

Consistent k-Clustering for General Metrics

Hendrik Fichtenberger* Silvio Lattanzi† Ashkan Norouzi-Fard‡ Ola Svensson§

Abstract

Given a stream of points in a metric space, is it possible to maintain a constant approximate clustering by changing the cluster centers only a small number of times during the entire execution of the algorithm?

This question received attention in recent years in the machine learning literature and, before our work, the best known algorithm performs $\tilde{O}(k^2)$ center swaps (the $\tilde{O}(\cdot)$ notation hides polylogarithmic factors in the number of points n and the aspect ratio Δ of the input instance). This is a quadratic increase compared to the offline case — the whole stream is known in advance and one is interested in keeping a constant approximation at any point in time — for which $\tilde{O}(k)$ swaps are known to be sufficient and simple examples show that $\Omega(k \log(n\Delta))$ swaps are necessary. We close this gap by developing an algorithm that, perhaps surprisingly, matches the guarantees in the offline setting. Specifically, we show how to maintain a constant-factor approximation for the k -median problem by performing an optimal (up to polylogarithmic factors) number $\tilde{O}(k)$ of center swaps. To obtain our result we leverage new structural properties of k -median clustering that may be of independent interest.

Keywords: k-median, k-clustering, consistency, approximation algorithms

*University of Vienna, Faculty of Computer Science, Austria. Email: hendrik.fichtenberger@univie.ac.at

†Google Research, Zurich. Email: silviol@google.com

‡Google Research, Zurich. Email: ashkannorouzi@google.com

§School of Computer and Communication Sciences, EPFL. Email: ola.svensson@epfl.ch. Supported by the Swiss National Science Foundation project 200021-184656 “Randomness in Problem Instances and Randomized Algorithms.”

1 Introduction

Detecting the clustering structure of real-world data is a basic primitive used in a wide range of data analysis tasks such as community detection, spam detection, computational biology and many others. Many different formulations of clustering problems have been proposed and studied throughout the years. Among these, the geometric versions of the problem has attracted a lot of attention for their theoretical and practical importance. In those problems we are given as input n points in a metric space and a distance oracle and we want to compute a clustering that minimizes an ℓ_p -objective function, such as the k -means, the k -median or the k -center objective.

Due to their relevance, the problems have been extensively studied and many algorithms [1, 3, 4, 6, 8, 28, 27, 31] and heuristics [32] have been proposed to solve the classic offline version of the problem. Furthermore they have been also extensively studied in the classic streaming and online setting where points are inserted in the instance sequentially. For those problems different techniques have been used in the Euclidean \mathbb{R}^d setting [11, 15, 14, 16, 24, 23] and in the general metric space setting [10, 17, 35]. In particular in the first case the algorithms are mostly based on coresets constructions. By contrast, in general metric spaces the algorithms are mostly based on adaptive sampling.

Inspired by real world applications where datasets continuously evolve in time, we study the consistent k -median problem in general metric spaces where points are inserted on the fly. In this problem we are interested in efficiently maintaining a valid solution that is a good approximation of the optimal solution and that is *consistent*. In particular given a stream of insertions, we are interested in designing an algorithm that maintains a constant-factor approximation at any time while minimizing the total number of changes in consecutive solutions.

Consistency of the solution is of importance from both a theoretical and a practical point of view and in recent years it received a lot of attention in the machine learning literature [12, 26, 30]. In fact, in various applications the computed cluster centers are used in other downstream tasks. In this case, changes in the solution might result in costly recomputations.

To formally capture the notion of consistency we use the same definition as in [30] where the consistency of a clustering is measured by the number of times the cluster centers are modified during the entire execution of the algorithm. Intuitively, this definition makes sense because it captures the number of times in which the clustering changes its underlying structure defined by its cluster centers.

For the consistent k -clustering problem Lattanzi and Vassilvitskii [30] presented a first constant-factor approximation algorithm that executes at most $k^2 \cdot \text{polylog}(n, \Delta)$ changes, where Δ is the aspect ratio of input points. A simplified version of their algorithm intuitively works as follows: First they run the Meyerson’s sketch [10, 35] to compress the stream down to $O(k \cdot \text{polylog}(n, \Delta))$ weighted point insertions and afterwards, for each point that is inserted in the weighted instance, they compute a solution using any constant-factor approximation algorithm. Moreover, they also show that, even if the stream of insertions is known in advance, any algorithm that maintains c -approximate solutions requires at least $\Omega(k \cdot (\log_c(\frac{n}{k}) + \log_c(\frac{\Delta}{k})))$ changes. Finally, they also prove that if the stream of insertions is known in advance, one can design a constant-factor approximation algorithm with $k \cdot \text{polylog}(n, \Delta)$ changes in the solution. Guo et al. [19] study consistent k -clustering with outliers. Their algorithm is based on local search and requires at most $O(k^2(\log n \Delta)^2)$ changes.

A natural open question left is to close the gap between the agnostic and the setting with knowledge of the future. In this paper we solve this question by providing the first consistent k -median algorithm that maintains a constant-factor approximation at any point in time by only changing the centers $k \cdot \text{polylog}(n, \Delta)$ times.

Our Results We introduce a novel approach for the consistent k -median problem in metric spaces. We present a constant-factor approximation algorithm with a total of $k \cdot \text{polylog}(n, \Delta)$ changes in the consecu-

tive solutions. Our approach is tight (up to a polylogarithmic factor) due to the above explained lower-bound. Moreover, surprisingly, our result shows that knowing the future is not a key information for this problem by achieving the same number of changes as in the offline setting with full knowledge of the future (up to a polylogarithmic factor). To obtain our result we leverage new structural properties of k -median clustering and we introduce a new concept of robust centers that may be of independent interest. Maintaining a robust solution enables us to change the solution a little even for adversarial point insertions. More precisely, in our algorithm, we first use the Meyerson’s sketch [35] and get a stream of $k \cdot \text{polylog}(n, \Delta)$ weighted points. Afterwards we show that for each insertion we change the center of the maintained robust solution at most $\text{polylog}(n, \Delta)$ times on average, instead of $k \cdot \text{polylog}(n, \Delta)$ of the previous work. In order to achieve this, we establish a strong relationship between the input points and the metric space. We change the position of the centers of the solution so that they have a better local coverage despite the fact that this might increase the cost of the clustering. This enables us to be more consistent in the future, whilst it increases the solution cost by a small factor.

Extensions to Other Problems. For simplicity, we state our result for the classic k -median problem, although it is not hard to extend our results for general p -norms for constant $p \geq 1$. In fact, it has already been noted that the Meyerson sketch can be extended to this setting [30] and our reductions only use basic geometric properties as triangle inequalities so all our proofs can be adapted to work in general p -norms. Furthermore, for the same reason, our result can also be extended to work in λ -metrics¹ for constant λ .

Further related works Interestingly, we note here that our notion of consistency is also closely related to the notion of recourse in online algorithms. In this setting, one seeks better online algorithms to compute optimal or approximate solutions for combinatorial problems by allowing the algorithm to make a limited number of changes to the online solution. The first problem studied in this setting was the classic online Steiner tree problem introduced by [25] for which it is possible to design better algorithms by allowing a small recourse as shown in several papers [18, 21, 29, 33]. After this result several other classic optimization problems have been studied in this setting as online scheduling [2, 13, 36, 37, 38], online flow [22, 39], online matching [5] and online set cover [20].

2 Overview of Our Approach

The starting point of our approach are two basic observations that allow us to reduce our problem to an equivalent simpler formulation that we will then solve efficiently. This reduction is not new, in fact it has been first observed in [30] (although in this paper we need to strengthen it slightly for our result).

The first observation is that we can divide the stream in $\log(n\Delta)$ phases so that in each phase the cost of an optimal solution does not change too much. In particular, after each insertion we compute an α -approximate solution (for $\alpha \leq 3$, using for example [6]). Now, when the current approximation cost increases by a factor 6 in comparison to the cost at the beginning of a phase, we know that the value of the optimum solution has increased by at least a factor 2 and at most a factor 18 so we restart a new phase. The second is that within a phase with high probability we can reduce the problem with n point insertions to a problem where only $k \cdot \text{polylog}(n, \Delta)$ weighted points are inserted to the instance. This can be achieved using a standard sketching tool, the Meyerson’s sketch [10, 35]. This technique results in losing a constant-factor in the approximation ratio but it provides a way to transform any constant-factor approximation algorithm that makes on average C changes to the cluster centers for each weighted point insertion to a constant approximation an algorithm that makes $Ck \cdot \text{polylog}(n, \Delta)$ changes in total. Note

¹ λ -metrics are metrics where the triangle inequality hold only with an additional multiplicative λ factor.

that this automatically results in a constant approximation algorithms that makes $k^2 \cdot \text{polylog}(n, \Delta)$ and it is the key idea behind the algorithm presented in [30].

Now we can turn our attention to the main technical contribution of the paper. Assuming that the cost of the optimal solution is stable we want to design an algorithm that maintains a constant approximation and makes on average $\text{polylog}(n, \Delta)$ changes to the cluster centers to handle insertions of weighted points.

Let \mathcal{V} be the set of centers in the optimal clustering at the end of the phase and let \mathcal{U} be the set of cluster centers computed by the algorithm. Now suppose that at any point in time during the execution of a phase we have that every center in \mathcal{U} can be paired with a center in \mathcal{V} so that the paired centers are close to each other but they are far away from all the other centers. Then in this setting we would expect future point insertions in the phase to modify the clustering structure only minimally. Intuitively, this is true because the cost of the optimal solution does not increase too much and so the clustering induced by the centers in \mathcal{U} is fairly stable and centers exchange only points in their peripheries. In fact, in this setting we can essentially reduce the problem to solving a set of disjoint 1-median problems.

Unfortunately this approach does not generalize to more complex instances where the clustering structure is not as neat. In fact, at any point in time we will have that only a (possibly empty) subset of points in \mathcal{U} can be paired with a center in \mathcal{V} so that the paired centers are close to each other but they are far away from all other centers. We refer to those pairs as well-separated pairs (for a formal definition please refer to Definition 7.2). To tackle this more challenging setting we prove a few key structural lemmas. From a high level perspective, our first key observation is that if the current set of centers is composed by $k - \ell$ well-separated pairs then it is possible to obtain another set of centers that has cost a constant time larger than the initial set of centers but that uses $\Omega(\ell)$ less centers. Informally, we prove such a statement by carefully constructing a fractional solution that opens all centers forming well-separated pairs and centers that are close to many optimal centers completely but opens only fractionally centers that are not close to multiple optimum centers (for more details refer to Section 7.1). Intuitively, this is true because we can show that centers that are not part of well-separated pairs have other centers in their proximity and so we can charge points in their cluster partially to other centers.

Thanks to the previous observation we can design an algorithm as follows: given a current solution reduce the number of centers as much as possible without increasing the cost of the induced clustering more than a constant factor. Suppose that in this way, we can reduce the number of centers by $\ell' \in \Omega(\ell)$. After this step, we can easily insert ℓ' new points and open them as new centers. In this way we can handle ℓ' insertions without increasing the cost of the solution. Unfortunately we cannot use this algorithm on its own to solve the problem because ℓ' may be 0 and the cost of the solution may increase too much in subsequent iterations. To tackle this issues we prove our second structural lemma where we show that after adding $\ell' + 1$ points to the current instance we can find a set of $O(\ell + 1)$ centers swaps such that the solution obtained after those swaps is a β -approximation for the optimal clustering for a fixed constant β . The main idea to prove such a lemma is to show that one can construct a β -approximation by keeping all the centers that are in well-separated pairs and for which no additional point has been added to their induced cluster, and by swapping all other centers. One key idea in this context is to identify a set of *robust* centers for the clusters. The main idea behind this notion is that we want our centers to be good centers for our clustering at different scales² so that small changes in the periphery of our clusters do not affect their quality (for a formal definition and more details please refer to Section 4). Interestingly, we can show in Section 4 that for any cluster it is possible to find a center that is a “good” center for the cluster at any scale. Furthermore in Section 6 we then show that those centers are robust meaning that they do not need to be updated often unless their clusters change significantly.

Combining these two observations we can define our update algorithm, `EpochAlgorithm`. It is called repeatedly in each phase until the phase is finished. `EpochAlgorithm` starts with a set of centers \mathcal{U} such

²Interestingly, we note that this notion is somehow related to the notion of center used for prefix clustering in [34].

that: i) \mathcal{U} has a constant approximation ratio for the current set of points P , ii) all centers in \mathcal{U} are *robust*. Then EpochAlgorithm removes from \mathcal{U} as many center as it can without significantly increasing the cost of the clustering. This set can be executed using an LP based algorithm (for more details please refer to Section 9.2). Suppose that the algorithm removes ℓ' centers. For the next ℓ' insertions, the algorithm inserts the inserted points as centers in the instance. Then an additional point is added, EpochAlgorithm finds $O(\ell)$ good centers swaps and performs them using another LP-based algorithm described in Section 9.1. In this way we obtain a new good solution and the last step of EpochAlgorithm is to make the new solution robust by robustifying the centers.

Roadmap. We start by introducing some basic notation in Section 3. Then we introduce the notion of robust center and few basic properties of robust centers in Section 4. After formally defining the notion of robust center we give a formal description of our EpochAlgorithm in Section 5. Then in Section 6 and Section 7 we bound the number of changes in cluster centers and the approximation factor for EpochAlgorithm, respectively. Then in Section 8 we present our complete algorithm (with the preprocessing steps) and in Section 9 our algorithms for reducing the number of centers and finding a good set of swaps. Finally, for completeness, in Appendix A we show how the Meyerson sketch can be used in our problem to reduce the size of the input instance.

3 Preliminaries and Formal Statement of the Main Result

In the consistent k -median problem, we are given a stream $\sigma = \langle \sigma_1, \sigma_2, \dots, \sigma_n \rangle$ consisting of n insertions of points and we are interested in maintaining a good and consistent solution to the k -median problem at any point in time. More formally, we want to maintain a solution that is a constant-factor approximation while the number of changes in the solution is minimized throughout the execution of the algorithm.³

Given any two points x_1 and x_2 , we assume to have access to a metric distance oracle that returns the distance between x_1 and x_2 . We denote this distance by $\text{dist}(x_1, x_2)$. By scaling, we assume that the smallest non-zero distance is 1 and we let Δ denote the largest distance. We thus have that Δ equals the ratio between the largest distance to the smallest non-zero distance, which is often referred to as the aspect ratio. For a set S of points and a point x , we also let $\text{dist}(x, S) = \min_{y \in S} \text{dist}(x, y)$ be the distance from x to the closest point in S . By convention, $\text{dist}(x, \emptyset) = \Delta$, i.e., we let the distance from a point to the empty set equal the largest possible distance. We remark that our algorithm does not need to know the value of n, Δ in advance although for the sake of simplicity, we assume that we know these values (we explain how one can remove such assumption in Section 8).

