



**HAL**  
open science

# Near-linear time approximations schemes for clustering in doubling metrics

Vincent Cohen-Addad, Andreas Emil Feldmann, David Saulpic

► **To cite this version:**

Vincent Cohen-Addad, Andreas Emil Feldmann, David Saulpic. Near-linear time approximations schemes for clustering in doubling metrics. FOCS'19, Nov 2019, Baltimore, United States. hal-02360768

**HAL Id: hal-02360768**

**<https://hal.science/hal-02360768v1>**

Submitted on 13 Nov 2019

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Near-linear time approximations schemes for clustering in doubling metrics

Vincent Cohen-Addad\*, Andreas Emil Feldmann†, David Saulpic\*‡

\*Sorbonne Université, CNRS, LIP6, F-75005 Paris, France

† Charles University in Prague, Czechia

‡ École Normale Supérieure, Paris, France

Email: vcohenad@gmail.com, feldmann.a.e@gmail.com, david.saulpic@lip6.fr

**Abstract**—We consider the classic Facility Location,  $k$ -Median, and  $k$ -Means problems in metric spaces of constant doubling dimension. We give the first nearly linear-time approximation schemes for each problem, making a significant improvement over the state-of-the-art algorithms.

Moreover, we show how to extend the techniques used to get the first efficient approximation schemes for the problems of prize-collecting  $k$ -Medians and  $k$ -Means, and efficient bicriteria approximation schemes for  $k$ -Medians with outliers,  $k$ -Means with outliers and  $k$ -Center.

## I. INTRODUCTION

The  $k$ -Median and  $k$ -Means problems are classic clustering problems that are highly popular for modeling the problem of computing a “good” partition of a set of points of a metric space into  $k$  parts so that points that are “close” should be in the same part. Since a good clustering of a dataset allows to retrieve information from the underlying data, the  $k$ -Median and  $k$ -Means problems are the cornerstone of various approaches in data analysis and machine learning. The design of efficient algorithms for these clustering problems has thus become an important challenge.

The input for the problems is a set of points in a metric space and the objective is to identify a set of  $k$  centers  $C$  such that the sum of the  $p$ th power of the distance from each point of the metric to its closest center in  $C$  is minimized. In the  $k$ -Median problem,  $p$  is set to 1 while in the  $k$ -Means problem,  $p$  is set to 2. In general metric spaces both problems are known to be APX-hard, and this hardness even extends to Euclidean spaces of any dimension  $d = \Omega(\log n)$  [?]. Both problems also remain NP-hard for points in  $\mathbb{R}^2$  [?]. For  $k$ -Center, the goal is to minimize the maximum distance from each point in the metric to its closest center. This problem is APX-hard even in Euclidean Spaces [?], and computing a solution with optimal cost but  $(1 + \varepsilon)k$  centers requires time at least  $\Omega(n^{\sqrt{1/\varepsilon}})$  [?]. Therefore, to get an efficient approximation scheme one needs to approximate both the number of centers and the cost. (See Section I-C for more related work).

To bypass these hardness of approximation results, researchers have considered low-dimensional inputs like Euclidean spaces of fixed dimension or more generally metrics of fixed doubling dimension. There has been a large body

of work to design good tools for clustering in metrics of fixed doubling dimension, from the general result of Talwar [?] to very recent coresets constructions for clustering problems [?]. In their seminal work, Arora et al. [?] gave a polynomial time approximation scheme (PTAS) for  $k$ -Median in  $\mathbb{R}^2$ , which generalizes to a quasi-polynomial time approximation scheme (QPTAS) for inputs in  $\mathbb{R}^d$ . This result was improved in two ways. First by Talwar [?] who generalized the result to any metric space of fixed doubling dimension. Second by Kolliopoulos and Rao [?] who obtained an  $f(\varepsilon, d) \cdot n \log^{d+6} n$  time algorithm for  $k$ -Median in  $d$ -dimensional Euclidean space. Unfortunately, Kolliopoulos and Rao’s algorithm relies on the Euclidean structure of the input and does not immediately generalize to low dimensional doubling metric. Thus, until recently the only result known for  $k$ -Median in metrics of fixed doubling dimension was a QPTAS. This was also the case for slightly simpler problems such as Uniform Facility Location. Moreover, as pointed out in [?], the classic approach of Arora et al. [?] cannot work for the  $k$ -Means problem. Thus no efficient algorithms were known for the  $k$ -Means problem, even in the plane.

Recently, Friggstad et al. [?] and Cohen-Addad et al. [?] showed that the classic local search algorithm for the problems gives a  $(1 + \varepsilon)$ -approximation in time  $n^{1/\varepsilon^{O(d)}}$  in Euclidean space, in time  $n^{O(1/\varepsilon^2)}$  for planar graphs (which also extends to minor-free graphs), and in time  $n^{(d/\varepsilon)^{O(d)}}$  in metrics of doubling dimension  $d$  [?]. More recently Cohen-Addad [?] showed how to speed up the local search algorithm for Euclidean space to obtain a PTAS with running time  $nk(\log n)^{(d/\varepsilon)^{O(d)}}$ .

Nonetheless, obtaining an efficient approximation scheme (namely an algorithm running in time  $f(\varepsilon, d)\text{poly}(n)$ ) for  $k$ -Median and  $k$ -Means in metrics of doubling dimension  $d$  has remained a major challenge.

The versatility of the techniques we develop to tackle these problems allows us to consider a broader setting, where the clients do not necessarily have to be served. In the prize-collecting version of the problems, every client has a penalty cost that can be paid instead of its serving cost. In the  $k$ -Median (resp.  $k$ -Means) with outliers problems, the goal is to serve all but  $z$  clients, and the cost is measured on

the remaining ones with the  $k$ -Median (resp.  $k$ -Means) cost. These objectives can help to handle some noise from the input: the  $k$ -Median objective can be dramatically perturbed by the addition of a few distant clients, which must then be discarded.

### A. Our Results

We solve this open problem by proposing the first near-linear time algorithms for the  $k$ -Median and  $k$ -Means problems in metrics of fixed doubling dimension. More precisely, we show the following theorems, where we let  $f(\varepsilon) = (1/\varepsilon)^{1/\varepsilon}$ .

**Theorem I.1.** *For any  $0 < \varepsilon < 1/3$ , there exists a randomized  $(1 + \varepsilon)$ -approximation algorithm for  $k$ -Median in metrics of doubling dimension  $d$  with running time  $f(\varepsilon)^{2^{O(d^2)}} n \log^4 n + 2^{O(d)} n \log^9 n$  and success probability at least  $1 - \varepsilon$ .*

**Theorem I.2.** *For any  $0 < \varepsilon < 1/3$ , there exists a randomized  $(1 + \varepsilon)$ -approximation algorithm for  $k$ -Means in metrics of doubling dimension  $d$  with running time  $f(\varepsilon)^{2^{O(d^2)}} n \log^5 n + 2^{O(d)} n \log^9 n$  and success probability at least  $1 - \varepsilon$ .*

Our results also extend to the Facility Location problem, in which no bound on the number of opened centers is given, but each center comes with an opening cost. The aim is to minimize the sum of the (1st power) of the distances from each point of the metric to its closest center, in addition to the total opening costs of all used centers.

**Theorem I.3.** *For any  $0 < \varepsilon < 1/3$ , there exists a randomized  $(1 + \varepsilon)$ -approximation algorithm for Facility Location in metrics of doubling dimension  $d$  with running time  $f(\varepsilon)^{2^{O(d^2)}} \cdot n + 2^{O(d)} n \log n$  and success probability at least  $1 - \varepsilon$ .*

In all these theorems, we make the common assumption to have access to the distances of the metric in constant time, as, e.g., in [?], [?], [?]. This assumption is discussed in Bartal et al. [?].

Note that the double-exponential dependence on  $d$  is unavoidable unless  $P = NP$ , since the problems are APX-hard in Euclidean space of dimension  $d = O(\log n)$ . For Euclidean inputs, our algorithms for the  $k$ -Means and  $k$ -Median problems outperform the ones of Cohen-Addad [?], removing in particular the dependence on  $k$ , and the one of Kolliopoulos and Rao [?] when  $d > 3$ , by removing the dependence on  $\log^{d+6} n$ . Interestingly, for  $k = \omega(\log^9 n)$  our algorithm for the  $k$ -Means problem is faster than popular heuristics like  $k$ -Means++ which runs in time  $O(nk)$  in Euclidean space.

We note that the success probability can be boosted to  $1 - \varepsilon^\delta$  by repeating the algorithm  $\log \delta$  times and outputting the best solution encountered.

After proving the three theorems above, we will apply the techniques to prove the following ones. We say an algorithm is an  $(\alpha, \beta)$ -approximation for  $k$ -Medians or  $k$ -Means with outliers if its cost is within an  $\alpha$  factor of the optimal one and the solution drops  $\beta z$  outliers. Similarly, an algorithm is an  $(\alpha, \beta)$ -approximation for  $k$ -Center if its cost is within an  $\alpha$  factor of the optimal one and the solution opens  $\beta k$  centers.

**Theorem I.4.** *For any  $0 < \varepsilon < 1/3$ , there exists a randomized  $(1 + \varepsilon)$ -approximation algorithm for Prize-Collecting  $k$ -Median (resp.  $k$ -Means) in metrics of doubling dimension  $d$  with running time  $f(\varepsilon)^{2^{O(d^2)}} n \log^4 n + 2^{O(d)} n \log^9 n$  and success probability at least  $1 - \varepsilon$ .*

**Theorem I.5.** *For any  $0 < \varepsilon < 1/3$ , there exists a randomized  $(1 + \varepsilon, 1 + O(\varepsilon))$ -approximation algorithm for  $k$ -Median (resp.  $k$ -Means) with outliers in metrics of doubling dimension  $d$  with running time  $f(\varepsilon)^{2^{O(d^2)}} n \log^6 n + T(n)$  and success probability at least  $1 - \varepsilon$ , where  $T(n)$  is the running time to construct a constant-factor approximation.*

We note as an aside that our proof of Theorem I.5 could give an approximation where at most  $z$  outliers are dropped, but  $(1 + O(\varepsilon))k$  centers are opened. For simplicity, we focused on the previous case.

**Theorem I.6.** *For any  $0 < \varepsilon < 1/3$ , there exists a randomized  $(1 + \varepsilon, 1 + O(\varepsilon))$ -approximation algorithm for  $k$ -Center in metrics of doubling dimension  $d$ , with running time  $f(\varepsilon)^{2^{O(d^2)}} n \log^6 n + n \log k$  and success probability at least  $1 - \varepsilon$ .*

As explained above, this bicriteria is necessary in order to get an efficient algorithm: it is APX-hard to approximate the cost [?], and achieving the optimal cost with  $(1 + \varepsilon)k$  centers requires a complexity  $\Omega(n^{1/\sqrt{\varepsilon}})$  [?].

### B. Techniques

To give a detailed insight on our techniques and our contribution we first need to quickly review previous approaches for obtaining approximation schemes on bounded doubling metrics. The general approach, due to Arora [?] and Mitchell [?], which was generalized to doubling metrics by Talwar [?], is the following.

1) *Previous Techniques:* The approach consists in randomly partitioning the metric into a constant number of regions, and applying this recursively to each region. The recursion stops when the regions contain only a constant number of input points. This leads to what is called a *split-tree decomposition*: a partition of the space into a finer and finer level of granularity. The reader who is not familiar with the split-tree decomposition of Talwar may refer to Section II-B for a more formal introduction.

*Portals:* The approach then identifies a specific set of points for each region, called *portals*, which allows to show

that there exists a near-optimal solution such that different regions “interplay” only through portals. For example, in the case of the Traveling Salesperson (TSP) problem, it is possible to show that there exists a near-optimal tour that enters and leaves a region only through its portals. In the case of the  $k$ -Median problem a client located in a specific region can be assigned to a facility in a different region only through a path that goes to a portal of the region. In other words, clients can “leave” a region only through the portals.

Proving the existence of such a structured near-optimal solution relies on the fact that the probability that two very close points end up in different regions of large diameter is very unlikely. Hence the expected *detour* paid by going through a portal of the region is small compared to the original distance between the two points, if the portals are dense enough.

For the sake of argument, we provide a proof sketch of the standard proof of Arora [?]. We will use a refined version of this idea in later sections. The split-tree recursively divides the input metric  $(V, \text{dist})$  into parts of smaller and smaller diameter. The root part consists of the entire point set and the parts at level  $i$  are of diameter roughly  $2^i$ . The set of portals of a part of level  $i$  is an  $\varepsilon_0 2^i$ -net for some  $\varepsilon_0$ , which is a small set such that every point of the metric is at distance at most  $\varepsilon_0 2^i$  to it. Consider two points  $u, v$  and let us bound the expected detour incurred by connecting  $u$  to  $v$  through portals. This detour is determined by a path that starting from  $u$  at the lowest level, in each step connects a vertex  $w$  to its closest net point of the part containing  $w$  on the next higher level. This is done until the lowest-level part  $R_{u,v}$  (i.e., the part of smallest diameter) is reached, which contains both  $u$  and  $v$ , from where a similar procedure leads from this level through portals of smaller and smaller levels all the way down to  $v$ . If the level of  $R_{u,v}$  is  $i$  then the detour, i.e., the difference between  $\text{dist}(u, v)$  and the length of the path connecting  $u$  and  $v$  through portals, is  $O(\varepsilon_0 2^i)$  by the definition of the net. Moreover, the proof shows that the probability that  $u$  and  $v$  are not in the same part on level  $i$  is at most  $\text{dist}(u, v)/2^i$ . Thus, the expected detour for connecting  $u$  to  $v$  is  $\sum_{\text{level } i} \Pr[R_{u,v} \text{ is at level } i] \cdot O(\varepsilon_0 2^i) = \sum_{\text{level } i} O(\varepsilon_0 \text{dist}(u, v))$ . Hence, setting  $\varepsilon_0$  to be some  $\varepsilon$  divided by the number of levels yields that the expected detour is  $O(\varepsilon \text{dist}(u, v))$ .

*Dynamic programming:* The portals now act as separators between different parts and allows to apply a dynamic programming (DP) approach for solving the problems. The DP consists of a DP-table entry for each part and for each *configuration* of the portals of the part. Here a configuration is a potential way the near-optimal solution interacts with the part. For example, in the case of TSP, a configuration is the information at which portal the near-optimal tour enters and leaves and how it connects the portals on the outside and inside of the part. For the  $k$ -Median problem, a configuration stores how many clients outside (respectively

inside) the part connect through each portal and are served by a center located inside (respectively outside). Then the dynamic program proceeds in a bottom-up fashion along the split-tree to fill up the DP table. The running time of the dynamic program depends exponentially on the number of portals.

*How many portals?:* The challenges that need to be overcome when applying this approach, and in particular to clustering problems, are two-fold. First the “standard” use of the split-tree requires  $O((\frac{\log n}{\varepsilon})^d)$  portals per part in order to obtain a  $(1 + \varepsilon)$ -approximation, coming from the fact that the number of levels can be assumed to be logarithmic in the number of input points. This often implies quasi-polynomial time approximation schemes since the running time of the dynamic program has exponential dependence on the number of portals. This is indeed the case in the original paper by Talwar [?] and in the first result on clustering in Euclidean space by Arora et al. [?]. However, in some cases, one can lower the number of portals per part needed. In Euclidean space for example, the celebrated “patching lemma” [?] shows that only a constant number (depending on  $\varepsilon$ ) of portals are needed for TSP. Similarly, Kolliopoulos and Rao [?] showed that for  $k$ -Median in Euclidean space only a constant number of portal are needed, if one uses a slightly different decomposition of the metric.

Surprisingly, obtaining such a result for doubling metrics is much more challenging. To the best of our knowledge, this work is the first one to reduce the number of portals to a constant.

A second challenge when working with split-tree decompositions and the  $k$ -Means problem is that because the cost of assigning a point to a center is the squared distance, the analysis of Arora, Mitchell, and Talwar does not apply. If two points are separated at a high level of the split-tree, then making a detour to the closest portal may incur an expected cost much higher than the cost of the optimal solution.

2) *Our Contributions:* Our contribution can be viewed as a “patching lemma” for clustering problems in doubling metrics. Namely, an approach that allows to solve the problems mentioned above: (1) it shows how to reduce the number of portals to a constant, similar to the one given by the patching lemma for TSP, (2) it works for any clustering objective which is defined as the sum of distances to some constant  $p$  (with  $k$ -Median and  $k$ -Means as prominent special cases), and (3) it works not only for Euclidean but also for doubling metrics.

