

Portland State University

PDXScholar

Computer Science Faculty Publications and
Presentations

Computer Science

6-1-2024

On Coresets for Fair Clustering in Metric and Euclidean Spaces and Their Applications

Sayan Bandyapadhyay
Portland State University

Fedor V. Fomin
University of Bergen, Norway

Kirill Simonov
Technische Universitat Wien, Austria

Follow this and additional works at: https://pdxscholar.library.pdx.edu/compsci_fac



Part of the [Computer Sciences Commons](#)

Let us know how access to this document benefits you.

Citation Details

Bandyapadhyay, S., Fomin, F. V., & Simonov, K. (2024). On coresets for fair clustering in metric and Euclidean spaces and their applications. *Journal of Computer and System Sciences*, 142, 103506.

This Article is brought to you for free and open access. It has been accepted for inclusion in Computer Science Faculty Publications and Presentations by an authorized administrator of PDXScholar. Please contact us if we can make this document more accessible: pdxscholar@pdx.edu.



On coresets for fair clustering in metric and Euclidean spaces and their applications ☆,☆☆

Sayan Bandyapadhyay^a, Fedor V. Fomin^{b,*}, Kirill Simonov^c

^a Portland State University, USA

^b University of Bergen, Norway

^c Technische Universität Wien, Austria

ARTICLE INFO

Article history:

Received 17 May 2021

Received in revised form 6 May 2023

Accepted 3 January 2024

Available online 11 January 2024

Keywords:

Fair clustering

Coresets

Approximation algorithms

ABSTRACT

Fair clustering is a constrained clustering problem where we need to partition a set of colored points. The fraction of points of each color in every cluster should be more or less equal to the fraction of points of this color in the dataset. The problem was recently introduced by Chierichetti et al. (2017) [1]. We propose a new construction of coresets for fair clustering for Euclidean and general metrics based on random sampling. For the Euclidean space \mathbb{R}^d , we provide the first coreset whose size does not depend exponentially on the dimension d . The question of whether such constructions exist was asked by Schmidt et al. (2019) [2] and Huang et al. (2019) [5]. For general metrics, our construction provides the first coreset for fair clustering. New coresets appear to be a handy tool for designing better approximation and streaming algorithms for fair and other constrained clustering variants.

© 2024 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

Given a set of n data points in a metric space and an integer k , *clustering* is the task of partitioning the points into k groups or clusters so that the points in each cluster are similar. In this paper, we consider clustering problems with *fairness* constraints. While there are many competing notions of fairness in the literature, here we consider clustering with fairness constraints or fair clustering as introduced by Chierichetti et al. [1] in their seminal work. The notion became widely popular within a short period triggering a large body of new work [2–10]. The idea of fair clustering is to enforce additional (fairness) constraints to remove the inherent bias or discrimination from vanilla (unconstrained) clustering. For example, suppose we have a sensitive feature (e.g., race or gender). We want to find a clustering where the fraction of points from a traditionally underrepresented group in every cluster is more or less equal to the fraction of points from this group in the dataset. Indeed, the work of Chierichetti et al. [1] shows that clustering computed by classical vanilla algorithms can lead to widely varied ratios for a particular group, especially when the number of clusters is large enough.

There are many settings where machine learning algorithms, trained on datasets of past instances, play a crucial role in decision-making [11–14]. These algorithms are sophisticated and time-efficient and produce accurate results most of the

☆ Funding: The research leading to these results has received funding from the Research Council of Norway via the project BWCA (grant no. 314528) and the European Research Council (ERC) [819416].

☆☆ A preliminary version of the article has been accepted in ICALP 2021.

* Corresponding author.

E-mail addresses: sayanb@pdx.edu (S. Bandyapadhyay), Fedor.Fomin@uib.no (F.V. Fomin), kirillsimonov@gmail.com (K. Simonov).

time. However, there has been a growing concern that these algorithms are biased or discriminatory towards traditionally underrepresented groups [15–17]. One example that stands out and has generated substantial controversy in recent years is concerning the COMPAS risk tool, which is a widely used statistical method for assigning risk scores in the criminal justice system. Angwin et al. argued that this tool was biased against African-American defendants [15,18]. Most of the automated decision-making systems are highly influenced by human players, especially during the training procedure. Importantly, clustering also plays a crucial role in this training part. For example, a widely used technique called feature engineering [19,20] labels samples with their cluster id to enhance the expressive power of learning methods. Hence, the study of biases and discriminatory practices in the context of clustering is well-motivated.

Over the past few years, researchers have put a lot of effort into understanding and resolving the issues of biases in machine learning. This research has led towards different notions of fairness [21–23]. Kleinberg et al. [24] formalized three fairness conditions and showed that it is not possible to satisfy them simultaneously, except in very special cases (see also [25] for a similar treatment). The notion of fairness studied by Chierichetti et al. [1] is based on the concept of *disparate impact* (DI) [26]. Roughly, the DI doctrine articulates that the protected attributes should not be explicitly used in decision-making, and the decisions taken should not be disproportionately different for members in different protected groups.

Following the DI doctrine, Chierichetti et al. [1] considered the model where there is a single sensitive or protected attribute called color that can take only two values: red and blue. The coordinates of the points are unprotected; that is, they do not take part in the fairness constraints. For any integer $t \geq 1$, Chierichetti et al. defined the (t, k) -fair clustering problem where in each cluster the ratio of the number of red points to the number of blue points must be at most t and at least $1/t$. Thus in their case, the notion of fairness is captured by the balance parameter t .

Rösner and Schmidt [27] studied a multicolored version of the above problem, where a clustering is fair if the ratios between points of different colors are the same in every cluster. Subsequently, Bercea et al. [4] and Bera et al. [3] independently formulated a model generalizing the problems studied in [1] and [27]. In this model, we are given ℓ groups P_1, \dots, P_ℓ of points in a metric space and balance parameters $\alpha_i, \beta_i \in [0, 1]$ for each group $1 \leq i \leq \ell$. A clustering is *fair* if the fraction of points from group i in every cluster is at least β_i and at most α_i . Additionally, in [3], the groups are allowed to overlap, i.e., a point can belong to multiple protected classes. Note that this assumption is needed to model many applications, e.g., consider clustering of individuals where a subset of the individuals are African-American women. In fact, the experiments in [3] show that imposing fairness concerning one sensitive attribute (say gender) might lead to unfairness to another (say race) if not protected. We refer to the fair clustering problem with overlapping groups as (α, β) -fair clustering. We note that this is the most general version of fair clustering considered in the literature, and this is the notion of fairness we adapt in this paper. Both [4] and [3] obtain polynomial time $O(1)$ -approximation for this problem that violates the fairness constraints by at most small additive factors. We denote by Γ the maximum number of distinct collections of groups to which a point may belong. If all the groups are disjoint, then $\Gamma = \ell$. Note that if a point can belong to at most Λ groups, then Γ is at most ℓ^Λ . As noted in [3] and [5], while Λ can very well be more than 1, it is usually a constant in most of the applications. Thus, in this case, $\Gamma = \ell^{O(1)}$, which is expected to be much smaller compared to n , the total number of points in the union of the groups.

Several works related to fair clustering have devoted to scalability [5,2,7,6]. Along this line, in a beautiful work, Schmidt et al. [2] defined coresets for fair clustering. Note that a coreset for a center-based vanilla clustering problem is roughly a summary of the data that for every set C of k centers approximately (within $(1 \pm \epsilon)$ factor) preserves the optimal clustering cost. Coresets have mainly two advantages: (1) they take lesser space compared to the original data, and (2) any clustering algorithm can be applied on a coreset to efficiently retrieve a clustering with guarantee almost the same as the one provided by the algorithm. Over the years, researchers have paid increasing attention to the design of coreset construction algorithms to optimize the coreset size. Indeed, finding smaller coresets continues to be an active research area in the context of vanilla k -median and k -means clustering. For general metric spaces, the best-known upper bound on coreset size is $O((k \log n)/\epsilon^2)$ [28] and the lower bound is known to be $\Omega((k \log n)/\epsilon^2)$ [29]. For the real Euclidean space of dimension d , it is possible to construct coresets of size $(k/\epsilon)^{O(1)}$ [30,31]. In particular, the size does not depend on n and d . We note that most of these small size coreset constructions are based on random sampling.

Motivated by the progress on coresets for vanilla clustering, Schmidt et al. [2] initiated the study of fair coresets. In the vanilla version of the clustering problems, given the cluster centers, clusters are formed by assigning each point to its nearest center. In contrast, in a constrained version, such an assignment might not lead to a clustering that satisfies the constraints. Hence, for fair clustering, we need a stronger definition of coreset. Suppose we want to cluster ℓ (possibly overlapping) groups of points. Consider any $k \times \ell$ matrix M with non-negative integer entries where the rows correspond to k clusters and columns to the ℓ groups of points. For every valid clustering \mathcal{I} , one can construct such a matrix: for the i -th cluster and the j -th group, set the number of points from group j in the i -th cluster to be $M[i][j]$. Note that each column of M defines a group's partition induced by the clustering \mathcal{I} . Such a *constraint matrix* M defines a set of cardinality constraints for every pair of a cluster and a group. In this case, we say that the clustering \mathcal{I} satisfies M . Informally, a weighted subset of points is a *fair coreset* if for every set of k centers and every constraint matrix M , the cost of an optimal clustering satisfying M is approximately preserved by the subset. Schmidt et al. [2] and subsequently Huang et al. [5] designed deterministic algorithms in \mathbb{R}^d that construct fair coresets whose sizes exponentially depend on d . To remove this exponential dependency on d , Schmidt et al. [2] proposed an interesting open question whether it is possible to use random sampling for construction of fair coresets. Huang et al. [5] also suggested the same open question. Besides, Huang et al. asked whether it is possible to achieve a similar size bound as in the vanilla setting.

Table 1
Previous and current coresets results in \mathbb{R}^d .

	k -median		k -means	
	size	construction time	size	construction time
[2]			$O(\Gamma k \epsilon^{-d-2} \log n)$	$O(k \epsilon^{-d-2} n \log n)$
[5]	$O(\Gamma k^2 \epsilon^{-d})$	$O(k \epsilon^{-d+1} n)$	$O(\Gamma k^3 \epsilon^{-d-1})$	$O(k \epsilon^{-d+1} n)$
Theorems 5.1 and 6.2	$O(\frac{\Gamma}{\epsilon^3} \cdot k^2 \log n (\log n + d \log(1/\epsilon)))$	$O(nd(k + \ell))$	$O(\frac{\Gamma}{\epsilon^5} \cdot k^7 (\log n)^6 (\log n + d \log(1/\epsilon)))$	$O(nd(k + \ell))$

Table 2

Approximation results for (α, β) -fair clustering in general metrics. “multi” denotes if the algorithm can handle overlapping groups. In “approx.” columns, the first (resp. second) value in a tuple is the approximation factor (resp. violation). [3] does not explicitly compute the $O(1)$ factor, but it is $> 3 + \epsilon$ (resp. $> 9 + \epsilon$) for k -median (resp. k -means), where ϵ is a sufficiently large constant.

	multi	k -median		k -means	
		approx.	time	approx.	time
[4]		(4.675, 1)	poly(n)	(62.856, 1)	poly(n)
[3]	✓	($O(1)$, $4\Lambda + 3$)	poly(n)	($O(1)$, $4\Lambda + 3$)	poly(n)
Theorem 9.1		≈ 3	$(k\ell)^{O(k\ell)} n \log n$	≈ 9	$(k\ell)^{O(k\ell)} n \log n$
Theorem 9.1	✓	≈ 3	$(k\Gamma)^{O(k\Gamma)} n \log n$	≈ 9	$(k\Gamma)^{O(k\Gamma)} n \log n$

1.1. Our results and contributions

We study fair clustering under the k -median and k -means objectives. Our first main result is the following theorem.

Theorem 1.1 (Informal). *There is an $O(n(k + \ell))$ time randomized algorithm that w.p. at least $1 - 1/n$, computes a coreset of size $O(\Gamma(k \log n)^2 / \epsilon^3)$ for (α, β) -fair k -median and $O(\Gamma(k \log n)^7 / \epsilon^5)$ for (α, β) -fair k -means, where Γ is the number of distinct collections of groups to which a point may belong. If the groups are disjoint, the algorithm runs in $O(nk)$ time. Moreover, in \mathbb{R}^d , the coreset sizes are $O(\frac{\Gamma}{\epsilon^3} \cdot k^2 \log n (\log n + d \log(1/\epsilon)))$ for (α, β) -fair k -median and $O(\frac{\Gamma}{\epsilon^5} \cdot k^7 (\log n)^6 (\log n + d \log(1/\epsilon)))$ for (α, β) -fair k -means.*

Theorem 1.1 provides the first coreset construction for fair clustering problem in general metric spaces. Our result is comparable to the best-known bound of $O_\epsilon(k \log n)$ [28] in the vanilla case. In particular, if the number of groups in our case is just 1, we obtain coresets of size $O_\epsilon(\text{poly}(k \log n))$, which matches with the best-known bound in the vanilla case, up to a small degree polynomial factor. We note, that this is the first sampling based coreset construction scheme for fair clustering, and in \mathbb{R}^d , the first coreset construction scheme where the size of the coreset does not depend exponentially on the dimension d . In fact, the dependency on d is only linear. Hence, our result solves the open question proposed in [2] and partly solves the open question proposed in [5]. As we already mentioned, in all of the previous results [2,5], coreset sizes depended exponentially on d (see Table 1). We note that the formal statement of Theorem 1.1 appears in Theorems 4.1, 5.1 and 6.2.

Actually, our coreset construction scheme, similar to [2,5], is much more general in the following sense. The coreset can preserve not only the cost of optimal fair clustering, but also the cost of any optimal clustering with group-cardinality constraints. In particular, for every set of k centers and constraint matrix M , our coreset approximately preserves the cost of optimal clustering that satisfies M . In fact, for any clustering problem with constraints that can be represented by a set of matrices, we obtain a small size coreset. This gives rise to coresets for a wide range of clustering problems including lower-bounded clustering [32,33,3]. Notably, in the case of lower-bounded clustering, the input consists of only one group of points, and thus M is a column matrix.

We further exploit the new coreset construction to design clustering algorithms in various settings. In general metrics, we obtain the first fixed-parameter tractable (FPT) constant-factor approximation for (α, β) -fair clustering with parameters k and Γ . That is, the running time of our algorithm is exponential only in the values k and Γ while polynomial in the size of the input. All previous constant-approximation algorithms were bicriteria and violated the fairness constraints by some additive factors. Hence, the study of FPT approximation is well-motivated. Our approximation factors are reasonably small and improve the best-known approximation factors of the existing bicriteria algorithms (see Table 2). Moreover, our coreset leads to improved constant FPT approximations for many other clustering problems. For example, we obtain an improved ≈ 3 -approximation algorithm for lower-bounded k -median [32,33,3] that is FPT parameterized by k . Previously, the best-known factor for FPT approximation for this problem was 3.736 [3].

Based on our coreset, we also obtain the first FPT $(1 + \epsilon)$ -approximation for (α, β) -fair clustering in \mathbb{R}^d with parameters k and Γ . Furthermore, the running time has a near-linear dependency on n and does not depend exponentially on

Table 3The running time of the $(1 + \epsilon)$ -approximations for fair clustering in \mathbb{R}^d .

	running time	version
[2]	$n^{O(k/\epsilon)}$	2-color, $(1, k)$ -fair clustering
[7]	$n^{\text{poly}(k/\epsilon)}$	ℓ -color, $(1, k)$ -fair clustering
Theorem 8.3	$2^{\tilde{O}(k/\epsilon^{O(1)})} (k\Gamma)^{O(k\Gamma)} nd \log n$	(α, β) -fair clustering

Table 4The running time of the $(1 + \epsilon)$ -approximations for capacitated clustering in \mathbb{R}^d .

	running time
[36]	$2^{\text{poly}(k/\epsilon)} n^2 (\log n)^{k+2d}$
[37]	$2^{\tilde{O}(k/\epsilon^{O(1)})} \cdot n^2 (\log n)^{2d}$
[34]	$(k\epsilon^{-1})^{k\epsilon^{-O(1)}} n^{O(1)}$
[35]	$n^{\epsilon^{-O(1)}} (d=2)$ $n^{(\log n/\epsilon)^{O(d)}} (d \geq 3)$
Theorem 10.4	$2^{\tilde{O}(k/\epsilon^{O(1)})} nd^{O(1)} + nk^2 \epsilon^{-O(1)} \log n$

d . A comparison with the running time of the previous $(1 + \epsilon)$ -approximation algorithms can be found in Table 3. The improvement on the approximation factor is possible in this special case, compared to the general metric case, because here we can efficiently enumerate a list of $(1 + \epsilon)$ -approximate canonical solutions. We also obtain FPT $(1 + \epsilon)$ -approximation algorithms with parameter k for the Euclidean version of several other problems including capacitated clustering [34,35] and lower-bounded clustering. We note that these are the first $(1 + \epsilon)$ -approximations for these problems with near-linear dependency on n . For Euclidean capacitated clustering, quadratic time FPT algorithms follow due to [36,37] (see Table 4). Also, the $(1 + \epsilon)$ -approximation for Euclidean capacitated clustering in [34] and [35] have running time $(k\epsilon^{-1})^{k\epsilon^{-O(1)}} n^{O(1)}$ and at least $n^{\epsilon^{-O(1)}}$ (see Table 4).

Our cores et also leads to small space $(1 + \epsilon)$ -approximation in streaming setting for (α, β) -fair clustering in \mathbb{R}^d when the groups are disjoint. We show how to maintain an $O(d^2 \ell \cdot \text{poly}(k \log n)/\epsilon^4)$ size cores et in each step. One can apply our $(1 + \epsilon)$ -approximation algorithm on the cores et to compute a near-optimal clustering. In the previous streaming algorithms [2], the space complexity depended exponentially on either d or k .

Our cores et construction scheme is based on the algorithm for the vanilla case due to Chen [38]. Chen showed that a small number of points can be chosen randomly from a metric space (in a clever way) that form a cores et for vanilla clustering. Our construction first divides the input points into a number of subsets, and for each subset, applies Chen's algorithm to compute a cores et for that subset. Our final cores et is the union of all these computed cores ets. Surprisingly, our algorithm for fair clustering is as simple as that. Our main technical contribution is the analysis of this algorithm, which shows that the returned weighted subset of the points is a cores et for fair clustering with high probability.

Although our algorithm follows the framework of Chen's, the proof that the algorithm correctly computes a fair cores et is much more complicated. Our analysis is strongly inspired by the analysis of Cohen-Addad and Li [34] of Chen's algorithm in the context of capacitated clustering. However, there are two major difficulties of applying the analysis technique of [34] directly to our problem. Firstly, in our case input points belong to groups which can overlap, and secondly, fairness constraints are much more general compared to capacity constraints. The novelty of our work lies in overcoming the two above-mentioned hurdles. To overcome the first hurdle, we divide the input points into equivalence classes based on their group membership and using a probabilistic argument prove that it is possible to consider these classes separately. To overcome the second hurdle, we prove that fairness constraints can be treated as a collection of independent constraints each of which is similar to a capacity constraint. All these ideas together help us generalize the approach of [34] to fair clustering.

Apart from the cores et construction, the novelty of our work lies in the design of an algorithm for computing the minimum cost fair assignment to given centers, based on mixed-integer linear programming. This is the heart of all our approximation algorithms.

1.2. Comparison with related work

Here we compare our results with closely related previous work. Schmidt et al. [2] defined the concept of fair cores ets and gave cores et of size $O(\ell k \epsilon^{-d-2} \log n)$ for the disjoint group case of Euclidean (α, β) -fair k -means. This can be extended to the overlapping case by replacing ℓ with Γ in the size bound. Using a sophisticated dimension reduction technique [39], they showed how to stream cores et whose size does not depend exponentially on d . Unfortunately, this cores et size depends exponentially on k . Schmidt et al. also gave an $n^{O(k/\epsilon)}$ time $(1 + \epsilon)$ -approximation for the two-color version of the

problem. Note that our work improves over all these results (see Tables 1 and 3). Using the framework in [40], Huang et al. [5] improved the coreset size bound of [2] by a factor of $\Theta\left(\frac{\log n}{k\epsilon^2}\right)$ and gave the first coreset for Euclidean (α, β) -fair k -median of size $O(\Gamma k^2 \epsilon^{-d})$. Both the coreset construction schemes in [2] and [5] use deterministic algorithms, and thus they proposed whether random sampling can be employed to remove the curse of dimensionality. Note that our result based on random sampling improves the bound (for k -median) in [5] by a factor of $\Theta\left(\frac{\epsilon^{-d+3}}{\log n(\log n+d)}\right)$ (see Table 1). Böhm et al. [7] considered $(1, k)$ -fair clustering with multiple colors. They designed near-linear time constant-approximation algorithms in this restricted setting. They also obtained an $n^{\text{poly}(k/\epsilon)}$ time $(1 + \epsilon)$ -approximation for the Euclidean version in the same setting. An FPT $(1 + \epsilon)$ -approximation follows from our work for this version (see Table 3).

Chierichetti et al. [1] gave a polynomial time $\Theta(t)$ -approximation for (t, k) -fair k -median with two groups (or colors). We improve their result by giving an FPT constant-approximation algorithm with parameters k and ℓ for (t, k) -fair clustering with arbitrary number of colors. Based on the framework implicitly mentioned in [41], Bera et al. [3] obtained polynomial time $O(1)$ -approximation for (α, β) -fair clustering that violates the fairness constraints by at most an additive factor of $4\Lambda + 3$. This framework first computes k centers using a ρ -approximation algorithm for vanilla clustering, and then finds an assignment of the points to these centers that satisfies the fairness constraints. They showed, e.g., for k -median, there is always such an assignment whose cost is at most $\rho + 2$ times the optimal cost of fair clustering. However, computing such an assignment is not an easy task. Indeed, this is a big hurdle one faces while studying fair clustering, which makes this problem substantially harder compared to other clustering problems like capacitated clustering. Based on the algorithm due to Király et al. [42], Bera et al. [3] showed that an optimal assignment can be computed by violating any fairness constraint by the mentioned factor. For the disjoint group case, their violation factor is only 3. Independently, Bercea et al. [4] obtained algorithms with the same approximation guarantees as in [3] for the disjoint version, but with at most 1 additive factor violation. We show that the above mentioned assignment problem for (α, β) -fair clustering can be solved exactly in FPT time parameterized by k and Γ . Plugging this in with our coreset, we obtain algorithms with better constant approximation factors compared to [3] and [4] that do not violate any constraint (see Table 2).

Ding and Xu [36] gave a unified framework with running time $2^{\text{poly}(k/\epsilon)} (\log n)^{k+1} nd$ that generates a collection of candidate sets of centers for clustering problems with constraints in \mathbb{R}^d . Subsequently, Bhattacharya et al. [37] and Feng et al. [43] designed similar frameworks having improved time complexity. None of these works study fair clustering. Our work can be viewed as an extension of these works to general metrics in the sense that we obtain constant-approximations for a range of constrained clustering problems. Furthermore, by applying the framework of [37] on our coreset, we obtain $(1 + \epsilon)$ -approximation algorithms with improved time complexity bounds for several clustering problems in \mathbb{R}^d .

1.3. Other related work

Fair clustering has received a huge amount of attention from both theory and practice. Most of the works considered the notion of fairness popularized by Chierichetti et al. [1]. Backurs et al. [6] studied Euclidean fair clustering with the goal of designing scalable algorithms. They followed a fairness notion very similar to the one in [1] and considered the two color case. Their main result is a near-linear time $O(d \log n)$ -approximation.

Fair version of k -center is also a well-studied problem [1,27,4,3,9]. In contrast to k -median and k -means, polynomial time true constant approximation is known for the multiple color generalization of (t, k) -fair clustering [27,4].

Fair clustering has been studied with different notions of fairness as well. Chen et al. [8] defined fairness as proportionality where any n/k points can form their own cluster if there is another center that is closer to all of these n/k points. Kleindessner et al. [10] considered the fair k -center problem where each center has a type and for each type, a fixed number of centers must be chosen. They gave a simple linear-time constant factor approximation for this problem. In a different work [44], they extended the fairness notion to spectral clustering.

Clustering problems have been studied in the literature with other constraints. One such popular problem is capacitated clustering. For capacitated k -center, polynomial time $O(1)$ -approximations are known both for the uniform [45,46] and non-uniform [47,48] versions. In contrast, for the capacitated version of k -median and k -means, no polynomial time $O(1)$ -approximation is known. However, bicriteria constant-approximations are known that violate either the capacity constraints or the constraint on the number of clusters, by an $O(1)$ factor [49–55]. Recently, Cohen-Addad and Li [34] designed FPT ≈ 3 - and ≈ 9 -approximation with parameter k for the capacitated version of k -median and k -means, respectively. Polynomial time constant-approximations for lower-bounded k -median follow from [32,33]. Also, an FPT $O(1)$ -approximation with parameter k is known for this problem [3]. Many other clustering constraints have been studied in the literature, e.g., matroid [56], fault tolerance [57], chromatic clustering [36] and diversity [58]. We will discuss more about the last two problems in Section 10.

Coresets have been used in the context of k -median and k -means clustering for obtaining near-optimal solutions, especially for points in the Euclidean spaces. Many different schemes have been proposed over the years for coreset construction. In the earlier works, standard techniques have been used that led to coresets whose size depend exponentially on the dimension d [59–61]. Chen [38] improved the dependence on d to be polynomial. Subsequently, this dependence has been further improved [62,28]. Finally, the dependence on d were removed for both of the problems [30,31]. See also [63–65,29,66] for recent improvements. For capacitated clustering, Cohen-Addad and Li [34] gave an $O_\epsilon(\text{poly}(k \log n))$ size

coreset in general metrics. In a very recent work, Braverman et al. [67] have improved our coreset bound in the Euclidean case to $O(\text{poly}(k/\epsilon))$.

Following our work, the coreset-size bound for k -median in general metrics has been improved to $\tilde{O}(\Gamma \cdot \frac{k^3}{\epsilon^6})$ [68] and subsequently to $\tilde{O}(\Gamma \cdot \frac{k^2}{\epsilon^4})$ [69]. Similar improvements have been shown for k -means as well.

Organization. In Section 2, we introduce the definitions and notation that we will use throughout the paper. The “stronger coreset” construction algorithm for k -median in the disjoint group case appears in Section 3 and is extended to the overlapping group case in Section 4. Section 5 describes the coreset construction for k -median in \mathbb{R}^d . The coreset constructions for k -means appear in Section 6. In the rest of the paper, we describe the applications of our coresets. In Section 7, we describe an algorithm for solving an assignment problem, which we will need to design our algorithms for (α, β) -fair clustering. In Section 8 and 9, we describe our approximation algorithms for the Euclidean and metric case of (α, β) -fair clustering, respectively. In Section 10, we apply our coreset to design improved algorithms for other constrained clustering problems. In Section 11, we show how to maintain our coreset in the streaming setting. Finally, in Section 12, we conclude with some open questions.