Given a set P of points and a subset $\mathcal{U} \subseteq P$ of at most k points, we refer to the points in \mathcal{U} as centers. A clustering is an assignment of points to centers. For a set of centers \mathcal{U} and a set P of points, we let the cost of the clustering induced by \mathcal{U} be

$$\text{cost}(\mathcal{U}, P) = \sum_{x \in P} \text{dist}(x, \mathcal{U}).$$

In the case that the input points P are weighted, by $w(x)$ we denote the weight of the point x for $x \in P$ and we define the cost of the clustering induced by \mathcal{U} to be

$$\text{cost}(\mathcal{U}, P) = \sum_{x \in P} w(x) \cdot \text{dist}(x, \mathcal{U}).$$

We also extend the notion of weights to sets, i.e., $w(P) = \sum_{x \in P} w(x)$. The k -median problem on input P asks for a set \mathcal{U} of k centers that minimizes $\text{cost}(\mathcal{U}, P)$ among all sets of k centers. We say that an algorithm

³We emphasize that the algorithm has no knowledge of the points $\sigma_{i+1}, \sigma_{i+2}, \dots, \sigma_n$ at time i .

ALG maintains an α -approximate clustering, if at any point in time the clustering induced by \mathcal{U} computed after an insertion has cost at most α -times the optimal solution to the k -median problem. Formally, let P_i be the set of points in our instance after the first i insertions of the stream, i.e., $P_i = \{\sigma_1, \dots, \sigma_i\}$. Then we say that ALG maintains an α -approximate solution if

$$\text{cost}(\mathcal{U}_i, P_i) \leq \alpha \cdot \text{OPT}(P_i) \quad \text{for all } i = 1, 2, \dots, n,$$

where \mathcal{U}_i denotes the centers computed by ALG after the i -th insertion and $\text{OPT}(Q)$ denotes the value of an optimal clustering of points in Q .

The main focus of this paper is in developing an algorithm that maintains a constant approximation at any point in time by keeping the set of centers as stable as possible. This notion of stability is formally captured by the consistency of the center set, which is the summation of the number of changes in the consecutive sets of centers produced by the clustering algorithm. Formally,

$$\sum_{1 \leq i \leq n-1} |\mathcal{U}_{i+1} \setminus \mathcal{U}_i|.$$

In this work we present a constant approximate algorithm with $k \cdot \text{polylog}(n, \Delta)$ number of changes in the solution, matching the lower bound on the number of changes up to polylogarithmic factors.

Theorem 3.1. *There exists an algorithm that given a stream of points $\sigma = \langle \sigma_1, \sigma_2, \dots, \sigma_n \rangle$, outputs a set of at most k centers \mathcal{U}_i after the i -th insertion for all $1 \leq i \leq n$ so that with probability at least $1 - \frac{1}{n^8}$*

1. \mathcal{U}_i induces a constant approximation clustering, i.e., for any $1 \leq i \leq n$

$$\text{cost}(\mathcal{U}_i, \{\sigma_1, \sigma_2, \dots, \sigma_i\}) \leq O(1) \cdot \text{OPT}(P_i).$$

2. The total number of changes in consecutive center sets is at most $k \cdot \text{polylog}(n\Delta)$.

$$\sum_{1 \leq i < n} |\mathcal{U}_{i+1} \setminus \mathcal{U}_i| \leq k \cdot \text{polylog}(n, \Delta).$$

4 Robust Centers

One of the key concept in our analysis is the notion of *robust* centers. Robust centers are not necessarily optimal centers, but they are robust to changes in the cluster structure. This notion is central in our algorithm, as we aim to find a stable set of k centers. To get an intuition behind our definition, consider a cluster in an approximately optimal solution that at a specific point in time has a rather blurry structure, e.g., many points in different location are good candidate centers. After few additional insertions the cluster structure may develop and concentrate, for example the cluster may grow, shrink, merge or split, but at its current state, it is not clear what will happen. In this setting an optimal solution with knowledge of future insertions would pick a center that is good for the current point set *and* good after inserting future points (see Fig. 1a). Unfortunately, our algorithm cannot predict how an optimal consistent clustering would look without knowing the future. Nevertheless, we can select our centers so that they are stable to small changes in the clusters. To this end, we say that a center is robust when it is locally optimal on different scales, i.e., we optimize centers so that they minimize the assignment cost of close points better than the assignment cost of points far away. Intuitively, this makes sense because points that are far away from the center are potentially prone to moving between clusters. Again, an algorithm cannot tell a priori what *far away* means because the scale of concentration of a cluster is unknown. Therefore, robust centers are optimized for (exponentially) growing distances.

Before formalizing this intuition we introduce few basic definitions.




Figure 1: (a) Consider the two initial choices for a single center in the left clustering (red circle, blue square). Both are good choices and, depending on the metric, the blue center might have the smaller cost, but the red center is robust. However, if points are inserted and points accumulate in the upper *and* the lower part, it may be necessary to split the cluster. In this case, the red center is better off in terms of consistency: while the blue solution needs to change two centers (the initial blue center and a new center, depicted in the middle), the red solution only needs to add a center at the top (depicted on the right); the original red center is already a good choice. (b) Consider a (sub-)sequence of a robust tuple: p_i (green triangle), p_{i-1} and p_{i-2} (blue square), and p_{i-3} (red circle). Since $\text{Ball}(p_{i-1}, 10^{i-1})$ (second inner blue circle) has large average cost, we have $p_{i-1} = p_{i-2}$. The $\text{Ball}(p_{i-2}, 10^{i-2})$ has small average cost, so p_{i-2} is chosen optimally inside it.

Definition 4.1. For a center u and a set of points P , we define the average cost of assigning the points in P to u by

$$\text{avgcost}(u, P) = \frac{\text{cost}(u, P)}{w(P)}.$$

Definition 4.2. For a set of points P , a point $u \in P$, and $r \geq 0$, we define $\text{Ball}_P(u, r) = \{p \in P : \text{dist}(u, p) \leq r\}$ to be the points of P contained in the ball of radius r centered at u .

When the set P of points is clear from the context, we simplify notation and write $\text{Ball}(u, r)$ instead of $\text{Ball}_P(u, r)$.

Let \mathcal{U} be a set of centers for a point set P , for $u \in \mathcal{U}$ consider a subset of points that is assigned to this center, say, $P' = \text{Ball}_P(u, \text{dist}(u, \mathcal{U} \setminus \{u\})/2)$. Now, imagine that after few insertion the set of centers changes and so the ball's radius decreases. Now we may need to move the center within the ball to a new optimal position, furthermore we can imagine that this could happen multiple times during the execution of the algorithm. Intuitively (e.g., for Euclidean distances), a center is attracted by accumulations of points in its current ball. To obtain a robust center we define our centers recursively so that they are good centers at different scale and so that they are not sensitive to accumulation of points in the boundary of the current ball. To do this we first discretize the shrinking process of the ball and consider exponential steps when decreasing its radius. Then we do not move the center when the average cost indicates that many of the points within the ball lie close to its boundary. This case indicates that, intuitively, there is no concentration of points close to the center that attracts it. See Fig. 1b for an example. More formally,

Definition 4.3. For an integer $t \geq 0$, we say that a tuple (p_0, p_1, \dots, p_t) of $t + 1$ points is t -robust with respect to a point set P if, for every $i \in \{1, 2, \dots, t\}$,

$$p_{i-1} = \begin{cases} p_i & \text{if } \text{avgcost}(p_i, \text{Ball}_P(p_i, 10^i)) \geq 10^i/5, \\ \arg \min_{p \in \text{Ball}_P(p_i, 10^i)} \text{cost}(p, \text{Ball}_P(p_i, 10^i)) & \text{otherwise.} \end{cases}$$

We also say that a point/center $u \in P$ is t -robust if there is a tuple (u, p_1, \dots, p_t) that is t -robust.

Interestingly toward the end of this section we will prove that p_0 is an approximately good center for any ball $\text{Ball}_P(p_i, 10^i)$.

For simplicity, we also write that (p_0, p_1, \dots, p_t) is t -robust without explicitly stating that it is with respect to P when there is no ambiguity. In particular, the point set will be fixed to be P throughout this section.

Before proving two basic properties of t -robust tuples, we describe a simple algorithm $\text{MakeRobust}(t, p)$ that, on input an integer $t \geq 0$ and a point $p \in P$, outputs a t -robust tuple (p_0, p_1, \dots, p_t) with $p_t = p$. The steps of $\text{MakeRobust}(t, p)$ are as follows:

- Let $p_t = p$.
- For i from t to 1, select p_{i-1} such that

$$p_{i-1} = \begin{cases} p_i & \text{if } \text{avgcost}(p_i, \text{Ball}(p_i, 10^i)) \geq 10^i/5, \\ \arg \min_{p \in \text{Ball}(p_i, 10^i)} \text{cost}(p, \text{Ball}(p_i, 10^i)) & \text{otherwise.} \end{cases}$$

- Output (p_0, p_1, \dots, p_t) .

By definition, we have that (p_0, p_1, \dots, p_t) is a t -robust tuple with $p = p_t$. We proceed to prove basic properties of t -robust tuples that will allow us to analyze our main algorithm that will repeatedly make calls to MakeRobust .

We start with a simple observation that bounds the distance from p_0 to the other points.

Lemma 4.4. Let (p_0, p_1, \dots, p_t) be a t -robust tuple and let $B_j = \text{Ball}(p_j, 10^j)$. For every $j = 1, \dots, t$,

$$\text{dist}(p_{j-1}, p_j) \leq 10^j/2, \quad B_{j-1} \subseteq B_j, \quad \text{and} \quad \text{dist}(p_0, p_j) \leq 10^j/2.$$

Proof. By the definition of a t -robust tuple, we have $\text{dist}(p_{j-1}, p_j) = 0$ if $\text{avgcost}(p_j, B_j) \geq 10^j/5$. Otherwise, by the selection of p_{j-1} , $\text{cost}(p_{j-1}, B_j) \leq \text{cost}(p_j, B_j)$ and so

$$\begin{aligned} \frac{10^j}{5} + \frac{10^j}{5} &\geq \text{avgcost}(p_j, B_j) + \text{avgcost}(p_{j-1}, B_j) \\ &= \sum_{p \in B_j} \frac{w(p)}{w(B_j)} (\text{dist}(p, p_j) + \text{dist}(p, p_{j-1})) \\ &\geq \text{dist}(p_j, p_{j-1}) \sum_{p \in B_j} \frac{w(p)}{w(B_j)} \end{aligned}$$

Hence, we have $\text{dist}(p_{j-1}, p_j) \leq \frac{2}{5} \cdot 10^j \leq 10^j/2$ for all $j = 1, 2, \dots, t$. It follows that any point $q \in B_{j-1}$ is also in B_j since by the triangle inequality $\text{dist}(q, p_j) \leq \text{dist}(q, p_{j-1}) + \text{dist}(p_{j-1}, p_j) \leq 10^{j-1} + 10^j/2 \leq 10^j$. Hence $B_{j-1} \subseteq B_j$. Finally, we again use the triangle inequality to conclude

$$\text{dist}(p_0, p_j) \leq \sum_{i=1}^j \text{dist}(p_{i-1}, p_i) \leq \frac{2}{5} \sum_{i=1}^j 10^i = \frac{2}{5} \frac{10^{j+1} - 2}{9} \leq 10^j/2.$$

□

The following lemma will be used to bound the cost of the center p_0 of a t -robust tuple in our analysis and it intuitively says that p_0 is a “good” center for any subset of P .

Lemma 4.5. *Let (p_0, p_1, \dots, p_t) be a t -robust tuple. For $i \in \{0, 1, \dots, t\}$ and a subset $P' \subseteq P$ of the points that contains $\text{Ball}(p_i, 10^i)$, i.e., $\text{Ball}(p_i, 10^i) \subseteq P'$, we have*

$$\text{cost}(p_0, P') \leq \frac{3}{2} \cdot \text{cost}(p_i, P').$$

Proof. For brevity, let $B_j = \text{Ball}(p_j, 10^j)$ for $j = 1, 2, \dots, t$. As (p_0, p_1, \dots, p_t) is t -robust, we have

$$\text{cost}(p_j, B_j) \geq \text{cost}(p_{j-1}, B_j) \quad \text{for } j = 1, \dots, t \quad (1)$$

Furthermore, a point $q \in P \setminus B_j$ has distance at least 10^j from p_j and we have $\text{dist}(p_0, p_j) \leq 10^j/2$ by Lemma 4.4. Therefore, (by the triangle inequality)

$$w(q)\text{dist}(p_j, q) \geq \frac{2}{3}w(q)\text{dist}(p_0, q) \quad \text{for any } q \in P \setminus B_j. \quad (2)$$

Repeatedly applying the Inequalities (1) and (2) and using $B_0 \subseteq B_1 \subseteq \dots \subseteq B_i \subseteq P'$ (by Lemma 4.4) give us

$$\begin{aligned} \text{cost}(p_i, P') &= \text{cost}(p_i, B_i) + \text{cost}(p_i, P' \setminus B_i) \\ &\geq \text{cost}(p_{i-1}, B_i) + \frac{2}{3}\text{cost}(p_0, P' \setminus B_i) \\ &= \text{cost}(p_{i-1}, B_{i-1}) + \text{cost}(p_{i-1}, B_i \setminus B_{i-1}) + \frac{2}{3}\text{cost}(p_0, P' \setminus B_i) \\ &\geq \text{cost}(p_{i-2}, B_{i-1}) + \frac{2}{3}\text{cost}(p_0, B_i \setminus B_{i-1}) + \frac{2}{3}\text{cost}(p_0, P' \setminus B_i) \\ &= \text{cost}(p_{i-2}, B_{i-1}) + \frac{2}{3}\text{cost}(p_0, P' \setminus B_{i-1}) \\ &\vdots \\ &\geq \text{cost}(p_0, B_1) + \frac{2}{3}\text{cost}(p_0, P' \setminus B_1) \end{aligned}$$

□

5 Description of EpochAlgorithm

In this section we describe EpochAlgorithm, the main new algorithm in our approach. We assume that

- the stream is compressed into $\tilde{O}(k)$ insertions of weighted points; and
- the insertions change the value of an optimum clustering by at most a factor of 18.

These assumptions are without loss of generality and follows from the application of known techniques as described in the overview. Indeed, by losing at most a constant-factor in the approximation guarantee and poly-logarithmic factors in the consistency, we can achieve the first assumption by an adaptation of Meyerson’s sketch (similar to what was previously used for this problem in [30]); and we can achieve the second assumption by restarting the algorithm every time the value of an optimum clustering increases by

a constant-factor (which can happen at most $O(\log(n\Delta))$ times). For a more formal description of this reduction please refer to Section 8.

The compressed stream is now divided into epochs. In each epoch we call EpochAlgorithm which takes as input the final clustering produced during the last epoch (or an initial solution if we consider the first epoch). In order to guarantee a constant approximate solution along with an amortized poly-logarithmic number of changes in the solution, we require that the set of input centers in each call to EpochAlgorithm has the following properties.

Definition 5.1. We call a set of centers \mathcal{U} bounded-robust if it has the following two properties.

1. \mathcal{U} is a 100-approximate solution.
2. Each center $u \in \mathcal{U}$ is t -robust where t is the smallest integer such that $10^t \geq \text{dist}(u, \mathcal{U} \setminus \{u\})/200$.

In the first epoch, we construct an initial bounded-robust solution to the point set as follows: we first obtain a 10-approximate solution \mathcal{W} to the initial point set using one of the known constant-factor approximation algorithms for k -median; and we then *robustify* the centers of \mathcal{W} as follows:

While there is a center $w \in \mathcal{W}$ that violates the second condition of bounded-robust, i.e., it is not t -robust for the smallest integer t such that $10^t \geq \text{dist}(w, \mathcal{W} \setminus \{w\})/200$:

1. Obtain t' -robust tuple $(w_0, w_1, \dots, w_{t'}) = \text{MakeRobust}(t', w)$ with $w_{t'} = w$ for smallest integer t' such that $10^{t'} \geq \text{dist}(w, \mathcal{W} \setminus w)/100$.
2. Remove w and add w_0 to the set of cluster centers.

