

Almost-Optimal Upper and Lower Bounds for Clustering in Low Dimensional Euclidean Spaces

Vincent Cohen-Addad 

Google Research, New York, NY, USA

Karthik C. S. 

Rutgers University, Piscataway, NJ, USA

David Saulpic 

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

Chris Schwiegelshohn 

Aarhus University, Denmark

Abstract

The k -median and k -means clustering objectives are classic objectives for modeling clustering in a metric space. Given a set of points in a metric space, the goal of the k -median (resp. k -means) problem is to find k representative points so as to minimize the sum of the distances (resp. sum of squared distances) from each point to its closest representative. Cohen-Addad, Feldmann, and Saulpic [JACM'21] showed how to obtain a $(1+\varepsilon)$ -factor approximation in low-dimensional Euclidean metric for both the k -median and k -means problems in near-linear time $2^{(1/\varepsilon)O(d^2)} n \cdot \text{polylog}(n)$ (where d is the dimension and n is the number of input points).

We improve this running time to $2^{\widetilde{O}(1/\varepsilon)^{d-1}} \cdot n \cdot \text{polylog}(n)$, and show an almost matching lower bound: under the Gap Exponential Time Hypothesis for 3-SAT, there is no $2^{o(1/\varepsilon^{d-1})} n^{O(1)}$ algorithm achieving a $(1+\varepsilon)$ -approximation for k -means.

2012 ACM Subject Classification Theory of computation $\rightarrow$ Facility location and clustering; Theory of computation $\rightarrow$ Computational geometry

Keywords and phrases k -means clustering, k -median clustering, Euclidean space, Fine-Grained Complexity

Digital Object Identifier 10.4230/LIPIcs.SoCG.2026.34

Related Version *Full Version*: <https://arxiv.org/abs/2603.09846>

Funding *Karthik C. S.*: This work was supported by the National Science Foundation under Grants CCF-2313372 and CCF-2443697, a grant from the Simons Foundation, Grant Number 825876, Awardee Thu D. Nguyen, and partially funded by the Ministry of Education and Science of Bulgaria's support for INSAIT, Sofia University "St. Kliment Ohridski" as part of the Bulgarian National Roadmap for Research Infrastructure.

Chris Schwiegelshohn: This work was partially supported by the Independent Research Fund Denmark (DFR) under a Sapere Aude Research Leader grant No 1051-00106B and by a Google Research Award.

1 Introduction

The k -median and k -means problems are minimization tasks of classic objectives for modeling clustering in a wide variety of applications arising in data mining and machine learning. Given a set of points $P \subset X$ and a set of *candidate centers* $\mathcal{S} \subset X$ in a metric space (X, Δ) ,



© Vincent Cohen-Addad, Karthik C. S., David Saulpic, and Chris Schwiegelshohn; licensed under Creative Commons License CC-BY 4.0


42nd International Symposium on Computational Geometry (SoCG 2026).

Editors: Hee-Kap Ahn, Michael Hoffmann, and Amir Nayyeri; Article No. 34; pp. 34:1–34:16

Leibniz International Proceedings in Informatics



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



the goal of the k -median problem¹ is to find a set of k points $S \subset \mathcal{S}$, called *centers*, so as to minimize the sum of the distances (given by Δ) from each point of P to its closest center in S (in the k -means problem, the goal is to instead minimize the sum of the squared distances from each point of P to its closest center in S).

The algorithmic study of k -median started in the early '60s in the operations research community [36, 46, 29], while the study of k -means was arguably initiated by the seminal work of Lloyd [39], which applied clustering to quantization of analog signals. Since then, both problems have received a tremendous amount of attention. Naturally, the complexity of the problems varies with the underlying metric space hosting the input points. On the negative side, both problems are known to be NP-hard in Euclidean space, even when the points lie in the Euclidean plane (and k is large) [44, 40], or when $k = 2$ [23, 2] (and the dimension is large). When both k and d are arbitrary, several groups of researchers have shown hardness of approximation results for both problems [6, 37, 17, 19]. Even stronger hardness results hold if the input lives in an arbitrary metric space (see Guha and Kuller [28] and [18]).

Fine-Grained Complexity. Despite those hardness results, the importance of Euclidean inputs in statistics and machine learning applications has led researchers to develop methods and algorithms, in particular through the study of the parameterized complexity of the problem, when the input lies in an Euclidean space. Both the dimensionality of the input, d , and the target number of clusters, k , have been studied as parameters. For exact algorithms, the seminal work of Inaba, Katoh and Imai [31] has shown that one can compute an exact solution to the k -median and k -means problems in time $O(n^{kd+1})$ – as opposed to the naive brute-force enumeration which runs in $\tilde{O}(k^n)$ time. On the negative side, Cohen-Addad, de Mesmay, Rotenberg and Roytman [11] much later showed that there is no $n^{o(k)}$ exact algorithm for (discrete) k -median or k -means problem, even when the dimension is four, and thus partially showing the optimality of the $O(n^{kd+1})$ algorithm.

To circumvent this hardness, researchers have turned to approximation algorithms. When the dimension d and the precision ε are taken as parameters, a line of work [5, 35, 27, 20, 10, 13] improved the running time of approximation schemes from quasi-polynomial time, $2^{(\log(n)/\varepsilon)^{d-1}}$, to $2^{(1/\varepsilon)^{O(d^2)}} n$. The doubly exponential dependency in d^2 stands in stark contrast with other similar results, e.g., for the Traveling Salesperson Problem (TSP). For this problem, Kisfaludi-Bak, Nederlof and Wegrzycki [34] recently settled the exact dependency on d , by showing an algorithm running essentially in time $2^{O(1/\varepsilon^{d-1})} n$, and a lower bound of $2^{\Omega((1/\varepsilon)^{d-1})}$ under the Gap-ETH hypothesis [25, 42]. In light of this result, we consider the following question:

► **Question 1.** *Is it possible to obtain a $2^{O(1/\varepsilon^{d-1})} n^{O(1)}$ approximation scheme for k -median and k -means problems in low-dimensional inputs? Would it be possible to get an even faster $2^{o(1/\varepsilon^{d-1})} n^{O(1)}$ approximation scheme for both problems?*

1.1 Our Results

In this paper, we make substantial progress towards answering Question 1. Our first contribution is the design of an improved algorithm:

► **Theorem 2.** *For every $\varepsilon > 0$ and dimension d , the k -median and k -means problems in $\mathbb{R}^d$ can both be approximated to a $(1 + \varepsilon)$ -factor in time $2^{\tilde{O}(1/\varepsilon^{d-1})} n \cdot \text{polylog}(n)$.²*

¹ This variant is sometimes referred to as the *discrete* k -median problem. In the continuous k -median problem we are allowed to pick the k centers S from anywhere in X .

² The $\tilde{O}$ notation hides an exponential dependency in d , and polynomial dependency in $\log(1/\varepsilon)$.

The above algorithm is our main contribution, and improves on the techniques of [13]. It also readily extends to the continuous k -median and k -means problems (see Remark 4).

On the technical side, this algorithm deepens our understanding of the *quadtree* decomposition for clustering. A long line of work based on this decomposition [5, 35, 27, 20, 10, 13] improved the running time of approximation schemes (i.e., $(1 + \varepsilon)$ -factor approximation algorithms for any small enough constant ε) for clustering from quasi-polynomial time, $2^{(\log(n)/\varepsilon)^{d-1}}$, to $2^{(1/\varepsilon)^{O(d^2)}} n$ [13]. The main conceptual insight of these works is to bound the number of *portals* necessary along the decomposition so that forcing a path to go through portals incurs at most a $(1 + \varepsilon)$ -factor distortion. The main contribution of [13] is that after processing a tiny fraction of the input, only a constant number of portals are sufficient at each level of the quadtree decomposition. However, this constant remains much larger than that required for other natural problems.

Indeed, carefully analyzing the performances of the quadtree dissection and the number of required portals to achieve a $(1 + \varepsilon)$ -factor approximation has been a fruitful direction of research to design approximation schemes in Euclidean space. The initial work of Arora [3, 4] showed approximation schemes for the Traveling Salesperson Problem (TSP) and other routing problems, first in quasi-polynomial time (for $d > 2$) and then with running time $n \log(n)^{O(1/\varepsilon)^{d-1}}$. For TSP, the running time was subsequently significantly improved: first by Rao and Smith [45] who achieved $(1/\varepsilon)^{1/\varepsilon^{d-1}} n \log n$, then by Bartal and Gottlieb [7] who brought it down to $2^{1/\varepsilon^{O(d)}} n$, linear in n , but suboptimal in d . Finally, [34] settled the exact dependency on d , by showing an algorithm running essentially in time $2^{O(1/\varepsilon^{d-1})} n$, and a lower bound of $2^{\Omega((1/\varepsilon)^{d-1})} \text{poly}(n)$ under the Gap-ETH hypothesis [25, 42].