Our starting point is the notion of *badly cut* vertices of Cohen-Addad [?] for the capacitated version of the above clustering problems. To provide some intuition on the definition, let us start with the following observation: consider a center  $f$  of the optimal solution and a client  $c$  assigned to  $f$ . If the diameter of the lowest-level part containing both  $f$  and  $c$  is of order  $\text{dist}(c, f)$  (say at most




Figure 1: Illustration of badly cut. The black point is  $c$  (resp.  $l$ ), the gray one is  $L(c)$  (resp.  $f_0$ ), and the blue point is  $q$ . The dashed line is the boundary of a part with “large” diameter.

$\text{dist}(c, f)/\varepsilon^2$ , then by taking a large enough but constant size net as a set of portals in each part (say an  $\varepsilon^3 2^i$ -net for a part of level  $i$ ), the total detour for the two points is at most  $O(\varepsilon \text{dist}(c, f))$ , which is acceptable.

The problematic scenario is when the lowest-level part containing  $f$  and  $c$  is of diameter much larger than  $\text{dist}(c, f)$ . In this case, it is impossible to afford a detour proportional to the diameter of the part in the case of the  $k$ -Medians and  $k$ -Means objective. To handle this case we first compute a constant approximation  $L$  (via some known algorithm) and use it to guide us towards a  $(1 + \varepsilon)$ -approximation.

*Badly cut clients and facilities:* Consider a client  $c$  and the center  $L(c)$  serving  $c$  in  $L$  (i.e.,  $L(c)$  is closest to  $c$  among the centers in  $L$ ), and call  $\text{OPT}(c)$  the facility of an optimum solution  $\text{OPT}$  that serves  $c$  in  $\text{OPT}$ . We say that  $c$  is *badly cut* if there is a point  $q$  in the ball centered at  $c$  of radius  $\text{dist}(c, L(c))/\varepsilon$  such that the highest-level part containing  $c$  and not  $q$  is of diameter much larger than  $\text{dist}(c, L(c))/\varepsilon$  (say greater than  $\text{dist}(c, L(c))/\varepsilon^3$ ). In other words, there is a point  $q$  in this ball such that paying a detour through the portal to connect  $c$  to  $q$  yields a detour larger than  $\varepsilon \text{dist}(c, q)$  (see Fig. 1).

Similarly, we say that a center  $l$  is *badly cut* if there is a point  $q$  in the ball centered at  $l$  of radius  $\text{dist}(l, f_0)$  (where  $f_0$  is the facility of  $\text{OPT}$  that is the closest to  $l$ ) such that the highest-level part containing  $l$  and not  $q$  is of diameter  $\text{dist}(l, f_0)/\varepsilon^2$ . The crucial property here is that any client  $c$  or any facility  $l$  is badly cut with probability  $O(\varepsilon^3)$ , as we will show.

*Using the notion of badly cut:* We now illustrate how this notion can help us. Assume for simplicity that  $\text{OPT}(c)$  is in the ball centered at a client  $c$  of radius  $\text{dist}(c, L(c))/\varepsilon$  (if this is not the case then  $\text{dist}(c, \text{OPT}(c))$  is much larger than  $\text{dist}(c, L(c))$ , so this is a less problematic scenario and a simple idea can handle it). If  $c$  is not badly cut, then the lowest-level part containing both  $c$  and  $\text{OPT}(c)$  is of diameter not much larger than  $\text{dist}(c, L(c))/\varepsilon$ . Taking a sufficiently fine net for each part (independent of the number of levels) allows to bound the detour through the portals to reach  $\text{OPT}(c)$  from  $c$  by at most  $\varepsilon \text{dist}(c, L(c))$ . Since  $L$  is

an  $O(1)$ -approximation, this is fine.

If  $c$  is badly cut, then we modify the instance by relocating  $c$  to  $L(c)$ . That is, we will work with the instance where there is no more client at  $c$  and there is an additional client at  $L(c)$ . We claim that any solution in the modified instance can be lifted to the original instance at an expected additional cost of  $O(\varepsilon^3 \text{OPT})$ . This comes from the fact that the cost increase for a solution is, by the triangle inequality, at most the sum of distances of the badly cut clients to their closest facility in the local solution. This is at most  $O(\varepsilon^3 \text{OPT})$  in expectation since each client is badly cut with probability at most  $O(\varepsilon^3)$  and  $L$  is an  $O(1)$ -approximation.

Here we should ask, what did we achieve by moving  $c$  to  $L(c)$ ? Note that  $c$  should now be assigned to facility  $f$  of  $\text{OPT}$  that is the closest to  $L(c)$ . So we can make the following observation: If  $L(c)$  is not badly cut, then the detour through the portals when assigning  $c$  to  $f$  is fine (namely at most  $\varepsilon$  times the distance from  $L(c)$  to its closest facility in  $\text{OPT}$ ). Otherwise, if  $L(c)$  is also badly cut, then we simply argue that there exists a near-optimal solution which contains  $L(c)$ , in which case  $c$  is now served optimally at a cost of 0 (in the new instance).

*From bicriteria to opening exactly  $k$  centers:* Since  $L(c)$  is badly cut with probability  $O(\varepsilon^3)$ , this leads to a solution opening  $(1 + O(\varepsilon^3))k$  centers. At first, it looks difficult to then reduce the number of centers to  $k$  without increasing the cost of the solution by a factor larger than  $(1 + \varepsilon)$ . However, and perhaps surprisingly, we show in Lemma IV.6 that this can be avoided: we show that there exists a near-optimal solution that contains the badly cut centers of  $L(c)$ .

We can then conclude that a near-optimal solution can be computed by a simple dynamic-programming procedure on the split-tree decomposition to identify the best solution in the modified instance.

Our result on Facility Location in Section III provides a simple illustration of these ideas — avoiding the bicriteria issue due to the hard bound on the number of opened facilities for the  $k$ -Median and  $k$ -Means problems. Our main result on  $k$ -Median and  $k$ -Means is described in Section IV-C. We discuss some extensions of the framework in Section V.

### C. Related work

*On clustering problems:* The clustering problems considered in this paper are known to be NP-hard, even restricted to inputs lying in the Euclidean plane (see Mahajan et al. [?] or Dasgupta and Freund [?] for  $k$ -Means, Megiddo and Supowit [?] for the problems with outliers, and Masuyama et al. [?] for  $k$ -Center). The problems of Facility Location and  $k$ -Median have been studied since a long time in graphs, see e.g. [?]. The current best approximation ratio for metric Facility Location is 1.488, due to Li [?], whereas it is 2.67 for  $k$ -Median, due to Byrka et al. [?].

The problem of  $k$ -Means in general graphs also received a lot of attention (see e.g., Kanungo et al. [?]) and the best approximation ratio is 6.357, due to Ahmadian et al. [?].

Clustering problems with outliers were first studied by Charikar et al. [?], who devised an  $(O(1), (1 + O(\varepsilon)))$ -approximation for  $k$ -Median with outliers and a constant factor approximation for prize-collecting  $k$ -Median. More recently, Friggstad et al. [?] showed that local search provides a bicriteria approximation, where the number of centers is approximate instead of the number of outliers. However, the runtime is  $n^{f(\varepsilon, d)}$ , and thus we provide a much faster algorithm. To the best of our knowledge, we present the first approximation scheme that preserves the number of centers.

The  $k$ -Center problem is known to be NP-hard to approximate within any factor better than 2, a bound that can be achieved by a greedy algorithm [?]. This is related to the problem of covering points with a minimum number of disks (see e.g. [?], [?]). Marx and Pilipczuk [?] proposed an exact algorithm running in time  $n^{\sqrt{k}+O(1)}$  to find the maximum number of points covered by  $k$  disks and showed a matching lower bound, whereas Liao et al. [?] presented an algorithm running in time  $O(mn^{O(1/\varepsilon^2 \log^2 1/\varepsilon)})$  to find a  $(1 + \varepsilon)$ -approximation to the minimal number of disks necessary to cover all the points (where  $m$  is the total number of disks and  $n$  the number of points). This problem is closely related to  $k$ -Center: the optimal value of  $k$ -Center on a set  $V$  is the minimal number  $L$  such that there exist  $k$  disks of radius  $L$  centered on points of  $V$  covering all points of  $V$ . Hence, the algorithm from [?] can be directly extended to find a solution to  $k$ -Center with  $(1 + \varepsilon)k$  centers and optimal cost. Loosing on the approximation allows us to present a much faster algorithm.

*On doubling dimension:* Despite their hardness in general metrics, these problems admit a PTAS when the input is restricted to a low dimensional metric space: Friggstad et al. [?] showed that local search gives a  $(1 + \varepsilon)$ -approximation. However, the running time of their algorithm is  $n^{(d/\varepsilon)^{O(d)}}$  in metrics with doubling dimension  $d$ .

A long line of research exists on filling the gap between results for Euclidean spaces and metrics with bounded doubling dimension. This started with the work of Talwar [?], who gave QPTASs for a long list of problems. The complexity for some of these problems was improved later on: for the Traveling Salesperson problem, Gottlieb [?] gave a near-linear time approximation scheme, Chan et al. [?] gave a PTAS for Steiner Forest, and Gottlieb [?] described an efficient spanner construction.

## II. PRELIMINARIES

### A. Definitions

Consider a metric space  $(V, \text{dist})$ . For a vertex  $v \in V$  and an integer  $r \geq 0$ , we let  $\beta(v, r) = \{w \in V \mid \text{dist}(v, w) \leq r\}$  be the *ball* around  $v$  with radius  $r$ . The *doubling dimension*

of a metric is the smallest integer  $d$  such that any ball of radius  $2r$  can be covered by  $2^d$  balls of radius  $r$ . We call  $\Delta$  the aspect-ratio (sometimes referred to as *spread* in the literature) of the metric, i.e., the ratio between the largest and the smallest distance.

Given a set of points called *clients* and a set of points called *candidate centers* in a metric space, the goal of the  $k$ -Median problem is to output a set of  $k$  centers (or *facilities*) chosen among the candidate centers that minimizes the sum of the distances from each client to its closest center. More formally, an instance to the  $k$ -Median problem is a 4-tuple  $(C, F, \text{dist}, k)$ , where  $(C \cup F, \text{dist})$  is a metric space and  $k$  is a positive integer. The goal is to find a set  $S \subseteq F$  such that  $|S| \leq k$  and  $\sum_{c \in C} \min_{f \in S} (\text{dist}(c, f))$  is minimized. Let  $n = |C \cup F|$ . The  $k$ -Means problem is identical except from the objective function which is  $\sum_{c \in C} \min_{f \in S} (\text{dist}(c, f))^2$ .

In the *Facility Location* problem, the number of centers in the solution is not limited but there is a cost  $w_f$  for each candidate center  $f$  and the goal is to find a solution  $S$  minimizing  $\sum_{c \in C} \min_{f \in S} (\text{dist}(c, f)) + \sum_{f \in S} w_f$ .

For those clustering problems, it is convenient to name the center serving a client. For a client  $c$  and a solution  $S$ , we denote  $S(c)$  the center closest to  $c$ , and  $S_c := \text{dist}(c, S(c))$  the distance to it.

In this paper, we consider the case where the set of candidate centers is part of the input. A variant of the  $k$ -Median and  $k$ -Means problems in Euclidean metrics allows to place centers anywhere in the space and specifies the input size as simply the number of clients. We note that up to losing a polylogarithmic factor in the running time, it is possible to reduce this variant to our setting by computing a set of candidate centers that approximate the best set of centers in  $\mathbb{R}^d$  [?].

A  $\delta$ -net of  $V$  is a set of points  $X \subseteq V$  such that for all  $v \in V$  there is an  $x \in X$  such that  $\text{dist}(v, x) \leq \delta$ , and for all  $x, y \in X$  we have  $\text{dist}(x, y) > \delta$ . A net is therefore a set of points not too close to each other, such that every point of the metric is close to a net point. The following lemma bounds the cardinality of a net in doubling metrics.

**Lemma II.1** (from Gupta et. al [?]). *Let  $(V, d)$  be a metric space with doubling dimension  $d$  and diameter  $\Delta$ , and let  $X$  be a  $\delta$ -net of  $V$ . Then  $|X| \leq 2^{d \lceil \log_2(\Delta/\delta) \rceil}$ .*

Another property of doubling metrics that will be useful for our purpose is the existence of low-stretch spanners with a linear number of edges. More precisely, Har-Peled and Mendel [?] showed that one can find a graph (called a *spanner*) in the input metric that has  $O(n)$  edges such that distances in the graph approximate the original distances up to a constant factor. This construction takes time  $2^{O(d)}n$ . We will make use of these spanners only for computing constant-factor approximations of our problems: for this purpose, we will therefore assume that the number of edges is  $m = 2^{O(d)}n$ .



We will also make use the following lemma.

**Lemma II.2** ([?]). *Let  $p \geq 0$  and  $1/2 > \varepsilon > 0$ . For any  $a, b, c \in A \cup F$ , we have  $\text{dist}(a, b)^p \leq (1 + \varepsilon)^p \text{dist}(a, c)^p + \text{dist}(c, b)^p (1 + 1/\varepsilon)^p$ .*

### B. Decomposition of Metric Spaces

As pointed out in our techniques section, we will make use of hierarchical decompositions of the input metric. We define a *hierarchical decomposition* (sometimes simply a decomposition) of a metric  $(V, \text{dist})$  as a collection of partitions  $\mathcal{D} = \{\mathcal{B}_0, \dots, \mathcal{B}_{|\mathcal{D}|}\}$  that satisfies the following:

- each  $\mathcal{B}_i$  is a partition of  $V$ ,
- $\mathcal{B}_i$  is a *refinement* of  $\mathcal{B}_{i+1}$ , namely for each part  $B \in \mathcal{B}_i$  there exists a part  $B' \in \mathcal{B}_{i+1}$  that contains  $B$ ,
- $\mathcal{B}_0$  contains a singleton set for each  $v \in V$ , while  $\mathcal{B}_{|\mathcal{D}|}$  is a trivial partition that contains only one set, namely  $V$ .

We define the  *$i$ th level* of the decomposition to be the partition  $\mathcal{B}_i$ , and call  $B \in \mathcal{B}_i$  a level- $i$  part. If  $B' \in \mathcal{B}_{i-1}$  is such that  $B' \subset B$ , we say that  $B'$  is a *subpart* of  $B$ .

For a given decomposition  $\mathcal{D} = \{\mathcal{B}_0, \dots, \mathcal{B}_{|\mathcal{D}|}\}$ , we say that a vertex  $u$  is *cut from  $v$  at level  $j$*  if  $j$  is the maximum integer such that  $v$  is in some  $B \in \mathcal{B}_j$  and  $u$  is in some  $B' \in \mathcal{B}_j$  with  $B \neq B'$ . For a vertex  $v \in \mathcal{F}$  we say that the ball  $\beta(v, 2^i)$  is *cut by  $\mathcal{D}$  at level  $j$*  if there is at least one vertex of the ball that is cut from  $v$  at level  $j$ .

A key ingredient for our result is the following lemma, that introduces some properties of the hierarchical decomposition (sometimes referred to as *split-tree*) proposed by Talwar [?] for low-doubling metrics.

**Lemma II.3** (Reformulation of [?], [?]). *For any metric  $(V, \text{dist})$  of doubling dimension  $d$  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  $2^{i+1}$ ,  $|\mathcal{D}| \leq \lceil \log_2(\text{diam}(V)) \rceil$ , and:*

- 1) **Scaling probability:** for any  $v \in V$ , radius  $r$ , and level  $i$ , we have

$$\Pr[\mathcal{D} \text{ cuts } \beta(v, r) \text{ at a level } i] \leq 2^{2d+2} r / 2^i.$$

- 2) **Concise and precise portal set:** For any set  $B \in \mathcal{B}_i$  where  $\mathcal{B}_i \in \mathcal{D}$ , there is a set of portals  $\mathcal{P}_B$  such that,

- a) **concise:**  $|\mathcal{P}_B| \leq 1/\rho^d$ ; and
- b) **precise:** for any ball  $\beta(v, r) \subseteq B$  cut by  $\mathcal{C}_T$  at level  $i$  and pair of distinct sets  $B_1, B_2 \in \mathcal{B}_{i-1}$  on level  $i-1$ , we have for any  $u \in B_1 \cap \beta(v, r)$ , and  $w \in B_2 \cap \beta(v, r)$ ,

$$\min_{p \in \mathcal{P}_B} \{\text{dist}(u, p) + \text{dist}(p, w)\} \leq \text{dist}(u, w) + O(\rho 2^i).$$

Moreover, this decomposition can be found in time  $(1/\rho)^{O(d)} n \log \Delta$ .

### C. Formal Definition of Badly Cut Vertices

As sketched in the introduction, the notion of badly cut lies at the heart of our analysis. We define it formally here. We denote  $\kappa(\varepsilon, p) = \frac{\varepsilon^{p+2}}{(p+1)^p}$  and  $\tau(\varepsilon, d) = 2d + 2 + \log \log(1/\varepsilon) + \log(1/\kappa(\varepsilon, p))$ , two parameters that are often used throughout this paper.