We refer to the above procedure that makes every center in \mathcal{W} robust as $\text{Robustify}(\mathcal{W})$. We prove that MakeRobust is called at most once for each center (see Lemma 6.1) and so the above procedure terminates with a set of cluster centers that satisfies the second condition of Definition 5.1. For the analysis, we may actually assume that MakeRobust is also called at least (i.e., exactly) once for each center. That the procedure satisfies the first condition follows from the fact that the cost is only increased by a factor $3/2$ by robustifying the cluster centers in \mathcal{W} (see Lemma 7.5).

Let $\mathcal{U}^{(0)}$ denote the bounded-robust cluster centers obtained for the initial point set $P^{(0)}$. EpochAlgorithm first detects the number of centers, ℓ , that can be removed from the current solution $\mathcal{U}^{(0)}$ without affecting the quality of the approximation too much; then it handles $\ell + 1$ insertions, i.e., produces $\ell + 1$ solutions $U^{(1)}, U^{(2)}, \dots, U^{(\ell+1)}$ for point sets $P^{(1)}, P^{(2)}, \dots, P^{(\ell+1)}$, where $P^{(i)}$ denotes the set of points obtained from $P^{(0)}$ after i weighted insertions. The last step ensures that the set of centers $\mathcal{U}^{(\ell+1)}$ is bounded-robust and is thus a valid initial solution in the next epoch where EpochAlgorithm is called with $U^{(\ell+1)}$ as the “ $\mathcal{U}^{(0)}$ ”-solution and $P^{(\ell+1)}$ as the “ $P^{(0)}$ ” point set. EpochAlgorithm is repeatedly called in this way until all insertions are considered in the compressed stream.

We proceed to describe the steps of EpochAlgorithm in detail when given as input a set $\mathcal{U}^{(0)}$ of bounded-robust centers of points $P^{(0)}$:

Step 1: Removing Centers. In this preprocessing step, we remove all the centers that are not necessary in our solution. Basically we remove ℓ centers while increasing the cost of the solution by a factor at most $O(1)$. More precisely, if we let ℓ^* be the largest value such that it is possible to remove ℓ^* centers from $\mathcal{U}^{(0)}$ while increasing the cost by at most a factor $c = 228000$. Then we find a solution of size at most $k - \ell$ with $\ell \geq \ell^*$ and cost at most $3c \cdot \text{cost}(\mathcal{U}^{(0)}, P^{(0)})$. The details of this step is described in Section 9.2 (Theorem 9.3). In short, we use the standard LP to find the value of ℓ and then we apply known algorithmic techniques for the k -median problem.

Step 2: Handling Insertions. From the preprocessing step, we know that the size of the current solution is at most $k - \ell$. This enables us to simply open the next ℓ point that are inserted as centers, so the cost of the solution does not increase. We refer to the clustering and the set of points after the i -th insertion as $\mathcal{U}^{(i)}$ and $P^{(i)}$ for $1 \leq i \leq \ell$, respectively. Then, we consider one more insertion, i.e., the point set $P^{(\ell+1)}$. After this insertion the cost of the solution may increase significantly and furthermore the current solution may not be t -robust. We address both issues in the next step.

Step 3: Swapping Center and Robustify. In this postprocessing step, the first goal is to find a set of $O(\ell + 1)$ swaps that minimizes the objective function. This enables us to bound the approximation ratio of our approach and prove that it is a 100-approximate solution. To that end, for $\ell' = 5\ell + 5$, we use an LP-rounding procedure that swaps $4\ell'$ elements from $\mathcal{U}^{(0)}$ and produces \mathcal{W} . We show that \mathcal{W} is a 13-approximate solution with respect to the cost of an optimum solution \mathcal{W}_{OPT} that swaps at most ℓ' centers, with probability at least $1 - (n + \Delta)^{10}$. The details of this procedure is described in Section 9.1 (Theorem 9.1).

Finally, we make the solution bounded-robust so that it can be used in the next EpochAlgorithm using MakeRobust. More precisely, we robustify the cluster centers in \mathcal{W} to obtain $U^{(\ell+1)}$ by calling Robustify(\mathcal{W}) as we did for the initial solution (described in the gray box above).

Having described our main algorithm we proceed to its analysis. In the next section, we first bound the number of changes to the solution we make, i.e., the consistency. Then in Section 7 we bound the cost of the maintained solution. In particular, we show that $\mathcal{U}^{(\ell+1)}$ is indeed a bounded-robust solution and thus a valid input to the next call to EpochAlgorithm.

6 Bound on the Number of Changes

We analyze the consistency, i.e., the number of changes made to the maintained solution. As we will see, the number of changes essentially boils down to analyzing the number of calls to MakeRobust in Robustify. We first observe that Robustify calls MakeRobust at most once for each center. This implies that the procedure terminates and it is also a fact that will be used in the cost analysis. We then proceed to the main part of this section, which is the consistency analysis.

6.1 Robustify Makes a Center Robust At Most Once

We start by showing that Robustify calls MakeRobust at most once for each center. This guarantees that Robustify terminates and it will later also be used to bound the cost of the clustering.

Lemma 6.1. *Consider a set \mathcal{W} of centers. If Robustify(\mathcal{W}) calls MakeRobust for center $w \in \mathcal{W}$ and thus replaces w by w_0 , then it makes no subsequent call to MakeRobust for center w_0 .*

Proof. The statement follows intuitively due to the following. When a call to MakeRobust on center w is made, it is with the smallest t such that $10^t \geq \text{dist}(w, \mathcal{W} \setminus \{w\})/100$. Center w is then replaced by a nearby center w_0 and, for w_0 to be selected in a subsequent iteration, it must be that at that point $\text{dist}(w_0, \mathcal{W} \setminus \{w_0\})/200 > 10^t$, i.e., other centers are now more than a factor 2 further away from w_0 than they were for w . This cannot happen since calls to MakeRobust only changes the position of centers relatively little, see Lemma 4.4.

We proceed with the formal proof. Suppose toward contradiction that there is a center $w \in \mathcal{W}$ such that Robustify(\mathcal{W}) calls MakeRobust for center w , which is replaced by w_0 , and then in a subsequent iteration makes a call to MakeRobust for center w_0 .

We consider the first pair w, w_0 for which this happens. Let \mathcal{W}' be the set of centers when MakeRobust is called for w and let $w' \in \mathcal{W}' \setminus \{w\}$ be a center such that $\text{dist}(w, \mathcal{W}' \setminus \{w\}) = \text{dist}(w, w')$. Then Robustify calls MakeRobust(t, w) where t is selected to be the smallest integer such that $10^t \geq \text{dist}(w, w')/100$. In particular, we have $10^t \leq \text{dist}(w, w')/10$ and so by Lemma 4.4

$$\text{dist}(w_0, w) \leq 10^t/2 \leq \text{dist}(w, w')/20.$$

Now let \mathcal{W}'' be the set of centers when MakeRobust is called for w_0 , then we have that either $w' \in \mathcal{W}''$ or \mathcal{W}'' contains the center w'_0 that replaced w' via a single call to MakeRobust. These two cases are exhaustive since w and w_0 was the first pair such that an additional call for w_0 was made. In the first case,

$$\text{dist}(w_0, \mathcal{W}'') \leq \text{dist}(w_0, w') \leq \text{dist}(w_0, w) + \text{dist}(w, w') \leq \frac{21}{20} \text{dist}(w, w')$$

which contradicts that w_0 was selected in the while-loop in Robustify since we have that w_0 is t -robust with $10^t \geq \text{dist}(w, w')/100$ and thus $10^t \geq \text{dist}(w_0, \mathcal{W}'')/200$.

The second case is similar but we need in addition to argue that $\text{dist}(w'_0, w') \leq \text{dist}(w, w')/10$. To see this note that Robustify calls MakeRobust(t', w') with a t' such that $10^{t'} \leq \text{dist}(w', w_0)/10$. Hence, Lemma 4.4 says that $\text{dist}(w'_0, w') \leq 10^{t'}/2 \leq \text{dist}(w', w_0)/20$ which in turn is upper bounded by $(\text{dist}(w_0, w) + \text{dist}(w, w'))/20 \leq \text{dist}(w, w')/10$. We thus have in the second case that

$$\begin{aligned} \text{dist}(w_0, \mathcal{W}'') &\leq \text{dist}(w_0, w'_0) \leq \text{dist}(w_0, w) + \text{dist}(w, w') + \text{dist}(w', w'_0) \\ &\leq \text{dist}(w, w')/20 + \text{dist}(w, w') + \text{dist}(w, w')/10 \\ &= \frac{23}{20} \text{dist}(w, w'), \end{aligned}$$

which contradicts that w_0 was selected in the while-loop in Robustify in the same way as in the previous case. □

6.2 Consistency Analysis

In this section we focus on analyzing the consistency of the algorithm. Recall that $m = \tilde{O}(k)$ denotes the number of weighted insertions in our compressed stream.

Theorem 6.2. *The number of changes to the maintained solution is at most $O(m(\log \Delta)^2)$.*

We use the following notation. Let E denote the total number of epochs. We use the convention that we subscript quantities in the call of EpochAlgorithm during the e -th epoch by e . So, for $1 \leq e \leq E$, the set of centers $\mathcal{U}_e^{(0)}$ denotes the initial solution to point set $P_e^{(0)}$ in the call to EpochAlgorithm at the start of the e -th epoch, and $\mathcal{U}_e^{(1)}, \mathcal{U}_e^{(2)}, \dots, \mathcal{U}_e^{(\ell_e+1)}$ denote the solutions produced to point sets $P_e^{(1)}, P_e^{(2)}, \dots, P_e^{(\ell_e+1)}$ during this run of EpochAlgorithm. With this notation, we have $\mathcal{U}_e^{(\ell_e+1)} = \mathcal{U}_{e+1}^{(0)}$ and $P_e^{(\ell_e+1)} = P_{e+1}^{(0)}$ for $e = 1, 2, \dots, E-1$. Moreover, the consistency equals

$$\sum_{e=1}^E \sum_{i=0}^{\ell_e} |\mathcal{U}_e^{(i+1)} \setminus \mathcal{U}_e^{(i)}|. \quad (3)$$

The consistency $\sum_{i=0}^{\ell_e} |\mathcal{U}_e^{(i+1)} \setminus \mathcal{U}_e^{(i)}|$ of epoch e is at most $\ell_e + |\mathcal{U}_e^{(\ell_e+1)} \setminus \mathcal{U}_e^{(\ell_e)}|$ since EpochAlgorithm opens up exactly one center for $i = 1, 2, \dots, \ell$. We further have

$$|\mathcal{U}_e^{(\ell_e+1)} \setminus \mathcal{U}_e^{(\ell_e)}| \leq |\mathcal{U}_e^{(\ell_e+1)} \setminus \mathcal{U}_e^{(0)}| + |\mathcal{U}_e^{(0)} \setminus \mathcal{U}_e^{(\ell_e)}| = |\mathcal{U}_e^{(\ell_e+1)} \setminus \mathcal{U}_e^{(0)}| + O(\ell_e).$$

This allows us to upper bound Eq. (3) by

$$\sum_{e=1}^E \left(|\mathcal{U}_e^{(\ell_e+1)} \setminus \mathcal{U}_e^{(0)}| + O(\ell_e) \right) \leq \sum_{e=1}^E |\mathcal{U}_e^{(\ell_e+1)} \setminus \mathcal{U}_e^{(0)}| + O(m),$$

where we used that $\sum_{e=1}^E (\ell_e + 1) = m$.

To analyze $|\mathcal{U}_e^{(\ell_e+1)} \setminus \mathcal{U}_e^{(0)}|$, recall that EpochAlgorithm constructs $\mathcal{U}_e^{(\ell_e+1)}$ in two steps: it first obtains an intermediate solution \mathcal{W}_e by swapping $4\ell'_e = 20(\ell_e + 1)$ centers from $\mathcal{U}_e^{(0)}$; and it then calls Robustify(\mathcal{W}_e) to obtain $\mathcal{U}_e^{(\ell_e+1)}$.

We distinguish two kind of centers in $\mathcal{U}_e^{(\ell_e+1)} \setminus \mathcal{U}_e^{(0)}$, the updated centers and the new centers. To this end, consider a center $w \in \mathcal{W}_e \cap \mathcal{U}_e^{(0)}$. If $w \notin \mathcal{U}_e^{(\ell_e+1)}$, then w was replaced by a new center w_0 via call to MakeRobust in the Robustify procedure. We will say that w is the *parent* of w_0 and we say that w_0 is an *updated* center. Those centers in $\mathcal{U}_e^{(\ell_e+1)} \setminus \mathcal{U}_e^{(0)}$ that are not updated (i.e., without a parent) are referred to as *new* centers. In addition the centers in the initial solution computer for the first epoch, are also classified as new centers.

It will be convenient to think of these centers as elements of chains formed as follows. Consider the graph that has a vertex for each center in $\mathcal{U}_e^{(\ell_e+1)} \setminus \mathcal{U}_e^{(0)}$ for each epoch $1 \leq e \leq E$ and there is an arc from u to u_0 if u is the parent of u_0 . We remark that the graph may have multiple vertices for a single center, if that center was added and removed from the solution multiple times. As each vertex has at most one parent and it is the parent of at most one center, we have the graph forms a collection of paths. We refer to these paths as chains. Moreover, if we consider a chain (u_1, u_2, \dots, u_s) then u_1 is a new center (without a parent), u_{i+1} replaced u_i via a call to MakeRobust for $i = 1, 2, \dots, s - 1$, and u_s is not the parent of any center.

By definition, we have that $\sum_{e=1}^E |\mathcal{U}_e^{(\ell_e+1)} \setminus \mathcal{U}_e^{(0)}|$ equals the number of vertices in the above-described graph. Or equivalently, and this is the viewpoint that we take, it equals the total lengths of the chains. We start to bound the number of chains.

Lemma 6.3. *The total number of chains (or equivalently, new centers) is at most $20 \cdot m$.*

Proof. Consider an epoch e . We have that a center in $\mathcal{W}_e \cap \mathcal{U}_e^{(0)}$ is either in $\mathcal{U}_e^{(\ell_e+1)}$ or it is the parent of an updated center in $\mathcal{U}_e^{(\ell_e+1)}$. It follows that the number of new centers in $\mathcal{U}_e^{(\ell_e+1)} \setminus \mathcal{U}_e^{(0)}$ is at most $|\mathcal{U}_e^{(\ell_e+1)}| - |\mathcal{W}_e \cap \mathcal{U}_e^{(0)}|$. Now \mathcal{W}_e is obtained from $\mathcal{U}_e^{(0)}$ by performing at most $4\ell'_e = 20(\ell_e + 1)$ swaps and so $|\mathcal{W}_e \cap \mathcal{U}_e^{(0)}| \geq k - 20(\ell_e + 1)$. Hence, we can upper bound the number of new centers in $\mathcal{U}_e^{(\ell_e+1)} \setminus \mathcal{U}_e^{(0)}$ by $20(\ell_e + 1)$. It follows that the number of new centers across all epochs is upper bounded by $20 \sum_{e=1}^E (\ell_e + 1) = 20 \cdot m$. Note that chains can start only in new centers by the definition of new centers. \square

We now give certain properties that must hold when Robustify calls MakeRobust for a center w and then replaces it. To this end, we use the following corollary of Definition 4.3 and Lemma 4.4, to identify points that get inserted and invalidate the robustness guarantees maintained by the algorithm.