2. Preliminaries

In all the clustering problems we study in this paper, we are given a set P of points in a metric space $(\mathcal{X}, d(\cdot, \cdot))$, that we have to cluster. We are also given a set F of cluster centers in the same metric space. We note that P and F are not-necessarily disjoint, and in fact, P may be equal to F . We assume that the distance function $d(\cdot, \cdot)$ is provided by an oracle that for any given $x, y \in \mathcal{X}$ in constant time returns $d(x, y)$. In the Euclidean version of a clustering problem, $P \subseteq \mathbb{R}^d$, $F \subseteq \mathbb{R}^d$ and $d(\cdot, \cdot)$ is the Euclidean metric. Due to the lack of better notations, we denote the dimension by d and distance function by $d(\cdot, \cdot)$. In the metric version, we assume that F is finite. Thus, strictly speaking, the Euclidean version is not a special case of the metric version. In the metric version, we denote $|P \cup F|$ by n and in the Euclidean version, $|P|$ by n . For any set S and a point p , $d(p, S) := \min_{q \in S} d(p, q)$. Also, for any integer $t \geq 1$, we denote the set $\{1, 2, \dots, t\}$ by $[t]$.

In the k -median problem, given an additional parameter k , the goal is to select a set of at most k centers $C \subseteq F$ such that the quantity $\sum_{p \in P} d(p, C)$ is minimized. k -means is identical to k -median, except here we would like to minimize $\sum_{p \in P} (d(p, C))^2$.

Next, we define our notion of fair clustering following the definition in [3].

Definition 2.1 (Definition 1, [3]). In the fair version of a clustering problem (k -median or k -means), one is additionally given ℓ many (not necessarily disjoint) groups of P , namely P_1, P_2, \dots, P_ℓ . One is also given two fairness vectors $\alpha, \beta \in [0, 1]^\ell$, $\alpha = (\alpha_1, \dots, \alpha_\ell)$, $\beta = (\beta_1, \dots, \beta_\ell)$. The objective is to select a set of at most k centers $C \subseteq F$ and an assignment $\varphi : P \rightarrow C$ such that φ satisfies the following fairness constraints:

$$|\{x \in P_i : \varphi(x) = c\}| \leq \alpha_i \cdot |\{x \in P : \varphi(x) = c\}|, \quad \forall c \in C, \forall i \in [\ell],$$

$$|\{x \in P_i : \varphi(x) = c\}| \geq \beta_i \cdot |\{x \in P : \varphi(x) = c\}|, \quad \forall c \in C, \forall i \in [\ell],$$

and $\text{cost}(\varphi)$ is minimized among all such assignments.

In the (α, β) -Fair k -median problem, $\text{cost}(\varphi) := \sum_{x \in P} d(x, \varphi(x))$, and in the (α, β) -Fair k -means problem, $\text{cost}(\varphi) := \sum_{x \in P} d(x, \varphi(x))^2$. To refer to these two problems together, we will use the term (α, β) -FAIR CLUSTERING. We call φ that satisfies the fairness constraints a fair assignment. We denote the minimum cost of a fair assignment of a set of points P to a set of k centers C by $\text{faircost}(P, C)$, and $\text{faircost}(P)$ denotes the minimum of $\text{faircost}(P, C')$ over all possible sets of k centers C' .

Next, we state our notion of coresets. We follow the definitions in [2,5]. For a clustering problem with k centers and ℓ groups P_1, \dots, P_ℓ , a coloring constraint is a $k \times \ell$ matrix M having non-negative integer entries. The entry of M corresponding to row i and column j is denoted by M_{ij} . Next, we have the following observation, which was also noted in [2,5].

Proposition 2.2. Given a set C of k centers, the assignment restriction required for (α, β) -FAIR CLUSTERING can be expressed as a collection of coloring constraints.

In our definition, a coreset is required to preserve the optimal clustering cost w.r.t. all coloring constraints, and hence it also preserves the optimal fair clustering cost. Next, we formally define the cost of a clustering w.r.t. a set of centers and a coloring constraint.

First, consider the k -median objective. Suppose we are given a weight function $w : P \rightarrow \mathbb{R}_{\geq 0}$ (non-negative reals). Let $W \subseteq P \times \mathbb{R}$ be the set of pairs $\{(p, w(p)) \mid p \in P \text{ and } w(p) > 0\}$. For a set of centers $C = \{c_1, \dots, c_k\}$ and a coloring constraint M , $\text{wcost}(W, M, C)$ is the minimum value $\sum_{p \in P, c_i \in C} \psi(p, c_i) \cdot d(p, c_i)$ over all assignments $\psi : P \times C \rightarrow \mathbb{R}_{\geq 0}$ such that

1. For each $p \in P$, $\sum_{c_i \in C} \psi(p, c_i) = w(p)$.
2. For each $c_i \in C$ and group $1 \leq j \leq \ell$, $\sum_{p \in P_j} \psi(p, c_i) = M_{ij}$.

For k -means, $wcost(W, M, C)$ is defined in the same way except it is the minimum value $\sum_{p \in P, c_i \in C} \psi(p, c_i) \cdot d(p, c_i)^2$. If there is no such assignment ψ , $wcost(W, M, C) = \infty$. When $w(p) = 1$ for all $p \in P$, we simply denote W by P and $wcost(W, M, C)$ by $cost(P, M, C)$. Now we define a coreset. We call it universal coreset, as it is required to preserve optimal clustering cost w.r.t. all coloring constraints.

Definition 2.3 (Universal coreset). For a given unweighted point set P and a clustering objective, a universal coreset is a set of weighted points $W \subseteq P \times \mathbb{R}$ such that for every set of centers C of size k and any coloring constraint M ,

$$(1 - \epsilon) \cdot cost(P, M, C) \leq wcost(W, M, C) \leq (1 + \epsilon) \cdot cost(P, M, C).$$

3. Coreset construction for k -median in the disjoint group case

In this section, we prove the following theorem.

Theorem 3.1. Given a set P of n points in a metric space along with a color function $c : P \rightarrow \{1, \dots, \ell\}$, there is an $O(nk)$ time randomized algorithm that w.p. at least $1 - 1/n$, computes a universal coreset for k -median clustering of size $O(\ell(k \log n)^2/\epsilon^3)$.

To prove this theorem, we design and analyze a coreset construction algorithm. Our algorithm is built upon the coreset construction algorithm for vanilla clustering due to Chen [38]. In our case, we have points from ℓ disjoint color classes. So, we apply Chen's algorithm for each color class independently. Note that Chen's algorithm was used to show that for any given set of centers C , the constructed coreset approximately preserves the optimal clustering cost. However, we would like to show that for any given set of centers C , the constructed coreset approximately preserves the optimal clustering cost corresponding to any given constraint M . At this stage, it is not clear why Chen's algorithm should work in such a generic setting. Our main technical contribution is to show that sampling based approaches like Chen's algorithm can be used even for such a stronger notion of universal coreset. We will try to give some intuition after describing our algorithm. Our algorithm is as follows.

Given the set of points P , first we apply the algorithm of Indyk [70] for computing a vanilla k -median clustering of P . This is a bicriteria approximation algorithm that uses $O(k)$ centers and runs in $O(nk)$ time. Let C^* be the set of computed centers, ν be the constant approximation factor and Π be the cost of the clustering. Also, let $\mu = \Pi/(\nu n)$ be a lower bound on the average cost of the points in any optimal k -median clustering. Note that for any point p , $d(p, C^*) \leq \Pi = \nu n \cdot \mu$.

For each center $c_i^* \in C^*$, let $P_i^* \subseteq P$ be the corresponding cluster of points assigned to c_i^* . We consider the ball $B_{i,j}$ centered at c_i^* and having radius $2^j \mu$ for $0 \leq j \leq N$, where $N = \lceil \log(\nu n) \rceil$. We note that any point at a distance $2^N \mu \geq \nu n \cdot \mu$ from c_i^* is in $B_{i,N}$, and thus all the points in P_i^* are also in $B_{i,N}$. Let $B'_{i,0} = B_{i,0}$ and $B'_{i,j} = B_{i,j} \setminus B_{i,j-1}$ for $1 \leq j \leq N$. We refer to each such $B'_{i,j}$ as a ring for $1 \leq i \leq k, 0 \leq j \leq N$. For each $0 \leq j \leq N$ and color $1 \leq t \leq \ell$, let $P'_{i,j,t}$ be the set of points in $B'_{i,j}$ of color t . Let $s = \Theta(k \log n / \epsilon^3)$ for a sufficiently large constant hidden in $\Theta(\cdot)$.

For each center $c_i^* \in C^*$, we perform the following steps.

Random sampling. For each color $1 \leq t \leq \ell$ and ring index $0 \leq j \leq N$, do the following. If $|P'_{i,j,t}| \leq s$, add all the points of $P'_{i,j,t}$ to $W_{i,j}$ and set the weight of each such point to 1. Otherwise, select s points from $P'_{i,j,t}$ independently and randomly (without replacement) and add them to $W_{i,j}$. Set the weight of each such point to $|P'_{i,j,t}|/s$.

The set $W = \cup_{i,j} W_{i,j}$ is the desired universal coreset. As the number of rings is $O(k \log n)$, the size of W is $O(\ell(k \log n)^2/\epsilon^3)$. From [38], it follows that for each color, the coreset points can be computed in time linear in the number of points of that color times $O(k)$. Thus, our coreset construction algorithm runs in $O(nk)$ time. Next, we show that W is indeed a universal coreset w.h.p.

Note that we need to show that for any set of centers C , the optimal clustering cost is approximately preserved w.r.t. all possible combinations of cluster sizes as defined by the constraint matrices. In Chen's analysis, it was sufficient to argue that for any set of centers C , the optimal clustering cost needs to be preserved. This seems much easier compared to our case. (Obviously, the details are much more complicated even in the vanilla case.) For example, in the vanilla case, let $p \in P$ be a point that is assigned to a center $c \in C$ in an optimal clustering. Note that c must be a closest center to p . For simplicity, suppose p has a unique closest center. Now, if p is chosen in the coreset, then the total weight of p must also be assigned to c in any optimal assignment w.r.t. C . Thus, the assignment function for original and coreset points remains same in the vanilla case. This fact is in the heart of their analysis. Note that this is not necessarily true in our case. We cannot just use the nearest neighbor assignment scheme, as in our case cluster sizes are predefined through M . Indeed, in our case we might very well need to assign the weight of a coreset point to multiple centers to satisfy M . In general, this is the main hurdle one faces while analyzing a sampling based approach for fair coreset construction.

For analyzing our algorithm, we follow an approach similar to the one by Cohen-Addad and Li in [34]. They considered the capacitated clustering problem, where for each center c a capacity value U_c is given, and if the center c is chosen, at most U_c points can be assigned to c . They analyzed Chen's algorithm and showed that for any center C , the coresot approximately preserves the optimal capacitated clustering cost. One crucial idea they use in their proof is representation of assignments through network flow. Suppose we are given a fixed set of centers and weighted input points, and we would like to compute a minimum cost assignment of the points to the centers such that the capacities are not violated. This problem can be modeled as a minimum cost network flow problem.

The first hurdle to adapt the approach in [34] is that it is not possible to represent the assignment problem for fair clustering as a simple flow computation problem. Thus it is not clear how to directly use their approach for fair clustering. Nonetheless, we show that for a fixed constraint M , the assignment problem can be modeled in the desired way. Thus, we can get high probability bound w.r.t. a fixed constraint M . However, to obtain a coresot for fair clustering we need to show this w.r.t. all such constraints (and this leads us towards a universal coresot). The number of such constraints can be as large as $n^{\Omega(k\ell)}$. Hence, to obtain the h.p. bound over all M , we need to show that for a fixed M the error probability is at most $1/n^{\Omega(k\ell)}$. However, it is not clear how to show such a strong bound, as we pick only $O(k \log n / \epsilon^3)$ size sample from each ring corresponding to each color. Nevertheless, we show that it is not necessary to consider all those choices of the constraints together – one can focus on a single color and the constraints w.r.t. that color only. Indeed, this is the reason that we apply Chen's algorithm to different color classes independently. Unfortunately, we pay a heavy toll for this: the coresot size is proportional to ℓ , unlike the vanilla coresot size, and it is not clear how to avoid this dependency. Anyway, this solves our problem, as now we have only $n^{\Omega(k)}$ constraints. It follows that, for a fixed color and a fixed constraint matrix, one can apply an approach similar to the one in [34] (the details are slightly different). This allows us to adapt the ideas from [34] and [38] to construct coresets for much more general clustering problems.

3.1. The analysis

One way to prove that W is a universal coresot is to show that w.h.p. for any fixed set of centers C of size k and any coloring constraint M ,

$$(1 - \epsilon) \cdot \text{cost}(P, M, C) \leq \text{wcost}(W, M, C) \leq (1 + \epsilon) \cdot \text{cost}(P, M, C).$$

Then, by taking union bound over all C and M , we obtain the desired bound. However, as we potentially have $n^{\Omega(k)}$ choices for C and $n^{\Omega(k\ell)}$ choices for M , we need this bound for fixed C and M w.p. $1 - 1/n^{\Omega(k\ell)}$. It is not clear how to prove such a bound, as we pick only $O(k \log n / \epsilon^3)$ size sample from each ring corresponding to each color. Instead, we prove that for any fixed C , and for all M , w.p. $1 - 1/n^{\Omega(k)}$, the above bound holds. In particular, we will show that for each ring $B'_{i,j}$ with points $P_{i,j}$ the error is bounded by $\epsilon |P_{i,j}| \cdot 2^j \mu$.

Lemma 3.2. *For any fixed set C of k centers and for all $k \times \ell$ matrices M , w.p. at least $1 - 1/n^{k+2}$, $|\text{cost}(P, M, C) - \text{wcost}(W, M, C)| \leq \sum_{(i,j)} \epsilon |P_{i,j}| \cdot 2^j \mu$.*

Now, consider all the rings $B'_{i,j}$ with $j = 0$. Then,

$$\sum_{(i,j): j=0} \epsilon |P_{i,j}| \cdot 2^j \mu \leq \epsilon n \cdot \mu \leq \epsilon \cdot \text{OPT}_v \leq \epsilon \cdot \text{cost}(P, M, C).$$

Here, OPT_v is the optimal cost of vanilla k -median clustering. The last inequality follows, as the optimal cost of vanilla clustering is at most the cost of any constrained clustering. Now, for any ring $B'_{i,j}$ with $j \geq 1$ and any point p in the ring, $d(p, c_i^*) \geq 2^{j-1} \mu$. Thus,

$$\sum_{(i,j): j \geq 1} \epsilon |P_{i,j}| \cdot 2^j \mu \leq \epsilon \sum_{p \in P} 2 \cdot d(p, c_{i_p}^*) \leq 2\epsilon \cdot \text{OPT}_v \leq 2\epsilon \cdot \text{cost}(P, M, C),$$

where for a point $p \in P$ by i_p we denote the index of a center such that p belongs to $B'_{i_p,j}$ for some j .

By taking union bound over all C and scaling ϵ down by a factor of 3, we obtain the desired result.

Lemma 3.3. *For every set C of k centers and every $k \times \ell$ matrices M , w.p. at least $1 - 1/n$, $|\text{cost}(P, M, C) - \text{wcost}(W, M, C)| \leq \epsilon \cdot \text{cost}(P, M, C)$.*

This completes the proof of Theorem 3.1. Now, we are left with the proof of Lemma 3.2.

3.2. Proof of Lemma 3.2

Let P_τ be the points in P of color τ . Also, let W_τ be the chosen samples of color τ . For $1 \leq t \leq \ell - 1$, let $W^t = (\sum_{\tau=1}^t W_\tau) \cup (\cup_{\tau=t+1}^\ell P_\tau)$. Also, let $W^\ell = \sum_{\tau=1}^\ell W_\tau$ be the coresot points of all colors. Recall that for any ring $B'_{i,j}$, $P'_{i,j,\tau}$ is the points of color τ in the ring. Also, $P_{i,j} = \cup_{\tau=1}^\ell P'_{i,j,\tau}$.

Note that in the above, W^t contains the sampled points for color 1 to t and original points of color $t+1$ to ℓ . We will prove the following lemma that gives a bound when the coresot contains sampled points of a fixed color t and original points of the other colors.

Lemma 3.4. Consider any color $1 \leq t \leq \ell$. For any fixed set C of k centers and for all $k \times \ell$ matrices M , w.p. at least $1 - 1/n^{k+4}$, $|\text{cost}(P, M, C) - \text{wcost}(W_t \cup (P \setminus P_t), M, C)| \leq \sum_{(i,j)} \epsilon |P'_{i,j,t}| \cdot 2^j \mu$.

Note that for a particular color class if we select all original points in the coresot, then there is no error corresponding to those coresot points. This is true, as one can use the corresponding optimal assignment for these points. Assuming that the above lemma holds, now, we prove Lemma 3.2. Consider the coresot W^1 . From the above lemma we readily obtain the following.

Corollary 3.5. For any fixed set C of k centers and for all $k \times \ell$ matrices M , w.p. at least $1 - 1/n^{k+4}$, $|\text{cost}(P, M, C) - \text{wcost}(W^1, M, C)| \leq \sum_{(i,j)} \epsilon |P'_{i,j,1}| \cdot 2^j \mu$.

Now, in W^1 consider replacing the points of P_2 by the samples in W_2 . We obtain the coresot W^2 . Note that the samples in W_1 and W_2 are chosen independent of each other. Thus, by taking union bound over color 1 and 2, from Lemma 3.4 we obtain, for all $k \times \ell$ matrices M , w.p. at least $1 - 2/n^{k+4}$, $|\text{cost}(P, M, C) - \text{wcost}(W^2, M, C)| \leq \sum_{(i,j)} \epsilon (|P'_{i,j,1}| + |P'_{i,j,2}|) \cdot 2^j \mu$. Similarly, by taking union bound over all $\ell \leq n$ colors and noting that $W^\ell = W$, Lemma 3.2 follows.

Next, we prove Lemma 3.4.

3.3. Proof of Lemma 3.4

Recall that P_t is the set of points of color t , and W_t is the coresot points of color t . C is the given set of centers. For any matrix M , let M^t be the t^{th} column of M . We have the following observation that implies that it is sufficient to consider the points only in P_t to give the error bound.

Observation 3.6. Suppose w.p. at least $1 - 1/n^{k+4}$, for all column matrix M' , $|\text{cost}(P_t, M', C) - \text{wcost}(W_t, M', C)| \leq \sum_{(i,j)} \epsilon |P'_{i,j,t}| \cdot 2^j \mu$. Then, with the same probability, for all $k \times \ell$ matrix M , $|\text{cost}(P, M, C) - \text{wcost}(W_t \cup (P \setminus P_t), M, C)| \leq \sum_{(i,j)} \epsilon |P'_{i,j,t}| \cdot 2^j \mu$.

Proof. Consider any $k \times \ell$ matrix M . Then,

$$\text{cost}(P, M, C) = \sum_{\tau=1}^{\ell} \text{cost}(P_\tau, M^\tau, C)$$

Also,

$$\text{wcost}(W_t \cup (P \setminus P_t), M, C) = \text{wcost}(W_t, M^t, C) + \sum_{\tau \in [\ell] \setminus \{t\}} \text{cost}(P_\tau, M^\tau, C)$$

It follows that,

$$|\text{cost}(P, M, C) - \text{wcost}(W_t \cup (P \setminus P_t), M, C)| = |\text{cost}(P_t, M^t, C) - \text{wcost}(W_t, M^t, C)|$$

Now, by our assumption, it follows that the probability of the event: for all M , $|\text{cost}(P_t, M^t, C) - \text{wcost}(W_t, M^t, C)|$ exceeds $\sum_{(i,j)} \epsilon |P'_{i,j,t}| \cdot 2^j \mu$ is at most $1/n^{k+4}$. Hence, the observation follows. \square

By the above observation, it is sufficient to prove that w.p. at least $1 - 1/n^{k+4}$, for all column matrix M , $|\text{cost}(P_t, M, C) - \text{wcost}(W_t, M, C)| \leq \sum_{(i,j)} \epsilon |P'_{i,j,t}| \cdot 2^j \mu$. The proof of this claim is similar to the analysis in [34]. In the rest of this section we prove this claim. For simplicity, we first do the analysis for a single ring. Later we will show how this idea in single ring case can be extended to obtain the h.p. bound for the multiple ring case.

3.3.1. Single ring case

We fix a ring $B'_{i,j}$ and rename the color t to γ . Note that we have points of only one color γ . For simplicity of notation, we rename P_γ to P . We do the analysis assuming that we sample points only from the ring $B'_{i,j}$. For simplicity, we denote this ring by B' . Let $P' = P'_{i,j,\gamma}$, $m = |P'|$, $\mu' = 2^j \mu$ and $c' = c_i^*$ for $1 \leq i \leq k$ and $0 \leq j \leq N$. Also, let S be the random sample chosen from P' . Thus in this case, our coreset W' consists of the points S , which have weight m/s and all the points in $P \setminus P'$, which have weight 1, i.e., $W' = S \cup (P \setminus P')$. We will show that the cost difference between P and W' is at most $\epsilon m \mu'$ w.h.p. Intuitively, for each point in P' , we allow at most $\epsilon \mu'$ error on average.

For the rest of the proof we fix a column matrix M such that $\text{cost}(P, M, C) < \infty$. We will prove the following theorem.

Theorem 3.7. *W.p. at least $1 - 1/n^{2k+10}$, it holds that $|\text{cost}(P, M, C) - \text{wcost}(W', M, C)| \leq \epsilon m \mu'$.*

By taking union bound over all (at most n^k) column matrices, we obtain the desired bound w.h.p. Towards this end, assume that $s < m$, otherwise $W' = P$ and the above theorem is trivially true.

An alternative way of sampling. Consider the points of P' and the following alternative way of sampling points from P' . For each $p \in P'$, select p w.p. s/m independently and set its weight to m/s . Otherwise, set its weight to 0. Let $X \in \mathbb{R}_{\geq 0}^m$ be the corresponding random vector such that $X[p] = m/s$ if p is selected, otherwise $X[p] = 0$.

We note two things here. First, for each p , $\mathbb{E}[X[p]] = 1$. Thus, $\mathbb{E}[X] = \mathbf{1}$, where $\mathbf{1}$ is the vector of length m whose entries are all 1. Intuitively, this shows that in expectation the chosen set of samples behave like the original points. We will heavily use this connection in our analysis. Second, this sampling is different from our sampling scheme in the sense that here we might end up selecting more (or less) than s samples. However, one can show that with sufficient probability, this sampling scheme selects exactly s points, as the expected number is $m \cdot (s/m) = s$.

Claim 3.8. [34] *Let n be a positive integer, and $p \in (0, 1)$ such that np is an integer. The probability that $\text{Bernoulli}(n, p) = np$ is at least \sqrt{p} .*

Using the above claim with $n = m$ and $p = s/m$, it follows that X contains exactly s non-zero entries w.p. $\Omega(1/\sqrt{n})$. Conditioned on this event, X accurately represents the outcome of our sampling process. Thus, both the sampling processes are same up to repetition. Henceforth, we assume that X contains exactly s non-zero entries.

Representing assignment by network flow. Given a vector $Y \in \mathbb{R}_{\geq 0}^m$ indexed by the points of P' we construct the following flow network G_Y . G_Y has two designated vertices s and t , which are called the source and the sink, respectively. For each point $p_j \in P$, there is a vertex u_j . For each center $c_i \in C$, there is a vertex v_i . There is also an auxiliary vertex w in G_Y corresponding to the center c' of the bicriteria solution. For each u_j , there is an edge between s and u_j , and also between w and u_j . s is also connected to w via an edge. w is connected to each v_i via an edge. Also, each v_i is connected to t via an edge. For each point p_j and center c_i , there is an edge between u_j and v_i . Formally, the vertex set V_Y of G_Y is defined as, $V_Y = \{s\} \cup \{t\} \cup \{w\} \cup \{u_j \mid 1 \leq j \leq n\} \cup \{v_i \mid 1 \leq i \leq k\}$. The set of edges $E_Y = \{(s, u_j) \mid 1 \leq j \leq n\} \cup \{(u_j, w) \mid 1 \leq j \leq n\} \cup \{(w, v_i) \mid 1 \leq i \leq k\} \cup \{(u_j, v_i) \mid 1 \leq j \leq n, 1 \leq i \leq k\}$. For each $p_j \in P \setminus P'$, (s, u_j) has a demand of 1. For each $p_j \in P'$, (s, u_j) has a demand of $Y[p_j]$. The demand of (s, w) is exactly $m - \sum_{p \in P'} Y[p]$, which can be negative. The capacity of each edge (v_i, t) is exactly $M[i]$, the i^{th} entry of M . Lastly, the cost of all the edges is 0 except the edges of $\{(u_j, v_i)\}$, $\{(u_j, w)\}$ and $\{(w, v_i)\}$. The cost of (u_j, v_i) is $d(p_j, c_i)$ and the cost of (u_j, w) is $d(p_j, c')$. The cost of (w, v_i) is $d(c', c_i)$.

We note that the assignment of points in P to the centers in C corresponding to an optimal clustering (with $\text{cost}(P, M, C) < \infty$) induces a flow for G_Y with $Y = \mathbf{1}$ that satisfies all the demands, which sum to $|P|$. Hence, for any $Y \in \mathbb{R}_{\geq 0}^m$, G_Y always has a feasible flow, as the sum of demands is exactly $|P \setminus P'| + \sum_{p \in P'} Y[p] + (m - \sum_{p \in P'} Y[p]) = |P|$.

For any $Y \in \mathbb{R}_{\geq 0}^m$, we denote by $f(Y)$ the cost of the minimum cost feasible flow in G_Y . Consider the random vector X defined before. We have the following important observation.

Observation 3.9. $f(X)$ and $\text{wcost}(W', M, C)$ are identically distributed. Moreover, $f(\mathbb{E}[X]) = \text{cost}(P, M, C)$.

Proof. Note that the total demand in G_X is $|P|$, as argued before. This demand must be routed to t through the edges $\{(v_i, t)\}$. Now, the capacity of (v_i, t) is $M[i]$. If M is a valid partition matrix, then $\sum_{i=1}^k M[i]$ must be $|P|$. Thus, any feasible flow in G_X , which satisfies all the demands, must saturate all the edges $\{(v_i, t)\}$. It follows that from this flow we can retrieve an assignment of the points in W' to the centers in C , such that exactly $M[i]$ weight is assigned to each center $c_i \in C$. Finally, as X contains exactly s non-zero entries, the cost of the minimum cost feasible flow in G_X and $\text{wcost}(W', M, C)$ must be identically distributed.

