilon^{-3} \log(k\varepsilon^{-1}))} \cdot n^{\mathcal{O}(1)}$ since $(\log n)^d \leq 2^{\mathcal{O}(d \log d)} \cdot n^{\mathcal{O}(1)}$: if $n \leq 2^{d^2}$ then $(\log n)^d \leq d^{2d} \leq 2^{\mathcal{O}(d \log d)}$, and if $n > 2^{d^2}$ then $(\log n)^d \leq 2^{\sqrt{\log n} \cdot \log \log n} \leq n^{\mathcal{O}(1)}$. ◀

Using Theorem 1 we can also give an efficient bicriteria PTAS for k -MEDIAN in planar graphs. This time, the proof is more involved and uses the local search techniques of [5].

► **Theorem 3.** *Given a k -Median instance (G, F, C, k) , where G is a planar graph, and an accuracy parameter $\varepsilon > 0$, one can in randomized time $2^{\mathcal{O}(\varepsilon^{-5} \log(\varepsilon^{-1}))} \cdot n^{\mathcal{O}(1)}$ compute a set $D \subseteq F$ of size at most $(1 + \varepsilon)k$ such that its connection cost is at most $(1 + \varepsilon)$ times the minimum possible connection cost for solutions of size k with constant probability.*

A direct corollary of Theorem 3 is an efficient PTAS for UNIFORM FACILITY LOCATION in planar graphs.

► **Theorem 4.** *Given a Uniform Facility Location instance (G, F, C, open) , where G is a planar graph, and an accuracy parameter $\varepsilon > 0$, one can in randomized time $2^{\mathcal{O}(\varepsilon^{-5} \log(\varepsilon^{-1}))} \cdot n^{\mathcal{O}(1)}$ compute a solution $D \subseteq F$ that has total cost at most $(1 + \varepsilon)$ times the optimum cost with constant probability.*

Proof. Iterate over all possible choices of k being the number of facilities opened by the optimum solution, and for every k invoke the algorithm of Theorem 3 for the k -MEDIAN instance (G, F, C, k) . From the obtained solutions output one with the smallest cost. ◀

Note that the approach presented above fails for the non-uniform case, where each facility has its own, distinct opening cost.

In this extended abstract we focus on proving the main result, Theorem 1. The proof of Theorem 3, on which Theorem 4 also relies, is fully deferred to the full version of the paper [10].

Our techniques. The first step in the proof of Theorem 1 is to reduce the number of relevant clients using the coreset construction of Feldman and Langberg [11]. By applying this technique, we may assume that there are at most $k \cdot \mathcal{O}(\varepsilon^{-2} \log n)$ clients in the instance, however they are weighted: every client c is assigned a nonnegative weight $\omega(c)$, and it contributes to the connection cost of any solution with $\omega(c)$ times the distance to the closest open facility in the solution.

We now examine the Voronoi diagram induced in the input graph G by the *clients*: vertices of G are classified into *cells* according to the closest client. This Voronoi diagram has one cell per every client, thus it can be regarded as a planar graph with $|C|$ faces, where each face accommodates one cell. To formally define the Voronoi diagram, and in particular the boundaries between neighboring cells, we use the framework introduced by Marx and Pilipczuk [19] and its extension used in [21].

Consider now all the *spokes* in the diagram, where a spoke is the shortest path connecting the center of a cell (i.e. a client) with a branching node of the diagram incident to the cell (which is a face of G). Removing all the spokes and all the branching nodes from the plane divides it into *diamonds*, where each diamond is delimited by four spokes, called further the *perimeter* of the diamond. See Figure 1 for an example. Since the diagram is a planar graph with $|C|$ faces, there are $\mathcal{O}(|C|) = k \cdot \mathcal{O}(\varepsilon^{-2} \log n)$ diamonds altogether. Moreover, since no diamond contains a client in its interior, whenever P is a path connecting a client with a facility belonging to some diamond Δ , P has to cross the perimeter of Δ .

Now comes the key and most technical part of the proof. We very carefully put $\mathcal{O}(\varepsilon^{-2} \log n)$ portals on the perimeter of each diamond. The idea of placement is similar to that of the resolution metric used in the QPTAS for FACILITY LOCATION. Namely, on a spoke Q starting at client c we put portals at distance $1, (1 + \varepsilon), (1 + \varepsilon)^2, \dots$ from c , so that the further we are on the spoke from c , the sparser the portals are. As a diamond is delimited by four spokes, we may thus use only $\mathcal{O}(\varepsilon^{-2} \log n)$ portals per diamond, while the cost of snapping a path crossing Q to the portal closest to the crossing point can be bounded by ε times the distance from the crossing point to c .

For a facility f in a diamond Δ , we define the *profile* of f as follows. For every spoke Q in the perimeter of Δ , we look at the closest portal π from f on Q . We record approximate (up to $(1 + \varepsilon)$ multiplicative error) distances from f to π and $\mathcal{O}(\varepsilon^{-3})$ neighboring portals, as well as the distance to the client endpoint of the spoke Q . The crux lies in the following fact: for every two facilities f, f' in Δ with the same profile, replacing f with f' increases the connection cost of any client c connected to f only by a multiplicative factor of $(1 + \varepsilon)$. Hence, for every profile in every diamond it suffices to keep just one facility with that profile. Since there are $k \cdot \mathcal{O}(\varepsilon^{-2} \log n)$ diamonds and $\mathcal{O}(\varepsilon^{-1} \log n)^{\mathcal{O}(\varepsilon^{-3})}$ possible profiles in each of them, we keep at most $k \cdot (\varepsilon^{-1} \log n)^{\mathcal{O}(\varepsilon^{-3})}$ facilities in total. This proves Theorem 1.