In light of this literature, our contribution is therefore an improved analysis of the quadtree dissection for clustering problems, that allows us to (almost) match the number of portals required for the well-studied TSP. This requires an analysis and techniques distinct from those for TSP, in particular because they fail if distances are squared – as what happens in the k -means problem. We note that, due to its simplicity, applications of the quadtree extend far beyond approximation algorithms in low-dimensional space. For clustering, quadtrees are used for instance in the streaming scenario, when memory has to be kept low [8], but also to design linear time algorithms in high-dimensional spaces [21] or to construct differentially-private algorithms [12]. Therefore, better understanding the quadtree's performance, even in the simplest scenario, is likely to have a far-reaching impact.

We complement Theorem 2 with a conditional lower bound, which shows the near-optimality of our algorithm under the Gap-ETH hypothesis:

► **Theorem 3.** *Assuming the Gap-ETH hypothesis, for every integer $d \geq 2$, there exists a constant $c > 0$ such that no approximation scheme can, given an instance of discrete k -means (or discrete k -median) with N points in Euclidean space $\mathbb{R}^d$ and a parameter $\varepsilon > 0$, compute a $(1 + \varepsilon)$ -factor approximation in time $2^{c(1/\varepsilon)^{d-1}} \text{poly}(N)$.*

This lower bound is obtained by amalgamating the technical ideas of [37, 11, 17] and the framework of de Berg, Bodlaender, Kisfaludi-Bak, Marx and van der Zanden [24]. Ideas to reduce the Vertex Cover problem to clustering problems were elaborated in [37, 11, 17], and [24] built a framework to prove fine-grained lower bounds for a host of geometric problems in low dimensions. We put them together to obtain Theorem 3. The proof is deferred to the full version.

Extensions. The framework developed in [13] applies to many variants of k -means (and k -median), namely prize-collecting k -means, k -means with outliers and Facility location. As our structure theorem is a direct improvement of theirs, we also improve the running time

for those problems. We give a sketch of the arguments in the full version. Our analysis also improves the result for doubling metric (a generalization of Euclidean space), however attaining a non-tight complexity $2^{\tilde{O}(1/\varepsilon^d)}$. We therefore focused this paper on the Euclidean setting, for which our analysis is almost tight.

1.2 Our Techniques

For the ease of presentation, we detail below our upper bound for the k -means problem in the Euclidean plane (i.e., with $d = 2$). Our input is a set of points $P \subset \mathbb{R}^2$, a set of candidate centers $\mathcal{C} \subset \mathbb{R}^2$, and the goal is to compute a set of centers $S \in \mathcal{C}^k$ ($\mathcal{C}^k$ is the k -wise partition product of $\mathcal{C}$) that minimizes the sum of squared distances from points of P to their closest center in S .

Our starting point is the quadtree dissection equipped with portals. We briefly recall in the next few paragraphs its construction and main property.

First, find an axis-aligned rectangle that contains all input points. Then, randomly split the rectangle into 4 axis-aligned rectangles of side length reduced by a factor of approximately 2. Continue this splitting recursively until each rectangle contains a single point: this defines a recursive decomposition of the input. We say that two points are *cut at level i* when the smallest rectangle containing both points is at the i^{th} level of the decomposition. Assuming that the leaves of the decomposition have diameter 1, the rectangles at the i^{th} level have diameter roughly 2^i .

Now, a set of *portals* is placed regularly along the boundary of each rectangle. Instead of connecting points to centers via straight lines, we consider only *portal-respecting* paths, defined as follows. To connect two points p and q cut at level i , a portal-respecting path consists of a sequence of segments connecting $p, p_1, \dots, p_i, q_i, \dots, q_1, q$, where p_j (resp. q_j) is a portal of the rectangle at level j containing p (resp. q). Considering only those paths enforces some detours, but a simple dynamic program can compute the best portal-respecting solution, whose complexity naturally depends (exponentially) on the number of portals. Furthermore, if $1/\rho$ many portals are placed along each boundary, it can be computed that a portal-respecting path between two points p, q cut at level i has length (essentially) $\|p - q\| + \rho 2^i$. Indeed, since the boundaries of the rectangles are geometrically decreasing, the detour to go through a portal at level i dominates all others; and since the rectangles at this level have diameter 2^i , by regularly placing $1/\rho$ many portals, we can ensure that there is a portal at distance at most $\rho 2^i$ from the point where the straight line path crosses the boundary. Therefore, the goal of the analysis is to show that, for a small number of portals, a portal-respecting solution with almost optimal cost exists.

The standard argument goes as follows. The main property of the quadtree decomposition is that, for any pair of points $p, q \in \mathbb{R}^2$, the probability that they are cut at level i is $\|p - q\|/2^i$. By using $1/\rho$ portals, the detour to connect p and q via portals in case they are cut at level i is $\rho 2^i$. Therefore, the expected detour at level i is $\|p - q\|/2^i \cdot \rho 2^i = \rho \|p - q\|$. Since there are $\log n$ levels in the decomposition, the expected detour is $\rho \|p - q\| \cdot \log n$, and one has to take $\rho = \varepsilon / \log n$ to get a tiny expected detour.

However, having $\log n / \varepsilon$ many portals does not yield an efficient algorithm. More importantly, this analysis does not extend to k -means problem: although the expected distance between p and q is small, the expected *squared* distance is not! Indeed, the expected squared detour at level i is $\|p - q\|^2 / 2^{2i} \cdot (\rho 2^i)^2 = \rho^2 \|p - q\|^2$, and the term 2^i cannot be easily controlled. To cope with this, [13] introduced a careful preprocessing of the input, to ensure that no point is cut at a level higher than $\|p - q\|/\varepsilon$. In that case, taking $1/\varepsilon^2$ portals is enough to ensure a small detour – and also small squared detour, as this is not an average-case argument (unlike the previously seen argument). This preprocessing is

essentially the following: start from a constant-factor approximate solution $\mathcal{A}$ (which is a set of k centers), and for each point p , let $\mathcal{A}_p$ be its distance to the closest center of $\mathcal{A}$. If the ball $B(p, \mathcal{A}_p/\varepsilon)$ is cut by the decomposition at a level higher than $\log(\mathcal{A}_p/\varepsilon^2)$, then replace p by a copy of its closest center in $\mathcal{A}$. They showed that each point is replaced with probability at most ε , therefore yielding an expected cost increase of $\varepsilon \text{cost}(P, \mathcal{A}) = O(\varepsilon) \text{OPT}$. Furthermore, [13] showed that placing $1/\varepsilon^3$ many portals ensured that, for any center of $\mathcal{A}$, either it can be connected to the optimal solution through portals with a small detour, or one can assume that the center is part of the optimal solution. Therefore, after preprocessing, each client can be connected through portals to the optimal solution, and dynamic programming allows to compute a near-optimal solution.

However, this analysis may be seen as “worst-case”, in the sense that it treats all the clients for which the ball is cut at a level below $\log(\mathcal{A}_p/\varepsilon^2)$ in the same way. Our contribution is to mix the average-case analysis with the techniques of [13] in order to reduce the number of necessary portals. We carefully define a budget for each point, according to the level where the point is cut from the approximate solution $\mathcal{A}$, as in [13], but now also according to the level where it is cut from the optimal solution. We first show that, with constant probability, this budget is very cheap, namely an ε -fraction of the optimal cost. We then show that this budget is enough to pay for the preprocessing of the input, and for the detour incurred by making a solution portal-respecting. The algorithm is the same as the one of [13]: our contribution lies in a much tighter analysis.

This analysis differs from [13] as the detour they tolerated was independent of the optimal solution: using this additional information allows us to be more precise, and reduce the number of required portals to $(\log(1/\varepsilon)/\varepsilon)^{d-1}$, instead of $1/\varepsilon^{O(d)}$. Although the global picture remains similar, finding the right balance between a budget (1) small enough to not blow up the cost, and (2) that is still enough to pay for making the optimal solution portal-respecting is a non-trivial extension of their proof. Indeed, it requires new insights on the way to connect a point p that is *badly cut*, namely cut from its optimal center s at a level higher than $\log(\|p - s\|/\varepsilon)$. For this, we show that our budget allows to connect the point in a portal-respecting way to the center $\tilde{s}$ closest to $\mathcal{A}(p)$, where $\mathcal{A}(p)$ is the closest center to p in the constant-factor approximate solution $\mathcal{A}$.