**Definition II.4.** *Let  $(V, \text{dist})$  be a doubling metric, let  $\mathcal{D}$  be a hierarchical decomposition on  $(V, \text{dist})$ , and a  $\varepsilon > 0$ . A client  $v$  is badly cut w.r.t.  $\mathcal{D}$  if there exists an integer  $i$  such that  $2^i \in [\varepsilon L_v, L_v/\varepsilon]$  and  $\beta(v, 2^i)$  is cut at some level  $j$  greater than  $i + \tau(\varepsilon, d)$ .*

*Similarly, a center  $f$  of  $L$  is badly cut w.r.t.  $\mathcal{D}$  if there exists an integer  $i$  such that  $2^i \in [\varepsilon \text{OPT}_f, \text{OPT}_f/\varepsilon]$  and  $\beta(f, 2^i)$  is cut at some level  $j$  greater than  $i + \tau(\varepsilon, d)$ , where  $\text{OPT}_f$  is the distance from  $f$  to the closest facility of OPT.*

In the following, when  $\mathcal{D}$  is clear from the context we simply say badly cut. The following lemma bounds the probability of being badly cut.

\*\*\* (David) Do we really need several rings? \*\*\*

←—1

**Lemma II.5.** *Let  $(C \cup F, \text{dist})$  be a metric, and  $\mathcal{D}$  a random hierarchical decomposition given by Lemma II.3. Let  $v$  be a vertex in  $C \cup F$ . The probability that  $v$  is badly cut is at most  $\kappa(\varepsilon, p)$ .*

*Proof:* Consider first a vertex  $v \in C$ . By Property 1, the probability that a ball  $\beta(v, 2^i)$  is cut at level at least  $j$  is at most  $2^{2d+2} 2^i / 2^j$ . Hence the probability that a ball  $\beta(v, 2^i)$ , where  $2^i \in [\varepsilon^2 L_v, L_v/\varepsilon^2]$ , is cut at a level  $j$  greater than  $i + \lceil \log \log(1/\varepsilon) + \log(1/\kappa(\varepsilon, p)) \rceil$  is at most  $\frac{\kappa(\varepsilon, p)}{4 \log(1/\varepsilon)}$ . Taking a union bound over all balls of radius  $2^i$  such that  $i$  is an integer and  $2^i \in [\varepsilon^2 L_v, L_v/\varepsilon^2]$  we have that the probability that  $v$  is badly cut is at most  $4 \log(1/\varepsilon) \cdot \frac{\kappa(\varepsilon, p)}{4 \log(1/\varepsilon)} = \kappa(\varepsilon, p)$ .

The proof for  $v \in F$  is identical. ■

### D. Preprocessing

In the following, we will work with the slightly more general version of the clustering problems where there is some demand on each vertex: there is a function  $\chi : C \mapsto \{1, \dots, n\}$  and the goal is to minimize  $\sum_{c \in C} \chi(c) \cdot \min_{f \in S} \text{dist}(c, f) + \sum_{f \in S} w_f$  for the Facility Location problem, or  $\sum_{c \in C} \chi(c) \cdot \min_{f \in S} \text{dist}(c, f)$  and  $\sum_{c \in C} \chi(c) \cdot \min_{f \in S} \text{dist}(c, f)^2$  for  $k$ -Median and  $k$ -Means respectively. This also extends to any  $\sum_{c \in C} \chi(c) \cdot \min_{f \in S} \text{dist}(c, f)^p$  with constant  $p$ .

We will preprocess the input instance to transform it into several instances of the more general clustering problem, ensuring that the aspect-ratio  $\Delta$  of each instance is polynomial. We defer this construction to Appendix A.

### III. A NEAR-LINEAR TIME APPROXIMATION SCHEME FOR NON-UNIFORM FACILITY LOCATION

To demonstrate the utility of the notion of badly cut, we show how to use it to get a near-linear time approximation

scheme for Facility Location in metrics of bounded doubling dimension. In this context we refer to centers in the set  $F$  of the input as facilities.

We first show a structural lemma that allows to focus on instances that do not contain any badly cut client. Then, we prove that these instance have *portal-respecting* solutions that are nearly optimal, and that can be computed with a dynamic program. We conclude by providing a fast dynamic program, that takes advantage of all the structure provided before.

### A. Structural Lemma

Let  $\varepsilon > 0$ , and consider a metric space  $(V, \text{dist})$  and an instance  $\mathcal{I}$  of the Facility Location problem on  $(V, \text{dist})$ . Namely, an instance whose client and candidate center sets are subsets of  $V$ . Our first step is to show that, given  $\mathcal{I}$ , a randomized decomposition  $\mathcal{D}$  of  $(V, \text{dist})$  and any solution  $L$  for  $\mathcal{I}$  on  $(V, \text{dist})$ , we can build an instance  $\mathcal{I}_{\mathcal{D}}$  such that any solution  $S$  has a similar cost in  $\mathcal{I}$  and in  $\mathcal{I}_{\mathcal{D}}$ , and more importantly  $\mathcal{I}_{\mathcal{D}}$  does not contain any badly cut client with respect to  $\mathcal{D}$ . The definition of  $\mathcal{I}_{\mathcal{D}}$  depends on the randomness of  $\mathcal{D}$ . Define  $B_{\mathcal{D}}$  be the set of badly cut facilities of  $L$  w.r.t  $\mathcal{D}$ .

Let  $\text{cost}_{\mathcal{I}_0} : V \rightarrow \mathbb{R}$  be a function that given a set of centers in an instance  $\mathcal{I}_0$  on  $(V, \text{dist})$ , returns the  $k$ -Median cost induced by the set of centers in  $\mathcal{I}_0$ . For any instance  $\mathcal{I}_{\mathcal{D}}$  on  $(V, \text{dist})$ , we let

$$nu_{\mathcal{I}_{\mathcal{D}}} = \max_{\text{solution } S} (\text{cost}_{\mathcal{I}}(S) - (1 + 3\varepsilon)\text{cost}_{\mathcal{I}_{\mathcal{D}}}(S), \\ (1 - 3\varepsilon)\text{cost}_{\mathcal{I}_{\mathcal{D}}}(S) - \text{cost}_{\mathcal{I}}(S)).$$

We say that an instance  $\mathcal{I}_{\mathcal{D}}$  has *small distortion w.r.t.  $\mathcal{I}$*  if  $\sum_{f \in B_{\mathcal{D}}} w_f \leq \varepsilon \text{cost}_{\mathcal{I}}(L)$  and  $\nu_{\mathcal{I}_{\mathcal{D}}} \leq \varepsilon \text{cost}_{\mathcal{I}}(L)$ . When  $\mathcal{I}$  is clear from the context we simply say that  $\mathcal{I}_{\mathcal{D}}$  has small distortion.

In the following, we will always work with a particular  $\mathcal{I}_{\mathcal{D}}$  constructed from  $\mathcal{I}$  and a precomputed approximate solution  $L$  as follows:  $\mathcal{I}$  is transformed such that every badly cut client  $c$  is moved to  $L(c)$ , namely, there is no more client at  $c$  in  $\mathcal{I}_{\mathcal{D}}$  but an extra client is added at  $L(c)$ . All the other clients stay as they are.

What we would like to prove is that the optimal solution in  $\mathcal{I}$  can be transformed to a solution in  $\mathcal{I}_{\mathcal{D}}$  with a small additional cost, and vice versa. The intuition behind this is the following: a client of the solution  $L$  is badly cut with probability  $\kappa(\varepsilon, p)$  (from Lemma II.5), hence every client contributes with  $\kappa(\varepsilon, p)L_c$  to transform any solution  $S$  for the instance  $\mathcal{I}$  to a solution for the instance  $\mathcal{I}_{\mathcal{D}}$ , and vice versa.

However, we will need to convert a particular solution in  $\mathcal{I}_{\mathcal{D}}$  (think of it as  $\text{OPT}_{\mathcal{I}_{\mathcal{D}}}$ ) to a solution in  $\mathcal{I}$ : this particular solution depends in the randomness of  $\mathcal{D}$ , and this short argument does not apply because of dependency issues. It is

nevertheless possible to prove that  $\mathcal{I}_{\mathcal{D}}$  has a small distortion, as done in the following lemma.

**Lemma III.1.** *Given an instance  $\mathcal{I}$  of Facility Location, a randomized decomposition  $\mathcal{D}$  and a solution  $L$ , let  $\mathcal{I}_{\mathcal{D}}$  be the instance obtained from  $\mathcal{I}$  by moving every badly cut client  $c$  to  $L(c)$  (as described above). The probability that  $\mathcal{I}_{\mathcal{D}}$  has small distortion is at least  $1 - \varepsilon$ .*

*Proof:* To show the lemma, we will show that  $\mathbb{E} \left[ \sum_{f \in B_{\mathcal{D}}} w_f \right] \leq \varepsilon^2 \text{cost}(L)/2$  and  $\mathbb{E}[\nu_{\mathcal{I}_{\mathcal{D}}}] \leq \varepsilon^2 \text{cost}(L)/2$ . Then, by Markov's inequality and taking a union bound over the probabilities of failure yields the lemma. Note that  $\mathbb{E} \left[ \sum_{f \in B_{\mathcal{D}}} w_f \right] = \sum_{f \in L} \Pr[f \text{ badly cut}] \cdot w_f \leq \varepsilon^2 \text{cost}(L)/2$  is immediate from Lemma II.5.

We thus aim at showing that  $\mathbb{E}[\nu_{\mathcal{I}_{\mathcal{D}}}] \leq \varepsilon^2 \text{cost}(L)/2$ . In the sake of lightening equations, we will note  $\sum_{\text{bcc. } c}$  the sum over all badly cut clients  $c$ .

By definition, we have that for any solution  $S$ ,

$$\text{cost}(S) - \text{cost}_{\mathcal{I}_{\mathcal{D}}}(S) \leq \sum_{\text{bcc. } c} \text{dist}(c, S)^p - \text{dist}(S, L(c))^p \\ \leq \sum_{\text{bcc. } c} ((1 + 3\varepsilon)\text{dist}(S, L(c))^p \\ + \frac{\text{dist}(c, L(c))^p}{(\varepsilon/(p+1))^p} - \text{dist}(S, L(c))^p),$$

using Lemma II.2 with parameter  $\varepsilon/p$ . This is equal to

$$\sum_{\text{bcc. } c} 3\varepsilon \cdot \text{dist}(S, L(c))^p + \frac{\text{dist}(c, L(c))^p}{(\varepsilon/(p+1))^p},$$

and so we have

$$\text{cost}(S) - (1 + 3\varepsilon)\text{cost}_{\mathcal{I}_{\mathcal{D}}}(S) \leq \sum_{\text{bcc. } c} \frac{\text{dist}(c, L(c))^p}{(\varepsilon/(p+1))^p}$$

Similarly, we have that

$$\text{cost}_{\mathcal{I}_{\mathcal{D}}}(S) - \text{cost}(S) \leq \sum_{\text{bcc. } c} \text{dist}(S, L(c))^p - \text{dist}(c, S)^p \\ \leq \sum_{\text{bcc. } c} ((1 + 3\varepsilon)\text{dist}(c, S)^p \\ + \frac{\text{dist}(c, L(c))^p}{(\varepsilon/(p+1))^p} - \text{dist}(c, S)^p) \\ \leq \sum_{\text{bcc. } c} 3\varepsilon \cdot \text{dist}(c, S)^p + \frac{\text{dist}(c, L(c))^p}{(\varepsilon/(p+1))^p}$$

and we conclude

$$(1 - 3\varepsilon)\text{cost}_{\mathcal{I}_{\mathcal{D}}}(S) - \text{cost}(S) \leq \sum_{\text{bcc. } c} \frac{\text{dist}(c, L(c))^p}{(\varepsilon/(p+1))^p}$$

Therefore, the expected value of  $\nu_{\mathcal{I}_{\mathcal{D}}}$  is

$$E[\nu_{\mathcal{I}_{\mathcal{D}}}] \leq \sum_{\text{client } c} \Pr[c \text{ badly cut}] \cdot \frac{\text{dist}(c, L(c))^p}{(\varepsilon/(p+1))^p}.$$



Applying [Lemma II.5](#) and using  $\kappa(\varepsilon, p) = \frac{\varepsilon^{p+2}}{(p+1)^p}$ , we conclude  $E[\nu_{\mathcal{I}_D}] \leq \varepsilon^2 \cdot \text{cost}(L)$ . The lemma follows for a sufficiently small  $\varepsilon$ . ■

### B. Portal Respecting Solution

In the following, we fix an instance  $\mathcal{I}$ , a decomposition  $\mathcal{D}$  and a solution  $L$ . By [Lemma III.1](#),  $\mathcal{I}_D$  has small distortion with probability at least  $1 - \varepsilon$  and so we condition on this event from now on.

We explore the structure that this conditioning gives to the solution. We will show that there exists a solution  $\text{OPT}'$  with small cost such that each client  $c$  is cut from its serving facility  $f$  at a level at most  $\log(5(L_c + \text{OPT}_c)) + \tau(\varepsilon, d)$ . This allows to consider *portal-respecting* solution, where every client to facility path goes in and out parts of the decomposition only at designated portals. Indeed, the detour incurred by making a path portal respecting depends on the level where its extremities are cut: more precisely, it is an epsilon fraction of the diameter at that level. Hence, ensuring that this level stays small implies that the detour made is small (in our case,  $O(\varepsilon(L_c + \text{OPT}_c))$ ). Such a solution can be computed by a dynamic program that we will present afterwards.

In the following, we consider the solution  $\text{OPT}' = \text{OPT} \cup B_D$  (where  $\text{OPT}$  is the optimal solution for the instance  $\mathcal{I}$ ). Recall that  $L_c$  and  $\text{OPT}_c$  are the distances from the *original* position of  $c$  to  $L$  and  $\text{OPT}$ , but  $c$  may have been moved to  $L(c)$  and  $B_D$  is the set of badly cut facilities of  $L$  w.r.t  $\mathcal{D}$ .

**Lemma III.2.** *Let  $\mathcal{I}$  be an instance of Facility Location with a randomized decomposition  $\mathcal{D}$ , and  $L$  be a solution for  $\mathcal{I}$ , such that  $\mathcal{I}_D$  has small distortion. For any client  $c$  in  $\mathcal{I}_D$ , let  $\text{OPT}'(c)$  be the closest facility to  $c$  in  $\text{OPT}'$ . Then  $c$  and  $\text{OPT}'(c)$  are separated in  $\mathcal{D}$  at level at most  $\log(5(L_c + \text{OPT}_c)) + \tau(\varepsilon, d)$ .*

*Proof:* Let  $c$  be a client. To find the level at which  $c$  and  $\text{OPT}'(c)$  are separated, we distinguish between two cases: either  $c$  in  $\mathcal{I}$  is badly cut w.r.t  $\mathcal{D}$ , or not.

If  $c$  is badly cut, then it is now located at  $L(c)$  in the instance  $\mathcal{I}_D$ . In that case, either:

- 1)  $L(c)$  is also badly cut, and therefore  $L(c) \in B_D \subset \text{OPT}'$  and so  $\text{OPT}'(c) = L(c)$ . It follows that  $c$  and  $\text{OPT}'(c)$  are never separated.
- 2)  $L(c)$  is not badly cut. Then  $\text{dist}(c, \text{OPT}'(c)) \leq \text{OPT}_{L(c)}$ . We bound the level at which  $c$  and  $\text{OPT}'(c)$  are separated. Since  $L(c)$  is not badly cut, [Definition II.4](#) implies that  $L(c)$  and  $\text{OPT}(L(c))$  are cut at a level at most  $\log(\text{OPT}_{L(c)}) + \tau(\varepsilon, d)$ . By triangle inequality,  $\text{OPT}_{L(c)} = \text{dist}(L(c), \text{OPT}(L(c))) \leq L_c + \text{OPT}_c$ , and thus  $c$  and  $\text{OPT}'(c)$  are also separated at level at most  $\log(L_c + \text{OPT}_c) + \tau(\varepsilon, d)$ .

We now turn to the case where  $c$  is not badly cut. In which case,  $c$  is not moved to  $L_c$  and the balls  $\beta(c, 2^i)$

with  $2^i \in [\varepsilon L_c, L_c/\varepsilon]$  are not badly cut. We make a case distinction according to  $\text{OPT}_c$ .