For the proof of Theorem 3, we first apply Theorem 1 to reduce the number of facilities to $k \cdot (\varepsilon^{-1} \log n)^{\mathcal{O}(\varepsilon^{-3})}$. Then we again inspect the Voronoi diagram, but now induced by the *facilities*. Having contracted every cell to a single vertex, we compute an r -division of the

obtained planar graph to cover it with regions of size $r = (\varepsilon^{-1} \log n)^{\mathcal{O}(\varepsilon^{-3})}$ so that only $\mathcal{O}(\varepsilon)k$ facilities are on boundaries of the regions. We open all the facilities in all the boundaries – thus exceeding the quota for open facilities by $\mathcal{O}(\varepsilon)k$ – run the PTAS of Cohen-Addad et al. [8] in each region independently, and at the end assemble regional solutions using a knapsack dynamic programming. Since within each region there are only polylogarithmically many facilities, each application of the PTAS actually works in time $f(\varepsilon) \cdot n^{\mathcal{O}(1)}$.

2 Preliminaries on Voronoi diagrams and coresets

In this section we recall some tools about Voronoi diagrams in planar graphs and coresets that will be used in the proof of Theorem 1. We will consider undirected graphs with positive edge lengths embedded in a sphere, with the standard shortest-paths metric $\text{dist}(u, v)$ for $u, v \in V(G)$. Contrary to the previous section, the metric is defined on the vertex set of G only, i.e., we do not consider G as a metric space with points in the interiors of edges. For $X, Y \subseteq V(G)$, we denote $\text{dist}(X, Y) = \min_{x \in X, y \in Y} \text{dist}(x, y)$ and similarly we define $\text{dist}(u, X)$ for $u \in V(G)$ and $X \subseteq V(G)$.

Recall that for a set $D \subseteq V(G)$ of *open facilities* and a set $C \subseteq V(G)$, we define the *connection cost* as

$$\text{conn}(D, C) = \sum_{v \in C} \text{dist}(v, D).$$

If the input is additionally equipped with opening costs $\text{open}: F \rightarrow \mathbb{R}_{\geq 0}$, then the *opening cost* of D is defined as $\sum_{w \in D} \text{open}(w)$.


2.1 Voronoi diagrams in planar graphs

We now recall the construction of Voronoi diagrams and related notions in planar graphs used by Marx and Pilipczuk [19]. The setting is as follows. Suppose G is an n -vertex simple graph embedded in a sphere Σ whose edges are assigned nonnegative real lengths. We consider the shortest path metric in G : for two vertices u, v , their *distance* $\text{dist}(u, v)$ is equal to the smallest possible total length of a path from u to v . We will assume that G is triangulated (i.e. every face of G is a cycle of length 3), for this may always be achieved by triangulating the graph using edges of infinite weight.

Further, we assume that shortest paths are unique in G and that finite distances between distinct vertices in G are pairwise different: for all vertices u, v, u', v' with $u \neq v$, $u' \neq v'$ and $\{u, v\} \neq \{u', v'\}$, we have $\text{dist}(u, v) \neq \text{dist}(u', v')$ or $\text{dist}(u, v) = \text{dist}(u', v') = +\infty$. This can be achieved by adding small perturbations to the edge lengths. Since we never specify degrees of polynomials in the running time of our algorithms, we may ignore the additional complexity cost incurred by the need of handling the perturbations in arithmetic operations.

Voronoi diagrams and their properties. Suppose that S is a subset of vertices¹ of G . First, define the *Voronoi partition*: for a vertex $p \in S$, the *Voronoi cell* $\text{Cell}_S(p)$ is the set of all those vertices $u \in V(G)$ whose distance from p is smaller than the distance from any other vertex $q \in S$; note that ties do not occur due to the distinctness of distances in G . Note that $\{\text{Cell}_S(p)\}_{p \in S}$ is a partition of the vertex set of G . For each $p \in S$, let $T(p)$ be the union of

¹ In [19] a more general setting is considered where objects inducing the diagram are connected subgraphs of G instead of single vertices. We will not need this generality here.



■ **Figure 1** A part of the Voronoi diagram with various features distinguished. Branching nodes of the diagram are grayed triangular faces, edges of the diagram are dashed. Solid paths of respective colors are spokes. (The interior of) one diamond is grayed in order to highlight it.

shortest paths from vertices of $\text{Cell}_S(p)$ to p ; recall here that shortest paths in G are unique. Note that, due to the distinctness of distances in G , $T(p)$ is a spanning tree of the subgraph of G induced by the cell $\text{Cell}_S(p)$.

The diagram Vor_S induced by G is a multigraph constructed as follows. First, take the dual G^* of G and remove all edges dual to the edges of all the trees $T(p)$, for $p \in S$. Then, exhaustively remove vertices of degree 1. Finally, for every maximal 2-path (i.e. path with internal vertices of degree 2), say with endpoints u and v , we replace this path by the edge uv ; note that this creates a loop at u in case $u = v$. The resulting multigraph Vor_S is the Voronoi diagram induced by S . Note that the vertices of Vor_S are faces of G ; for clarity we shall call them *branching nodes*. Furthermore, Vor_S inherits an embedding in Σ from the dual G^* , where an edge uv that replaced a maximal 2-path P is embedded precisely as P , i.e., as the concatenation of (the embeddings of) the edges comprising P . From now on we will assume this embedding of Vor_S .