1.3 Further Related Work

We covered previously the hardness results for k -means and k -medians; in terms of upper-bound, for general metrics a recent breakthrough from [15] achieves a $(2 + \varepsilon)$ -approximation for k -median, and a 5.6-approximation for k -means [9]. Those two works improve a long line of works, based on primal-dual techniques and the so-called Lagrangian-multiplier preserving algorithms for Facility Location, see e.g. [33, 32, 38].

To improve those approximation ratio, one can resort to algorithms parameterized by the number of clusters, k : one can get FPT algorithm with approximation matching the lower bound of [32], namely $1 + 2/e$ for k -median and $1 + 8/e$ for k -means [16].

Other techniques, specific to Euclidean Space, have been developed in the past decades. Dimension reduction [41] and coresset [22] allow to reduce the dimension of the input to $O(\log k/\varepsilon^{-2})$ and the number of distinct points to $\tilde{O}(k\varepsilon^{-O(1)})$, hence leading to simple FPT algorithms based on naive enumeration. Recent papers improved this simple approach, reducing the dependency in k (in the exponent) [1].

We covered in the introduction most of the literature related on low-dimensional clustering – which is the study of the problems parameterized by d . To compute a $(1 + \varepsilon)$ -approximation in time polynomial in n and k , any algorithm must have a running time at least doubly exponential in d , as the problem is APX-hard in dimension $\Omega(\log n)$. The best of these algorithms is from [13], with a near-linear running time of $f(\varepsilon, d) \cdot n \text{polylog } n$.

Finally, tractability of clustering has also been studied with different parameters, e.g., the cost by Fomin, Golovach and Simonov [26], who showed a $D^D \text{poly}(nd)$ exact algorithm for k -median, where D is the optimal cost.

We present in the full version how to adapt the argument for k -median. The main body contains the proof of the structure theorem that allows to reduce the number of portals required by the quadtree. We extend the algorithmic results to other objectives in the full version.

2 Preliminaries

2.1 Definitions

We consider the Euclidean space $\mathbb{R}^d$, with the ℓ_2 metric: $\text{dist}(p, q) := \sqrt{\sum_{i=1}^d (p_i - q_i)^2}$. For any point p and $r \geq 0$, the ball centered at p of radius r is $B(p, r) := \{x \in \mathbb{R}^d : \text{dist}(p, x) \leq r\}$. The (discrete) k -median and k -means problems are defined as follows: we are given a set of *clients* P and of *candidate centers* $\mathcal{C}$. A *solution* is any subset of $\mathcal{C}$ with size k . The goal is to compute the solution $\mathcal{S}$ that minimizes $\text{cost}(P, \mathcal{S}) := \sum_{p \in P} \text{dist}(p, \mathcal{S})^z$, with $z = 1$ for k -median and $z = 2$ for k -means. We say a solution is an α -approximation if its cost is at most α times the minimal cost.

► **Remark 4.** We consider in this paper the *discrete* Euclidean clustering. In the continuous version, centers can be placed arbitrarily in the space. A classical result shows how to reduce continuous to discrete, by computing a set of $1/\varepsilon^d \cdot |P|$ candidate centers that contain a $(1 + \varepsilon)$ -approximation [43].

For some set of points $\mathcal{S}$, and any point $p \in \mathbb{R}^d$, we define $\mathcal{S}(p) = \arg \min_{s \in \mathcal{S}} \text{dist}(p, s)$ be the closest point in $\mathcal{S}$ to p ; and we let $\mathcal{S}_p := \text{dist}(p, \mathcal{S}(p))$ be the distance from p to its closest point in $\mathcal{S}$.

We assume that the minimal distance between two points is 1, and denote $\text{diam}(P)$ the maximum distance between two points of P . We also assume that $P \subset B(0, \text{diam}(P))$. Using standard arguments for clustering problems, we can assume $\text{diam}(P) = \text{poly}(n)$ [14].

A hierarchical decomposition $\mathcal{D}$ is a sequence $\mathcal{B}_1, \mathcal{B}_2, \dots$ such that for any i , $\mathcal{B}_i$ is a partition of P , and $\mathcal{B}_i$ is a refinement of $\mathcal{B}_{i+1}$: namely, every part of $\mathcal{B}_i$ is fully contained in a part of $\mathcal{B}_{i+1}$ (note that in this definition, parts are increasing with i).

2.2 The Hierarchical Decomposition and Its Properties

► **Lemma 5.** *For any set of points $P \subset \mathbb{R}^d$ and any $\rho > 0$, there is a randomized hierarchical decomposition $\mathcal{D}$ such that the diameter of a part $B \in \mathcal{B}_i$ is at most 2^{i+1} , $|\mathcal{D}| \leq \lceil \log(\text{diam}(P)) \rceil$, each part $B \in \mathcal{B}_i$ is refined in at most $2^{O(d)}$ parts at level $i - 1$, and:*

1. *for any $p \in \mathbb{R}^d$, radius r , and level i , we have*

$$\Pr[\mathcal{D} \text{ cuts } B(p, r) \text{ at a level } i] \leq dr/2^i.$$

2. *for every part $B \in \mathcal{B}_i$ of the decomposition, there is a set of at most $1/\rho^{d-1}$ many portals such that for every points $p \in B, q \notin B$ there is a portal $x \in \mathcal{P}_B$ close to the line connecting p to q , i.e., $\|p - x\| + \|x - q\| \leq \|p - q\| + \rho 2^{i+1}$. Furthermore, any portal of level $i + 1$ that lies in B is also a portal of B .*

This decomposition can be found in time $(1/\rho)^{O(d)} n \log(\text{diam}(P))$.

In Euclidean space, this decomposition is precisely the quadtree dissection. The previous statement is an adaptation of Lemma 2.3 in [13] to Euclidean space. The improved Property 1 is shown, e.g., in Lemma 11.3 of [30]. Since $\text{diam}(P) = \text{poly}(n)$, the decomposition can be computed in near-linear time.

We denote $\tau(\varepsilon, d) = \log(d) + \log(1/\varepsilon)$, a parameter often used throughout this paper.

A *portal-respecting* path between two points p and q cut at level i consists of a sequence of segments connecting $p, p_1, \dots, p_i, q_i, \dots, q_1, q$, where p_j (resp. q_j) is a portal of the rectangle at level j containing p (resp. q). By extension, a solution to k -means is portal-respecting if all paths connecting clients to center are portal-respecting. Recall that $\mathcal{A}_p$ is the distance from p to the closest point in a solution $\mathcal{A}$. An important definition we will use is that of badly cut, slightly modified from [13]:

► **Definition 6.** Let $P \subset \mathbb{R}^d$, $\mathcal{D}$ be a hierarchical decomposition of P , and $\varepsilon > 0$.

For any ball $B(x, r)$, we define $\text{lvl}(x, r)$ the level at which the ball $B(x, r)$ is cut by $\mathcal{D}$.

We say a ball $B(x, r)$ is badly cut w.r.t. $\mathcal{D}$ if $\text{lvl}(x, r) \geq \log(3r) + \tau(\varepsilon, d)$. We say a point p is badly cut w.r.t. $\mathcal{D}$ and $\mathcal{A}$ if $B(p, 3\mathcal{A}_p)$ is badly cut w.r.t. $\mathcal{D}$ (namely, $\text{lvl}(p, 3\mathcal{A}_p) \geq \log(3\mathcal{A}_p) + \tau(\varepsilon, d)$).

In the following, we will consider badly cut points w.r.t. two different sets: $\mathcal{A}$ will be a constant factor approximation, known to the algorithm, and $\mathcal{S}$ a slight modification of the optimal solution, unknown to the algorithm, that we will specify later; those two solutions being independent of the decomposition $\mathcal{D}$. For input points in P , we will examine whether they are badly cut w.r.t. $\mathcal{D}$ and $\mathcal{A}$. For centers of $\mathcal{A}$, it will be w.r.t. $\mathcal{D}$ and $\mathcal{S}$. Furthermore, $\mathcal{D}$ will always be the same, hence we say for simplicity badly cut w.r.t. $\mathcal{A}$ (or w.r.t. $\mathcal{S}$).

The following lemma bounds the probability of being badly cut, and is a direct consequence of Property 1.

► **Lemma 7.** Let $P, S \subset \mathbb{R}^d$ and $\mathcal{D}$ a random hierarchical decomposition given by Lemma 5, with $\rho = O(\varepsilon/\log(1/\varepsilon))$. For any point $p \in \mathbb{R}^d$, the probability that p is badly cut w.r.t. S is at most ε .


Given a random decomposition, each point is assigned a *budget*. In the following definition, one can think of $\mathcal{A}$ as an $O(1)$ -approximation known by the algorithm, and $\mathcal{S}$ as the optimal solution. We show in Lemma 12 that this budget is small, and that one can compute a solution with cost at most the optimal cost plus the budget.