- 1) If  $L_c \leq \varepsilon \text{OPT}_c$ , then we have the following. If  $L(c)$  is badly cut,  $L(c)$  is open and therefore  $\text{OPT}'_c = L_c$ . Moreover, since  $c$  is not badly cut the ball  $\beta(c, L_c)$  is cut at level at most  $\log(L_c) + \tau(\varepsilon, d)$ . Therefore  $c$  and  $\text{OPT}'(c)$  are separated at level at most  $\log(L_c) + \tau(\varepsilon, d)$ . In the case where  $L(c)$  is not badly cut, both  $c$  and  $\text{OPT}'(c)$  lie in the ring centered at  $L(c)$  and of diameter  $2\text{OPT}_{L(c)}$ . Indeed,

$$\begin{aligned} \text{dist}(c, L(c)) &\leq \varepsilon \text{dist}(c, \text{OPT}(c)) \leq \varepsilon \text{dist}(c, \text{OPT}(L(c))) \\ &\leq \varepsilon \text{dist}(c, L(c)) + \varepsilon \text{dist}(L(c), \text{OPT}(L(c))) \end{aligned}$$

And therefore, for any  $\varepsilon \leq 2/3$ ,

$$\text{dist}(c, L(c)) \leq \frac{\varepsilon}{1-\varepsilon} \text{OPT}_{L(c)} \leq 2\text{OPT}_{L(c)}.$$

On the other hand,

$$\begin{aligned} \text{dist}(\text{OPT}'(c), L(c)) &\leq \text{dist}(\text{OPT}'(c), c) + \text{dist}(c, L(c)) \\ &\leq \text{dist}(c, \text{OPT}(L(c))) + \text{dist}(c, L(c)) \\ &\leq 2\text{dist}(c, L(c)) + \text{dist}(L(c), \text{OPT}(L(c))) \\ &\leq \left(1 + \frac{2\varepsilon}{1-\varepsilon}\right) \text{OPT}_{L(c)}, \end{aligned}$$

which is smaller than  $2\text{OPT}_{L(c)}$  for any  $\varepsilon \leq 1/3$ . Hence we have  $c, \text{OPT}'(c) \in \beta(L(c), 2\text{OPT}_{L(c)})$ . To apply the definition of badly cut, we need to consider rings with radius power of 2: let us therefore consider  $i$  such that  $2\text{OPT}_{L(c)} \in (2^{i-1}, 2^i]$  (note that  $2^i \leq 4\text{OPT}_{L(c)}$ ). Since  $L(c)$  is not badly cut, this ring is not cut by a too high level part either. Therefore  $c$  and  $\text{OPT}'(c)$  are separated at level at most  $i + \tau(\varepsilon, d)$ , which is at most  $\log(4\text{OPT}_{L(c)}) + \tau(\varepsilon, d)$ .

Now, since  $\text{OPT}_{L(c)} \leq \text{dist}(L(c), \text{OPT}(c)) \leq \text{dist}(L(c), c) + \text{dist}(c, \text{OPT}(c)) \leq (1+\varepsilon)\text{OPT}_c$ , we have that  $\log(4\text{OPT}_{L(c)}) \leq \log(5\text{OPT}_c)$ , and hence  $c$  and  $\text{OPT}'(c)$  are separated at level at most  $\log(5\text{OPT}_c) + \tau(\varepsilon, d)$ .

- 2) If  $L_c \geq \text{OPT}_c/\varepsilon$  then, since  $c$  is not badly cut, the ball centered at  $c$  and of radius  $\varepsilon L_c$  is not badly cut. Since we have  $\text{dist}(c, \text{OPT}'(c)) \leq \text{OPT}_c \leq \varepsilon L_c$ ,  $c$  and  $\text{OPT}'(c)$  lie in the ball  $\beta(c, \varepsilon L_c)$  and are thus cut at level at most  $\log(\varepsilon L_c) + \tau(\varepsilon, d)$ .
- 3) If  $\varepsilon L_c \leq \text{OPT}_c \leq L_c/\varepsilon$ , then since  $c$  is not badly cut the ball  $\beta(c, \text{OPT}(c))$  is cut at level at most  $\log(2\text{OPT}_c) + \tau(\varepsilon, d)$ . Moreover,  $\text{OPT}'(c)$  lies in this ball.

This concludes the proof. ■

A *path* between two nodes  $u$  and  $v$  is a sequence of nodes  $w_1, \dots, w_k$  with  $u = w_1, v = w_k$ , and its length is  $\sum \text{dist}(w_j, w_{j+1})$ . A solution to the problem can therefore be seen as a set of facility, together with a path for each client that connects it to a facility. We say a path is *portal-respecting* if it enters and leaves parts of the decomposition  $\mathcal{D}$  only at portals. More precisely, for every pair  $w_j, w_{j+1}$  of

the sequence, if  $w_j$  and  $w_{j+1}$  lie in different parts of some level  $i$ , then these nodes are also portals at this level (note that such a path is guaranteed to exist, since we assume that portals are nested; cf. Lemma II.3). We define a *portal-respecting* solution to be a solutions such that each path from a client to its closest facility in the solution is portal-respecting.

The dynamic program will compute an optimal portal-respecting solution. Therefore, we need to prove that the optimal portal-respecting solution is close to the optimal solution. Let  $u$  and  $v$  be two vertices separated at level  $i$  by the decomposition  $\mathcal{D}$ . We note a property of the decomposition that will simplify our calculations. For the path between  $u$  and  $v$  to be portal-respecting, there needs to be a detour at every level below  $i$ , with an error of at most  $\sum_{j \leq i} O(\rho 2^j) \leq O(\rho 2^{i+2})$ . This error comes from the preciseness property in Lemma II.3. In the remainder of the paper, we will thus bound the total error incurred across all levels by  $O(\rho \cdot 2^i)$ , where  $i$  is the level at which  $u$  and  $v$  are separated. We let  $\rho = \varepsilon 2^{-\tau(\varepsilon, d)}$ , and for a solution  $S$  define  $B(S) := \sum_{c, i : c \text{ and } S(c) \text{ cut at level } i} \varepsilon 2^i$ . One can see  $B(S)$  as a *budget*, given by the fact that vertices are not badly cut.

**Lemma III.3.** *Given an instance  $\mathcal{I}$  and a solution  $L$ , it holds with probability  $1 - \varepsilon$  (over  $\mathcal{D}$ ) that there exists a portal-respecting solution  $S$  in  $\mathcal{I}_{\mathcal{D}}$  is such that  $\text{cost}_{\mathcal{I}_{\mathcal{D}}}(S) + B(S) = (1 + O(\varepsilon))\text{cost}_{\mathcal{I}}(\text{OPT}) + O(\varepsilon \text{cost}_{\mathcal{I}}(L))$ .*

*Proof:* From Lemma III.1, with probability  $1 - \varepsilon$  it holds that the instance  $\mathcal{I}_{\mathcal{D}}$  has small distortion. We condition now on this event. Consider solution  $\text{OPT}'$ . Since  $\mathcal{I}_{\mathcal{D}}$  has small distortion, we have that the facility cost of  $\text{OPT}'$  is at most the facility cost of  $\text{OPT}$  plus  $\varepsilon \text{cost}(L)$ . Furthermore, again since  $\mathcal{I}_{\mathcal{D}}$  has small distortion we have that  $\text{cost}_{\mathcal{I}_{\mathcal{D}}}(\text{OPT}') \leq (1 + O(\varepsilon))\text{cost}_{\mathcal{I}}(\text{OPT}) + O(\varepsilon \text{cost}_{\mathcal{I}}(L))$ .

We now bound the cost of making  $\text{OPT}'$  portal respecting by applying Lemma III.2. Since each client  $c$  of  $\mathcal{I}_{\mathcal{D}}$  is separated from  $\text{OPT}'(c)$  at level at most  $\log(5(L_c + \text{OPT}_c)) + \tau(\varepsilon, d)$ , we have that the detour for making the assignment of  $c$  to  $\text{OPT}'(c)$  portal-respecting is at most  $\rho 2^{\tau(\varepsilon, d)} 5(L_c + \text{OPT}_c)$ . Choosing  $\rho = \varepsilon 2^{-\tau(\varepsilon, d)}$  ensures that the detour is at most  $O(\varepsilon(L_c + \text{OPT}_c))$ . This also bounds  $B(\text{OPT}') \leq O(\varepsilon)(\text{cost}_{\mathcal{I}}(L) + \text{cost}_{\mathcal{I}}(\text{OPT}))$ .

Therefore, taking  $S = \text{OPT}'$  ensures that

$$\begin{aligned} \text{cost}_{\mathcal{I}_{\mathcal{D}}}(S) &\leq \text{cost}_{\mathcal{I}_{\mathcal{D}}}(\text{OPT}') + 20\varepsilon(\text{cost}_{\mathcal{I}}(\text{OPT}) + \text{cost}_{\mathcal{I}}(L)) \\ &\leq (1 + O(\varepsilon))\text{cost}_{\mathcal{I}}(\text{OPT}) + O(\varepsilon \text{cost}_{\mathcal{I}}(L)) \end{aligned}$$

■

### C. The Algorithm

Using Lemmas A.1 and A.2, we can assume that the aspect-ratio of the instance is  $O(n^5/\varepsilon)$ . Our algorithm starts by computing a constant-factor approximation  $L$ , using Meyerson's algorithm [?]. It then computes a hierarchical

decomposition  $\mathcal{D}$ , as explained in the Section II-B, with parameter  $\rho = \varepsilon 2^{-\tau(\varepsilon, d)}$ .

Given  $L$  and the decomposition  $\mathcal{D}$ , our algorithm finds all the badly cut clients as follows. For each client  $c$ , to determine whether  $c$  is badly cut or not, only  $O(\log(1/\varepsilon))$  balls have to be considered, namely the balls centered at  $c$  and with exponentially growing radius in  $[\varepsilon L_c, L_c/\varepsilon]$ . For each such ball  $\beta$ , the algorithm checks whether the decomposition cuts  $\beta$  at a level that is too high, making  $c$  badly cut. This can be done efficiently by verifying whether  $c$  is at distance smaller than  $L_c/\varepsilon$  to such a part of too high level. Thus, the algorithm finds all the badly cut clients in near-linear time.

The next step of the algorithm is to compute instance  $\mathcal{I}_{\mathcal{D}}$  by moving every badly cut client  $c$  to its facility in  $L$ . This can also be done in linear time.

*A first attempt at a dynamic program.:* We now turn to the description of the dynamic program (DP) for obtaining the best portal-respecting solution of  $\mathcal{I}_{\mathcal{D}}$ . This is the standard dynamic program for Facility Location and we only describe it for the sake of completeness, the reader familiar with this can skip to the analysis.

There is a table entry for each part of the decomposition, and two vectors of length  $n_p$ , where  $n_p$  is the number of portals in the part (we call such a triplet a configuration). Each configuration given by a part  $R$  and vectors  $\langle \ell_1, \dots, \ell_{n_p} \rangle$  and  $\langle s_1, \dots, s_{n_p} \rangle$  (called the *portal parameters*), encodes a possible interface between part  $R$  and a solution for which the  $i$ th portal has approximate distance  $\ell_i$  to the closest facility inside of  $R$ , and approximate distance  $s_i$  to its closest facility outside of  $R$ . The value stored for such a configuration in a table entry is the minimal cost for a solution with facilities respecting the constraints induced by the vectors on the distances between the solution and the portals inside the part (as described below).

To fill the table, we use a dynamic program following the lines of Arora et al. [?] or Kolliopoulos and Rao [?]. If a part has no descendant (meaning the part contains a single point), computing the solution given the configuration is straightforward: either a center is opened on this point or not, and it is easy to check the consistency with the configuration, where only the distances to portals inside the part need to be verified. At a higher level of the decomposition, a solution is simply obtained by going over all the sets of parameter values for all the children parts. It is immediate to see whether sets of parameter values for the children can lead to a consistent solution:

- for each portal  $p_1$  of the parent part, there must be one portal  $p_2$  of a child part such that the distance from  $p_1$  to a center inside the part prescribed by the configuration corresponds to  $\text{dist}(p_1, p_2)$  plus the distance from  $p_2$  to a center inside the child part;
- for each portal  $p_2$  of a child part, there must exist either:
  - a portal  $p_1$  of the parent part such that the distance

from  $p_2$  to a center outside its part prescribed by the configuration is  $\text{dist}(p_1, p_2)$  plus the distance from  $p_1$  to a center outside of the part,

- or a portal  $p_1$  of another child part such that this distance is  $\text{dist}(p_1, p_2)$  plus the distance from  $p_1$  to a center inside the child part.

Recall that the aspect ratio is  $n^{O(1)}$ , and so this dynamic program has a complexity polylogarithmic in  $n$ , since there are  $O(\log n)$  possible values for a rounded distance. However, using the budget given by [Lemma III.3](#), one can shave off the logarithmic factors.

*A faster dynamic program.:* We now describe a faster dynamic program. Consider a level where the diameter of the parts is say  $\Delta$ . Each configuration is again given by a part  $R$  and portal parameters  $\langle \ell_1, \dots, \ell_{n_p} \rangle$  and  $\langle s_1, \dots, s_{n_p} \rangle$ , but with the restriction that  $\ell_i$  and  $s_i$  are multiples of  $\varepsilon\Delta$  in the range  $[0, \Delta/\varepsilon + \Delta]$ . A flag is additionally attached to the configuration (whose meaning will be explained shortly).

We sketch here the intuition behind this restriction. Since the diameter of the part is  $\Delta$  we can afford a detour of  $\varepsilon\Delta$ , that is taken into account in the budget  $B(S)$ . Hence, distances can be rounded to the closest multiple of  $\varepsilon\Delta$ .

Now, suppose that the closest facility outside the part is at distance greater than  $\Delta/\varepsilon$ , and that there is no facility inside the part. Then, since the diameter is  $\Delta$ , up to losing an additive  $\varepsilon\text{OPT}$  in the cost of the solution computed, we may assume that all the points of the part are assigned to the same facility. So the algorithm is not required to have the precise distance to the closest center outside the part, and it uses the flag to reflect that it is in this regime. We can then treat this whole part as a single client (weighted by the number of clients inside the part) to be considered at higher levels. Assuming that the closest facility is at distance less than  $\Delta/\varepsilon$ , we have that for any portal of the part the closest facility is at distance at most  $\Delta/\varepsilon + \Delta$  (since  $\Delta$  is the diameter of the part).

On the other hand, if there is some facility inside the part and the closest facility outside the part is at distance at least  $\Delta/\varepsilon$ , then each client of the part should be served by a facility inside the part in any optimal assignment. Thus it is not necessary that the algorithm iterates over configurations where the distances outside the part are more than  $\Delta/\varepsilon$ : it is enough to do it once and use the answer for all other queries.

*Analysis – Proof of [Theorem I.3](#).:* The two following lemmas show that the solution computed by this algorithm is a near-optimal one, and that the complexity is near-linear: this proves [Theorem I.3](#).

**Lemma III.4.** *Let  $S$  be as in [Lemma III.3](#). The algorithm computes a solution  $S^*$  with cost at most  $\text{cost}_{\mathcal{I}_D}(S^*) \leq (1 + O(\varepsilon))\text{cost}_{\mathcal{I}_D}(S) + B(S)$ .*

*Proof:* We show that the solution  $S$  can be adapted to a configuration of the DP with an extra cost  $B(S)$ . For this,

let  $c$  be a client served by a facility  $S(c)$ , and let  $p_1, \dots, p_k$  be the portal-respecting path from  $c$  to  $S(c)$ ,  $p_i$  being a portal at level  $l_i$ , with  $p_1 = c$  and  $p_k = S(c)$ . The cost of  $c$  in  $S$  is therefore  $\sum d(p_i, p_{i+1})$ .

The distance between  $c$  and  $S(c)$  is approximated at several places of the DP:

- When  $d(p_i, S(c)) \leq 2^{l_i}/\varepsilon + 2^{l_i}$ , the distance between  $p_i$  and  $S(c)$  is rounded to the closest multiple of  $\varepsilon 2^{l_i}$ , incurring a cost difference of  $\varepsilon 2^{l_i}$ .
- When  $d(p_i, S(c)) \geq 2^{l_i}/\varepsilon + 2^{l_i}$ , the whole part is contracted and served by a single facility at distance at least  $2^{l_i}/\varepsilon$ . The cost difference for client  $c$  is therefore  $2^{l_i} \leq \varepsilon d(p_i, S(c))$ . Since the diameters of the parts are geometrically increasing, the total cost difference for such operations is bounded by  $2\varepsilon d(p_j, S(c))$ , where  $l_j$  is the highest level where  $d(p_j, S(c)) \geq 2^{l_j}/\varepsilon + 2^{l_j}$ . This cost verifies  $2\varepsilon d(p_j, S(c)) \leq 2\varepsilon \sum d(p_i, p_{i+1})$ .