We recall several properties of Vor_S , observed in [19]

► **Lemma 5** (Lemmas 4.4 and 4.5 of [19]). *The diagram Vor_S is a connected and 3-regular multigraph embedded in Σ , which has exactly $|S|$ faces, $2|S| - 4$ branching nodes, and $3|S| - 6$ edges. The faces of Vor_S are in one-to-one correspondence with vertices of S : each $p \in S$ corresponds to a face of Vor_S that contains all vertices of $\text{Cell}_S(p)$ and no other vertex of G .*

Spokes and diamonds. We now introduce further structural elements that can be distinguished in the Voronoi diagram, see Figure 1 for reference. The definitions and basic observations presented below are taken from Pilipczuk et al. [21], and were inspired by the Euclidean analogues due to Har-Peled [14].

An *incidence* is a triple $\tau = (p, u, f)$ where $p \in S$, f is a branching node of the diagram Vor_S , and u is a vertex of G that lies on f (recall that f is a triangular face of G) and belongs to $\text{Cell}_S(p)$. The *spoke* of the incidence τ , denoted $\text{Spoke}(\tau)$, is the shortest path in G between p and u . Note that all the vertices of $\text{Spoke}(\tau)$ belong to $\text{Cell}_S(p)$.

Let $e = f_1 f_2$ be an edge of the diagram Vor_S , where f_1, f_2 are branching nodes (possibly $f_1 = f_2$ if e is a loop in Vor_S). Further, let p_1 and p_2 be the vertices from S that correspond to faces of Vor_S incident to e (possibly $p_1 = p_2$ if e is a bridge in Vor_S). Suppose for a moment that $f_1 \neq f_2$. Then, out of the three edges of f_1 (these are edges in G) there is

exactly one that crosses the edge e of Vor_S ; say it is the edge $u_{1,1}u_{1,2}$ where $u_{1,1} \in \text{Cell}_S(p_1)$ and $u_{1,2} \in \text{Cell}_S(p_2)$. Symmetrically, there is one edge of f_2 that crosses e , say it is $u_{2,1}u_{2,2}$ where $u_{2,1} \in \text{Cell}_S(p_1)$ and $u_{2,2} \in \text{Cell}_S(p_2)$. In case $f_1 = f_2$, the edge e crosses two different edges of $f_1 = f_2$ and we define $u_{1,1}, u_{1,2}, u_{2,1}, u_{2,2}$ analogously for these two crossings; note that then, provided p_1 corresponds to the face enclosed by the loop e , we have $u_{1,1} = u_{1,2}$. For all $i, j \in \{1, 2\}$, consider the incidence $\tau_{i,j} = (p_i, u_{i,j}, f_j)$.

Consider removing the following subsets from the sphere Σ : interiors of faces f_1, f_2 and spokes $\text{Spoke}(\tau_{i,j})$ for all $i, j \in \{1, 2\}$. After this removal the sphere breaks into two regions, out of which exactly one, say R , intersects (the embedding of) e . Let the *diamond* of e , denoted $\text{Diam}(e)$, be the subgraph of G consisting of all features (vertices and edges) embedded in $R \cup \bigcup_{i,j \in \{1,2\}} \text{Spoke}(\tau_{i,j})$. The region R as above is the *interior* of the diamond $\text{Diam}(e)$. Note that in particular, the spokes $\text{Spoke}(\tau_{i,j})$ for $i, j \in \{1, 2\}$ and the edges $u_{1,1}u_{1,2}$ and $u_{2,1}u_{2,2}$ belong to $\text{Diam}(e)$. The *perimeter* of the diamond of e is the closed walk obtained by concatenating spokes $\text{Spoke}(\tau_{i,j})$ for $i, j \in \{1, 2\}$ and edges $u_{1,1}u_{1,2}, u_{2,1}u_{2,2}$ in the natural order around $\text{Diam}(e)$. The following observation is immediate:

► **Proposition 6.** *Consider removing all the spokes (considered as curves on Σ) and all the branching nodes (considered as interiors of faces on Σ) of the diagram Vor_S from the sphere Σ . Then Σ breaks into $3|S| - 6$ regions that are in one-to-one correspondence with edges of Vor_S : a region corresponding to the edge e is the interior of the diamond $\text{Diam}(e)$. Consequently, the intersection of diamonds of two different edges of Vor_S is contained in the intersection of their perimeters.*

Finally, we note that the perimeter of a diamond separates it from the rest of the graph. Since vertices of S are never contained in the interior of a diamond, this yields the following.

► **Lemma 7.** *Let $p \in S$ and u be a vertex of G belonging to the diamond $\text{Diam}(e)$ for some edge e of Vor_S . Then every path in G connecting u and p intersects the perimeter of $\text{Diam}(e)$.*

2.2 Coresets

In most our algorithms, the starting point is the notion of a *coreset* and a corresponding result of Feldman and Langberg [11]. To this end, we need to slightly generalize the notion of a client set in a k -MEDIAN instance. A *client weight function* is a function $\omega: C \rightarrow \mathbb{R}_{\geq 0}$. Given a set $D \subseteq F$ of open facilities, the (weighted) connection cost is defined as

$$\text{conn}(D, \omega) = \sum_{v \in C} \text{dist}(v, D) \cdot \omega(v).$$

That is, every client v is assigned a weight $\omega(v)$ with which it contributes to the objective function. The *support* of a weight function ω is defined as $\text{supp}(\omega) = \{v \in C \mid \omega(v) > 0\}$. From now on, whenever we speak about a k -MEDIAN instance without specified client weight function, we assume the standard function assigning each client weight 1.

