

2	Modeling and Definitions	7
2.1	General Assignment Constraints	7
2.2	Coresets for Clustering with General Assignment Constraints	8
2.3	Handling Capacitated and Fair Clustering	8
3	Coresets for Clustering with General Assignment Constraints	9
3.1	Overview of the Proof of Theorem 3.1	11
3.1.1	Improved Hierarchical Uniform Sampling Framework	11
3.1.2	Handling Assignment Structure Constraints	13
3.2	The Coreset Construction Algorithm	13
3.3	Proof of Theorem 3.1: Performance Analysis of Algorithm 1	15
3.4	Proof of Lemma 3.3: Error Analysis for Rings	18
3.5	Proof of Lemma 3.4: Error Analysis for Groups	28
4	Bounding the Lipschitz Constant $\text{Lip}(\mathcal{B})$	32
4.1	Lipschitz Constant for (Laminar) Matroid Basis Polytopes	33
4.1.1	Lipschitz Constant for Matroid Basis Polytopes	36
4.1.2	Lipschitz Constant for Laminar Matroid Basis Polytopes	39
4.2	Lipschitz Constant $\text{Lip}(\mathcal{B})$ can be Unbounded	40
5	Bounding the Covering Exponent $\Lambda_\varepsilon(\mathcal{X})$	41
A	Application of Theorem 3.1: Simultaneous Coresets for Multiple \mathcal{B}'s	46
B	Missing Proofs in Section 2	47
B.1	Proof of Claim 1: Capacitated Clustering	47
B.2	Proof of Claim 2: Fair Clustering	47
C	Proof of Lemma 3.7: Relation between Two Covering Notions	49

1 Introduction

We study coresets for clustering with general assignment constraints. Clustering is a fundamental data analysis task that receives significant attention in various areas. In the (center-based) clustering problem, given as input a metric space (\mathcal{X}, d) and a finite point set $P \subseteq \mathcal{X}$, the clustering cost is defined for a set of clustering centers C of k points from \mathcal{X} , and an assignment function $\sigma : P \times C \rightarrow \mathbb{R}_+$ that assigns each data point p to the centers (such that $\|\sigma(p, \cdot)\|_1 = 1$ which ensures the point p is fully assigned¹), as an ℓ_z aggregation ($z \geq 1$) of the distances from the data points to the centers C weighted by σ , i.e.,

$$(1.1) \quad \text{cost}_z^\sigma(P, C) := \sum_{x \in P} \sum_{c \in C} \sigma(x, c) \cdot (d(x, c))^z.$$

The goal of the clustering problem is to find a center set C and assignment σ that minimizes the clustering cost. If there is no additional constraint on σ , then it is optimal to assign the data points to the (unique) nearest center, which is the case of vanilla (k, z) -CLUSTERING. However, the general formulation above allows a point to be assigned fractionally to multiple centers (which is sometimes called soft clustering) if we impose some additional constraint on σ .

General Assignment Constraints In this paper, we introducing the *general assignment constraints*, which impose various constraints on the assignment function σ . The formal definition can be found in Section 2. The general assignment constraint is very general and unifies several important clustering problems as follows:

1. Capacitated clustering: The problem is an important class of constrained clustering problems and has been studied extensively (see e.g., [CGTS02, LSS12, CHK12, Li16]). In this problem, we impose a *capacity upper bound* on every center $c \in C$, i.e., $\|\sigma(\cdot, c)\|_1 \leq u_c$.

¹Here, for some function of the form $\sigma : X \times Y \rightarrow \mathbb{R}_{\geq 0}$, we write $\sigma(x, \cdot)$ as the vector $u \in \mathbb{R}^Y$ such that $\forall y \in Y, u_y = \sigma(x, y)$. The notation $\sigma(\cdot, y)$ is defined similarly.

2. Fair clustering: Recently, a group-membership fairness definition, called fair clustering, has received significant interest [CKLV17, BCFN19, BIO⁺19]. Each point has a color indicating its group membership, and we need to find an assignment subject to the constraint that each cluster has each color represented within some pre-specified proportions.
3. Clustering with outliers: In some application domains, some data points may be regarded as outliers and does not need to be clustered. The problem has also been studied extensively for various objectives [CKMN01, Che08, GKL⁺17, SRF20, BK18, MNP⁺14]. This can be captured by imposing total capacity constraints $\|\sigma(\cdot, \cdot)\|_1 = n - m$ (i.e., the number of outliers is m) and $\|\sigma(p, \cdot)\|_1 \leq 1$ for every point p .
4. Fault-tolerant clustering: Another significant extension of clustering is the *fault tolerant* clustering problem, which requires each point to be assigned to at least l centers ($l \geq 1$) (see e.g., [KPS00, SS08, HHL⁺16]). This can be captured by the *assignment structure constraints* which requires that the vector $\sigma(p, \cdot)$ lies in a constrained convex set $\mathcal{B} = \{x \in \Delta_k : x_i \leq \frac{1}{l}, \forall i \in [k]\}$.

Coresets We focus on coresets for constrained clustering with general assignment constraints. Roughly speaking, an ε -coreset S is a tiny proxy of the data set P , such that the cost evaluated on both S and P are within $(1 \pm \varepsilon)$ factor for all potential centers C and assignments satisfying assignment constraints. The coreset is a powerful technique that can be used to compress large datasets, speed up existing algorithms, and design efficient approximation schemes [CL19, BFS21, BJKW21a]. The concept of coreset has also found many applications in modern sublinear models, such as streaming algorithms [HM04], distributed algorithms [BEL13] and dynamic algorithms [HK20], by combining it with the so-called merge-and-reduce framework [HM04].

The study of coreset size bounds for clustering has been very fruitful, especially the case of vanilla (k, z) -CLUSTERING (i.e., without constraints). A series of works focus on the Euclidean spaces [HM04, HK07, FL11, FSS20, SW18, HV20, CSS21, CLSS22] and near-optimal size bounds have been obtained in more recent works for k -MEDIAN and k -MEANS [CSS21, CLSS22, CALS⁺22, HLW24]. Recently, coreset size bounds in small dimensional Euclidean spaces have also been investigated [HHHW23]. Going beyond Euclidean spaces, another series of works provide coresets of small size in other types of metrics, such as doubling metrics [HJLW18, CSS21] and graph shortest-path metrics [BBH⁺20, BJKW21a]. More interestingly, throughout the line of research, various fundamental techniques have been proposed, including importance sampling [FL11, FSS20] which can be applied to several problems including clustering, and a more recently developed hierarchical sampling framework [Che09, CSS21, BCJ⁺22] that employs uniform sampling in a novel way.

Coreset for Constrained Clustering Unfortunately, coresets for constrained clustering has been less understood. In particular, only the capacity constraints were considered, and the research focus was on coresets for fair clustering [CKLV17] and capacitated clustering [CGTS02]. Earlier works [SSS19, HJV19, CL19, BFS21] achieved coresets of size either depending on n or exponential in d (which is the Euclidean dimension). Recently, a breakthrough was made in [BCJ⁺22] where the first coresets for both fair and capacitated k -MEDIAN in Euclidean \mathbb{R}^d with size $\text{poly}(k\varepsilon^{-1})$ were obtained, via an improved hierarchical uniform sampling framework. The framework has also been adapted to the outlier setting in a more recent work [HJLW23]. This framework certainly provides a good starting point, but several fundamental issues still remain. One issue is that coresets obtained through this framework are still somewhat ad-hoc, and it is unclear if the result can be adapted to more general assignment constraints such as the aforementioned structure constraints, and/or other metrics such as graph shortest-path metrics. Indeed, a perhaps more fundamental question is that, a systematic characterization of what types of assignment constraints allow small coresets, is still missing in the literature. In addition, the framework and analysis in [BCJ⁺22] only lead to $\text{poly}(k\varepsilon^{-1})$ size bound with high degree polynomial, which is also sub-optimal.

1.1 Our Contributions Our main contribution is two-fold. (1) We propose a very general model of assignment constraints (including capacity constraint, outliers constraint, and the aforementioned structure constraint), and provide a characterization of families of assignment constraints that admit small coresets. (2) Our new analysis leads to improved coreset size upper bounds, even for the important cases of fair/capacitated clustering and clustering with outliers (without additional structure constraints), achieving state-of-the-art bounds for these problems. Next, we discuss our contributions in more detail.

A General Model of Assignment Constraints Our new model for the assignment constraints, called the *general assignment constraints*, is a combination of three types of constraints on the assignment function

$\sigma(\cdot, \cdot)$: (1) the *assignment structure constraint* (see Definition 2.3) which is a new notion proposed in this work, (2) the standard *capacity constraint* (Definition 2.1) that constrains the total weight assigned to each center, and (3) the *total capacity constraint* (Definition 2.2), which can be used to capture the number of outliers. The new structure constraint (Definition 2.3) specifies a convex body $\mathcal{B} \subseteq \Delta_k$ where $\Delta_k := \{x \in \mathbb{R}_{\geq 0}^k : \|x\|_1 = 1\}$ is the k -dimensional simplex, and it requires that for every point p the assignment vector $\sigma(p, \cdot)$ must lie in \mathcal{B} . Note that the case $\mathcal{B} = \Delta_k$ corresponding to the standard constraint that a point is fully assigned to the centers. We focus on the well-known cases for convex body \mathcal{B} , such as matroid basis polytopes and knapsack polytopes. Indeed, these types of \mathcal{B} are already general enough to capture many of the above-mentioned constraints, including the fault-tolerance constraint (in which case $\mathcal{B} = \{x \in \Delta_k : x_i \leq \frac{1}{l}, \forall i \in [k]\}$). See Definition 2.3 and the subsequent discussions. In addition, to capture the outliers in clustering, we introduce the *total capacity constraint* which requires $\|\sigma\|_1 = n - m$ (i.e., the number of outliers is m).

Main Theorem In our main theorem (Theorem 1.1), we show that a small coresets exists, as long as the structure constraints have bounded complexity $\text{Lip}(\mathcal{B}) = O_k(1)$ (see Definitions 3.2 and 3.3), which only depends on the convex body \mathcal{B} , and the *covering exponent* $\Lambda_\varepsilon(\mathcal{X})$ (see Definition 3.4) can be bounded by a number that is independent of n . Hence, the main theorem systematically reduces the problem of constructing coresets of size $\text{poly}(k\varepsilon^{-1})$, to the mathematical problems of bounding the parameters $\text{Lip}(\mathcal{B})$ and $\Lambda_\varepsilon(\mathcal{X})$. The parameter $\text{Lip}(\mathcal{B})$ is new, and is defined as the Lipschitz constant of a certain transportation procedure inside the convex body \mathcal{B} . On the other hand, the covering exponent $\Lambda_\varepsilon(\mathcal{X})$ (also known as the log covering number or the metric entropy) of metric space (\mathcal{X}, d) is closely related to several combinatorial dimension notions, such as the VC-dimension and the (fat) shattering dimension of the set system formed by all metric balls which have been extensively studied in previous works, e.g. [LS10, FL11, FSS20, HJLW18, BJKW21a]. Bounds for $\Lambda_\varepsilon(\mathcal{X})$ are known for multiple metric spaces, including Euclidean metrics, doubling metrics, general discrete metrics, and shortest-path metrics (see Remark 1). Now, we state our main theorem.

THEOREM 1.1. (Informal; see Theorem 3.1) *We consider (k, z) -CLUSTERING with capacity upper/lower bound constraint for each center, assignment structure constraint for each point (specified by convex body $\mathcal{B} \subseteq \Delta_k$), and a total capacity constraint $\|\sigma\|_1 = n - m$ (i.e., m outliers). For any $0 < \varepsilon < 1$, there is a near linear-time algorithm that computes an ε -coreset of size $O(m) + \tilde{O}_z(\text{Lip}(\mathcal{B})^2 \cdot (\Lambda_\varepsilon(\mathcal{X}) + k + \varepsilon^{-1}) \cdot k^2 \varepsilon^{-2z})$.²*

Moreover, if there is no additional structure constraint (i.e., $\mathcal{B} = \Delta_k$), the coresets size bound can be improved to $O(m) + \tilde{O}_z((\Lambda_\varepsilon(\mathcal{X}) + \varepsilon^{-1}) \cdot k^2 \varepsilon^{-2z})$.

We provide an overview of the proof in Section 3.1, and the complete proof appears in Section 3. The theorem is general and provides improved coresets size bounds for a variety of constrained clustering problems. Unlike many previous works that use ad-hoc methods to deal with constraints in specific metric spaces (such as Euclidean \mathbb{R}^d) [SSS19, HJV19, CL19, BFS21, BCJ⁺22], we completely decouple the parameter of constraints $\text{Lip}(\mathcal{B})$ and the complexity $\Lambda_\varepsilon(\mathcal{X})$ of the metric space, so that they can be dealt with independently. Moreover, our coresets size bound is optimal (up to constant factor) in the dependence of m , due to an $\Omega(m)$ lower bound for coresets for clustering with m outliers [HJLW23], and the dependence in k and ε improves over the bounds in [BCJ⁺22] by a factor of $k\varepsilon^{-z}$ (which works for only the capacity constraint). Hence, even without any constraint on \mathcal{B} , i.e., $\mathcal{B} = \Delta_k$ where we have $\text{Lip}(\mathcal{B}) = 1$, and only considering the capacity constraints, we can already obtain several new/improved coresets results, by simply using known bounds of $\Lambda_\varepsilon(\mathcal{X})$ (see Remark 1). We list more concrete implications of our general theorem, which can be found at the end of this section and Table 1.

New Coresets Results and Improved Bounds We list several notable results followed by our main theorem (Theorem 1.1), using new or existing upper bounds on $\text{Lip}(\mathcal{B})$ and $\Lambda_\varepsilon(\mathcal{X})$. Also, see Table 1 for a summary.

- Coresets for fair/capacitated clustering with outliers under various metrics. This corresponds to the case $\mathcal{B} = \Delta_k$ and $\text{Lip}(\mathcal{B}) = 1$. We take capacitated clustering as an example. In the context of fair clustering, the coresets size includes an additional factor Γ compared to that of capacitated clustering. For metric spaces with bounded doubling dimension, shortest-path metrics of planar graphs, or more generally graphs that exclude a fixed minor, the covering exponent $\Lambda_\varepsilon(\mathcal{X})$ can also be bounded independent of n (see Remark 1). Hence, we obtain the first $O(m + \text{poly}(k\varepsilon^{-1}))$ -sized coresets of capacitated k -MEDIAN with m outliers under the

²Throughout, the notation $\tilde{O}_z(f)$ hides factors $\text{poly} \log f$ and $2^{\text{poly}(z)}$

Variants of constrained clustering in \mathbb{R}^d		Prior results	Our results
(k, z) -CLUSTERING	with outliers	$O(m) + \tilde{O}(k^3 \varepsilon^{-3z-2})$ [HHHW23]	$O(m) + \tilde{O}(k^2 \varepsilon^{-2z-2})$
Capacitated (k, z) -CLUSTERING	without outliers	$O(k^2 \varepsilon^{-3} \log^2 n)$ ($z = 1$) [CL19] $\tilde{O}(k^3 \varepsilon^{-6})$ ($z = 1$) [BCJ ⁺ 22] $O(k^3 \varepsilon^{-3} \log^5 n)$ ($z = 2$) [CL19] $\text{poly}(k, \varepsilon^{-1})$ ($z = 2$) [BCJ ⁺ 22]	$\tilde{O}(k^2 \varepsilon^{-2z-2})$
	with outliers	/	$O(m) + \tilde{O}(k^2 \varepsilon^{-2z-2})$
Fair (k, z) -CLUSTERING	without outliers	$O(\Gamma k^2 \varepsilon^{-3} \log^2 n)$ ($z = 1$) [CL19] $\tilde{O}(\Gamma k^3 \varepsilon^{-6})$ ($z = 1$) [BCJ ⁺ 22] $O(\Gamma k^5 \varepsilon^{-3} \log^5 n)$ ($z = 2$) [CL19] $\Gamma \cdot \text{poly}(k, \varepsilon^{-1})$ ($z = 2$) [BCJ ⁺ 22]	$\tilde{O}(k^2 \varepsilon^{-2z-2})$
	with outliers	/	$O(\Gamma m) + \tilde{O}(\Gamma k^2 \varepsilon^{-2z-2})$
Fault-tolerant (k, z) -CLUSTERING	without outliers	/	$\tilde{O}(k^2 \varepsilon^{-2z}(k + \varepsilon^{-2}))$
	with outliers	/	$O(m) + \tilde{O}(k^2 \varepsilon^{-2z}(k + \varepsilon^{-2}))$

Table 1: Comparison of the state-of-the-art coresets sizes and our results for different variants of constrained (k, z) -CLUSTERING in Euclidean spaces \mathbb{R}^d . The factor Γ in fair clustering denotes the number of distinct collections of groups that a point may belong to; also see Claim 1. We assume that $z \geq 1$ is a constant and ignore $2^{O(z)}$ or $z^{O(z)}$ factors in the coreset size, and also ignore the logarithmic terms in $\tilde{O}(\cdot)$.

above metric spaces. Previously, for capacitated clustering, even without outliers, coresets of size $\text{poly}(k\varepsilon^{-1})$ are known only for Euclidean k -MEDIAN [BCJ⁺22]³ and our size bound is already better by a factor of $k\varepsilon^{-2}$ for this special case (e.g., the bound in [BCJ⁺22] for capacitated k -MEDIAN is $\tilde{O}(k^3 \varepsilon^{-6})$ and our bound is $\tilde{O}(k^2 \varepsilon^{-4})$). Our result also generalizes the recent work [HJLW23] which provides coresets for clustering with outliers (but cannot handle e.g., fairness constraints). Our bound also achieves a tight linear dependence in m , and the other term is a factor of $k\varepsilon^{-z}$ better (the bound in [HJLW23] is $O(m) + \tilde{O}(k^3 \varepsilon^{-3z-2})$ and our new bound is $O(m) + \tilde{O}(k^2 \varepsilon^{-2z-2})$ in this setting).