Corollary 6.4. *Consider a point set P , and suppose that the center $u \in P$ is t -robust. Then u is t -robust for every superset $P' \supseteq P$ of points satisfying $\text{Ball}_P(u, 2 \cdot 10^t) = \text{Ball}_{P'}(u, 2 \cdot 10^t)$, i.e., $P' \setminus P$ does not contain any point within distance $2 \cdot 10^t$ from u .*

Proof. Let $p_0 = u$ and let (p_0, p_1, \dots, p_t) be a t -robust sequence with respect to P . Then Definition 4.3 and Lemma 4.4 imply that (p_0, p_1, \dots, p_t) is also a t -robust sequence with respect to P' if $\text{Ball}_P(p_t, 10^t) = \text{Ball}_{P'}(p_t, 10^t)$. The statement now follows from Lemma 4.4 which says that $\text{dist}(p_0, p_t) \leq 10^t/2$ and thus any point in a ball of radius 10^t around p_t is contained in a ball of radius $2 \cdot 10^t$ around $p_0 = u$. \square

For a center u that is in one of the maintained solutions, define the integer $t(u)$ as follows. If u is a new center, let $t(u) = 0$. Otherwise if u is an updated center, then we let $t(u)$ be the integer t used in the call to $\text{MakeRobust}(t, u')$ by Robustify when it replaced u 's parent u' by u .

Now consider an updated center u_0 with parent u . Let f be the epoch when u joined the solution for the last time before u_0 replaced u and let e be the epoch when it was replaced. As Robustify replaced u with u_0 , at least one of the following two cases must hold:

- Center u is not $t(u)$ -robust with respect to the bigger point sets $P_e^{(\ell_e+1)}$. By the above corollary, there must then be a point in $p \in P_e^{(\ell_e+1)} \setminus P_f^{(\ell_f+1)}$ within distance $2 \cdot 10^{t(u)}$ from u . We say that the point p invalidates the center u in this case.
- The integer $t(u)$ is too small: When $w = u$ is selected in the while-loop of Robustify , we have $10^{t(u)} < \text{dist}(w, \mathcal{W} \setminus \{w\})/200$. Then u_0 is obtained via a call to $\text{MakeRobust}(t', u)$ where t' is the smallest integer such that $10^{t'} > \text{dist}(w, \mathcal{W} \setminus \{w\})/100$. In particular, $t(u_0) = t' \geq t + 1$. In this case, we say that the updated center u_0 *increased its robustness*.

We first bound the number of updates that increases the robustness in terms of the number of new centers and other updated centers. We then bound the number of updated centers that were invalidated by points.

Lemma 6.5. *The number of updated centers that increased their robustness is at most $O(\log \Delta)$ times the number of new centers plus the number of updated centers that were invalidated by points.*

Proof. Consider a chain (u_0, u_1, \dots, u_s) . We have that u_0 is a new center. The statement now follows by observing that we cannot have a subsequence of updated centers u_i, u_{i+1}, \dots, u_j so that everyone increases their robustness $j - i > \log \Delta$. Indeed, in that case, we have $t(u_j) \geq j - i > \log \Delta$ which is a contradiction because for any center u we have that $10^{t(u)}$ is at most the maximum distance Δ . Therefore a chain of length s must include at least $\lfloor s / \log(\Delta) \rfloor$ centers that are invalidated by points and for which the robustness is not increased. \square

Lemma 6.6. *The number of updated centers that were invalidated by points is at most $m \cdot \log \Delta$.*

Proof. The stream consists of m point insertions so it is sufficient to prove that each point p invalidates at most $\log \Delta$ many updated centers. We bound the number of centers that are invalidated by a point p introduced in epoch e , i.e., $p \in P_e^{(\ell+1)} \setminus P_{e-1}^{(\ell+1)}$. Let μ be the number of centers in $\mathcal{W}_e \cap \mathcal{U}_e^{(0)}$ that are invalidated by p , and let $(u_i)_{i \in \mu}$ be the sequence of these centers ordered in decreasing order by the epoch that they were added to the maintained solution: letting e_i equal the first epoch when $u_i \in \mathcal{U}_{e_i}^{(\ell+1)}$, we order the sequence so that $e_i > e_{i+1}$ for $i = 1, 2, \dots, \mu - 1$.

This ordering guarantees that when u_i replaced its parent u'_i via a call to MakeRobust then u_{i+1} was already in the solution and so

$$10^{t(u_i)} \leq \text{dist}(u_i, u_{i+1})/10 \quad \text{for } i \in \{1, 2, \dots, \mu - 1\}.$$

In addition, as p invalidates u_i , we have

$$\text{dist}(p, u_i) \leq 2 \cdot 10^{t(u_i)} \quad \text{for } i \in \{1, 2, \dots, \mu\}.$$

It follows that for $i \in \{1, 2, \dots, \mu - 1\}$

$$10^{t(u_{i+1})} \geq \text{dist}(p, u_{i+1}) \geq \text{dist}(u_i, u_{i+1}) - \text{dist}(p, u_i) \geq 10 \cdot 10^{t(u_i)} - 2 \cdot 10^{t(u_i)} \geq 8 \cdot 10^{t(u_i)}.$$

We can now conclude the proof by observing that due to the bounded aspect ratio Δ , $\mu \leq \log \Delta$. \square

Equipped with the above lemmas, we can complete the proof of Theorem 6.2. Lemma 6.3 says that there are $O(m)$ new centers and Lemma 6.6 says that there are at most $O(m \log \Delta)$ number of updated centers that were invalidated by points. Combining this with Lemma 6.5 yields that the total number of new and updated centers is at most $O(m(\log \Delta)^2)$. As the number of new and updated centers is equal to the consistency, this completes the proof of the theorem.

7 Cost Analysis of EpochAlgorithm

Recall that EpochAlgorithm runs during $\ell+1$ steps and maintains sets of centers $\mathcal{U}^{(0)}, \mathcal{U}^{(1)}, \dots, \mathcal{U}^{(\ell+1)}$ for the point sets $P^{(0)}, P^{(1)}, \dots, P^{(\ell+1)}$. We bound their costs as follows, recall that $c = 228000$.

Theorem 7.1. *On input a bounded-robust set of centers $\mathcal{U}^{(0)}$ of point set $P^{(0)}$, EpochAlgorithm produces sets of centers $\mathcal{U}^{(1)}, \mathcal{U}^{(2)}, \dots, \mathcal{U}^{(\ell+1)}$ of the point sets $P^{(0)}, P^{(1)}, \dots, P^{(\ell+1)}$ satisfying:*

- *the clustering induced by $\mathcal{U}^{(i)}$ is a $(6 \cdot 100 \cdot c)$ -approximate clustering for $P^{(i)}$ for $i = 1, 2, \dots, \ell$; and*
- *$\mathcal{U}^{(\ell+1)}$ is a bounded-robust set of centers for $P^{(\ell+1)}$.*

The first part of the theorem follows easily: As $\mathcal{U}^{(0)}$ is a bounded-robust set of centers of $P^{(0)}$ it is a 100-approximation. Now EpochAlgorithm drops ℓ centers from $\mathcal{U}^{(0)}$ while increasing the cost by at most a factor $3 \cdot c$. The cost of the solutions $\mathcal{U}^{(1)}, \mathcal{U}^{(2)}, \dots, \mathcal{U}^{(\ell)}$ does not change (since the newly arrived points are opened as centers) and thus they remain $(2 \cdot 3 \cdot 100 \cdot c)$ -approximate solutions. The factor 2 is due to the fact that the value of an optimum solution can at most go down by a factor of 2 by introducing new points.

The remaining part of this section is thus devoted to proving the second part of the theorem, i.e., that $\mathcal{U}^{(\ell+1)}$ is a bounded-robust set of centers for $P^{(\ell+1)}$. Recall that Step 3 of EpochAlgorithm calls Robustify(\mathcal{W}) procedure which repeatedly call MakeRobust until every center in $u \in \mathcal{W}$ is t -robust with $10^t \geq \text{dist}(u, \mathcal{W} \setminus \{u\})/100$. Therefore it suffice to showing that $\mathcal{U}^{(\ell+1)}$ induce a 100-approximate clustering of $P^{(\ell+1)}$.

An important concept for bounding the cost of $\mathcal{U}^{(\ell+1)}$ is the notion of well-separated pairs. To simplify notation, we let \mathcal{U} denote $\mathcal{U}^{(0)}$ throughout this section. We also let \mathcal{V} be a fixed optimal solution to $P^{(\ell+1)}$.

Definition 7.2. *Select $\gamma = 2000$. We say that centers $u \in \mathcal{U}, v \in \mathcal{V}$ form a well-separated pair if*

$$\text{dist}(u, \mathcal{U} \setminus \{u\}) \geq \gamma \cdot \text{dist}(u, v) \quad \text{and} \quad \text{dist}(v, \mathcal{V} \setminus \{v\}) \geq \gamma \cdot \text{dist}(u, v).$$

Informally, $u \in \mathcal{U}$ and $v \in \mathcal{V}$ form a well-separated pair if they are isolated in the sense that they are much closer to each other than the distance to any other center. In particular, by definition, if u and v form a well-separated pair, then v must be the closest center to u in \mathcal{V} and u must be the closest center to v in \mathcal{U} (see Fig. 2a). We can therefore partition the centers in $\mathcal{U} \cup \mathcal{V}$ into the well-separated pairs $(u_1, v_1), (u_2, v_2), \dots, (u_{k-m}, v_{k-m})$ and the remaining centers $u_{k-m+1}, \dots, u_k, v_{k-m+1}, \dots, v_k$ that do not form a well-separated pair with any center. In Section 7.1 we prove the following lemma which relates the number of well-separated pairs with the number of centers we can drop from \mathcal{U} without increasing the cost too much.

Lemma 7.3. *Suppose that the number of well-separated pairs is $k - m$. Then there exists a clustering $\mathcal{U}' \subseteq \mathcal{U}$ with at most $k - \lfloor m/4 \rfloor$ centers and whose cost is bounded by*

$$\text{cost}(\mathcal{U}', P^{(0)}) \leq 6\gamma \left(\text{cost}(\mathcal{U}, P^{(0)}) + \text{cost}(\mathcal{V}, P^{(0)}) \right).$$

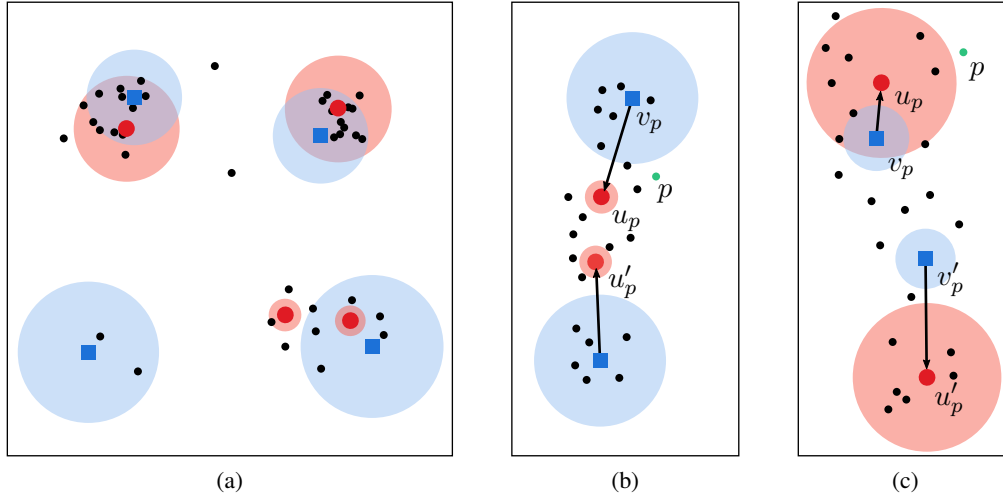


Figure 2: Each figure shows two sets of centers \mathcal{U} (red circles) and \mathcal{V} (blue squares). For each center $u \in \mathcal{U}$, $\text{Ball}(u, \text{dist}(u, \mathcal{U} \setminus \{u\})/\gamma)$ is depicted, and similarly for centers $v \in \mathcal{V}$ (for this figure only, we assume $\gamma = 4$). Therefore $(u, v) \in \mathcal{U} \times \mathcal{V}$ is a well-separated pair if and only if u lies within the ball of v and vice versa. (a) The upper two pairs are well-separated, the lower centers are not well-separated. (b) Fractional assignment: the case that $y_{u_p} = 1/2$ and there is a center $u'_p \in \mathcal{U} \setminus \{u_p\}$ such that $\text{dist}(u_p, u'_p) \leq \gamma \cdot \text{dist}(u_p, v_p)$. (c) Fractional assignment: the case that $y_{u_p} = 1/2$ and the previous case (b) does not hold.

By the assumption that the value of an optimum clustering increases by at most a factor 18 during an epoch, we have, by the optimality of \mathcal{V} , that $\text{cost}(\mathcal{V}, P^{(0)})$ is within a factor of 18 of the cost of an optimal clustering of $P^{(0)}$. It follows that $6\gamma (\text{cost}(\mathcal{U}, P^{(0)}) + \text{cost}(\mathcal{V}, P^{(0)})) \leq (6\gamma \cdot (18 + 1)) \text{cost}(\mathcal{U}, P^{(0)})$. As $c = 228000 = 6\gamma \cdot 19$, this implies that Step 1 of EpochAlgorithm removes $\ell \geq \lfloor m/4 \rfloor$ centers where m is the number of centers of \mathcal{U} that does not form a well-separated pair with a center in \mathcal{V} .

During Step 3, EpochAlgorithm creates an intermediate clustering \mathcal{W} whose cost is 13-approximate with respect to the cost of an optimum solution \mathcal{W}_{OPT} that swaps at most $\ell' = 4\ell + 4 + \ell + 1 \geq m + \ell + 1$ centers. We now bound the cost of such an optimum solution \mathcal{W}_{OPT} to be at most 3 times the cost of an optimum clustering of $P^{(\ell+1)}$. Let $\mathcal{V}(v)$ be the points in $P^{(\ell+1)}$ that are closest to center $v \in \mathcal{V}$ in the clustering \mathcal{V} . We say that a center $u \in \mathcal{U}$ is *good* if

- it forms a well-separated pair with a center in \mathcal{V} ; and
- $\mathcal{V}(v) \subseteq P^{(0)}$, i.e., the center v is not closest to any of the new points in the considered epoch.

The number of centers of \mathcal{U} that are bad, i.e., not good, is at most the number of centers that do not form well-separated pairs (m many) plus the number of centers that form a well-separated pair with a center $v \in \mathcal{V}$ such that $\mathcal{V}(v) \cap (P^{(\ell+1)} \setminus P^{(0)}) \neq \emptyset$ (at most $\ell + 1$ many). Denote the good centers of \mathcal{U} by u_1, u_2, \dots, u_g and index the centers of $\mathcal{V} = \{v_1, v_2, \dots, v_k\}$ so that u_i form a well-separated pair with v_i for $i \in \{1, \dots, g\}$. Then the clustering $\{u_1, \dots, u_g, v_{g+1}, \dots, v_k\}$ is obtained from \mathcal{U} by doing $k - g \leq m + \ell + 1 \leq \ell'$ swaps.