The moreover part follows by noting that $\mathbb{E}[X] = \mathbf{1}$. \square

From the above observation it follows that to prove Theorem 3.7, it is sufficient to prove that w.p. $1 - 1/n^{\Omega(k)}$, $|f(X) - f(\mathbb{E}[X])| \leq \epsilon m \mu'$. Now, we have another observation which will be useful later.

Observation 3.10. The function f is μ' -Lipschitz w.r.t. the ℓ_1 distance in $\mathbb{R}_{\geq 0}^m$.

Proof. Consider two vectors $Y, Y' \in \mathbb{R}_{\geq 0}^m$ such that $Y' = Y + \delta \mathbb{1}_p$, where $\mathbb{1}_p$ is the m -dimensional vector which has a single non-zero entry 1 corresponding to $p \in P'$. Suppose we are given a minimum cost flow in G_Y . We can route δ additional flow from the vertex of p to w , which incurs $\delta\mu'$ cost. The modified flow is a feasible flow in $G_{Y'}$. Thus, $f(Y') \leq f(Y) + \delta\mu'$.

Similarly, suppose we are given a minimum cost flow in $G_{Y'}$. We can route δ additional flow from w to the vertex of p , which incurs $\delta\mu'$ cost. The modified flow is a feasible flow in G_Y . Thus, $f(Y) \leq f(Y') + \delta\mu'$. Together these show that f is μ' -Lipschitz. \square

Towards this end, we state the following concentration bound, which will be useful in the analysis.

Lemma 3.11. W.p. at least $1 - 1/n^{2k+20}$, $|f(X) - \mathbb{E}[f(X)]| \leq \epsilon m\mu'/2$.

The proof of this lemma is very similar to the proof of Lemma 15 in [34], which essentially follows from the fact that f is μ' -Lipschitz and from the following Chernoff type bound.

Theorem 3.12. [34] Let x_1, \dots, x_n be independent random variables taking value b w.p. p and value 0 w.p. $1 - p$, and let $g : [0, 1]^n \rightarrow \mathbb{R}$ be an L -Lipschitz function in ℓ_1 norm. Define $X := (x_1, \dots, x_n)$ and $\mu := \mathbb{E}[g(X)]$. Then, for $0 \leq \epsilon \leq 1$:

$$\Pr[|g(X) - \mathbb{E}[g(X)]| \geq \epsilon pnbL] \leq 2 \exp(-\epsilon^2 pn/3).$$

We apply the above theorem with $p = s/m$, $n = m$, $b = m/s$, $g = f$ and $L = \mu'$. Then,

$$\begin{aligned} \Pr[|f(X) - \mathbb{E}[f(X)]| \geq \epsilon m\mu'/2] &= \Pr[|f(X) - \mathbb{E}[f(X)]| \geq (\epsilon/2)(s/m) \cdot m \cdot (m/s) \cdot \mu'] \\ &= \Pr[|f(X) - \mathbb{E}[f(X)]| \geq (\epsilon/2) \cdot pnbL] \\ &\leq 2 \exp(-(\epsilon/2)^2 pn/3) \\ &= 2 \exp(-(\epsilon/2)^2 s/3) \\ &= 2 \exp(-(\epsilon^2/12)\Theta(k \log n/\epsilon^3)) \\ &\leq 1/n^{2k+20} \end{aligned}$$

The last inequality follows due to the sufficiently large constant hidden in the Θ notation. Now, we proceed towards the proof of Theorem 3.7. We will show the desired bound in two steps. Here we take a slightly different way than [34] for our convenience. First, we show that w.p. at least $1 - 1/n^{2k+20}$, $f(\mathbb{E}[X]) \leq f(X) + \epsilon m\mu'$. Then, we show that w.p. at least $1 - 1/n^{2k+20}$, $f(X) \leq f(\mathbb{E}[X]) + \epsilon m\mu'$.

The first step. From Lemma 3.11 it follows that it is sufficient to prove $f(\mathbb{E}[X]) \leq \mathbb{E}[f(X)]$. Now, $\mathbb{E}[X] = \mathbb{1}$. Let Y be any outcome of X and $X = Y$ w.p. $p(Y)$. Let y_i be the value in Y corresponding to $p_i \in P'$. Then, there is a feasible flow in G_Y , where for each p_i , at least y_i demand is satisfied. Now, consider the flow ϕ obtained by summing, for each Y , the minimum cost feasible flow in G_Y scaled by $p(Y)$. Note that the cost of ϕ is $\sum_Y p(Y)f(Y) = \mathbb{E}[f(X)]$. Also, this flow does not violate any capacity, as the sum of the probabilities is 1. Now, in each flow corresponding to Y scaled by $p(Y)$, for each p_i , $p(Y) \cdot y_i$ demand is satisfied. Hence, in ϕ , for each p_i , at least $\sum_Y p(Y) \cdot y_i = 1$ demand is satisfied, as the expected value of y_i is 1. It follows that, $f(\mathbb{1}) = f(\mathbb{E}[X])$ is at most the cost of ϕ and we obtain the desired bound.

The second step. Here we will show that w.p. at least $1 - 1/n^{2k+20}$, $f(X) \leq f(\mathbb{E}[X]) + \epsilon m\mu'$. First, we prove that it is sufficient to show that w.p. at least $1 - 1/n^3$, $f(X) \leq f(\mathbb{E}[X]) + \epsilon m\mu'/3$.

Lemma 3.13. If $f(X) \leq f(\mathbb{E}[X]) + \epsilon m\mu'/3$ holds w.p. at least $1 - 1/n^3$, then w.p. $1 - 1/n^{2k+20}$, $f(X) \leq f(\mathbb{E}[X]) + \epsilon m\mu'$.

Proof. Here we will prove that $\mathbb{E}[f(X)] \leq f(\mathbb{E}[X]) + \epsilon m\mu'/2$. Then, by Lemma 3.11, it follows that w.p. $1 - 1/n^{2k+20}$, $f(X) \leq \mathbb{E}[f(X)] + \epsilon m\mu'/2 \leq f(\mathbb{E}[X]) + \epsilon m\mu'$.

First, note that $X \in [0, m/s]^m$. As the function f is μ' -Lipschitz by Observation 3.10, the values of $f(X)$ must lie in an interval of length at most $m/s \cdot m\mu' \leq m^2\mu'$. Similarly, $\mathbb{E}[X] = \mathbb{1} \in [0, m/s]^m$, and thus $f(\mathbb{E}[X])$ is also contained in that interval. Hence, $f(X) \leq f(\mathbb{E}[X]) + m^2\mu'$. Now,

$$\begin{aligned} \mathbb{E}[f(X)] &\leq (1 - 1/n^3) \cdot (f(\mathbb{E}[X]) + \epsilon m\mu'/3) + (1/n^3) \cdot (f(\mathbb{E}[X]) + m^2\mu') \\ &\leq f(\mathbb{E}[X]) + \epsilon m\mu'/3 + (1/n^2) \cdot m\mu' \\ &\leq f(\mathbb{E}[X]) + \epsilon m\mu'/2 \quad \square \end{aligned}$$

The following lemma completes the proof of Theorem 3.7.

Lemma 3.14. *W.p. at least $1 - 1/n^3$, $f(X) \leq f(\mathbb{1}) + \epsilon m\mu'/3$.*

The proof of this lemma is very similar to the proof of Lemma 20 in [34]. For completeness, we describe the proof here.

We consider a minimum cost feasible flow ϕ in G_Y for $Y = \mathbb{1}$. We can assume that this flow is integral, as all the demands and capacities are integral. We compute a feasible flow ϕ' in G_X modifying the flow ϕ whose cost is at most $f(\mathbb{1}) + \epsilon m\mu'/3$ w.p. at least $1 - 1/n^3$.

The construction of the modified flow is as follows. For each point $p \in P \setminus P'$, we route the demand of p in ϕ' in the same way as in ϕ . Now consider the points in P' . Let P'_i be the subset of points of P' that are assigned to the center $c_i \in C$. Also, let Q'_i be the subset of points of P'_i that are sampled, and hence are contained in W' .

For each center $c_i \in C$, we have two cases. The first case is $|Q'_i| \leq |P'_i| \cdot s/m$. In this case, we route m/s amount of flow from each vertex u_j corresponding to the point p_j of Q'_i to the vertex v_i corresponding to c_i . We also route $|P'_i| - |Q'_i| \cdot m/s$ amount of flow from w to v_i . Note that the total amount of flow routed to v_i in these above two steps is exactly $|P'_i|$ and does not depend on $|Q'_i|$ as long as $|Q'_i| \leq |P'_i| \cdot s/m$. In the second case, $|Q'_i| > |P'_i| \cdot s/m$. In this case, first we select a random sample Q''_i from Q'_i of size $\lfloor |P'_i| \cdot s/m \rfloor$ and apply the same steps in the first case with Q''_i instead of Q'_i . Finally, route m/s amount of flow from the vertex corresponding to each point in $Q'_i \setminus Q''_i$ to w .

We note that the computed flow ϕ' in the above satisfies all the demands. Also, none of the capacities are violated, as the flow in and out for each vertex v_i remain the same as in ϕ . In the following we give a bound on the cost of ϕ' . To unify the analysis, in the first case, we set $Q''_i = Q'_i$. We consider two cases depending on the value of $\mathbb{E}[|Q'_i|]$ for every $c_i \in C$. In the first case, $\mathbb{E}[|Q'_i|] \geq \epsilon s/(100k)$ and this implies that the size of P'_i is sufficiently large. The other case is complimentary, where the size of P'_i is small. We handle these two cases separately to prove the lemma. The second case is relatively easier, as here upper-bounding some costs becomes trivial due to the small size of P'_i (and hence of Q'_i).

Case 1. $\mathbb{E}[|Q'_i|] \geq \epsilon s/(100k)$. As Q'_i is distributed as Bernoulli($|P'_i|, s/m$), $\mathbb{E}[|Q'_i|] = |P'_i| \cdot s/m$. Thus, $|P'_i| \cdot s/m \geq \epsilon s/(100k)$, or $|P'_i| \geq \epsilon m/(100k)$. We have the following observation.

Observation 3.15. W.p. at least $1 - 1/n^{10}$, $||Q'_i| - |P'_i| \cdot s/m| \leq \epsilon |P'_i| \cdot s/(50m)$.

Proof. Using the Chernoff bound, $\Pr[||Q'_i| - |P'_i| \cdot s/m| > \epsilon |P'_i| \cdot s/(50m)] \leq \exp(-\Theta(\epsilon^2 \cdot |P'_i| \cdot s/m)) \leq \exp(-\Theta(\epsilon^2 \cdot \epsilon(m/k) \cdot (s/m))) \leq \exp(-\Theta(\log n)) \leq 1/n^{10}$, for sufficiently large constant hidden in $\Theta(\cdot)$ in the definition of s . \square

From the above observation and considering the fact that $|Q''_i| \geq |P'_i| \cdot s/m - 1$, we have the following bound.

Observation 3.16. W.p. at least $1 - 1/n^{10}$, $|Q'_i| - |Q''_i| \leq \epsilon |P'_i| \cdot s/(40m)$.

From the above two observations, we have the following observation.

Observation 3.17. W.p. at least $1 - 1/n^9$, $|P'_i| \cdot s/m - |Q''_i| \leq \epsilon |P'_i| \cdot s/(20m)$.

Now, we give bound on the cost of the computed flow. Note that we route m/s flow for each point in Q''_i to c_i whose total cost is $\sum_{p \in Q''_i} (m/s) \cdot d(p, c_i)$. To give bound on this cost we need the following lemma from [38].

Lemma 3.18. (Lemma 3.2. of [38]) *Let $T \geq 0$ and η be fixed constants, and let $h(\cdot)$ be a function defined on a set V such that $\eta \leq h(p) \leq \eta + T$ for all $p \in V$. Let $U = \{p_1, \dots, p_r\}$ be a set of r samples drawn independently and uniformly from V , and let $\delta > 0$ be a parameter. If $r \geq (T^2/2\delta^2) \ln(2/\lambda)$, then $\Pr[|\frac{h(U)}{|U|} - \frac{h(V)}{|V|}| \geq \delta] \leq \lambda$, where $h(U) = \sum_{u \in U} h(u)$ and $h(V) = \sum_{v \in V} h(v)$.*

Fix any integer $r \in [1 - \epsilon/20, 1] \cdot |P'_i| \cdot s/m$ and consider the event that $|Q''_i| = r$. Conditioned on this event Q''_i is a set of r samples drawn independently and uniformly from P'_i . We apply Lemma 3.18 setting $T = 2\mu'$, $V = P'_i$, $U = Q''_i$, $h(p) = d(p, c_i)$, $\delta = \epsilon\mu'/20$ and $\lambda = 1/n^{10}$. Note that,

$$r \geq (1 - \epsilon/20) \cdot |P'_i| \cdot s/m \geq (1 - \epsilon/20) \cdot \frac{\epsilon m}{100k} \cdot \frac{s}{m} \geq \Theta(\log n / \epsilon^2) \geq (T^2/2\delta^2) \ln(2/\lambda)$$

The last inequality follows assuming a sufficiently large constant is hidden in $\Theta(\cdot)$ in the definition of s .

We obtain, w.p. at least $1 - 1/n^{10}$,

$$\left| \frac{h(P'_i)}{|P'_i|} - \frac{h(Q''_i)}{|Q''_i|} \right| \leq \delta$$

$$\begin{aligned}
& \text{Or, } \left| \frac{h(P'_i) \cdot |Q''_i|}{|P'_i|} - h(Q''_i) \right| \leq \delta |Q''_i| \\
& \text{Or, } h(Q''_i) \leq \frac{h(P'_i) \cdot |Q''_i|}{|P'_i|} + \delta \cdot (|P'_i| \cdot s/m) \\
& \text{Or, } h(Q''_i) \cdot (m/s) \leq \frac{h(P'_i) \cdot |Q''_i|}{|P'_i|} \cdot (m/s) + \epsilon |P'_i| \cdot \mu' / 20
\end{aligned}$$

Note that $h(Q''_i) \cdot (m/s)$ is exactly the cost of flow routing for the points in Q''_i to c_i . Taking union bound over all possible values of $r = |Q''_i|$, we obtain the above bound w.p. at least $1 - 1/n^9$.

Now, we give a bound on the cost of flow routing from w to c_i . The cost is $(|P'_i| - |Q''_i| \cdot m/s) \cdot d(c', c_i)$. Now, for any $p \in P'_i$, $d(c', c_i) \leq d(p, c_i) + \mu'$. Averaging gives, $d(c', c_i) \leq h(P'_i)/|P'_i| + \mu'$. Thus, the cost is at most,

$$\begin{aligned}
(|P'_i| - |Q''_i| \cdot m/s) \cdot (h(P'_i)/|P'_i| + \mu') & \leq h(P'_i) - \frac{h(P'_i) \cdot |Q''_i|}{|P'_i|} \cdot (m/s) + (|P'_i| - |Q''_i| \cdot m/s) \cdot \mu' \\
& \leq h(P'_i) - \frac{h(P'_i) \cdot |Q''_i|}{|P'_i|} \cdot (m/s) + (\epsilon |P'_i| / 20) \mu' \\
& = h(P'_i) - \frac{h(P'_i) \cdot |Q''_i|}{|P'_i|} \cdot (m/s) + \epsilon |P'_i| \mu' / 20
\end{aligned}$$

The second inequality follows from Observation 3.17. Next, we bound the third and the last type of cost, which corresponds to flow routing from points in $Q'_i \setminus Q''_i$ to w . This cost is at most,

$$\begin{aligned}
\sum_{p \in (Q'_i \setminus Q''_i)} (m/s) \cdot d(p, c') & \leq (|Q'_i| - |Q''_i|) \cdot (m/s) \cdot \mu' \\
& \leq (\epsilon |P'_i| \cdot s / (40m)) \cdot (m/s) \cdot \mu' \\
& \leq \epsilon |P'_i| \mu' / 40
\end{aligned}$$

The second inequality follows from Observation 3.16. Thus, in this case, the total cost is bounded by,

$$h(P'_i) + \epsilon |P'_i| \mu' / 20 + \epsilon |P'_i| \mu' / 20 + \epsilon |P'_i| \mu' / 40 \leq h(P'_i) + \epsilon |P'_i| \mu' / 8.$$

Case 2. $\mathbb{E}[|Q'_i|] < \epsilon s / (100k)$. Note that in this case, $|P'_i| < \epsilon m / (100k)$. First, we have the following observation.

Observation 3.19. W.p. at least $1 - 1/n^{10}$, $|Q'_i| \leq \epsilon s / (50k)$.

Proof. We use the Chernoff bound: $\Pr[|Q'_i| \geq \epsilon s / (50k)] \leq \exp(-\Theta(\epsilon s / k)) \leq \exp(-\Theta(\log n)) \leq 1/n^{10}$. \square

The cost of flow routing from points in Q''_i to c_i is,

$$\begin{aligned}
\sum_{p \in Q''_i} (m/s) \cdot d(p, c_i) & \leq \sum_{p \in Q''_i} (d(p, c') + d(c', c_i)) \cdot (m/s) \\
& \leq |Q''_i| \cdot \mu' \cdot (m/s) + |Q''_i| \cdot (m/s) \cdot d(c', c_i)
\end{aligned}$$

The cost of flow routing from w to c_i is,

$$\begin{aligned}
(|P'_i| - |Q''_i| \cdot m/s) \cdot d(c', c_i) & \leq \sum_{p \in P'_i} d(c', c_i) - |Q''_i| \cdot (m/s) \cdot d(c', c_i) \\
& \leq \sum_{p \in P'_i} (d(c', p) + d(p, c_i)) - |Q''_i| \cdot (m/s) \cdot d(c', c_i) \\
& \leq |P'_i| \cdot \mu' + \sum_{p \in P'_i} d(p, c_i) - |Q''_i| \cdot (m/s) \cdot d(c', c_i)
\end{aligned}$$

The cost of flow routing from points in $Q'_i \setminus Q''_i$ to w is,

$$\begin{aligned} & \sum_{p \in (Q'_i \setminus Q''_i)} (m/s) \cdot d(p, c') \\ & \leq |Q'_i| \cdot (m/s) \cdot \mu' \end{aligned}$$

The total cost in this case is at most,

$$\begin{aligned} & |Q''_i| \cdot \mu' \cdot (m/s) + |Q''_i| \cdot (m/s) \cdot d(c', c_i) + |P'_i| \cdot \mu' + \sum_{p \in P'_i} d(p, c_i) - |Q''_i| \cdot (m/s) \cdot d(c', c_i) + |Q'_i| \cdot (m/s) \cdot \mu' \\ & \leq (\epsilon s/(50k)) \cdot \mu' \cdot (m/s) + (\epsilon m/(100k)) \cdot \mu' + \sum_{p \in P'_i} d(p, c_i) + (\epsilon s/(50k)) \cdot \mu' \cdot (m/s) \\ & \leq (\epsilon m/(20k)) \cdot \mu' + \sum_{p \in P'_i} d(p, c_i) \end{aligned}$$

The first inequality follows from Observation 3.19 and by noting that $|Q''_i| \leq |Q'_i|$.

General upper bound on the cost. By merging the cost in both cases, we obtain the common upper bound, $\sum_{p \in P'_i} d(p, c_i) + (\epsilon |P'_i| \cdot \mu'/8) + (\epsilon m/(20k)) \cdot \mu'$. Summing over all the centers in C , we obtain,

$$f(X) \leq f(\mathbb{1}) + \epsilon |P'| \cdot \mu'/8 + (\epsilon m/20) \cdot \mu' \leq f(\mathbb{1}) + \epsilon |P'| \cdot \mu'/3.$$

It is not hard to see that this bound holds w.p. at least $1 - 1/n^3$. This completes the proof of Lemma 3.14.

3.3.2. Multiple ring case

In the previous section, we have shown how to bound the error for a fixed ring. Here we extend the ideas to the multiple ring case. Intuitively, we use a union bound over all rings to obtain the desired high probability bound. However, we need to consider the samples from all the rings corresponding to the color γ together. Let W' be the corresponding coreset.

We consider any arbitrary ordering of all the rings, and for any two rings $B'_{i,j}$ and $B'_{i',j'}$, we say $(i, j) < (i', j')$ if $B'_{i,j}$ precedes $B'_{i',j'}$ in this ordering. Consider any ring $B'_{i,j}$. We define a function $f_{i,j}$ corresponding to this ring similar to the function f . Let $P'_{i,j} = P'_{i,j,\gamma}$. Also, let $W'_{i,j}$ be the samples chosen from $P'_{i,j}$. The input to the function $f_{i,j}$ is a vector $Y \in \mathbb{R}_{\geq 0}^{|P'_{i,j}|}$ that is indexed by the points of the ring. We construct a network G_Y as before. But, as we consider samples from all the rings, the demands of the points are defined in a different way than before. For each point in $p \in P'_{i,j}$, its demand is $Y[p]$. Set the demand of w to $|P'_{i,j}| - \sum_{p \in P'_{i,j}} Y[p]$. For each ring $B'_{i',j'} \neq B'_{i,j}$, and for each point $p \in W'_{i',j'}$, set its demand to $|P'_{i',j'}|/s$. Note that the total demand corresponding to $B'_{i',j'}$ is $s \cdot |P'_{i',j'}|/s = |P'_{i',j'}|$. Thus, in G_Y we fix the samples of all the rings except $B'_{i,j}$. $f_{i,j}(Y)$ is the cost of the minimum cost flow in G_Y .

Let $\mathbb{E}_{W'_{i,j}:(i',j') > (i,j)}[f_{i,j}(Y)|W'_{i_1,j_1} : (i_1, j_1) < (i, j)]$ ($\mathbb{E}_{>(i,j)}[f_{i,j}(Y)]$ in short) be the expectation of $f_{i,j}(Y)$ over all samples $W'_{i',j'}$ for $(i', j') > (i, j)$ given fixed samples W'_{i_1,j_1} for all $(i_1, j_1) < (i, j)$. Similarly, define

$$\mathbb{E}_{W'_{i,j}:(i',j') \geq (i,j)}[f_{i,j}(Y)|W'_{i_1,j_1} : (i_1, j_1) < (i, j)]$$

or $\mathbb{E}_{\geq(i,j)}[f_{i,j}(Y)]$ in short. Recall that in the single ring case we showed that w.p. at least $1 - 1/n^{2k+10}$, $|f(X) - f(\mathbb{E}[X])| \leq \epsilon m \mu^7$. Similarly, here we obtain the following lemma.

Lemma 3.20. *W.p. at least $1 - 1/n^{2k+10}$, for any ring $B'_{i,j}$,*

$$|\mathbb{E}_{>(i,j)}[f_{i,j}(Y)] - \mathbb{E}_{\geq(i,j)}[f_{i,j}(Y)]| \leq \epsilon |P'_{i,j}| \cdot 2^j \mu.$$

Note that we would like to show the bound in terms of multiple rings together instead of just one ring $B'_{i,j}$. In particular, we would like to give a bound w.r.t. $\text{wcost}(W', M, C)$, where M is a column matrix. Correspondingly we define $\mathbb{E}_{>(i,j)}[\text{wcost}(W', M, C)]$ and $\mathbb{E}_{\geq(i,j)}[\text{wcost}(W', M, C)]$. From Lemma 3.20, we readily obtain the following lemma.

Lemma 3.21. *W.p. at least $1 - 1/n^{2k+8}$,*

$$|\mathbb{E}_{>(i,j)}[\text{wcost}(W', M, C)] - \mathbb{E}_{\geq(i,j)}[\text{wcost}(W', M, C)]| \leq \epsilon |P'_{i,j}| \cdot 2^j \mu.$$

Now consider going over all the rings in the ordering and applying the above lemma. Let (i_1, j_1) and (i', j') be the indexes of the first and last ring, respectively. Then the total deviation between $\mathbb{E}_{\geq(i_1, j_1)}[\text{wcost}(W', M, C)]$ and $\mathbb{E}_{>(i', j')}[\text{wcost}(W', M, C)]$ is at most $\sum_{(i,j) \in P'_{i,j}} |P'_{i,j}| \cdot 2^j \mu$. But, $\mathbb{E}_{\geq(i_1, j_1)}[\text{wcost}(W', M, C)] = \text{cost}(P, M, C)$ and $\mathbb{E}_{>(i', j')}[\text{wcost}(W', M, C)] = \text{wcost}(W', M, C)$, and hence by taking union bound over all $O(k \log n)$ rings, we obtain the following lemma.

Lemma 3.22. *For any fixed set C of k centers and any fixed column matrix M , w.p. at least $1 - 1/n^{2k+5}$, $|\text{cost}(P, M, C) - \text{wcost}(W', M, C)| \leq \sum_{(i,j) \in P'_{i,j}} |P'_{i,j}| \cdot 2^j \mu$.*

By taking union bound over all column matrices M , we obtain the desired bound.

Lemma 3.23. *For any fixed set C of k centers and for all column matrices M , w.p. at least $1 - 1/n^{k+4}$, $|\text{cost}(P, M, C) - \text{wcost}(W', M, C)| \leq \sum_{(i,j) \in P'_{i,j}} |P'_{i,j}| \cdot 2^j \mu$.*

4. Coreset construction for k -median in the overlapping group case

In this section, we prove the following theorem.

Theorem 4.1. *Given a collection of ℓ possibly overlapping groups consisting of n points in total in a metric space, there is an $O(n(k+\ell))$ time randomized algorithm that w.p. at least $1 - 1/n$, computes a universal coreset for k -median clustering of size $O(\Gamma(k \log n)^2 / \epsilon^3)$.*

To prove the above theorem, we design a sampling based algorithm. Note that the algorithm in the disjoint case does not work here. This is because we sample points from each group separately and independently, and thus it is not clear how to assign the weight of a point that belongs to multiple groups. One might think of the following trivial modification of this algorithm. Assign each point to a single group to which it belongs. Based on this assignment, now we have disjoint groups, and we can apply our previous algorithm. But, the new algorithm can have a very large error bound. For example, suppose a point p belongs to two groups i and j , and it is assigned to group i . Also, suppose p was not chosen in the sampling process. Note that the weight of p is represented by some other chosen point p' , which was also assigned to group i . However, now we have lost the information that this weight of p was also contributing towards fairness of group j . Thus, the constructed coreset might not preserve any optimal fair clustering with a small error. In the overlapping case, it is not clear how to obtain a coreset whose size depends linearly in ℓ – we design a new coreset construction algorithm where the size depends linearly on Γ .