- **Coresets for fault-tolerant clustering.** In fault-tolerant clustering, we require each point to be assigned to at least $l \geq 1$ centers. In this case, $\mathcal{B} = \{x \in \Delta_k : x_i \leq 1/l, \forall i \in [k]\}$ is a (scaled) uniform matroid basis polytope. By Theorem 4.2, $\text{Lip}(\mathcal{B})$ is bounded by 2 since uniform matroid is a laminar matroid of depth 1. Hence, we obtain the first coreset of size $\text{poly}(k\varepsilon^{-1})$ for the fault-tolerant k -median in Euclidean space (and other metric spaces with bounded covering exponents).
- **Coresets for clustering with more general fault-tolerance requirements.** In the variant of clustering defined in [BDHR05], suppose the points are partitioned into several groups based on some geographical regions. The goal is to choose k centers subject to the constraint that k_i center are chosen from the i -th group ($\sum_i k_i = k$). In addition, it is required that each point is connected to centers in at least l different groups. For this variant, \mathcal{B} corresponds to a Laminar matroid basis polytope of depth 2 (see the discussion after Definition 2.3), and by Theorem 4.2, $\text{Lip}(\mathcal{B})$ is bounded by 3. Hence, we obtain the the first coreset of size $\text{poly}(k\varepsilon^{-1})$ for this clustering problem.
- **Simultaneous coresets.** As another corollary of the main theorem, we obtain coresets that hold simultaneously for a set of m structure constraints, which only requires enlarging the coreset by a $\log m$ factor (see Appendix A). This is particularly useful when the parameter \mathcal{B} is to be picked from a family that is not known in advance and is subject to experiment. In this scenario, the same coreset can be re-used, which avoids recomputing a new coreset every time a new \mathcal{B} is tested.

Upper Bound of $\text{Lip}(\mathcal{B})$ when \mathcal{B} is Matroid Basis Polytope In light of Theorem 1.1, one can see that proving upper bound of $\text{Lip}(\mathcal{B})$ is crucial for bounding the coreset size. In fact, one can easily construct polytope \mathcal{B} , which is simply defined by some linear constraints, such that $\text{Lip}(\mathcal{B})$ is unbounded (See Section 4). In this paper, we focus on an important class of polytopes, called matroid basis polytopes (see e.g., [Sch03]). A matroid basis polytope is the convex hull of all 0/1 indicator vectors of the basis of a matroid. Matroids generalize many combinatorial structures and are popular ways to model the structure of an

³We remark that Braverman et al. [BCJ⁺22] also provided coresets for various metric spaces, but only for vanilla clustering. For constrained clustering (such as fair/capacitated clustering), only Euclidean \mathbb{R}^d was considered and it turns out to be nontrivial to generalize to other metrics. See Section 3.1 for a more detailed discussion.

assignment/solution in various contexts, such as online computation [BIK07, BIKK18], matching [Law01, LSV13], diversity maximization [AMT13] and variants of clustering problems [CLLW16, KKN⁺11, KKN⁺15, Swa16]. For a general matroid basis polytope \mathcal{B} , we prove $\text{Lip}(\mathcal{B}) \leq k - 1$ (see Theorem 4.1). We also provide an improved $\text{Lip}(\mathcal{B}) \leq \ell + 1$ bound for the special case of laminar matroids of depth ℓ (see Theorem 4.2). This readily implies coresets of size $\text{poly}(k\varepsilon^{-1})$ for any general assignment constraints with a matroid basis polytope \mathcal{B} , under various types of metric spaces.

1.2 Related Work Approximation algorithms have been extensively studied for constrained clustering problems. We focus on the k -median case for several notable problems in the following discussion. For fair k -median, [BCFN19] provided a bi-criteria $O(1)$ -approximation in general metric spaces, but the solution may violate the capacity/fair constraint by an additive error. [BIO⁺19] gave $O(\log n)$ -approximation without violating the constraints in Euclidean spaces. For capacitated k -median, $O(1)$ -approximation were known in general metrics, but they either need to violate the capacity constraint [DL16, BRU16] or the number of centers k [Li17, Li16] by a $(1 + \varepsilon)$ factor. Both problems admit polynomial-time algorithms that have better approximation and/or no violation of constraints when k is not considered a part of the input [ABM⁺19, CL19, FZH⁺20, BFS21]. For fault-tolerant k -median, a constant approximation was given in [SS08], and $O(1)$ -approximation also exists even when the number of centers that each data point needs to connect can be different [HHL⁺16]. Finally, we mention a variant of clustering called matroid k -median which also admit $O(1)$ -approximation in general metrics [KKN⁺11]. In this problem, a matroid is defined on the vertices of the graph, and only the center set that form independent sets may be chosen. While this sounds different from the constraints that we consider, our coreset actually captures this trivially since our coreset preserves the cost for all centers (not only those that form independent sets).

Apart from constrained clustering, coresets were also considered for clustering with generalized objectives (but without constraints). Examples include projective clustering [FL11, FSS20, TWZ⁺22], clustering with missing values [BJKW21b], ordered-weighted clustering [BJKW19] and clustering with panel data [HSV21].

1.3 Technical Overview Our approach is a generalization and improvement over a recent hierarchical uniform sampling framework developed in [BCJ⁺22]. Our contribution is two-fold: 1) We improve the coreset size of the framework of [BCJ⁺22], even without the new assignment structure constraints (i.e., in the same setting as in [BCJ⁺22]), and this is achieved by employing a new adaptive sampling step; 2) We incorporate the additional assignment structure constraints into the framework, which can handle more general clustering problems such as fault-tolerant clustering.

Size Improvement At a high level, the framework of [BCJ⁺22] decomposes the dataset P into disjoint rings R , and takes a uniform sample S_R on every ring with a *uniform* size. This uniform size bound on rings directly affects the size of the coreset, and our idea is to improve this sample size for rings. As observed in [BCJ⁺22], this simple uniform sampling is very powerful and can preserve the coreset error $|\text{cost}_z(R, C, \Delta_k, h) - \text{cost}_z(S_R, C, \Delta_k, h)|$ incurred on the ring R , for every center set $C \subset \mathbb{R}^d$ and every capacity constraint h , by charging to a certain additive error $\text{err}(R)$ that only depends on R (Inequality (3.6)). This charging is worst-case optimal over the choice of C , and hence, we cannot expect to improve the error analysis for a single ring. However, we find that such additive error $\text{err}(R)$ is only incurred when the ring R is “close enough” to center set C , while the number of such rings is always small (say $\tilde{O}(k)$) for every choice of C (Lemma 3.2). This novel geometric observation enables us to adaptively tune the sample size for each ring, which leads to a bounded total error of rings (Lemma 3.3) and significantly improves the coreset size.

Handling Assignment Structure Constraints To handle the assignment structure constraints, the main technical step is still to bound the coreset error $|\text{cost}_z(R, C, \mathcal{B}, h) - \text{cost}_z(S_R, C, \mathcal{B}, h)|$ incurred on a ring R . As a central step, we need to show it is possible to modify the optimal assignment from R to C , to an assignment from S_R to C , with small additional cost subject to the constraint \mathcal{B} . We reduce the problem of bounding the extra cost of such conversion to a so-called *optimal assignment transportation* (OAT) problem, which aims to bound the total transportation cost from a given assignment σ to any assignment σ' that is consistent with (\mathcal{B}, h') . We define a new notion $\text{Lip}(\mathcal{B})$ as the universal upper bound of the OAT cost, and prove nontrivial upper bounds of $\text{Lip}(\mathcal{B})$ for various polytopes, including the general matroid basis polytopes (Theorem 4.1), and laminar matroid polytopes with a better bound (Theorem 4.2). These bounds imply that our algorithm produces coreset for fault-tolerant clustering and even more general assignment structure constraints. Bounding $\text{Lip}(\mathcal{B})$ for other convex

set \mathcal{B} , as well as the efficient computation of it (which we do not need) may be of independent interest for future research.

2 Modeling and Definitions

Throughout, we are given an underlying metric space (\mathcal{X}, d) , integer $k \geq 1$, constant $z \geq 1$, and precision parameter $\varepsilon \in (0, 1)$. For integer $n \geq 1$, let $[n] := \{1, \dots, n\}$. For a function $\sigma : X \times Y \rightarrow \mathbb{R}$, for $i \in X$ we write $\sigma(i, \cdot)$ as the vector $u_i \in \mathbb{R}^Y$ such that $u_i(j) := \sigma(i, j)$, and define $\sigma(\cdot, j)$ similarly. We say σ is an assignment function if $\|\sigma(i, \cdot)\|_1 \leq 1$ for every point $i \in X$. Sometimes, we also interpret σ as a vector in $\mathbb{R}^{X \times Y}$, so $\|\sigma\|_1 = \sum_{i \in X, j \in Y} |\sigma(i, j)|$.⁴ We use Δ_k to denote the simplex $\{x \in \mathbb{R}_{\geq 0}^k : \|x\|_1 = 1\}$. Given a point $a \in \mathcal{X}$ and a radius $r > 0$, we define $\text{Ball}(a, r) := \{x \in \mathcal{X}, d(x, a) \leq r\}$ to be the ball of radius r centered at a . Moreover, for two positive real numbers $r_1, r_2 > 0$, define $\text{ring}(a, r_1, r_2) := \text{Ball}(a, r_2) \setminus \text{Ball}(a, r_1)$. Throughout this paper, we assume there exists an oracle that answers $d(p, q)$ in $O(1)$ time for any $p, q \in \mathcal{X}$.

2.1 General Assignment Constraints We consider three types of assignment constraints in clustering literature: (1) *capacity constraints* on cluster centers and (2) *assignment structure constraints* on $\sigma(p, \cdot)$ for points p , and (3) *total capacity constraints* which can capture outliers in clustering.

First, we model the capacity constraints in a way similar to previous works [HJV19, CL19, BFS21, BCJ⁺22]. We simply consider a vector $h \in \mathbb{R}_{\geq 0}^k$ (k is the number of centers), and require that the mass assigned to each center c equals h_c .

DEFINITION 2.1. (Capacity constraint) *Given a set C of k centers, a capacity constraint can be specified by a vector $h \in \mathbb{R}_{\geq 0}^k$. We say an assignment function $\sigma : P \times C \rightarrow \mathbb{R}_{\geq 0}$ is consistent with h , denoted as $\sigma \sim h$, if $\|\sigma(\cdot, c)\|_1 = h_c$ for every center $c \in C$, which means the total assignment to a center c is exactly h_c .⁵ Equivalently, we can write $\sum_p \sigma(p, \cdot) = h$ (both sides are k -dimensional vectors).*

The above capacity constraint consists of equality constraints, which seem different from inequality constraints for capacitated clustering or fair clustering. However, we can show that capacitated clustering and fair clustering can be captured by such equality constraint; see Section 2.3 (Claims 1 and 2) for the detailed reductions. In addition, we further allow the total capacity to be less than the total weight n , and this is useful for dealing with outliers in clustering.

DEFINITION 2.2. (Total capacity constraint) *Given an integer $0 \leq m \leq n$, we can impose a total capacity constraint of the form $\|\sigma\|_1 = \|h\|_1 = n - m$. In an integral assignment, the total capacity constraint says that we can exclude m points as outliers.*

In fact, if the capacity vector $h \in \mathbb{R}_{\geq 0}^k$ is given, the total capacity constraint is already determined by h . Due to the special meaning of m (i.e., the number of outliers), we make this parameter explicit and we will introduce new ideas to analyze the $m > 0$ case in Lemmas 3.3 and 3.4.

On the other hand, an assignment structure constraint concerns the range of the assignment vector $\sigma(p, \cdot)$ for each point p . We model an assignment structure constraint by a convex body $\mathcal{B} \subseteq \Delta_k$, where $\Delta_k := \{x \in \mathbb{R}_{\geq 0}^k : \sum_{i \in [k]} x_i = 1\}$ denotes the simplex in \mathbb{R}^k .

DEFINITION 2.3. (Assignment structure constraint) *Given convex body $\mathcal{B} \subseteq \Delta_k$, we say an assignment function $\sigma : P \times C \rightarrow \mathbb{R}_{\geq 0}$ is consistent with \mathcal{B} , denoted as $\sigma \sim \mathcal{B}$, if*

$$\sigma(p, \cdot) \in \mathcal{B} \text{ for every } p \in P.$$

We define \mathcal{B}° as $\mathcal{B} \cup \{0\}$ (0 is used to capture the assignment for the outlier). Given that $\|\sigma(p, \cdot)\|_1 \leq 1$, we can infer that $\sigma(p, \cdot) \in \text{conv}(\mathcal{B}^\circ)$ as per the definition above.

We list a few examples of assignment structure constraints as follows.

⁴Here, we abuse the notation by using $\sigma \in \mathbb{R}^{X \times Y}$ to represent a vector instead of a matrix such that $\|\sigma\|_1$ is well defined.

⁵We require the capacity to be exactly h_c instead of placing a lower and/or upper bound, since we would preserve the cost for all h simultaneously in our coresets. See Definition 2.4.

1. $\mathcal{B} = \{x \in \Delta_k : x_i \leq \frac{1}{l}, \forall i \in [k]\}$ for some integer $l \in [k]$. If $l = 1$, $\mathcal{B} = \Delta_k$, and in this case a point is assigned to its nearest center. If $l > 1$, the cheapest way to assign p is to connect it to the l nearest centers, which captures fault-tolerant clustering. Such fault-tolerant constraints have been studied extensively in a variety of clustering problems [KPS00, SS08, HHL⁺16].
2. \mathcal{B} is a (scaled) matroid basis polytope. This is a significant generalization of the above constraint which corresponds to a uniform matroid polytope. Motivated by the clustering variants studied in [BDHR05, ZG03, BRS08], we consider more advanced fault-tolerance requirements, which can be captured by *laminar matroid* basis polytope. For example, suppose P_1, \dots, P_g is a partition of $[k]$ such that $|P_j| = k_j$ (thus $\sum_j k_j = k$). Consider the partition matroid polytope $\mathcal{B} = \left\{x \in \Delta_k : \sum_{i \in P_j} x_i \leq \frac{1}{l}, \forall j \in [g], \sum_i x_i = 1\right\}$. This captures the clustering problem with advanced fault-tolerant constraints.

General Assignment Constraints: In this paper, we consider the case where all points p are subject to the same assignment structure constraint \mathcal{B} . When an assignment function σ satisfies both capacity constraint h and assignment structure constraint \mathcal{B} , we denote it as $\sigma \sim (\mathcal{B}, h)$, and call it a *general assignment constraint*. Given a general assignment constraint (\mathcal{B}, h) , we define the cost of center set C for point set P as

$$(2.2) \quad \text{cost}_z(P, C, \mathcal{B}, h) := \min_{\sigma \sim (\mathcal{B}, h)} \text{cost}_z^\sigma(P, C).$$

That is, for the center set C , the cost is computed via the min-cost assignment σ that is consistent with both constraints \mathcal{B} and h .

2.2 Coresets for Clustering with General Assignment Constraints In our paper, a coreset $S \subset P$ is a weighted subset with weight $w_S(\cdot) : S \rightarrow \mathbb{R}_{\geq 0}$. In particular, the weight should satisfy $w_S(S) := \sum_{p \in S} w_S(p) = n$. We need to extend the general assignment constraint to handle weighted points. In particular, for a point set S with weight $w_S(\cdot)$, if we require a point p fully assigned, the corresponding constraint is $\|\sigma(p, \cdot)\|_1 = w_S(p)$. For a center set C and a capacity constraint h , the capacity and total capacity constraints are the same: we still require that $\sum_p \sigma(p, \cdot) = h$ and $\|\sigma\|_1 = \|h\|_1 = n - m$. For the assignment structure constraint, given convex body $\mathcal{B} \subseteq \Delta_k$, we say an assignment function $\sigma : P \times C \rightarrow \mathbb{R}_{\geq 0}$ is consistent with \mathcal{B} , denoted as $\sigma \sim \mathcal{B}$, if $\sigma(p, \cdot) \in w_S(p) \cdot \mathcal{B}$ for every $p \in S$ (note that this implies $\|\sigma(p, \cdot)\|_1 \leq w_S(p)$).

Now, we provide the definition of coreset for our problem.

DEFINITION 2.4. (Coreset) For a dataset $P \subseteq \mathcal{X}$ of size n , an outlier number $0 \leq m \leq n$ and an assignment structure constraint \mathcal{B} , an $(\varepsilon, \mathcal{B}, m)$ -coreset for clustering with general assignment constraints is a (weighted) set $S \subseteq P$ that satisfies

$$(2.3) \quad \forall C \in \mathcal{X}^k, h \in (n - m) \cdot \mathcal{B}, \quad \text{cost}_z(S, C, \mathcal{B}, h) \in (1 \pm \varepsilon) \cdot \text{cost}_z(P, C, \mathcal{B}, h),$$

where $\text{cost}_z(P, C, \mathcal{B}, h)$ is defined as in (2.2).

Note that for coreset S , the cost $\text{cost}_z(S, C, \mathcal{B}, h)$ is also computed according to (2.2), with the understanding of $\sigma \sim (\mathcal{B}, h)$ defined for weighted points. The coreset guarantee states that (2.3) should hold for *all* capacity vectors $h \in (n - m) \cdot \mathcal{B}$. Requiring $h \in (n - m) \cdot \mathcal{B}$ is necessary since we only need to focus on *feasible* capacity vector h for which there exists σ such that $\sigma \sim (h, \mathcal{B})$.⁶

2.3 Handling Capacitated and Fair Clustering It is important to note that our notion of coreset preserves the cost for all centers C and all feasible capacity constraints h *simultaneously*, and for an assignment structure constraint \mathcal{B} given in advance. Hence, as also noted in previous works [BFS21, BCJ⁺22], the guarantee that the cost is preserved for all h simultaneously implies that such a coreset simultaneously captures all types of upper/lower bound capacity constraints of the form $\ell_c \leq \|\sigma(\cdot, c)\|_1 \leq u_c$; summarized by the following claim.