It follows that the cost of \mathcal{W}_{OPT} is bounded by the cost of this solution and so

$$\begin{aligned}
\text{cost}(\mathcal{W}, P^{(\ell+1)}) &\leq 13 \cdot \text{cost}(\mathcal{W}_{\text{OPT}}, P^{(\ell+1)}) \\
&\leq 13 \cdot \left(\sum_{i=1}^g \text{cost}(u_i, \mathcal{V}(v_i)) + \sum_{i=g+1}^k \text{cost}(v_i, \mathcal{V}(v_i)) \right) \\
&\leq 13 \cdot \left(\sum_{i=1}^g 3 \cdot \text{cost}(v_i, \mathcal{V}(v_i)) + \sum_{i=g+1}^k \text{cost}(v_i, \mathcal{V}(v_i)) \right) \\
&= 39 \cdot \text{cost}(\mathcal{V}, P^{(\ell+1)}),
\end{aligned}$$

where the first inequality follows from Theorem 9.1 and the last inequality is due to the following lemma which we prove in Section 7.2.

Lemma 7.4. *If $u \in \mathcal{U}, v \in \mathcal{V}$ form a well-separated pair and $\mathcal{V}(v) \subseteq P^{(0)}$,*

$$\text{cost}(u, \mathcal{V}(v)) \leq 3 \cdot \text{cost}(v, \mathcal{V}(v)).$$

Now EpochAlgorithm obtains $\mathcal{U}^{(\ell+1)}$ by calling MakeRobust on centers in \mathcal{W} . We bound the cost of $\mathcal{U}^{(\ell+1)}$ by the following lemma.

Lemma 7.5. *For any clustering \mathcal{W} , let \mathcal{U} be the result of Robustify(\mathcal{W}) procedure described in Section 5 for a set of points P . We have*

$$\text{cost}(\mathcal{U}, P) \leq \frac{3}{2} \text{cost}(\mathcal{W}, P).$$

Proof. By Lemma 6.1, Robustify makes at most one call to MakeRobust per center. Let w_1, \dots, w_k denote the centers of \mathcal{W} and let w'_1, \dots, w'_k denote the centers of \mathcal{U} where w'_j was obtained by a call to MakeRobust on w_j or $w'_j = w_j$ if no such call was made. Further index the centers in the order in which the calls to MakeRobust were made, putting those centers last for which no call was made.

With this notation, EpochAlgorithm calls MakeRobust on center w_j with input parameter t_j selected to be the smallest integer so that

$$10^{t_j} \geq \text{dist}(w_j, \{w'_1, \dots, w'_{j-1}, w_{j+1}, \dots, w_k\})/100.$$

This implies that $10^{t_j} \leq \text{dist}(w_j, w'_i)/10$ for any $i < j$ and $10^{t_j} \leq \text{dist}(w_j, w_i)/10$ for any $i > j$. Lemma 4.4 says that $\text{dist}(w_j, w'_j) \leq 10^{t_j}/2$. Hence, for two centers w_i and w_j with $i < j$,

$$10 \cdot 10^{t_j} \leq \text{dist}(w'_i, w_j) \leq \text{dist}(w_i, w_j) + \text{dist}(w_i, w'_i) \leq \text{dist}(w_i, w_j) + 10^{t_i}/2 \leq \frac{11}{10} \text{dist}(w_i, w_j),$$

where the last inequality follows as by the properties of the ordering, i.e., $t_i \leq t_j$. If we let $\mathcal{W}(w_j)$ be the points of P assigned to center w_j in the clustering $\mathcal{W} = \{w_1, w_2, \dots, w_k\}$, it follows that $\text{Ball}(w_j, 10^{t_j}) \subseteq \mathcal{W}(w_j)$ since $10^{t_j} < \text{dist}(w_j, w_i)/2$ for any $i \neq j$. We can therefore apply Lemma 4.5 (with $P' = \mathcal{W}(w_j)$) to bound the increase of the cost by

$$\text{cost}(\mathcal{U}, P) = \sum_{j=1}^k \text{cost}(w'_j, \mathcal{W}(w_j)) \leq \frac{3}{2} \sum_{j=1}^k \text{cost}(w_j, \mathcal{W}(w_j)) = \frac{3}{2} \text{cost}(\mathcal{W}, P).$$

□

By the above lemma, we conclude

$$\text{cost} \left(U^{(\ell+1)}, P^{(\ell+1)} \right) \leq \frac{3}{2} \cdot \text{cost} \left(\mathcal{W}, P^{(\ell+1)} \right) \leq 100 \cdot \text{cost} \left(\mathcal{V}, P^{(\ell+1)} \right),$$

i.e., that $U^{(\ell+1)}$ is a 100-approximate solution. As aforementioned this implies that $U^{(\ell+1)}$ is bounded-robust which in turn implies Theorem 7.1. It remains to prove Lemma 7.3 and Lemma 7.4 which we do in Section 7.1 and Section 7.2, respectively.

7.1 Proof of Lemma 7.3: Relating Non-Well-Separated and Removed Centers

Let $k - m$ be the number of well-separated pairs in \mathcal{U} and \mathcal{V} . We prove Lemma 7.3, i.e., that there is a clustering $\mathcal{U}' \subseteq \mathcal{U}$ obtained from \mathcal{U} by removing at least $\lfloor m/4 \rfloor$ centers and whose cost is bounded by

$$\text{cost} \left(\mathcal{U}', P^{(0)} \right) \leq 6\gamma \left(\text{cost} \left(\mathcal{U}, P^{(0)} \right) + \text{cost} \left(\mathcal{V}, P^{(0)} \right) \right).$$

On a very high level, the reason why the above holds is that if centers are not too far away from each other then we can safely close some of them. The formal proof strategy is as follows. We define a feasible solution (y, x) to the standard LP relaxation with the potential center locations \mathcal{U} and point set P such that

- $\sum_{u \in \mathcal{U}} y_u \leq k - m/4$;
- the cost is bounded by $\sum_{u \in \mathcal{U}, p \in P} x_{up} \text{dist}(u, p) \leq 2\gamma \left(\text{cost} \left(\mathcal{U}, P^{(0)} \right) + \text{cost} \left(\mathcal{V}, P^{(0)} \right) \right)$.

To this end, let us recall the standard LP relaxation for the weighted k -median problem. In the following LP, $w(p)$ represents the weight of the point p for any $p \in P$,

$$\min \sum_{u \in \mathcal{U}, p \in P} x_{up} \cdot w(p) \cdot \text{dist}(u, p) \tag{4}$$

$$\text{s.t. } x_{up} \leq y_u \quad \forall u \in \mathcal{U}, p \in P \tag{5}$$

$$\sum_{u \in \mathcal{U}} x_{up} \geq 1 \quad \forall p \in P \tag{6}$$

$$\sum_{u \in \mathcal{U}} y_u \leq k - \frac{m}{4} \tag{7}$$

$$x_{up}, y_u \geq 0 \quad \forall u \in \mathcal{U}, p \in P \tag{8}$$

Lemma 7.3 then follows because the integrality gap of the standard LP relaxation is known to be at most 3 [7]. We proceed to describe the fractional opening of the centers (the setting of y), the fractional assignment of points to centers (the setting of x), and we then analyze the cost of the solution.

Fractional Opening of Centers. Let $\pi : \mathcal{V} \rightarrow \mathcal{U}$ be the function that maps each center \mathcal{V} to its closest center in \mathcal{U} (breaking ties arbitrarily). Partition \mathcal{U} into \mathcal{U}_I and \mathcal{U}_F , where \mathcal{U}_I contains those centers $u \in \mathcal{U}$ that satisfies at least one of the following conditions:

- u forms a well-separated pair with one center in \mathcal{V} ;
- $|\pi^{-1}(u)| \geq 2$.

So \mathcal{U}_F contains those centers $u \in \mathcal{U}$ that does *not* form a well-separated pair and $|\pi^{-1}(u)| \leq 1$. We define a solution to the standard LP in which every center $u \in \mathcal{U}_I$ is integrally opened $y_u = 1$ and each center $u \in \mathcal{U}_F$ is fractionally opened $y_u = 1/2$.

We now show that $\sum_{u \in \mathcal{U}} y_u \leq k - m/4$. Each center $u \in \mathcal{U}$ that forms a well-separated pair with a center $v \in \mathcal{V}$ has $|\pi^{-1}(u)| \geq 1$ since u must be the closest center to v in \mathcal{U} . In addition, as there are $k - m$ well-separated pairs in total, we get

$$k \geq \sum_{u \in \mathcal{U}_I} |\pi^{-1}(u)| \geq k - m + 2 \cdot (|\mathcal{U}_I| - (k - m)) ,$$

and so $|\mathcal{U}_I| \leq k - m/2$. We thus have $\sum_{u \in \mathcal{U}} y_u \leq k - m/2 + 1/2 \cdot m/2 = k - m/4$ as required.

Fractional Assignment of Points. Each point $p \in P$ is assigned as follows. Let $v_p \in \mathcal{V}$ be the closest center to p in \mathcal{V} and let $u_p = \pi(v_p)$ be the closest center in \mathcal{U} to v_p . We distinguish three cases:

- If $y_{u_p} = 1$ then assign point p to u_p , i.e., set $x_{u_p p} = 1$. This assignment costs $w(p) \cdot \text{dist}(u_p, p)$.
- If $y_{u_p} = 1/2$ and there is a center $u'_p \in \mathcal{U} \setminus \{u_p\}$ such that $\text{dist}(u_p, u'_p) \leq \gamma \cdot \text{dist}(u_p, v_p)$, then assign p equally to u_p and u'_p , i.e., set $x_{u_p p} = x_{u'_p p} = 1/2$ (see Fig. 2b). This assignment costs

$$\begin{aligned} \frac{1}{2} w(p) (\text{dist}(p, u_p) + \text{dist}(p, u'_p)) &\leq \frac{1}{2} w(p) (\text{dist}(p, u_p) + \text{dist}(p, u_p) + \text{dist}(u_p, u'_p)) \\ &\leq w(p) \left(\text{dist}(p, u_p) + \frac{\gamma}{2} \text{dist}(u_p, v_p) \right) . \end{aligned}$$

- If $y_{u_p} = 1/2$ and the previous case does not hold, then since u_p, v_p does not form a well-separated pair, there is a center $v'_p \in \mathcal{V} \setminus \{v_p\}$ such that $\text{dist}(v_p, v'_p) \leq \gamma \cdot \text{dist}(u_p, v_p)$ (see Fig. 2c). We assign p equally to u_p and $u'_p = \pi(v'_p)$, i.e., we set $x_{u_p p} = x_{u'_p p} = 1/2$. This assignment costs

$$\begin{aligned} \frac{1}{2} (\text{dist}(p, u_p) + \text{dist}(p, u'_p)) &\leq \frac{1}{2} w(p) (\text{dist}(p, u_p) + \text{dist}(p, u_p) + \text{dist}(u_p, v'_p) + \text{dist}(v'_p, u'_p)) \\ &\leq \frac{1}{2} w(p) (\text{dist}(p, u_p) + \text{dist}(p, u_p) + \text{dist}(u_p, v'_p) + \text{dist}(v'_p, u_p)) \\ &= w(p) (\text{dist}(p, u_p) + \text{dist}(u_p, v'_p)) \\ &\leq w(p) (\text{dist}(p, u_p) + \text{dist}(u_p, v_p) + \text{dist}(v_p, v'_p)) \\ &\leq w(p) (\text{dist}(p, u_p) + (\gamma + 1) \text{dist}(u_p, v_p)) . \end{aligned}$$

In the third case, we have $u_p \neq u'_p$ since $\pi(v_p) = u_p$ and $|\pi^{-1}(u_p)| \leq 1$ since $y_{u_p} = 1/2$ and therefore $u_p \in \mathcal{U}_F$. We have thus that each point is fractionally assigned one unit and that $x_{u_p} \leq y_u$ for all $u \in \mathcal{U}$ and $p \in P$. In other words, x is a feasible fractional assignment of the points.

Bounding the Cost. By the above calculations, the cost of assigning a point $p \in \mathcal{V}(v)$ is always at most $w(p) (\text{dist}(p, u_p) + (\gamma + 1) \text{dist}(u_p, v_p))$, with equality in the third case. We thus have that the total cost is upper bounded by

$$\begin{aligned} \sum_{p \in P} w(p) (\text{dist}(p, u_p) + (\gamma + 1) \text{dist}(u_p, v_p)) &\leq \sum_{p \in P} w(p) (\text{dist}(p, v_p) + (\gamma + 2) \text{dist}(u_p, v_p)) \\ &= \text{cost}(\mathcal{V}, P^{(0)}) + (\gamma + 2) \sum_{p \in P} w(p) \text{dist}(u_p, v_p) . \end{aligned}$$

To bound the last term, for a point $p \in P$, let u_p^* denote its closest center in \mathcal{U} . Then by the definition of $u_p = \pi(v_p)$ and the triangle inequality,

$$\text{dist}(u_p, v_p) \leq \text{dist}(u_p^*, v_p) \leq \text{dist}(u_p^*, p) + \text{dist}(p, v_p).$$

We thus have

$$\sum_{p \in P} w(p) \text{dist}(u_p, v_p) \leq \sum_{p \in P} w(p) (\text{dist}(u_p^*, p) + \text{dist}(p, v_p)) = \text{cost}(\mathcal{U}, P^{(0)}) + \text{cost}(\mathcal{V}, P^{(0)})$$

and the cost of the fractional solution is at most

$$\text{cost}(\mathcal{V}, P^{(0)}) + (\gamma + 2) \left(\text{cost}(\mathcal{U}, P^{(0)}) + \text{cost}(\mathcal{V}, P^{(0)}) \right) \leq 2\gamma \left(\text{cost}(\mathcal{U}, P^{(0)}) + \text{cost}(\mathcal{V}, P^{(0)}) \right),$$

as required.

7.2 Proof of Lemma 7.4: Cost of Well-Separated Pairs

In this section we prove Lemma 7.4, i.e., the statement that if $u \in \mathcal{U}, v \in \mathcal{V}$ form a well-separated pair and $\mathcal{V}(v) \subseteq P^{(0)}$, then

$$\text{cost}(u, \mathcal{V}(v)) \leq 3 \cdot \text{cost}(v, \mathcal{V}(v)).$$

The statement is trivial if $\text{dist}(u, v) = 0$ so we assume throughout that $\text{dist}(u, v) \geq 1$.

Since \mathcal{U} is a bounded-robust solution, we have that u is t -robust where t is the smallest integer such that

$$10^t \geq \text{dist}(u, \mathcal{U} \setminus \{u\})/100 \geq (\gamma/100) \cdot \text{dist}(u, v) = 20 \cdot \text{dist}(u, v), \quad (9)$$

where the second inequality holds because u and v form a well-separated pair and the equality is due to the selection of $\gamma = 2000$. Select t^* to be the integer that satisfies $20 \cdot \text{dist}(u, v) \leq 10^{t^*} < 10 \cdot 20 \cdot \text{dist}(u, v)$. As $t^* \leq t$, there is a t^* -robust tuple $(p_0, p_1, \dots, p_{t^*})$ with $t^* \geq 1$ and $p_0 = u$. Let $B_i = \text{Ball}(p_i, 10^i)$ for $i = 0, 1, \dots, t^*$. We shall use the following simple fact:

Claim 7.6. For every $i = 0, 1, \dots, t^*$, $B_i \subseteq \mathcal{V}(v)$.