Hence, summing over all clients, the additional cost incurred by the DP compared to  $S$  is at most  $B(S) + 2\varepsilon \text{cost}_{\mathcal{I}_D}(S)$ . Since it computes a solution with minimal cost, it holds that  $\text{cost}_{\mathcal{I}_D}(S^*) \leq (1 + 2\varepsilon)\text{cost}_{\mathcal{I}_D}(S) + B(S)$ . ■

**Corollary III.5.** *Let  $S^*$  be the solution computed by the algorithm. With probability  $1 - \varepsilon$ , it holds that  $\text{cost}_{\mathcal{I}}(S^*) = (1 + O(\varepsilon))\text{cost}_{\mathcal{I}}(\text{OPT}) + O(\varepsilon \text{cost}_{\mathcal{I}}(L))$*

*Proof:* [Lemma III.3](#) ensures that, with probability  $1 - \varepsilon$ , the cost of  $S$  is at most  $(1 + O(\varepsilon))\text{cost}_{\mathcal{I}}(\text{OPT}) + O(\varepsilon \text{cost}_{\mathcal{I}}(L))$ . Since  $L$  is a constant-factor approximation, this cost turns out to be  $(1 + O(\varepsilon))\text{cost}_{\mathcal{I}}(\text{OPT})$ . Combining this with [Lemma III.4](#) concludes the proof:

$$\begin{aligned} \text{cost}_{\mathcal{I}}(S^*) &= (1 + O(\varepsilon))\text{cost}_{\mathcal{I}_D}(S^*) \\ &= (1 + O(\varepsilon))(\text{cost}_{\mathcal{I}_D}(S) + B(S)) \\ &\leq (1 + O(\varepsilon))\text{cost}_{\mathcal{I}}(\text{OPT}) + O(\varepsilon \text{cost}_{\mathcal{I}}(L)) \end{aligned}$$

■

**Lemma III.6.** *This algorithm has complexity  $(\frac{1}{\varepsilon})^{2^{O(d \log d)}/\varepsilon} \cdot n + 2^{O(d)} n \log n$ .*

*Proof:* The preprocessing step (computing  $L$ , the hierarchical decomposition  $\mathcal{D}$ , and the instance  $\mathcal{I}_D$ ) has a running time  $O(n \log n)$ , as all the steps can be done with this complexity: a fast implementation of Meyerson’s algorithm [?] tailored to graphs can compute  $L$  in time  $O(m \log n)$ . Using it on the spanner computed with [?] gives a  $O(1)$ -approximation in time  $O(n \log n)$ . As explained earlier, the hierarchical decomposition  $\mathcal{D}$  and the instance  $\mathcal{I}_D$  can also be computed with this complexity.

The DP has a linear time complexity: in a part of diameter  $\Delta$ , the portal set is a  $\varepsilon 2^{-\tau(\varepsilon, d)} \Delta$ -net, and hence has size  $2^{d \log(2^{\tau(\varepsilon, d)}/\varepsilon)}$  by [Lemma II.1](#). Since  $\tau(\varepsilon, d) = O(\log d) + 2 \log \frac{(p+1)^p}{\varepsilon^{p+2}}$ , this number can be simplified to  $2^{O(d \log(d))}/\varepsilon$ . Since each portal stores a distance that can take only

$1/\varepsilon^2$  values, there are at most  $E = (1/\varepsilon^2)^{2^{O(d^2)}/\varepsilon} = (1/\varepsilon)^{2^{O(d^2)}/\varepsilon}$  possible table entries for a given part.

To fill the table, notice that a part has at most  $2^{O(d)}$  children, due to the properties of the hierarchical decomposition. Going over all the sets of parameter values for all the children parts therefore takes time  $E^{2^{O(d)}} = (1/\varepsilon)^{2^{O(d^2)}/\varepsilon}$ . This dominates the complexity of the dynamic program, which is therefore  $n(1/\varepsilon)^{2^{O(d^2)}/\varepsilon}$ .

The total complexity of the algorithm is thus

$$\left(\frac{1}{\varepsilon}\right)^{2^{O(d^2)}/\varepsilon} \cdot n + 2^{O(d)} n \log n$$

■

#### IV. THE $k$ -MEDIAN AND $k$ -MEANS PROBLEMS

We aim at using the same approach as for Facility Location. We focus the presentation on  $k$ -Median, and only later show how to adapt the proof for  $k$ -Means.

We will work with the more general version of  $k$ -Median as defined in Section II-D, where the instance consists of a set of clients  $C$ , a set of candidate centers  $F$ , an integer  $k$ , and a function  $\chi : C \mapsto \{1, \dots, n\}$  and the goal is to minimize  $\sum_{c \in C} \chi(c) \cdot \min_{f \in S} \text{dist}(c, f)$ .

The road-map is as for Facility Location: we show in Lemma IV.2 that an instance  $\mathcal{I}_{\mathcal{D}}$  has a *small distortion* with good probability, and then in Lemma IV.5 that if an instance has small distortion then there exists a near-optimal portal-respecting solution. We finally present a dynamic program that compute such a solution.

A key ingredient of the proof for Facility Location was our ability to add all badly-cut facilities to the solution  $\text{OPT}'$ . As the number of facilities is fixed, this is not directly possible in the case of  $k$ -Median and  $k$ -Means. Hence, the first step of our proof is to show that one can make some room in  $\text{OPT}$ , by removing a few centers without increasing too much the cost.

##### A. Towards a Structured Near-Optimal Solution

**\*\*\* (Vincent) What lemma are we trying to prove in this section? It is very unclear what is going on in here... \*\*\*** **\*\*\* (David) better with previous paragraph? \*\*\***

**\*\*\*** Let  $\text{OPT}$  be an optimal solution to  $\mathcal{I}$ . We consider the mapping of the facilities of  $\text{OPT}$  to  $L$  defined as follows: for any  $f \in \text{OPT}$ , let  $L(f)$  denote the facility of  $L$  that is the closest to  $f$ . Recall that for a client  $c$ ,  $L(c)$  is the facility serving  $c$  in  $L$ .

For any facility  $\ell$  of  $L$ , define  $\psi(\ell)$  to be the set of facilities of  $\text{OPT}$  that are mapped to  $\ell$ , namely,  $\psi(\ell) = \{f \in \text{OPT} \mid L(f) = \ell\}$ . Define  $L^1$  to be the set of facilities  $\ell$  of  $L$  for which there exists a unique  $f \in \text{OPT}$  such that  $L(f) = \ell$ , namely  $L^1 = \{\ell \mid |\psi(\ell)| = 1\}$ . Let  $L^0 = \{\ell \mid |\psi(\ell)| = 0\}$ , and  $L^{\geq 2} = L - (L^1 \cup L^0)$ . Similarly, define  $\text{OPT}^1 = \{f \in \text{OPT} \mid L(f) \in L^1\}$

and  $\text{OPT}^{\geq 2} = \{f \in \text{OPT} \mid L(f) \in L^{\geq 2}\}$ . Note that  $|\text{OPT}^{\geq 2}| = |L^0| + |L^{\geq 2}|$ , since  $|\text{OPT}^1| = |L^1|$  and, w.l.o.g.,  $|\text{OPT}| = |L| = k$ .

The construction of a structured near-optimal solution is made in 3 steps. The first one defines a solution  $\text{OPT}'$  as follows. Start with  $\text{OPT}' = \text{OPT}$ .

- **Step 1.** Among the facilities of  $\text{OPT}^{\geq 2}$  that are not the closest of their corresponding facility in  $L^{\geq 2}$ , remove from  $\text{OPT}'$  the subset  $\widehat{\mathcal{H}}$  of size  $\lfloor \varepsilon \cdot |\text{OPT}^{\geq 2}|/2 \rfloor$  that yields the smallest cost increase.

This step makes room to add the badly cut facilities without violating the constraint on the maximum number of centers, while at the same time ensures that  $S^*$  has near-optimal cost, as the following lemma shows.

**Lemma IV.1.** *After step 1,  $\text{OPT}'$  has cost  $(1 + O(\varepsilon))\text{cost}(\text{OPT}) + O(\varepsilon)\text{cost}(L)$*

*Proof:* We show that the cost increase is at most  $O(\varepsilon)(\text{cost}(\text{OPT}) + \text{cost}(L))$ .

Let  $\mathcal{H} \subseteq \text{OPT}^{\geq 2}$  be the set of facility of  $\text{OPT}^{\geq 2}$  that are not the closest to their corresponding facility in  $L^{\geq 2}$ , i.e.,  $f \in \mathcal{H}$  if and only if  $f \in \psi(\ell)$  for some  $\ell \in L^{\geq 2}$  and  $\text{dist}(f, \ell) > \min_{f' \in \psi(\ell)} \text{dist}(f', \ell)$  (breaking ties arbitrarily). The only elements in  $\text{OPT}^{\geq 2} - \mathcal{H}$  are the ones closest to their corresponding facilities. Hence for every facility of  $L^{\geq 2}$  such that  $|\psi(f)| \geq 2$  there is therefore exactly one facility in  $\text{OPT}^{\geq 2} - \mathcal{H}$ , and at least two in  $\text{OPT}^{\geq 2}$ ; and if  $|\psi(f)| = 0$  then  $f$  does not correspond to any facility at all in  $\text{OPT}^{\geq 2}$ . Therefore  $|\mathcal{H}| \geq |\text{OPT}^{\geq 2}|/2$ .

We claim that for a client  $c$  served by  $f \in \mathcal{H}$  in the optimum solution  $\text{OPT}$ , i.e.,  $f = \text{OPT}(c)$ , the detour entailed by the deletion of  $f$  is  $O(\text{OPT}_c + L_c)$ . Indeed, let  $f'$  be the facility of  $\text{OPT}$  that is closest to  $L(f)$ , and recall that  $L(c)$  is the facility that serves  $c$  in the solution  $L$ . Since  $f' \notin \mathcal{H}$ , the cost to serve  $c$  after the removal of  $f$  is at most  $\text{dist}(c, f')$ , which can be bounded by  $\text{dist}(c, f') \leq \text{dist}(c, f) + \text{dist}(f, L(f)) + \text{dist}(L(f), f')$ . But by definition of  $f'$ ,  $\text{dist}(f', L(f)) \leq \text{dist}(L(f), f)$ , and by definition of the function  $L$  we have  $\text{dist}(L(f), f) \leq \text{dist}(L(c), f)$ , so that  $\text{dist}(c, f') \leq \text{dist}(c, f) + 2\text{dist}(f, L(c))$ . Using the triangle inequality finally gives  $\text{dist}(c, f') \leq 3\text{dist}(c, f) + 2\text{dist}(c, L(c))$  which is  $O(\text{OPT}_c + L_c)$ . For a facility  $f$  of  $\text{OPT}$ , we denote  $C(f)$  the set of clients served by  $f$ , i.e.  $C(f) = \{c \in C \mid \text{OPT}(c) = f\}$ . The total cost incurred by the removal of  $f$  is then  $O(\text{cost}_{\text{OPT}}(C(f)) + \text{cost}_L(C(f)))$ .

Recall that in Step 1 we remove the set  $\widehat{\mathcal{H}}$  of size  $\lfloor \varepsilon |\text{OPT}^{\geq 2}| \rfloor$  from  $\mathcal{H}$ , such that  $\widehat{\mathcal{H}}$  minimizes the cost increase. We use an averaging argument to bound the cost increase: the sum among all facilities  $f \in \widehat{\mathcal{H}}$  of the cost of removing the facility  $f$  is less than  $O(\text{cost}(\text{OPT}) + \text{cost}(L))$ , and  $|\widehat{\mathcal{H}}| = O(1/\varepsilon) \cdot \lfloor \varepsilon |\text{OPT}^{\geq 2}| \rfloor$ . Therefore removing  $\widehat{\mathcal{H}}$  increases the cost by at most  $O(\varepsilon)(\text{cost}(\text{OPT}) + \text{cost}(L))$ , so that Step 1 is not too expensive. ■

← 2

← 3



We can therefore use this solution as a proxy for the optimal solution, and henceforth we will denote this solution by OPT. In particular, the badly cut facilities are defined for this solution and not the original OPT.

### B. Structural Lemma

As in Section III, the algorithm computes a randomized hierarchical decomposition  $\mathcal{D}$ , and transforms the instance of the problem. Every badly cut client  $c$  is moved to  $L(c)$ , namely, there is no more client at  $c$  and we add an extra client at  $L(c)$ . Again, we let  $\mathcal{I}_{\mathcal{D}}$  denote the resulting instance and note that  $\mathcal{I}_{\mathcal{D}}$  is a random variable that depends on the randomness of  $\mathcal{D}$ .

Moreover, as for Facility Location, we let  $B_{\mathcal{D}}$  be the set of badly cut centers of  $L$ . We call  $\text{cost}_{\mathcal{I}}(S)$  the cost of a solution  $S$  in the original instance  $\mathcal{I}$ , and  $\text{cost}_{\mathcal{I}_{\mathcal{D}}}(S)$  its cost in  $\mathcal{I}_{\mathcal{D}}$ . We let  $\nu_{\mathcal{I}_{\mathcal{D}}} = \max_{\text{solution } S} (\text{cost}_{\mathcal{I}}(S) - (1 + 3\varepsilon)\text{cost}_{\mathcal{I}_{\mathcal{D}}}(S), (1 - 3\varepsilon)\text{cost}_{\mathcal{I}_{\mathcal{D}}}(S) - \text{cost}_{\mathcal{I}}(S))$ . We say that an instance  $\mathcal{I}_{\mathcal{D}}$  has *small distortion* if  $\nu_{\mathcal{I}_{\mathcal{D}}} \leq \varepsilon \text{cost}(L)$  and there exists a solution  $S$  that contains  $B_{\mathcal{D}}$  with  $\text{cost}_{\mathcal{I}}(S) \leq (1 + \varepsilon)\text{cost}_{\mathcal{I}}(\text{OPT}) + \varepsilon \text{cost}_{\mathcal{I}}(L)$ .

We go on with the next two steps of our construction, defining a solution  $S^*$ . We start with  $S^* = \text{OPT}$  (and recall our convention  $\text{OPT} = \text{OPT}'$ ).

- **Step 2.** For each badly-cut facility  $f \in L$  for which  $\psi(f) \neq \emptyset$ , let  $f' \in \psi(f)$  be the closest to  $f$ . Replace  $f'$  by  $f$  in  $S^*$ .
- **Step 3.** Add all badly cut facility  $f'$  of  $L^0$  to  $S^*$ .

We show next that  $S^*$  satisfies the conditions for  $\mathcal{I}_{\mathcal{D}}$  to have small distortion with good probability.

**Lemma IV.2.** *The probability that  $\mathcal{I}_{\mathcal{D}}$  has small distortion is at least  $1 - \varepsilon$ .*

*Proof:* The proof that  $\nu_{\mathcal{I}_{\mathcal{D}}} \leq \varepsilon \text{cost}(L)$  with probability at least  $1 - \varepsilon/2$  is identical to the one in Lemma III.1. We thus turn to bound the probability that solution  $S^*$  satisfies the cardinality and cost requirements. Our goal is to show that this happens with probability at least  $1 - \varepsilon/2$ . Then, taking a union bound over the probabilities of failure yields the proposition.

By definition, we have that  $S^*$  contains  $B_{\mathcal{D}}$ . We prove in the two following claims some properties on  $S^*$ .

**Claim IV.3.** *With probability at least  $1 - \varepsilon/4$ , the set  $S^*$  is an admissible solution, i.e.,  $|S^*| \leq k$ .*

*Proof:* We let  $b$  be the number of facilities of  $L^0$  that are badly cut. By Lemma II.5, we have that  $\mathbb{E}[b] \leq \varepsilon^2 |L|/4$ . By Markov's inequality, the probability that  $b$  is such that  $b > \varepsilon |L^0|/2$  is at most  $\varepsilon/2$ . Now, condition on the event that  $b \leq \varepsilon |L^0|/2$ . Since  $|L^0| + |L^{\geq 2}| = |\text{OPT}^2|$ , we have that  $b \leq \varepsilon |\text{OPT}^2|/2$ . Moreover, the three steps converting OPT into  $S^*$  ensure that  $|S^*| \leq k + b - \varepsilon \lfloor |\text{OPT}^2|/2 \rfloor$ . Combining the two inequalities gives  $|S^*| \leq k$ . ■

**Claim IV.4.** *With probability at least  $1 - \varepsilon/4$ ,  $\text{cost}(S^*) \leq (1 + O(\varepsilon))\text{cost}(\text{OPT}) + O(\varepsilon \cdot \text{cost}(L))$*

*Proof:* We showed in Lemma IV.1 that the cost increase due to Step 1 is at most  $O(\varepsilon)(\text{cost}(\text{OPT}) + \text{cost}(L))$ . We show now that Step 2 leads to a cost increase of  $O(\varepsilon \cdot (\text{cost}(\text{OPT}) + \text{cost}(L)))$  with good probability. For that, let  $\text{OPT}_{\text{close}} := \{f \in \text{OPT} : f \text{ is the closest facility to } L(f)\}$ . We show that the cost of replacing all  $f \in \text{OPT}_{\text{close}}$  by  $L(f) \in L$  is  $O(\text{cost}(\text{OPT}) + \text{cost}(L))$ . In order to prove that, we call the *mixed solution* the solution with facilities where every facility of  $f \in \text{OPT}_{\text{close}}$  is replaced by  $L(f)$ . Note that  $L(\text{OPT}_{\text{close}}) = L - L^0$ .