► **Definition 8.** Let $P \subset \mathbb{R}^d$, $\mathcal{D}$ be a hierarchical decomposition of P , $\varepsilon > 0$, $\mathcal{A}$ and $\mathcal{S}$ be two solutions to k -means on P .

For a ball $B(x, r)$, its *detour* with respect to B and $\mathcal{D}$ is $\text{det}_{\mathcal{D}}(x, r) = \varepsilon 2^{\text{lvl}(x, r)} r + \varepsilon^2 2^{2 \text{lvl}(x, r)}$.

A point $p \in P$ has *budget* with respect to $\mathcal{A}$ and $\mathcal{S}$ $b(p, \varepsilon) = b_1(p, \varepsilon) + b_2(p, \varepsilon) + b_3(p, \varepsilon)$, with

$$\begin{aligned} b_1(p, \varepsilon) &= \begin{cases} \text{det}_{\mathcal{D}}(p, 3\mathcal{A}_p) & \text{if } \text{lvl}(p, 3\mathcal{A}_p) \leq \log(3\mathcal{A}_p) + \tau(\varepsilon, d) \\ 0 & \text{otherwise} \end{cases} \\ b_2(p, \varepsilon) &= \begin{cases} \text{det}_{\mathcal{D}}(p, 3(\mathcal{A}_p + \mathcal{S}_p)) & \text{if } \text{lvl}(p, 3(\mathcal{A}_p + \mathcal{S}_p)) \leq \log(3(\mathcal{A}_p + \mathcal{S}_p)) + \tau(\varepsilon, d) \\ 36d\mathcal{A}_p^2 + 16d\mathcal{S}_p^2 & \text{otherwise} \end{cases} \\ b_3(p, \varepsilon) &= \begin{cases} \text{det}_{\mathcal{D}}(\mathcal{A}(p), 3\mathcal{S}_{\mathcal{A}(p)}) & \text{if } \text{lvl}(\mathcal{A}(p), 3\mathcal{S}_{\mathcal{A}(p)}) \leq \log(3\mathcal{S}_{\mathcal{A}(p)}) + \tau(\varepsilon, d) \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$



■ **Figure 1** Illustration of one case of the distinction. $B(p, 3(\mathcal{A}_p + \text{OPT}_p))$ is badly cut by the thick orange-dashed line, so p cannot be connected via portals to $\text{OPT}(p)$. However, p and $\mathcal{A}(p)$ are not badly cut, so p can be connected to $\text{OPT}(\mathcal{A}(p))$ instead, making a detour through the thin orange-dashed line. The cost of this reassignment is charged to $b_2(p)$.

3 Structure Theorem and the Approximation Algorithm

3.1 Roadmap

Now that all definitions are in place, we give a more precise high-level overview of the proof of Theorem 2. Our goal is to show that the decomposition of Lemma 5 with $\rho = \varepsilon / \log(1/\varepsilon)$ is precise enough for a dynamic program to compute a portal-respecting solution that is a $(1 + \varepsilon)$ -approximation.

To do this, we show that one can find a solution $\mathcal{S}^*$ such that: (1) it is a $(1 + \varepsilon)$ -approximation, and (2) the detour from each point to $\mathcal{S}^*$ is at most the budget of the point. To conclude, it is enough to bound the budget, which is a mere application of Lemma 5 and first-moment analysis. To find such a solution $\mathcal{S}^*$, we start from the optimal solution (which respects (1) but not (2)), and transform it such that each client can be connected through the portals.

First, consider the easy case where a client p is not badly cut, and $B(p, 3(\mathcal{A}_p + \text{OPT}_p))$ is not badly cut. Then, the detour from p to OPT_p can be charged to $b_2(p, \varepsilon)$, as the closest optimal center to p lies in the latter ball. When this ball is badly cut, instead of serving p by its closest center $\text{OPT}(p)$, we serve it by $\text{OPT}(\mathcal{A}(p))$. We show that this can always be charged to the budget of p . Indeed, in the sub-case where $\mathcal{A}(p)$ itself is not badly cut, we can show that the detour between p and $\text{OPT}(\mathcal{A}(p))$ is affordable compared to the distance between them; and this distance can be charged to the budget $b_2(p, \varepsilon)$. This sub-case is illustrated in Figure 1. However, in the sub-case where $\mathcal{A}(p)$ is badly cut as well, we have no choice but to add the center to OPT . This results in a solution with slightly too many centers – in expectation $(1 + \varepsilon)k$, since each center is badly cut with probability ε – and we will need to remove some centers of the optimal solution without increasing the cost too much. This step is already presented in [13], and we recall it in Section 3.2.

This essentially concludes the case where p is not badly cut. When it is badly cut, the decomposition $\mathcal{D}$ approximates its distances poorly; however, this happens with tiny probability, and we can afford to move p to $\mathcal{A}(p)$ (charging this movement to $b_1(p, \varepsilon)$). A case analysis, similar to the previous one, then shows that the detour between $\mathcal{A}(p)$ to OPT can be charged to $b_2(p, \varepsilon)$ and $b_3(p, \varepsilon)$.

In the remainder of this section, we show how to remove some centers from the optimal solution (to then add the badly cut centers of $\mathcal{A}$), then prove the structure theorem that shows the existence of a solution satisfying (1) and (2). We explain in Section 3.4 the dynamic program that relies on this structure theorem and enables the computation of a good approximation, to conclude the proof of Theorem 2.

3.2 Preparing the Instance

We consider the mapping of the centers of OPT to $\mathcal{A}$ defined as follows: for any $f \in \text{OPT}$, let $\mathcal{A}(f)$ denote the center of $\mathcal{A}$ that is the closest to f . Recall that for a client c , $\mathcal{A}(c)$ is the center serving c in $\mathcal{A}$.

For any center ℓ of $\mathcal{A}$, define $\psi(\ell)$ to be the set of centers of OPT that are mapped to ℓ , namely, $\psi(\ell) = \{f \in \text{OPT} \mid \mathcal{A}(f) = \ell\}$. Define $\mathcal{A}^1$ to be the set of centers ℓ of $\mathcal{A}$ for which there exists a unique $f \in \text{OPT}$ such that $\mathcal{A}(f) = \ell$, namely $\mathcal{A}^1 = \{\ell \in \mathcal{A} \mid |\psi(\ell)| = 1\}$. Let $\mathcal{A}^0 = \{\ell \in \mathcal{A} \mid |\psi(\ell)| = 0\}$, and $\mathcal{A}^{\geq 2} = \mathcal{A} \setminus (\mathcal{A}^1 \cup \mathcal{A}^0)$. Similarly, define $\text{OPT}^1 = \{f \in \text{OPT} \mid \mathcal{A}(f) \in \mathcal{A}^1\}$ and $\text{OPT}^{\geq 2} = \{f \in \text{OPT} \mid \mathcal{A}(f) \in \mathcal{A}^{\geq 2}\}$. Note that $|\text{OPT}^{\geq 2}| = |\mathcal{A}^0| + |\mathcal{A}^{\geq 2}|$, since $|\text{OPT}^1| = |\mathcal{A}^1|$ and, w.l.o.g., $|\text{OPT}| = |\mathcal{A}| = k$.

As shown in [13], one can remove some of the centers from $\text{OPT}^{\geq 2}$ to obtain the following result:

► **Lemma 9** (See Step 1. and Lemma 4.1 in [13]). *There exists a set $\mathcal{H} \subseteq \text{OPT}^{\geq 2}$ such that:*

- $|\mathcal{H}| = \lfloor \varepsilon \cdot |\text{OPT}^{\geq 2}|/2 \rfloor$,
- for any $\ell \in \mathcal{A}$, the center in $\psi(\ell)$ closest to ℓ is not in $\mathcal{H}$,
- let $\text{OPT}' = \text{OPT} \setminus \mathcal{H}$: OPT' has cost at most $(1 + O(\varepsilon))\text{cost}(\text{OPT}) + O(\varepsilon)\text{cost}(\mathcal{A})$.

Furthermore, we transform the input set P to give it more structure, as follows. Given a randomized hierarchical decomposition $\mathcal{D}$ and a solution $\mathcal{A}$, we define the set $\tilde{P}$ as follows. For every point p ,

$$\tilde{p} = \begin{cases} \mathcal{A}(p) & \text{if } p \text{ is badly cut w.r.t. } \mathcal{D} \text{ and } \mathcal{A} \\ p & \text{otherwise.} \end{cases}$$

We note that $\tilde{P} := \{\tilde{p} : p \in P\}$ is a random variable that depends on the randomness of $\mathcal{D}$.