Proof of claim. Consider a point $q \in B_i$. We have

$$\begin{aligned} \text{dist}(q, v) &\leq \text{dist}(q, p_i) + \text{dist}(p_i, u) + \text{dist}(u, v) \\ &= \text{dist}(q, p_i) + \text{dist}(p_i, p_0) + \text{dist}(u, v) \\ &\leq 10^i + 10^i/2 + 10^{t^*}/20 \\ &\leq 2 \cdot 10^{t^*}, \end{aligned}$$

where, for the penultimate inequality, we used $q \in B_i$, Lemma 4.5, and the selection of t^* . That $q \in \mathcal{V}(v)$ and thus $B_i \subseteq \mathcal{V}(v)$ now follows because u and v form a well-separated pair and so by the selection of t^* (specifically, that $10^{t^*} < 200 \text{dist}(u, v)$)

$$\text{dist}(v, \mathcal{V} \setminus \{v\}) \geq 2 \cdot 10^3 \cdot \text{dist}(u, v) > 10^{t^*+1} \geq 2 \cdot \text{dist}(q, v).$$

□

We divide the remaining part of the proof of the lemma into two cases. Suppose first that $\text{avgcost}(p_{t^*}, B_{t^*}) \geq 10^{t^*}/5$. We then have $p_{t^*-1} = p_{t^*}$ and so

$$\begin{aligned} \text{dist}(v, p_{t^*}) &= \text{dist}(v, p_{t^*-1}) \leq \text{dist}(v, p_0) + \text{dist}(p_0, p_{t^*-1}) \\ &= \text{dist}(u, v) + \text{dist}(p_0, p_{t^*-1}) \\ &\leq 10^{t^*}/20 + 10^{t^*}/20 \\ &\leq 10^{t^*}/10, \end{aligned} \tag{10}$$

where the penultimate inequality holds by the selection of t^* and because $\text{dist}(p_0, p_{t^*-1}) \leq 10^{t^*}/20$ by Lemma 4.5. Hence, since $\text{avgcost}(p_{t^*}, B_{t^*}) \geq 10^{t^*}/5 \geq 2 \cdot \text{dist}(v, p_{t^*})$ and $B_{t^*} \subseteq \mathcal{V}(v)$ (by the above claim),

$$\begin{aligned} \text{cost}(v, \mathcal{V}(v)) &= \text{cost}(v, B_{t^*}) + \text{cost}(v, \mathcal{V}(v) \setminus B_{t^*}) \\ &= w(B_{t^*}) \cdot \text{avgcost}(v, B_{t^*}) + \text{cost}(v, \mathcal{V}(v) \setminus B_{t^*}) \\ &\geq w(B_{t^*}) \cdot (\text{avgcost}(p_{t^*}, B_{t^*}) - \text{dist}(v, p_{t^*})) + \text{cost}(v, \mathcal{V}(v) \setminus B_{t^*}) \\ &\geq w(B_{t^*}) \cdot \frac{\text{avgcost}(p_{t^*}, B_{t^*})}{2} + \text{cost}(v, \mathcal{V}(v) \setminus B_{t^*}) \\ &\geq \frac{\text{cost}(p_{t^*}, \mathcal{V}(v))}{2}, \end{aligned}$$

where the last inequality holds because any point $q \in \mathcal{V}(v) \setminus B_{t^*}$ has

$$\text{dist}(q, v) \geq \text{dist}(p_{t^*}, q) - \text{dist}(p_{t^*}, v) \geq 10^{t^*} - \text{dist}(p_{t^*}, v) \geq \text{dist}(p_{t^*}, v)$$

where the first inequality follows by triangle inequality, second inequality follows since $q \in \mathcal{V}(v) \setminus B_{t^*}$, and last inequality holds due to Eq. (10). We can therefore conclude the proof of the lemma in this case by applying Lemma 4.5 which says that $\text{cost}(p_0, \mathcal{V}(v)) \leq (3/2) \cdot \text{cost}(p_{t^*}, \mathcal{V}(v))$.

It remains to consider when $\text{avgcost}(p_{t^*}, B_{t^*}) < 10^{t^*}/5$. We then have

$$p_{t^*-1} = \arg \min_{p \in B_{t^*}} \text{cost}(p, B_{t^*}).$$

Therefore, since $B_{t^*} \subseteq \mathcal{V}(v)$ (by the above claim)

$$\begin{aligned} \text{cost}(p_{t^*-1}, \mathcal{V}(v)) &= \text{cost}(p_{t^*-1}, B_{t^*}) + \text{cost}(p_{t^*-1}, \mathcal{V}(v) \setminus B_{t^*}) \\ &\leq \text{cost}(v, B_{t^*}) + \text{cost}(p_{t^*-1}, \mathcal{V}(v) \setminus B_{t^*}), \end{aligned}$$

where the last inequality follows since $v \in B_{t^*}$ from Eq. (9). To bound the second term, we use Lemma 4.5 to obtain

$$\text{dist}(q, p_{t^*-1}) \geq \text{dist}(q, p_{t^*}) - \text{dist}(p_{t^*}, p_{t^*-1}) \geq 10^{t^*} - 10^{t^*}/2 \quad \text{for any } q \in \mathcal{V}(v) \setminus B_{t^*},$$

and

$$\begin{aligned} \text{dist}(v, p_{t^*-1}) &\leq \text{dist}(v, u) + \text{dist}(u, p_{t^*-1}) \\ &= \text{dist}(v, u) + \text{dist}(p_0, p_{t^*-1}) \\ &\leq 10^{t^*}/20 + 10^{t^*}/20 \leq 10^{t^*}/10. \end{aligned}$$

Hence,

$$\frac{\text{dist}(p_{t^*-1}, q)}{\text{dist}(v, q)} \leq \frac{\text{dist}(p_{t^*-1}, q)}{\text{dist}(p_{t^*-1}, q) - \text{dist}(p_{t^*-1}, v)} \leq \frac{10^{t^*}/2}{10^{t^*}/2 - 10^{t^*}/10} \leq 2$$

and so

$$\text{cost}(p_{t^*-1}, \mathcal{V}(v)) \leq \text{cost}(v, B_{t^*-1}) + \text{cost}(p_{t^*-1}, \mathcal{V}(v) \setminus B_{t^*-1}) \leq 2 \cdot \text{cost}(v, \mathcal{V}(v)).$$

We can thus also conclude this case by applying Lemma 4.5 which give us $\text{cost}(u, \mathcal{V}(v)) \leq 3 \cdot \text{cost}(v, \mathcal{V}(v))$.

8 Consistent k -Clustering Algorithm and Analysis

In this section we present our consistent algorithm. One of the ingredients used in this algorithm is Meyerson's sketch. In this work we use a modified version of the Meyerson's sketch presented in [35]. Intuitively, Multi-Meyerson's sketch is an algorithm that produces a weighted instance of size $k \cdot \text{polylog}(n, \Delta)$ that any constant approximate solution of this instance is also a constant approximate solution for the original instance. These properties are explained more precisely in the following theorem. Notice that in this theorem, the elements are inserted one by one and there is no need to have access to the entire stream at the beginning.

Theorem 8.1. *[Multi-Meyerson's procedure] Given a set of initial points P' and a stream of insertions of points $\sigma' = \langle \sigma'_1, \sigma'_2, \dots, \sigma'_{t'} \rangle$ such that*

$$18\text{OPT}(P') \geq \text{OPT}(P' \cup \{\sigma'_1, \sigma'_2, \dots, \sigma'_{t'}\}),$$

Multi-Meyerson's procedure finds a weighted initial solution P and a stream of insertions of weighted points $\sigma = \langle \sigma_1, \sigma_2, \dots, \sigma_t \rangle$. The stream is built on the fly, and satisfies the following properties with probability at least $1 - \frac{1}{(n+\Delta)^{10}}$.

1. *The length of the produced stream σ is small,*

$$t \leq k \cdot \text{polylog}(n, \Delta).$$

2. *For any $1 \leq j \leq t'$, let $\langle \sigma_1, \sigma_2, \dots, \sigma_i \rangle$ be the points on the stream σ after the j -th insertion of stream σ' . Let \mathcal{U} be a set of centers inducing an α -approximate solution for the weighted instance defined by $P \cup \{\sigma_1, \sigma_2, \dots, \sigma_i\}$, then \mathcal{U} is also a $O(\alpha + 1)$ -approximate solution for $P' \cup \{\sigma'_1, \sigma'_2, \dots, \sigma'_j\}$. I.e.,*

$$\text{cost}(\mathcal{U}, P' \cup \{\sigma'_1, \sigma'_2, \dots, \sigma'_j\}) \leq O(\alpha + 1) \cdot \text{OPT}(P' \cup \{\sigma'_1, \sigma'_2, \dots, \sigma'_j\}).$$

3. *The weights are positive integers and the total sum of the weights is at most $2n$.*

A full description of the Multi-Meyerson procedure along with a formal proof for this theorem is presented in Appendix A.

Our algorithm first runs the Multi-Meyerson procedure, therefore we need to detect when the cost of the optimum solution increases. To that end, independent from Meyerson's sketch we compute an α -approximate solution (for $\alpha \leq 3$) for the entire instance after each insertion. Afterwards when it increases by a factor 6, we get that the value of the optimum solution has increased by a factor at least 2 and at most 18. In this case restart the Multi-Meyerson procedure. This results in at most $\log(n\Delta)$ restarts of this algorithm and guarantees the condition in Theorem 8.1 is satisfied.

At the beginning of each run of the Multi-Meyerson procedure we compute a bounded-robust solution by simply finding a 3-approximate solution for the the initial points produced by the Multi-Meyerson procedure. Then we make it robust, as explained in Section 5. By Lemma 7.5 we lose a factor $\frac{3}{2}$ in this step, therefore the resulting solution is a $\frac{9}{2}$ -approximate solution. We pass it to EpochAlgorithm as the initial solution. We keep running EpochAlgorithm until the stream σ is finished, where each call to it uses the last solution created by the previous call. Notice that we know that the final solution created by EpochAlgorithm is bounded-robust. Let us now analyze the approximation ratio and the total number of the changes in the solution.

Approximation Ratio. Notice that the only place that we output a solution is in `EpochAlgorithm` which we showed that is a constant approximate algorithm in Theorem 7.1. Combining it with Theorem 8.1, we get that the solution produced is indeed a constant approximate solution throughout the algorithm.

Consistency of the Solution. At the beginning of each run of `Multi-Meyerson`, we compute a new solution which might result in k changes compared to the previous solution. Therefore the total changes caused by it is at most $k \log(n\Delta)$. The rest of the changes happens during the `EpochAlgorithm`, and by Theorem 6.2 we get that the total number of changes is at most a factor $O(\log^2 \Delta)$ more than the total length of the streams which is $k \cdot \text{polylog}(n, \Delta)$ by Theorem 8.1, therefore we conclude that the total number of the changes is at most $k \cdot \text{polylog}(n, \Delta)$.

Success Probability. The algorithm is successful if no error happens in the following three steps.

1. Swapping the centers: The probability that this step fails is also at most $\frac{1}{(n+\Delta)^{10}}$ from Theorem 9.1 and is called at most $O(n \log(n\Delta))$ times.
2. Multi-Meyerson procedure: The probability that this step fails is at most $\frac{1}{(n+\Delta)^{10}}$ from Theorem 8.1 and is called at most $O(\log(n\Delta)^2)$ times.

Putting together these two properties results in

Theorem 3.1. *There exists an algorithm that given a stream of points $\sigma = \langle \sigma_1, \sigma_2, \dots, \sigma_n \rangle$, outputs a set of at most k centers \mathcal{U}_i after the i -th insertion for all $1 \leq i \leq n$ so that with probability at least $1 - \frac{1}{n^8}$*

1. \mathcal{U}_i induces a constant approximation clustering, i.e., for any $1 \leq i \leq n$

$$\text{cost}(\mathcal{U}_i, \{\sigma_1, \sigma_2, \dots, \sigma_i\}) \leq O(1) \cdot \text{OPT}(P_i).$$

2. The total number of changes in consecutive center sets is at most $k \cdot \text{polylog}(n\Delta)$.

$$\sum_{1 \leq i < n} |\mathcal{U}_{i+1} \setminus \mathcal{U}_i| \leq k \cdot \text{polylog}(n, \Delta).$$

Removing the Assumption that Δ and n are Known in Advance. Throughout this paper, we assumed that we know the value of the Δ and n in advance. This knowledge is not necessary in order to run our algorithm. Notice that we can keep track of both n and Δ during the execution of the algorithm. When one of them increases by a constant factor, we restart our algorithm from the scratch. This results losing a constant factor in approximation guarantee, a $\text{polylog}(n, \Delta)$ factor in total number of changes in consecutive center sets, and the success probability becomes $1 - O(\frac{1}{n^8})$.

9 LP-based Algorithm for Swapping and Removing Centers

In this section we design two LP-based algorithms for the following tasks:

Swapping Centers. Compute a set of k centers that differs from an input set of centers in at most $O(\ell)$ centers and such that the cost of their induced clustering is a constant-factor approximation to the best possible clustering whose centers differ in at most ℓ from the input set of centers.

Removing Centers. Compute a minimal subset of the input centers that achieves a constant-factor approximation to the input instance.

Throughout this section we assume that the instance is not weighted and we copy the points according to their weights. Notice that from Theorem 8.1 we know that the weights of the points are positive integers with total summation of at most $2n$.

9.1 LP-based Algorithm for Swapping Centers

In this section we show

Theorem 9.1. *There exists an algorithm that given a set of centers \mathcal{U} , a set of points P and a parameter ℓ , finds a new set of centers \mathcal{V} with probability at least $1 - \frac{1}{(n+\Delta)^{10}}$, such that*

1. *These two sets of centers differ in at most 4ℓ centers, i.e.,*

$$|\mathcal{U} \setminus \mathcal{V}| \leq 4\ell.$$

2. *The cost of \mathcal{V} is at most a factor 13 larger than the best set of centers \mathcal{W} that differs with V in at most ℓ centers.*

$$\text{cost}(\mathcal{V}, P) \leq 13 \cdot \min_{\mathcal{W}: |\mathcal{U} \setminus \mathcal{W}| \leq \ell} \text{cost}(\mathcal{W}, p).$$

To this end, we use the following LP-relaxation, where y_0 is equal to one if $i \in \mathcal{U}$ and zero otherwise.

$$\min \sum_{i,j \in P} x(i,j) \text{dist}(i,j) \tag{11}$$

$$\text{s.t. } x(i,j) \leq y(i) \quad \forall i,j \in P \tag{12}$$

$$\sum_{i \in P} x(i,j) \geq 1 \quad \forall j \in P \tag{13}$$

$$\sum_{i \in P} y(i) \leq k \tag{14}$$

$$v(i) \geq y(i) - y_0(i) \quad \forall i \in P \tag{15}$$

$$\sum_{i \in P} v(i) \leq \ell \tag{16}$$

$$v(i), x(i,j), y(i) \geq 0 \quad \forall i,j \in P \tag{17}$$

Intuitively, Eq. (11) minimizes the cost of the solution. Eq. (12) ensures that a center is open before a point is assigned to it. Eq. (13) ensures that all the point are assigned to a center. Eq. (14) ensures that there are at most k open centers. Eq. (15) and Eq. (16) ensure that the solution found is at most a value ℓ different from the initial solution y_0 . Together with Eq. (17) that ensure all the variables are positive, gives that

$$\sum_{i \in P} \max\{0, y(i) - y_0(i)\} \leq \ell.$$

We refer to the optimum solution of the above LP as y^* , x^* and w^* and let OPT denote its cost. One can observe that for any set of centers \mathcal{W} such that $|\mathcal{U} \setminus \mathcal{W}| \leq \ell$, we have

$$\text{cost}(\mathcal{W}, P) \geq \text{OPT}. \tag{18}$$

Afterwards we use the rounding procedure described in [9] over the y^* and x^* fractional solution and obtain a random solution y and x . Therefore, we get that

Lemma 9.2 (Restated from [9]). *There is a rounding procedure that opens at most k centers and the probability that i is open is exactly $y^*(i)$. Moreover, in expectation, the cost of the solution is at most a factor 3.25 more than the fractional solution x^* .*

Using Markov's inequality, we get that with probability at least $3/4$

$$\sum_{i \in P} \max\{0, y(i) - y_0(i)\} \leq 4\ell,$$

and similarly we get that with probability at least $3/4$

$$\text{cost}(y, P) \leq 3.25 \cdot 4 \sum_{i, j \in P} x^*(i, j) \text{dist}(i, j) \leq 13 \cdot \text{OPT}.$$

Combining with Eq. (18) we get that for any set of centers \mathcal{W} such that $|\mathcal{U} \setminus \mathcal{W}| \leq \ell$,

$$\text{cost}(y, P) \leq 13 \text{cost}(\mathcal{W}, P).$$

Therefore, by applying the union bound we get that the solution has both above properties with probability at least $1/2$. By repeating the rounding algorithm $O(\log(n + \Delta))$ times, we can boost the probability to $1 - \frac{1}{(n+\Delta)^{10}}$.