The main idea of our algorithm is to divide the points into equivalence classes based on their group membership and sample points from each equivalence class. Let $P = \bigcup_{i=1}^{\ell} P_i$. For each point $p \in P$, let $J_p \subseteq [\ell]$ be the set of indexes of the groups to which p belongs. Let I be the distinct collection of these sets $\{J_p \mid p \in P\}$ and $|I| = \Gamma$. In particular, let I_1, \dots, I_{Γ} be the distinct sets in I . Now, we partition the points in P based on these sets. For $1 \leq i \leq \Gamma$, let $P^i = \{p \in P \mid I_i = J_p\}$. Thus, $\{P^i \mid 1 \leq i \leq \Gamma\}$ defines equivalence classes for P such that two points $p, p' \in P$ belong to the same equivalence class if they are in exactly the same set of groups.

In the overlapping case, we will work with an even stronger definition of coresets. This is for the ease of computation of an optimal cost assignment of the points in the coreset. Here instead of $k \times \ell$ matrices, coloring constraints are defined by $k \times \Gamma$ matrices. The rows still correspond to k centers, but the columns now correspond to the Γ equivalence classes. Thus, for such a matrix M , M_{ij} denotes the number of points from P^j that are in cluster i . Thus, the entries of M define a partition of the points in P . We note that Proposition 2.2 continues to hold, as any fair assignment of the points in P defines such a matrix M . Now, the definition of universal coreset remains same, except here $\text{wcost}(W, M, C)$ is defined in the following natural way.

Suppose we are given a weight function $w : P \rightarrow \mathbb{R}_{\geq 0}$. Let $W \subseteq P \times \mathbb{R}$ be the set of pairs $\{(p, w(p)) \mid p \in P \text{ and } w(p) > 0\}$. For a set of centers $C = \{c_1, \dots, c_k\}$ and a coloring constraint M , $\text{wcost}(W, M, C)$ is the minimum value $\sum_{p \in P, c_i \in C} \psi(p, c_i) \cdot d(p, c_i)$ over all assignments $\psi : P \times C \rightarrow \mathbb{R}_{\geq 0}$ such that

1. For each $p \in P$, $\sum_{c_i \in C} \psi(p, c_i) = w(p)$.
2. For each $c_i \in C$ and class $1 \leq j \leq \Gamma$, $\sum_{p \in P^j} \psi(p, c_i) = M_{ij}$.

If there is no such assignment ψ , $\text{wcost}(W, M, C) = \infty$. When $w(p) = 1$ for all $p \in P$, we simply denote W by P and $\text{wcost}(W, M, C)$ by $\text{cost}(P, M, C)$. Note that for a fixed matrix M , an optimal assignment ψ must be integral due to integrality of flow. This was not-necessarily true with our previous definition in the overlapping case. We will compute a coreset that satisfies this even stronger definition.

With the above definitions, our algorithm in the overlapping case is a natural extension of the one in the disjoint case. The main idea of our algorithm is to divide the points into disjoint equivalence classes based on their group membership and sample points from each equivalence class. We compute the disjoint classes $\{P^i \mid 1 \leq i \leq \Gamma\}$ defined above. Then, apply our algorithm in the disjoint case on these disjoint sets of points P^1, \dots, P^{Γ} . Let W be the constructed coreset.

4.1. The analysis

Recall that $P_{i,j}$ is the total number of points in each ring $B'_{i,j}$. We will prove the following lemma.

Lemma 4.2. *For any fixed set C of k centers and for all $k \times \Gamma$ matrix M , w.p. at least $1 - 1/n^{k+2}$, $|\text{cost}(P, M, C) - \text{wcost}(W, M, C)| \leq \sum_{(i,j) \in |P_{i,j}|} 2^j \mu$.*

Like before, by taking union bound over all C , we obtain the desired result. This completes the proof of Theorem 4.1. Next, we prove Lemma 4.2.

4.2. Proof of Lemma 4.2

Note that P^τ is the points in P from class τ for $1 \leq \tau \leq \Gamma$. Let W_τ be the chosen samples from class τ . For any ring $B'_{i,j}$, let $P'_{i,j,\tau}$ be the points from class τ in the ring. Also, let $P_{i,j} = \bigcup_{\tau=1}^{\Gamma} P'_{i,j,\tau}$.

Like in the disjoint case, here also we will prove the following lemma that gives a bound when the coreset contains sampled points from a fixed class t and original points from the other classes.

Lemma 4.3. *Consider any class $1 \leq t \leq \Gamma$. For any fixed set C of k centers and for all $k \times \Gamma$ matrix M , w.p. at least $1 - 1/n^{k+4}$, $|\text{cost}(P, M, C) - \text{wcost}(W_t \cup (P \setminus P^t), M, C)| \leq \sum_{(i,j) \in |P'_{i,j,t}|} 2^j \mu$.*

By using expectation argument similar to the one in the disjoint-group-multiple-ring case and taking union bound over all $\Gamma < n$ classes, Lemma 4.2 follows. Next, we prove Lemma 4.3.

4.3. Proof of Lemma 4.3

We have the following lemma that implies that it is sufficient to consider the points only in P^t to give the error bound.

Lemma 4.4. *Suppose w.p. at least $1 - 1/n^{k+4}$, for all $k \times 1$ matrix M' such that the sum of the entries in each column is exactly $|P^t|$, $|\text{cost}(P^t, M', C) - \text{wcost}(W_t, M', C)| \leq \sum_{(i,j) \in |P'_{i,j,t}|} 2^j \mu$. Then, with the same probability, for all $k \times \Gamma$ matrix M , $|\text{cost}(P, M, C) - \text{wcost}(W_t \cup (P \setminus P^t), M, C)| \leq \sum_{(i,j) \in |P'_{i,j,t}|} 2^j \mu$.*

Proof. Consider any $k \times \Gamma$ matrix M . Also consider a clustering C_1, \dots, C_k of P that has cost $\text{cost}(P, M, C)$. We construct two $k \times \Gamma$ matrices M_1 and M_2 from M . For $j \neq t$, and for $1 \leq i \leq k$, $M_1[i][j] = 0$ and $M_2[i][j] = M[i][j]$. For $1 \leq i \leq k$, $M_1[i][t] = |C_i \cap P^t|$ and $M_2[i][t] = M[i][t] - |C_i \cap P^t|$.

$$\text{cost}(P, M, C) = \text{cost}(P^t, M_1, C) + \text{cost}(P \setminus P^t, M_2, C)$$

Also, as $W_t \subseteq P^t$ and the sum of the weights of the points in W_t is $|P^t|$,

$$\text{wcost}(W_t \cup (P \setminus P^t), M, C) = \text{wcost}(W_t, M_1, C) + \text{cost}(P \setminus P^t, M_2, C)$$

It follows that,

$$|\text{cost}(P, M, C) - \text{wcost}(W_t \cup (P \setminus P^t), M, C)| = |\text{cost}(P^t, M_1, C) - \text{wcost}(W_t, M_1, C)|$$

Let M'_1 be the t^{th} column of M_1 . Now, considering the fact that P^t does not contain any points from any other classes,

$$\text{cost}(P^t, M_1, C) - \text{wcost}(W_t, M_1, C) = \text{cost}(P^t, M'_1, C) - \text{wcost}(W_t, M'_1, C).$$

Also, by the definition of M_1 , the sum of the entries in M'_1 is $\sum_{i=1}^k |C_i \cap P^t| = |P^t|$.

Now, by our assumption, it follows that the probability of the event: for all M , $|\text{cost}(P^t, M'_1, C) - \text{wcost}(W_t, M'_1, C)|$ exceeds $\sum_{(i,j) \in |P'_{i,j,t}|} 2^j \mu$ is at most $1/n^{k+4}$. Hence, the lemma follows. \square

By the above observation, it is sufficient to prove that w.p. at least $1 - 1/n^{k+4}$, for all $k \times 1$ matrix M such that the sum of the entries in each column is exactly $|P^t|$, $|\text{cost}(P^t, M, C) - \text{wcost}(W_t, M, C)| \leq \sum_{(i,j) \in |P'_{i,j,t}|} 2^j \mu$. Now, as we select samples from P^t separately and independently, this claim boils down to the corresponding claim in the disjoint case. Recall that we proved this claim for a single ring first, and then extended to multiple rings. The proof of our claim here is very similar, and thus we omit it.

5. Coreset construction for k -median in \mathbb{R}^d

In this section, we prove the following theorem.

Theorem 5.1. *Given a collection of ℓ possibly overlapping groups consisting of n points in total in \mathbb{R}^d , there is an $O(nd(k + \ell))$ time randomized algorithm that w.p. at least $1 - 1/n$, computes a universal coreset for Euclidean k -median clustering of size $O\left(\frac{\Gamma}{\epsilon^3} \cdot k^2 \log n (\log n + d \log(1/\epsilon))\right)$.*

The algorithm in the Euclidean case is the same as for general metrics, except we set s to $\Theta(k \log(nb)/\epsilon^3)$ instead of $\Theta(k \log n/\epsilon^3)$, where $b = \Theta(k \log(n/\epsilon)/\epsilon^d)$. The analysis for general metrics holds in this case, but the assumption that the number of distinct sets of centers is at most n^k is no longer true. Here any point in \mathbb{R}^d is a potential center. Nevertheless, we show that for every set $C \subseteq \mathbb{R}^d$ of k centers and constraint M , the optimal cost is preserved approximately w.h.p. The idea is to use a discretization technique to obtain a finite set of centers so that if instead we draw centers from this set, the cost of any clustering is preserved approximately.

In the following, we analyze the coreset construction algorithm in the overlapping case. First, we construct a set of points F that we will use as the center set. Recall that C^* is the set of centers computed by the bicriteria approximation algorithm. ν is the constant approximation factor and Π is the cost of clustering. Also, $\mu = \Pi/(\nu n)$. Note that for any point p , $d(p, C^*) \leq \Pi = \nu n \cdot \mu$.

For each center $c_i^* \in C^*$, we consider the d -dimensional axis-parallel hypercubes $R_{i,j}$ having sidelength $2^j \mu$ and centered at c_i^* for $0 \leq j \leq N$, where $N = \lceil \log(14\nu n/\epsilon) \rceil$. We note that any point at a distance $(2^N \mu)/2 \geq 7\nu n \cdot \mu/\epsilon$ from c_i^* is in $R_{i,N}$. Let $R'_{i,0} = R_{i,0}$ and $R'_{i,j} = R_{i,j} \setminus R_{i,j-1}$ for $1 \leq j \leq N$. For each $0 \leq j \leq N$, we divide $R'_{i,j}$ into gridcells of sidelength $(\epsilon 2^j \mu)/(10\nu)$. Let Q_i be the exponential grid for $R'_{i,0}, \dots, R'_{i,N}$, i.e., Q_i is the amalgamation of the gridcells in $R'_{i,0}, \dots, R'_{i,N}$. For each gridcell in the exponential grid Q_i , we select any arbitrary point and add it to F_i .

We repeat the above process for all $c_i^* \in C^*$. Let $F = \cup_i F_i$. Note that the total number of gridcells of Q_i is at most $O(\log(n/\epsilon)/\epsilon^d)$. Now, from each such gridcell, we pick at most 1 point. As C^* contains $O(k)$ centers, the size of F is $O(k \log(n/\epsilon)/\epsilon^d)$.

Note that if the centers can only be chosen from F , then by the analysis for general metrics, we obtain the following lemma.

Lemma 5.2. *For any fixed set $C \subseteq F$ of k centers and for all $k \times \Gamma$ matrices M , w.p. at least $1 - 1/(bn)^{k+2}$, $|\text{cost}(P, M, C) - \text{wcost}(W, M, C)| \leq \epsilon \cdot \text{cost}(P, M, C)$.*

This lemma is similar to Lemma 4.2. The error probability is now $1/(bn)^{k+2}$ as s is set to the larger value $\Theta(k \log(nb)/\epsilon^3)$ instead of $\Theta(k \log n/\epsilon^3)$. Now the number of distinct sets of k centers from F is at most $|F|^k \leq b^k$. Thus, by taking union bound over all such sets, we obtain the bound in the above lemma for every $C \subseteq F$ w.h.p.

Lemma 5.3. *For every set $C \subseteq F$ of k centers and for all $k \times \Gamma$ matrices M , w.p. at least $1 - 1/n^2$, $|\text{cost}(P, M, C) - \text{wcost}(W, M, C)| \leq \epsilon \cdot \text{cost}(P, M, C)$.*

Next, we show that if in a clustering a center c is chosen that is not in any of the exponential grids considered before, then W preserves the cost of such clustering.

Lemma 5.4. *Consider any set $C \subseteq \mathbb{R}^d$ of k centers containing a center \hat{c} such that a point $\hat{p} \in P$ is assigned to \hat{c} in a clustering that satisfies a constraint M . Moreover, suppose \hat{c} is not in $\cup_i Q_i$. Then, $|\text{cost}(P, M, C) - \text{wcost}(W, M, C)| \leq \epsilon \cdot \text{cost}(P, M, C)$.*

Proof. Consider any class P^t and a ring $B'_{i,j}$. Let $P'_{i,j,t}$ be the points in $B'_{i,j}$ from P^t and $W_{i,j,t}$ be the points of $P'_{i,j,t}$ that are in W . Then, there is an assignment $\phi: P'_{i,j,t} \rightarrow W_{i,j,t}$ such that exactly $|P'_{i,j,t}|/|W_{i,j,t}|$ points are assigned to each point $q \in W_{i,j,t}$. Note that $d(p, \phi(p)) \leq d(p, c_i^*) + d(c_i^*, \phi(p)) \leq 2^j \mu + 2^j \mu = 2^{j+1} \mu$. Now, consider an optimal assignment ψ corresponding to $\text{cost}(P, M, C)$. We compute the following assignment for each $1 \leq t \leq \Gamma$ and ring $B'_{i,j}$. Assign 1 weight of each point $\phi(p) \in W_{i,j,t}$ to the center of C where p is assigned in ψ . (WLOG, one can assume that the weights of our coreset points are integral.) Note that for each point in $W_{i,j,t}$ exactly $|P'_{i,j,t}|/|W_{i,j,t}|$ amount of weight has been assigned. The new assignment for coreset points induces a valid clustering and satisfies M . By triangle inequality it follows that,

$$|\text{cost}(P, M, C) - \text{wcost}(W, M, C)| \leq \sum_{t=1}^{\Gamma} \sum_{(i,j)} \sum_{p \in P'_{i,j,t}} d(p, \phi(p))$$

$$\begin{aligned}
&\leq \sum_{t=1}^{\Gamma} \sum_{(i,j)} \sum_{p \in P'_{i,j,t}} 2^{j+1} \mu \\
&\leq \sum_{(i,j)} \sum_{p \in P_{i,j}} 2^{j+1} \mu \\
&= \sum_{i=1}^k \sum_{p \in P_{i,0}} 2\mu + \sum_{p \in P_{i,j} | j \geq 1} 2^{j+1} \mu \\
&\leq 2n\mu + 4 \sum_{i=1}^k \sum_{p \in P_i^*} d(p, c_i^*) \\
&\leq 2 \cdot \text{OPT}_v + 4 \cdot \Pi \leq 6 \cdot \Pi
\end{aligned}$$

Here $P_i^* \subseteq P$ is the set of points assigned to c_i^* . The second last inequality follows, as for each point $p \in P_{i,j}$ with $j \geq 1$, $d(p, c_i^*) \geq 2^{j-1} \mu$. Now there is a point \hat{p} that is assigned to $\hat{c} \in C$ such that \hat{c} is not in $\cup_i Q_i$. Let $\hat{p} \in P_i^*$. It follows that,

$$\text{cost}(P, M, C) \geq d(\hat{c}, \hat{p}) \geq d(\hat{c}, c_i^*) - d(c_i^*, \hat{p}) \geq 7vn \cdot \mu/\epsilon - vn\mu \geq 6vn \cdot \mu/\epsilon = 6 \cdot \Pi/\epsilon$$

The third inequality follows, as $d(\hat{c}, c_i^*) > (2^N \mu)/2 \geq vn \cdot \mu/\epsilon$ and $d(c_i^*, \hat{p}) \leq \Pi = vn\mu$. Thus, $\Pi \leq \epsilon \cdot \text{cost}(P, M, C)/6$. Hence,

$$|\text{cost}(P, M, C) - \text{wcost}(W, M, C)| \leq 6 \cdot \Pi \leq \epsilon \cdot \text{cost}(P, M, C). \quad \square$$

Next, we consider the case when all points in C are in $\cup_i Q_i$. Let $C' \subseteq F$ be the set of centers constructed by replacing each point c in C , by the representative of the gridcell that contains c . Then, we have the following observation.

Observation 5.5. $|\text{cost}(P, M, C) - \text{cost}(P, M, C')| \leq \epsilon \cdot \text{OPT}_v$ and $|\text{cost}(W, M, C) - \text{cost}(W, M, C')| \leq \epsilon \cdot \text{OPT}_v$.

Lemma 5.6. For every set $C \subseteq \mathbb{R}^d$ of k centers such that all centers are contained in $\cup_i Q_i$ and for all constraint M , w.p. at least $1 - 1/n^2$, $|\text{cost}(P, M, C) - \text{wcost}(W, M, C)| \leq \epsilon \cdot \text{cost}(P, M, C)$.

Proof. Define the set C' from C as above. It follows that,

$$\begin{aligned}
&|\text{cost}(P, M, C) - \text{wcost}(W, M, C)| \\
&\leq |\text{cost}(P, M, C) - \text{cost}(P, M, C') + \text{cost}(P, M, C') - \text{wcost}(W, M, C') + \text{wcost}(W, M, C') - \text{wcost}(W, M, C)| \\
&\leq |\text{cost}(P, M, C) - \text{cost}(P, M, C')| + |\text{cost}(P, M, C') - \text{wcost}(W, M, C')| + |\text{wcost}(W, M, C') - \text{wcost}(W, M, C)| \\
&\leq 2\epsilon \cdot \text{OPT}_v + |\text{cost}(P, M, C') - \text{wcost}(W, M, C')|
\end{aligned}$$

The last inequality follows from Observation 5.5. By Lemma 5.3, we obtain for every C and all M , w.p. at least $1 - 1/n^2$,

$$\begin{aligned}
|\text{cost}(P, M, C) - \text{wcost}(W, M, C)| &\leq 2\epsilon \cdot \text{OPT}_v + \epsilon \cdot \text{cost}(P, M, C') \\
&\leq \epsilon \cdot \text{cost}(P, M, C) + 3\epsilon \cdot \text{OPT}_v \\
&\leq 4\epsilon \cdot \text{cost}(P, M, C)
\end{aligned}$$

By scaling ϵ by a factor of 4, the lemma follows. \square

Now, $\Theta(\log(nb)) = \Theta(\log n + \log k + \log \log(n/\epsilon) + d \log(1/\epsilon)) = \Theta(\log n + d \log(1/\epsilon))$. Thus $s = \Theta\left(\frac{1}{\epsilon^3} \cdot k(\log n + d \log(1/\epsilon))\right)$. By Lemmas 5.6 and 5.4, Theorem 5.1 follows.

6. Coreset construction for k -means clustering

Here we describe the changes needed to extend the coreset construction scheme for k -median to k -means. Apart from the technical differences, one major feature of k -means compared to k -median is that well-known dimensionality reduction techniques allow readily to reduce to a small-dimensional instance in the Euclidean case. Thus, we can obtain a compression that is independent of d , however we postpone this discussion until Subsection 8.2, as this does not fit precisely with our definition of a coreset.

First, we consider the disjoint group case. The cores set construction algorithm is identical except here from each ring and for each color, we select a sample of size $O(k \log n / \epsilon^5)$. The analysis remains almost the same except in places we obtain worse bounds due to squaring of the distances.

In the single ring-single color case, instead of Theorem 3.7, we have the following modified theorem.

Theorem 6.1. *W.p. at least $1 - 1/n^{2k+10}$, it holds that $|\text{cost}(P, M, C) - \text{wcost}(W', M, C)| \leq \epsilon m \mu'^2 + O(\epsilon) \cdot \text{cost}(P, M, C)$.*

The network G_Y is defined in a different way in this case to deal with the square of distances. In particular, we adapt a bipartite matching framework. The points (sources) have positive demands and are placed on the left side, and centers (sinks) have negative demands and are placed on the right. If the demand $m - \sum_{p \in P'} Y[p]$ corresponding to the bicriteria center c' is non-negative, it is placed on the left as a source. Otherwise, it is placed on the right as a sink. The costs of the edges are now set to square of the corresponding distances.

Lemma 3.11 continues to hold even in this case. Thus for the same reason we readily obtain, w.p. $1 - 1/n^{2k+20}$, $f(\mathbb{E}[X]) \leq f(X) + \epsilon m \mu'$. To prove, w.p. $1 - 1/n^{2k+20}$, $f(X) \leq f(\mathbb{E}[X]) + \epsilon m \mu'$, we need to show, w.p. at least $1 - 1/n^3$, $f(X) \leq f(\mathbb{E}[X]) + \epsilon m \mu'/3$. Here we need significant amount of changes in the analysis. Again we have two cases based on the expectation of $|Q'_i|$. Here we need a slightly different bound on the expectation $\epsilon^3 s / (100k)$ instead of $\epsilon s / (100k)$.

Case 1. $\mathbb{E}[|Q'_i|] \geq \epsilon^3 s / (100k)$. In this case, $|P'_i| \cdot s / m \geq \epsilon^3 s / (100k)$, or $|P'_i| \geq \epsilon^3 m / (100k)$. Note that Observation 3.15 continues to hold, as s is set to $\Theta(k \log n / \epsilon^5)$, and thus Observations 3.16 and 3.17 as well.

Now, we give bound on the cost of the computed flow. Note that we route m/s flow for each point in Q'_i to c_i whose total cost is $\sum_{p \in Q'_i} (m/s) \cdot d(p, c_i)$.

For points $p \in P'_i$, the distances $d(p, c_i)$ lie in an interval of length at most $2\mu'$. Thus the average of these distances must also lie in this interval. It follows that,

$$\begin{aligned} d(p, c_i)^2 &\leq \left(\frac{1}{|P'_i|} \cdot \sum_{p \in P'_i} d(p, c_i) + 2\mu' \right)^2 \\ &\leq 2 \left(\frac{1}{|P'_i|} \cdot \sum_{p \in P'_i} d(p, c_i) \right)^2 + 8\mu'^2 \\ &\leq \frac{2}{(|P'_i|)^2} \cdot |P'_i| \cdot \sum_{p \in P'_i} d(p, c_i)^2 + 8\mu'^2 \\ &\leq \frac{2}{|P'_i|} \cdot \sum_{p \in P'_i} d(p, c_i)^2 + 8\mu'^2 \end{aligned}$$

The second last inequality follows from Cauchy-Schwarz's inequality. Now, we can apply Lemma 3.18 setting $T = 8\mu'^2$, $V = P'_i$, $U = Q'_i$, $h(p) = d(p, c_i)$, $\delta = \epsilon \mu'^2 / 20$ and $\lambda = 1/n^{10}$. Note that,

$$r \geq (1 - \epsilon/20) \cdot |P'_i| \cdot s / m \geq (1 - \epsilon/20) \cdot \frac{\epsilon^3 m}{100k} \cdot \frac{s}{m} \geq \Theta(\log n / \epsilon^2) \geq (T^2 / 2\delta^2) \ln(2/\lambda)$$

The last inequality follows assuming a sufficiently large constant is hidden in $\Theta(\cdot)$ in the definition of s .

We obtain, w.p. at least $1 - 1/n^{10}$,

$$\begin{aligned} h(Q'_i) &\leq \frac{h(P'_i) \cdot |Q'_i|}{|P'_i|} + \delta \cdot (|P'_i| \cdot s / m) \\ \text{Or, } h(Q'_i) \cdot (m/s) &\leq \frac{h(P'_i) \cdot |Q'_i|}{|P'_i|} \cdot (m/s) + \epsilon |P'_i| \cdot \mu'^2 / 20 \\ \text{Or, } h(Q'_i) \cdot (m/s) &\leq (1 + \frac{\epsilon}{50}) \cdot h(P'_i) + \epsilon |P'_i| \cdot \mu'^2 / 20 \end{aligned}$$

The last inequality follows from Observation 3.15 considering both cases in the flow construction. Next we compute the additional costs. We have two cases. In the first case, $|Q'_i| \leq |P'_i| \cdot s / m$ and we need to route $(|P'_i| - |Q'_i| \cdot m/s)$ amount of flow from c' to c_i . The cost is at most,

$$(|P'_i| - |Q'_i| \cdot m/s) \cdot d(c', c_i) \leq \epsilon \cdot \frac{|P'_i|}{50} \cdot \left(\frac{2}{|P'_i|} \cdot \sum_{p \in P'_i} d(p, c_i)^2 + 8\mu'^2 \right)$$

$$\leq \frac{\epsilon}{25} \cdot \sum_{p \in P'_i} d(p, c_i)^2 + \frac{4\epsilon \cdot |P'_i|}{25} \mu'^2$$

The first inequality follows from Observation 3.15 and from the fact that c' is the ring center.