3.3 Construction of a Structured Solution

Using Lemma 9, we can now show the existence of a near-optimal portal-respecting solution. Let $\mathcal{A}$ be a fixed solution, and OPT' the solution from Lemma 9 for P . Let $\mathcal{D}$ be a randomized hierarchical decomposition, and $B_{\mathcal{D}}$ be the set of centers of $\mathcal{A}$ that are badly cut w.r.t. OPT' , i.e., $c \in B_{\mathcal{D}}$ when the ball $B(c, 3\text{OPT}'_c)$ is cut at some level greater than $\log(3\text{OPT}'_c) + \tau(\varepsilon, d)$.

► **Definition 10.** For $\varepsilon > 0$, we say that $\mathcal{D}$ has ε -small distortion if

1. the total budget w.r.t. $\mathcal{A}$ and OPT' is bounded: $\sum_{p \in P} b(p, \varepsilon) \leq \varepsilon(\text{cost}(P, \text{OPT}') + \text{cost}(P, \mathcal{A}))$.
2. there exists a solution $\mathcal{S}^*$ that contains $B_{\mathcal{D}}$ with $\sum_{p \in P} (\text{dist}(p, \tilde{p}) + \text{dist}(\tilde{p}, \mathcal{S}^*))^2 \leq (1 + \varepsilon)\text{cost}(P, \text{OPT}') + \varepsilon\text{cost}(P, \mathcal{A})$,
3. for each point $p \in P$, its detour to $\mathcal{S}^*$ is at most $b(p, \varepsilon)$, namely there is a center $s \in \mathcal{S}^*$ cut from $\tilde{p}$ at level i such that $(\text{dist}(p, \tilde{p}) + \text{dist}(\tilde{p}, s) + \varepsilon 2^i)^2 \leq (\text{dist}(p, \tilde{p}) + \text{dist}(\tilde{p}, \mathcal{S}^*))^2 + b(p, \varepsilon)$.

As explained in the introduction, we will use those properties as follows: a standard dynamic program will compute the best *portal-respecting* solution that minimizes $\widetilde{\text{cost}}(P, \mathcal{S}) := \sum_p (\text{dist}(p, \widehat{p}) + \text{dist}(\widehat{p}, \mathcal{S}))^2$. Making the solution portal-respecting incurs an error $\varepsilon 2^i$ for two points cut at level i : Property 3 therefore bounds this error by $b(p, \varepsilon)$. Property 1 of Definition 10 ensures that the total budget is bounded, and Property 2 relates $\widetilde{\text{cost}}$ to cost .

Our main structural theorem for k -means in low dimension is the following:

► **Theorem 11** (Structure Theorem). *For any $\varepsilon > 0$, $\mathcal{D}$ has ε -small-distortion with probability at least $2/3$.*

The rest of the paper is dedicated to the proof of this theorem.

3.3.1 Bounding the Budget

► **Lemma 12.** *Fix two solutions $\mathcal{A}$ and $\mathcal{S}$, and let $b(p, \varepsilon)$ be the budget of p w.r.t. solutions $\mathcal{A}$ and $\mathcal{S}$. With probability $2/3$ (over the randomness of $\mathcal{D}$), $\sum_p b(p, \varepsilon) = O(d^2 \log(1/\varepsilon) \cdot \varepsilon)(\text{cost}(P, \mathcal{A}) + \text{cost}(P, \mathcal{S}))$*

Proof. We bound the expectation of the budget and then use Markov's inequality. Recall that any ball $B(x, r)$ is cut at level i with probability $dr/2^i$, and is badly cut with probability ε . Furthermore, when $B(x, r)$ is cut at level i , then $\det_{\mathcal{D}}(x, r) = \varepsilon r 2^i + \varepsilon^2 2^{2i}$.

To simplify the analysis slightly, we first analyze the expectation of the variable

$$\beta(x, r) := \begin{cases} \det_{\mathcal{D}}(x, r) & \text{when } \text{lvl}(x, r) \leq \log(r) + \tau(\varepsilon, d) \\ 0 & \text{otherwise.} \end{cases}$$

To do so, we distinguish between the cases where $\text{lvl}(x, r)$ is at most $\log(r)$, between $\log(r)$ and $\log(r) + \tau(\varepsilon, d)$, and at least $\log(r) + \tau(\varepsilon, d)$. This gives the following upper bound:

$$\begin{aligned} \mathbb{E}[\beta(x, r)] &\leq \sum_{i \leq \log(r)} \det_{\mathcal{D}}(x, r) + \sum_{i=\log(r)}^{\log(r)+\tau(\varepsilon, d)} \det_{\mathcal{D}}(x, r) \cdot \Pr[\mathcal{D} \text{ cuts } B(p, r) \text{ at a level } i] \\ &\leq \sum_{i \leq \log(r)} (\varepsilon 2^i r + \varepsilon^2 2^{2i}) + \sum_{i=\log(r)}^{\log(r)+\tau(\varepsilon, d)} (\varepsilon 2^i r + \varepsilon^2 2^{2i}) \cdot \frac{dr}{2^i} \\ &\leq 3\varepsilon r^2 + \sum_{i=\log(r)}^{\log(r)+\tau(\varepsilon, d)} (\varepsilon dr^2 + \varepsilon^2 2^i \cdot dr) \\ &\leq 3\varepsilon r^2 + \tau(\varepsilon, d) \cdot \varepsilon dr^2 + \varepsilon^2 \cdot dr \cdot 2^{\log(r)+\tau(\varepsilon, d)} \\ &= O(d^2 \log(1/\varepsilon)) \varepsilon r^2, \end{aligned}$$

where the last line uses $2^{\tau(\varepsilon, d)} = 2^{\log(d)+\log(1/\varepsilon)} = O(d/\varepsilon)$.

Since $b_1(p, \varepsilon) = \beta(p, 3\mathcal{A}_p)$ and $b_3(p, \varepsilon) = \beta(\mathcal{A}(p), 3\mathcal{S}_{\mathcal{A}(p)})$, we directly get $\mathbb{E}[b_1(p, \varepsilon) + b_3(p, \varepsilon)] \leq O(d^2 \log(1/\varepsilon) \varepsilon)(\mathcal{A}_p^2 + \mathcal{S}_p^2)$. Bounding $b_2(p, \varepsilon)$ is only slightly more complicated: $b_2(p, \varepsilon) = \beta(p, \mathcal{A}_p + \mathcal{S}_p) + X(p, \varepsilon)$, where $X(p, \varepsilon) := 36d\mathcal{A}_p^2 + 16d\mathcal{S}_p^2$ when $\text{lvl}(p, 3(\mathcal{A}_p + \mathcal{S}_p)) > \log(3(\mathcal{A}_p + \mathcal{S}_p)) + \tau(\varepsilon, d)$. Therefore, it is enough to compute $\mathbb{E}[X(p, \varepsilon)]$.

Properties of the decomposition in Lemma 5 ensure that $\text{lvl}(p, 3(\mathcal{A}_p + \mathcal{S}_p)) > \log(3(\mathcal{A}_p + \mathcal{S}_p)) + \tau(\varepsilon, d)$ with probability at most ε , and therefore $\mathbb{E}[X(p, \varepsilon)] \leq \varepsilon \cdot (36d\mathcal{A}_p^2 + 16d\mathcal{S}_p^2) = O(\varepsilon d(\mathcal{A}_p^2 + \mathcal{S}_p^2))$.

Therefore, $\mathbb{E}[b_2(p, \varepsilon)] = \mathbb{E}[\beta(p, 3(\mathcal{A}_p + \mathcal{S}_p))] + \mathbb{E}[X(p, \varepsilon)] \leq O(d^2 \log(1/\varepsilon) \cdot \varepsilon)(\mathcal{A}_p^2 + \mathcal{S}_p^2)$.

Summing these inequalities for b_1, b_2, b_3 , and for all $p \in P$ we get $\mathbb{E}[\sum_{p \in P} b(p, \varepsilon)] \leq O(d^2 \log(1/\varepsilon) \cdot \varepsilon)(\text{cost}(P, \mathcal{A}) + \text{cost}(P, \mathcal{S}))$. Markov's inequality concludes the lemma. ◀

3.3.2 Construction of $\mathcal{S}^*$

To construct $\mathcal{S}^*$, we start from the solution OPT' obtained from Lemma 9 and apply the two following steps.

- For any $\ell \in \mathcal{A}$, define f_ℓ to be its closest center in $\psi(\ell)$, breaking ties arbitrarily. We start with $\mathcal{S}^* = \text{OPT}'$ obtained from Lemma 9. Note that for every $\ell \in \mathcal{A}^1 \cup \mathcal{A}^{\geq 2}$, the second property of Lemma 9 ensures that f_ℓ is a center of OPT' . $\mathcal{S}^*$ is then transformed as follows:
- **Step 1.** For each center $\ell \in B_{\mathcal{D}} \setminus \mathcal{A}^0$ (i.e., ℓ is badly cut w.r.t. $\mathcal{D}$ and OPT' , and $\psi(\ell) \neq \emptyset$), replace f_ℓ by ℓ in $\mathcal{S}^*$.
 - **Step 2.** Add all centers of $\mathcal{A}^0$ badly cut w.r.t. $\mathcal{D}$ and OPT' to $\mathcal{S}^*$.