9.2 LP-based Algorithm for Polynomial Time Removing Centers

In this section we show that

Theorem 9.3. *There exists an algorithm that given a set of centers \mathcal{U} , a set of points P and a parameter $c > 1$, finds a set of centers \mathcal{V} such that*

1. $\mathcal{V} \subseteq \mathcal{U}$ such that

$$\text{cost}(\mathcal{V}, P) \leq 3c \cdot \text{cost}(\mathcal{U}, P).$$

2. For any set of centers $\mathcal{W} \subseteq \mathcal{U}$ such that $\text{cost}(\mathcal{W}, P) \leq c \cdot \text{cost}(\mathcal{U}, P)$, we have

$$|\mathcal{V}| \leq |\mathcal{W}|.$$

We use the classic LP for the k-median problem with parameter ℓ that controls the size of the solution.

$$\min \sum_{i \in \mathcal{U}, j \in P} x(i, j) \text{dist}(i, j) \tag{19}$$

$$\text{s.t. } x(i, j) \leq y(i) \quad \forall i \in \mathcal{U}, \forall j \in P \tag{20}$$

$$\sum_{i \in \mathcal{U}} x(i, j) \geq 1 \quad \forall j \in P \tag{21}$$

$$\sum_{i \in \mathcal{U}} y(i) \leq k - \ell \tag{22}$$

$$x(i, j), y(i) \geq 0 \quad \forall i, j \in P \tag{23}$$

Intuitively, Eq. (19) minimizes the cost of the solution. Eq. (20) ensures that a center is open before a point is assigned to it. Eq. (21) ensures that all the point are assigned to a center. Eq. (22) ensures that there are at most $k - \ell$ open centers. Eq. (23) ensures that all the variables are positive.

Given a value c , we try all possible k values of ℓ and return the largest one such that the cost of the above LP is at most a factor c away from the cost of the optimum solution (denoted by OPT). After the value ℓ is fixed, we can use any algorithm for the k -median problem with the number of centers equal to $k - \ell$, set of point P and set of potential centers \mathcal{U} . Therefore we get a 3-approximate solution \mathcal{V} such that, $|\mathcal{V}| \leq k - \ell$ and $\text{cost}(\mathcal{V}, P) \leq 3\text{OPT} \leq 3c \cdot \text{cost}(\mathcal{U}, P)$.

We know that there exist no fractional solution that has size less than $k - \ell$ and cost less than $c \cdot \text{cost}(\mathcal{U}, P)$, therefore there exist no such a integral solution as well. This means that any solution $\mathcal{W} \subseteq \mathcal{U}$ such that $\text{cost}(\mathcal{W}, P) \leq c \cdot \text{cost}(\mathcal{U}, P)$, we get that $|\mathcal{W}| \geq k - \ell$. Therefore,

$$|\mathcal{V}| \leq k - \ell \leq |\mathcal{W}|.$$

References

- [1] S. Ahmadian, A. Norouzi-Fard, O. Svensson, and J. Ward. Better guarantees for k-means and euclidean k-median by primal-dual algorithms. In *Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on*, pages 61–72. Ieee, 2017.
- [2] M. Andrews, M. X. Goemans, and L. Zhang. Improved bounds for on-line load balancing. *Algorithmica*, 23(4):278–301, 1999.
- [3] D. Arthur and S. Vassilvitskii. k-means++: The advantages of careful seeding. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1027–1035. Society for Industrial and Applied Mathematics, 2007.
- [4] V. Arya, N. Garg, R. Khandekar, A. Meyerson, K. Munagala, and V. Pandit. Local search heuristics for k-median and facility location problems. *SIAM Journal on computing*, 33(3):544–562, 2004.
- [5] A. Bernstein, J. Holm, and E. Rotenberg. Online bipartite matching with amortized $o(\log 2n)$ replacements. *Journal of the ACM (JACM)*, 66(5):1–23, 2019.
- [6] J. Byrka, T. Pensyl, B. Rybicki, A. Srinivasan, and K. Trinh. An improved approximation for k-median, and positive correlation in budgeted optimization. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 737–756. SIAM, 2014.
- [7] T. Carnes and D. B. Shmoys. Primal-dual schema and lagrangian relaxation for the k-location-routing problem. In L. A. Goldberg, K. Jansen, R. Ravi, and J. D. P. Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 14th International Workshop, APPROX 2011, and 15th International Workshop, RANDOM 2011, Princeton, NJ, USA, August 17-19, 2011. Proceedings*, volume 6845 of *Lecture Notes in Computer Science*, pages 99–110. Springer, 2011.
- [8] M. Charikar and S. Guha. Improved combinatorial algorithms for facility location problems. *SIAM Journal on Computing*, 34(4):803–824, 2005.
- [9] M. Charikar and S. Li. A dependent lp-rounding approach for the k-median problem. In A. Czumaj, K. Mehlhorn, A. M. Pitts, and R. Wattenhofer, editors, *Automata, Languages, and Programming - 39th International Colloquium, ICALP 2012, Warwick, UK, July 9-13, 2012, Proceedings, Part I*, volume 7391 of *Lecture Notes in Computer Science*, pages 194–205. Springer, 2012.

- [10] M. Charikar, L. O’Callaghan, and R. Panigrahy. Better streaming algorithms for clustering problems. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 30–39. ACM, 2003.
- [11] K. Chen. On coresets for k-median and k-means clustering in metric and euclidean spaces and their applications. *SIAM Journal on Computing*, 39(3):923–947, 2009.
- [12] V. Cohen-Addad, N. O. D. Hjuler, N. Parotsidis, D. Saulpic, and C. Schwiegelshohn. Fully dynamic consistent facility location. In *Advances in Neural Information Processing Systems*, pages 3255–3265, 2019.
- [13] L. Epstein and A. Levin. Robust algorithms for preemptive scheduling. *Algorithmica*, 69(1):26–57, 2014.
- [14] D. Feldman and M. Langberg. A unified framework for approximating and clustering data. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 569–578. ACM, 2011.
- [15] D. Feldman, M. Monemizadeh, and C. Sohler. A ptas for k-means clustering based on weak coresets. In *Proceedings of the twenty-third annual symposium on Computational geometry*, pages 11–18. ACM, 2007.
- [16] D. Feldman, M. Schmidt, and C. Sohler. Turning big data into tiny data: Constant-size coresets for k-means, pca and projective clustering. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, pages 1434–1453. Society for Industrial and Applied Mathematics, 2013.
- [17] D. Fotakis. Online and incremental algorithms for facility location. *ACM SIGACT News*, 42(1):97–131, 2011.
- [18] A. Gu, A. Gupta, and A. Kumar. The power of deferral: maintaining a constant-competitive steiner tree online. *SIAM Journal on Computing*, 45(1):1–28, 2016.
- [19] X. Guo, J. Kulkarni, S. Li, and J. Xian. Consistent k -median: Simpler, better and robust. *arXiv preprint arXiv:2008.06101*, 2020.
- [20] A. Gupta, R. Krishnaswamy, A. Kumar, and D. Panigrahi. Online and dynamic algorithms for set cover. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 537–550, 2017.
- [21] A. Gupta and A. Kumar. Online steiner tree with deletions. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 455–467. SIAM, 2014.
- [22] A. Gupta, A. Kumar, and C. Stein. Maintaining assignments online: Matching, scheduling, and flows. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 468–479. SIAM, 2014.
- [23] S. Har-Peled and A. Kushal. Smaller coresets for k-median and k-means clustering. *Discrete & Computational Geometry*, 37(1):3–19, 2007.
- [24] S. Har-Peled and S. Mazumdar. On coresets for k-means and k-median clustering. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 291–300. ACM, 2004.

- [25] M. Imase and B. M. Waxman. Dynamic steiner tree problem. *SIAM Journal on Discrete Mathematics*, 4(3):369–384, 1991.
- [26] M. R. K. Jaghargh, A. Krause, S. Lattanzi, and S. Vassilvitskii. Consistent online optimization: Convex and submodular. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2241–2250, 2019.
- [27] K. Jain, M. Mahdian, E. Markakis, A. Saberi, and V. V. Vazirani. Greedy facility location algorithms analyzed using dual fitting with factor-revealing lp. *Journal of the ACM (JACM)*, 50(6):795–824, 2003.
- [28] T. Kanungo, D. M. Mount, N. S. Netanyahu, C. D. Piatko, R. Silverman, and A. Y. Wu. A local search approximation algorithm for k-means clustering. *Computational Geometry*, 28(2-3):89–112, 2004.
- [29] J. Łacki, J. Oćwieja, M. Pilipczuk, P. Sankowski, and A. Zych. The power of dynamic distance oracles: Efficient dynamic algorithms for the steiner tree. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 11–20, 2015.
- [30] S. Lattanzi and S. Vassilvitskii. Consistent k-clustering. In *International Conference on Machine Learning*, pages 1975–1984, 2017.
- [31] S. Li. A 1.488 approximation algorithm for the uncapacitated facility location problem. In *International Colloquium on Automata, Languages, and Programming*, pages 77–88. Springer, 2011.
- [32] S. Lloyd. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137, 1982.
- [33] N. Megow, M. Skutella, J. Verschae, and A. Wiese. The power of recourse for online mst and tsp. In *International Colloquium on Automata, Languages, and Programming*, pages 689–700. Springer, 2012.
- [34] R. R. Mettu and C. G. Plaxton. The online median problem. *SIAM Journal on Computing*, 32(3):816–832, 2003.
- [35] A. Meyerson. Online facility location. In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, pages 426–431, 2001.
- [36] S. Phillips and J. Westbrook. Online load balancing and network flow. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 402–411, 1993.
- [37] P. Sanders, N. Sivadasan, and M. Skutella. Online scheduling with bounded migration. *Mathematics of Operations Research*, 34(2):481–498, 2009.
- [38] M. Skutella and J. Verschae. A robust ptas for machine covering and packing. In *European Symposium on Algorithms*, pages 36–47. Springer, 2010.
- [39] J. Westbrook. Load balancing for response time. *Journal of Algorithms*, 35(1):1–16, 2000.

A Multi-Meyerson Procedure

We first present Meyerson’s sketch which depends on an estimate of the optimum solution. Our approach is logically the same as [35] but not identical. Then we prove that it has the desired guarantees in expectation,

and use Markov inequality to achieve that the properties are preserved with constant probability. The guarantees needed here are slightly stronger than the ones used in previous works, but the proof techniques are similar. Next, we boost the probability to $1 - 1/n^{10}$ by running copies of Meyerson's sketch. In this step we need to be extra careful with the weights of the points in the instance as it is important that the weights are only increasing. Specifically, we prove the following:

Theorem 8.1. [Multi-Meyerson's procedure] *Given a set of initial points P' and a stream of insertions of points $\sigma' = \langle \sigma'_1, \sigma'_2, \dots, \sigma'_{t'} \rangle$ such that*

$$18\text{OPT}(P') \geq \text{OPT}(P' \cup \{\sigma'_1, \sigma'_2, \dots, \sigma'_{t'}\}),$$

Multi-Meyerson's procedure finds a weighted initial solution P and a stream of insertions of weighted points $\sigma = \langle \sigma_1, \sigma_2, \dots, \sigma_t \rangle$. The stream is built on the fly, and satisfies the following properties with probability at least $1 - \frac{1}{(n+\Delta)^{10}}$.

1. *The length of the produced stream σ is small,*

$$t \leq k \cdot \text{polylog}(n, \Delta).$$

2. *For any $1 \leq j \leq t'$, let $\langle \sigma_1, \sigma_2, \dots, \sigma_i \rangle$ be the points on the stream σ after the j -th insertion of stream σ' . Let \mathcal{U} be a set of centers inducing an α -approximate solution for the weighted instance defined by $P \cup \{\sigma_1, \sigma_2, \dots, \sigma_i\}$, then \mathcal{U} is also a $O(\alpha + 1)$ -approximate solution for $P' \cup \{\sigma'_1, \sigma'_2, \dots, \sigma'_j\}$. I.e.,*

$$\text{cost}(\mathcal{U}, P' \cup \{\sigma'_1, \sigma'_2, \dots, \sigma'_j\}) \leq O(\alpha + 1) \cdot \text{OPT}(P' \cup \{\sigma'_1, \sigma'_2, \dots, \sigma'_j\}).$$

3. *The weights are positive integers and the total sum of the weights is at most $2n$.*

We describe the Multi-Meyerson Procedure and prove its correctness in the following subsections.

A.1 Meyerson's Sketch

Meyerson's sketching operates on a guess of an optimum solution denoted by GOPT. We can use many of the known algorithms for the k -median problem to achieve a GOPT such that

$$\text{OPT}(P') \leq \text{GOPT} \leq 3\text{OPT}(P').$$

Let us start by describing Meyerson's sketch and its analysis. For sake of simplicity, we add the point in the set $P' = \{p'_1, \dots, p'_{|P'|}\}$ to the beginning of stream σ' in arbitrarily order resulting in

$$\sigma' = \langle p'_1, \dots, p'_{|P'|}, \sigma'_1, \dots, \sigma'_{t'} \rangle.$$

Description of Meyerson's Sketch. The algorithm maintains a solution S , i.e., a set of centers S where initially $S = \emptyset$. Then the Meyerson's sketch as follows:

Insertion of point x : We open x with probability $h_x \leftarrow \min\{1, \text{dist}(x, S) \frac{k(\log \Delta + 1)}{\text{GOPT}}\}$ (where $\text{dist}(x, \emptyset)$ is defined to be ∞). If x is opened, we add it to S and assign it to itself. Otherwise we assign x to the closest open center in S denoted by f_x .