CLAIM 1. (Capacitated Clustering) Consider the (k, z) -CLUSTERING problem with capacity upper/lower bound constraint for each center c of the form $\ell_c \leq \|\sigma(\cdot, c)\|_1 \leq u_c$, assignment structure constraint $\mathcal{B} \subseteq \Delta_k$,

⁶To see this, note that we require $\sigma(p, \cdot) \in \frac{n-m}{n} \cdot \mathcal{B}$, hence $h = \sum_{p \in P} \sigma(p, \cdot) \in (n - m) \cdot \mathcal{B}$.

and a total capacity constraint $\|\sigma\|_1 = n - m$. An $(\varepsilon, \mathcal{B}, m)$ -coreset S is an ε -coreset for this constrained clustering problem, i.e., for any center set $C \in \mathcal{X}^k$,

$$\min_{\substack{\sigma: \|\sigma\|_1 = n - m, \sigma \sim \mathcal{B} \\ \ell_c \leq \|\sigma(\cdot, c)\|_1 \leq u_c, \forall c \in C}} \text{cost}_z^\sigma(S, C) \in (1 \pm \varepsilon) \cdot \min_{\substack{\sigma: \|\sigma\|_1 = n - m, \sigma \sim \mathcal{B} \\ \ell_c \leq \|\sigma(\cdot, c)\|_1 \leq u_c, \forall c \in C}} \text{cost}_z^\sigma(P, C).$$

Proof. The proof can be found in Appendix B.1. \square

This claim implies that we do not need to specify capacity upper/lower bounds in the definition of our coreset.

Similarly, the following claim shows that our coreset works for fair clustering via a reduction of [HJV19, Theorem 4.2]. Recall that in the fair clustering problem, there are s groups $G_1, \dots, G_s \subseteq P$ (groups can be non-disjoint). We constrain that the proportion of points from a group G_j contained in a cluster centered at c , $\frac{|\sigma^{-1}(c) \cap G_j|}{|\sigma^{-1}(c)|}$, has a group-wise upper bound u_i and lower bound ℓ_i with $0 \leq \ell_i \leq u_i \leq 1$.

CLAIM 2. (Fair Clustering) Let $P \subseteq \mathcal{X}$ be a dataset and $G_1, \dots, G_s \subseteq P$ be s groups (groups can be non-disjoint). We denote $\mathcal{G}_p = \{j \in [s] : p \in G_j\}$ to be the collection of groups that point $p \in P$ belongs to, and let $\Gamma := |\{\mathcal{G}_p : p \in P\}|$ denote the number of distinct \mathcal{G}_p 's. Consider the (k, z) -CLUSTERING problem with fairness upper/lower bound constraint for each group G_i and each center c of the form $\ell_i \leq \frac{|\sigma^{-1}(c) \cap G_i|}{|\sigma^{-1}(c)|} \leq u_i$, and a total capacity constraint $\|\sigma\|_1 = n - m$. Suppose there exists S of size at most A , which is an $(\varepsilon, \Delta_k, m')$ -coreset for clustering with general assignment constraints for all $0 \leq m' \leq m$. Then there exists an ε -coreset of size at most ΓA for this fair clustering problem with outliers, i.e., for any center set $C \in \mathcal{X}^k$,

$$\min_{\substack{\sigma: \|\sigma\|_1 = n - m \\ \ell_i \leq \|\sigma(\cdot, c)\|_1 \leq u_i, \forall i \in [s], c \in C}} \text{cost}_z^\sigma(S, C) \in (1 \pm \varepsilon) \cdot \min_{\substack{\sigma: \|\sigma\|_1 = n - m \\ \ell_i \leq \|\sigma(\cdot, c)\|_1 \leq u_i, \forall i \in [s], c \in C}} \text{cost}_z^\sigma(P, C).$$

Proof. The proof can be found in Appendix B.2. \square

Intuitively, we begin by partitioning the set P into Γ subsets P_1, \dots, P_Γ , where each subset P_j consists of points p sharing the same group label $\mathcal{G}_p = \mathcal{G}_j$. Let $h^{(j)}$ denote the capacity vector constrained to subset P_j . For every $i \in [s]$ and $c \in C$, we can establish that

$$|\sigma^{-1}(c)| = \sum_{j=1}^{\Gamma} h_c^{(j)} \quad \text{and} \quad |\sigma^{-1}(c) \cap G_i| = \sum_{j=1}^{\Gamma} \mathbb{I}[i \in \mathcal{G}_j] \cdot h_c^{(j)},$$

where $\mathbb{I}[\cdot]$ is the indicator function. Thus, the capacity vectors $h^{(1)}, \dots, h^{(\Gamma)}$ determine whether an assignment satisfies fairness constraints, allowing us to represent fairness constraints through a set of capacity constraints on $h^{(1)}, \dots, h^{(\Gamma)}$. This reduction approach introduces an additional factor Γ in the size of the coreset.

The claim indicates that the coreset size for fair clustering only contains an additional multiplicative factor Γ . Different from Claim 1, we do not handle assignment structure constraints and fairness constraints simultaneously since the reduction of [HJV19, Theorem 4.2] relies on a decomposition of dataset P into Γ pieces while it is unknown how to decompose assignment constraints accordingly. By our coreset construction (Algorithm 1), if S is an $(\varepsilon, \Delta_k, m)$ -coreset, then S is always an $(\varepsilon, \Delta_k, m')$ -coreset for all $0 \leq m' \leq m$. Thus, combining with our coreset results for capacitated clustering, Claim 2 leads to the results for fair clustering in Table 1.

Throughout, our goal is to obtain an $(\varepsilon, \mathcal{B}, m)$ -coreset for fixed \mathcal{B} and m . If the context is clear, we also use the shorthand ε -coreset in replacement of $(\varepsilon, \mathcal{B}, m)$ -coreset.

3 Coresets for Clustering with General Assignment Constraints

Before stating our main theorem, we first introduce a special case, the (k, z) -CLUSTERING with m outliers problem. The (k, z) -CLUSTERING with m outliers problem is the (k, z) -CLUSTERING problem with a total capacity constraint $\|\sigma\|_1 = n - m$, and the goal is to find a center set $C^* \in \mathcal{X}^k$ and an assignment function $\sigma^* : P \times C \rightarrow \mathbb{R}_{\geq 0}$ with $\|\sigma^*\|_1 = n - m$ that solve the following problem:

$$\min_{C \in \mathcal{X}^k, \sigma: \|\sigma\|_1 = n - m} \text{cost}_z^\sigma(P, C).$$

The problem reduces to vanilla (k, z) -CLUSTERING when $m = 0$. Similar to prior studies on this problem [BCJ⁺22, HJLW23], we need to introduce a tri-criteria approximation algorithm for (k, z) -CLUSTERING with m outliers for constructing coresets.

DEFINITION 3.1. ((α, β, γ) -Approximation for (k, z) -Clustering with m outliers) Let $P \subset \mathbb{R}^d$ be a dataset and $\alpha, \beta, \gamma \geq 1$ be constants. An (α, β, γ) -approximation of P for (k, z) -CLUSTERING with m outliers is a center set $C^* \subset \mathbb{R}^d$ of size at most βk such that $\text{cost}_z^{(\gamma m)}(P, C) \leq \alpha \cdot \text{OPT}_{k,z,m}(P)$, where $\text{cost}_z^{(m)}(P, C) := \min_{\sigma: \|\sigma\|_1 = n-m} \text{cost}_z^\sigma(P, C)$ and $\text{OPT}_{k,z,m}$ denotes the optimal value of (k, z) -CLUSTERING with m outliers.

Please refer to [HJLW23, Appendix A] for more discussions of such approximation algorithms. For instance, a $(2^{O(z)}, O(1), O(1))$ -approximation can be constructed in near-linear time [BVX19].

Now, we are ready to state our main theorem.

THEOREM 3.1. (Coresets for clustering with general assignment constraints) Let (\mathcal{X}, d) be a metric space, $k \geq 1, m \geq 0$ be integers, and $z \geq 1$ be a constant. Let $\varepsilon, \delta \in (0, 1)$ and $\mathcal{B} \subseteq \Delta_k$ be a convex body specifying the assignment structure constraint. There exists a randomized algorithm that given a dataset $P \subseteq \mathcal{X}$ of size $n \geq 1$ and an $(2^{O(z)}, O(1), O(1))$ -approximation $C^* \in \mathcal{X}^k$ of P for (k, z) -CLUSTERING with m outliers, constructs an $(\varepsilon, \mathcal{B}, m)$ -coreset for (k, z) -CLUSTERING with general assignment constraints of size

$$(3.4) \quad O(m) + 2^{O(z \log z)} \cdot \tilde{O}(\text{Lip}(\mathcal{B})^2 \cdot (\Lambda_\varepsilon(\mathcal{X}) + k + \varepsilon^{-1}) \cdot k^2 \varepsilon^{-2z}) \cdot \log \delta^{-1},$$

in $O(nk)$ time with probability at least $1 - \delta$, where $\text{Lip}(\mathcal{B})$ is the Lipschitz constant of \mathcal{B} defined in Definition 3.3, $\Lambda_\varepsilon(\mathcal{X})$ is the covering exponent of \mathcal{X} defined in Definition 3.4, and \tilde{O} hides a poly $\log(\text{Lip}(\mathcal{B}) \cdot \Lambda_\varepsilon(\mathcal{X}) \cdot k \varepsilon^{-1})$ term. Moreover, when $\mathcal{B} = \Delta_k$ (in this case, $\text{Lip}(\mathcal{B}) = 1$), the coreset size can be further improved to

$$(3.5) \quad O(m) + 2^{O(z \log z)} \cdot \tilde{O}((\Lambda_\varepsilon(\mathcal{X}) + \varepsilon^{-1}) \cdot k^2 \varepsilon^{-2z}) \cdot \log \delta^{-1}.$$

Our theorem provides the first coreset construction for capacitated and fair (k, z) -CLUSTERING with m outliers, improves the previous coreset size for capacitated/fair/robust (k, z) -CLUSTERING by at least a factor of $k \varepsilon^{-z}$, and establishes the first coreset construction for fault-tolerant clustering in various metric spaces; see “Notable Concrete Results” in Section 1.1 for more details.

For ease of analysis, we assume the given dataset P is unweighted. This assumption can be removed by a standard process of scaling the point weights to large integers,⁷ and treating each weight as a multiplicity of a point; details can be found in [CSS21, Corollary 2.3] and [CLSS22, Section 6.1]. By this theorem, the coreset size is decided by the Lipschitz constant $\text{Lip}(\mathcal{B})$ and the covering exponent $\Lambda_\varepsilon(\mathcal{X})$. If both $\Lambda_\varepsilon(\mathcal{X})$ and $\text{Lip}(\mathcal{B})$ are independent of n , e.g., upper bounded by poly(k, ε^{-1}), our coreset size is at most poly(k, ε^{-1}). When there are no structure constraints ($\mathcal{B} = \Delta_k$ and $\text{Lip}(\mathcal{B}) = 1$), we highlight that the coreset size in Equation (3.5) is better than that in Equation (3.4), by reducing the factor of $(\Lambda_\varepsilon(\mathcal{X}) + k + \varepsilon^{-1})$ to $(\Lambda_\varepsilon(\mathcal{X}) + \varepsilon^{-1})$.

The coreset construction algorithm of Theorem 3.1 is shown in Section 3.2, and the proof of Theorem 3.1 can be found in Section 3.3. Now we provide the formal definitions of $\text{Lip}(\mathcal{B})$ and $\Lambda_\varepsilon(\mathcal{X})$ that appear in the statement of Theorem 3.1. We discuss the Lipschitz constant $\text{Lip}(\mathcal{B})$ in Section 4, and show how to bound $\Lambda_\varepsilon(\mathcal{X})$ in Section 5.

Lipschitz constant $\text{Lip}(\mathcal{B})$. We first define optimal assignment transportation (OAT). Since this notion may be of independent interest, we present it in a slightly abstract way and we explain how it connects to our problem after the definition. Then, the key notion $\text{Lip}(\mathcal{B})$ is defined (in Definition 3.3) as the Lipschitz constant of this OAT procedure.

DEFINITION 3.2. (Optimal assignment transportation) Given a convex body $\mathcal{B} \subseteq \Delta_k$, two k -dimensional vectors $h, h' \in \mathcal{B}$, and a function $\sigma : [n] \times [k] \rightarrow \mathbb{R}_{\geq 0}$ such that $\|\sigma\|_1 = 1$ and $\sigma \sim (\mathcal{B}, h)$, the optimal assignment transportation is defined as the minimum total mass transportation $\|\sigma - \sigma'\|_1$ from σ to σ' , over functions $\sigma' : [n] \times [k] \rightarrow \mathbb{R}_{\geq 0}$ such that $\sigma' \sim (\mathcal{B}, h')$ (see Definition 2.1, 2.3) and $\forall p \in [n], \|\sigma'(p, \cdot)\|_1 = \|\sigma(p, \cdot)\|_1$. Namely,

$$\text{OAT}(\mathcal{B}, h, h', \sigma) := \min_{\substack{\sigma' \sim (\mathcal{B}, h') \\ \forall p \in [n], \|\sigma'(p, \cdot)\|_1 = \|\sigma(p, \cdot)\|_1}} \|\sigma - \sigma'\|_1.$$

⁷We suppose all weights are rational numbers such that we can round them to integers. If not, we can always replace it with a sufficiently close rational number such that the slight difference does not affect the clustering cost.

To see how our problem is related to Definition 3.2, we view σ as an assignment function, and $[n]$ and $[k]$ are interpreted as a data set P of n points and a center set C , respectively. Without loss of generality, we can assume $n = 1$ by normalization, and the requirement $\|\sigma\|_1 = n = 1$ in Definition 3.2 is satisfied. In fact, we choose to use $[n]$ and $[k]$ since the representation/identity of a point in the metric does not affect transportation; in other words, OAT is oblivious to the metric space. The requirement of $h, h' \in \mathcal{B}$ is to ensure the feasibility of OAT. Intuitively, this OAT aims to find a transportation plan, that transports the minimum mass to turn an initial capacity vector h into a target capacity vector h' . However, due to the presence of the assignment structure constraint \mathcal{B} , not all transportation plans are allowed. In particular, we are in addition required to start from a given assignment $\sigma \in \mathcal{B}$ that is consistent with h , and the way we reach h' must be via another assignment $\sigma' \in \mathcal{B}$ (which we optimize).

DEFINITION 3.3. (Lipschitz constant of OAT on \mathcal{B}) Let $\mathcal{B} \subseteq \Delta_k$ be a convex body. We define the Lipschitz constant of OAT on \mathcal{B} as $\text{Lip}(\mathcal{B}) := \max_{\substack{h, h' \in \mathcal{B}, \\ \sigma \sim (\mathcal{B}, h): \|\sigma\|_1 = 1}} \frac{\text{OAT}(\mathcal{B}, h, h', \sigma)}{\|h - h'\|_1}$.

Note that $\text{Lip}(\mathcal{B}) \geq 1$ since $\sum_{p \in [n]} (\sigma(p, \cdot) - \sigma'(p, \cdot)) = h - h'$. It is not hard to see that the Lipschitz constant is scale-invariant, i.e., $\text{Lip}(\mathcal{B}) = \text{Lip}(c \cdot \mathcal{B})$ for any \mathcal{B} and $c > 0$.

Covering Exponent $\Lambda_\epsilon(\mathcal{X})$. We now introduce the notion of covering for all center points $c \in \mathcal{X}$. Furthermore, we also define a quantity called *covering number* to measure the size/complexity of the covering, and we use a more convenient *covering exponent* to capture the worst-case size of the covering, which can serve as a parameter of the complexity for the metric space.

DEFINITION 3.4. (Covering, covering number and covering exponent) Let $P \subset \mathcal{X}$ be an unweighted set of $n \geq 1$ points and $a \in \mathcal{X}$ be a point. Let $r_{\max} = \max_{p \in P} d(p, a)$ such that $P \subseteq \text{Ball}(a, r_{\max})$. We say a collection $\mathcal{C} \subseteq \text{Ball}(a, 48zr_{\max}\epsilon^{-1})$ of points is an α -covering of P if for every point $c \in \text{Ball}(a, 48zr_{\max}\epsilon^{-1})$, there exists some $c' \in \mathcal{C}$ such that $\max_{p \in P} |d(p, c) - d(p, c')| \leq \frac{\alpha r_{\max}}{12z}$.

Define $N_{\mathcal{X}}(P, \alpha)$ to be the minimum cardinality $|\mathcal{C}|$ of any α -covering \mathcal{C} of P . Define the α -covering number of \mathcal{X} to be $N_{\mathcal{X}}(n, \alpha) := \max_{P \subseteq \mathcal{X}: |P|=n} N_{\mathcal{X}}(P, \alpha)$, i.e., the maximum $N_{\mathcal{X}}(P, \alpha)$ over all possible unweighted sets P of size n .

Moreover, we define the α -covering exponent of the metric space (\mathcal{X}, d) , denoted by $\Lambda_\alpha(\mathcal{X})$, to be the least integer $\gamma \geq 1$ such that $N_{\mathcal{X}}(n, \alpha) \leq O(n^\gamma)$ holds for any $n \geq 2$.

Roughly speaking, an α -covering \mathcal{C} is an discretization of the continuous space of all possible centers within $\text{Ball}(a, 48zr\epsilon^{-1})$ of a large radius w.r.t. the ℓ_∞ -distance differences from points $p \in P$ to $c \in \text{Ball}(a, 48zr\epsilon^{-1})$. For those center $c \notin \text{Ball}(a, 48zr\epsilon^{-1})$, we can verify that the distances between every point $p \in P$ and c are very close, i.e., $d^z(p, c) \in (1 \pm \epsilon) \cdot d^z(q, c)$ for any two points $p, q \in P$. This observation enables us to safely “ignore” the complexity of these remote centers in coresets construction, and hence we only need to consider centers within $\text{Ball}(a, 48zr\epsilon^{-1})$.

Since $|\mathcal{C}| \leq |\mathcal{X}|$, we have that the covering exponent $\Lambda_\alpha(\mathcal{X}) \leq \log |\mathcal{X}|$. This notion is closely related to other combinatorial dimension notions such as the VC-dimension and the (fat) shattering dimension of the set system formed by all metric balls which have been extensively considered in previous works, e.g. [LS10, FL11, FSS20, HJLW18, BJKW21a]. We will show the relations between the covering exponent in our setting and two well-studied dimension notions, the shattering dimension and the doubling dimension, in Section 5.

3.1 Overview of the Proof of Theorem 3.1

We provide an overview of the proof of Theorem 3.1.

3.1.1 Improved Hierarchical Uniform Sampling Framework We take Euclidean k -MEDIAN as an example, whose idea can be extended to (k, z) -CLUSTERING via the generalized triangle inequality (Lemma 3.5). The following analysis is for capacity constraints, which can be extended to general assignment constraints in Section 3.1.2. For simplicity, we use cost to represent cost_1 .