The two key properties of that procedure are the following:

► **Fact 13** (Claim 4.3 and 4.4 in [13]). *With probability $4/5$, $|\mathcal{S}^*| \leq k$, so $\mathcal{S}^*$ is a valid solution, and $\text{cost}(P, \mathcal{S}^*) \leq (1 + \varepsilon)\text{cost}(P, \text{OPT}) + \varepsilon\text{cost}(P, \mathcal{A})$. Furthermore, for any point p , it holds that $\text{dist}(p, \mathcal{S}^*) \leq 2 \text{dist}(p, \text{OPT}') + \text{dist}(p, \mathcal{A})$.*

In addition, the construction of $\mathcal{S}^*$ ensures the below whose proof is in the full version:

► **Fact 14.** *For any point p ,*

$$\begin{aligned} \text{dist}(p, \tilde{p}) + \text{dist}(\tilde{p}, \mathcal{S}^*) &\leq 3\mathcal{A}_p + 2\text{OPT}'_p \\ \text{dist}(p, \mathcal{S}^*(\mathcal{A}(p))) &\leq 3\mathcal{A}_p + 2\text{OPT}'_p. \end{aligned}$$

Furthermore, for any center $\ell \in \mathcal{A}$, it holds that $\text{dist}(\ell, \mathcal{S}^) \leq 2 \text{dist}(\ell, \text{OPT}')$.*

We can now use this solution $\mathcal{S}^*$ to show $\mathcal{D}$ has ε -small distortion.

3.3.3 Bounding the cost of $\mathcal{S}^*$

► **Lemma 15.** *$\mathcal{S}^*$ contains $B_{\mathcal{D}}$, and with probability at least $4/5$ it holds that*

$$\sum_{p \in P} (\text{dist}(p, \tilde{p}) + \text{dist}(\tilde{p}, \mathcal{S}^*))^2 \leq (1 + \varepsilon)\text{cost}(P, \text{OPT}) + \varepsilon\text{cost}(P, \mathcal{A})$$

Proof. By construction, $\mathcal{S}^*$ does contain $B_{\mathcal{D}}$. To simplify the equations, we write p b.c. when p is badly cut w.r.t. $\mathcal{A}$ (in which case $\tilde{p} = \mathcal{A}(p)$). It holds that

$$\sum_{p \in P} (\text{dist}(p, \tilde{p}) + \text{dist}(\tilde{p}, \mathcal{S}^*))^2 \leq \sum_{p \in P} \text{dist}(p, \mathcal{S}^*)^2 + \sum_{p \text{ b.c.}} (\text{dist}(p, \tilde{p}) + \text{dist}(\tilde{p}, \mathcal{S}^*))^2. \quad (1)$$

By Fact 13, the first term is at most $(1 + \varepsilon)\text{cost}(P, \text{OPT}) + \varepsilon\text{cost}(P, \mathcal{A})$: thus, we focus only on the second term. For this, we use Fact 14:

$$\begin{aligned} \sum_{p \text{ b.c.}} (\text{dist}(p, \tilde{p}) + \text{dist}(\tilde{p}, \mathcal{S}^*))^2 &\leq \sum_{p \text{ b.c.}} (3\mathcal{A}_p + 2\text{OPT}'_p)^2 \\ &\leq 18 \sum_{p \text{ b.c.}} \mathcal{A}_p^2 + \text{OPT}'_p^2 \end{aligned}$$

Each point p is badly cut with probability ε : therefore, in expectation, it holds that

$$\begin{aligned} \mathbb{E}[18 \sum_{p \text{ b.c.}} \mathcal{A}_p^2 + \text{OPT}'_p^2] &\leq 18\varepsilon \sum_{p \in P} \mathcal{A}_p^2 + \text{OPT}'_p^2 \\ &\leq 18\varepsilon(\text{cost}(P, \mathcal{A}) + \text{cost}(P, \text{OPT}')) \end{aligned}$$

Using Markov's inequality, it holds with probability $4/5$ that

$$\sum_{p \text{ b.c.}} (\text{dist}(p, \tilde{p}) + \text{dist}(\tilde{p}, \text{OPT}'))^2 \leq 18 \cdot 5 \cdot \varepsilon (\text{cost}(P, \mathcal{A}) + \text{cost}(P, \text{OPT}')).$$

Plugging this bound into Equation (1) and using Lemma 9 to bound the cost of OPT' , we conclude:

$$\begin{aligned} & \sum_{p \in P} (\text{dist}(p, \tilde{p}) + \text{dist}(\tilde{p}, \mathcal{S}^*))^2 \\ & \leq (1 + \varepsilon) \text{cost}(P, \text{OPT}) + \varepsilon \text{cost}(P, \mathcal{A}) + 18 \cdot 5 \cdot \varepsilon (\text{cost}(P, \mathcal{A}) + \text{cost}(P, \text{OPT}')) \\ & = (1 + O(\varepsilon)) \text{cost}(P, \text{OPT}) + O(\varepsilon) \text{cost}(P, \mathcal{A}). \end{aligned} \quad \blacktriangleleft$$

3.3.4 Bounding the detour

We now show that the detour incurred by connecting $\tilde{P}$ to $\mathcal{S}^*$ through portals of the decomposition $\mathcal{D}$ is within the budget defined in Definition 8.

► **Lemma 16.** *Let $P \subset \mathbb{R}^d$, $\mathcal{A}$ be a solution for k -means on P and $\mathcal{D}$ be a randomized decomposition. Let $\mathcal{S}^*$ be the solution from Lemma 15.*

Then, for any $p \in P$, there is a center $s \in \mathcal{S}^$ cut from $\tilde{p}$ at some level i such that $(\text{dist}(p, \tilde{p}) + \text{dist}(\tilde{p}, s) + \varepsilon 2^i)^2 \leq (\text{dist}(p, \tilde{p}) + \text{dist}(\tilde{p}, \mathcal{S}^*))^2 + b(p, \varepsilon)$.*

Proof. We recall that any point $x \in \mathbb{R}^d$ is badly cut w.r.t. $\mathcal{D}$ and a set S when the ball $B(x, 3 \text{dist}(x, S))$ is cut at level higher than $\log(9 \text{dist}(x, S)) + \tau(\varepsilon, d)$. We recall our construction of $\tilde{P}$ and $\mathcal{S}^*$: When a point p is badly cut w.r.t. $\mathcal{D}$ and $\mathcal{A}$, $\tilde{p} = \mathcal{A}(p)$; when a center $\ell \in \mathcal{A}$ is badly cut w.r.t. $\mathcal{D}$ and OPT' , $\ell \in \mathcal{S}^*$.

To show the lemma, we make a case distinction, first according to whether p is badly cut or not w.r.t. $\mathcal{D}$ and $\mathcal{A}$. In the case where it is badly cut, it holds that $\tilde{p} = \mathcal{A}(p)$. Then, we have the following cases.

1. If $\mathcal{A}(p)$ is badly cut w.r.t. $\mathcal{D}$ and OPT' , then $\mathcal{A}(p) \in \mathcal{S}^*$: therefore, $\tilde{p} = \mathcal{A}(p) = \mathcal{S}^*(p)$, and $\tilde{p}$ and $\mathcal{S}^*(p)$ are not cut, e.g. $i = -\infty$. Therefore, $s = \mathcal{A}(p)$ satisfies the lemma, with $(\text{dist}(p, \tilde{p}) + \text{dist}(\tilde{p}, s) + \varepsilon 2^i)^2 = \text{dist}(p, \tilde{p})^2$.
2. $\mathcal{A}(p)$ is not badly cut w.r.t. $\mathcal{D}$ and OPT' . Let $s = \mathcal{S}^*(\mathcal{A}(p))$, and i be the level at which $\mathcal{A}(p)$ and s are cut. We write $(\text{dist}(p, \tilde{p}) + \text{dist}(\tilde{p}, s) + \varepsilon 2^i)^2 = (\mathcal{A}_p + \text{dist}(\tilde{p}, \mathcal{S}^*))^2 + 2\varepsilon 2^i \cdot \mathcal{A}_p + 2\varepsilon 2^i \cdot \mathcal{S}_{\mathcal{A}(p)}^* + \varepsilon^2 2^{2i}$.