Analysis of Meyerson's Sketch. For completeness we start by analyzing the performance of Meyerson's Sketch. Note that the analysis is similar to the analysis presented in [35]. We first bound the size of the dynamic set S after all insertions, i.e., after insertion of σ'_t . We let OPT_i denote the value of the optimum solution after i -th insertion for $1 \leq i \leq m$ and where m is the length of the stream, i.e., $m = t' + |P'|$. Moreover, we let P'_i and S_i denote the set of points and the solution maintained by the Meyerson's sketch after i -th insertion for $1 \leq i \leq m$.

Lemma A.1. *We have*

$$\mathbb{E}[|S_m|] \leq \beta k, \quad \text{where } \beta = \left(1 + 4 \cdot \frac{\text{OPT}_m}{\text{GOPT}}\right) k(\log \Delta + 1).$$

Proof. Let C^* indicate the clusters of an optimum solution after insertions, i.e., of an optimal clustering P'_m . We partition the points according to C^* . Consider a center $c^* \in C^*$ and let $P(c^*)$ be the subset of points that is served by c^* in the optimum solution. Moreover, recall that $\text{avgcost}(c^*, P(c^*)) = \text{cost}(c^*, P(c^*)) / |P(c^*)|$. Now partition the points $P(c^*)$ as follows: Let $P(c^*, 0) = \{x \in P(c^*) \mid \text{dist}(x, c^*) \leq \text{avgcost}(c^*, P(c^*))\}$ be those points of $P(c^*)$ whose distance to c^* is at most the average cost of the cluster, and similarly, let $P(c^*, i) = \{x \in P(c^*) \mid 2^{i-1} \text{avgcost}(c^*, P(c^*)) < \text{dist}(x, c^*) \leq 2^i \text{avgcost}(c^*, P(c^*))\}$. Notice that the total number of partitions that we have is $k(\log \Delta + 1)$, since all the non zero pairwise distances are in $[1, \Delta]$. Now consider one of the cluster centers c^* and a partition $P(c^*, i)$. The expected summation of h_x value for $x \in P(c^*, i)$ before we open a center is one. Moreover, if a point $y_{(c^*, i)}$ is opened in this partition then for any point afterwards, $h_x \leq \text{dist}(x, y_{(c^*, i)}) \frac{k(\log \Delta + 1)}{\text{GOPT}}$. Let us bound $\text{dist}(x, y_{(c^*, i)})$. By triangular inequality we have:

$$\begin{aligned} \text{dist}(x, y_{(c^*, i)}) &\leq \text{dist}(x, c^*) + \text{dist}(y_{(c^*, i)}, c^*) \\ &\leq \text{dist}(x, c^*) + 2\text{dist}(x, c^*) + \text{avgcost}(c^*, P(c^*)) \\ &= 3\text{dist}(x, c^*) + \text{avgcost}(c^*, P(c^*)), \end{aligned} \tag{24}$$

where the last inequality follows since by assumption x and $y_{(c^*, i)}$ are always in the same partition, so if $x \in P(c^*, 0)$ then $\text{dist}(y_{(c^*, i)}, c^*) \leq \text{avgcost}(c^*, P(c^*))$ otherwise $\text{dist}(y_{(c^*, i)}, c^*) \leq 2\text{dist}(x, c^*)$. Hence,

$$\begin{aligned} \mathbb{E}[|S|] &= \mathbb{E}\left[\sum_{x \in P} h_x\right] \\ &\leq \sum_{x \in P} \mathbb{E}[h_x] \\ &\leq \sum_{(c^*, i)} \left(1 + \sum_{x \in P(c^*, i)} \text{dist}(x, y_{(c^*, i)}) \frac{k(\log \Delta + 1)}{\text{GOPT}}\right) \\ &\leq \sum_{(c^*, i)} \left(1 + \sum_{x \in P(c^*, i)} (3\text{dist}(x, c^*) + \text{avgcost}(c^*, P(c^*))) \frac{k(\log \Delta + 1)}{\text{GOPT}}\right) \\ &\leq (k(\log \Delta + 1) + 4\text{OPT}_m \frac{k(\log \Delta + 1)}{\text{GOPT}}) \\ &= \beta k. \end{aligned}$$

□

We proceed to analyze the expected cost of the solution that is maintained by Meyerson's sketch.

Lemma A.2. For $1 \leq j \leq m$ we have

$$\mathbb{E}\left[\sum_{x \in P'_j} \text{dist}(x, f_x)\right] \leq 4\text{OPT}_j + \text{GOPT},$$

where f_x is set when x is inserted as described above.

Proof. We also inherit the notations from the previous proof Lemma A.1. Now consider one of the cluster centers c^* and the partition $P(c^*, i)$. Similar to the previous lemma the expected summation of the costs of points before we open a center is $\frac{\text{GOPT}}{k(\log \Delta + 1)}$. Moreover, let $y_{(c^*, i)}$ be the first point opened in this partition then for any point inserted afterwards in this partition, $\text{dist}(x, f_x) \leq \text{dist}(x, y_{(c^*, i)})$. Hence,

$$\begin{aligned} \mathbb{E}\left[\sum_{x \in P'_j} \text{dist}(x, f_x)\right] &\leq \sum_{(c^*, i)} \sum_{x \in P(c^*, i)} \mathbb{E}[\text{dist}(x, f_x)] \\ &\leq \sum_{(c^*, i)} \sum_{x \in P(c^*, i)} \mathbb{E}[\text{dist}(x, y_{(c^*, i)})] \\ &\leq \sum_{(c^*, i)} \left(\frac{\text{GOPT}}{k(\log \Delta + 1)} + \sum_{x \in P(c^*, i)} (3\text{dist}(x, c^*) + \text{avgcost}(c^*, P(c^*))) \right) \\ &\leq \text{GOPT} + (3\text{OPT}_j + \text{OPT}_j), \end{aligned}$$

where the third inequality follows from Eq. (24). □

A.2 Extending Meyerson's Sketch to Multi-Meyerson

In this section we show that we can achieve similar bounds as the previous section with high probability. To this end, we first show the properties that one can achieve by running copies of Meyerson's sketch. Afterwards we use this approach and present the Multi-Meyerson procedure.

Recall that $\text{OPT}_{|P'|} \leq \text{GOPT} \leq 3\text{OPT}_{|P'|}$ and from the condition of Theorem 8.1 we have $\text{OPT}_m \leq 18\text{OPT}_{|P'|}$. In what follows, we assume that this condition is satisfied. Therefore we get the following two bounds on GOPT.

Observation A.3. For $|P'| \leq j \leq m$, we have

$$\text{GOPT} \leq 6\text{OPT}_j.$$

Proof. This follows since adding new points can decrease the cost of the solution by at most a factor 2, since it gives new options to open as centers. Therefore

$$\text{GOPT} \leq 3\text{OPT}_{|P'|} \leq 3 \cdot 2\text{OPT}_j \leq 6\text{OPT}_j. \quad \square$$

Observation A.4. $\text{GOPT} \geq \frac{1}{18}\text{OPT}_m$.

Proof. Follows by combining $\text{GOPT} \geq \text{OPT}_{|P'|}$ and the assumption from Theorem 8.1. □

From Observation A.3 and Lemma A.2 we get

Lemma A.5. For $|P'| \leq j \leq m$ with probability at least $3/4$

$$\sum_{x \in P'_j} \text{dist}(x, f_x) \leq 4(4\text{OPT}_j + \text{GOPT}) \leq 40\text{OPT}_j.$$

Proof. The first inequality follows by applying markov inequality on Lemma A.2, and the second inequality follows from Observation A.3 \square

Similarly, from Observation A.4 and Lemma A.1 by applying markov bound we get

Lemma A.6. With probability at least $3/4$,

$$|S| \leq 4\beta k, \quad \text{where } \beta = (1 + 4 \cdot 18)k(\log \Delta + 1).$$

In order to achieve high probability, we run $q = O(\log(n + \Delta))$ copies of Meyerson's sketch in parallel. Let $S_j^{(0)}, S_j^{(1)}, \dots, S_j^{(q)}$ denote the set maintained by the copies after the j -th insertion of the stream σ' . Also let for any point x the point that it's assigned to in the i -th copy be $f_x^{(i)}$.

Using Lemma A.6 and Lemma A.5

Lemma A.7. With probability $1 - \frac{1}{(n+\Delta)^{10}}$ for any $|P| \leq j \leq m$, there exist a $1 \leq i \leq q$ such that

1. $\sum_{x \in P'_j} \text{dist}(x, f_x^{(i)}) \leq 40\text{OPT}_j.$
2. $|S_j^{(i)}| \leq 292k(\log \Delta + 1).$

Proof. The first property follows from Lemma A.5. The second property follows from Lemma A.6 and the observation that $|S_j^{(i)}| \leq |S_m^{(i)}|$. The probability follows from the number of repetitions (simple product of complementary events). \square

We are ready to explain the construction of the output set P and the stream σ and analyze its performance.

Multi-Meyerson Procedure. We start with an empty stream σ and additional vector of weights $w(x)$ where $w(x)$ is the weight of the point x . Notice that the weight of point x when inserted to the stream is not necessary equal to $w(x)$, since $w(x)$ is dynamic and changes over the time but the weight of an inserted point cannot be changed. Moreover we might insert a point to the stream multiple times with different weights. For each point x that is insert (on the stream σ')⁴, we append it to σ with weight 1 if it is added to $S^{(i)}$ for at least one $1 \leq i \leq q$ for which the current $|S^{(i)}| \leq 292k(\log \Delta + 1)$. In this case we let $w(x) = 1$ as well. Otherwise, we find the closest point $y \in \sigma$ to x and increase the weight of y by one. We assign x to y and let $f_x = y$. Now if $w(y) = 2^\ell$ for some value ℓ , we insert y to the stream with weight 2^ℓ . When the first $|P'|$ insertions are done, we let $P = \sigma$, clear σ and continue the algorithm as before. This step does not logically effect the algorithm. Basically it is dividing the created stream into initial solution and output stream.

⁴Recall that we added the point P' to the beginning of the σ' .

Proof of Theorem 8.1: Let us start by bounding the length of the stream. We have q copies of the Meyerson's sketch and each inserts at most $292k(\log \Delta + 1)$ centers with weight 1 to the stream. The reason is that we skip adding the points added by the sketches with size more than $292k(\log \Delta + 1)$. Therefore in total we have $O(qk(\log \Delta + 1))$ points with weight 1. Moreover, each point that is inserted can be re-inserted at most $\log n + 1$ times, since it is inserted at most once for each power of two and the highest power of two is at most $\log n$. Therefore we get that the length of the stream is upper bounded by

$$k \cdot \text{polylog}(n, \Delta).$$

We now analyze the quality of the solution. Let us focus on a value r , such that, $|P'| \leq r \leq m$ and fix the state of the algorithm and the variables after the r -th insertion of σ' stream. We first observe that, for any point y on the stream, the number of points that are assigned to it is equal to $w(y)$, as we increase $w(y)$ anytime we assign a point to y . We let $v(y)$ be the summation of weights of the point y currently on the stream. Select integer $\ell \geq 1$ such that $2^{\ell-1} \leq w(y) \leq 2^\ell - 1$, by the description of the algorithm we get that

$$v(y) = \sum_{1 \leq i < \ell} 2^i = 2^\ell - 1.$$

Therefore,

$$w(y) \leq v(y) \leq 2w(y). \quad (25)$$

This results in the third property in this theorem.

From Lemma A.7 we get that with probability $1 - \frac{1}{(n+\Delta)^{10}}$, there exist a copy of the Meyerson's sketch that has a good approximation ratio and a small size. Without loss of generality we assume that the copy is $S^{(0)}$ and for any x , the point that it's assigned to is $f_x^{(0)}$. Observe that size $|S^{(0)}| \leq 292k(\log \Delta + 1)$, so all the points in $S^{(0)}$ are part of the stream. Therefore, $\text{dist}(x, f_x) \leq \text{dist}(x, f_x^{(0)})$. The cost of each point x is the cost to the closest center in \mathcal{U} which is by triangular inequality at most

$$\text{dist}(x, \mathcal{U}) \leq \text{dist}(x, f_x) + \text{dist}(f_x, \mathcal{U}).$$

Therefore the total cost is

$$\begin{aligned} \text{cost}(\mathcal{U}, P'_r) &= \sum_{x \in P'_r} \text{dist}(x, \mathcal{U}) \\ &\leq \sum_{x \in P'_r} \text{dist}(x, f_x) + \text{dist}(f_x, \mathcal{U}) \\ &\leq \sum_{x \in P'_r} \text{dist}(x, f_x^{(0)}) + \text{dist}(f_x, \mathcal{U}) \\ &\leq \sum_{x \in P'_r} \text{dist}(x, f_x^{(0)}) + \sum_{x \in P'_r} \text{dist}(f_x, \mathcal{U}) \\ &\leq 40\text{OPT}_r + \sum_{x \in P'_r} \text{dist}(f_x, \mathcal{U}) \end{aligned}$$

where the last inequality follows from Lemma A.7, now we use the weights to simplify the second term. To this end let $P_r = P \cup \sigma$, denote the union of the output set P and current output stream σ . Recall that for

any set Q , $\text{OPT}(Q)$ is the optimum solution for the instance Q .

$$\begin{aligned}
\text{cost}(\mathcal{U}, P'_r) &= \sum_{x \in P'_r} \text{dist}(x, \mathcal{U}) \\
&\leq 40\text{OPT}_r + \sum_{x \in P'_r} \text{dist}(f_x, \mathcal{U}) \\
&\leq 40\text{OPT}_r + \sum_{y \in P_r} \text{dist}(y, \mathcal{U}) \cdot w(y) \\
&\leq 40\text{OPT}_r + \sum_{y \in P_r} \text{dist}(y, \mathcal{U}) \cdot v(y) \\
&\leq 40\text{OPT}_r + \alpha \text{cost}(\text{OPT}(P_r), P_r) \\
&\leq 40\text{OPT}_r + \alpha \text{cost}(\text{OPT}(P'_r), P_r) \\
&\leq 40\text{OPT}_r + \alpha \sum_{y \in P_r} \text{dist}(y, \text{OPT}(P'_r)) \cdot v(y) \\
&\leq 40\text{OPT}_r + \alpha \sum_{y \in P_r} \text{dist}(y, \text{OPT}(P'_r)) \cdot 2w(y) \\
&\leq 40\text{OPT}_r + \alpha \sum_{y \in P_r} \sum_{x | f_x=y} 2(\text{dist}(f_x, x) + \text{dist}(x, \text{OPT}(P'_r))) \\
&\leq 40\text{OPT}_r + \alpha \sum_{x \in P'_r} 2(\text{dist}(f_x, x) + \text{dist}(x, \text{OPT}(P'_r))) \\
&\leq (2\alpha + 1)40\text{OPT}_r + \alpha \sum_{x \in P'_r} 2\text{dist}(x, \text{OPT}(P'_r)) \\
&\leq (2\alpha + 1)40\text{OPT}_r + 2\alpha\text{OPT}_r \\
&\leq 82\alpha + 40\text{OPT}_r,
\end{aligned}$$

where the second inequality follows by rearranging the terms, the third inequality follows from Eq. (25), the fourth inequality follows from Theorem 8.1 since the clustering \mathcal{U} is α -approximate, the fifth inequality follows since the cost of any clustering for points in P_r is more than the cost of an optimum clustering $\text{OPT}(P_r)$ for it, the sixth inequality follows by definition of cost, seventh inequality follows by Eq. (25), eighth inequality follows by triangle inequality and rearranging the terms. This concludes the analysis of the quality of the solution and therefore the entire proof.