For that, let  $c$  be a client that is served in OPT by a facility  $f$  of  $\text{OPT}_{\text{close}}$ . If  $c$  is served in  $L$  by a facility of  $L - L^0$ , then the facility appears in the mixed solution and the serving cost of  $c$  is  $\text{dist}(c, L)$ . On the other hand, if  $c$  is served by a facility  $f_0$  of  $L^0$  in  $L$ , then it is possible to serve it by the  $L(f)$  that replaces  $f$  in the mixed solution. The serving cost is therefore  $\text{dist}(c, L(f)) \leq \text{dist}(c, f) + \text{dist}(f, L(f)) \leq \text{dist}(c, f) + \text{dist}(f, f_0)$ , using the definition of  $L(f)$  for the last inequality. Using again the triangle inequality, this cost is at most  $2\text{dist}(c, f) + \text{dist}(f, f_0)$ . Moreover, any client served by a facility of  $\text{OPT} - \text{OPT}_{\text{close}}$  is served by its optimal facility in the mixed solution, with cost  $\text{dist}(c, \text{OPT})$ . Hence the cost of the mixed solution is at most  $2\text{cost}(\text{OPT}) + \text{cost}(L)$ .

Moreover, the probability of replacing  $f \in \text{OPT}_{\text{close}}$  by  $L(f) \in L - L^0$  in Step 2 is the probability that  $L(f)$  is badly cut, which is  $\kappa(\varepsilon, p)$  by Lemma II.5. Finally, with linearity of expectation, the expected cost to add the badly cut facilities  $f \in L - L^0$  instead of their closest facility of OPT in Step 2 is  $O(\kappa(\varepsilon, p)(\text{cost}(\text{OPT}) + \text{cost}(L)))$ . Markov's inequality thus implies that the cost of the first step is at most  $O(\varepsilon \cdot (\text{cost}(\text{OPT}) + \text{cost}(L)))$  with probability  $1 - \frac{O(\kappa(\varepsilon, p))}{\varepsilon} \geq 1 - \varepsilon/4$ , since  $\kappa(\varepsilon, p) \leq \varepsilon^2/4$  in the case of  $k$ -Median. ■

Lemma IV.2 follows from taking a union bound over the probabilities of failure of Claim IV.3 and IV.4. ■

Condition now on  $\mathcal{I}_{\mathcal{D}}$  having small distortion, and let  $\text{OPT}'$  be the solution containing  $B_{\mathcal{D}}$  with cost  $(1 + \varepsilon)\text{cost}_{\mathcal{I}}(\text{OPT}) + \varepsilon \text{cost}_{\mathcal{I}}(L)$ . We have to prove the same structural lemma as for Facility Location, to say that there exists a portal-respecting solution with cost close to  $\text{cost}(\text{OPT}')$ .

Recall that  $L_c$  and  $\text{OPT}_c$  are the distances from the original position of  $c$  to  $L$  and OPT; but  $c$  may have been moved to  $L(c)$ . Recall also that OPT is defined after removing some centers in Step 1.

**Lemma IV.5.** *Condition on  $\mathcal{I}_{\mathcal{D}}$  having small distortion. For any client  $c$  in  $\mathcal{I}_{\mathcal{D}}$ , let  $\text{OPT}'(c)$  be the closest facility to  $c$  in  $\text{OPT}'$ . Then  $c$  and  $\text{OPT}'(c)$  are separated in  $\mathcal{D}$  at level at most  $\log(7(L_c + \text{OPT}_c)) + \tau(\varepsilon, d)$ .*

*Proof:* The proof of this lemma is very similar to the one of Lemma III.2. However, since some facilities of OPT were removed in Step 2, we need to adapt the proof carefully.



Let  $c$  be a client. If  $\text{OPT}(c)$  was removed in Step 2, it was replaced by a facility  $f$  such that  $\text{dist}(\text{OPT}(c), f) \leq \text{dist}(\text{OPT}(c), L(c))$  (because  $L(\text{OPT}(c)) = f$  means that  $f$  is the facility of  $L$  closest to  $\text{OPT}(c)$ ). Therefore

$$\text{dist}(c, f) \leq 2\text{dist}(c, \text{OPT}(c)) + \text{dist}(c, L(c)). \quad (1)$$

To find the level at which  $c$  and  $\text{OPT}'(c)$  are separated, we distinguish between two cases: either  $c$  is badly cut, or not.

If  $c$  is badly cut, then it is now located at  $L(c)$  in the instance  $\mathcal{I}_{\mathcal{D}}$ . In that case, either:

- 1)  $L(c)$  is also badly cut, and therefore  $L(c) \in \text{OPT}'$  and so  $\text{OPT}'(c) = L(c)$ . It follows that  $c$  and  $\text{OPT}'(c)$  are never separated.
- 2)  $L(c)$  is not badly cut. Then  $\text{dist}(c, \text{OPT}'(c)) = \text{dist}(L(c), \text{OPT}'(L(c)))$ .  $\text{OPT}(L(c))$  is not necessarily in  $\text{OPT}'$ : in that case, it was replaced by a facility  $f$  that verifies  $\text{dist}(c, f) \leq 2\text{dist}(c, \text{OPT}(c)) + \text{dist}(c, L(c))$ , by Property (1). Since  $\text{dist}(c, L(c)) = 0$ , we have (either if  $\text{OPT}(L(c)) \in \text{OPT}'$  or not) that  $\text{dist}(c, \text{OPT}'(c)) \leq 2\text{OPT}_{L(c)}$ .

Since  $L(c)$  is not badly cut, the ball  $\beta(L(c), 2\text{OPT}_{L(c)})$  is cut at level at most  $\log(4\text{OPT}_{L(c)}) + \tau(\varepsilon, d)$ . By triangle inequality,  $\text{OPT}_{L(c)} = \text{dist}(L(c), \text{OPT}(L(c))) \leq L_c + \text{OPT}_c$ , and thus  $c$  and  $\text{OPT}'(c)$  are also separated at level at most  $\log(4L_c + 4\text{OPT}_c) + \tau(\varepsilon, d)$ .

In the other case where  $c$  is not badly cut, all of the balls  $\beta(c, 2^i)$  where  $2^i \in [\varepsilon L_c, L_c/\varepsilon]$  are not badly cut, and  $c$  is not moved to  $L_c$ . We make a case distinction according to  $\text{OPT}_c$ .

- 1) If  $L_c \leq \varepsilon \text{OPT}_c$ , then we have the following. If  $L(c)$  is badly cut,  $L(c)$  is open and therefore  $\text{OPT}'(c) = L_c$ . Moreover, since  $c$  is not badly cut the ball  $\beta(c, L_c)$  is cut at level at most  $\log L_c + \tau(\varepsilon, d)$ . Therefore  $c$  and  $\text{OPT}'(c)$  are separated at level at most  $\log L_c + \tau(\varepsilon, d)$ . In the case where  $L(c)$  is not badly cut, both  $c$  and  $\text{OPT}'(c)$  lie in the ring centered at  $L(c)$  and of diameter  $3\text{OPT}_{L(c)}$ . Indeed,

$$\begin{aligned} \text{dist}(c, L(c)) &\leq \varepsilon \text{dist}(c, \text{OPT}(c)) \leq \varepsilon \text{dist}(c, \text{OPT}(L(c))) \\ &\leq \varepsilon \text{dist}(c, L(c)) + \varepsilon d(L(c), \text{OPT}(L(c))) \end{aligned}$$

And therefore, for any  $\varepsilon \leq 3/4$ ,

$$\text{dist}(c, L(c)) \leq \frac{\varepsilon}{1-\varepsilon} \text{OPT}_{L(c)} \leq 3\text{OPT}_{L(c)}. \text{ On the other hand,}$$

$$\begin{aligned} \text{dist}(\text{OPT}'(c), L(c)) &\leq \text{dist}(\text{OPT}'(c), c) + \text{dist}(c, L(c)) \\ &\leq 2\text{dist}(c, \text{OPT}(c)) + 2\text{dist}(c, L(c)) \\ &\quad (\text{using Property (1)}) \\ &\leq 2\text{dist}(c, \text{OPT}(L(c))) + 2\text{dist}(c, L(c)) \\ &\leq 4\text{dist}(c, L(c)) + 2\text{dist}(L(c), \text{OPT}(L(c))) \\ &\leq \left(2 + \frac{4\varepsilon}{1-\varepsilon}\right) \text{OPT}_{L(c)}, \end{aligned}$$

which is smaller than  $3\text{OPT}_{L(c)}$  for any  $\varepsilon \leq 1/2$ . Hence we have  $c, \text{OPT}'(c) \in \beta(L(c), 3\text{OPT}_{L(c)})$ . To apply the definition of badly cut, we need to consider rings with radius power of 2: let us therefore pick  $i$  such that  $3\text{OPT}_{L(c)} \in (2^{i-1}, 2^i]$  (note that  $2^i \leq 6\text{OPT}_{L(c)}$ ). Since  $L(c)$  is not badly cut, this ring is not badly cut either and thus  $c$  and  $\text{OPT}'(c)$  are separated at level at most  $i + \tau(\varepsilon, d)$ . Since  $\text{dist}(L(c), \text{OPT}(L(c))) \leq \text{dist}(L(c), \text{OPT}(c)) \leq \text{dist}(L(c), c) + \text{dist}(c, \text{OPT}(c)) \leq (1 + \varepsilon)\text{OPT}_c$ , we have that  $i \leq \log(6\text{OPT}_{L(c)}) \leq \log(7\text{OPT}_c)$ , which is smaller than what we want.

- 2) If  $L_c \geq \text{OPT}_c/\varepsilon$  then, since  $c$  is not badly cut, the ball centered at  $c$  and of radius  $\varepsilon L_c$  is not badly cut. Since we have  $\text{dist}(c, \text{OPT}'(c)) \leq 2\text{OPT}_c + L_c \leq 2L_c$ ,  $c$  and  $\text{OPT}'(c)$  lie in the ball  $\beta(c, 2L_c)$  and are thus cut at level at most  $\log(4L_c) + \tau(\varepsilon, d)$ .
- 3) If  $\varepsilon L_c \leq \text{OPT}_c \leq L_c/\varepsilon$  then, since  $c$  is not badly cut, the ball  $\beta(c, 2\text{OPT}_c + L_c)$  is cut at level at most  $\log(4\text{OPT}_c + 2L_c) + \tau(\varepsilon, d)$ . Moreover,  $\text{OPT}'(c)$  lies in this ball, which concludes the lemma. ■

Equipped with these two lemmas, we can prove the following lemma, which concludes the section:

**Lemma IV.6.** *Condition on  $\mathcal{I}_{\mathcal{D}}$  having small distortion. There exists a portal-respecting solution  $S$  in  $\mathcal{I}_{\mathcal{D}}$  is such that  $\text{cost}_{\mathcal{I}}(S) + B(S) \leq (1 + O(\varepsilon))\text{cost}_{\mathcal{I}}(\text{OPT}) + O(\varepsilon \text{cost}_{\mathcal{I}}(L))$ .*

*Proof:* The proof follows exactly the one of Lemma III.3, using the definition of small distortion,  $\text{OPT}'$ , Lemma IV.2, and Lemma IV.5. ■

*Extension to  $k$ -Means:* The adaptation to  $k$ -Means can be essentially captured by the following inequality:  $(x + y)^2 \leq 2(x^2 + y^2)$ . Indeed, taking the example of Claim IV.4, the detour  $\text{dist}(c, f') \leq 3\text{dist}(c, f) + 2\text{dist}(c, l)$  gives a cost  $\text{dist}(c, f')^2 = O(\text{dist}(c, f)^2 + \text{dist}(c, l)^2 + \text{dist}(c, f) \cdot \text{dist}(c, l)) = O(\text{dist}(c, f)^2 + \text{dist}(c, l)^2)$ . This follows through all the other lemmas, and therefore the structural lemma holds also for  $k$ -Means.

*C. Finding a Near-Optimal Portal-Respecting Solution to  $k$ -Median and  $k$ -Means*

\*\*\* (David) I kept the subsection to maintain the numbering, but to be removed \*\*\* ←4

*D. The Algorithm*

The algorithm follows the lines of the one for Facility Location, in Section III-C. It first computes a constant-factor approximation  $L$ , then the hierarchical decomposition  $\mathcal{D}$  (with parameter  $\rho = \varepsilon 2^{-\tau(\varepsilon, d)}$ ) and constructs instance  $\mathcal{I}_{\mathcal{D}}$ . A dynamic program is then used to solve efficiently the problem, providing a solution  $S$  of cost at most  $(1 + \varepsilon)\text{cost}_{\mathcal{I}}(\text{OPT})$  – conditioned on the event that the instance  $\mathcal{I}_{\mathcal{D}}$  has small distortion.

*Dynamic programming.*: The algorithm proceeds bottom up along the levels of the decomposition. We give an overview of the dynamic program which is a slightly refined version of the one presented for Facility Location in Section III-C. We make use of two additional ideas.

To avoid the dependency on  $k$  we proceed as follows. In the standard approach, a cell of the dynamic program is defined by a part of the decomposition  $\mathcal{D}$ , the portal parameters (as defined in Section III-C), and a value  $k_0 \in [k]$ . The value of an entry in the table is then the cost of the best solution that uses  $k_0$  centers, given the portal parameters.

For our dynamic program for the  $k$ -Median and  $k$ -Means problems, we define a cell of the dynamic program by a part  $B$ , the portal parameters  $\langle \ell_1, \dots, \ell_{n_p} \rangle$  and  $\langle s_1, \dots, s_{n_p} \rangle$  and a value  $c_0$  in  $[\text{cost}(L)/n; (1 + \varepsilon)\text{cost}(L)]$ . The entry of the cell is equal to the minimum number  $k_0$  of centers that need to be placed in part  $B$  in order to achieve cost at most  $c_0$ , given the portal parameters. Moreover, we only consider values for  $c_0$  that are powers of  $(1 + \varepsilon/\log n)$ . The output of the algorithm is the minimum value  $c_0$  such that the root cell has value at most  $k$  (i.e., the minimum value such that at most  $k$  centers are needed to achieve it).

The DP table can be computed the following way. For the parts that have no descendant, namely the base cases, computing the best clustering given a set of parameters can be done easily: there is at most one client in the part, and verifying that the parameter values for the centers inside the part are consistent can be done easily. At a higher level of the decomposition, a solution is obtained by going over all the sets of parameter values for all the children parts. It is immediate to see whether sets of parameter values for the children can lead to a consistent solution (similar to [?], [?]). Since there are at most  $2^{O(d)}$  children parts, this gives a running time of  $q^{2^{O(d)}}$ , where  $q$  is the total number of parameter values.

This strategy would lead to a running time of  $f(\varepsilon, d)n \log^{2^{O(d)}} n$ . We can however treat the children in order, instead of naively testing all parameter values for them. We use a classical transformation of the dynamic program, in which the first table is filled using an auxiliary dynamic program. A cell of this auxiliary DP is a value  $c_0$  in  $[\text{cost}(L)/n; (1 + \varepsilon)\text{cost}(L)]$ , a part  $C$ , one of its children  $C_i$ , and the portal parameters for the portals of  $C$  and all its children before  $C_i$  in the given order. The entry of the cell is equal to the minimum number of centers  $k_0$  that need to be placed in the children parts following  $C_i$  to achieve a cost of  $c_0$  given the portal parameters. To fill this table, one can try all possible sets of parameters for the following children, see whether they can lead to a consistent solution, and compute the minimum value among them.

*Analysis – proof of Theorem I.1 and Theorem I.2.*: We first show that the solution computed by the algorithm gives a  $(1 + O(\varepsilon))$ -approximation, and then prove the claim on the complexity.

**Lemma IV.7.** *Let  $S^*$  be the solution computed by the algorithm. With probability  $1 - \varepsilon$ , it holds that  $\text{cost}_{\mathcal{I}}(S^*) = (1 + O(\varepsilon))\text{cost}_{\mathcal{I}}(\text{OPT}) + O(\varepsilon\text{cost}_{\mathcal{I}}(L))$*

*Proof:* With probability  $1 - \varepsilon$ ,  $\mathcal{I}_{\mathcal{D}}$  has small distortion (Lemma IV.2). Following Lemma IV.6, let  $S$  be a portal-respecting solution such that  $\text{cost}_{\mathcal{I}}(S) + B(S) \leq (1 + O(\varepsilon))\text{cost}_{\mathcal{I}}(\text{OPT}) + O(\varepsilon\text{cost}_{\mathcal{I}}(L))$ .