In the second case, $|Q'_i| > |P'_i| \cdot s/m$. Note that in this case we need to route at least $|P'_i| - |Q'_i| \cdot m/s$ flow from one point p to c_i , as $|Q''_i| = \lfloor |P'_i| \cdot s/m \rfloor$, and m/s flow for each point in $Q'_i \setminus Q''_i$ to w . The first cost is at most,

$$\begin{aligned} (|P'_i| - |Q''_i| \cdot m/s) \cdot d(p, c_i) &\leq \epsilon \cdot \frac{|P'_i|}{20} \cdot \left(\frac{2}{|P'_i|} \cdot \sum_{p \in P'_i} d(p, c_i)^2 + 8\mu'^2 \right) \\ &\leq \frac{\epsilon}{10} \cdot \sum_{p \in P'_i} d(p, c_i)^2 + \frac{2\epsilon \cdot |P'_i|}{5} \mu'^2 \end{aligned}$$

The first inequality follows from Observation 3.17. The second cost can be bounded by,

$$\begin{aligned} \sum_{p \in (Q'_i \setminus Q''_i)} (m/s) \cdot d(p, c')^2 &\leq (|Q'_i| - |Q''_i|) \cdot (m/s) \cdot \mu'^2 \\ &\leq (\epsilon |P'_i| \cdot s/(40m)) \cdot (m/s) \cdot \mu'^2 \\ &\leq \epsilon |P'_i| \mu'^2 / 40 \end{aligned}$$

The second inequality follows from Observation 3.16. Thus, in this case, the total cost is bounded by,

$$\left(1 + \frac{3\epsilon}{25}\right) \cdot \sum_{p \in P'_i} d(p, c_i)^2 + \frac{19\epsilon \cdot |P'_i|}{40} \mu'^2.$$

Case 2. $\mathbb{E}[|Q'_i|] < \epsilon^3 s/(100k)$. We give separate bounds for the two cases. In the first case, $|Q'_i| \leq |P'_i| \cdot s/m$. In this case, we need to route m/s flow from points in Q'_i to c_i and $(|P'_i| - |Q'_i| \cdot m/s)$ amount of flow from c' to c_i . Let p_{\min} and p_{\max} be the nearest and farthest points in P'_i from c_i . The total cost is,

$$\begin{aligned} \sum_{p \in Q'_i} (m/s) \cdot d(p, c_i)^2 + (|P'_i| - |Q'_i| \cdot m/s) \cdot d(c', c_i)^2 &\leq \sum_{p \in Q'_i} (m/s) \cdot d(p_{\max}, c_i)^2 + (|P'_i| - |Q'_i| \cdot m/s) \cdot d(c', c_i)^2 \\ &\leq |P'_i| \cdot \max\{d(p_{\max}, c_i)^2, d(c', c_i)^2\} \\ &\leq |P'_i| \cdot (\mu' + d(c', c_i))^2 \\ &\leq |P'_i| \cdot (2\mu' + d(p_{\min}, c_i))^2 \end{aligned}$$

The first inequality follows by replacing the squares of the distances by their maximum. The third inequality follows by noting that $d(p_{\max}, c_i) \leq d(p_{\max}, c') + d(c', c_i) \leq \mu' + d(c', c_i)$. The last inequality follows by noting that $d(c', c_i) \leq d(c', p_{\min}) + d(p_{\min}, c_i)$.

Next, we upper bound the above expression. We consider two subcases. The first one is $d(p_{\min}, c_i) \leq 2\mu'/\epsilon$. In this subcase,

$$\begin{aligned} |P'_i| \cdot (2\mu' + d(p_{\min}, c_i))^2 &\leq \frac{\epsilon^3 m}{100k} \cdot (2\mu' + 2\mu'/\epsilon)^2 \\ &\leq \frac{\epsilon^3 m}{25k} \cdot \mu'^2 (1 + 1/\epsilon)^2 \\ &= \frac{O(\epsilon m)}{k} \cdot \mu'^2. \end{aligned}$$

In the other subcase $d(p_{\min}, c_i) > 2\mu'/\epsilon$.

$$\begin{aligned} |P'_i| \cdot (2\mu' + d(p_{\min}, c_i))^2 &\leq |P'_i| \cdot (\epsilon d(p_{\min}, c_i) + d(p_{\min}, c_i))^2 \\ &\leq |P'_i| \cdot d(p_{\min}, c_i)^2 \cdot (1 + \epsilon)^2 \\ &= (1 + O(\epsilon)) \sum_{p \in P'_i} d(p, c_i)^2 \end{aligned}$$

The last inequality follows, as $d(p, c_i) \leq d(p_{\min}, c_i)$ for all $p \in P'_i$. Now, we consider the second case: $|Q'_i| > |P'_i| \cdot s/m$. We need to route the flow from points in Q''_i to c_i . Additionally, we need to route at least $|P'_i| - |Q'_i| \cdot m/s$ flow from one point p^* to c_i , as $|Q''_i| = \lfloor |P'_i| \cdot s/m \rfloor$, and m/s flow for each point in $Q'_i \setminus Q''_i$ to w . The sum of the first two costs is at most,

$$\begin{aligned} & \sum_{p \in Q''_i} (m/s) \cdot d(p, c_i)^2 + (|P'_i| - |Q'_i| \cdot m/s) \cdot d(p^*, c_i)^2 \\ & \leq \sum_{p \in Q''_i} (m/s) \cdot d(p_{\max}, c_i)^2 + (|P'_i| - |Q'_i| \cdot m/s) \cdot d(p_{\max}, c_i)^2 \\ & \leq |P'_i| \cdot d(p_{\max}, c_i)^2 \\ & = \frac{O(\epsilon m)}{k} \cdot \mu'^2 + (1 + O(\epsilon)) \sum_{p \in P'_i} d(p, c_i)^2 \end{aligned}$$

The last equality follows in the same way as in the first case. The remaining cost in the second case can be bounded by,

$$\begin{aligned} \sum_{p \in (Q'_i \setminus Q''_i)} (m/s) \cdot d(p, c_i)^2 & \leq (|Q'_i| - |Q''_i|) \cdot (m/s) \cdot \mu'^2 \\ & \leq |Q'_i| \cdot (m/s) \cdot \mu'^2 \\ & \leq (\epsilon s / (50k)) \cdot (m/s) \cdot \mu'^2 \\ & = (\epsilon m / (50k)) \cdot \mu'^2 \end{aligned}$$

Thus, the total cost in both the cases is bounded by,

$$\frac{O(\epsilon m)}{k} \cdot \mu'^2 + (1 + O(\epsilon)) \sum_{p \in P'_i} d(p, c_i)^2$$

General upper bound on the cost. By merging the cost in both cases, we obtain the common upper bound,

$$\begin{aligned} & (1 + O(\epsilon)) \sum_{p \in P'_i} d(p, c_i)^2 + \frac{19\epsilon \cdot |P'_i|}{40} \mu'^2 + \frac{O(\epsilon m)}{k} \cdot \mu'^2 \\ & = (1 + O(\epsilon)) \sum_{p \in P'_i} d(p, c_i)^2 + O(\epsilon \cdot |P'_i|) \cdot \mu'^2 + \frac{O(\epsilon m)}{k} \cdot \mu'^2 \end{aligned}$$

Summing over all the centers in C , we obtain,

$$\begin{aligned} \text{wcost}(W', C, M) & \leq (1 + O(\epsilon)) \cdot \text{cost}(P, C, M) + O(\epsilon \cdot |P'|) \cdot \mu'^2 + O(\epsilon) \cdot m \cdot \mu'^2 \\ & = (1 + O(\epsilon)) \cdot \text{cost}(P, C, M) + O(\epsilon) \cdot m \cdot \mu'^2. \end{aligned}$$

Summing the cost over all rings gives us,

$$\begin{aligned} |\text{cost}(P, C, M) - \text{wcost}(W', C, M)| & \leq \sum_{(i,j)} O(\epsilon) \cdot |P'_{i,j}| \cdot 2^j \mu^2 + O(\epsilon k \log n) \cdot \text{cost}(P, C, M) \\ & = O(\epsilon k \log n) \cdot \text{cost}(P, C, M) \end{aligned}$$

Note that the coreset size for each ring and for each color was $O(k \log n / \epsilon^5)$. To obtain the desired ϵ error, we need to scale ϵ by a factor of $\Theta(k \log n)$. Thus, the required size of the coreset becomes $O((k \log n)^6 / \epsilon^5)$. Summing over all rings and colors we obtain the desired bound of $O(\ell(k \log n)^7 / \epsilon^5)$ on our coreset size.

This proves the disjoint case of Theorem 1.1 for k -means. The coreset construction algorithm for k -means in the overlapping group case is again the same as that for k -median, except the bound on sample size. From the above analysis and the analysis for k -median, we obtain the desired result. This proves the overlapping case of Theorem 1.1 for k -means.

In the Euclidean case, the extension of the analysis for k -median to k -means is trivial. We obtain the following generic theorem.

Theorem 6.2. *Given a collection of ℓ possibly overlapping groups consisting of n points in total in a metric space, there is an $O(n(k + \ell))$ time randomized algorithm that w.p. at least $1 - 1/n$, computes a universal coreset for k -means clustering of size $O(\Gamma(k \log n)^7 / \epsilon^5)$. In the Euclidean case, the size of the coreset is $O\left(\frac{\Gamma}{\epsilon^5} \cdot k^7 (\log n)^6 (\log n + d \log(1/\epsilon))\right)$, and the running time is $O(nd(k + \ell))$.*

7. Assignment problem for (α, β) -fair clustering

Recall that we are given ℓ groups $\{P_i\}$ of P , and P^1, \dots, P^Γ are the point equivalence classes. Also, I_t is the set of indexes of the groups corresponding to P^t , for each $t \in [\Gamma]$. We aim to solve (α, β) -FAIR CLUSTERING on our coresot W instead of on the original points. Suppose we are given the optimal set of centers C for (α, β) -FAIR CLUSTERING. Let \mathcal{M} be the collection of coloring constraints that express the assignment restriction of (α, β) -FAIR CLUSTERING. Since W is a universal coresot, computing the minimum $\text{wcost}(W, M, C)$ over all $k \times \Gamma$ matrix $M \in \mathcal{M}$ would give us the optimal cost of fair clustering, modulo a $(1 \pm \epsilon)$ factor. Now, recall that, for k -median, $\text{wcost}(W, M, C)$ is the minimum value $\sum_{x \in P, c_j \in C} \psi(x, c_j) \cdot d(x, c_j)$ over all assignments $\psi : P \times C \rightarrow \mathbb{R}_{\geq 0}$ such that

1. For each $x \in P$, $\sum_{c_j \in C} \psi(x, c_j) = w(x)$.
2. For each $c_j \in C$ and class $1 \leq t \leq \Gamma$, $\sum_{x \in P^t} \psi(x, c_j) = M_{jt}$.

Thus, given an M , we can compute $\text{wcost}(W, M, C)$ by solving a minimum cost flow problem. But, as the size of \mathcal{M} can be sufficiently large, we cannot try out all possible M . Note that as the optimal $M \in \mathcal{M}$ represents a fair partition of the equivalence classes $\{P^t\}$ between the centers, ψ automatically satisfies the fairness properties:

$$\begin{aligned} \sum_{x \in P_i} \psi(x, c_j) &\leq \alpha_i \cdot \sum_{x \in P} \psi(x, c_j), \quad \forall c_j \in C, \forall i \in [\ell], \\ \sum_{x \in P_i} \psi(x, c_j) &\geq \beta_i \cdot \sum_{x \in P} \psi(x, c_j), \quad \forall c_j \in C, \forall i \in [\ell]. \end{aligned}$$

Now, as the optimal M has all integer entries, the optimal cost assignment ψ must also be integral. Here we assume that the coresot points have integer weights. We note that our construction can be slightly modified to obtain coresot with integer weights (e.g., see Chen's adaptation [38]). Thus, given W and C it is sufficient to compute a minimum cost integral assignment that satisfies the above two inequalities and the constraint: For each $x \in P$, $\sum_{c_j \in C} \psi(x, c_j) = w(x)$. We refer to this assignment problem as WEIGHTED FAIR ASSIGNMENT. Our main theorem of this section provides an algorithm with running time $(k\Gamma)^{O(k\Gamma)} |W|^{O(1)}$ for this problem. The general idea is to reduce the assignment problem to a linear programming problem. The unknown optimal assignment can be naturally expressed in terms of linear inequalities, along with the condition that the assignment is fair. However, the issue is that in general the optimal fractional solution to this linear programming problem is not integral, and the integrality gap could be arbitrarily large. Thus, an optimal fractional solution does not yield the desired assignment. And indeed, it was observed already by Bera et al. [3] that the assignment problem for (α, β) -FAIR CLUSTERING is NP-hard, so there is no hope to have a polynomial time assignment algorithm.

We cannot afford to make all variables integral and solve an integer linear program (ILP) instead, as the number of variables is large, of order $|W|k$, and in our construction $|W|$ is polylogarithmic in n . However, note that the optimal assignment has the property that for each $c_j \in C$ and class $1 \leq t \leq \Gamma$, $\sum_{x \in P^t} \psi(x, c_j) = M_{jt}$. Thus the amount of weight assigned from each class to each center is an integer. Using this observation, we reduce our problem to a mixed-integer linear programming problem and force only $k \cdot \Gamma$ variables to be integral. These variables correspond exactly to the entries of the constraint matrix M . Then, we show that this automatically ensures that all the other variables are integral as well, in the optimal solution.

Next, we state one of the equivalent formulations of the MIXED-INTEGER LINEAR PROGRAMMING problem. The input to the problem is a matrix $A \in \mathbb{R}^{m \times d}$, a vector $b \in \mathbb{R}^m$, a vector $c \in \mathbb{R}^d$, and a parameter p , $0 \leq p \leq d$. The goal is to find a vector $x = (x_1, \dots, x_d) \in \mathbb{R}^d$ such that $x_1, \dots, x_p \in \mathbb{Z}$, $A \cdot x \leq b$, and the value $c \cdot x$ is minimized across all vectors satisfying the above.

By the celebrated result of Lenstra [71], MIXED-INTEGER LINEAR PROGRAMMING is solvable in FPT time when parameterized by the number of integer variables p . We use the following commonly employed version of this result, following the improvements to the original Lenstra's algorithm given by Kannan [72], and Frank and Tardos [73].

Proposition 7.1 ([71], [72], [73]). *There is an algorithm solving MIXED-INTEGER LINEAR PROGRAMMING in time $O(p^{2.5p+o(p)} d^4 L)$ and space polynomial in L , where L is the bitsize of the given instance.*

Now we present the assignment algorithm itself. Note that it is sufficient to consider only the points in W for the purpose of computing an assignment, as the other points in P have zero weights. For simplicity, we denote $|W|$ by n . There is practically no difference between the cases of k -median and k -means concerning the assignment problem, and thus we state it for both cases.

Theorem 7.2. *There is an algorithm that given an instance of WEIGHTED FAIR ASSIGNMENT, i.e., a weighted set W of n points and a set $C = \{c_1, \dots, c_k\}$ of k centers, computes an optimal assignment of W with the set of centers C . That is, the output is a minimum cost assignment $\psi : P \times C \rightarrow \mathbb{Z}_{\geq 0}$ that corresponds to (α, β) -FAIR CLUSTERING. The running time of the algorithm is $(k\Gamma)^{O(k\Gamma)} n^{O(1)} L$, where L is the total number of bits in the encoding of distances and weights in the instance.*

Proof. We reduce WEIGHTED FAIR ASSIGNMENT to MIXED-INTEGER LINEAR PROGRAMMING. The formulation of our problem itself follows the natural way of treating a clustering assignment problem as a flow problem. Let $W = \{(p_1, w(p_1)), \dots, (p_n, w(p_n))\}$. For every point p_i and center c_j introduce a variable f_{ij} corresponding to how much weight from the i -th point is assigned to the j -th center. Also, for every center c_j and point equivalence class $t \in \{1, \dots, \Gamma\}$ introduce a variable g_{tj} , corresponding to how much weight from points of the class t the j -th center gets. The following constraints express that $\{f_{ij}\}$ and $\{g_{tj}\}$ define a fair clustering:

$$f_{ij} \geq 0 \quad \forall i \in \{1, \dots, n\}, j \in \{1, \dots, k\}, \quad (1)$$

$$g_{tj} \in \mathbb{Z}_{\geq 0} \quad \forall j \in \{1, \dots, k\}, t \in \{1, \dots, \Gamma\}, \quad (2)$$

$$\sum_{1 \leq j \leq k} f_{ij} = w(p_i) \quad \forall i \in \{1, \dots, n\}, \quad (3)$$

$$\sum_{i \in [n]: p_i \in P^t} f_{ij} = g_{tj} \quad \forall j \in \{1, \dots, k\}, t \in \{1, \dots, \Gamma\}, \quad (4)$$

$$\sum_{i \in [n]: p_i \in P_q} f_{ij} \geq \beta_q \sum_{i \in [n]} f_{ij} \quad \forall j \in \{1, \dots, k\}, \forall q \in \{1, \dots, \ell\}, \quad (5)$$

$$\sum_{i \in [n]: p_i \in P_q} f_{ij} \leq \alpha_q \sum_{i \in [n]} f_{ij} \quad \forall j \in \{1, \dots, k\}, \forall q \in \{1, \dots, \ell\}. \quad (6)$$

Note that for a color $q \in \{1, \dots, \ell\}$, $\sum_{i \in [n]: p_i \in P_q} f_{ij}$ is precisely the weight assigned from points of color q to the center j , and $\sum_{i \in [n]} f_{ij}$ is the total weight assigned to the center j . Thus Constraints (5) and (6) ensure that the assignment is indeed fair. Finally, the objective function is

$$\text{Minimize} \quad \sum_{i=1}^n \sum_{j=1}^k d_{ij} f_{ij}, \quad (7)$$

where $d_{ij} = d(p_i, c_j)$ in the case of k -median, and $d_{ij} = d(p_i, c_j)^2$ in the case of k -means.

We solve the MIXED-INTEGER LINEAR PROGRAMMING defined above by using Proposition 7.1. We require that the variables $\{g_{tj}\}$ take integral values, while we do not impose this restriction on the variables $\{f_{ij}\}$. Thus, in time $(k\Gamma)^{O(k\Gamma)} n^{O(1)} L$ we find the optimal solution $\{f_{ij}\}, \{g_{tj}\}$.

Clearly, Constraints (1)–(6) ensure that the assignment defined by $\{f_{ij}\}$ corresponds to WEIGHTED FAIR ASSIGNMENT, except for the fact that some of $\{f_{ij}\}$ might be fractional. We now show that the integrality of $\{g_{tj}\}$ guarantees that there exists an optimal solution to (1)–(7) that is integral. For every equivalence class $t \in \{1, \dots, \Gamma\}$ consider the following flow network. The network is essentially a restriction of (1)–(4) to the class t assuming that the values $\{g_{tj}\}$ are fixed. There is a node associated with every point $p_i \in P^t$ that has a supply of $w(p_i)$, and there is a node associated with every $c_j \in C$ that has a demand of g_{tj} . There is an edge e_{ij} between each point $p_i \in P^t$ and every center $c_j \in C$ that has an unlimited capacity and the cost d_{ij} . In this network, there is a maximum flow of minimal cost $\{f'_{ij}\}$ that has only integral values, since all the supplies, demands and capacities in the network are integers. Now we replace the respective values of $\{f_{ij}\}$ with the obtained $\{f'_{ij}\}$ that are integral and still satisfy (1)–(4). The cost is unchanged since $\{f_{ij}\}$ induces a maximum flow in the network as well. Thus the old cost can only be larger, but also $\{f_{ij}\}$ is an optimal solution to the MIXED-INTEGER LINEAR PROGRAMMING instance, so the new cost cannot be smaller. After we perform the above for every class, the whole assignment is integral, now satisfying the statement of the theorem completely. Finally, note that $\{f'_{ij}\}$ can be found in $n^{O(1)} L$ time with the known values of $\{g_{tj}\}$ by any polynomial time minimum-cost flow algorithm. \square

The algorithm in Theorem 7.2 allows us to solve (α, β) -FAIR CLUSTERING on the original points as well, as long as we know a suitable set of k centers. However, the running time would have a heavy dependence on n , roughly n^5 . So to obtain a near-linear time algorithm, we cannot use Theorem 7.2 directly on the original points, even if we know the centers. Instead, in the approximation algorithms we present, we first compute a universal coreset of the original set of points, and then solve all the arising instances of the assignment problem on the coreset, thus inflicting only polylogarithmic in n time. Still, at the end we have to output a low-cost fair assignment of the original points, and again we cannot directly use Theorem 7.2. So we show how to compute the assignment in near-linear time with the help of the coreset. The idea is to run Theorem 7.2 on the coreset and then use the optimal solution there to find a good assignment of the original points in a simpler way. Namely, knowing how many points from each equivalence class are assigned to each center, the assignment problem boils down to finding a minimum-cost flow in a bipartite network where one of the parts is small. First, we recall a suitable minimum-cost flow result by Ahuja et al. [74].

Proposition 7.3 (Theorem 7.3 in [74]). *The minimum-cost flow problem on a bipartite network is solvable in time $O((n_1m + n_1^3)\log(n_1D))$, where n_1 is the size of the smaller part in the network, m is the number of edges, and D is the maximum cost of an edge in the network.*

Now we prove a general lemma that allows us to transfer any fair assignment from the coreset to a fair assignment on the original points in polynomial time, while losing only a factor of $(1 + \epsilon)$ in the cost.

Lemma 7.4. *There is an algorithm that given a set of points P with the ℓ groups P_1, \dots, P_ℓ , a coreset W of P , a set of k centers C , a fair assignment $\psi : P \times C \rightarrow \mathbb{Z}_{\geq 0}$, and a value $0 < \epsilon \leq 1$, computes a fair assignment of the points of P to the centers of C with the cost at most $(1 + \epsilon) \cdot \text{cost}(\psi)$ in time $O(\Gamma \cdot k^3 / \epsilon^{O(1)} \cdot n \log n)$. This holds for both (α, β) -Fair k -median and (α, β) -Fair k -means in general metric, provided that W satisfies*

$$\text{cost}(P^t, M, C) \leq (1 + \epsilon/3) \text{wcost}(W^t, M, C),$$

for every column constraint matrix $M \in \mathbb{Z}^k$, where by W^t we denote the restriction of W to the points of the equivalence class P^t . In the Euclidean case, the running time is multiplied by d .

Proof. For the assignment ψ , consider the values $\{g_{tj}\}_{t \in [\Gamma], j \in [k]}$, using the notation in Theorem 7.2, where g_{tj} denotes how many points from the t -th class are assigned to the j -th center by ψ , and the values $\{A_t\}_{t \in [\Gamma]}$, where A_t is the cost of ψ restricted to the t -th class. Now for each class P^t in the original point set P , solve the following assignment problem: assign points of P^t to centers in C such that there are exactly g_{tj} points assigned to the j -th center, and the cost of the assignment is minimum among all such assignments. We naturally view this problem as a minimum-cost flow problem, and we solve it by running the algorithm given by Proposition 7.3. Note that the resulting network has $O(k)$ vertices in the part corresponding to the centers C , and $O(nk)$ edges in total. Finally, the resulting fair assignment φ from P to C is the union of assignments from P^t to C for all $t \in [\Gamma]$. Clearly, the obtained assignment is fair, since the fairness condition is completely determined by the numbers $\{g_{tj}\}$. This is true, as in the Constraints (5) and (6), $\sum_{i \in [n]: p_i \in P_q} f_{ij}$ can be expressed by $\sum_{t \in [\Gamma]: q \in I_t} g_{tj}$ and $\sum_{i \in [n]} f_{ij}$ by $\sum_{t=1}^{\Gamma} g_{tj}$. We now argue about the cost. By construction, the cost of the resulting assignment is $\sum_{t=1}^{\Gamma} \text{cost}(P^t, (g_{tj})_{j=1}^k, C)$. Now,

$$\text{cost}(P^t, (g_{tj})_{j=1}^k, C) \leq (1 + \epsilon/3) \text{wcost}(W^t, (g_{tj})_{j=1}^k, C) \leq (1 + \epsilon/3) A_t, \quad \forall t \in [\Gamma].$$

Summing over all $t \in [\Gamma]$, we obtain

$$\text{cost}(\varphi) = \sum_{t=1}^{\Gamma} \text{cost}(P^t, (g_{tj})_{j=1}^k, C) \leq (1 + \epsilon/3) \sum_{t=1}^{\Gamma} A_t = (1 + \epsilon/3) \text{cost}(\psi).$$

By Proposition 7.3 it takes time $O(\Gamma \cdot k^3 / \epsilon^2 \cdot n \log n)$ to run the minimum-cost flow algorithm Γ times, where we assume that $\log D = O(\log n / \epsilon^2)$. Finally, we justify the latter by a standard argument reducing the ratio of maximum distance in the instance to the minimum distance. In the network flow instance that we construct from P^t and C , tweak slightly the costs on the edges. Set $\epsilon_0 = \epsilon/6$, if an edge costs more than $D_{\max} := 2A_t$, replace its cost by D_{\max} , and if an edge costs less than $D_{\min} := \epsilon_0 A_t / (2n)$, replace its cost by D_{\min} . For all the other edges, round up their cost to the closest value of the form $(1 + \epsilon_0)^q D_{\min}$, where q is an integer. In the modified network, the cost scaling part then induces a factor of $\log_{1+\epsilon_0}(D_{\max}/D_{\min}) = \log_{1+\epsilon_0}(4n/\epsilon_0) = O(\log n / \epsilon^2)$, instead of $O(\log D)$.

Now we argue about how this change influences the cost. Consider an optimal assignment $\varphi : P^t \rightarrow C$ in the modified network, obtained by the network flow algorithm. Its cost is at most $(1 + 2\epsilon_0) \cdot (1 + \epsilon/3) A_t$, since the cost of an optimal assignment φ^* in the original network is at most $(1 + \epsilon/3) A_t$ by the argument above, and the cost of φ^* in the new network is at most $(1 + 2\epsilon_0)$ times the cost in the original network. The latter holds since φ^* never uses edges of cost more than $D_{\max} = 2A_t$, for the edges between D_{\min} and D_{\max} the cost increase is at most a factor of $(1 + \epsilon_0)$, and for the edges with the cost less than D_{\min} , their total contribution in the new network is at most $n \cdot \epsilon_0 A_t / n = \epsilon_0 A_t$. The algorithm outputs the optimal assignment φ in the modified network, and its cost in the original network is at most its cost in the modified network, since edges with cost at least D_{\max} are never used, and the cost of all the other edges is less in the original network. Thus, we have shown that the cost of the assignment we constructed is at most $(1 + 2\epsilon_0)(1 + \epsilon/3) A_t = (1 + \epsilon/3)(1 + \epsilon/3) A_t \leq (1 + \epsilon) A_t$. From this point, the cost analysis above proceeds, and summing over all $t \in [\Gamma]$ we obtain $\text{cost}(\varphi) \leq (1 + \epsilon) \text{cost}(\psi)$.

Observe that in the Euclidean case we compute distances between the points from their respective d -dimensional vectors, thus taking an extra factor of d in the running time. \square

Note that the condition on W in Lemma 7.4 is satisfied by the coresets obtained from Theorem 4.1 and Theorem 6.2 with a suitable error parameter, since the coreset construction samples points in each equivalence class independently, and thus approximately preserves the cost with respect to any column matrix constraint on each of them. Now we show that

any instance of the assignment problem can be approximately solved in near-linear FPT time with the help of our cores et construction, Theorem 7.2, and Lemma 7.4.