By Fact 14, $s \in B(\mathcal{A}(p), 2\text{OPT}'_{\mathcal{A}(p)})$, and so $i \leq i' := \text{lvl}(\mathcal{A}(p), 3\text{OPT}'_{\mathcal{A}(p)})$. Combined with the fact that $\mathcal{A}(p)$ is not badly cut, it therefore holds that $2\varepsilon 2^i \cdot \mathcal{S}_{\mathcal{A}(p)}^* + \varepsilon^2 2^{2i} \leq 2\varepsilon 2^{i'} \cdot 2\text{OPT}'_{\mathcal{A}(p)} + \varepsilon^2 2^{2i'} \leq b_3(p)$. Thus, it is enough to bound $2\varepsilon 2^i \cdot \mathcal{A}_p$. For this, we refine our case distinction, according to whether the ball $B(p, 3(\mathcal{A}_p + \text{OPT}'_p))$ is badly cut or not, and use the budget from $b_2(p)$.

- a. If the ball is not badly cut, then we remark that Fact 14 ensures that both s and $\mathcal{A}(p)$ are in $B(p, 3(\mathcal{A}_p + \text{OPT}'_p))$, and therefore $i \leq i'' := \text{lvl}(p, 3(\mathcal{A}_p + \text{OPT}'_p))$. Therefore, $2\varepsilon 2^i \cdot \mathcal{A}_p \leq \varepsilon 2^{i''} 3(\mathcal{A}_p + \text{OPT}'_p) \leq b_2(p)$.
- b. If this ball is badly cut, then we use $i \leq i' \leq \log(3\text{OPT}'_{\mathcal{A}(p)}) + \tau(\varepsilon, d)$ in addition to $\text{OPT}'_{\mathcal{A}(p)} \leq \mathcal{A}_p + \text{OPT}'_p$ as follows: $2\varepsilon 2^i \cdot \mathcal{A}_p \leq 6\varepsilon \cdot \text{OPT}'_{\mathcal{A}(p)} \cdot 2^{\tau(\varepsilon, d)} \mathcal{A}_p \leq 6d(\mathcal{A}_p + \text{OPT}'_p) \mathcal{A}_p \leq b_2(p)$.

Combining these cases, we get that when $\mathcal{A}(p)$ is not badly cut, $(\text{dist}(p, \tilde{p}) + \text{dist}(\tilde{p}, s) + \varepsilon 2^i)^2 \leq (\mathcal{A}_p + \text{dist}(\tilde{p}, \mathcal{S}^*))^2 + b_2(p) + b_3(p)$.

We now turn to the case where p is not badly cut. In that case, $\tilde{p} = p$: we therefore only need to bound $(\text{dist}(p, \mathcal{S}^*) + \varepsilon 2^i)^2$ and the details of this are presented in the full version. ◀

3.4 The Algorithm: how to use the Structure Theorem

Suppose that $\mathcal{D}$ has ρ -small-distortion, for $\rho = O(\varepsilon / \log(1/\varepsilon))$ (where the O hides dependencies in d), and let $\mathcal{S}^*$ be the solution provided by Definition 10. As explained previously, our algorithm will be a standard dynamic program working as follows. The decomposition $\mathcal{D}$ will be augmented with a set of portals, as defined in Lemma 5, with $\rho = \varepsilon / \log(1/\varepsilon)$. The dynamic program (described later in greater detail) computes the best portal-respecting solution, namely, the solution where each path connecting a client to a center crosses boundaries of $\mathcal{D}$ only at portals. More precisely, it computes the portal-respecting solution with (almost) smallest $\widetilde{\text{cost}}(P, \mathcal{S}) := \sum_p (\text{dist}(p, \tilde{p}) + \text{dist}(\tilde{p}, \mathcal{S}))^2$

► **Theorem 17.** *There exists an algorithm running in time $(2/\varepsilon)^{O(\log(1/\varepsilon)/\varepsilon)^{d-1}} \cdot n \log^3 n$ that computes a portal-respecting solution $\mathcal{S}$ such that $\widetilde{\text{cost}}(P, \mathcal{S}) \leq (1 + \varepsilon) \min_{\mathcal{S}'} \widetilde{\text{cost}}(P, \mathcal{S}')$.*

The preciseness and nestedness of portal sets ensure that, when two points x and y are cut at level i , then the portal-respecting path between x and y has length $\|x - y\| + \rho 2^{i+1}$. Furthermore, the algorithm from Theorem 17 computes a solution at least as good as $\mathcal{S}^*$: by Property 3 of Definition 10, the portal-respecting $\widetilde{\text{cost}}$ of $\mathcal{S}^*$ is at most $\sum_p (\text{dist}(p, \tilde{p}) + \text{dist}(\tilde{p}, \mathcal{S}^*))^2 + b(p, \rho)$. Property 1 ensures that the total budget is at most $\rho(\text{cost}(P, \text{OPT}) + \text{cost}(P, \mathcal{A})) = O(\rho)\text{cost}(P, \text{OPT})$, since $\mathcal{A}$ is an $O(1)$ -approximation. Combined with Property 2 of Definition 10, this shows that the portal-respecting $\widetilde{\text{cost}}$ of the solution computed by the algorithm is at most $(1 + O(\varepsilon))\text{cost}(P, \text{OPT})$.

To wrap up, we note that the portal respecting $\widetilde{\text{cost}}$ is an upper bound on cost : going in straight lines instead of going through portals only improves the cost, and for any point p , $\text{dist}(p, \mathcal{S}) \leq \text{dist}(p, \tilde{p}) + \text{dist}(\tilde{p}, \mathcal{S})$. Therefore, the solution computed by the dynamic program is a $(1 + O(\varepsilon))$ -approximation for k -means.

The algorithm used to show Theorem 17 is a quite standard dynamic program (DP), described in detail in [13] and is deferred to the full version.

References

- 1 Fateme Abbasi, Sandip Banerjee, Jaroslav Byrka, Parinya Chalermsook, Ameet Gadekar, Kamyar Khodamoradi, Dániel Marx, Roohani Sharma, and Joachim Spoerhase. Parametrized approximation schemes for clustering with general norm objectives. In *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023, Santa Cruz, CA, USA, November 6-9, 2023*, pages 1377–1399. IEEE, 2023. doi:10.1109/FOCS57990.2023.00085.
- 2 Enver Aman, Karthik C. S., and Sharath Punna. On connections between k -coloring and Euclidean k -means. In Timothy M. Chan, Johannes Fischer, John Iacono, and Grzegorz Herman, editors, *32nd Annual European Symposium on Algorithms, ESA 2024, September 2-4, 2024, Royal Holloway, London, United Kingdom*, volume 308 of *LIPIcs*, pages 9:1–9:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.ESA.2024.9.
- 3 Sanjeev Arora. Polynomial time approximation schemes for Euclidean TSP and other geometric problems. In *37th Annual Symposium on Foundations of Computer Science, FOCS '96, Burlington, Vermont, USA, 14-16 October, 1996*, pages 2–11. IEEE Computer Society, 1996. doi:10.1109/SFCS.1996.548458.
- 4 Sanjeev Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *J. ACM*, 45(5):753–782, 1998. doi:10.1145/290179.290180.