The essence of coresets is that one can find weight functions with small support that well approximate the original instance. Given a k -MEDIAN instance (G, F, C, k) (without weights) and an accuracy parameter $\varepsilon > 0$, a *coreset* is a weight function ω such that for every set $D \subseteq F$ of size at most k , it holds that

$$|\text{conn}(D, C) - \text{conn}(D, \omega)| \leq \varepsilon \cdot \text{conn}(D, C).$$

We rely on the following result of Feldman and Langberg [11].

► **Theorem 8** (Theorem 15.4 of [11]). *Given a k -Median instance (G, F, C, k) with $n = |V(G)|$ and accuracy parameter $\varepsilon > 0$, one can in randomized polynomial time find a weight function ω with support of size $\mathcal{O}(k\varepsilon^{-2} \log n)$ that is a coreset with constant probability.*

We note that Ke Chen [4] gave a construction of a strong coreset with support of size $\mathcal{O}(k^2\varepsilon^{-2} \log n)$ that is much simpler than the later construction of Feldman and Langberg [11]. By using this construction instead, we would obtain a weaker version of Theorem 1, with a bound on $|F_0|$ that is quadratic in k instead of linear. This would be perfectly sufficient to derive an FPT approximation scheme as in Corollary 2, but for Theorem 3 we will vitally use the stronger statement. A construction of coresets with similar size guarantees, but maintainable in the streaming model, has been proposed by Braverman et al. [2].

3 Facility coreset for k -Median in planar graphs

In this section we give a coreset for centers for the k -MEDIAN problem, that is, we prove Theorem 1. We shall focus on the following lemma, which in combination with Theorem 8 yields Theorem 1.

► **Lemma 9.** *Given a k -Median instance (G, F, C, k) with a weight function ω and an accuracy parameter $\varepsilon > 0$, one can in polynomial time compute a set $F_0 \subseteq F$ of size $|\text{supp}(\omega)| \cdot (\varepsilon^{-1} \log |V(G)|)^{\mathcal{O}(\varepsilon^{-3})}$ satisfying the following condition with constant probability: there exists a set $D_0 \subseteq F_0$ of size at most k such that for every set $D \subseteq F$ of size at most k it holds that $\text{conn}(D_0, \omega) \leq (1 + \varepsilon) \cdot \text{conn}(D, \omega)$.*

Before we proceed, let us verify that Theorem 8 and Lemma 9 together imply Theorem 1. Given an instance (G, F, C, k) of k -MEDIAN, we first apply Theorem 8 to obtain a coreset ω with support of size $\mathcal{O}(k\varepsilon^{-2} \log n)$. Next, we pass this coreset to Lemma 9, thus obtaining a set $F_0 \subseteq F$ of size $k \cdot (\varepsilon^{-1} \log n)^{\mathcal{O}(\varepsilon^{-3})}$. Let D_0 be the subset of F_0 of size at most k that minimizes $\text{conn}(D_0, \omega)$. Then using the approximation guarantees of Theorem 8 and Lemma 9, for any $D \subseteq F$ we have

$$\text{conn}(D_0, C) \leq (1 + \varepsilon) \text{conn}(D_0, \omega) \leq (1 + \varepsilon)^2 \text{conn}(D, \omega) \leq (1 + \varepsilon)^3 \text{conn}(D, C).$$

It remains to rescale ε . Hence, for the rest of this section we focus on proving Lemma 9.

Let $\mathcal{I} = (G, F, C, k)$ be an input k -MEDIAN instance with a weight function ω , where G is planar. Let $\varepsilon > 0$ be an accuracy parameter and without loss of generality assume that $\varepsilon < 1/4$. Let $n = |V(G)|$ and $m = |E(G)|$. Without loss of generality assume that $n = \Theta(m)$.

We assume that G is embedded in a sphere Σ and apply the necessary modifications explained in the beginning of Section 2.1 to fit into the framework of Voronoi diagrams. Denote $S = \text{supp}(\omega)$. We compute the Voronoi partition Cell_S induced by S and the Voronoi diagram Vor_S induced by S . By Proposition 6, Vor_S has $\mathcal{O}(|S|)$ vertices, faces, and edges.

Distance levels. We first compute an $\mathcal{O}(1)$ -approximate solution $\tilde{D} \subseteq F$ using the algorithm given by Feldman and Langberg [11, Theorem 15.1]; this algorithm outputs an $\mathcal{O}(1)$ -approximate solution with constant probability. Let us scale all the edge lengths in G by the same ratio so that

$$\text{conn}(\tilde{D}, \omega) = |S|/\varepsilon. \tag{1}$$

Next, we assign length $+\infty$ to every edge of length larger than $\text{conn}(\tilde{D}, \omega)$; clearly, they are not used in the computation of the connection cost of an optimum solution. Without loss of

generality we assume that all the distances between vertices in G are finite: otherwise we can split the instance into a number of independent ones, compute a suitable set F_0 for each of them and take the union.