Review: Hierarchical Uniform Sampling Framework [BCJ⁺22, HJLW23] We briefly review the coresets algorithm in [BCJ⁺22] now. We first compute an $O(1)$ -approximation $C^* = \{c_1^*, \dots, c_k^*\}$ for the k -MEDIAN problem and partition dataset P into k clusters P_1, \dots, P_k . Then adopting the idea of [BCJ⁺22] (Theorem 3.2),

we decompose every P_i into a collection \mathcal{R}_i of $\tilde{O}(k\varepsilon^{-1})$ disjoint *rings* and a collection \mathcal{G}_i of $\tilde{O}(k\varepsilon^{-1})$ disjoint *groups* centered at c_i^* , and reduce the problem to constructing coresets for \mathcal{R}_i and \mathcal{G}_i (Theorem 3.2). To this end, Braverman et al. [BCJ⁺22] takes a uniform sample S_R on every ring $R \in \mathcal{R}_i$ with a uniform size $\Gamma = \tilde{O}(k\varepsilon^{-4})$ of samples⁸ such that for every center set $C \subset \mathbb{R}^d$ and every capacity constraint h ,

$$(3.6) \quad |\text{cost}(R, C, h) - \text{cost}(S_R, C, h)| \leq \varepsilon(\text{cost}(R, C, h) + \text{cost}(R, c_i^*));$$

and construct a two-point coreset S_G on every group $G \in \mathcal{G}_i$. Later on, Huang et al. [HJLW23] showed this framework also works for k -MEDIAN with m outliers (the only difference is that we need to compute an $O(1)$ -approximation $C^* = \{c_1^*, \dots, c_k^*\}$ for the k -MEDIAN problem with m outliers in the first step). In previous work [BCJ⁺22, HJLW23], it is unknown whether Inequality (3.6) still work if we allow m outliers.

New Idea: Adaptive Sample Sizes for Rings The main improvement of our algorithm (Algorithm 1) compared to [BCJ⁺22, HJLW23] is a significant decrease in the number of samples for the rings. The key observation is that we only need the total error of rings to be upper bounded, i.e.,

$$(3.7) \quad \sum_{R \in \mathcal{R}_i} |\text{cost}(R, C, h) - \text{cost}(S_R, C, h)| \leq \varepsilon \left(\sum_{R \in \mathcal{R}_i} \text{cost}(R, C, h) \right) + \varepsilon \text{cost}(P, c_i^*),$$

which is intuitively much easier to be satisfied than Inequality (3.6). Concretely, we regard \mathcal{R}_i as a whole instead of independent rings and focus on ensuring Inequality (3.7). This point of view enables us to adaptive select the sample size Γ_R for every ring $R \in \mathcal{R}_i$ (Line 6 of Algorithm 1), say $\Gamma_R = \tilde{O}(k\varepsilon^{-4}) \cdot \lambda_R$ where $\lambda_R = \text{cost}(R, c_i^*) / \text{cost}(P_i, c_i^*)$ is the relative contribution of R to P_i . Then our coreset size is dominated by $\sum_{i \in [k], R \in \mathcal{R}_i} \Gamma_R \leq \tilde{O}(k^2\varepsilon^{-4})$, saving a factor of $k\varepsilon^{-1}$ compared to that of [BCJ⁺22, HJLW23].

Handling Rings The most technical part is to show that Inequality (3.7) always holds by our construction, and we introduce the ideas now. We remark that the estimation error $|\text{cost}(R, C, h) - \text{cost}(S_R, C, h)|$ of $R = \text{ring}(c_i^*, r, 2r)$ is likely to be decided by those centers $c \in C$ that are both “not too far” from R and “not too close” to c_i^* . This intuition motivates us to consider the number of “effective centers” to R , denoted as the level $t_R(C) := |C \cap \text{ring}(c_i^*, \frac{\varepsilon r}{48}, \frac{48r}{\varepsilon})|$ of R w.r.t. C (Definition 3.6). The idea of excluding centers outside $B(c_i^*, \frac{48r}{\varepsilon})$ has been applied in [BCJ⁺22] to handle rings, and of excluding centers within $B(c_i^*, \frac{\varepsilon r}{48})$ is new. An immediate geometric observation is the following bound of the total level for any center set $C \subset \mathbb{R}^d$ (Lemma 3.2):

$$(3.8) \quad \sum_{R \in \mathcal{R}_i} t_R(C) \leq O(k \log \varepsilon^{-1}),$$

due to the ring structure of \mathcal{R}_i . This property is somewhat surprising since the positions of C seem to be arbitrary, and is a key for our improvement. We remark that the analysis of [BCJ⁺22] simply bounds $t_R(C)$ by k , and hence, leads to a bound $\sum_{R \in \mathcal{R}_i} t_R(C) \leq \tilde{O}(k^2\varepsilon^{-1})$, which is a factor of $k\varepsilon^{-1}$ larger than Inequality (3.8). This factor also matches our improvement in the coreset size. Our key lemma (Lemma 3.3) shows that the estimation error of R is “proportional to” $t_R(C)$ such that the total error is “proportional to” $\sum_{R \in \mathcal{R}_i} t_R(C)$, which is small by the above observation.

Overall, our geometric observation of $t_R(C)$ (Inequality (3.8)) enables us to tolerant a larger estimation error for every ring that is captured by λ_R , and the total estimation error is under control due to a combined consideration of $t_R(C)$ and λ_R .

Handling Groups We expand upon the analysis presented in [HJLW23] for the k -MEDIAN problem with m outliers, specifically addressing the inclusion of capacity constraints. A crucial geometric observation made in [HJLW23] is that, in the context of any center set $C \subset \mathbb{R}^d$, there are at most two “special uncolored groups” which intersect the outliers in a partial manner. However, this property does not hold when considering capacity constraints with outliers. Fortunately, we can overcome this limitation by dividing the groups into k equivalent classes based on their corresponding remote centers (Lemma 3.14). Consequently, the number of special uncolored groups is bounded by $O(k)$, thereby achieving the desired error bound for the groups (Lemma 3.4).

⁸In their original paper, [BCJ⁺22] claims a bound $\Gamma = \tilde{O}(k\varepsilon^{-5})$ and in a follow-up work [HJLW23] it is shown that $\tilde{O}(k\varepsilon^{-4})$ is already a sufficient choice

Extension to General Metric Spaces We apply a high-level idea in [BCJ⁺22] to discretize the hyperparameter space $\mathcal{X}^k \times (|R| \cdot \mathcal{B})$ into a small number of representative pairs (Lemma 3.7), and show that the error is preserved for every representative pair (C, h) with very high probability (using e.g., concentration bounds). Our discretization of the center space \mathcal{X}^k directly relates to the *covering exponent* $\Lambda_\varepsilon(\mathcal{X})$ of the metric space, which is a complexity measure of general metric spaces, instead of a more geometric discretization based on the notion of metric balls which are more specific to Euclidean spaces. Notably, this feature enables us to handle capacitated and fair clustering on any metric that admits a small shattering dimension or doubling dimension, while the analysis of [BCJ⁺22] is specific to metric space with bounded doubling dimension due to the requirement of the existence of a small ε -net.

3.1.2 Handling Assignment Structure Constraints Next, we provide some explanations on why the algorithm can handle assignment structure constraints. We discuss rings and show how the additional term $\text{Lip}(\mathcal{B})$ appears in the coreset size. More detailed analysis for rings and groups can be found in Sections 3.4 and 3.5, respectively.

Complexity Measure of Assignment Structure Constraint: $\text{Lip}(\mathcal{B})$ We start with a more detailed discussion on the new notion $\text{Lip}(\mathcal{B})$ due to its conceptual and technical importance. The definition of $\text{Lip}(\mathcal{B})$ (Definitions 3.2 and 3.3) may be interpreted in several ways. We start with an explanation from a technical perspective of coreset construction. A natural way of building a coreset which we also use, is to draw independent samples S from data points P (and re-weight). To analyze S , let us fix some capacity constraint h and some center C . Let σ^* be an assignment of h , i.e., $\text{cost}_z(P, C, \mathcal{B}, h) = \text{cost}_{z^*}^{\sigma^*}(P, C)$. Then, one can convert σ^* to $\sigma : S \times C \rightarrow \mathbb{R}_{\geq 0}$ for the sample $S \subseteq P$, by setting $\sigma(p, \cdot) := w(p) \cdot \sigma^*(p, \cdot)$ and $w(p) = \frac{n}{|S|}$ for $p \in S$ (where we can guarantee that $w(S) = n$). Even though this assignment σ may slightly violate the capacity constraint h , the violation, denoted as $\|h - h'\|_1$ where h' is the capacity induced by σ , is typically very small (by concentration inequalities), and it may be charged to $\text{cost}_z(P, C, \mathcal{B}, h)$. However, we still need to transport σ to σ' so that it is consistent with h' . More precisely, we need to find $\sigma' : S \times C \rightarrow \mathbb{R}_{\geq 0}$ for the sample S that satisfies $\sigma' \sim (h', \mathcal{B})$, such that the total transportation $\|\sigma - \sigma'\|_1$ ⁹ (which relates to the cost change), is minimized. This minimum transportation plan is exactly defined by the *optimal assignment transportation* (OAT) (Definition 3.2). Since we eventually wish to bound the OAT cost against the mentioned $\|h - h'\|_1$, our $\text{Lip}(\mathcal{B})$ is defined as the universal upper bound of the OAT cost relative to $\|h - h'\|_1$ over all h, h', σ , which can also be viewed as the Lipschitz constant of OAT.

Another perspective of interpreting the notion of OAT is via the well-known *optimal transportation*. Specifically, the minimum L_1 transportation cost for turning h to h' without any constraint is exactly $\|h - h'\|_1$, which equals to the minimum transportation cost from any assignment $\sigma \sim h$ to some $\sigma' \sim h'$. Hence, compared with optimal transportation, our notion adds additional requirements that the transportation plan must be inside of \mathcal{B} , and from a given starting assignment σ with $\sigma \sim h$. We care about the worst-case relative ratio $\text{Lip}(\mathcal{B})$, which measures how many times more expensive the constrained optimal transportation cost OAT than the optimal transportation cost $\|h - h'\|_1$. Even though notions of constrained optimal transportation were also considered in the literature, see e.g., [TT13, KM13, KM15, KMS15] and a survey [PC⁺19, Chapter 10.3] for more, we are not aware of any previous work that studies exactly the same problem.

3.2 The Coreset Construction Algorithm We introduce the algorithm used in Theorem 3.1. The algorithm improves the coreset algorithm proposed very recently by [BCJ⁺22, HJLW23].

We need the following decomposition that is defined with respect to a dataset $Q \subseteq \mathcal{X}$ and a given center $c \in \mathcal{X}$.

THEOREM 3.2. (Decomposition into rings and groups, [BCJ⁺22]) *Let $Q \subseteq \mathcal{X}$ be a weighted dataset and $c \in \mathcal{X}$ be a center point. There exists an $O(|Q|)$ time algorithm $\text{Decom}(Q, c)$ that outputs $(\mathcal{R}^*, \mathcal{G}^*)$ as a partition of Q , where \mathcal{R}^* and \mathcal{G}^* are two disjoint collections of sets such that $Q = (\cup_{R \in \mathcal{R}^*} R) \cup (\cup_{G \in \mathcal{G}^*} G)$. Moreover, \mathcal{R}^* is a collection of disjoint rings satisfying*

1. $\forall R \in \mathcal{R}^*$, R is a ring of form $R = R_i(Q, c)$ for some integer $i \in \mathbb{Z} \cup \{-\infty\}$, where $R_i(Q, c) := Q \cap \text{ring}(c, 2^{i-1}, 2^i)$ for $i \in \mathbb{Z}$ and $R_{-\infty}(Q, c) := Q \cap \{c\}$;

⁹Here, we interpret σ, σ' as vectors on $\mathbb{R}^{S \times C}$.

2. $\forall R \in \mathcal{R}^*, \text{cost}_z(R, c) \geq \left(\frac{\varepsilon}{6z}\right)^z \cdot \frac{\text{cost}_z(Q, c)}{k \cdot \log(48z\varepsilon^{-1})}$;
3. $|\mathcal{R}^*| \leq 2^{O(z \log z)} \cdot \tilde{O}(k\varepsilon^{-z})$;

and \mathcal{G}^* is a collection of disjoint groups satisfying

1. $\forall G \in \mathcal{G}^*, G$ is the union of a few consecutive rings of (Q, c) and all these rings are disjoint. Formally, $\forall G \in \mathcal{G}^*,$ there exists $l_G, r_G \in \mathbb{Z}^*, l_G \leq r_G$ such that $G = \cup_{i=l_G}^{r_G} R_i(Q, c)$ and the intervals $\{[l_G, r_G], G \in \mathcal{G}^*\}$ are disjoint;
2. $\forall G \in \mathcal{G}^*, \text{cost}_z(G, c) \leq \left(\frac{\varepsilon}{6z}\right)^z \cdot \frac{\text{cost}_z(Q, c)}{k \cdot \log(48z\varepsilon^{-1})}$;
3. $|\mathcal{G}^*| \leq 2^{O(z \log z)} \cdot \tilde{O}(k\varepsilon^{-z})$.

Roughly speaking, we decompose Q into rings w.r.t. c . The collection \mathcal{R}^* contains those rings with “heavy” costs, say $\text{cost}_z(R, c) > \left(\frac{\varepsilon}{6z}\right)^z \cdot \frac{\text{cost}_z(Q, c)}{k \cdot \log(48z\varepsilon^{-1})}$ for $R \in \mathcal{R}^*$.¹⁰ They also gather the remaining “light” rings and form the collection of groups \mathcal{G}^* (see [BCJ⁺22, Lemma 3.4]), and ensure that the cardinality $|\mathcal{G}^*|$ is upper bounded by $2^{O(z \log z)} \cdot \tilde{O}(k\varepsilon^{-z})$.

For each group G , we provide the following data structure.

DEFINITION 3.5. (Two-point coreset, Line 5 of Algorithm 1 in [BCJ⁺22]) For a weighted dataset/group $G \subset \mathbb{R}^d$ and a center point $c \in \mathbb{R}^d$, let p_{far}^G and p_{close}^G denote the furthest and closest point to c in G . For every $p \in G$, compute the unique $\lambda_p \in [0, 1]$ such that $d^z(p, c) = \lambda_p \cdot d^z(p_{\text{close}}^G, c) + (1 - \lambda_p) \cdot d^z(p_{\text{far}}^G, c)$. Let $D = \{p_{\text{far}}^G, p_{\text{close}}^G\}$, $w_D(p_{\text{close}}^G) = \sum_{p \in G} \lambda_p$, and $w_D(p_{\text{far}}^G) = \sum_{p \in G} (1 - \lambda_p)$. D is called the two-point coreset of G with respect to c .

By definition, we know that $w(D) = |G|$ and $\text{cost}_z(D, c) = \text{cost}_z(G, c)$, which are useful for upper bounding the error induced by such two-point coresets.

Algorithm 1 HUS(P, k, z, Γ)

Input: An unweighted dataset $P \subseteq \mathcal{X}$ of size $n \geq 1$, an integer $k \geq 1$, an integer $0 \leq m \leq n$, constant $z \geq 1$, a sampling size $\Gamma \geq 1$, and an $(2^{O(z)}, O(1), O(1))$ -approximate solution $C^* = \{c_1^*, \dots, c_k^*\} \subseteq \mathcal{X}$ of P for (k, z) -CLUSTERING with m outliers

Output:

- 1: Let $L^* \leftarrow \arg \min_{L \subseteq P: |L|=m} \text{cost}_z(P \setminus L, C^*)$ denote the set of m outliers of P w.r.t. C^*
 - 2: Decompose $P \setminus L^*$ into k clusters P_1, \dots, P_k such that each P_i contains all points in P whose closest center in C^* is c_i^* (breaking ties arbitrarily).
 - 3: For each $i \in [k]$, apply the decomposition of Theorem 3.2 to (P_i, c_i^*) and obtain $(\mathcal{R}_i, \mathcal{G}_i) \leftarrow \text{Decom}(P_i, c_i^*)$, where \mathcal{R}_i is a collection of disjoint rings and \mathcal{G}_i is a collection of disjoint groups.
 - 4: For each $i \in [k]$ and each ring $R \in \mathcal{R}_i$, set $\lambda_R \leftarrow \frac{\text{cost}_z(R, c_i^*)}{\text{cost}_z(P_i, c_i^*)}$ and take a uniform sample S_R of size $\Gamma_R \leftarrow \lceil \Gamma \cdot \lambda_R \rceil$ from R , and set $w_{S_R}(x) = \frac{|R|}{\Gamma_R}$ for each point $x \in S_R$.
 - 5: For each $i \in [k]$, for each group $G \in \mathcal{G}_i$ and center c_i^* , construct a two-point coreset D_G of G by Definition 3.5.
 - 6: Return $S \leftarrow L^* \cup (\cup_R S_R) \cup (\cup_G D_G)$.
-

Hierarchical Uniform Sampling Coreset Framework We are now ready to introduce the hierarchical uniform sampling framework HUS(P, k, z, Γ) (Algorithm 1). To simplify our analysis, we slightly abuse the notation and consider C^* as an $(2^{O(z)}, 1, 1)$ -approximation instead of a tri-approximation, e.g., a $(2^{O(z)}, O(1), O(1))$ -approximation by [BVX19]. This simplification allows the algorithm to decompose the inliers $P \setminus L^*$ into simply k clusters rather than $O(k)$ clusters, which only results in a factor of $O(1)$ difference in the coreset size. We first compute an outlier set L^* w.r.t. C^* in Line 1 (as in [HJLW23, Algorithm 1]). Note that we apply Theorem 3.2 for each P_i ($i \in [k]$) in Line 3, and obtain collections $\mathcal{R}_1, \dots, \mathcal{R}_k$ and $\mathcal{G}_1, \dots, \mathcal{G}_k$. Let $\mathcal{R} = \cup_{i \in [k]} \mathcal{R}_i$ be the collection of all rings in different clusters and $\mathcal{G} = \cup_{i \in [k]} \mathcal{G}_i$ be the collection of all groups. By Theorem 3.2, we have the following observation.

¹⁰In [BCJ⁺22], they call these rings heavy rings or marked rings.