As in Lemma III.4,  $S$  can be adapted to a configuration of the DP with a small extra cost. The cost incurred to the rounding of distances can be charged either to  $B(S)$  or is a  $O(\varepsilon)\text{cost}_{\mathcal{I}_{\mathcal{D}}}(S)$ , as in Lemma III.4. The cost to round the value  $c_0$  is a  $(1 + \varepsilon/\log n)$  factor at every level of the decomposition. Since there are  $O(\log n)$  of them, the total factor is  $(1 + \varepsilon/\log n)^{O(\log n)} = 1 + O(\varepsilon)$ . Hence, we have the following:

$$\begin{aligned} \text{cost}_{\mathcal{I}}(S^*) &= (1 + O(\varepsilon))\text{cost}_{\mathcal{I}_{\mathcal{D}}}(S^*) \\ &\quad (\text{Since } \mathcal{I}_{\mathcal{D}} \text{ has small distortion}) \\ &= (1 + O(\varepsilon))(\text{cost}_{\mathcal{I}_{\mathcal{D}}}(S) + B(S)) \\ &\quad (\text{Following the previous paragraph}) \\ &\leq (1 + O(\varepsilon))\text{cost}_{\mathcal{I}}(\text{OPT}) + O(\varepsilon\text{cost}_{\mathcal{I}}(L)) \\ &\quad (\text{By definition of } S) \end{aligned}$$

■

**Lemma IV.8.** *The running time of the DP is  $n \cdot (1/\varepsilon)^{2^{O(d^2)}/\varepsilon} \cdot \log^4 n$ .*

*Proof:* The number of cells in the auxiliary DP is given by the number of parts ( $O(n)$ ), the number of children of a part ( $2^{O(d)}$ ), the number of portal parameters ( $((1/\varepsilon)^{2^{O(d^2)}/\varepsilon})$ ) and the possible values for  $c_0$  ( $O(\log^2 n)$ ): it is therefore  $n \cdot 2^{O(d)} \cdot (1/\varepsilon)^{2^{O(d^2)}/\varepsilon} \cdot \log^2 n$ .

The complexity to fill the table adds a factor  $(1/\varepsilon)^{2^{O(d^2)}/\varepsilon} \cdot \log^2 n$ , to try all possible combination of portal parameters and value of  $c_0$ . Hence, the overall running time of the DP is  $n \cdot (1/\varepsilon)^{2^{O(d^2)}/\varepsilon} \cdot \log^4 n$ . ■

The proof of Theorem I.1 and Theorem I.2 are completed by the following lemma, which bounds the running time of the preprocessing steps.

**Lemma IV.9.** *For  $k$ -Median and  $k$ -Means, the total running time of the algorithms are respectively  $2^{O(d)}n \log^9 n + f(\varepsilon, d)n \log^4 n$  and  $2^{O(d)}n \log^{10} n + f(\varepsilon, d)n \log^5 n$ , where  $f(\varepsilon, d) = (1/\varepsilon)^{2^{O(d^2)}/\varepsilon}$ .*

*Proof:* We need to bound the running time of three steps: computing an approximation, computing the hierarchical decomposition, and running the dynamic program.

For  $k$ -Median, a constant-factor approximation can be computed in  $O(m \log^9 n) = 2^{O(d)}n \log^9 n$  time with Thorup's algorithm [?]. The split-tree decomposition can be found in  $2^{O(d)}n \log n$  time as explained in Section II. More-

over, as explained in Lemma IV.8, the dynamic program runs in time  $f(\varepsilon, d)n \log^4 n$ , ending the proof of the Theorem I.1.

Another step is required for  $k$ -Means. It is indeed not known how to find a constant-factor approximation in near-linear time. However, one can notice that a  $c$ -approximation for  $k$ -Median is an  $nc$ -approximation for  $k$ -Means, using the Cauchy-Schwarz inequality. Moreover, notice that starting from a solution  $S$ , our algorithm finds a solution with cost  $(1 + O(\varepsilon))\text{cost}(\text{OPT}) + O(\varepsilon)\text{cost}(S)$  in time  $f(\varepsilon, d)n \log^4 n$ , as for  $k$ -Median.

Repeating this algorithm  $N$  times, using in step  $i + 1$  the solution given at step  $i$ , gives thus a solution of cost  $(1 + O(\varepsilon))\text{cost}(\text{OPT}) + O(\varepsilon^N)\text{cost}(S)$ . Starting with  $\text{cost}(S) = O(n)\text{cost}(\text{OPT})$  and taking  $N = O(\log n)$  ensures to find a solution for  $k$ -Means with cost  $(1 + O(\varepsilon))\text{cost}(\text{OPT})$ . The complexity for  $k$ -Means is therefore the same as for  $k$ -Median, with an additional  $\log n$  factor. This concludes the proof of Theorem I.2. ■

## V. OTHER APPLICATIONS OF THE FRAMEWORK

Our techniques can be generalized to variants of the clustering problems where *outliers* are taken into account. We consider here two of them:  $k$ -Median with Outliers and its Lagrangian relaxation, Prize-Collecting  $k$ -Median. It can also be used to find a bicriteria approximation to  $k$ -Center.

### A. Prize-Collecting $k$ -Median

In the ‘‘prize-collecting’’ version of the problems, it is possible not to serve a client  $c$  by paying a penalty  $p_c$  (these problems are also called clustering ‘‘with penalties’’). For a solution  $S$ , we call an *outlier for  $S$*  a client that is not served by  $S$ . Formally, an instance is a quintuple  $(C, F, \text{dist}, p, k)$  where  $(C \cup F, \text{dist})$  is a metric,  $k$  is an integer and  $p : C \rightarrow \mathbb{R}^+$  the penalty function, and the goal is to find  $S \subseteq F$  and  $O \subseteq C$  such that  $|S| = k$  and  $\sum_{c \in C-O} \text{dist}(c, F) + \sum_{c \in O} p_c$  is minimized. \*\*\* (David) put it rather in the intro? Or keep the emphasize of intro on standard pbs and introduce the others only here ? \*\*\*

Looking at the Prize-Collecting  $k$ -Median problem, we aim at applying the framework from Section IV. Let  $L = (L_C, L_O)$  be an approximate solution:  $L_C$  is a set of center,  $L_O$  a set of outliers. We define *badly cut* for outliers as we did for centers: an outlier  $c$  of  $L_O$  is *badly cut w.r.t.  $\mathcal{D}$*  if there exists an integer  $i$  such that  $2^i \in [\varepsilon \text{OPT}_c, \text{OPT}_c/\varepsilon]$  and the ball  $\beta(v, 2^i)$  is cut at some level  $j$  greater than  $i + \tau(\varepsilon, d)$ , where  $\text{OPT}_c$  is the distance from  $c$  to the closest facility of the optimum solution  $\text{OPT}$ . Hence, Lemma II.5 extends directly, and the probability that an outlier in  $L_O$  is badly cut is  $\kappa(\varepsilon, p)$ .

We now turn to the previous framework, showing how to construct a near-optimal solution containing all badly-cut centers of  $L$ . For that we transfer the definitions of the mappings  $L_C, \phi$  ( $L_C$  maps a client to its closest center of  $L$ , and  $\phi(l) = \{f \in \text{OPT} \mid L(f) = l\}$ ) and of the

sets  $L^0, L^1, L^{\geq 2}, \text{OPT}^1$ , and  $\text{OPT}^{\geq 2}$ . We will show that this framework, with only a few modifications, leads to an approximation scheme for Prize-Collecting  $k$ -Median. Let  $S^* = \text{OPT}$ . As in Section IV, we start by removing a few centers from the optimal solution, without increasing the cost too much:

- **Step 1.** Among the facilities of  $\text{OPT}^{\geq 2}$  that are not the closest of their corresponding facility in  $L^{\geq 2}$ , remove from  $S^*$  the subset  $\hat{\mathcal{H}}$  of size  $\lfloor \varepsilon \cdot |\text{OPT}^{\geq 2}|/2 \rfloor$  that yields the smallest cost increase, i.e. the smallest value of  $\sum_{c \in C-L_O: \text{OPT}(c) \in \hat{\mathcal{H}}} d(c, S^* - \hat{\mathcal{H}}) + \sum_{c \in L_O: \text{OPT}(c) \in \hat{\mathcal{H}}} p_c$ .

The function minimized by  $\hat{\mathcal{H}}$  corresponds to redirecting all clients served in the local solution to a center of  $S^* - \hat{\mathcal{H}}$  and paying the penalty for clients  $c \in L_O$  such that  $\text{OPT}(c) \in \hat{\mathcal{H}}$ . Those clients are thus considered as outliers in the constructed solution.

**Lemma V.1.** *After step 1,  $S^*$  has cost  $(1 + O(\varepsilon))\text{cost}(\text{OPT}) + O(\varepsilon)\text{cost}(L)$*

*Proof sketch:* The proof is essentially the same as Lemma IV.1, with an averaging argument: for a client  $c$ , the cost of removing  $\text{OPT}(c)$  from  $S^*$  is  $O(\text{OPT}_c + L_c)$ : if  $c \notin L_O$ , the argument is the same as in Lemma IV.1, and if  $c \in L_O$  the cost is  $p_c = L_c$ . Hence the proof follows. ■

Again, we denote now by  $\text{OPT}$  this solution  $S^*$  and define the instance  $\mathcal{I}_{\mathcal{D}}$  according to this solution. Recall that  $B_{\mathcal{D}}$  is the set of badly cut centers of  $L$ , and denote  $O_{\mathcal{D}}$  the set of badly cut outliers of  $L$ . We say that an instance  $\mathcal{I}_{\mathcal{D}}$  has *small distortion* if  $\nu_{\mathcal{I}_{\mathcal{D}}} \leq \varepsilon \text{cost}(L)$  and there exists a solution  $S$  that contains  $B_{\mathcal{D}}$  as centers and  $O_{\mathcal{D}}$  as outliers with  $\text{cost}_{\mathcal{I}}(S) \leq (1 + \varepsilon)\text{cost}_{\mathcal{I}}(\text{OPT}) + \varepsilon \text{cost}_{\mathcal{I}}(L)$ .

To deal with the badly cut centers, there is only one hurdle to be able to apply the proof of Lemma IV.6. Indeed, when the algorithm deletes a center of  $\text{OPT}$  that serves a client  $c$ , it is possible to bound the cost of reassigning  $c$  using  $\text{dist}(c, S)$ . However this is not possible to do when  $c$  is an outlier for  $S$ : there is no control on the cost  $\text{dist}(c, S)$ , and hence one has to pay the penalty  $p_c$ . It is thus necessary to find a mechanism that ensures to pay this penalty only with a probability  $\varepsilon$  for each client  $c$ . Similar to Section IV, this is achieved with the following three steps:

- **Step 2.** For each badly-cut facility  $f \in L$  for which  $\psi(f) \neq \emptyset$ , let  $f' \in \psi(f)$  be the closest to  $f$ . Replace  $f'$  by  $f$  in  $S^*$ . For all clients  $c \in L_O$  such that  $\text{OPT}(c) = f'$ , add  $c$  as outliers.
- **Step 3.** Add all badly cut facility  $f'$  of  $L^0$  to  $S^*$
- **Step 4.** Add all badly cut outliers of  $L$  to the outliers of  $S^*$ .

We show next that  $S^*$  satisfies the conditions for  $\mathcal{I}_{\mathcal{D}}$  to have small distortion with good probability.

**Lemma V.2.** *The probability that  $\mathcal{I}_{\mathcal{D}}$  has small distortion is at least  $1 - \varepsilon$ .*

*Proof:* When bounding the cost increase due to Step 2, it is necessary to add as outliers all clients served by  $f'$  that are outliers in  $L$ . Since  $f'$  is deleted from  $S^*$  with probability  $\kappa(\varepsilon, p)$ , the expected cost due to this is  $\sum_{c \in L_O} \kappa(\varepsilon, p) \cdot p_c \leq \kappa(\varepsilon, p) \text{cost}_{\mathcal{I}}(L)$ . Using Markov's inequality, this is at most  $\varepsilon/3 \text{cost}_{\mathcal{I}}(L)$  with probability  $1 - \varepsilon/3$ .

Step 3 does not involve outliers at all. Hence, [Claim IV.3](#) and [IV.4](#) are still valid. Combined with the previous observation about Step 2, this proves that after Step 3,  $S^*$  contains at most  $k$  centers - including the ones in  $B_{\mathcal{D}}$  - and has cost at most  $(1 + \varepsilon) \text{cost}_{\mathcal{I}}(\text{OPT}) + \varepsilon/3 \text{cost}_{\mathcal{I}}(L)$  with probability at least  $1 - \varepsilon/3$ .

Step 4 implies ensures that all outliers in  $O_{\mathcal{D}}$  are also outliers in the constructed solution. Moreover, since an outlier of  $L$  is badly cut with probability  $\kappa(\varepsilon, p)$ , the expected cost increase due to this step is at most  $\kappa(\varepsilon, p) \text{cost}_{\mathcal{I}}(L)$ . Using again Markov's inequality, this cost is at most  $\varepsilon/3 \text{cost}_{\mathcal{I}}(L)$  with probability  $1 - \varepsilon/3$ .

By union-bound, the solution  $S^*$  has cost at most  $(1 + \varepsilon) \text{cost}_{\mathcal{I}}(\text{OPT}) + \varepsilon \text{cost}_{\mathcal{I}}(L)$  with probability  $1 - \varepsilon$ . Hence,  $\mathcal{I}_{\mathcal{D}}$  has small distortion with probability  $1 - \varepsilon$ . ■

Given an instance with low distortion, it is again possible to prove that there exists a near optimal portal-respecting solution, and the same DP as for  $k$ -Median can find it.

Therefore, using the polynomial time algorithm of Charikar et al. [?] to compute a constant-factor approximation, the algorithm presented in [Section IV-C](#) can be straightforwardly adapted, concluding the proof of [Theorem I.4](#).

### B. $k$ -Median with Outliers

In the  $k$ -Median with Outliers problem, the number of outliers allowed is bounded by some given integer  $z$ . We do not manage to respect this bound together with having at most  $k$  facilities and a near-optimal solution: we need to relax it a little bit, and achieve a bicriteria approximation, with  $k$  facilities and  $(1 + O(\varepsilon))z$  outliers. For this, our framework applies nearly without a change.

The first step in the previous construction does not apply directly: the ‘‘cost’’ of removing a center is not well defined. In order to fix this part, Step 1 is randomized: among the facilities of  $\text{OPT}^{\geq 2}$  that are not the closest of their corresponding facility in  $L^{\geq 2}$ , remove from  $S^*$  a random subset  $\hat{\mathcal{H}}$  of size  $\lfloor \varepsilon \cdot |\text{OPT}^{\geq 2}|/2 \rfloor$ .

**Lemma V.3.** *After the randomized Step 1,  $S^*$  has expected cost  $(1 + O(\varepsilon)) \text{cost}(\text{OPT}) + O(\varepsilon) \text{cost}(L)$*

*Proof:* Since there are at least  $|\text{OPT}^{\geq 2}|/2$  facilities of  $\text{OPT}^{\geq 2}$  that are not the closest of their corresponding facility in  $L^{\geq 2}$ , the probability to remove one of them is  $O(\varepsilon)$ . Hence, every outlier of  $L$  that is served in  $\text{OPT}$  must be added as an outlier in  $S^*$  with probability  $O(\varepsilon)$  – when its serving center in  $\text{OPT}$  is deleted. Hence, the expected number of outliers added is  $O(\varepsilon z)$ .

Moreover, the proof of [Lemma IV.1](#) shows that the sum of the cost of deleting all possible facilities is at most  $O(\text{cost}(\text{OPT}) + \text{cost}(L))$  (adding a point as outlier whenever it is necessary). Removing each one of them with probability  $O(\varepsilon)$  ensures that the expected cost of  $S^*$  after step 1 is  $(1 + O(\varepsilon)) \text{cost}(\text{OPT}) + O(\varepsilon) \text{cost}(L)$ . ■

The three following steps are the same as in the previous section, and the proof follows: with constant probability, the instance  $\mathcal{I}_{\mathcal{D}}$  has small distortion (defined as for  $k$ -Median with penalties), and one can use a dynamic program to solve the problem on it. The DP is very similar to the one for  $k$ -Median. The only difference is the addition of a number  $x$  to each table entry, which is a power of  $(1 + \varepsilon/\log n)$ , and represents the (rounded) number of outliers allowed in the subproblem. This adds a factor  $\log^2 n/\varepsilon$  to the complexity.

It is possible to compute a constant factor approximation  $S$  in polynomial time (using Krishnaswamy et al. [?]). Hence, this algorithm is a polynomial time bicriteria approximation scheme for  $k$ -Median with outliers. As in [Section IV-C](#), this directly extends to  $k$ -Means with outliers.

This concludes the proof of [Theorem I.5](#).

### C. $k$ -Center

In the  $k$ -Center problem, the goal is to place  $k$  centers such as to minimize the largest distance from a point to its serving center. We propose a bicriteria approximation, allowing the algorithm to open  $(1 + O(\varepsilon))k$  centers.