The next step is to assign levels to distances in the graph. For any $c \in [0, +\infty)$, define the *level* of c , denoted $\text{level}(c)$, to be the smallest nonnegative integer ℓ such that $c < (1 + \varepsilon)^\ell$. Note that $\text{level}(c) = 0$ if and only if $c < 1$. Let $L = 1 + \text{level}(m \cdot \text{conn}(\tilde{D}, \omega))$, then we have

$$\text{level}(\text{dist}(u, v)) \in \{0, 1, \dots, L - 1\} \quad \text{for all } u, v \in V(G).$$

Observe that since $m = \Theta(n)$, by (1) we have

$$L \leq \mathcal{O}(\varepsilon^{-1} \log(m|S|/\varepsilon)) \leq \mathcal{O}(\varepsilon^{-2} \log n). \quad (2)$$

Portals and profiles. Let $\tau = (p, u, f)$ be an incidence in Vor_S . Let $d(\tau) = \text{dist}(p, u)$ and let $\ell(\tau) = \text{level}(d(\tau))$; note that $\text{Spoke}(\tau)$ has length exactly $d(\tau)$. For every integer $\iota \in \{1, \dots, \ell(\tau)\}$, we define the *portal* $\mathbf{p}\langle\tau, \iota\rangle$ as a vertex on $\text{Spoke}(\tau)$ at distance exactly $(1 + \varepsilon)^{\iota-1}$ from p ; we subdivide an edge and create a new vertex to accommodate $\mathbf{p}\langle\tau, \iota\rangle$ if necessary. Furthermore, we add also a portal $\mathbf{p}\langle\tau, 0\rangle = p$. Since $\ell(\tau) = \text{level}(d(\tau)) < L$, there are at most L portals on the spoke $\text{Spoke}(\tau)$.

Consider a diamond $\text{Diam}(e)$ induced by some edge e of Vor_S , and a vertex w in $\text{Diam}(e)$. Recall that the perimeter of $\text{Diam}(e)$ consists of spokes $\text{Spoke}(\tau_{i,j})$ for four incidence $\tau_{i,j}$, where $i, j \in \{1, 2\}$. The *profile* of a vertex w belonging to the diamond $\text{Diam}(e)$ consists of the following information, for all $\tau \in \{\tau_{i,j} : i, j \in \{1, 2\}\}$:

1. The minimum index $\lambda \in \{0, 1, \dots, \ell(\tau)\}$ satisfying

$$\text{dist}(\mathbf{p}\langle\tau, \lambda\rangle, p) > \varepsilon \cdot \text{dist}(\mathbf{p}\langle\tau, \lambda\rangle, w),$$

where $p = \mathbf{p}\langle\tau, 0\rangle$ is the vertex of S involved in τ . If no such index exists, we set $\lambda = \ell(\tau)$.

2. Letting

$$I = (\{0\} \cup \{\iota : |\iota - \lambda| \leq 1/\varepsilon^3\}) \cap \{0, 1, \dots, \ell(\tau)\},$$

the profile records the value of $\text{level}(\text{dist}(w, \mathbf{p}\langle\tau, \iota\rangle))$ for all $\iota \in I$.

Whenever speaking about a vertex w and incidence τ , we use $\lambda(\tau, w)$ and $I(\tau, w)$ to denote λ and I as above. We note that in total there are only few possible profiles.

▷ **Claim 10.** The number of possible different profiles of vertices in $\text{Diam}(e)$ is $L^{\mathcal{O}(\varepsilon^{-3})}$.

Proof. Since $0 \leq \ell(\tau) < L$ for every incidence τ , there are at most L^4 choices for the four values $\lambda(\tau_{i,j}, w)$ for $i, j \in \{1, 2\}$. Further, we have $|I(\tau_{i,j}, w)| \leq \mathcal{O}(\varepsilon^{-3})$, so there are at most $L^{\mathcal{O}(\varepsilon^{-3})}$ choices for the values $\text{level}(\text{dist}(w, \mathbf{p}\langle\tau_{i,j}, \iota\rangle))$ for $i, j \in \{1, 2\}$ and $\iota \in I(\tau_{i,j}, w)$. ◁

For future reference, we state the key property of profiles: having the same profile implies having approximately same distances to the profiles with indices in I .

▷ **Claim 11.** Suppose w and w' are two vertices of $\text{Diam}(e)$ that have the same profile. Then for each $\tau \in \{\tau_{i,j} : i, j \in \{1, 2\}\}$ and $\iota \in I(\tau, w)$, we have

$$\text{dist}(w', \mathbf{p}\langle\tau, \iota\rangle) \leq (1 + \varepsilon) \cdot \text{dist}(w, \mathbf{p}\langle\tau, \iota\rangle) + 1.$$

Proof. Let $\ell = \text{level}(\text{dist}(w, \mathbf{p}\langle\tau, \iota\rangle)) = \text{level}(\text{dist}(w', \mathbf{p}\langle\tau, \iota\rangle))$, as recorded in the common profile. If $\ell = 0$, then $\text{dist}(w', \mathbf{p}\langle\tau, \iota\rangle) < 1$ and we are done. Otherwise, $\text{dist}(w, \mathbf{p}\langle\tau, \iota\rangle)$ and $\text{dist}(w', \mathbf{p}\langle\tau, \iota\rangle)$ are both contained in the interval $[(1 + \varepsilon)^{\ell-1}, (1 + \varepsilon)^\ell]$. This interval has length $\varepsilon \cdot (1 + \varepsilon)^{\ell-1} \leq \varepsilon \cdot \text{dist}(w, \mathbf{p}\langle\tau, \iota\rangle)$, hence the claim follows. ◁

33:10 Efficient Approximation Schemes for Planar Clustering

Construction of the set F_0 . We now construct the set F_0 as follows: for every diamond $\text{Diam}(e)$ and every possible profile in $\text{Diam}(e)$, include in F_0 one facility with that profile (if one exists). Since there are $\mathcal{O}(|S|)$ diamonds, by Claim 10 and (10) we have

$$|F_0| \leq \mathcal{O}(|S|) \cdot L^{\mathcal{O}(\varepsilon^{-3})} = |\text{supp}(\omega)| \cdot (\varepsilon^{-1} \log n)^{\mathcal{O}(\varepsilon^{-3})},$$

as claimed. It remains to prove that F_0 has the claimed approximation properties.