OBSERVATION 1. (**Bound for $|\mathcal{R}|$ and $|\mathcal{G}|$)** $|\mathcal{R}|, |\mathcal{G}| \leq 2^{O(z \log z)} \cdot \tilde{O}(k^2 \varepsilon^{-z})$.

The primary improvement of Algorithm 1 lies in Line 4, where we selectively choose the coreset size Γ_R that is proportional to the relative contribution λ_R of each ring R . Moreover, we demonstrate that our algorithm, denoted as $\text{HUS}(P, k, z, \Gamma)$, can generate an output (S, w) that is a coreset for (k, z) -CLUSTERING with general assignment constraints in general metric spaces, provided that we carefully choose the sample number Γ .

We remark that Algorithm 1 presents a hierarchical uniform sampling framework instead of purely sampling. However, even in vanilla clustering, achieving state-of-the-art coreset size purely through sampling, without any pre-processing, has only been recently established [BCP⁺24]. It is interesting to study whether this purely sampling framework can be extended to constrained clustering.

3.3 Proof of Theorem 3.1: Performance Analysis of Algorithm 1 We use the following well-known generalized triangle inequalities for (k, z) -CLUSTERING; see e.g., [MSSW18, BCJ⁺22] for more variants. The key is to bound the induced errors for rings and groups for every k -center set C respectively (Lemmas 3.3 and 3.4). For groups, we provide a unified error bound that does not depend on the choice of C . In contrast, our error bounds for rings R depend on the relative location of C to R , captured by a new notion called the level $t_R(C)$ of rings R w.r.t. C (Definition 3.6). Bounding the total error induced by rings relies on a new geometric observation that bounds the total levels $\sum_R t_R(C)$ for all C (Lemma 3.2).

The following lemma is useful for bounding the total induced error of rings.

LEMMA 3.1. (**Hölder’s Inequality**) *Assume $p, q > 1$ such that $\frac{1}{p} + \frac{1}{q} = 1$ then for every integer $n \geq 1$ and two sequences of numbers $a_1, a_2, \dots, a_n, b_1, b_2, \dots, b_n > 0$,*

$$\sum_{i=1}^n a_i b_i \leq \left(\sum_{i=1}^n a_i^p \right)^{\frac{1}{p}} \cdot \left(\sum_{i=1}^n b_i^q \right)^{\frac{1}{q}}$$

Now we are ready to prove Theorem 3.1. Like [BCJ⁺22], we prove the theorem by analyzing rings and groups separately; summarized by Lemmas 3.3 and 3.4 respectively.

Given a ring $R \subset \text{ring}(a, r, 2r)$ and a center set $C \in \mathcal{X}^k$, we denote $C_{\text{far}}^R := \{c \in C : d(c, a) \geq 48z\varepsilon^{-1}r\}$ to be the centers that are remote to R and $C_{\text{close}}^R := \{c \in C : d(c, a) \leq \varepsilon r / 48z\}$ to be the centers that are close to a . We first introduce the following important notion.

DEFINITION 3.6. (**Level of rings**) *Given a ring $R \subset \text{ring}(a, r, 2r)$ for some $a \in \mathbb{R}^d$ and radius $r > 0$ and a center set $C \in \mathcal{X}^k$, we denote the level of ring R w.r.t. C to be $t_R(C) := |C \setminus C_{\text{far}}^R \setminus C_{\text{close}}^R|$.*

Note that $0 \leq t_R(C) \leq k$. We will see that the level $t_R(C)$ heavily affects the induced error of samples S_R w.r.t. C . By Theorem 3.2, we directly have the following observation.

LEMMA 3.2. (**Bounding total levels**) *Let $\varepsilon \in (0, \frac{1}{4})$. Given a center set $C \subseteq \mathcal{X}^k$, for every $i \in [k]$, we have that $\sum_{R \in \mathcal{R}_i} t_R(C) \leq 10zk \log \varepsilon^{-1}$.*

Proof. We say center $c \in C$ is interesting to ring R if $c \in C \setminus C_{\text{far}}^R \setminus C_{\text{close}}^R$. It suffices to prove that every $c \in C$ can be interesting to at most $10z \log \varepsilon^{-1}$ rings in \mathcal{R}_i . For the sake of contradiction, suppose c is interesting to more than $8z \log \varepsilon^{-1} + 1$ rings. Among these rings, let $R_1 = P_i \cap \text{ring}(c_i^*, r_1, 2r_1)$ and $R_2 = P_i \cap \text{ring}(c_i^*, r_2, 2r_2)$ denote rings with the largest and smallest radii respectively. Recall that all rings are disjoint, thus we have $r_1/r_2 > 2^{8z \log \varepsilon^{-1}} = \varepsilon^{-8z}$. However, as c is interesting to both R_1 and R_2 , by Definition 3.6, we know that

$$\frac{\varepsilon r_1}{48z} \leq d(c, c_i^*) \leq \frac{48z r_2}{\varepsilon}$$

which implies $r_1/r_2 \leq 2304z^2 \varepsilon^{-2} < \varepsilon^{-8z}$ since $\varepsilon < \frac{1}{4}$.

So we conclude with a contradiction and have proved Lemma 3.2. □

We have the following lemma that relates the induced error of rings with their levels.

LEMMA 3.3. (**Error analysis for rings**) For each $i \in [k]$ and ring $R \in \mathcal{R}_i$, suppose

$$\Gamma_R \geq 2^{O(z \log z)} \cdot \lambda_R \cdot \text{Lip}(\mathcal{B})^2 \cdot (\Lambda_\varepsilon(\mathcal{X}) + k + \varepsilon^{-1}) \cdot k\varepsilon^{-2z} \cdot \log(\text{Lip}(\mathcal{B}) \cdot \Lambda_\varepsilon(\mathcal{X})\delta^{-1}) \log^7(k\varepsilon^{-1}),$$

and when $\mathcal{B} = \Delta_k$, suppose

$$\Gamma_R \geq 2^{O(z \log z)} \cdot \lambda_R \cdot (\Lambda_\varepsilon(\mathcal{X}) + \varepsilon^{-1}) \cdot k\varepsilon^{-2z} \cdot \log(\Lambda_\varepsilon(\mathcal{X})\delta^{-1}) \log^7(k\varepsilon^{-1}),$$

where Γ_R is the sample size of S_R as in Line 4 of Algorithm 1. With probability at least $1 - \frac{\delta}{|\mathcal{R}|}$, for every k -center set $C \in \mathcal{X}^k$ and capacity constraint $h \in |R| \cdot \text{conv}(\mathcal{B}^\circ)$,

$$\begin{aligned} & |\text{cost}_z(R, C, \mathcal{B}, h) - \text{cost}_z(S_R, C, \mathcal{B}, h)| \\ & \leq \varepsilon (\text{cost}_z(R, C, \mathcal{B}, h) + \text{cost}_z(R, c_i^*)) + \left(\frac{t_R(C)}{10zk\lambda_R \log \varepsilon^{-1}}\right)^{\frac{1}{2}} \cdot \varepsilon \text{cost}_z(R, c_i^*). \end{aligned}$$

Proof. The proof can be found in Section 3.4. \square

For groups, we have the following lemma.

LEMMA 3.4. (**Error analysis for groups**) For each $i \in [k]$, let $G[i] = \bigcup_{G \in \mathcal{G}_i} G$ be the union of all groups $G \in \mathcal{G}_i$ and $D[i] = \bigcup_{G \in \mathcal{G}_i} D_G$ be the union of all two-point coresets D_G with $G \in \mathcal{G}_i$. For every k -center set $C \in \mathcal{X}^k$ and capacity constraint $h \in |G[i]| \cdot \text{conv}(\mathcal{B}^\circ)$,

$$|\text{cost}_z(G[i], C, \mathcal{B}, h) - \text{cost}_z(D[i], C, \mathcal{B}, h)| \leq O(\varepsilon) \cdot (\text{cost}_z(G[i], C, \mathcal{B}, h) + \text{cost}_z(P_i, c_i^*)).$$

Proof. The proof can be found in Section 3.5. \square

Note that for groups, the induced error of two-point coresets $D[i]$ is deterministically upper bounded, which is not surprising since there is no randomness in the construction of $D[i]$. This property is quite powerful since we do not need to consider the complexity of center sets in different metric spaces when analyzing the performance of two-point coresets. Now we are ready to prove Theorem 3.1.

Proof. [Proof of Theorem 3.1] By Lemma 3.3, we can select

$$\Gamma = 2^{O(z \log z)} \cdot \tilde{O}(\text{Lip}(\mathcal{B})^2 \cdot (\Lambda_\varepsilon(\mathcal{X}) + k + \varepsilon^{-1}) \cdot k\varepsilon^{-2z}) \cdot \log \delta^{-1}$$

for general assignment structure constraint \mathcal{B} and select

$$\Gamma = 2^{O(z \log z)} \cdot \tilde{O}((\Lambda_\varepsilon(\mathcal{X}) + \varepsilon^{-1}) \cdot k\varepsilon^{-2z}) \cdot \log \delta^{-1}$$

when $\mathcal{B} = \Delta_k$, and apply HUS(P, k, z, Γ) that outputs (S, w) . We verify that S is the desired $O(\varepsilon)$ -coreset.

For the coreset size $|S|$, we first note that $\sum_{R \in \mathcal{R}_i} \Gamma_R \leq \Gamma$ for every $i \in [k]$ by the definition of λ_R . Also by Observation 1, $\sum_{G \in \mathcal{G}} \Gamma_G \leq O(k^2 \varepsilon^{-z})$. Hence, the size $|S|$ is dominated by $|L^*| + \Gamma \cdot k$, which matches the coreset size in Theorem 3.1. For correctness, we first have the following claim by Lemma 3.3.

CLAIM 3. With probability at least $1 - \delta$, for every $i \in [k]$, for every center set $C \in \mathcal{X}^k$ and capacity constraints $\{h_R\}_{R \in \mathcal{R}_i}$ satisfying $\forall R \in \mathcal{R}_i, h_R \in |R| \cdot \text{conv}(\mathcal{B}^\circ)$, we have

$$\sum_{R \in \mathcal{R}_i} |\text{cost}_z(R, C, \mathcal{B}, h_R) - \text{cost}_z(S_R, C, \mathcal{B}, h_R)| \leq \varepsilon \sum_{R \in \mathcal{R}_i} \text{cost}_z(R, C, \mathcal{B}, h_R) + 2\varepsilon \text{cost}_z(P_i, c_i^*).$$

Proof. Assume Lemma 3.3 holds for all rings $R \in \mathcal{R}$, whose success probability is at least $1 - \delta$ by the union bound. Fix a center set $C \in \mathcal{X}^k$ and capacity constraints $\{h_R\}_{R \in \mathcal{R}}$. By Lemma 3.3 and $\lambda_R = \frac{\text{cost}_z(R, c_i^*)}{\text{cost}_z(P_i, c_i^*)}$, we have that for every ring $R \in \mathcal{R}_i$,

$$\begin{aligned}
 & |\text{cost}_z(R, C, \mathcal{B}, h_R) - \text{cost}_z(S_R, C, \mathcal{B}, h_R)| \\
 \leq & \varepsilon (\text{cost}_z(R, C, \mathcal{B}, h_R) + \text{cost}_z(R, c_i^*)) + \left(\frac{t_R(C)}{10zk\lambda_R \cdot \log \varepsilon^{-1}}\right)^{\frac{1}{2}} \cdot \varepsilon \text{cost}_z(R, c_i^*) \\
 = & \varepsilon (\text{cost}_z(R, C, \mathcal{B}, h_R) + \text{cost}_z(R, c_i^*)) + \left(\frac{t_R(C)}{10zk \log \varepsilon^{-1}}\right)^{\frac{1}{2}} \lambda_R^{\frac{1}{2}} \cdot \varepsilon \text{cost}_z(P_i, c_i^*)
 \end{aligned}$$

Summing over all $R \in \mathcal{R}_{i,0}$ we have

$$\begin{aligned}
 & \sum_{R \in \mathcal{R}_i} |\text{cost}_z(R, C, \mathcal{B}, h_R) - \text{cost}_z(S_R, C, \mathcal{B}, h_R)| \\
 \leq & \varepsilon \sum_{R \in \mathcal{R}_i} (\text{cost}_z(R, C, \mathcal{B}, h_R) + \text{cost}_z(R, c_i^*)) + \left(\frac{t_R(C)}{10zk \log \varepsilon^{-1}}\right)^{\frac{1}{2}} \lambda_R^{\frac{1}{2}} \cdot \varepsilon \text{cost}_z(P_i, c_i^*) \\
 \leq & \varepsilon \sum_{R \in \mathcal{R}_i} \text{cost}_z(R, C, \mathcal{B}, h_R) + 2\varepsilon \text{cost}_z(P_i, c_i^*)
 \end{aligned}$$

where for the last inequality, we are using Holder's inequality (Lemma 3.1) and the fact

$$\sum_{R \in \mathcal{R}_i} t_R(C) \leq 10zk \log \varepsilon^{-1} \quad \text{and} \quad \sum_{R \in \mathcal{R}_i} \lambda_R \leq 1$$

to obtain

$$\begin{aligned}
 & \sum_{R \in \mathcal{R}_i} \left(\frac{t_R(C)}{10zk \log \varepsilon^{-1}}\right)^{\frac{1}{2}} \lambda_R^{\frac{1}{2}} \\
 \leq & \left(\sum_{R \in \mathcal{R}_i} \frac{t_R(C)}{10zk \log \varepsilon^{-1}}\right)^{\frac{1}{2}} \cdot \left(\sum_{R \in \mathcal{R}_i} \lambda_R\right)^{\frac{1}{2}} \\
 \leq & 1.
 \end{aligned}$$

Thus, we prove Claim 3. \square

Fix a k -center set $C \in \mathcal{X}^k$, and a capacity constraint $h \in (n - m) \cdot \mathcal{B}$. Suppose a collection $h^L \cup \{h^R \in |R| \cdot \text{conv}(\mathcal{B}^o) : R \in \mathcal{R}\} \cup \{h^{(i)} \in |G[i]| \cdot \text{conv}(\mathcal{B}^o) : G \in \mathcal{G}\}$ of capacity constraints satisfy that

$$h^L + \sum_{R \in \mathcal{R}} h^R + \sum_{i \in [k]} h^{(i)} = h,$$

and

$$(3.9) \quad \text{cost}_z(P, C, \mathcal{B}, h) = \text{cost}_z(L^*, C, \mathcal{B}, h^L) + \sum_{R \in \mathcal{R}} \text{cost}_z(R, C, \mathcal{B}, h^R) + \sum_{i \in [k]} \text{cost}_z(G[i], C, \mathcal{B}, h^{(i)}).$$

We have

$$\begin{aligned}
 & \text{cost}_z(S, C, \mathcal{B}, h) \\
 \leq & \text{cost}_z(L^*, C, \mathcal{B}, h^L) + \sum_{R \in \mathcal{R}} \text{cost}_z(S_R, C, \mathcal{B}, h^R) + \sum_{i \in [k]} \text{cost}_z(D[i], C, \mathcal{B}, h^{(i)}) \quad (\text{by optimality}) \\
 \leq & \text{cost}_z(L^*, C, \mathcal{B}, h^L) \\
 & + (1 + O(\varepsilon)) \cdot \sum_{R \in \mathcal{R}} \text{cost}_z(R, C, \mathcal{B}, h^R) + O(\varepsilon) \cdot \sum_{i \in [k]} \text{cost}_z(P_i, c_i^*) \quad (\text{Claim 3}) \\
 & + \sum_{i \in [k]} (1 + O(\varepsilon)) \cdot \text{cost}_z(G[i], C, \mathcal{B}, h^{(i)}) + O(\varepsilon) \cdot \text{cost}_z(P_i, c_i^*) \quad (\text{Lemma 3.4}) \\
 \leq & (1 + O(\varepsilon)) \cdot \text{cost}_z(P, C, \mathcal{B}, h) + O(\varepsilon) \cdot \text{cost}_z(P, C^*) \quad (\text{Ineq. (3.9)}) \\
 \leq & (1 + O(\varepsilon)) \cdot \text{cost}_z(P, C, \mathcal{B}, h). \quad (\text{Defn. of } C^*)
 \end{aligned}$$

Similarly, we also have that $\text{cost}_z(P, C, \mathcal{B}, h) \leq (1 + O(\varepsilon)) \cdot \text{cost}_z(S, C, \mathcal{B}, h)$. Thus, (S, w) is indeed an $O(\varepsilon)$ -coreset.

For the running time, Line 1 costs $O(nk)$ time. Line 2 costs $\sum_{i \in [k]} O(|P_i|) = O(n)$ time by Theorem 3.2. Line 3 costs $O(n)$ time. Line 4 costs $\sum_{G \in \mathcal{G}_i} O(|G|) = O(n)$ time by Definition 3.5. Overall, the total time is $O(nk)$. \square

3.4 Proof of Lemma 3.3: Error Analysis for Rings Now we analyze the errors induced by rings. We first provide an overview of the proof. Motivated by [BCJ⁺22], we first observe that error induced by centers in $C_{\text{far}}^R \setminus C_{\text{close}}^R$ can be well controlled and we only need to bound the error induced by centers in $C \setminus C_{\text{far}}^R \setminus C_{\text{close}}^R$ (Lemma 3.6). We also define another notion of covering $N_{\mathcal{X}}(R, t, \beta, \mathcal{B})$ (Definition 3.7), and relate it to $\Lambda_{\beta}(\mathcal{X})$ (Lemma 3.7). Then we prove the main technical lemma that bounds the estimation error $|\text{cost}_z(R, C, \mathcal{B}, h) - \text{cost}_z(S_R, C, \mathcal{B}, h)|$ for every ring R : Lemma 3.8, which is a weak version of Lemma 3.3. Compared to Lemma 3.3, the sample size in Lemma 3.8 contains an additional term $\log n_R$, which can be removed by applying the iterative size reduction approach of [BJKW21a].