Lemma 7.5. *Given a set of points P with the ℓ groups P_1, \dots, P_ℓ , a set of k centers C , and a parameter $0 < \epsilon \leq 1$, a fair assignment of the points of P to the centers of C with the cost at most $(1 + \epsilon) \text{faircost}(P, C)$ can be computed in time $(k\Gamma)^{O(k\Gamma)} (\log n / \epsilon)^{O(1)} + O(\Gamma \cdot k^3 / \epsilon^2 \cdot n \log n)$ w.h.p. This holds for both (α, β) -Fair k -median and (α, β) -Fair k -means in general metric. In the Euclidean case, the running time is multiplied by d .*

Proof. The algorithm proceeds as follows. First, we compute a cores et W from the point set P using Theorem 4.1 or Theorem 6.2, depending on the problem, with the error parameter ϵ_0 to be defined later. Then we compute an optimal fair assignment ψ from W to the centers C by applying Theorem 7.2. Finally, we invoke Lemma 7.4 on the assignment ψ to obtain a fair assignment $\varphi : P \rightarrow C$ with the cost at most $(1 + 3\epsilon_0) \text{cost}(\psi)$. The algorithm returns φ , and in what follows we bound the cost of this assignment. Denote by $M \in \mathbb{Z}^{k \times \ell}$ the constraint matrix corresponding to the assignment ψ , i.e. M_{ij} is equal to how many points from the j -th group ψ sends to the i -th center, and by M^* the constraint matrix corresponding to an optimal fair assignment from P to C . By the choice of M and M^* , and the fact that W is a universal cores et of P , we obtain

$$\begin{aligned} \text{cost}(\psi) &= \text{wcost}(W, M, C) \leq \text{wcost}(W, M^*, C) \leq (1 + \epsilon_0) \text{cost}(P, M^*, C) \\ &= (1 + \epsilon_0) \text{faircost}(P, C) \end{aligned}$$

Thus, $\text{cost}(\varphi)$ is at most $(1 + 3\epsilon_0)(1 + \epsilon_0) \text{faircost}(P, C)$, and setting ϵ_0 such that $(1 + 3\epsilon_0)(1 + \epsilon_0) \leq (1 + \epsilon)$ finishes the proof.

As for the running time, the $O((k + \ell) \cdot n)$ is for the cores et construction, $O((k\Gamma)^{O(k\Gamma)} (k \log n / \epsilon)^{O(1)})$ is for solving the assignment problem on the cores et, and $O(\Gamma \cdot k^3 / \epsilon^{O(1)} \cdot n \log n)$ is for restoring φ by Lemma 7.4. Note that the cores et construction time is dominated by the last term.

Finally, in the Euclidean case we compute distances between the points from their respective d -dimensional vectors, thus taking an extra factor of d in the running time. Note that we still use the general metric case in Theorems 4.1 and 6.2 for cores et construction, since we only need to preserve the objective with respect to the given set of centers C . \square

8. $(1 + \epsilon)$ -approximation in \mathbb{R}^d

In this section, we present a near-linear time $(1 + \epsilon)$ -approximation algorithm for Euclidean (α, β) -Fair k -median and (α, β) -Fair k -means. For that purpose, we combine our cores et construction (Theorem 4.1 and Theorem 6.2), our assignment algorithm (Theorem 7.2), and the linear-time constrained clustering algorithm of Bhattacharya et al. [37].

We denote the cost of clustering C_1, \dots, C_k with the centers $C = (c_1, \dots, c_k)$ by $\text{cost}_C(C_1, \dots, C_k)$. By $\text{cost}(C_1, \dots, C_k)$ we denote

$$\min_C \text{cost}_C(C_1, \dots, C_k)$$

where the minimum is over all possible k centers C . It is well-known that in the case of k -means the optimal center for a cluster C_i is its mean $\mu(C_i) := 1/|C_i| \sum_{x \in C_i} x$, thus $\text{cost}(C_1, \dots, C_k) = \text{cost}_{(\mu(C_1), \dots, \mu(C_k))}(C_1, \dots, C_k)$.

Next, we formally restate the result of Bhattacharya et al.

Proposition 8.1 ([37], Theorem 1). *Given a set of n points $P \subset \mathbb{R}^d$, parameters k and ϵ , there is a randomized algorithm that outputs a list \mathcal{L} of $2^{\tilde{O}(k/\epsilon)}$ sets of centers of size k such that for any clustering $\{C_1^*, \dots, C_k^*\}$ of P , the following event happens with probability at least $1/2$: there is a set $C \in \mathcal{L}$ such that*

$$\text{cost}_C(C_1^*, \dots, C_k^*) \leq (1 + \epsilon) \text{cost}(C_1^*, \dots, C_k^*),$$

where cost is with respect to the k -means clustering objective. The running time of the algorithm is $nd \cdot 2^{\tilde{O}(k/\epsilon)}$, where \tilde{O} notation hides a $O(\log \frac{k}{\epsilon})$ factor. The same statement holds for k -median, except the size of the list \mathcal{L} becomes $2^{\tilde{O}(k/\epsilon^{O(1)})}$, and the running time becomes $nd \cdot 2^{\tilde{O}(k/\epsilon^{O(1)})}$.

Note that Proposition 8.1 together with our assignment algorithm from Theorem 7.2 already implies $(1 + \epsilon)$ -approximation algorithm, as stated in the next claim.

Claim 8.2. *There exists a $(1 + \epsilon)$ -approximation algorithm solving (α, β) -FAIR CLUSTERING in \mathbb{R}^d in time $2^{\tilde{O}(k/\epsilon^{O(1)})} (k\Gamma)^{O(k\Gamma)} n^{O(1)} d$ with high probability. The algorithm also extends to the weighted version of the problem.*

Proof. The proof is by solving the assignment problem with the help of Theorem 7.2 on each set of centers in the list returned by Proposition 8.1. We run Proposition 8.1 $\Theta(\log n)$ times to succeed with high probability, and thus run Theorem 7.2 on $2^{\tilde{O}(k/\epsilon^{O(1)})} \log n$ candidate sets of centers.

For the weighted version, observe that the algorithm of Proposition 8.1 trivially extends to the case where the input points have weight, since the only step where all the input points are used is to perform D^2 -sampling, and there the sampling probabilities just need to be multiplied by the respective weights. Theorem 7.2 holds in the weighted case by definition. \square

However, the running time of Claim 8.2 has a high-degree polynomial dependency on n . To achieve a near-linear time algorithm, we use the help of our coresets construction. Observe that Proposition 8.1 together with Lemma 7.5 already imply an algorithm of this form. Nevertheless, we proceed with a variation of this scheme that leads to a slightly better running time, in particular, avoiding a $n \log^2 n$ factor.

The general idea of our algorithm is as follows. First, we obtain a list of candidate sets of centers by Proposition 8.1. Then we compute a universal coreset from the input points such that the objective is preserved with respect to all the computed sets of centers. For each set of k centers in the list we run our assignment algorithm on the coreset to determine the set of centers with the best cost. The algorithm of Bhattacharya et al. and the coreset computation take linear time, and the assignment problem is solved on the coreset, thus taking time polylogarithmic in n . Finally, we run Lemma 7.4 on the best set of centers to construct a fair assignment on the original points. We state and prove the theorem formally next.

Theorem 8.3. *There is a randomized algorithm that given an instance P of (α, β) -FAIR CLUSTERING and a parameter $0 < \epsilon \leq 1$ outputs a set of k centers C and a fair assignment $\varphi : P \rightarrow C$ satisfying $\text{cost}(\varphi) \leq (1 + \epsilon) \text{faircost}(P)$ with high probability. The running time of the algorithm is*

$$2^{\tilde{O}(k/\epsilon^{O(1)})} (k\Gamma)^{O(k\Gamma)} n d \log n.$$

Proof. First, we run the algorithm given by Proposition 8.1 to obtain a list \mathcal{L} of $2^{\tilde{O}(k/\epsilon_0^{O(1)})}$ candidate sets of centers, using the error parameter $\epsilon_0 < \epsilon$ to be defined later. To increase the probability of success, we repeat this $\Theta(\log n)$ times concatenating all the obtained lists, to form a list \mathcal{L} of $2^{\tilde{O}(k/\epsilon_0^{O(1)})} \log n$ candidate sets of centers.

For (α, β) -Fair k -median, we then compute a universal coreset W of size

$$O\left(\frac{\Gamma}{\epsilon_0} k^2 (\log(n + k 2^{\tilde{O}(k/\epsilon_0^{O(1)})} \log n))^2\right) = \Gamma(k/\epsilon_0 \log n)^{O(1)},$$

using Theorem 4.1, again with the error parameter ϵ_0 . We use the general metric case of the theorem with respect to the points P and the possible centers contained in the list \mathcal{L} . For (α, β) -Fair k -means, we employ instead Theorem 6.2 to obtain a universal coreset W , its size is also $\Gamma(k/\epsilon_0 \log n)^{O(1)}$. For the rest of the proof, there is no difference between the two problems.

For each set of k centers in \mathcal{L} we run the assignment algorithm given by Theorem 7.2, and select the set of centers C with the best cost. Now we bound $\text{faircost}(P, C)$. Denote by M the color constraint matrix that corresponds to an optimal fair assignment from W to C , it holds that

$$\text{faircost}(P, C) \leq \text{cost}(P, M, C) \leq \frac{1}{1 - \epsilon_0} \text{wcost}(W, M, C),$$

where the last inequality is by the definition of a universal coreset. By Proposition 8.1 with probability $1 - (1/2)^{\Theta(\log n)} = 1 - (1/n)^{\Theta(1)}$ there is a set \tilde{C} in \mathcal{L} such that $\text{faircost}(P, \tilde{C}) \leq (1 + \epsilon_0) \text{faircost}(P)$. Denote by \tilde{M} the color constraint matrix that corresponds to an optimal fair assignment from W to \tilde{C} , since C achieves the lowest cost of fair clustering for W among \mathcal{L} , $\text{wcost}(W, M, C) \leq \text{wcost}(W, \tilde{M}, \tilde{C})$. Denote by \tilde{M}^* the constraint matrix achieving $\text{faircost}(P, \tilde{C}) = \text{cost}(P, \tilde{M}^*, \tilde{C})$, by the choice of \tilde{M} we have that

$$\text{wcost}(W, \tilde{M}, \tilde{C}) \leq \text{wcost}(W, \tilde{M}^*, \tilde{C}) \leq (1 + \epsilon_0) \text{cost}(P, \tilde{M}^*, \tilde{C}),$$

where the last inequality is because W is a universal coreset of P . And since $\text{cost}(P, \tilde{M}^*, \tilde{C}) = \text{faircost}(P, \tilde{C}) \leq (1 + \epsilon_0) \text{faircost}(P)$ by the choice of \tilde{M}^* and \tilde{C} , we have the following bound:

$$\text{faircost}(P, C) \leq \frac{1 + \epsilon_0}{1 - \epsilon_0} \text{cost}(P, \tilde{M}^*, \tilde{C}) \leq \frac{(1 + \epsilon_0)^2}{1 - \epsilon_0} \text{faircost}(P).$$

Finally, we compute a fair assignment from P to C running the algorithm from Lemma 7.5, using the error parameter ϵ_0 . The computed assignment has cost at most $(1 + \epsilon_0) \text{faircost}(P, C)$, which by the above is at most $\frac{(1 + \epsilon_0)^3}{1 - \epsilon_0} \text{faircost}(P)$. Setting ϵ_0 such that $1 + \epsilon \geq \frac{(1 + \epsilon_0)^3}{1 - \epsilon_0}$ concludes the proof.

The running time of the algorithm is the sum of the $2^{\tilde{O}(k/\epsilon^{O(1)})}nd$ running time of Proposition 8.1 multiplied by $O(\log n)$, the $O((k+l)nd)$ running time given by Theorem 4.1/Theorem 6.2, $2^{\tilde{O}(k/\epsilon^{O(1)})}$ times the $(k\Gamma)^{O(k\Gamma)}(\log n)^{O(1)}d$ running time of the assignment algorithm given by Theorem 7.2 on the coreset, and finally the running time of Lemma 7.4. All of these terms are dominated by $2^{\tilde{O}(k/\epsilon^{O(1)})}(k\Gamma)^{O(k\Gamma)}nd \log n$. \square

Note that the algorithm in Theorem 8.3 is in a sense a non-typical use of a coreset: we first do the heavy part of running Proposition 8.1 on the original points, and only then use the coreset to speed up the assignment problem. We can also devise a true reductive algorithm, where we first construct a universal coreset from the input data, and then do everything on the coreset, both Proposition 8.1 and selection of the best centers. We show this algorithm in the next subsection.

8.1. Reduction to a small-sized instance

In fact, we show a general reduction result: that the original instance of (α, β) -FAIR CLUSTERING could be replaced by a small-sized one, such that any approximate solution could be lifted from the reduced instance with an extra error factor of $(1 + \epsilon)$. Moreover, it can be done in polynomial time that is near-linear in n and linear in d . Essentially, this result is a combination of the universal coreset property and Lemma 7.4, and it shows that our universal coreset construction can indeed be used for data compression wrt. (α, β) -FAIR CLUSTERING. In the next theorem, we state and prove the result formally.

Theorem 8.4. *There is a randomized algorithm that given an instance P of (α, β) -FAIR CLUSTERING in \mathbb{R}^d outputs a reduced weighted instance W of size $d(k/\epsilon \log n)^{O(1)}$ in the same space. W.h.p. it holds that for any $\gamma \geq 1$, and for any set of k centers C in \mathbb{R}^d and a fair assignment ψ from W to C such that $\text{cost}(\psi) \leq \gamma \text{faircost}(W)$, there exists a fair assignment $\varphi : P \rightarrow C$ with the cost at most $(1 + \epsilon)\gamma \text{faircost}(P)$ that can be restored from ψ and C . Both constructing W from P and restoring φ from ψ and C take time $O(\Gamma k^3/\epsilon^2 nd \log n)$.*

Proof. The algorithm to construct W from P is simply the algorithm from Theorem 4.1 constructing a universal coreset for (α, β) -Fair k -median in the Euclidean case (Theorem 6.2 for (α, β) -Fair k -median). We invoke the coreset construction algorithm with the error parameter $\epsilon_0 < \epsilon$ to be defined later, the $O((k+l)nd)$ running time is dominated by $O(\Gamma k^3/\epsilon^2 nd \log n)$. The reduced weighted instance W is exactly the obtained coreset. Its size is $d(k/\epsilon \log n)^{O(1)}$, and w.h.p. for any set C of k centers in \mathbb{R}^d and any constraint matrix $M \in \mathbb{Z}^{k \times \Gamma}$ it holds that

$$(1 - \epsilon) \cdot \text{cost}(P, M, C) \leq \text{wcost}(W, M, C) \leq (1 + \epsilon) \cdot \text{cost}(P, M, C).$$

Now consider a particular $\gamma > 1$, a set of k centers C and a fair assignment $\psi : P \times C \rightarrow \mathbb{Z}_{\geq 0}$ of the coreset W such that $\text{cost}(\psi) \leq \gamma \text{faircost}(W)$. Observe that $\text{faircost}(W) \leq (1 + \epsilon_0) \text{faircost}(P)$ since for the set of centers C^* and the constraint matrix M^* achieving $\text{faircost}(P) = \text{cost}(P, M^*, C^*)$, it holds that

$$\text{faircost}(W) \leq \text{wcost}(W, M^*, C^*) \leq (1 + \epsilon_0) \text{cost}(P, M^*, C^*) = (1 + \epsilon_0) \text{faircost}(P).$$

Thus, $\text{cost}(\psi) \leq (1 + \epsilon_0)\gamma \text{faircost}(P)$. To construct the fair assignment φ , we invoke Lemma 7.4 on the assignment ψ . By Lemma 7.4, the cost of φ is at most

$$(1 + 3\epsilon_0) \text{cost}(\psi) \leq (1 + 3\epsilon_0)(1 + \epsilon_0)\gamma \text{faircost}(P).$$

Finally, we set ϵ_0 such that $(1 + 3\epsilon_0)(1 + \epsilon_0) \leq (1 + \epsilon)$ to obtain the desired bound.

The running time of Lemma 7.4 is exactly $O(\Gamma k^3/\epsilon^2 nd \log n)$, and this dominates the $O((k+l)nd)$ running time required by the coreset construction. \square

Theorem 8.4 allows for any exact or approximate algorithm for (α, β) -FAIR CLUSTERING to be run on the small-sized coreset instead of the original points. By plugging in the $(1 + \epsilon)$ -approximation algorithm given by Claim 8.2, we obtain the following theorem.

Theorem 8.5. *There is an algorithm solving (α, β) -FAIR CLUSTERING in time*

$$O(\Gamma k^3/\epsilon^2 nd \log n) + 2^{\tilde{O}(k/\epsilon^{O(1)})}(k\Gamma)^{O(k\Gamma)}(d \log n)^{O(1)}$$

with high probability, for any given $0 < \epsilon \leq 1$.

Proof. Set $\epsilon_0 = \epsilon/3$. Invoke Theorem 8.4 with the error parameter ϵ_0 to obtain a reduced weighted instance W from the input points P . Run the algorithm from Claim 8.2 on W to obtain the set of k centers C and a fair assignment ψ from W to C of cost at most $(1 + \epsilon_0) \text{faircost}(W)$. Finally, by the second part of Theorem 8.4 compute a fair assignment $\varphi : P \rightarrow C$. The assignment φ is the output of the algorithm, and its cost is at most $(1 + \epsilon_0)^2 \text{faircost}(P) \leq (1 + \epsilon) \text{faircost}(P)$.

Both algorithms from Theorem 8.4 run in time $O(\Gamma k^3/\epsilon^2 nd \log n)$, and running the algorithm from Claim 8.2 on the input of size $n' := d(k/\epsilon \log n)^{O(1)}$ amounts to the time complexity of

$$2^{\tilde{O}(k/\epsilon^{O(1)})} (k\Gamma)^{O(k\Gamma)} (n')^{O(1)} d = 2^{\tilde{O}(k/\epsilon^{O(1)})} (k\Gamma)^{O(k\Gamma)} (d \log n)^{O(1)}. \quad \square$$

In the running time of Theorem 8.5, observe that the exponential term is just polylogarithmic in n , compared to Theorem 8.3. However, the dependency on d in Theorem 8.5 is a high-degree polynomial. That is since we invoke the Euclidean case of Theorem 4.1/Theorem 6.2 so that coresot preserves the objective with respect to every k points in \mathbb{R}^d , and that requires an additional factor of d in the coresot size.

8.2. Dimensionality reduction

In the case of k -means, we show how to apply the recent dimensionality reduction tools to effectively replace the dimension d by $O(k/\epsilon)$, thus making the algorithm from Theorem 8.5 linear in d too, and independent of d after the computation of the coresot. At the end, dimensionality reduction and our coresot construction effectively compress the instance to just $(k/\epsilon \log n)^{O(1)}$ real numbers, providing a stronger variant of Theorem 8.4.

In what follows, we employ the results and notation of Cohen et al. [39]. First, we define a *projection-cost preserving sketch*.

Definition 8.6 (Definition 2 in [39]). $\tilde{A} \in \mathbb{R}^{n \times d'}$ is a rank k projection-cost preserving sketch of $A \in \mathbb{R}^{n \times d}$ with one-sided error $0 \leq \epsilon < 1$ if, for all rank k orthogonal projection matrices $M \in \mathbb{R}^{n \times n}$,

$$\|A - MA\|_F^2 \leq \|\tilde{A} - M\tilde{A}\|_F^2 + c \leq (1 + \epsilon)\|A - MA\|_F^2,$$

for some fixed non-negative constant c that may depend on A and \tilde{A} but is independent of M .

We will employ a dimensionality reduction scheme based on approximate singular value decomposition. Note that any other projection-cost preserving sketch can be used too, with the appropriate change in running time and dimension.

Proposition 8.7 (Theorem 8 in [39]). Let $m = \lceil k/\epsilon \rceil$. For any $A \in \mathbb{R}^{n \times d}$ and any orthonormal matrix $Z \in \mathbb{R}^{d \times m}$ satisfying $\|A - AZZ^T\|_F^2 \leq (1 + \epsilon')\|A - A_m\|_F^2$, the sketch $\tilde{A} = AZ$ satisfies the conditions of Definition 8.6 with error $(\epsilon + \epsilon')$. Here A_m is A projected onto its top m singular vectors.

There is a long line of work providing algorithms to compute this sort of relative approximation to the SVD, we use the algorithm Boutsidis et al. [75], stating the version appearing in [76].

Proposition 8.8 (Lemma 4 in [76]). Given $A \in \mathbb{R}^{n \times d}$ of rank ρ , a target rank $2 \leq m < \rho$, and $0 < \epsilon < 1$, there exists a randomized algorithm that computes an orthonormal matrix $Z \in \mathbb{R}^{d \times m}$ such that

$$\mathbb{E}\|A - AZZ^T\|_F^2 \leq (1 + \epsilon)\|A - A_m\|_F^2.$$

The proposed algorithm runs in time $O(ndm/\epsilon)$.

Now we employ these results to strengthen Theorem 8.4 in the case of k -means.

Theorem 8.9. There is a randomized algorithm that given an instance P of (α, β) -FAIR CLUSTERING in \mathbb{R}^d outputs a reduced weighted instance W of size $(k/\epsilon \log n)^{O(1)}$ in a low-dimensional space \mathbb{R}^m , where $m = O(k/\epsilon)$. W.h.p. for any $\gamma \geq 1$, and for any set of k centers \tilde{C} in \mathbb{R}^m and a fair assignment ψ from W to \tilde{C} such that $\text{cost}(\psi) \leq \gamma \text{faircost}(W)$, there exists a set of k centers C in \mathbb{R}^d and a fair assignment $\varphi : P \rightarrow C$ with the cost at most $(1 + \epsilon)\gamma \text{faircost}(P)$ that can be restored from ψ and \tilde{C} . Both constructing W from P and restoring (C, φ) from (\tilde{C}, ψ) take time $O(\Gamma k^3/\epsilon^2 nd \log n)$.

Proof. Fix a value $0 < \epsilon_0 < \epsilon$ to be defined later. Represent the given points P as a matrix $A \in \mathbb{R}^{n \times d}$, where each row corresponds to a point. Set $m = \lceil k/\epsilon_0 \rceil$ and run the algorithm from Proposition 8.8 on the matrix A and the parameter m to obtain a matrix $Z \in \mathbb{R}^{d \times m}$. By Markov inequality, it holds with probability at least $1 - \frac{1+\epsilon_0}{1+2\epsilon_0} = \Omega(\epsilon_0)$ that

$$\|A - AZZ^T\|_F^2 \leq (1 + 2\epsilon_0)\|A - A_m\|_F^2.$$

By invoking Proposition 8.8 $O(\epsilon_0^{-1} \log n)$ times and picking Z with the smallest value of $\|A - AZZ^T\|_F^2$, we achieve that the bound above holds with high probability. Then, by Proposition 8.7, the sketch $\tilde{A} = AZ$ is a projection-cost preserving sketch, i.e. it holds that

$$\|A - MA\|_F^2 \leq \|\tilde{A} - M\tilde{A}\|_F^2 + c \leq (1 + 3\epsilon_0)\|A - MA\|_F^2,$$

for any rank k orthogonal projection matrix $M \in \mathbb{R}^{n \times n}$ and some constant c independent of M . Consider the corresponding to \tilde{A} set of points \tilde{P} in \mathbb{R}^m . It is well-known that any k -means clustering of the rows of A may be represented by a particular orthogonal projection matrix M , such that the cost of the clustering is equal to $\|A - MA\|_F^2$, see e.g. Section 2.3 in [39] for an in-depth explanation. Thus, for any clustering C_1, \dots, C_k of P and the corresponding clustering $\tilde{C}_1, \dots, \tilde{C}_k$ of \tilde{P} , it holds that

$$\text{cost}(C_1, \dots, C_k) \leq \text{cost}(\tilde{C}_1, \dots, \tilde{C}_k) + c \leq (1 + 3\epsilon_0)\text{cost}(C_1, \dots, C_k). \quad (8)$$

In particular, if we equip \tilde{P} with the same l groups as P , (8) holds for any fair clustering.

We run the algorithm given by Theorem 8.4 on \tilde{P} to obtain a reduced weighted instance W in \mathbb{R}^m , using the error parameter ϵ_0 . Now, consider a set of k centers \tilde{C} in \mathbb{R}^m , and an assignment ψ from W to \tilde{C} that has the cost of at most $\gamma \text{faircost}(W)$. By Theorem 8.4, ψ can be lifted to a fair assignment $\tilde{\varphi}$ from \tilde{P} to \tilde{C} with the cost of at most $(1 + \epsilon_0)\gamma \text{faircost}(\tilde{P})$. Consider the clustering $\{\tilde{C}_1, \dots, \tilde{C}_k\}$ of \tilde{P} that corresponds to the assignment $\tilde{\varphi}$. Consider also the clustering $\{C_1, \dots, C_k\}$ of P that corresponds to $\{\tilde{C}_1, \dots, \tilde{C}_k\}$, i.e. for each $i \in [k]$, C_i contains exactly the preimages of points in \tilde{C}_i under sketching. The resulting set of centers $C = \{c_1, \dots, c_k\}$ is the set of means of the clusters C_1, \dots, C_k , that is, for each $i \in [k]$, $c_i = \mu(C_i)$. The resulting assignment φ sends C_i to c_i , for each $i \in [k]$. Clearly, φ is a fair assignment since it clusters together exactly the same points as $\tilde{\varphi}$, and $\tilde{\varphi}$ is a fair assignment by Theorem 8.4. Now we bound the cost of φ , by (8),

$$\text{cost}(\varphi) = \text{cost}(C_1, \dots, C_k) \leq \text{cost}(\tilde{C}_1, \dots, \tilde{C}_k) + c \leq (1 + \epsilon_0)\gamma \text{faircost}(\tilde{P}) + c.$$

To bound $\text{faircost}(\tilde{P})$ in terms of $\text{faircost}(P)$, consider an optimal clustering C_1^*, \dots, C_k^* of P , and the corresponding clustering $\tilde{C}_1^*, \dots, \tilde{C}_k^*$ of \tilde{P} . By (8),

$$\begin{aligned} \text{faircost}(\tilde{P}) + c &\leq \text{cost}(\tilde{C}_1^*, \dots, \tilde{C}_k^*) + c \leq (1 + \epsilon_0)\text{cost}(C_1^*, \dots, C_k^*) \\ &= (1 + \epsilon_0) \text{faircost}(P) \end{aligned}$$

Combining it with the earlier bound on $\text{cost}(\varphi)$, we obtain

$$\text{cost}(\varphi) \leq (1 + \epsilon_0)^2 \gamma \text{faircost}(P).$$

Finally, setting ϵ_0 such that $(1 + \epsilon_0)^2 \leq (1 + \epsilon)$ shows that φ satisfies the statement of the theorem.

The running time of the algorithm reducing P to W is $O((k/\epsilon^2)nd \log n)$ from Proposition 8.8 and $O(\Gamma k^3/\epsilon^2 nd \log n)$ from Theorem 8.4. The algorithm computing (C, φ) from (\tilde{C}, ψ) runs in time $O(\Gamma k^3/\epsilon^2 nd \log n)$ by Theorem 8.4, plus an additional $O(ndk)$ time required to compute ϕ and C from $\tilde{\varphi}$. Clearly, $O(\Gamma k^3/\epsilon^2 nd \log n)$ dominates the total running time. \square

As Theorem 8.4, Theorem 8.9 allows to speed up any approximate algorithm for weighted (α, β) -Fair k -means by running it on the small-sized coreset in the low-dimensional space instead of the original points. In particular, we obtain an analogue of Theorem 8.5.