- 5 Sanjeev Arora, Prabhakar Raghavan, and Satish Rao. Approximation schemes for Euclidean k -medians and related problems. In Jeffrey Scott Vitter, editor, *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 106–113. ACM, 1998. doi:10.1145/276698.276718.
- 6 Pranjal Awasthi, Moses Charikar, Ravishankar Krishnaswamy, and Ali Kemal Sinop. The hardness of approximation of Euclidean k -means. In Lars Arge and János Pach, editors, *31st International Symposium on Computational Geometry, SoCG 2015, June 22-25, 2015, Eindhoven, The Netherlands*, volume 34 of *LIPICs*, pages 754–767. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPICs.SOCG.2015.754.
- 7 Yair Bartal and Lee-Ad Gottlieb. A linear time approximation scheme for Euclidean TSP. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 698–706. IEEE Computer Society, 2013. doi:10.1109/FOCS.2013.80.
- 8 Vladimir Braverman, Gereon Frahling, Harry Lang, Christian Sohler, and Lin F. Yang. Clustering high dimensional dynamic data streams. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 576–585. PMLR, 2017. URL: <http://proceedings.mlr.press/v70/braverman17a.html>.
- 9 Moses Charikar, Vincent Cohen-Addad, Ruiquan Gao, Fabrizio Grandoni, Euiwoong Lee, and Ernest van Wijland. An improved greedy approximation for (metric) k -means. In *66th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2025, Sydney, Australia, December 14-17, 2025*, pages 233–240. IEEE, 2025. doi:10.1109/FOCS63196.2025.00016.
- 10 Vincent Cohen-Addad. A fast approximation scheme for low-dimensional k -means. In Artur Czumaj, editor, *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 430–440. SIAM, 2018. doi:10.1137/1.9781611975031.29.
- 11 Vincent Cohen-Addad, Arnaud de Mesmay, Eva Rotenberg, and Alan Roytman. The bane of low-dimensionality clustering. In *Proceedings of the 2018 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2018. doi:10.1137/1.9781611975031.30.
- 12 Vincent Cohen-Addad, Alessandro Epasto, Silvio Lattanzi, Vahab Mirrokni, Andres Muñoz Medina, David Saulpic, Chris Schwiegelshohn, and Sergei Vassilvitskii. Scalable differentially private clustering via hierarchically separated trees. In Aidong Zhang and Huzefa Rangwala, editors, *KDD '22: The 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, August 14 - 18, 2022*, pages 221–230. ACM, 2022. doi:10.1145/3534678.3539409.
- 13 Vincent Cohen-Addad, Andreas Emil Feldmann, and David Saulpic. Near-linear time approximation schemes for clustering in doubling metrics. *J. ACM*, 68(6):44:1–44:34, 2021. doi:10.1145/3477541.
- 14 Vincent Cohen-Addad, Andreas Emil Feldmann, and David Saulpic. Near-linear time approximation schemes for clustering in doubling metrics. In *J. ACM*, volume 68, 2021. doi:10.1145/3477541.
- 15 Vincent Cohen-Addad, Fabrizio Grandoni, Euiwoong Lee, Chris Schwiegelshohn, and Ola Svensson. A $(2+\epsilon)$ -approximation algorithm for metric k -median. In Michal Koucký and Nikhil Bansal, editors, *Proceedings of the 57th Annual ACM Symposium on Theory of Computing, STOC 2025, Prague, Czechia, June 23-27, 2025*, pages 615–624. ACM, 2025. doi:10.1145/3717823.3718299.
- 16 Vincent Cohen-Addad, Anupam Gupta, Amit Kumar, Euiwoong Lee, and Jason Li. Tight FPT approximations for k -median and k -means. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece*, volume 132 of *LIPICs*, pages 42:1–42:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.ICALP.2019.42.

- 17 Vincent Cohen-Addad and Karthik C. S. Inapproximability of clustering in l_p metrics. In David Zuckerman, editor, *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 519–539. IEEE Computer Society, 2019. doi:10.1109/FOCS.2019.00040.
- 18 Vincent Cohen-Addad, Karthik C. S., and Euiwoong Lee. On approximability of clustering problems without candidate centers. In Dániel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 2635–2648. SIAM, 2021. doi:10.1137/1.9781611976465.156.
- 19 Vincent Cohen-Addad, Karthik C. S., and Euiwoong Lee. Johnson coverage hypothesis: Inapproximability of k -means and k -median in l_p -metrics. In Joseph (Seffi) Naor and Niv Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022*, pages 1493–1530. SIAM, 2022. doi:10.1137/1.9781611977073.63.
- 20 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. *SIAM J. Comput.*, 48(2):644–667, 2019. doi:10.1137/17M112717X.
- 21 Vincent Cohen-Addad, Silvio Lattanzi, Ashkan Norouzi-Fard, Christian Sohler, and Ola Svensson. Fast and accurate k -means++ via rejection sampling. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL: <https://proceedings.neurips.cc/paper/2020/hash/babcff88f8be8c4795bd6f0f8cccca61-Abstract.html>.
- 22 Vincent Cohen-Addad, David Saulpic, and Chris Schwiegelshohn. A new coresets framework for clustering. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 169–182. ACM, 2021. doi:10.1145/3406325.3451022.
- 23 Sanjoy Dasgupta and Yoav Freund. Random projection trees for vector quantization. *IEEE Transactions on Information Theory*, 55(7):3229–3242, 2009. doi:10.1109/TIT.2009.2021326.
- 24 Mark de Berg, Hans L. Bodlaender, Sándor Kisfaludi-Bak, Dániel Marx, and Tom C. van der Zanden. A framework for exponential-time-hypothesis-tight algorithms and lower bounds in geometric intersection graphs. *SIAM J. Comput.*, 49(6):1291–1331, 2020. doi:10.1137/20M1320870.
- 25 Irit Dinur. Mildly exponential reduction from gap 3sat to polynomial-gap label-cover. *Electron. Colloquium Comput. Complex.*, TR16-128, 2016. URL: <https://eccc.weizmann.ac.il/report/2016/128>.
- 26 Fedor V. Fomin, Petr A. Golovach, and Kirill Simonov. Parameterized k -clustering: Tractability island. *J. Comput. Syst. Sci.*, 117:50–74, 2021. doi:10.1016/J.JCSS.2020.10.005.
- 27 Zachary Friggstad, Mohsen Rezapour, and Mohammad R. Salavatipour. Local search yields a PTAS for k -means in doubling metrics. *SIAM J. Comput.*, 48(2):452–480, 2019. doi:10.1137/17M1127181.
- 28 Sudipto Guha and Samir Khuller. Greedy strikes back: Improved facility location algorithms. *J. Algorithms*, 31(1):228–248, 1999. doi:10.1006/jagm.1998.0993.
- 29 S Louis Hakimi. Optimum distribution of switching centers in a communication network and some related graph theoretic problems. *Operations research*, 13(3):462–475, 1965.
- 30 Sarel Har-Peled. *Geometric approximation algorithms*. Number 173 in Mathematical Surveys and Monographs. American Mathematical Soc., 2011.
- 31 Mary Inaba, Naoki Katoh, and Hiroshi Imai. Applications of weighted voronoi diagrams and randomization to variance-based k -clustering (extended abstract). In *Proceedings of the Tenth Annual Symposium on Computational Geometry, Stony Brook, New York, USA, June 6-8, 1994*, pages 332–339, 1994. doi:10.1145/177424.178042.

- 32 Kamal Jain, Mohammad Mahdian, and Amin Saberi. A new greedy approach for facility location problems. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 731–740, 2002. doi:10.1145/509907.510012.
- 33 Kamal Jain and Vijay 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. doi:10.1145/375827.375845.
- 34 Sándor Kisfaludi-Bak, Jesper Nederlof, and Karol Wegrzycki. A gap-ETH-tight approximation scheme for Euclidean TSP. *Journal of the ACM*, 72(6):1–48, 2025.
- 35 Stavros G. Kolliopoulos and Satish Rao. A nearly linear-time approximation scheme for the Euclidean k -median problem. *SIAM J. Comput.*, 37(3):757–782, 2007. doi:10.1137/S0097539702404055.
- 36 Alfred A Kuehn and Michael J Hamburger. A heuristic program for locating warehouses. *Management science*, 9(4):643–666, 1963.
- 37 Euiwoong Lee, Melanie Schmidt, and John Wright. Improved and simplified inapproximability for k -means. *Inf. Process. Lett.*, 120:40–43, 2017. doi:10.1016/j.ipl.2016.11.009.
- 38 Shi Li and Ola Svensson. Approximating k -median via pseudo-approximation. *SIAM J. Comput.*, 45(2):530–547, 2016. doi:10.1137/130938645.
- 39 Stuart Lloyd. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137, 1982. doi:10.1109/TIT.1982.1056489.
- 40 Meena Mahajan, Prajakta Nimbhorkar, and Kasturi R. Varadarajan. The planar k -means problem is np-hard. *Theor. Comput. Sci.*, 442:13–21, 2012. doi:10.1016/J.TCS.2010.05.034.
- 41 Konstantin Makarychev, Yury Makarychev, and Ilya P. Razenshteyn. Performance of johnson-lindenstrauss transform for k -means and k -medians clustering. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 1027–1038, 2019. doi:10.1145/3313276.3316350.
- 42 Pasin Manurangsi and Prasad Raghavendra. A birthday repetition theorem and complexity of approximating dense csps. In Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl, editors, *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland*, volume 80 of *LIPIcs*, pages 78:1–78:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPIcs.ICALP.2017.78.
- 43 Jirí Matousek. On approximate geometric k -clustering. *Discret. Comput. Geom.*, 24(1):61–84, 2000. doi:10.1007/s004540010019.
- 44 Nimrod Megiddo and Kenneth J Supowit. On the complexity of some common geometric location problems. *SIAM journal on computing*, 13(1):182–196, 1984. doi:10.1137/0213014.
- 45 Satish B Rao and Warren D Smith. Approximating geometrical graphs via “spanners” and “banyans”. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 540–550, 1998.
- 46 John Fred Stollsteimer. *The effect of technical change and output expansion on the optimum number, size, and location of pear marketing facilities in a California pear producing region*. University of California, Berkeley, 1961.