For every facility $w \in F$, pick a diamond $\text{Diam}(e)$ containing w and let $f(w)$ to be the facility $f(w) \in F_0 \cap \text{Diam}(e)$ that has the same profile as w . Fix a solution $D^* \subseteq F$ with $|D^*| \leq k$ minimizing $\text{conn}(D^*, \omega)$. Let $D_0 = \{f(w) : w \in D^*\}$. Clearly, $|D_0| \leq |D^*| \leq k$. To finish the proof of Lemma 9 it suffices to show that

$$\text{conn}(D_0, \omega) \leq (1 + \mathcal{O}(\varepsilon))\text{conn}(D^*, \omega). \quad (3)$$

To this end, consider any client $v \in S = \text{supp}(\omega)$ and let $w \in D^*$ be the facility in D^* serving v , that is, $\text{dist}(v, w) = \text{dist}(v, D^*)$. To show (3), it suffices to prove that

$$\text{dist}(v, f(w)) \leq (1 + \mathcal{O}(\varepsilon))\text{dist}(v, w) + \mathcal{O}(1). \quad (4)$$

Indeed, by summing (4) through all $v \in S$ and using (1) we obtain

$$\begin{aligned} \text{conn}(D_0, \omega) &\leq (1 + \mathcal{O}(\varepsilon))\text{conn}(D^*, \omega) + \mathcal{O}(1) \cdot |S| \\ &\leq (1 + \mathcal{O}(\varepsilon))\text{conn}(D^*, \omega) + \mathcal{O}(\varepsilon) \cdot \text{conn}(\tilde{D}, \omega) \leq (1 + \mathcal{O}(\varepsilon))\text{conn}(D^*, \omega), \end{aligned}$$

where the last inequality is due to \tilde{D} being an $\mathcal{O}(1)$ -approximate solution.

Hence, from now on we focus on proving (4). Let $\text{Diam}(e)$ be the diamond containing w and $f(w)$. Consider the shortest path P from w to v in G . By Lemma 7, the path P intersects the perimeter of the diamond $\text{Diam}(e)$. Let u be the vertex on the perimeter of $\text{Diam}(e)$ that lies on P and, among such, is closest to w on P . Since P is a shortest path, the length of the subpaths of P between v and u and between u and w equal $\text{dist}(v, u)$ and $\text{dist}(u, w)$, respectively, and in particular $\text{dist}(v, w) = \text{dist}(v, u) + \text{dist}(u, w)$.

We now observe that to prove (4), it suffices show the following.

$$\text{dist}(u, f(w)) \leq \text{dist}(u, w) + \mathcal{O}(\varepsilon)\text{dist}(v, w) + \mathcal{O}(1). \quad (5)$$

Indeed, assuming (5) we have

$$\begin{aligned} \text{dist}(v, f(w)) &\leq \text{dist}(v, u) + \text{dist}(u, f(w)) \\ &\leq \text{dist}(v, u) + \text{dist}(u, w) + \mathcal{O}(\varepsilon)\text{dist}(v, w) + \mathcal{O}(1) \\ &= \text{dist}(v, w) + \mathcal{O}(\varepsilon)\text{dist}(v, w) + \mathcal{O}(1). \end{aligned}$$

Hence, from now on we focus on proving (5).

Let $\tau_{i,j}$ for $i, j \in \{1, 2\}$, be the four incidences involved in the diamond $\text{Diam}(e)$. Since u lies on the perimeter of $\text{Diam}(e)$, actually u lies on $\text{Spoke}(\tau)$, where $\tau = \tau_{i,j}$ for some $i, j \in \{1, 2\}$. Let $p = \mathbf{p}\langle \tau, 0 \rangle$ be the vertex of S involved in the incidence τ . Since $u \in \text{Cell}_S(p)$ while $v \in S$, we have that $\text{dist}(u, p) \leq \text{dist}(u, v)$. Consequently, we have $\text{dist}(u, w) \leq \text{dist}(v, w)$ and $\text{dist}(u, p) \leq \text{dist}(v, w)$, so to prove (5) it suffices to prove the following:

$$\text{dist}(u, f(w)) \leq \text{dist}(u, w) + \mathcal{O}(\varepsilon)(\text{dist}(u, w) + \text{dist}(u, p)) + \mathcal{O}(1). \quad (6)$$

Let $\mathbf{p}_u = \mathbf{p}\langle \tau, \iota \rangle$ be the portal on the subpath of $\text{Spoke}(\tau)$ between u and p that is closest to u . Intuitively, \mathbf{p}_u is a good approximation of u and distances from \mathbf{p}_u are almost the same as distances from u . As this idea will be repeatedly used in this sequel, we encapsulate it in a single claim.