For this, we change slightly the definition of badly-cut. Given a solution  $L$  with cost  $\gamma$  and a hierarchical decomposition  $\mathcal{D}$ , a center  $f$  of  $L$  is *badly cut w.r.t  $\mathcal{D}$*  if the ball  $\beta(f, 2^i)$  is cut at some level  $j$  greater than  $i + \tau(\varepsilon, d)$ , for  $i$  such that  $2^{i-1} \leq 2\gamma \leq 2^i$ .

Note that [Lemma II.5](#) still holds with this definition : a center  $f$  is badly cut with probability at most  $\kappa(\varepsilon, p)$ . Let  $B_{\mathcal{D}}$  be the set of badly cut centers. We assume in the following that  $L$  is a 2-approximation, i.e.  $\gamma \leq 2\text{OPT}$ .

We make the crucial following observation, using the doubling property of the metric. Let  $f$  be a center of  $L$ . By definition of doubling dimension, the ball  $\beta(f, \gamma)$  can be covered by  $2^d$  balls of radius  $\gamma/2 \leq \text{OPT}$ . Let  $\mathcal{C}_c$  be the set of centers of such balls, such that  $\beta(f, \gamma) \subseteq \bigcup_{f' \in \mathcal{C}_c} \beta(f', \gamma/2)$ .

Given an instance  $\mathcal{I}$ , we construct  $\mathcal{I}_{\mathcal{D}}$  the following way: for each badly cut facility  $f$ , open all the facilities in  $\mathcal{C}_f$ , and remove all the clients in  $\beta(f, \gamma)$  from the instance. We let  $\mathcal{C} = \bigcup_{f \text{ badly cut}} \mathcal{C}_f$ . The structural lemma of this section is the following:

**Lemma V.4.** *It holds that for all solution  $S$ :*

- $\text{cost}_{\mathcal{I}_{\mathcal{D}}}(S) \leq \text{cost}_{\mathcal{I}}(S)$
- $\text{cost}_{\mathcal{I}}(S \cup \mathcal{C}) \leq \max(\text{cost}_{\mathcal{I}_{\mathcal{D}}}(S), \text{OPT})$

*Proof:* Since the instance  $\mathcal{I}_{\mathcal{D}}$  contains a subset of clients of  $\mathcal{I}$ , it holds that  $\text{cost}_{\mathcal{I}_{\mathcal{D}}}(S) \leq \text{cost}_{\mathcal{I}}(S)$ .



Let  $S$  be a solution in  $\mathcal{I}_{\mathcal{D}}$ . It serves all client in  $\mathcal{I}$  but the one removed: these ones are served by  $\mathcal{C}$  at a cost  $\gamma/2 \leq \text{OPT}$ . Hence, the cost of  $S \cup \mathcal{C}$  is at most  $\max(\text{cost}_{\mathcal{I}_{\mathcal{D}}}(S), \text{OPT})$ . ■

We now show, in a similar fashion as Lemma III.2, that the clients in  $\mathcal{I}_{\mathcal{D}}$  are cut from their serving facility of OPT at a controlled level. Recall that OPT is defined for instance  $\mathcal{I}$ .

**Lemma V.5.** *Let  $c$  be a client in  $\mathcal{I}_{\mathcal{D}}$  and  $\text{OPT}(c)$  its serving facility in OPT.  $C$  and  $\text{OPT}(c)$  are cut at level at most  $\log(2\gamma) + \tau(\varepsilon, d)$ .*

*Proof:* Let  $c$  be a client,  $L(c)$  its serving center in  $L$  and  $\text{OPT}(c)$  its serving center in OPT. If  $c$  is still a client in  $\mathcal{I}_{\mathcal{D}}$ , it means that  $L(c)$  is not badly cut. Observe that  $\text{dist}(L(c), \text{OPT}(c)) \leq \text{dist}(c, L(c)) + \text{dist}(c, \text{OPT}(c)) \leq \gamma + \text{OPT} \leq 2\gamma$

Let  $i$  such that  $2^{i-1} \leq 2\gamma \leq 2^i$ . Since  $L(c)$  is not badly cut, the ball  $\beta(L(c), 2^i)$  is not badly cut neither: hence,  $c$  and  $\text{OPT}(c)$  (that are in this ball) are cut at level at most  $i + \tau(\varepsilon, d) \leq \log(2\gamma) + \tau(\varepsilon, d)$ . ■

This lemma is stronger than Lemmas III.2 and IV.5: it allows us to consider only levels of the decomposition with diameter less than  $2^{1+\tau(\varepsilon, d)}\gamma$ .

Since the set  $\mathcal{C}$  has expected size  $\kappa(\varepsilon, p)k$ , Markov's inequality ensures that with probability  $1 - \varepsilon$  this set has size  $O(\varepsilon)k$ . If every part with diameter  $\Delta$  of the hierarchical decomposition is equipped with a  $\rho\Delta$ -net (for  $\rho = \varepsilon 2^{-\tau(\varepsilon, d)}$ ), Lemma V.5 ensure that there exists a portal-respecting solution  $S$  with cost  $\text{cost}_{\mathcal{I}_{\mathcal{D}}}(S) \leq \text{OPT} + O(\varepsilon)\gamma = (1 + O(\varepsilon))\text{OPT}$ . Lemma V.4 ensures that lifting this solution back to  $\mathcal{I}$  and adding  $\mathcal{C}$  as centers gives a near-optimal solution.

Using the same algorithm as for  $k$ -Medians to compute a good portal-respecting solution, and computing a 2-approximation with a simple greedy algorithm (see e.g. [?]), that runs in time  $O(n \log k)$  concludes the proof of Theorem I.6.

## VI. ACKNOWLEDGMENT

Andreas Emil Feldman is supported by the Czech Science Foundation GAČR (grant #17-10090Y), and by the Center for Foundations of Modern Computer Science (Charles Univ. project UNCE/SCI/004).

## APPENDIX

### A. Proof of Section II

*Proof of Lemma II.3:* We present the algorithm constructing the hierarchical decomposition and proves the lemma as a second step.

Without loss of generality, assume that the smallest distance in the metric is 1: the aspect-ratio  $\Delta$  is therefore the diameter of the metric. Start from a hierarchy of nets  $Y_0 := V, \dots, Y_{\log(\Delta)}$  such that  $Y_i$  is a  $2^{i-2}$ -net of  $Y_{i-1}$ . Moreover,

pick a random order on the points  $V$  and a random number  $\tau \in [1/2, 1)$ . The hierarchical decomposition  $\mathcal{D}$  is defined inductively, starting from  $\mathcal{B}_{\log \Delta} = V$ . To partition a part  $B$  at level  $i$  into subpart at level  $i-1$ , do the following: for each  $y \in Y_{i-1} \cap B$  taken in the random order, define  $B \cap \beta(y, \tau 2^i)$  to be a part at level  $i-1$  and remove  $B \cap \beta(y, \tau 2^i)$  from  $C$ .

When we assume access to the distances through an oracle, it is possible to construct this hierarchy an augment it with the set of portals in time  $(1/\rho)^{O(d)}n \log(\Delta)$ . Moreover, these portals can be made *nested*, meaning that portals at level  $i+1$  are also portals at level  $i$  [?], [?].

We prove now that this hierarchical decomposition has the required properties. The diameter of each part is bounded by  $2^{i+1}$  by construction; therefore to have Property 2 it is enough to make  $\mathcal{P}_i$  an  $(\varepsilon 2^{i+1})$ -net of  $V$ . The Lemma II.1 ensures the conciseness, and the definition of a net ensures that every point is at distance  $\varepsilon 2^{i+1}$  of  $\mathcal{P}_i$ , which implies the preciseness. Proving the scaling property requires a bit more work.

The two ingredients needed for this part stem from the construction of the decomposition: the diameter of any part at level  $i$  is at most  $2^{i+1}$ , and the minimum distance between two points of  $Y_i$  is bigger than  $2^{i-2}$ .

These two properties are enough in order to prove our lemma. Let  $i$  be a level such that  $2^i \leq r$ : then  $r/2^i = \Omega(1)$  so there is nothing to prove. Otherwise, we proceed in two steps. First, let us count the number of level  $i$  parts that could possibly cut a ball  $\beta(x, r)$ . A level  $i$  part is included in a ball  $\beta(y, 2^i)$  for some  $y \in Y_i$ ; therefore if  $\text{dist}(x, y) > r + 2^i$  then  $y$ 's part cannot cut  $\beta(x, r)$ . So it is required that  $\text{dist}(x, y) \leq r + 2^i \leq 2 \cdot 2^i$ . But since the minimum distance between two points of  $Y_i$  is  $2^{i-2}$ , and  $Y_i$  has doubling dimension  $d$ , we have  $|Y_i \cap \beta(x, 2 \cdot 2^i)| = 2^{d \log(2^i/2^{i-2})} = 2^{2d}$ . Thus there is only a bounded number of parts to consider.

We prove for each of them that the probability that it cuts  $\beta(x, r)$  is  $O(r/2^i)$ . A union-bound on all the possible parts is then enough to conclude. Let therefore  $y \in Y_i \cap \beta(x, 2 \cdot 2^i)$ , and  $x_m$  and  $x_M$  be the respective closest and farthest point of  $\beta(x, r)$  from  $y$ . A necessary condition for  $y$ 's part to cut  $\beta(x, r)$  is that the diameter of the part is in the open interval  $(d(y, x_m), d(y, x_M))$ . Since  $x_m, x_M \in \beta(x, r)$  this interval has size  $2r$ , and the radius of the part is picked uniformly in  $[2^i/2, 2^i)$ . Therefore the probability that the radius of the part falls in  $(d(y, x_m), d(y, x_M))$  is at most  $4r/2^i$ . And finally, the probability that  $y$ 's part cuts  $\beta(x, r)$  is indeed  $4r/2^i$ .

By a union-bound over all the parts that could possibly cut  $\beta(x, r)$  we obtain the claimed probability  $\Pr[\mathcal{C} \text{ cuts } \beta(x, r) \text{ at a level } i] = 2^{2d+2}r/2^i$ . ■

**Lemma A.1.** *Let  $P$  be a problem among Facility Location,  $k$ -Medians or  $k$ -Means. Given an instance  $(\mathcal{I}, \text{dist})$  with  $n$*



points,  $\varepsilon > 0$  and a constant-factor approximation for  $P$  on  $\mathcal{I}$ , there exists a linear-time algorithm that outputs a set of instances  $(\mathcal{I}_1, \text{dist}_1), \dots, (\mathcal{I}_m, \text{dist}_m)$  such that

- $\mathcal{I}_1, \dots, \mathcal{I}_m$  is a partition of  $\mathcal{I}$
- for all  $i$ ,  $\mathcal{I}_i$  has aspect-ratio  $O(n^4/\varepsilon)$ ,
- if  $(\bigcup \mathcal{I}_i, \min \text{dist}_i)$  is the instance where distances between points of the same part  $\mathcal{I}_i$  are given by  $\text{dist}_i$  and distances between points of different parts is set to  $\infty$ , then
  - there exists a solution on  $\bigcup \mathcal{I}_i$  with cost  $(1 + \varepsilon/n)\text{cost}(\text{OPT})$ , and
  - any solution on  $\bigcup \mathcal{I}_i$  of cost  $X$  induces a solution of cost at most  $X + \varepsilon\text{cost}(\text{OPT})/n$  on  $\mathcal{I}$ .

*Proof:* The cost of the constant-factor approximation is an estimate  $\gamma$  on the cost of the optimum solution  $\text{OPT}$ :  $\gamma = \Theta(\text{cost}(\text{OPT}))$ . It is then possible to replace all distances longer than  $2\gamma$  by  $\infty$ : distances longer than  $\gamma$  will indeed never be used by solution with cost better than  $\gamma$ , so the cost of these solutions is preserved after this transformation. We say that two vertices are *connected* if their distance is not  $\infty$ , and call a connected component any maximal set of connected vertices. The transformation ensured that any connected component has diameter at most  $2\text{OPT}$ , and that every cluster of  $\text{OPT}$  is contained inside a single connected component. Moreover, any connected component has doubling dimension  $2d$ : indeed, a subspace of a metric with doubling dimension  $d$  has a doubling dimension at most  $2d$ . Note also that this transformation can be made implicitly: every time the algorithm queries an edge, it can replace the result by  $\infty$  if necessary.

To identify the connected component, the algorithm builds a spanner with the algorithm of [?]: the connected components of the spanner are exactly the ones of our metric, and can be found in linear time.

Then, for each connected component, the algorithm defines an instance of the more general version of the clustering problem by the following way. It first sets  $\chi(v) = 1$  for all vertex  $v$ . Then, it iterates over all edges, it contracts every edge  $(u, v)$  with length less than  $(\varepsilon \cdot \gamma/n^3)$  to form a new vertex  $w$  and sets  $\chi(w) = \chi(u) + \chi(v)$ .

Now, we aim at reconstructing a metric from this graph. We will do it in an approximate way: for all connected points  $u, v$  of connected component  $i$ , we set  $\text{dist}_i(u, v)$  to be 0 if  $u$  and  $v$  are merged in the graph, and otherwise  $\text{dist}(u, v)$ . This ensures that  $\varepsilon \cdot \gamma/n^3 \leq \text{dist}_i(u, v) \leq 2\gamma$ , hence the aspect-ratio of  $\mathcal{I}_i$  is  $O(n^4/\varepsilon)$ . Moreover, every distance is preserved up to an additive  $O(\varepsilon \cdot \text{cost}(\text{OPT})/n^2)$ .

Since every cluster of  $\text{OPT}$  is contained inside a single connected component, this ensures that  $\text{OPT}$  induces a solution of cost  $(1 + \varepsilon/n)\text{cost}(\text{OPT})$  on  $\bigcup \mathcal{I}_i$ . Moreover, lifting a solution in  $\bigcup \mathcal{I}_i$  to  $\mathcal{I}$  costs at most  $\varepsilon\text{cost}(\text{OPT})/n^2$  per pair (client, center) and therefore  $\varepsilon\text{cost}(\text{OPT})/n$  in total. ■

If the problem considered is Facility Location, it is easy to merge the solutions on subinstances: since there is no cardinality constraint, the global solution is simply the union of all the solutions. The hard constraint on  $k$  makes things a bit harder. Note that the dynamic program presented in Section IV-C naturally handles it without any increase in its complexity: however, for completeness we present now a direct reduction.

**Lemma A.2.** *Given a problem  $P$  among  $k$ -Medians or  $k$ -Means, a set of instances  $(\mathcal{I}_1, \text{dist}_1), \dots, (\mathcal{I}_m, \text{dist}_m)$  given by Lemma A.1 and an algorithm running in time  $n_i(\log n_i)^{\alpha}t(\Delta)$  to solve  $P$  on instances with  $n_i$  points and aspect-ratio  $\Delta$ , there exists an algorithm that runs in time  $O(n(\log n)^{\alpha+2}t(O(n^4/\varepsilon)))$  to solve  $P$  on  $\bigcup \mathcal{I}_i$ .*

*Proof:* First, note that the optimal solution in  $\bigcup \mathcal{I}_i$  is  $O(n^5/\varepsilon)$ , since the maximal distance in any of  $\mathcal{I}_1, \dots, \mathcal{I}_m$  is  $n^4/\varepsilon$ . Using this fact, we build a simple dynamic program to prove the lemma. For all  $i \leq m$  and  $j \leq \log_{1+\varepsilon/\log n}(n^5/\varepsilon)$ , let  $k_{i,j}$  be the minimal  $k'$  such that the cost of  $P$  with  $k'$  centers in  $\mathcal{I}_i$  is at most  $(1 + \varepsilon/\log n)^j$ .  $k_{i,j}$  can be computed with a simple binary search, using the fact that the cost of a solution is decreasing with  $k'$ .

Given all the  $k_{i,j}$ , a simple dynamic program can compute  $k_{\geq i,j}$ , the minimal number of centers needed to have a cost at most  $(1 + \varepsilon)^j$  on  $\mathcal{I}_i, \dots, \mathcal{I}_m$  (the  $\varepsilon/\log n$  becomes a simple  $\varepsilon$  because of the accumulation of errors). The solution for our problem is  $(1 + \varepsilon)^j$ , where  $j$  is the minimal index such that  $k_{\geq 1,j} \leq k$ .

The complexity of computing  $k_{i,j}$  is  $O(\log k \cdot n_i(\log n)^{\alpha}t(O(n^4/\varepsilon)))$ , hence the complexity of computing all the  $k_{i,j}$  is  $O(n(\log n)^{\alpha+2}t(O(n^4/\varepsilon)))$ . The complexity of the dynamic program computing  $k_{\geq i,j}$  is then simply  $O(m \log n) = O(n \log n)$ , which concludes the proof. ■

Hence, in the following, we only focus on solving problems on instances where the aspect-ratio is polynomial in  $n$ .