To prove Lemma 3.8, we consider two cases: $t_R(C) = 0$ and $t_R(C) \in [k]$, and discuss the induced errors respectively (Lemmas 3.9 and 3.10). The case of $t_R(C) = 0$ is easier and can be shown by McDiarmid’s Inequality (Theorem 3.3). The case of $t_R(C) \in [k]$ is much more involved. We extend the flow idea in [CL19] to handle assignment constraints. Our idea is to show that $\text{cost}_z(S_R, C, \mathcal{B}, h)$ concentrates on its expectation $\mathbb{E}_S[\text{cost}_z(S_R, C, \mathcal{B}, h)]$ (Lemma 3.11) and to show the expectation is very close to $\text{cost}_z(R, C, \mathcal{B}, h)$ (Lemma 3.12). The proof of the former about the expectation follows easily from concentration inequalities, but that of the latter one is much more difficult and constitutes a major part of our analysis. Let σ^* be an optimal assignment of h on R , i.e., $\text{cost}_z(R, C, \mathcal{B}, h) = \text{cost}_z^{\sigma^*}(R, C)$. We convert σ^* to $\sigma : S_R \times C \rightarrow \mathbb{R}_{\geq 0}$, by setting $\sigma(p, \cdot) := w_S(p) \cdot \sigma^*(p, \cdot)$ for $p \in S$ (as mentioned in Section 3.1.2), and can show that $\text{cost}_z^{\sigma}(S_R, C) \approx \text{cost}_z^{\sigma^*}(R, C)$ ((3.20) in the proof of Claim 6). We are done if $\sigma \sim h$, but unfortunately, this generally does not hold. Hence, we turn to show the existence of an assignment $\sigma' \sim (\mathcal{B}, h)$ on S such that $|\text{cost}_z^{\sigma}(S_R, C) - \text{cost}_z^{\sigma'}(S_R, C)|$ is small enough. This existence of such σ' is shown in Claim 6, and here we sketch the main technical ideas. We reduce the problem of bounding $|\text{cost}_z^{\sigma}(S_R, C) - \text{cost}_z^{\sigma'}(S_R, C)|$ to bounding the mass movement $\|\sigma - \sigma'\|_1$ from σ to σ' , based on a novel idea that we can safely ignore the total assignment cost of R and S to remote centers (denoted by C_{far}), and the difference between the assignment cost induced by σ and σ' to $C \setminus C_{\text{far}}$ is proportional to $\|\sigma - \sigma'\|_1$, due to the generalized triangle inequality (Lemma 3.5). Now it suffices to require that $\|\sigma - \sigma'\|_1 \leq \tilde{O}_z(\varepsilon^{z+1}|R|)$ for bounding $|\text{cost}_z^{\sigma}(S_R, C) - \text{cost}_z^{\sigma'}(S_R, C)|$. By the definition of $\text{Lip}(\mathcal{B})$, we only need to make sure that $\|h - h'\|_1$ (h' is induced by σ) is as small up to $\tilde{O}_z(\varepsilon^{z+1}|R|)/\text{Lip}(\mathcal{B})$ (Claim 8), which is again guaranteed by McDiarmid’s Inequality (Theorem 3.3). The additional factor $1/\text{Lip}(\mathcal{B})$ results in the term $\text{Lip}(\mathcal{B})^2$ in our coreset size.

For preparation, we introduce the following generalized triangle inequality.

LEMMA 3.5. (Generalized triangle inequality [BCJ⁺22, Lemma 2.1]) *Let $a, b, c \in \mathcal{X}$ and $z \geq 1$. For every $t \in (0, 1]$, the following inequalities hold:*

$$d^z(a, b) \leq (1 + t)^{z-1} d^z(a, c) + \left(1 + \frac{1}{t}\right)^{z-1} d^z(b, c),$$

and

$$|d^z(a, c) - d^z(b, c)| \leq t \cdot d^z(a, c) + \left(\frac{3z}{t}\right)^{z-1} \cdot d^z(a, b).$$

We also use the following concentration inequality for analysis.

THEOREM 3.3. (McDiarmid’s Inequality [vH14, Theorem 3.11]) *Let E be a ground set and $n \geq 1$ be an integer. Let $g : E^n \rightarrow \mathbb{R}$ be a function satisfying that for any sequence $(x_1, x_2, \dots, x_n) \in E^n$, there exists a universal constant $\delta_i > 0$ for each $i \in [n]$ such that*

$$\delta_i \geq \sup_{y \in E} g(x_1, \dots, x_{i-1}, y, x_{i+1}, \dots, x_n) - \inf_{y \in E} g(x_1, \dots, x_{i-1}, y, x_{i+1}, \dots, x_n).$$

Then for independent random variables X_1, \dots, X_n , we have for every $t > 0$,

$$\Pr[g(X_1, \dots, X_n) - \mathbb{E}_{S_R}[g(X_1, \dots, X_n)] \geq t] \leq e^{-\frac{2t^2}{\sum_{i \in [n]} \delta_i^2}}.$$

Suppose $R \subseteq \text{ring}(c_i^*, r, 2r)$ for some $r > 0$. For preparation, we have the following observation that shows that $\text{cost}_z(P, C, \mathcal{B}, h)$ has some Lipschitz property. The proof is the same as that in [HJLW23, Lemma 3.6].

OBSERVATION 2. (Lipschitz property of $\text{cost}_z(P, C, \mathcal{B}, h)$ on P) Let $P, Q \subseteq \mathcal{X}$ be two weighted sets with $w_P(P) = w_Q(Q)$ and $c \in \mathcal{X}$ be a center point. For any k -center set $C \in \mathcal{X}^k$, any assignment constraint (\mathcal{B}, h) and any $\varepsilon \in (0, 1]$, we have

$$|\text{cost}_z(P, C, \mathcal{B}, h) - \text{cost}_z(Q, C, \mathcal{B}, h)| \leq \varepsilon \cdot \text{cost}_z(P, C, \mathcal{B}, h) + \left(\frac{6z}{\varepsilon}\right)^{z-1} \cdot (\text{cost}_z(P, c) + \text{cost}_z(Q, c)).$$

This observation is useful for providing an upper bound for $|\text{cost}_z(P, C, \mathcal{B}, h) - \text{cost}_z(Q, C, \mathcal{B}, h)|$. For instance, if $z = 1$ and $\text{cost}_1(P, c) + \text{cost}_1(Q, c) \ll \text{cost}_1(P, C, \mathcal{B}, h)$, we may have

$$|\text{cost}_1(P, C, \mathcal{B}, h) - \text{cost}_1(Q, C, \mathcal{B}, h)| \leq O(\varepsilon) \cdot \text{cost}_1(P, C, \mathcal{B}, h).$$

For any center set $C \in \mathcal{X}^k$, consider a mapping $\nu : C \rightarrow \mathbb{R}^d$ defined as follows: $\nu(c) = c_i^*$ for every $c \in C_{\text{far}}^R \cup C_{\text{close}}^R$, and $\nu(c) = c$ for the remaining centers $c \in C \setminus C_{\text{far}}^R \setminus C_{\text{close}}^R$. By definition, we know that $\nu(C)$ contains at most $t_R(C) + 1$ distinct centers. We first have the following lemma that enables us to only focus on the concentration for center sets $C \subset B(c_i^*, \frac{48zr}{\varepsilon})$.

LEMMA 3.6. (APPROXIMATION OF cost_z) For every weighted set $Q \subseteq R$ with total weight n_R , we have

$$\text{cost}_z(Q, C, \mathcal{B}, h) \in (1 \pm \frac{\varepsilon}{4}) \cdot (\text{cost}_z(Q, \nu(C), \mathcal{B}, h) + \phi(C, h)),$$

where

$$\phi(C, h) := \sum_{c \in C_{\text{far}}^R} h_c \cdot d^z(c, c_i^*),$$

which is independent of the choice of Q .¹¹

Proof. The lemma is implied by the proof of Lemma 4.4 of [BCJ+22]. For completeness, we provide proof here.

By the construction of $\nu(C)$, we can check that for every $p \in Q$ and $c \in C$,

1. if $c \in C_{\text{far}}^R$, $d^z(p, c) \in (1 \pm \frac{\varepsilon}{4}) \cdot (d^z(p, \nu(c)) + d^z(c, \nu(c)))$ by the definition of C_{far}^R ;
2. if $c \in C_{\text{close}}^R$, $d^z(p, c) \in (1 \pm \frac{\varepsilon}{4}) \cdot d^z(p, \nu(c))$ by the definition of C_{close}^R ;
3. if $c \in C \setminus C_{\text{far}}^R \setminus C_{\text{close}}^R$; $d^z(p, c) = d^z(p, \nu(c))$.

Let $\sigma \sim (\mathcal{B}, h)$ be an arbitrary assignment function. Combining the above properties, we have

$$\text{cost}^\sigma(Q, C) \in (1 \pm \frac{\varepsilon}{4}) \cdot (\text{cost}^\sigma(Q, \nu(C)) + \phi(C, h)),$$

which implies the lemma. \square

We also need to define another notion of covering for ring R , which aims to cover both the metric space \mathcal{X} and the hyper-parameter space of the feasible capacity constraints induced by \mathcal{B} (Definition 3.7). A similar idea appeared in [BCJ+22, Lemma 4.4] but it concerns capacity constraints only in Euclidean spaces. Recall that $\text{conv}(\mathcal{B}^o) = \text{conv}(\mathcal{B} \cup \{0\})$. Let Φ denote the collection of all pairs $(C, h) \in \mathcal{X}^k \times (n_R \cdot \text{conv}(\mathcal{B}^o))$. We partition Φ into $k + 1$ sub-collections Φ_t for $t \in \{0, 1, \dots, k\}$ where Φ_t is the collection of $(C, h) \in \Phi$ such that R is a t -level ring w.r.t. C ($t_R(C) = t$).

¹¹ $\phi(C, h)$ plays the same role as $\Delta(C)$ defined in [BCJ+22, Lemma 4.4], that captures the clustering cost of points to remote centers in C_{far} .

DEFINITION 3.7. (Coverings and covering numbers with assignment structure constraints) Let $\mathcal{B} \subseteq \Delta_k$ be an assignment structure constraint. Let $R \subset \text{ring}(c_i^*, r, 2r)$ be a ring of $n_R \geq 1$ points. Let $t \in [k]$ be an integer. We say a collection $\mathcal{F} \subset \Phi_t$ is a (t, α) -covering w.r.t. (R, \mathcal{B}) if for every $(C, h) \in \Phi_t$, there exists $(C', h') \in \mathcal{F}$ such that for every weighted set $Q \subseteq R$ with $w_Q(Q) = n_R$,

$$\text{cost}_z(Q, C, \mathcal{B}, h) \in (1 \pm (\beta + \varepsilon)) \cdot (\text{cost}_z(Q, C', \mathcal{B}, h') + \phi(C, h)) \pm \beta n_R r^z.$$

Define $N_{\mathcal{X}}(R, t, \beta, \mathcal{B})$ to be the minimum cardinality $|\mathcal{F}|$ of any (t, β) -covering \mathcal{F} w.r.t. (R, \mathcal{B}) .

As the covering number $N_{\mathcal{X}}(R, t, \beta, \mathcal{B})$ becomes larger, the family Φ_t is likely to induce more types of $\text{cost}_z(Q, C, \mathcal{B}, h)$'s. Although the definition of $N_{\mathcal{X}}(R, t, \beta, \mathcal{B})$ is based on the clustering cost, which looks quite different from Definition 3.4, we have the following lemma that relates the two notions of covering numbers.

LEMMA 3.7. (Relating two types of covering numbers) Let \mathcal{B} be an assignment structure constraint. For every $\beta > 0$ and $t \in [k]$, we have

$$\begin{aligned} \log N_{\mathcal{X}}(R, t, \beta, \mathcal{B}) &\leq O(t \cdot \log N_{\mathcal{X}}(n_R, \beta)) + zk \cdot \log(\text{Lip}(\mathcal{B}) \cdot z\varepsilon^{-1}\beta^{-1}) \\ &\leq O(\Lambda_{\beta}(\mathcal{X}) \cdot t \log n_R + zk \cdot \log(\text{Lip}(\mathcal{B}) \cdot z\varepsilon^{-1}\beta^{-1})). \end{aligned}$$

Moreover, when $\mathcal{B} = \Delta_k$, we have

$$\log N_{\mathcal{X}}(R, t, \beta, \mathcal{B}) \leq O(\Lambda_{\beta}(\mathcal{X}) \cdot t \log n_R + zk \cdot \log(zk\varepsilon^{-1}\beta^{-1})).$$

Proof. The proof can be found in Section C. \square

Now we are ready to prove Lemma 3.3. We first prove the following weak version.

LEMMA 3.8. (A weak version of Lemma 3.3) For each $i \in [k]$ and ring $R \in \mathcal{R}_i$ of size $n_R \geq 1$, suppose for every $t \in [k+1]$,

$$(3.10) \quad \Gamma_R \geq 2^{O(z \log z)} \cdot \lambda_R \cdot \text{Lip}(\mathcal{B})^2 \cdot \varepsilon^{-2z} \cdot (k\varepsilon^{-1} \log \delta^{-1} + \frac{k}{t} \log(N_{\mathcal{X}}(R, t, \varepsilon, \mathcal{B}) \cdot 2^k \delta^{-1})) \cdot \log^3(k\varepsilon^{-1}),$$

and specifically, when $\mathcal{B} = \Delta_k$,

$$(3.11) \quad \Gamma_R \geq 2^{O(z \log z)} \cdot \lambda_R \cdot \varepsilon^{-2z} \cdot (k\varepsilon^{-1} \log \delta^{-1} + \frac{k}{t} \log(N_{\mathcal{X}}(R, t, \varepsilon, \mathcal{B}) \cdot 2^t \delta^{-1})) \cdot \log^3(k\varepsilon^{-1}).$$

With probability at least $1 - \frac{\delta}{|\mathcal{R}|}$, for every k -center set $C \in \mathcal{X}^k$ and capacity constraint $h \in n_R \cdot \text{conv}(\mathcal{B}^o)$, we have that

$$\begin{aligned} &|\text{cost}_z(R, C, \mathcal{B}, h) - \text{cost}_z(S_R, C, \mathcal{B}, h)| \\ &\leq \varepsilon (\text{cost}_z(R, C, \mathcal{B}, h) + \text{cost}_z(R, c_i^*)) + \left(\frac{t_R(C)}{10zk\lambda_R \log \varepsilon^{-1}}\right)^{\frac{1}{2}} \cdot \varepsilon \text{cost}_z(R, c_i^*). \end{aligned}$$

Combining with Lemma 3.7, we know that the required sample number in Lemma 3.8 is at most

$$2^{O(z \log z)} \cdot \text{Lip}(\mathcal{B})^2 \cdot (\Lambda_{\varepsilon}(\mathcal{X}) + k + \varepsilon^{-1}) \cdot k\varepsilon^{-2z} \cdot \log(n_R \cdot \text{Lip}(\mathcal{B}) \cdot \delta^{-1}) \log^4(k\varepsilon^{-1})$$

for general assignment structure constraint \mathcal{B} , and is at most

$$2^{O(z \log z)} \cdot (\Lambda_{\varepsilon}(\mathcal{X}) + \varepsilon^{-1}) \cdot k\varepsilon^{-2z} \cdot \log(n_R \cdot \text{Lip}(\mathcal{B}) \cdot \delta^{-1}) \log^4(k\varepsilon^{-1})$$

when $\mathcal{B} = \Delta_k$. Compared to Theorem 3.1, there is an additional term $\log n_R$ in the coreset size, which can be as large as $O(\log n)$. We will show how to remove this term later.

We first give the following lemma that solves the case of $(C, h) \in \Phi_0$ for Lemma 3.8.

LEMMA 3.9. (**Lemma 3.8 for Φ_0**) With probability at least $1 - \frac{\delta}{2|\mathcal{R}|}$, for all $(C, h) \in \Phi_0$, the following inequality holds:

$$|\text{cost}_z(R, C, \mathcal{B}, h) - \text{cost}_z(S_R, C, \mathcal{B}, h)| \leq \varepsilon (\text{cost}_z(R, C, \mathcal{B}, h) + \text{cost}_z(R, c_i^*)).$$

Proof. We claim that

$$(3.12) \quad |\text{cost}_z(R, \nu(C), \mathcal{B}, h) - \text{cost}_z(S_R, \nu(C), \mathcal{B}, h)| \leq \frac{\varepsilon}{4} (\text{cost}_z(R, \nu(C), \mathcal{B}, h) + \text{cost}_z(R, c_i^*)),$$

which implies that

$$\begin{aligned} \text{cost}_z(S_R, C, \mathcal{B}, h) &\in (1 \pm \frac{\varepsilon}{4}) \cdot (\text{cost}_z(S_R, \nu(C), \mathcal{B}, h) + \phi(C, h)) && \text{(Lemma 3.6)} \\ &\in (1 \pm \frac{\varepsilon}{2}) \cdot (\text{cost}_z(R, \nu(C), \mathcal{B}, h) + \phi(C, h)) \pm \frac{\varepsilon}{2} \cdot \text{cost}_z(R, c_i^*) && \text{(Ineq. (3.12))} \\ &\in (1 \pm \varepsilon) \cdot \text{cost}_z(R, C, \mathcal{B}, h) \pm \varepsilon \cdot \text{cost}_z(R, c_i^*). && \text{(Lemma 3.6)} \end{aligned}$$

Hence, it suffices to prove Inequality (3.12). Since $(C, h) \in \Phi_0$ implies that $t_R(C) = 0$, we have that $\nu(C) = c_i^*$, which implies that $\text{cost}_z(Q, \nu(C), \mathcal{B}, h) = \text{cost}_z^{(n_R - \|h\|_1)}(Q, c_i^*)$. Then it is equivalent to prove that for all $0 \leq m_R \leq n_R$,

$$|\text{cost}_z^{(m_R)}(R, c_i^*) - \text{cost}_z^{(m_R)}(S_R, c_i^*)| \leq \frac{\varepsilon}{4} (\text{cost}_z^{(m_R)}(R, c_i^*) + \text{cost}_z(R, c_i^*)).$$

Since $\text{cost}_z^{(m_R)}(R, c_i^*) \geq 0$, it suffices to prove that for all $0 \leq m_R \leq n_R$,

$$|\text{cost}_z^{(m_R)}(R, c_i^*) - \text{cost}_z^{(m_R)}(S_R, c_i^*)| \leq \frac{\varepsilon}{4} \text{cost}_z(R, c_i^*).$$

Note that $\text{cost}_z(R, c_i^*) \geq n_R r^z$. Also note that for every $p \in R$, we have $d^z(p, c_i^*) \leq 2^z r^z$, which implies that for any $0 \leq m \leq m' \leq n_R$ with $m' - m \leq \frac{\varepsilon n_R}{2^z + 4}$,

$$\text{cost}_z^{(m)}(R, c_i^*) \leq \text{cost}_z^{(m')}(R, c_i^*) + \frac{\varepsilon}{16} \text{cost}_z(R, c_i^*),$$

and

$$\text{cost}_z^{(m)}(S_R, c_i^*) \leq \text{cost}_z^{(m')}(S_R, c_i^*) + \frac{\varepsilon}{16} \text{cost}_z(R, c_i^*).$$

Then it suffices to prove that for $m_R = 0, \frac{\varepsilon n_R}{2^z + 4}, \frac{2\varepsilon n_R}{2^z + 4}, \dots, n_R$, the following inequality holds:

$$(3.13) \quad |\text{cost}_z^{(m_R)}(R, c_i^*) - \text{cost}_z^{(m_R)}(S_R, c_i^*)| \leq \frac{\varepsilon}{8} \cdot n_R r^z,$$

i.e., we only need to consider $O(2^z \varepsilon^{-1})$ different values of m_R .