▷ Claim 12. Suppose for some vertices x and y we have

$$\text{dist}(\mathbf{p}_u, x) \leq A \cdot \text{dist}(\mathbf{p}_u, y) + B,$$

for some A, B . Then

$$\text{dist}(u, x) \leq A \cdot \text{dist}(u, y) + B + (A + 1)\varepsilon \cdot \text{dist}(u, p) + (A + 1).$$

Proof. By the choice of \mathbf{p}_u we have


$$\text{dist}(\mathbf{p}_u, p) \leq \text{dist}(u, p) \leq (1 + \varepsilon)\text{dist}(\mathbf{p}_u, p) + 1,$$

so $\text{dist}(u, \mathbf{p}_u) \leq \varepsilon \cdot \text{dist}(u, p) + 1$. Therefore, we have

$$\begin{aligned} \text{dist}(u, x) &\leq \text{dist}(\mathbf{p}_u, x) + \text{dist}(\mathbf{p}_u, u) \leq A \cdot \text{dist}(\mathbf{p}_u, y) + B + \text{dist}(\mathbf{p}_u, u) \\ &\leq A \cdot (\text{dist}(u, y) + \text{dist}(\mathbf{p}_u, u)) + B + \text{dist}(\mathbf{p}_u, u) \\ &= A \cdot \text{dist}(u, y) + B + (A + 1) \cdot \text{dist}(\mathbf{p}_u, u) \\ &\leq A \cdot \text{dist}(u, y) + B + (A + 1) \cdot \varepsilon \cdot \text{dist}(u, p) + (A + 1), \end{aligned}$$

as claimed. ◁

Since w and $f(w)$ have the same profile, we may denote $\lambda = \lambda(\tau, w) = \lambda(\tau, f(w))$ and $I = I(\tau, w) = I(\tau, f(w))$. Further, let $\mathbf{p}_\lambda = \mathbf{p}(\tau, \lambda)$. We now consider a number of cases depending on the relative values of ι and λ , with the goal on proving that (6) holds in each case. See Figure 2 for an illustration.



■ **Figure 2** The diamond $\text{Diam}(e)$ with vertices v, u, w, p , and $f(w)$. Red vertices are clients, black squares are portals. The case distinction in the proof corresponds to relative order of \mathbf{p}_λ and \mathbf{p}_ι .

Middle case: $\iota \in I$. As profiles of w and $f(w)$ are the same and $\iota \in I$, by Claim 11 we have

$$\text{dist}(\mathbf{p}_u, f(w)) \leq (1 + \varepsilon)\text{dist}(\mathbf{p}_u, w) + 1.$$

It suffices now to apply Claim 12 to infer inequality (6).

Close case: $\iota < \lambda - 1/\varepsilon^3$. Since $\iota < \lambda - 1/\varepsilon^3$, by the choice of λ we have

$$\text{dist}(\mathbf{p}_u, p) \leq \varepsilon \cdot \text{dist}(\mathbf{p}_u, w).$$

By applying Claim 12, we infer that

$$\text{dist}(u, p) \leq \varepsilon \cdot \text{dist}(u, w) + (1 + \varepsilon)\varepsilon \cdot \text{dist}(u, p) + \mathcal{O}(1) \leq \varepsilon \cdot \text{dist}(u, w) + \text{dist}(u, p)/2 + \mathcal{O}(1),$$

33:12 Efficient Approximation Schemes for Planar Clustering

which entails

$$\text{dist}(u, p) \leq 2\varepsilon \cdot \text{dist}(u, w) + \mathcal{O}(1).$$

Since $0 \in I$ and the profiles of w and $f(w)$ are equal, by Claim 11 we have

$$\text{dist}(p, f(w)) \leq (1 + \varepsilon) \cdot \text{dist}(p, w) + 1.$$

By combining the last two inequalities, we conclude that

$$\begin{aligned} \text{dist}(u, f(w)) &\leq \text{dist}(u, p) + \text{dist}(p, f(w)) \\ &\leq \text{dist}(u, p) + (1 + \varepsilon)\text{dist}(p, w) + 1 \\ &\leq \text{dist}(u, p) + (1 + \varepsilon)(\text{dist}(u, p) + \text{dist}(u, w)) + 1 \\ &= (2 + \varepsilon)\text{dist}(u, p) + (1 + \varepsilon)\text{dist}(u, w) + 1 \\ &\leq (1 + \mathcal{O}(\varepsilon))\text{dist}(u, w) + \mathcal{O}(1). \end{aligned}$$

Thus, inequality (6) holds in this case.

Far case: $\iota > \lambda + 1/\varepsilon^3$. By the definition of λ and since $\iota > \lambda + 1/\varepsilon^3$, we have in particular $\lambda < \ell(\tau)$ and hence

$$\varepsilon \cdot \text{dist}(\mathbf{p}_\lambda, w) < \text{dist}(\mathbf{p}_\lambda, p).$$

On the other hand, we have

$$\text{dist}(p, \mathbf{p}_\lambda) = (1 + \varepsilon)^{\lambda - \iota} \text{dist}(p, \mathbf{p}_u) \leq (1 + \varepsilon)^{-1/\varepsilon^3} \text{dist}(p, \mathbf{p}_u) \leq \varepsilon^2 \cdot \text{dist}(p, \mathbf{p}_u),$$

where the last step follows from Bernoulli's inequality. By combining the above two inequalities we obtain

$$\text{dist}(\mathbf{p}_\lambda, w) \leq \varepsilon \cdot \text{dist}(p, \mathbf{p}_u),$$

implying

$$\text{dist}(p, w) \leq \text{dist}(p, \mathbf{p}_\lambda) + \text{dist}(\mathbf{p}_\lambda, w) \leq (\varepsilon + \varepsilon^2) \cdot \text{dist}(p, \mathbf{p}_u) \leq 2\varepsilon \cdot \text{dist}(p, \mathbf{p}_u) \leq 2\varepsilon \cdot \text{dist}(p, u).$$

As before, by Claim 11 we have $\text{dist}(p, f(w)) \leq (1 + \varepsilon)\text{dist}(p, w) + 1$ due to $0 \in I$, hence

$$\begin{aligned} \text{dist}(u, f(w)) &\leq \text{dist}(u, p) + \text{dist}(p, f(w)) \\ &\leq (\text{dist}(u, w) + \text{dist}(p, w)) + (1 + \varepsilon)\text{dist}(p, w) + 1 \\ &\leq \text{dist}(u, w) + (2 + \varepsilon) \cdot (2\varepsilon) \cdot \text{dist}(u, p) + 1 \\ &\leq \text{dist}(u, w) + \mathcal{O}(\varepsilon)\text{dist}(u, p) + \mathcal{O}(1). \end{aligned}$$

We conclude that inequality (6) holds in this case.