Theorem 8.10. *There is a randomized algorithm that given an instance P of (α, β) -Fair k -means and a parameter $0 < \epsilon \leq 1$ outputs a set of k centers C and a fair assignment $\varphi : P \rightarrow C$ such that $\text{cost}(\varphi) \leq (1 + \epsilon) \text{faircost}(P)$ with high probability. The running time of the algorithm is*

$$O(\Gamma k^3/\epsilon^2 nd \log n) + 2^{\tilde{O}(k/\epsilon^{O(1)})} (k\Gamma)^{O(k\Gamma)} (\log n)^{O(1)}.$$

Proof. The proof is identical to the proof of Theorem 8.5, the only difference is that Theorem 8.9 is used to reduce the instance, instead of Theorem 8.4. \square

The benefit of the algorithm in Theorem 8.10 compared to Theorem 8.3 is that only the “simple” steps like sketching, sampling the coreset, and running the flow to restore the assignment, are applied to the “big” original data. While the “heavy” part of the algorithm that has an exponential dependency on the parameters, deals exclusively with the compressed instance, with the size independent of the dimension d , and polylogarithmic in the number of points n . It might be said that the combination of the dimensionality reduction and our coreset construction in the proof of Theorem 8.9 obtains a coreset of size $O((k \log n/\epsilon)^{O(1)})$ for fair k -means in the Euclidean case. However, since after reducing the dimension the points lie in a different low-dimensional space, our definition of a universal coreset could not be applied to the coreset with respect to the original points. Therefore we do not state Theorem 8.9 as a coreset result, but rather as a reduction procedure.

Finally, note that we only implement the dimensionality reduction for k -means, since for k -median the reduction techniques are more limiting. In particular, the correspondence between clusterings and particular orthogonal projection operators does not hold. It is an open question whether it is possible to achieve the analogue of Theorem 8.9 for k -median.

9. $(3 + \epsilon)$ - and $(9 + \epsilon)$ -approximations in general metric

In this section, we show a $(3 + \epsilon)$ -approximation algorithm for fair k -median in general metric, and $(9 + \epsilon)$ -approximation for fair k -means in general metric. After computing the coreset by Theorem 4.1, the strategy is essentially identical to that used in [77] and [34]: from each of the clusters in an optimal solution on the coreset we guess the closest point to the center, called a *leader* of that cluster. We also guess a suitably discretized distance from each leader to the center of the corresponding cluster. Finally, selecting any center that has roughly the guessed distance to the leader provides us with a $(3 + \epsilon)$ -approximation, in the case of k -median. Now we state formally the main result of the section.

Theorem 9.1. *For any $1 \geq \epsilon > 0$, there exists a $(3 + \epsilon)$ -approximation algorithm for (α, β) -Fair k -median, and $(9 + \epsilon)$ -approximation algorithm for (α, β) -Fair k -means. Both algorithms run in time*

$$(k\Gamma)^{O(k\Gamma)} / \epsilon^{O(k)} \cdot n + \Gamma k^3 / \epsilon^2 n \log n.$$

Note that the distance guessing step of our algorithm requires that the *aspect ratio* of the instance is bounded by a polynomial in n , where the aspect ratio is the ratio of the maximum distance between the points to the minimum distance. As opposed to the case of capacitated clustering studied in [34], achieving polynomial aspect ratio is less straightforward for fair clustering, since there was no previously known true approximation algorithm for the general version of fair clustering. We refer the reader to the introduction for the discussion on assumptions and limitations in previous works. So, for the ease of presentation, we first prove Theorem 9.1 under the polynomial aspect ratio assumption, and later show how to achieve this assumption for any instance.

Claim 9.2. *The statement of Theorem 9.1 holds in the case when the aspect ratio of the input instance is bounded by $n^{O(1)}$.*

Proof. For now, focus on the case of k -median. Fix a small positive number $\epsilon_0 < \epsilon$ that will be defined later. We start by computing a universal coreset W of size $O(\Gamma(k \log n)^2 \epsilon_0^{-3})$ by Theorem 4.1, applied with the error parameter ϵ_0 . Then we try all possible sets of k points l_1, \dots, l_k out of the points in the coreset W . We also try all possible sets of k values R_1, \dots, R_k , where each R_i ranges from the minimum distance between the points in the space to the maximum distance, taking values that are powers of $(1 + \epsilon_0)$ times the minimum distance. Thus, there are $|W|^k$ choices of l_1, \dots, l_k , and $(\log n / \epsilon_0)^{O(k)}$ choices for R_1, \dots, R_k , since the ratio of maximum distance to minimum distance is at most $n^{O(1)}$. Now, for every choice of l_1, \dots, l_k and R_1, \dots, R_k , we take a tuple of k centers $C = (c_1, \dots, c_k)$ such that $d(l_i, c_i) \in [R_i, (1 + \epsilon_0)R_i]$ for every $i \in [k]$. If for $i \in [k]$ there are multiple choice of c_i , we take any one of them. If for some $i \in [k]$ there is no suitable c_i , we continue to the next choice of l_1, \dots, l_k , and R_1, \dots, R_k . After the centers are fixed, we run the assignment algorithm given by Theorem 7.2 on the coreset W and the centers C . Out of all considered tuples of centers, we select the one with the lowest cost of the assignment. Then we compute a fair assignment from P to these centers with the help of Lemma 7.5, and return the assignment and the centers. This concludes the algorithm.

For the proof of correctness, consider an optimal solution $C^* = \{c_1^*, \dots, c_k^*\}$. Since W is a universal coreset of P , $\text{faircost}(P, C^*) \leq (1 + \epsilon_0) \text{faircost}(W, C^*)$. Consider an optimal assignment φ from W to C^* achieving the cost of $\text{faircost}(W, C^*)$. Take l_1^*, \dots, l_k^* such that l_i^* is the closest point to c_i^* among the points in $\varphi^{-1}(c_i^*)$, for each $i \in [k]$. Here by $\varphi^{-1}(c_i^*)$ we mean the set of points in W such that φ sends positive weight from them to c_i^* . Take R_1^*, \dots, R_k^* such that for each $i \in [k]$, $R_i^* = (1 + \epsilon_0)^t m$ for a certain nonnegative integer t , where m is the minimum distance between the points, and $R_i^* \leq d(l_i^*, c_i^*) < (1 + \epsilon_0)R_i^*$. At some point, the algorithm considers the choice of l_1^*, \dots, l_k^* and R_1^*, \dots, R_k^* , take the tuple of centers $C = (c_1, \dots, c_k)$ obtained by the algorithm at this iteration. We know that C exists since (c_1^*, \dots, c_k^*) is one of the possible choices for C . Consider the assignment ψ from W to C that behaves in the same way as φ : for each $i \in [k]$, ψ sends to c_i exactly the same weight from the same points in W , as φ does to c_i^* . Clearly, ψ is a fair assignment since the composition of each cluster is exactly the same as for φ . Now we bound the cost of ψ , for each point x in the coreset W and each center c_i such that a positive weight is assigned from x to c_i by ψ , it holds that

$$d(x, c_i) \leq d(x, l_i^*) + d(l_i^*, c_i) \leq d(x, c_i^*) + d(c_i^*, l_i^*) + d(l_i^*, c_i) \leq d(x, c_i^*) + (2 + \epsilon_0)d(c_i^*, l_i^*).$$

The first two inequalities are by triangle inequality, and the last is since $d(l_i^*, c_i^*)$ is at least R_i^* , and $d(l_i^*, c_i)$ is at most $(1 + \epsilon_0)R_i^*$. Moreover, l_i^* is chosen in a way that $d(l_i^*, c_i^*) \leq d(x, c_i^*)$, thus $d(x, c_i) \leq (3 + \epsilon_0)d(x, c_i^*)$. Now, the total cost of ψ is

$$\begin{aligned} \sum_{x \in W} \sum_{i=1}^k \psi(x, c_i) \cdot d(x, c_i) &\leq \sum_{x \in W} \sum_{i=1}^k (3 + \epsilon_0) \psi(x, c_i) \cdot d(x, c_i^*) \\ &= (3 + \epsilon_0) \sum_{x \in W} \sum_{i=1}^k \varphi(x, c_i^*) \cdot d(x, c_i^*) \end{aligned}$$

$$\begin{aligned}
&= (3 + \epsilon_0) \cdot \text{faircost}(W, C^*) \\
&\leq (3 + \epsilon_0)(1 + \epsilon_0) \text{faircost}(P, C^*).
\end{aligned}$$

Observe that $\text{faircost}(P, C) \leq \frac{1}{1-\epsilon_0} \text{faircost}(W, C)$ since W is a universal coresets. The assignment ψ is a particular fair assignment from W to C , thus its cost is at least $\text{faircost}(W, C)$, and finally we get

$$\text{faircost}(P, C) \leq \frac{1}{1-\epsilon_0} (3 + \epsilon_0)(1 + \epsilon_0) \text{faircost}(P, C^*).$$

Recall that $\text{faircost}(P, C^*)$ is the cost of an optimal solution, and that Lemma 7.5 returns a fair assignment of cost at most $(1 + \epsilon_0) \text{faircost}(P, C)$. Thus setting ϵ_0 small enough such that $(3 + \epsilon_0) \frac{(1+\epsilon_0)^2}{1-\epsilon_0}$ is at most $(3 + \epsilon)$, provides the desired approximation.

For the running time, recall that first we compute the coresets in time $O(n(k+l))$, and then we consider

$$|W|^k (\log n / \epsilon_0)^{O(k)} = (\Gamma(k \log n)^2 / \epsilon_0^3)^k (\log n / \epsilon_0)^{O(k)} = (k \Gamma \log n / \epsilon_0)^{O(k)}$$

tuples of k centers, and for each of them we run the assignment algorithm in time $(k \Gamma)^{O(k \Gamma)} (\log n / \epsilon_0)^{O(1)}$. Thus, the total running time is $n(k+l) + (k \Gamma)^{O(k \Gamma)} (\log n)^{O(k)} / \epsilon^{O(k)}$. Note that for any constant $c > 0$, $(\log n)^{O(k)}$ might be upper-bounded by $n^c + k^{O_c(k)}$, and we can bound the total running time required to find the best centers by $(k \Gamma)^{O(k \Gamma)} / \epsilon^{O(k)} \cdot n$. Finally, an additional term of $O(\Gamma k^3 / \epsilon^2 n \log n)$ is from Lemma 7.4.

Now to the case of fair k -means. The algorithm and analysis are essentially the same, up to a few minor details. For the coresets construction here we use Theorem 6.2 that constructs a universal coresets with respect to the k -means objective. The size of the coresets is still bounded by $\Gamma(k \log n / \epsilon)^{O(1)}$. Now the only difference is the bound on the cost of the assignment ψ . It becomes

$$\begin{aligned}
\sum_{x \in W} \sum_{i=1}^k \psi(x, c_i) \cdot d(x, c_i)^2 &\leq \sum_{x \in W} \sum_{i=1}^k (3 + \epsilon_0)^2 \psi(x, c_i) \cdot d(x, c_i^*)^2 \\
&= (3 + \epsilon_0)^2 \sum_{x \in W} \sum_{i=1}^k \varphi(x, c_i^*) \cdot d(x, c_i^*)^2 \\
&= (3 + \epsilon_0)^2 \cdot \text{faircost}(W, C^*) \\
&\leq (3 + \epsilon_0)^2 (1 + \epsilon_0) \text{faircost}(P, C^*).
\end{aligned}$$

Analogously, we obtain

$$\text{faircost}(P, C) \leq \frac{1}{1-\epsilon_0} (3 + \epsilon_0)^2 (1 + \epsilon_0) \text{faircost}(P, C^*),$$

and we set ϵ_0 small enough such that $(3 + \epsilon_1)^2 \frac{(1+\epsilon_0)^2}{1-\epsilon_0} \leq 9 + \epsilon$ to finally get the desired $(9 + \epsilon)$ approximation. \square

9.1. Polynomial aspect ratio

We follow the standard trick to reduce the aspect ratio of the instance, see e.g. [34]. For that, we require an estimate of the cost of the optimal solution. So we start with showing a $O(n)$ -factor approximation algorithm for (α, β) -FAIR CLUSTERING. This algorithm combines the simple linear time $O(n)$ -approximation to the vanilla clustering problem, then the argument due to Bera et al. [3] that a set of centers that provides a good approximation w.r.t. vanilla clustering objective is also good enough for the purpose of fair clustering, and finally our assignment algorithm given by Theorem 7.2. We state a slight modification of the result of Bera et al. [3] first.

Proposition 9.3 (Lemma 3 in [3]). Assume we are given a ρ -approximation algorithm \mathcal{A} for k -median. Run \mathcal{A} on the input set of points P and denote by C the returned set of centers. It holds that $\text{faircost}(P, C)$ is at most $(\rho + 2)$ times the cost of an optimal solution to (α, β) -Fair k -median on P . The same holds for k -means and (α, β) -Fair k -means, only the cost factor is $(\rho + 2)^2$.

Proof. The statement for k -median is exactly a special case of Lemma 3 in [3] where we only restrict to k -median and k -means, and the assignment algorithm has no violation of the constraints. For k -means, Lemma 3 in [3] holds for the same k -means and (α, β) -Fair k -means we consider in this paper, with the only difference that their objective function is the square root of the sum of squared distances. Thus, from their lemma we immediately get that the square root of the cost of the approximate solution is at most $(\rho + 2)$ times the square root of the cost of the optimal solution. Squaring both sides provides the approximation factor of $(\rho + 2)^2$. \square

To achieve a linear-time algorithm, we are rather restricted in what kind of algorithm we can use to get the initial approximation for the vanilla clustering. Thus we use a simple $O(n)$ -approximation given by the classical k -center algorithm that is enough for our purposes. We also need to use the coresset construction as an intermediate step, so that computing the fair assignment takes time sublinear in n . Note that in this result, we do not aim to return the actual fair assignment, just the approximation to the cost.

Lemma 9.4. *There exists a $O(n)$ -factor approximation algorithm for computing the optimal cost in both (α, β) -Fair k -median and (α, β) -Fair k -means, with the running time of $(k\Gamma)^{O(k\Gamma)} \cdot n$.*

Proof. For k -median, we start with computing the initial approximation using the min-max algorithm for k -center [78] in time $O(nk)$. It is well-known that this gives a $O(n)$ -approximation of the k -median objective. For the obtained set C of k centers, by Proposition 9.3 it holds that $\text{faircost}(P, C)$ is at most $O(n)$ times the optimal fair clustering cost of P . So it only remains to run the assignment algorithm. First, we compute a universal coresset W of P of size $O(\Gamma(k \log n)^{O(1)})$ by Theorem 4.1, using a constant error parameter. Then we run the assignment algorithm given by Theorem 7.2 on the weighted points W and the centers C . By definition of a universal coresset, the cost is changed by at most a constant factor, thus the cost of the fair assignment achieved by the algorithm is an $O(n)$ -approximation of the optimal solution on the original points P . The total running time of the algorithm is $O(nk + n(k + l) + (k\Gamma)^{O(k\Gamma)} \cdot (\log n)^{O(1)})$.

For k -means, the only difference is that we run the k -center min-max algorithm using distances given by $d'(u, v) = d(u, v)^2$ for all $u, v \in \mathcal{X}$. This is not a metric, but it holds that $d'(u, v) \leq 2(d'(u, w) + d'(w, v))$ for all $u, v, w \in \mathcal{X}$, since $d(u, v)^2 \leq (d(u, w) + d(w, v))^2 \leq 2(d(u, w)^2 + d(w, v)^2)$. The min-max algorithm still obtains a $O(1)$ -approximation for the k -center objective with such a relaxed triangle inequality, and thus a $O(n)$ -approximation for the k -means objective with the original distances given by d . The rest is the same, but we invoke Theorem 6.2 to obtain the coresset. \square

Finally, we show how to reduce an arbitrary instance of (α, β) -FAIR CLUSTERING to an equivalent one that has polynomial aspect ratio by modifying the distances.

Lemma 9.5. *Given an instance of (α, β) -FAIR CLUSTERING in the metric space with the distance function d , we can construct a distance function d' such that the cost of any n^{10} -approximate solution changes by at most a factor of $1 + 1/n$. The distance function d' has polynomial aspect ratio. This requires the preprocessing time of $(k\Gamma)^{O(k\Gamma)} \cdot n$.*

Proof. We state first how d' is obtained from d . By Lemma 9.4 compute D that is a $O(n)$ -approximation of the cost of an optimal solution. Set $D_{\max} = 2n^{10}D$ and $D_{\min} = \alpha D/n^3$, for a sufficiently small constant α . For all the distances that are larger than D_{\max} , set them to D_{\max} . Then increase the distance between every pair of points by D_{\min} . Clearly, the distances still form a metric.

No solution to the instance of (α, β) -FAIR CLUSTERING that has the cost of at most n^{10} times the optimal cost uses distances that are set to D_{\max} , since $D_{\max} = 2n^{10}D$ and D is at least the cost of the optimal solution. Thus the decreasing large distances to D_{\max} does not affect the instance. Due to increasing by D_{\min} , the cost of any solution is increased by at most a factor of $(1 + 1/n)$, since $n \cdot D_{\min} = \alpha D/n^2$. This is at most $1/n$ times the cost of the optimal solution, since D is a $O(n)$ -approximation. Now the aspect ratio is at most $D_{\max}/D_{\min} = O(n^{13})$.

Note that running Lemma 9.4 incurs the preprocessing time of $(k\Gamma)^{O(k\Gamma)} \cdot n$. After D_{\max} and D_{\min} are computed, the distance oracle for d' is obtained from the distance oracle for d by an extra constant time per query. \square

Finally, we prove Theorem 9.1 by combining Lemma 9.4 and Claim 9.2.

Proof of Theorem 9.1. Apply Lemma 9.4 to obtain the new instance of (α, β) -FAIR CLUSTERING with polynomial aspect ratio. For every solution that is $3 + \epsilon$ (or $9 + \epsilon$ in the case of k -means), the cost w.r.t. the new instance is at least the cost w.r.t. the old instance, and at most $(1 + 1/n)$ of that cost. Observe that $1/\epsilon = o(n)$, otherwise all possible sets of centers can be trivially enumerated in time $n^k = (1/\epsilon)^{O(k)}$. Thus $(1 + 1/n) \leq (1 + \epsilon/3)$, and invoking Claim 9.2 with the error parameter $\epsilon/3$ finishes the proof. \square

10. Algorithms for other clustering problems

We note that the algorithms for fair clustering in general metrics suggest a generic algorithm for any clustering problem with constraints, such that the constraints can be represented by a set of matrices. Here we state this algorithm. Let \mathcal{D} be the set of all possible distinct distances. Also, let D_{\min} and D_{\max} be the minimum and maximum distances in \mathcal{D} , respectively. This algorithm has the following steps.

- Compute a universal coresset W .
- For every pair of tuples (l_1, \dots, l_k) and (R_1, \dots, R_k) such that $l_i \in W$ for all i and $R_j \in \mathcal{D}$ for all j , do the following.

- Select a set $C = \{c_1, \dots, c_k\}$ of centers such that $c_i \in F$ and $d(l_i, c_i) \in [R_i, (1 + \epsilon)R_i]$ for all i . If no such set C exists, probe the next choice.
- Find an assignment of the points in P to the centers in C of the minimum cost that satisfies the respective clustering constraints (assignment problem).
- Return the set of centers and the assignment that minimizes the cost over all choices.

From the analysis for fair clustering, we have the following theorem.

Theorem 10.1. *Consider any clustering problem with constraint K , such that the constraint can be represented by a set of matrices, and suppose the aspect ratio D_{\max}/D_{\min} is bounded by Δ for all instances. Moreover, suppose the universal coreset can be computed in $T_1(n, k, \ell)$ time and the assignment problem for the clustering problem can be solved in $T_2(n, k)$ time. Then one can obtain, w.h.p., a $(3 + \epsilon)$ - (resp. $(9 + \epsilon)$ -) approximation for k -median (resp. k -means) with constraint K in time $T_1(n, k, \ell) + (\epsilon^{-1} k \Gamma \log(n + \Delta))^{O(k)} \cdot T_2(n, k)$.*

From the above theorem, it is sufficient to (i) show that the aspect ratios of instances are bounded and (ii) design an efficient algorithm for the assignment problem, to obtain constant approximations for a clustering problem with constraints. We will use this theorem on various clustering problems.

For the Euclidean version of clustering problems with constraints, we have the following generic algorithm.

- Compute a universal coreset W .
- Apply the algorithm mentioned in Proposition 8.1 to find the list \mathcal{L} of candidate sets of centers.
- For every set $C = \{c_1, \dots, c_k\}$ of centers in \mathcal{L} , do the following.
 - Find an assignment of the points in P to the centers in C of the minimum cost that satisfies the respective clustering constraints (assignment problem).
- Let C' be the set of centers that minimizes the cost over all choices. Find an assignment of the points in P to the centers in C' of the minimum cost that satisfies the respective clustering constraints.

From the analysis for fair clustering and Proposition 8.1, we know that C' is an approximately optimal set of centers with constant probability. By repeating step 3 of the above algorithm $O(\log n)$ times, we obtain this w.h.p. Hence, we have the following theorem.

Theorem 10.2. *Consider any Euclidean clustering problem with constraint K , such that the constraint can be represented by a set of matrices. Moreover, suppose the universal coreset can be computed in $T_1(n, k, \ell, d)$ time and the assignment problem for the clustering problem can be solved in $T_2(n, k)$ time. Then one can obtain, w.h.p., a $(1 + \epsilon)$ -approximation for k -median (resp. k -means) with constraint K in time $T_1(n, k, \ell, d) + 2^{\tilde{O}(k/\epsilon^{O(1)})} \cdot (nd + \log n \cdot T_2(|W|, k)) + T_2(n, k)$.*

From the above theorem, it is sufficient to design an efficient algorithm for the assignment problem, to obtain constant approximations for a clustering problem with constraints. In the following, we will use this theorem on various clustering problems.

10.1. Lower-bounded clustering

In the lower-bounded clustering problem, we are given a lower bound parameter L and the size of each cluster must be at least L . In the Euclidean version of the problem the set of points $P \subset \mathbb{R}^d$ and the set of centers $F = \mathbb{R}^d$.

First, we note that the lower bound constraint can be represented by a set of $k \times 1$ column matrices such that each of the entries is at least L and at most n .

To apply Theorem 10.1, we need to show two things as mentioned before. To show the bounded aspect ratio, we can argue in the same way as we did for fair clustering. The assignment problem for lower-bounded clustering can be modeled as a minimum cost network flow problem that can be solved in polynomial time. Indeed, the modeling is similar to the one in the proof of Lemma 7.4. We construct a bipartite network where on one side we have the n points of P and on the other side the k centers of C and an additional node w . Source s is connected to all points of P . Sink t is connected to all centers through edges of capacity L . t is also connected to w through an edge of capacity $n - kL$. w is connected to all points through an edge. The cost of the edges between points and centers are their respective distances. The cost between a point p and w is $d(p, C)$. The idea is to route L flow to each center and $n - kL$ flow to w . This is equivalent to assigning at least L points to each center using the minimum cost and assigning the remaining points to their closest neighbor in C . We also scale the distances, as mentioned in the proof of Lemma 7.4. Then, the cost of the minimum cost flow in this network (with n flow) is a $(1 + \epsilon)$ -approximation of the minimum assignment cost. Hence, by Proposition 7.3, the assignment problem in this case can be solved in time $nk^2\epsilon^{-O(1)}\log n$. Hence, the constant approximations follow for this problem in general metrics in time

$$O(nk) + (\epsilon^{-1}k \log n)^{O(k)} nk^2 \epsilon^{-O(1)} \log n = (\epsilon^{-1}k \log n)^{O(k)} n = (k/\epsilon)^{O(k)} n^{O(1)}.$$

From the above discussion, one can also apply Theorem 10.2. However, to solve the assignment problem on coreset, we do not want to spend $nk^2 \epsilon^{-O(1)} \log n$ time. We would like to obtain an algorithm that is polynomial in the size of the coreset. We do the same as we did for fair clustering. We solve a mixed ILP that has $|W| \times k$ unconstrained variables and k integer variables. The construction is much easier compared to fair clustering. The running time is $k^{O(k)} |W|^{O(1)}$. Hence, we obtain $(1 + \epsilon)$ -approximation for the Euclidean version in time

$$\begin{aligned} & 2^{\tilde{O}(k/\epsilon^{O(1)})} (nd + \log n \cdot k^{O(k)} (k \log n / \epsilon)^{O(1)}) + nk^2 \epsilon^{-O(1)} \log n \\ &= 2^{\tilde{O}(k/\epsilon^{O(1)})} (nd + (d \log n)^{O(1)}) + nk^2 \epsilon^{-O(1)} \log n. \end{aligned}$$

Theorem 10.3. For any $\epsilon > 0$, there exists a $(3 + \epsilon)$ - and a $(9 + \epsilon)$ -approximation algorithm for lower-bounded k -median and lower-bounded k -means, respectively, that runs in time $(k/\epsilon)^{O(k)} n^{O(1)}$. In \mathbb{R}^d , there are improved $(1 + \epsilon)$ -approximation algorithms for both of the problems that run in time $2^{\tilde{O}(k/\epsilon^{O(1)})} (nd + (d \log n)^{O(1)}) + nk^2 \epsilon^{-O(1)} \log n$.

10.2. Capacitated clustering

Here we study the Euclidean capacitated clustering where the set of points $P \subset \mathbb{R}^d$ and the set of centers $F = \mathbb{R}^d$. Additionally, we are given a capacity parameter U and the size of each cluster must be at most U .

First, we note that the capacity constraint can be represented by a set of $k \times 1$ column matrices such that each of the entries is at least 0 and at most U .

Now, the assignment problem for capacitated clustering can be modeled as a minimum cost network flow problem that can be solved in polynomial time. Again, the modeling is similar to the one in the proof of Lemma 7.4. We construct a complete bipartite network where on one side we have the n points of P and on the other side the k centers of C . Source s is connected to all points of P . Sink t is connected to all centers with edges of capacity U . The cost of the edges between points and centers are their respective distances. We also scale the distances, as mentioned in the proof. Then, the cost of the minimum cost flow in this network (with n flow) is a $(1 + \epsilon)$ -approximation of the minimum assignment cost. We note that one can assume that the aspect ratio of the input instance is bounded by $(n/\epsilon)^{O(1)}$. The assumption can be removed in the same way as in the case of fair clustering. Hence, by Proposition 7.3, the assignment problem in this case can be solved in time $nk^2 \epsilon^{-O(1)} \log n$.