By Theorem 3.2 and the definition of λ_R , we know that $\lambda_R \geq (\frac{\varepsilon}{6z})^z \cdot \frac{1}{k \cdot \log(48z\varepsilon^{-1})}$. Consequently, we have

$$(3.14) \quad \begin{aligned} \Gamma_R &\geq 2^{O(z \log z)} \cdot \lambda_R \cdot \varepsilon^{-2z} \cdot (k\varepsilon^{-1} \log \delta^{-1} + \log(N_{\mathcal{X}}(R, k, \varepsilon, \mathcal{B}) \cdot 2^k \delta^{-1})) \cdot \log^3(k\varepsilon^{-1}) && \text{(Ineq. (3.10))} \\ &\geq 2^{O(z \log z)} \cdot \varepsilon^{-z-1} \cdot \log(|\mathcal{R}| \cdot \delta^{-1} \varepsilon^{-1}) \\ &\geq 2^{O(z \log z)} \cdot \varepsilon^{-z-1} \cdot \log(|\mathcal{R}| \cdot \delta^{-1} \varepsilon^{-1}). && (z \geq 1) \end{aligned}$$

Suppose $S_R = S \cup \{q\}$ and $S'_R = S \cup \{q'\}$ where $S \subseteq R$, $|S| = \Gamma_R - 1$ and $q \neq q' \in R$. We know that $|\text{cost}_z^{(m)}(S_R, c_i^*) - \text{cost}_z^{(m)}(S'_R, c_i^*)| \leq \frac{2n_R}{\Gamma_R} \cdot 2^z r^z$. By construction, S_R consists of Γ_R i.i.d. uniform samples. Hence, we can apply Theorem 3.3 (McDiarmid's Inequality) and obtain that for every $t > 0$,

$$\Pr \left[\left| \text{cost}_z^{(m_R)}(S_R, c_i^*) - \mathbb{E}_{S_R} \left[\text{cost}_z^{(m_R)}(S_R, c_i^*) \right] \right| \geq t \right] \leq e^{-\frac{2t^2}{\Gamma_R \cdot (\frac{2n_R}{\Gamma_R} \cdot 2^z r^z)^2}}.$$

Since $\mathbb{E}_{S_R} \left[\text{cost}_z^{(m_R)}(S_R, c_i^*) \right] = \text{cost}_z^{(m_R)}(R, c_i^*)$, we conclude that Inequality (3.13) holds with probability at least $1 - \frac{\delta\varepsilon}{2^z + 10|\mathcal{R}|}$, which can be verified by letting $t = \frac{\varepsilon}{8} \cdot n_R r^z$ and the bound of Γ_R in Inequality (3.14). \square

Next, we consider the case of Φ_t for $t \in [k]$. The proof idea is to first show the concentration property of S_R w.r.t. a fixed pair $(C, h) \in \Phi_t$, and then discretize the parameter space Φ_t and use the union bound to handle all pairs $(C, h) \in \Phi_t$. We first propose the following key lemma for a specific pair $(C, h) \in \Phi_t$, which is a generalization of [BCJ⁺22, Lemma 4.3] based on [CL19, BFS21] by considering general $z \geq 1$, introducing assignment structure constraints \mathcal{B} , and carefully analyzing the induced error of rings w.r.t. different levels $t \in [k]$.

LEMMA 3.10. (Error analysis for rings w.r.t. a pair $(C, h) \in \Phi_t$) Fix $t \in [k]$ and $(C, h) \in \Phi_t$. Let $\alpha = (\frac{t}{10zk\lambda_R \log \varepsilon^{-1}})^{\frac{1}{2}} \varepsilon$. With probability at least $1 - \frac{\delta}{2kN_{\mathcal{X}}(R, t, \varepsilon, \mathcal{B}) \cdot |\mathcal{R}|}$, the following holds:

$$|\text{cost}_z(R, C, \mathcal{B}, h) - \text{cost}_z(S_R, C, \mathcal{B}, h)| \leq \varepsilon \text{cost}_z(R, C, \mathcal{B}, h) + \alpha \text{cost}_z(R, c_i^*).$$

Proof. By a similar argument as in Lemma 3.9, it suffices to prove that

$$|\text{cost}_z(R, \nu(C), \mathcal{B}, h) - \text{cost}_z(S_R, \nu(C), \mathcal{B}, h)| \leq \alpha \text{cost}_z(R, c_i^*).$$

We define a function g that takes a weighted set $Q \subseteq R$ with $w_Q(Q) = n_R$ as input, and outputs $g(Q) = \text{cost}_z(Q, \nu(C), \mathcal{B}, h)$. Thus, it suffices to prove that with probability at least $1 - \frac{\delta}{2kN_{\mathcal{X}}(R, t, \varepsilon, \mathcal{B}) \cdot |\mathcal{R}|}$,

$$(3.15) \quad |g(R) - g(S_R)| \leq O(\alpha n_R r^z).$$

Actually, the construction and the following analysis of g is the key for extending the k -MEDIAN results in [CL19] to general $z \geq 1$. Our idea is to prove the following two lemmas, where the first one shows the concentration property of $g(S_R)$ and the second one shows the closeness of $\mathbb{E}_{S_R} [g(S_R)]$ and $g(R)$. The main difficulty is handling the additional assignment structure constraint \mathcal{B} .

LEMMA 3.11. (Concentration of $g(S_R)$) With probability at least $1 - \frac{\delta}{2kN_{\mathcal{X}}(R, t, \varepsilon, \mathcal{B}) \cdot |\mathcal{R}|}$,

$$|g(S_R) - \mathbb{E}_{S_R} [g(S_R)]| \leq \alpha n_R r^z.$$

LEMMA 3.12. (Closeness of $\mathbb{E}_{S_R} [g(S_R)]$ and $g(R)$) The following holds:

$$g(R) \leq \mathbb{E}_{S_R} [g(S_R)] \leq g(R) + \alpha n_R r^z.$$

Inequality (3.15) is a direct corollary of the above lemmas since with probability at least $1 - \frac{\delta}{2kN_{\mathcal{X}}(R, t, \varepsilon, \mathcal{B}) \cdot |\mathcal{R}|}$, we have that

$$|g(R) - g(S_R)| \leq |g(R) - \mathbb{E}_{S_R} [g(S_R)]| + |\mathbb{E}_{S_R} [g(S_R)] - g(S_R)| \leq 2\alpha n_R r^z.$$

Hence, it remains to prove these two lemmas. Note that $\alpha = (\frac{t}{10zk\lambda_R \log \varepsilon^{-1}})^{\frac{1}{2}} \varepsilon \leq 2^{O(z \log z)} k^{1/2} \varepsilon^{1-z/2}$ since $\lambda_R \geq (\frac{\varepsilon}{6z})^z \cdot \frac{1}{k \cdot \log(48z\varepsilon^{-1})}$. Then by Inequality (3.10), we note that

$$(3.16) \quad \Gamma_R \geq 2^{O(z \log z)} \cdot \text{Lip}(\mathcal{B})^2 \cdot \varepsilon^{-2z+2} \alpha^{-2} \cdot \log(N_{\mathcal{X}}(R, t, \varepsilon, \mathcal{B}) \cdot 2^k \delta^{-1}) \cdot \log(k\alpha^{-1} \varepsilon^{-1}),$$

and when $\mathcal{B} = \Delta_k$,

$$(3.17) \quad \Gamma_R \geq 2^{O(z \log z)} \cdot \varepsilon^{-2z+2} \alpha^{-2} \cdot \log(N_{\mathcal{X}}(R, t, \varepsilon, \mathcal{B}) \cdot 2^t \delta^{-1}) \cdot \log(k\alpha^{-1} \varepsilon^{-1}).$$

For ease of analysis, we slightly abuse the notation by using C to replace $\nu(C)$ in the following. Then we know that $C \subset B(c_i^*, \frac{48zr}{\varepsilon})$ and there are $k - t$ centers $c \in C$ located at c_i^* .

Proof. [Proof of Lemma 3.11] We first show the following Lipschitz property of g .

CLAIM 4. (Lipschitz property of g) For every two realizations S_R and S'_R of size Γ_R that differ by one sample, we have

$$|g(S_R) - g(S'_R)| \leq \frac{n_R}{\Gamma_R} \cdot (62z)^z \varepsilon^{-z+1} r^z.$$

Proof. Suppose $S_R = S \cup \{q\}$ and $S'_R = S \cup \{q'\}$ where $S \subseteq R$, $|S| = \Gamma_R - 1$ and $q \neq q' \in R$. Let $\sigma : S_R \times C \rightarrow \mathbb{R}_{\geq 0}$ with $\sigma \sim (\mathcal{B}, h)$ be an optimal assignment function such that $g^\sigma(S_R) = g(S_R)$. We define $\sigma' : S'_R \times C \rightarrow \mathbb{R}_{\geq 0}$ with $\sigma' \sim (\mathcal{B}, h)$ to be:

$$\sigma'(p, \cdot) = \sigma(p, \cdot), \forall p \in S, \text{ and } \sigma'(q', \cdot) = \sigma(q, \cdot).$$

We have

$$\begin{aligned} & g(S'_R) \\ & \leq g^{\sigma'}(S'_R) && \text{(by optimality)} \\ & = g^\sigma(S_R) + \sum_{c \in C: c \neq c_i^*} \sigma(q, c) \cdot (d^z(q', c) - d^z(q, c)) && \text{(Defns. of } \sigma' \text{ and } g) \\ & \leq g^\sigma(S_R) + \sum_{c \in C: c \neq c_i^*} \sigma(q, c) \cdot \left(\varepsilon \cdot d^z(q, c) + \left(\frac{3z}{\varepsilon}\right)^{z-1} d^z(q, q') \right) && \text{(Lemma 3.5)} \\ & \leq g^\sigma(S_R) + \sum_{c \in C: c \neq c_i^*} \sigma(q, c) \cdot \left(\varepsilon \cdot (d(c_i^*, c) + d(q, c_i^*))^z + \left(\frac{3z}{\varepsilon}\right)^{z-1} (4r)^z \right) && \text{(triangle ineq.)} \\ & \leq g^\sigma(S_R) + \sum_{c \in C: c \neq c_i^*} \sigma(q, c) \cdot \left(\varepsilon \cdot \left(\frac{48zr}{\varepsilon} + 2r\right)^z + \left(\frac{3z}{\varepsilon}\right)^{z-1} (4r)^z \right) && \text{(Defn. of } \bar{C}_{\text{far}}) \\ & \leq g^\sigma(S_R) + \sum_{c \in C: c \neq c_i^*} \sigma(q, c) \cdot (62z)^z \varepsilon^{-z+1} r^z \\ & \leq g^\sigma(S_R) + \frac{n_R}{\Gamma_R} \cdot (62z)^z \varepsilon^{-z+1} r^z. && (w_{S_R}(q) = \frac{n_R}{\Gamma_R}) \end{aligned}$$

By symmetry, we complete the proof. \square

By construction, S_R consists of Γ_R i.i.d. uniform samples. Hence, we can apply Theorem 3.3 (McDiarmid's Inequality) and obtain that for every $t > 0$,

$$\Pr [|g(S_R) - \mathbb{E}_{S_R} [g(S_R)]| \geq t] \leq e^{-\frac{2t^2}{\Gamma_R \cdot \left(\frac{n_R}{\Gamma_R} \cdot (62z)^z \varepsilon^{-z+1} r^z\right)^2}}.$$

Lemma 3.11 can be verified by letting $t = \alpha n_R r^z$ and Inequalities (3.16) and (3.17). \square

Proof. [Proof of Lemma 3.12] We first show the easy direction, say $g(R) \leq \mathbb{E}_{S_R} [g(S_R)]$, which is guaranteed by the following convexity of g . For each possible realization of S_R , let μ_{S_R} denote the realized probability of S_R and $\sigma_{S_R} : S_R \times C \rightarrow \mathbb{R}_{\geq 0}$ with $\sigma_{S_R} \sim (\mathcal{B}, h)$ be an optimal assignment function such that $g^{\sigma_{S_R}}(S_R) = g(S_R)$. For the convenience of the argument, we extend the domain of σ_{S_R} to be $\sigma_{S_R} : R \times C \rightarrow \mathbb{R}_{\geq 0}$ where $\sigma_{S_R}(p, \cdot) = 0$ if $p \in R \setminus S_R$. By definition, we know that

$$(3.18) \quad \mathbb{E}_{S_R} [g(S_R)] = \mathbb{E}_{S_R} [g^{\sigma_{S_R}}(S_R)] = \sum_{S_R} \mu_{S_R} \cdot g(S_R).$$

Now consider an assignment function $\sigma : R \times C \rightarrow \mathbb{R}_{\geq 0}$ obtained by adding up σ_{S_R} , i.e., $\sigma = \mathbb{E}_{S_R} [\sigma_{S_R}] = \sum_{S_R} \mu_{S_R} \cdot \sigma_{S_R}$. We first show that $\sigma \sim (\mathcal{B}, h)$ is a feasible assignment function on R . On one hand, since $\sigma_{S_R} \sim \mathcal{B}$, we know that for every $p \in S_R$, $\sigma_{S_R}(p, \cdot) \in w_{S_R}(p) \cdot \mathcal{B}$. Thus, we have for every $p \in R$,

$$\sigma(p, \cdot) = \sum_{S_R} \mu_{S_R} \cdot \sigma_{S_R}(p, \cdot) \in \sum_{S_R} \mu_{S_R} \cdot w_{S_R}(p) \cdot \mathcal{B} = \mathcal{B},$$

since $\mathbb{E}_{S_R} [w_{S_R}(p)] = 1$. On the other hand, since $\sigma \sim h$, it is obvious that $\sigma \sim h$ since $\sum_{S_R} \mu_{S_R} = 1$. Thus, we obtain the following inequality

$$\begin{aligned} g(R) & \leq g^\sigma(R) && \text{(by optimality)} \\ & = \sum_{S_R} \mu_{S_R} \cdot g(S_R) && \text{(by linearity)} \\ & = \mathbb{E}_{S_R} [g(S_R)]. && \text{(Eq. (3.18))} \end{aligned}$$

Next, we show the difficult direction that $\mathbb{E}_{S_R} [g(S_R)] \leq (1 + \varepsilon) \cdot g(R) + \alpha n_R r^z$. We first have the following claim showing that the maximum difference between $g(S_R)$ and $g(R)$ is bounded.

CLAIM 5. (Uniform upper bound of $g(S_R)$) *The following holds*

$$g(S_R) \leq g(R) + (62z)^z \varepsilon^{-z+1} n_R r^z.$$

Proof. The proof is almost identical to that of Claim 4. The only difference is that R and S_R may differ by n points instead of one point with weight $\frac{n_R}{\Gamma_R}$, which results in a total difference on weights

$$\sum_{p \in R} |1 - w_{S_R}(p)| \leq 2n_R,$$

where we let $w_{S_R}(p) = 0$ for $p \notin S_R$. □

Let $\sigma^* : R \times C \rightarrow \mathbb{R}_{\geq 0}$ with $\sigma^* \sim (\mathcal{B}, h)$ be an optimal assignment function such that $g^{\sigma^*}(R) = g(R)$. We first construct $\pi : S_R \times C \rightarrow \mathbb{R}_{\geq 0}$ as follows: for every $p \in S_R$,

$$\pi(p, \cdot) = \frac{n_R}{\Gamma_R} \cdot \sigma^*(p, \cdot).$$

Note that $\pi \sim \mathcal{B}$ always holds, but $\pi \sim h$ may not hold. We then construct another assignment function $\pi' : S_R \times C \rightarrow \mathbb{R}_{\geq 0}$ with $\|\pi'(p, \cdot)\|_1 = \|\pi(p, \cdot)\|_1$ for every $p \in S_R$ and $\pi' \sim (\mathcal{B}, h)$ such that the total mass movement $\sum_{p \in S_R} \|\pi(p, \cdot) - \pi'(p, \cdot)\|_1$ from π to π' is minimized. We have the following claim.