As the case investigation presented above covers all the possibilities, the proof of Lemma 9 is complete, so we have also proved Theorem 1.

References

- 1 Vijay Arya, Naveen Garg, Rohit Khandekar, Adam Meyerson, Kamesh Munagala, and Vinayaka Pandit. Local Search Heuristics for k -Median and Facility Location Problems. *SIAM J. Comput.*, 33(3):544–562, 2004.
- 2 Vladimir Braverman, Dan Feldman, and Harry Lang. New Frameworks for Offline and Streaming Coreset Constructions. *CoRR*, abs/1612.00889, 2016. [arXiv:1612.00889](#).
- 3 Jarosław Byrka, Thomas Pensyl, Bartosz Rybicki, Aravind Srinivasan, and Khoa Trinh. An Improved Approximation for k -Median and Positive Correlation in Budgeted Optimization. *ACM Trans. Algorithms*, 13(2):23:1–23:31, 2017.
- 4 Ke Chen. On Coresets for k -Median and k -Means Clustering in Metric and Euclidean Spaces and Their Applications. *SIAM J. Comput.*, 39(3):923–947, 2009.
- 5 Vincent Cohen-Addad. A Fast Approximation Scheme for Low-dimensional k -means. In *SODA 2018*, pages 430–440. SIAM, 2018.
- 6 Vincent Cohen-Addad, Andreas Emil Feldmann, and David Saulpic. Near-Linear Time Approximation Schemes for Clustering in Doubling Metrics. *CoRR*, abs/1812.08664, 2018. [arXiv:1812.08664](#).
- 7 Vincent Cohen-Addad, Anupam Gupta, Amit Kumar, Euiwoong Lee, and Jason Li. Tight FPT Approximations for k -Median and k -Means. *ICALP’19*, 2019.
- 8 Vincent Cohen-Addad, Philip N. Klein, and Claire Mathieu. Local Search Yields Approximation Schemes for k -Means and k -Median in Euclidean and Minor-Free Metrics. In *FOCS 2016*, pages 353–364. IEEE Computer Society, 2016.
- 9 Vincent Cohen-Addad, Marcin Pilipczuk, and Michał Pilipczuk. A Polynomial-Time Approximation Scheme for Facility Location on Planar Graphs. In *FOCS 2019*, 2019.
- 10 Vincent Cohen-Addad, Marcin Pilipczuk, and Michał Pilipczuk. Efficient approximation schemes for uniform-cost clustering problems in planar graphs. *CoRR*, abs/1905.00656, 2019. [arXiv:1905.00656](#).
- 11 Dan Feldman and Michael Langberg. A Unified Framework for Approximating and Clustering Data. *CoRR*, abs/1106.1379, 2011. Conference version presented at STOC 2011. [arXiv:1106.1379](#).
- 12 Wenceslas Fernandez de la Vega, Marek Karpinski, Claire Kenyon, and Yuval Rabani. Approximation schemes for clustering problems. In *STOC 2003*, pages 50–58, 2003.
- 13 Sudipto Guha and Samir Khuller. Greedy Strikes Back: Improved Facility Location Algorithms. *J. Algorithms*, 31(1):228–248, 1999.
- 14 Sarel Har-Peled. Quasi-Polynomial Time Approximation Scheme for Sparse Subsets of Polygons. In *SoCG 2014*. ACM, 2014.
- 15 Dorit S. Hochbaum. Heuristics for the fixed cost median problem. *Math. Program.*, 22(1):148–162, 1982.
- 16 K. Jain and V. Vazirani. Approximation algorithms for metric facility location and k -Median problems using the primal-dual schema and Lagrangian relaxation. *J. ACM*, 48(2):274–296, 2001.
- 17 Amit Kumar, Yogish Sabharwal, and Sandeep Sen. Linear-time approximation schemes for clustering problems in any dimensions. *J. ACM*, 57(2), 2010.
- 18 Shi Li. A 1.488 approximation algorithm for the uncapacitated facility location problem. *Inf. Comput.*, 222:45–58, 2013.
- 19 Dániel Marx and Michał Pilipczuk. Optimal Parameterized Algorithms for Planar Facility Location Problems Using Voronoi Diagrams. In *ESA 2015*, volume 9294 of *Lecture Notes in Computer Science*, pages 865–877. Springer, 2015. See [arXiv:1504.05476](#) for the full version.
- 20 Jiří Matoušek. On Approximate Geometric k -Clustering. *Discrete & Computational Geometry*, 24(1):61–84, 2000.

33:14 Efficient Approximation Schemes for Planar Clustering

- 21 Michał Pilipczuk, Erik Jan van Leeuwen, and Andreas Wiese. Quasi-Polynomial Time Approximation Schemes for Packing and Covering Problems in Planar Graphs. In *ESA 2018*, volume 112 of *LIPICs*, pages 65:1–65:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018.
- 22 David B Shmoys, Éva Tardos, and Karen Aardal. Approximation algorithms for facility location problems. In *STOC 1997*, pages 265–274. ACM, 1997.