We solve the assignment problem on coreset in the same way mentioned for lower-bounded clustering. A $(1 + \epsilon)$ -approximation follows for the Euclidean version in time

$$\begin{aligned} & 2^{\tilde{O}(k/\epsilon^{O(1)})} (nd + \log n \cdot k^{O(k)} (k \log n / \epsilon)^{O(1)}) + nk^2 \epsilon^{-O(1)} \log n \\ &= 2^{\tilde{O}(k/\epsilon^{O(1)})} (nd + (d \log n)^{O(1)}) + nk^2 \epsilon^{-O(1)} \log n. \end{aligned}$$

Theorem 10.4. For any $\epsilon > 0$, there exists $(1 + \epsilon)$ -approximation algorithms for capacitated k -median and capacitated k -means that run in time $2^{\tilde{O}(k/\epsilon^{O(1)})} (nd + (d \log n)^{O(1)}) + nk^2 \epsilon^{-O(1)} \log n$.

10.3. ℓ -diversity clustering

In the ℓ -Diversity clustering problem, $P = \bigcup_{i=1}^{\tilde{n}} P_i$ is a set of n colored points such that all points in P_i have the same color, and each cluster must have no more than a fraction $1/\ell$ (for some constant $\ell > 1$) of its points sharing the same color. Thus, for each cluster A and $i \in [\ell]$, $|A \cap P_i| \leq |A|/\ell$. We note that each point can have only one color. Ding and Xu [36] gave a $(1 + \epsilon)$ -approximation for this problem in \mathbb{R}^d with time complexity $O(n^2 (\log n)^{k+2} (t+1)^k d)$, where $t = \max_{1 \leq i \leq \tilde{n}} |P_i|$.

We note that ℓ -Diversity clustering is a special case of (α, β) -fair clustering without the lower bound constraints involving parameter β , and $\alpha_i = 1/\ell$ for all i . Thus, we obtain algorithms for this problem with bounds same as for (α, β) -fair clustering, including a $(1 + \epsilon)$ -approximation that significantly improves the time complexity of the one in [36].

Theorem 10.5. For any $\epsilon > 0$, there exists a $(3 + \epsilon)$ - and a $(9 + \epsilon)$ -approximation algorithm for ℓ -Diversity k -median and ℓ -Diversity k -means, respectively, that runs in time $(k\ell)^{O(k\ell)} / \epsilon^{O(k)} \cdot n + \ell k^3 / \epsilon^2 n \log n$. In \mathbb{R}^d , there are improved $(1 + \epsilon)$ -approximation algorithms for both of the problems that run in time $2^{\tilde{O}(k/\epsilon^{O(1)})} (k\ell)^{O(k\ell)} nd \log n$.

10.4. Chromatic clustering

In chromatic clustering, again $P = \bigcup_{i=1}^{\tilde{n}} P_i$ is a set of n colored points such that all points in P_i have the same color, and each cluster contains at most one point from P_i for each i . Ding and Xu [36] obtained a linear time $(1 + \epsilon)$ -approximation for this problem in \mathbb{R}^d . To the best of our knowledge the metric version of the problem was not studied before. Thus, we give the first constant approximation for this version.

First, we note that the chromatic constraint can be represented by a set of $k \times \ell$ matrices with 0/1 entries.

Ding and Wu [36] showed that the assignment problem for chromatic clustering can be modeled as a bipartite matching problem that can be solved in $O(k^3 n)$ time. However, it is not clear how to bound the aspect ratio for this problem. Hence, the constant approximations follow for this problem in general metrics in time $O(nk) + (\epsilon^{-1} k \ell \log(n + \Delta))^{O(k)} n$.

Theorem 10.6. *For any $\epsilon > 0$, there exists a $(3 + \epsilon)$ - and a $(9 + \epsilon)$ -approximation algorithm for chromatic k -median and chromatic k -means, respectively, that runs in time $O(nk) + (\epsilon^{-1} k \ell \log(n + \Delta))^{O(k)} n$.*

11. Streaming universal coresets

Here we describe a streaming algorithm for maintaining universal coreset for k -median. The algorithm can be trivially extended to k -means with a slightly larger space complexity. Our algorithm is based on the merge and reduce framework of Bentley and Saxe [79], which was first applied in the context of clustering in [80]. Indeed, in streaming setting this is a standard technique, which have been applied in many works [81,38,82]. We mainly follow the approach of Har-Peled and Mazumdar [81], which was further refined by Chen [38] for randomized coresets. In the following we describe how to maintain a small size coreset in each step. Clustering is frequently used for training ML models, where the clusters put labels on the training data. One can use the coreset as the training data instead of the original data for the purpose of training.

Let us refer to a universal coreset as an ϵ -coreset, where ϵ is the corresponding error parameter. Our approach is based on composability of coresets.

Lemma 11.1. *Suppose S_1 and S_2 are the ϵ -coresets of the points in P_1 and P_2 , respectively. Then, $S_1 \cup S_2$ is an ϵ -coreset of the points in $P_1 \cup P_2$.*

The proof of this lemma follows by definition of universal coresets and can be found in [2]. The proof assumes that the coreset points have integer weights. We note that our construction can be slightly modified to obtain coreset with integer weights (e.g., see Chen's adaptation [38]). Next, we have an observation that again follows by the definition of coresets.

Observation 11.2. *Suppose S_1 is an ϵ -coreset of the points in S_2 and S_2 is an δ -coreset of the points in S_3 . Then, S_1 is an $((1 + \epsilon)(1 + \delta) - 1)$ -coreset of the points in S_3 .*

Let λ be the confidence probability parameter for the coreset we want to construct. Suppose the points arrive in the order p_1, p_2, \dots and let $P = (p_1, \dots, p_n)$ be the set of points arrived so far. We partition P into $t + 1$ subsets P_0, P_1, \dots, P_t such that $|P_i| = 2^i T$, where $T = \lceil \ell k^2 / \epsilon^3 \rceil$.

Let $\rho_j = \epsilon / (b \cdot (j + 1)^2)$ for a sufficiently large constant b , and $1 + \delta_j = \prod_{i=0}^j (1 + \rho_i)$ for $j = 1, \dots, \lceil \log n \rceil$. It is not hard to verify that $1 + \delta_j \leq 1 + \epsilon$ for all j . We maintain a δ_j -coreset Q_j for the points in P_j , where $Q_0 = P_0$. Thus, by composability of coresets, $\cup_{j \geq 0} Q_j$ is an ϵ -coreset for points in P .

When a new point p_m arrives, we add it to Q_0 . If Q_0 contains fewer than T points, we are done. Otherwise, let $r \geq 1$ be the minimum index such that Q_r is empty. We compute a ρ_r -coreset Q'_r of $\cup_{j=0}^{r-1} Q_j$ with confidence parameter $\lambda_m = \lambda / m^2$ and set Q'_r to be Q_r . We also make all the sets Q_j empty for $1 \leq j \leq r - 1$. It is not hard to verify that the total weight of the points in Q_r is $2^{r-1} T$.

Note that here we need to compute coresets of weighted points. We can trivially extend our coreset construction algorithm in the sequential setting to handle points with integer weights. For example, for a point p with weight w , now we treat it as the point p with w copies. Instead of using the algorithm of Indyk at the start, we use our algorithm in Section 8. Thus the algorithm can be implemented where the space complexity is linear in the number of points.

Next we claim that $Q = \cup_{i \geq 0} Q_i$ is an ϵ -coreset of the points received so far w.p. at least $1 - \lambda$. First, note that Q_r is constructed by computing a ρ_r -coreset of $\cup_{j=0}^{r-1} Q_j$. By applying Observation 11.2 repetitively, Q_r is a $(\prod_{i=0}^r (1 + \rho_i) - 1)$ -coreset of the corresponding subset of the input points w.p. at least $1 - \lambda / m^2$. Now for $m \geq T$, when p_m arrives, our computation fails w.p. at most λ / m^2 . The failure probability over all iterations is at most $\sum_{m=T}^n \lambda / m^2 \leq \lambda$ for $T \geq 2$. As $\prod_{i=0}^r (1 + \rho_i) - 1 = 1 + \delta_r \leq 1 + \epsilon$, the claim follows by composability.

Now, we bound the size of individual coreset Q_i . Note that $|Q_0| \leq T = \lceil \ell k^2 / \epsilon^3 \rceil$. Also for $i \geq 1$, Q_i is constructed for a subset of input points of size $2^{i-1} T$ and with error parameter $\rho_i = \epsilon / (b \cdot (i + 1)^2)$. Thus by Theorem 4.1, the size of Q_i is

$$O(\ell k^2 \log(2^i + T)(\log(2^i + T) + d \log(1/\epsilon)) \cdot i^6 / \epsilon^3) = O(d \ell k^2 (\log n)^8 / \epsilon^4).$$

Hence, the total size of the coreset $Q = \cup_{i=0}^{\lceil \log n \rceil} Q_i$ is bounded by $O(d \ell k^2 (\log n)^9 / \epsilon^4)$. In \mathbb{R}^d , we need $O(d)$ space for storing each point, and thus we obtain the following theorem.

Theorem 11.3. *In one pass streaming model, a universal coreset for k -median clustering of size $O(d^2 \ell k^2 (\log n)^9 / \epsilon^4)$ can be computed w.h.p.*

12. Conclusions and open questions

In this paper, we studied the widely popular fair clustering problem with k -median and k -means objectives. Our universal coresets construction allows us to obtain the first coreset for fair clustering in general metric spaces. The coreset size is comparable to the best-known bound in the vanilla case. For Euclidean spaces, we obtain the first coreset for this problem whose size does not depend exponentially on the dimension. In the vanilla case, it is possible to construct coresets of size $(k/\epsilon)^{O(1)}$. Thus, an interesting open question is to remove the dependence on d and $(\log n)^{O(1)}$ completely from our coreset size.

The new coreset construction helps to design improved FPT constant-approximations for a wide range of problems including fair clustering in general metrics and $(1 + \epsilon)$ -approximations in the Euclidean metric. However, for fair clustering, it is not trivial to find an optimal solution on a coreset like in the case of other popular clustering problems. This is true, as the assignment problem is not easy to solve in this case. We give a novel algorithm for this problem that runs in time FPT parameterized by k and Γ . We note that for (t, k) -fair clustering the factor of $(k\ell)^{O(k\ell)}$ in the running time of our algorithms can be improved to only $(k\ell)^{O(k)}$, as in this case the assignment problem can be solved in time FPT parameterized by only k . Designing a polynomial time constant-approximation for fair clustering still remains an open question. Our $(1 + \epsilon)$ -approximation algorithms in the Euclidean case run in near-linear time. It would be interesting to see if one can obtain similar $(1 + \epsilon)$ -approximation in linear time matching the bound of the vanilla case.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

References

- [1] F. Chierichetti, R. Kumar, S. Lattanzi, S. Vassilvitskii, Fair clustering through fairlets, in: *Advances in Neural Information Processing Systems*, 2017, pp. 5029–5037.
- [2] M. Schmidt, C. Schwiegelshohn, C. Sohler, Fair coresets and streaming algorithms for fair k -means, in: *International Workshop on Approximation and Online Algorithms*, Springer, 2019, pp. 232–251.
- [3] S. Bera, D. Chakrabarty, N. Flores, M. Negahbani, Fair algorithms for clustering, in: *Advances in Neural Information Processing Systems*, 2019, pp. 4954–4965.
- [4] I.O. Bercea, M. Groß, S. Khuller, A. Kumar, C. Rösner, D.R. Schmidt, M. Schmidt, On the cost of essentially fair clusterings, in: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2019)*, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
- [5] L. Huang, S. Jiang, N. Vishnoi, Coresets for clustering with fairness constraints, in: *Advances in Neural Information Processing Systems*, 2019, pp. 7589–7600.
- [6] A. Backurs, P. Indyk, K. Onak, B. Schieber, A. Vakilian, T. Wagner, Scalable fair clustering, in: *International Conference on Machine Learning*, 2019, pp. 405–413.
- [7] M. Böhm, A. Fazzzone, S. Leonardi, C. Schwiegelshohn, Fair clustering with multiple colors, arXiv preprint arXiv:2002.07892, 2020.
- [8] X. Chen, B. Fain, L. Lyu, K. Munagala, Proportionally fair clustering, in: *International Conference on Machine Learning*, 2019, pp. 1032–1041.
- [9] S. Ahmadian, A. Epasto, R. Kumar, M. Mahdian, Clustering without over-representation, in: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019, pp. 267–275.
- [10] M. Kleindessner, P. Awasthi, J. Morgenstern, Fair k -center clustering for data summarization, in: *36th International Conference on Machine Learning, ICML 2019, International Machine Learning Society (IMLS)*, 2019, pp. 5984–6003.
- [11] J. Dischler, Putting machine learning into the hands of every advertiser, Google Blog 10 (2018) 2018.
- [12] A.E. Khandani, A.J. Kim, A.W. Lo, Consumer credit-risk models via machine-learning algorithms, *J. Bank. Finance* 34 (11) (2010) 2767–2787.
- [13] R. Malhotra, D.K. Malhotra, Evaluating consumer loans using neural networks, *Omega* 31 (2) (2003) 83–96.
- [14] C. Perlich, B. Dalessandro, T. Raeder, O. Stitelman, F. Provost, Machine learning for targeted display advertising: transfer learning in action, *Mach. Learn.* 95 (1) (2014) 103–127.
- [15] J. Angwin, J. Larson, S. Mattu, L. Kirchner, Machine bias: there's software used across the country to predict future criminals. And it's biased against blacks, in: *ProPublica*, May 23, 2016.
- [16] A. Datta, M.C. Tschantz, A. Datta, Automated experiments on ad privacy settings: a tale of opacity, choice, and discrimination, *Proc. Priv. Enh. Technol.* 2015 (1) (2015) 92–112.
- [17] H.N. Garb, Race bias, social class bias, and gender bias in clinical judgment, *Clin. Psychol.: Sci. Pract.* 4 (2) (1997) 99–120.
- [18] J. Larson, S. Mattu, L. Kirchner, J. Angwin, How we analyzed the compas recidivism algorithm, *ProPublica* 9 (1) (2016), (2016).
- [19] S.-Y. Jiang, Q. Zheng, Q.-S. Zhang, Clustering-based feature selection, *Acta Electron. Sin.* 36 (12) (2008) 157–160.
- [20] E.L. Glassman, R. Singh, R.C. Miller, Feature engineering for clustering student solutions, in: *Proceedings of the First ACM Conference on Learning @ Scale Conference*, 2014, pp. 171–172.
- [21] T. Calders, S. Verwer, Three naive Bayes approaches for discrimination-free classification, *Data Min. Knowl. Discov.* 21 (2) (2010) 277–292.
- [22] C.S. Crowson, E.J. Atkinson, T.M. Therneau, Assessing calibration of prognostic risk scores, *Stat. Methods Med. Res.* 25 (4) (2016) 1692–1706.
- [23] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, R. Zemel, Fairness through awareness, in: *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, 2012, pp. 214–226.
- [24] J. Kleinberg, S. Mullainathan, M. Raghavan, Inherent trade-offs in the fair determination of risk scores, in: *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.

- [25] A. Chouldechova, Fair prediction with disparate impact: a study of bias in recidivism prediction instruments, *Big Data* 5 (2) (2017) 153–163.
- [26] M. Feldman, S.A. Friedler, J. Moeller, C. Scheidegger, S. Venkatasubramanian, Certifying and removing disparate impact, in: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2015, pp. 259–268.
- [27] C. Rösner, M. Schmidt, Privacy preserving clustering with constraints, in: *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [28] D. Feldman, M. Langberg, A unified framework for approximating and clustering data, in: *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, 2011, pp. 569–578.
- [29] V. Cohen-Addad, K.G. Larsen, D. Saulpic, C. Schwiegelshohn, Towards optimal lower bounds for k-median and k-means coresets, in: S. Leonardi, A. Gupta (Eds.), *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing*, Rome, Italy, June 20–24, 2022, ACM, 2022, pp. 1038–1051.
- [30] D. Feldman, M. Schmidt, C. Sohler, Turning big data into tiny data: constant-size coresets for k-means, pca, and projective clustering, *SIAM J. Comput.* 49 (3) (2020) 601–657.
- [31] C. Sohler, D.P. Woodruff, Strong coresets for k-median and subspace approximation: goodbye dimension, in: *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE, 2018, pp. 802–813.
- [32] Z. Svitkina, Lower-bounded facility location, *ACM Trans. Algorithms* 6 (4) (2010) 1–16.
- [33] S. Ahmadian, C. Swamy, Improved approximation guarantees for lower-bounded facility location, in: *International Workshop on Approximation and Online Algorithms*, Springer, 2012, pp. 257–271.
- [34] V. Cohen-Addad, J. Li, On the fixed-parameter tractability of capacitated clustering, in: C. Baier, I. Chatzigiannakis, P. Flocchini, S. Leonardi (Eds.), *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9–12, 2019, Patras, Greece*, in: *LIPICs*, vol. 132, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019, 41.
- [35] V. Cohen-Addad, Approximation schemes for capacitated clustering in doubling metrics, in: S. Chawla (Ed.), *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5–8, 2020*, SIAM, 2020, pp. 2241–2259.
- [36] H. Ding, J. Xu, A unified framework for clustering constrained data without locality property, *Algorithmica* 82 (4) (2020) 808–852.
- [37] A. Bhattacharya, R. Jaiswal, A. Kumar, Faster algorithms for the constrained k-means problem, *Theory Comput. Syst.* 62 (1) (2018) 93–115, <https://doi.org/10.1007/s00224-017-9820-7>.
- [38] K. Chen, On coresets for k-median and k-means clustering in metric and Euclidean spaces and their applications, *SIAM J. Comput.* 39 (3) (2009) 923–947.
- [39] M.B. Cohen, S. Elder, C. Musco, C. Musco, M. Persu, Dimensionality reduction for k-means clustering and low rank approximation, in: *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, 2015, pp. 163–172.
- [40] S. Har-Peled, A. Kushal, Smaller coresets for k-median and k-means clustering, *Discrete Comput. Geom.* 37 (1) (2007) 3–19.
- [41] D. Chakrabarty, C. Swamy, Facility location with client latencies: Lp-based techniques for minimum-latency problems, *Math. Oper. Res.* 41 (3) (2016) 865–883.
- [42] T. Király, L.C. Lau, M. Singh, Degree bounded matroids and submodular flows, *Combinatorica* 32 (6) (2012) 703–720.
- [43] Q. Feng, J. Hu, N. Huang, J. Wang, Improved ptas for the constrained k-means problem, *J. Comb. Optim.* 37 (4) (2019) 1091–1110.
- [44] M. Kleindessner, S. Samadi, P. Awasthi, J. Morgenstern, Guarantees for spectral clustering with fairness constraints, in: *International Conference on Machine Learning*, PMLR, 2019, pp. 3458–3467.
- [45] J. Bar-Ilan, G. Kortsarz, D. Peleg, How to allocate network centers, *J. Algorithms* 15 (3) (1993) 385–415.
- [46] S. Khuller, Y.J. Sussmann, The capacitated K-center problem, *SIAM J. Discrete Math.* 13 (3) (2000) 403–418.
- [47] H. An, A. Bhaskara, C. Chekuri, S. Gupta, V. Madan, O. Svensson, Centrality of trees for capacitated k-center, *Math. Program.* 154 (1–2) (2015) 29–53.
- [48] M. Cygan, M. Hajiaghayi, S. Khuller, LP rounding for k-centers with non-uniform hard capacities, in: *FOCS*, 2012, pp. 273–282.
- [49] J. Byrka, B. Rybicki, S. Uniyal, An approximation algorithm for uniform capacitated k-median problem with $1 + \epsilon$ capacity violation, in: Q. Louveaux, M. Skutella (Eds.), *Integer Programming and Combinatorial Optimization - 18th International Conference, IPCO 2016, Proceedings, Liège, Belgium, June 1–3, 2016*, in: *Lecture Notes in Computer Science*, vol. 9682, Springer, 2016, pp. 262–274.
- [50] J. Byrka, K. Fleszar, B. Rybicki, J. Spoerhase, Bi-factor approximation algorithms for hard capacitated k-median problems, in: P. Indyk (Ed.), *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4–6, 2015*, SIAM, 2015, pp. 722–736.
- [51] M. Charikar, S. Guha, É. Tardos, D.B. Shmoys, A constant-factor approximation algorithm for the k-median problem, *J. Comput. Syst. Sci.* 65 (1) (2002) 129–149.
- [52] J. Chuzhoy, Y. Rabani, Approximating k-median with non-uniform capacities, in: *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2005, Vancouver, British Columbia, Canada, January 23–25, 2005*, SIAM, 2005, pp. 952–958.
- [53] H.G. Demirci, S. Li, Constant approximation for capacitated k-median with $(1+\epsilon)$ -capacity violation, in: *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11–15, 2016, Rome, Italy, 2016*, 73.
- [54] S. Li, On uniform capacitated k-median beyond the natural LP relaxation, in: *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4–6, 2015*, pp. 696–707.
- [55] S. Li, On uniform capacitated k-median beyond the natural LP relaxation, *ACM Trans. Algorithms* 13 (2) (2017) 22.
- [56] D.Z. Chen, J. Li, H. Liang, H. Wang, Matroid and knapsack center problems, *Algorithmica* 75 (1) (2016) 27–52.
- [57] S. Khuller, R. Pless, Y.J. Sussmann, Fault tolerant k-center problems, *Theor. Comput. Sci.* 242 (1–2) (2000) 237–245.
- [58] J. Li, K. Yi, Q. Zhang, Clustering with diversity, in: *International Colloquium on Automata, Languages, and Programming*, Springer, 2010, pp. 188–200.
- [59] S. Har-Peled, S. Mazumdar, On coresets for k-means and k-median clustering, in: L. Babai (Ed.), *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, Chicago, IL, USA, June 13–16, 2004, ACM, 2004, pp. 291–300.
- [60] S. Har-Peled, A. Kushal, Smaller coresets for k-median and k-means clustering, *Discrete Comput. Geom.* 37 (1) (2007) 3–19.
- [61] G. Frahling, C. Sohler, Coresets in dynamic geometric data streams, in: *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*, 2005, pp. 209–217.
- [62] M. Langberg, L.J. Schulman, Universal ϵ -approximators for integrals, in: *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, SIAM, 2010, pp. 598–607.
- [63] L. Becchetti, M. Bury, V. Cohen-Addad, F. Grandoni, C. Schwiegelshohn, Oblivious dimension reduction for k-means: beyond subspaces and the Johnson-Lindenstrauss lemma, in: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, 2019, pp. 1039–1050.
- [64] V. Braverman, D. Feldman, H. Lang, New frameworks for offline and streaming coreset constructions, *arXiv preprint arXiv:1612.00889*, 2016.
- [65] L. Huang, N.K. Vishnoi, Coresets for clustering in Euclidean spaces: importance sampling is nearly optimal, in: *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, 2020, pp. 1416–1429.
- [66] V. Cohen-Addad, D. Saulpic, C. Schwiegelshohn, A new coreset framework for clustering, in: S. Khuller, V.V. Williams (Eds.), *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing*, Virtual Event, Italy, June 21–25, 2021, ACM, 2021, pp. 169–182.
- [67] V. Braverman, V. Cohen-Addad, S.H. Jiang, R. Krauthgamer, C. Schwiegelshohn, M.B. Tofturp, X. Wu, The power of uniform sampling for coresets, in: *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 – November 3, 2022*, IEEE, 2022, pp. 462–473.

- [68] V. Braverman, V. Cohen-Addad, H.-C.S. Jiang, R. Krauthgamer, C. Schlegelshohn, M.B. Toft, X. Wu, The power of uniform sampling for coresets, in: 2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2022, pp. 462–473.
- [69] L. Huang, S.H. Jiang, J. Li, X. Wu, Coresets for clustering with general assignment constraints, CoRR, arXiv:2301.08460 [abs], 2023, <https://doi.org/10.48550/arXiv.2301.08460>.
- [70] P. Indyk, Sublinear time algorithms for metric space problems, in: Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing, 1999, pp. 428–434.
- [71] H.W. Lenstra, Integer programming with a fixed number of variables, Math. Oper. Res. 8 (4) (1983) 538–548, <https://doi.org/10.1287/moor.8.4.538>.
- [72] R. Kannan, Minkowski's convex body theorem and integer programming, Math. Oper. Res. 12 (3) (1987) 415–440, <https://doi.org/10.1287/moor.12.3.415>.
- [73] A. Frank, É. Tardos, An application of simultaneous Diophantine approximation in combinatorial optimization, Combinatorica 7 (1) (1987) 49–65, <https://doi.org/10.1007/bf02579200>.
- [74] R.K. Ahuja, J.B. Orlin, C. Stein, R.E. Tarjan, Improved Algorithms for Bipartite Network Flow, 1994.
- [75] C. Boutsidis, P. Drineas, M. Magdon-Ismail, Near-optimal column-based matrix reconstruction, Comput. Res. Repos. - CORR 43 (2011), <https://doi.org/10.1109/FOCS.2011.21>.
- [76] C. Boutsidis, A. Zouzias, M.W. Mahoney, P. Drineas, Randomized dimensionality reduction for k -means clustering, IEEE Trans. Inf. Theory 61 (2) (2014) 1045–1062.
- [77] V. Cohen-Addad, A. Gupta, A. Kumar, E. Lee, J. Li, Tight FPT approximations for k -median and k -means, in: C. Baier, I. Chatzigiannakis, P. Flocchini, S. Leonardi (Eds.), 46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9–12, 2019, Patras, Greece, in: LIPIcs, vol. 132, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019, 42.
- [78] T.F. Gonzalez, Clustering to minimize the maximum intercluster distance, Theor. Comput. Sci. 38 (1985) 293–306.
- [79] J.L. Bentley, J.B. Saxe, Decomposable searching problems I: static-to-dynamic transformation, J. Algorithms 1 (4) (1980) 301–358.
- [80] P.K. Agarwal, S. Har-Peled, K.R. Varadarajan, Approximating extent measures of points, J. ACM 51 (4) (2004) 606–635.
- [81] S. Har-Peled, S. Mazumdar, On coresets for k -means and k -median clustering, in: Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing, 2004, pp. 291–300.
- [82] M.R. Ackermann, M. Mörtens, C. Raupach, K. Swierkot, C. Lammersen, C. Sohler, Streamkm++ a clustering algorithm for data streams, ACM J. Exp. Algorithmics 17 (2012) 1–2.