CLAIM 6. (Upper bound of $g(S_R)$ w.h.p.) *With probability at least $1 - \frac{\varepsilon^{z-1}\alpha}{2 \cdot (62z)^z}$, we have*

$$g(S_R) \leq g^{\pi'}(S_R) \leq g(R) + 0.5\alpha n_R r^z.$$

Proof. $g(S_R) \leq g^{\pi'}(S_R)$ holds by optimality, and it remains to prove $g^{\pi'}(S_R) \leq g^{\sigma^*}(R) + 0.5\alpha n_R r^z$. Since $g^{\sigma^*}(R) = g(R)$ by definition, note that

$$\begin{aligned} & g^{\pi'}(S_R) - g^{\sigma^*}(R) \\ &= \sum_{c \in C: c=c_i^*} \left(\sum_{p \in S_R} \pi'(p, c) \cdot d^z(c, c_i^*) - \sum_{p \in R} \sigma^*(p, c) \cdot d^z(c, c_i^*) \right) \\ &+ \sum_{c \in C: c \neq c_i^*} \left(\sum_{p \in S_R} \pi'(p, c) \cdot d^z(p, c) - \sum_{p \in R} \sigma^*(p, c) \cdot d^z(p, c) \right) \quad (\text{Defn. of } g) \\ &= \sum_{c \in C: c \neq c_i^*} \left(\sum_{p \in S_R} \pi'(p, c) \cdot d^z(p, c) - \sum_{p \in R} \sigma^*(p, c) \cdot d^z(p, c) \right). \quad (\pi', \sigma^* \sim h) \end{aligned}$$

Hence, it is equivalent to proving that with probability at least $1 - \frac{\varepsilon^{z-1}\alpha}{2 \cdot (62z)^z}$,

$$(3.19) \quad \sum_{c \in C: c \neq c_i^*} \left(\sum_{p \in S_R} \pi'(p, c) \cdot d^z(p, c) - \sum_{p \in R} \sigma^*(p, c) \cdot d^z(p, c) \right) \leq 0.5\alpha n_R r^z.$$

The idea is to use π as an intermediary and we define another helper function ϕ that takes S_R as input, and outputs

$$\phi(S_R) := \sum_{c \in C: c \neq c_i^*} \sum_{p \in S_R} \pi(p, c) \cdot d^z(p, c) + \sum_{c \in C: c \neq c_i^*} (h_c - \sum_{p \in S_R} \pi(p, c)) \cdot d^z(c, c_i^*).$$

To prove Inequality (3.19), it suffices to prove that each of the following two inequalities holds with probability at least $1 - \frac{\varepsilon^{z-1}\alpha}{4 \cdot (62z)^z}$.

$$(3.20) \quad \phi(S_R) - \sum_{c \in C: c \neq c_i^*} \sum_{p \in R} \sigma^*(p, c) \cdot d^z(p, c) \leq 0.25\alpha n_R r^z,$$

and

$$(3.21) \quad \sum_{c \in C: c \neq c_i^*} \sum_{p \in S_R} \pi'(p, c) d^z(p, c) - \phi(S_R) \leq 0.25\alpha n_R r^z.$$

Proof of Inequality (3.20) By linearity and the definition of π , we know that $\mathbb{E}_{S_R}[\phi(S_R)] = \sum_{c \in C: c \neq c_i^*} \sum_{p \in R} \sigma^*(p, c) \cdot d^z(p, c)$. Similar to g , we show the following Lipschitz of ϕ .

CLAIM 7. (Lipschitz property of ϕ) For every two realizations S_R and S'_R of size Γ_R that differ by one sample, we have

$$|\phi(S_R) - \phi(S'_R)| \leq \frac{2n_R}{\Gamma_R} \cdot (62z)^z \varepsilon^{-z+1} r^z.$$

Proof. Suppose the only different points are $p \in S_R$ and $q \in S'_R$. By definition, we know that

$$\begin{aligned} & |\phi(S_R) - \phi(S'_R)| \\ &= \frac{n_R}{\Gamma_R} \cdot \left| \sum_{c \in C: c \neq c_i^*} \sigma^*(p, c) \cdot (d^z(p, c) - d^z(c, c_i^*)) - \sigma^*(q, c) \cdot (d^z(q, c) - d^z(c, c_i^*)) \right| \quad (\text{Defns. of } \phi \text{ and } \pi') \\ &\leq \frac{2n_R}{\Gamma_R} \cdot (62z)^z \varepsilon^{-z+1} r^z, \quad (\text{Proof of Claim 4}) \end{aligned}$$

which completes the proof. \square

Now similar to Lemma 3.11, we can apply Theorem 3.3 (McDiarmid’s Inequality) and obtain that for every $t > 0$,

$$\Pr [|\phi(S_R) - \mathbb{E}_{S_R}[\phi(S_R)]| \geq t] = \Pr [|\phi(S_R)| \geq t] \leq e^{-\frac{2t^2}{\Gamma_R \cdot (\frac{2n_R}{\Gamma_R} \cdot (62z)^z \varepsilon^{-z+1} r^z)^2}}.$$

Inequality (3.20) can be verified by letting $t = 0.25\alpha n_R r^z$ and Inequalities (3.16) and (3.17).

Proof of Inequality (3.21) We first consider the general assignment structure constraint \mathcal{B} . For analysis, we construct a function ψ that takes S_R as input, and outputs

$$\psi(S_R) := \sum_{c \in C} \left| \sum_{p \in S_R} \pi(p, c) - h_c \right|,$$

i.e., the total capacity difference between π and σ^* . We have the following claim showing that $\psi(S_R)$ is likely to be small.

CLAIM 8. ($\psi(S_R)$ is tiny w.h.p.) With probability at least $1 - \frac{\varepsilon^{z-1}\alpha}{4 \cdot (62z)^z}$, we have

$$\psi(S_R) \leq \frac{\varepsilon^{z-1}\alpha n_R}{4 \cdot \text{Lip}(\mathcal{B}) \cdot (62z)^z}.$$

Proof. For every string $s \in \{+1, -1\}^k$, we define a function ψ_s that takes S_R as input, and outputs

$$\psi_s(S_R) := \sum_{c \in C} s_c \cdot \left(\sum_{p \in S_R} \pi(p, c) - h_c \right).$$

Note that $\psi(S_R) = \max_{s \in \{+1, -1\}^k} \psi_s(S_R)$. It is equivalent to proving that with probability at least $1 - \frac{\varepsilon^{z-1}\alpha}{4 \cdot (62z)^z}$, for all $s \in \{+1, -1\}^k$,

$$(3.22) \quad \psi_s(S_R) \leq \frac{\varepsilon^{z-1}\alpha n_R}{4 \cdot \text{Lip}(\mathcal{B}) \cdot (62z)^z}.$$

By the union bound, it suffices to prove that for any $s \in \{+1, -1\}^k$, Inequality (3.22) holds with probability at least $1 - \frac{\varepsilon^{z-1}\alpha}{2^{k+2} \cdot (62z)^z}$.

Fix a string $s \in \{+1, -1\}^k$ and note that $\mathbb{E}_{S_R} [\psi_s(S_R)] = 0$. Also noting that for any two realizations S_R and S'_R of size Γ_R that differ by one sample, we have

$$|\psi_s(S_R) - \psi_s(S'_R)| \leq \frac{2n_R}{\Gamma_R},$$

i.e., $\psi_s(S_R)$ is $\frac{2n_R}{\Gamma_R}$ -Lipschitz. We again apply Theorem 3.3 (McDiarmid's Inequality) and obtain that for every $t > 0$,

$$\Pr[|\psi_s(S_R) - \mathbb{E}_{S_R} [\psi_s(S_R)]| \geq t] = \Pr[|\psi_s(S_R)| \geq t] \leq e^{-\frac{2t^2}{\Gamma_R \cdot (2n_R/\Gamma_R)^2}}.$$

Inequality (3.22) can be verified by letting $t = \frac{\varepsilon^{z-1}\alpha n_R}{4 \cdot \text{Lip}(\mathcal{B}) \cdot (62z)^z}$ and Inequality (3.16), which ensures the success probability to be at least $1 - \frac{\varepsilon^{z-1}\alpha}{2^{k+2} \cdot (62z)^z}$. \square

Let $h' \in n_R \cdot \mathcal{B}$ be the capacity constraint with $h'_c = \sum_{p \in S_R} \pi(p, c)$. We have that $\pi \sim (\mathcal{B}, h')$ and $\psi(S_R) = \|h' - h\|_1$. Let $\tau = \sum_{p \in S_R} \|\pi(p, \cdot) - \pi'(p, \cdot)\|_1$. Recall that $\pi(p, \cdot), \pi'(p, \cdot) \in \frac{n_R}{\Gamma_R} \cdot \mathcal{B}$ for every $p \in S_R$. We have that with probability at least $1 - \frac{\varepsilon^{z-1}\alpha}{4 \cdot (62z)^z}$,

$$(3.23) \quad \begin{aligned} \tau &\leq \text{Lip}\left(\frac{n_R}{\Gamma_R} \cdot \mathcal{B}\right) \cdot \|h' - h\|_1 && \text{(Defns. of Lip}(\mathcal{B}) \text{ and } \pi') \\ &= \text{Lip}(\mathcal{B}) \cdot \psi(S_R) && \text{(scale-invariant of Lip)} \\ &\leq \frac{\varepsilon^{z-1}\alpha n_R}{4 \cdot (62z)^z}, && \text{(Claim 8)} \end{aligned}$$

which implies that

$$\begin{aligned} &\sum_{c \in C: c \neq c_i^*} \sum_{p \in S_R} \pi'(p, c) \cdot d^z(p, c) - \phi(S_R) \\ &= \sum_{c \in C: c \neq c_i^*} \sum_{p \in S_R} (\pi'(p, c) - \pi(p, c)) \cdot d^z(p, c) - \sum_{c \in C: c \neq c_i^*} (h_c - \sum_{c \in C: c \neq c_i^*} \pi(p, c)) \cdot d^z(c, c_i^*) && \text{(Defn. of } \phi) \\ &= \sum_{c \in C: c \neq c_i^*} \sum_{p \in S_R} (\pi'(p, c) - \pi(p, c)) \cdot (d^z(p, c) - d^z(c, c_i^*)) && (\pi' \sim h) \\ &\leq \left| \sum_{c \in C: c \neq c_i^*} \sum_{p \in S_R} (\pi'(p, c) - \pi(p, c)) \right| \cdot \max_{p \in R, c \in C: c \neq c_i^*} (d^z(p, c) - d^z(c, c_i^*)) \\ &\leq \tau \cdot \max_{p \in R, c \in C: c \neq c_i^*} (d^z(p, c) - d^z(c, c_i^*)) && \text{(Defn. of } \tau) \\ &\leq \tau \cdot (62z)^z \varepsilon^{-z+1} r^z && \text{(Defn. of } \bar{C}_{\text{far}}) \\ &\leq 0.25\alpha n_R r^z, && \text{(Ineq. (3.23))} \end{aligned}$$

i.e., Inequality (3.21) holds for general assignment structure constraint \mathcal{B} .

When $\mathcal{B} = \Delta_k$, we can safely regard (C, h) as $(\widehat{C}, \widehat{h})$ where

- when $t \in [k-1]$, $\widehat{C} = \{c \in C : c \neq c_i^*\} \cup c_i^*$, and when $t = k$, $\widehat{C} = C$.

- $\widehat{h}_c = h_c$ for $c \neq c_i^*$, and $\widehat{h}_{c_i^*} = \sum_{c \in C: c=c_i^*} h_c$.

The only difference is that we only need to consider at most 2^{t+1} different functions ψ_s in the proof of Claim 8. This difference enables us to set Γ_R as in Inequality (3.17), which contains a factor of 2^t instead of 2^k in Inequality (3.16). The remaining proofs are the same, and Inequality (3.21) holds when $\mathcal{B} = \Delta_k$.

Thus, we complete the proof of Claim 6. \square

Now we are ready to prove $\mathbb{E}_{S_R} [g(S_R)] \leq g(R) + \alpha n_R r^z$. By Claims 5 and 6,

$$\begin{aligned} & \mathbb{E}_{S_R} [g(S_R)] \\ & \leq \frac{\varepsilon^{z-1} \alpha}{2 \cdot (62z)^z} \cdot (g(R) + (62z)^z \varepsilon^{-z+1} n_R r^z) + \left(1 - \frac{\varepsilon^{z-1} \alpha}{2 \cdot (62z)^z}\right) (g(R) + 0.5 \alpha n_R r^z) \\ & \leq g(R) + \alpha n_R r^z. \end{aligned}$$

Thus, we complete the proof of Lemma 3.12. \square

Overall, we complete the proof of Lemma 3.10. \square

Combining with Definition 3.7, we are ready to prove Lemma 3.8.

Proof. [Proof of Lemma 3.8] Lemma 3.9 shows the correctness of Lemma 3.8 for all $(C, h) \in \Phi_0$. Hence, it suffices to prove that for every $t \in [k]$, with probability at least $1 - \frac{\delta}{2k \cdot |\mathcal{R}|}$, the following holds for all $(C, h) \in \Phi_t$:

$$(3.24) \quad \begin{aligned} & |\text{cost}_z(R, C, \mathcal{B}, h) - \text{cost}_z(S_R, C, \mathcal{B}, h)| \\ & \leq \varepsilon (\text{cost}_z(R, C, \mathcal{B}, h) + \text{cost}_z(R, c_i^*)) + \left(\frac{t}{10zk\lambda_R \log \varepsilon^{-1}}\right)^{\frac{1}{2}} \cdot \varepsilon \text{cost}_z(R, c_i^*). \end{aligned}$$

Lemma 3.8 is a direct corollary by the union bound.

Fix $t \in [k]$ and let $\mathcal{F}_t \subset \Phi_t$ be a (t, ε) -covering of (R, \mathcal{B}) of size at most $N_{\mathcal{X}}(R, t, \varepsilon, \mathcal{B})$. Let $\alpha = \left(\frac{t}{10zk\lambda_R \log \varepsilon^{-1}}\right)^{\frac{1}{2}} \varepsilon$. By Lemma 3.10, with probability at least $1 - \frac{\delta}{2k \cdot |\mathcal{R}|}$, for every $(C, h) \in \mathcal{F}$, the following holds:

$$(3.25) \quad |\text{cost}_z(R, C, \mathcal{B}, h) - \text{cost}_z(S_R, C, \mathcal{B}, h)| \leq \varepsilon \text{cost}_z(R, C, \mathcal{B}, h) + \alpha \text{cost}_z(R, c_i^*).$$

For every $(C, h) \in \Phi_t$, there must exist $(C', h') \in \mathcal{F}_t$ such that

$$\text{cost}_z(R, C, \mathcal{B}, h) \in (1 \pm 2\varepsilon) \cdot (\text{cost}_z(R, C', \mathcal{B}, h') + \phi(C, h)) \pm \varepsilon n_R r^z,$$

and

$$\text{cost}_z(S_R, C, \mathcal{B}, h) \in (1 \pm 2\varepsilon) \cdot (\text{cost}_z(S_R, C', \mathcal{B}, h') + \phi(C, h)) \pm \varepsilon n_R r^z.$$

Combining the above two inequalities with Inequality (3.25), we conclude that

$$|\text{cost}_z(R, C, \mathcal{B}, h) - \text{cost}_z(S_R, C, \mathcal{B}, h)| \leq O(\varepsilon) (\text{cost}_z(R, C, \mathcal{B}, h) + \text{cost}_z(R, c_i^*)) + O(\alpha) \cdot \text{cost}_z(R, c_i^*),$$

which completes the proof of Inequality (3.24). \square

Finally, we show how to prove Lemma 3.3 by shaving off the term $\log n_R$ in Lemma 3.8. The main idea is to use the iterative size reduction approach of [BJKW21a] that has been widely employed in recent works [CSS21, CLSS22, BCJ⁺22] to obtain coresets of size independent of n_R . The technique is somewhat standard and we omit the details.

Proof. [Proof of Lemma 3.3] We can interpret S_R as an $L = O(\log^* |R|)$ -steps uniform sampling. Specifically, set $T = 2^{O(z \log z)} \cdot \lambda_R \cdot \text{Lip}(\mathcal{B})^2 \cdot \varepsilon^{-2z} \cdot \log^3(k\varepsilon^{-1}\delta^{-1})$, and for $i = 1, \dots, L$, let $S_R^{(i)}$ be a uniform sample of size $\tilde{O}((\Gamma_R + k)T \cdot (\log^{(i)}(|R|))^3)$ from $S_R^{(i-1)}$ where $\log^{(i)}$ denotes the i -th iterated logarithm and $S_R^{(0)} = R$. We remark that S_R has the same distribution as $S_R^{(L)}$.

Identical to the proof of Theorem 3.1 in [BJKW21a], we can obtain that with probability at least $1 - \frac{\delta}{|\mathcal{R}| \cdot |S_R^{(i-1)}|}$, for every (C, h) ,

$$(3.26) \quad \left| \text{cost}_z(S_R^{(i)}, C, \mathcal{B}, h) - \text{cost}_z(S_R^{(i-1)}, C, \mathcal{B}, h) \right| \leq O\left(\frac{\varepsilon \cdot (\text{cost}_z(R, C, \mathcal{B}, h) + \text{cost}_z(R, c_i^*))}{(\log^{(i)} |R|)^{\frac{1}{2}}}\right).$$

Summing (3.26) over $i = 1, 2, \dots, L$ and applying the union bound, we finish the proof. \square

3.5 Proof of Lemma 3.4: Error Analysis for Groups Finally, we analyze the error induced by groups. Throughout this section, we fix $i \in [k]$ and a k -center set $C \in \mathcal{X}^k$. Technically, we prove that the inclusion of assignment structure constraints \mathcal{B} does not pose additional challenges compared to [BCJ+22]. This is due to the subdivision of uncolored groups into at most k equivalent classes according to C_{far}^G (Lemma 3.14) and the projection of all centers to c_i^* for examination (Lemma 3.16). It is important to note that the introduction of \mathcal{B} does not impact the cost function when all centers are positioned at the same location (c_i^*). These new geometric observations conclude Lemma 3.4.

We recall the following group decomposition as in [BCJ+22].

DEFINITION 3.8. (Colored and uncolored groups [BCJ+22]) Fix $i \in [k]$ and a k -center set $C \in \mathcal{X}^k$. The collection of groups \mathcal{G}_i can be decomposed into colored groups and uncolored groups w.r.t. C such that

1. There are at most $O(k \log(z\varepsilon^{-1}))$ colored groups;
2. For every uncolored group $G \in \mathcal{G}_i$, center set C can be decomposed into two parts $C = C_{\text{close}}^G \cup C_{\text{far}}^G$ such that
 - For any $c \in C_{\text{close}}^G$ and $p \in G$, $d(c, c_i^*) < \frac{\varepsilon}{9z} \cdot d(p, c_i^*)$;
 - For any $c \in C_{\text{far}}^G$ and $p \in G$, $d(c, c_i^*) > \frac{24z}{\varepsilon} \cdot d(p, c_i^*)$.

Intuitively, for every uncolored group $G \in \mathcal{G}_i$, every center $c \in C$ is either very “close” or very “far” from G .

For every colored group $G \in \mathcal{G}_i$, we use the following lemma to upper bound the induced error, which is a corollary of Theorem 3.2 and Observation 2.

LEMMA 3.13. (Error analysis of colored groups) Let $G \in \mathcal{G}_i$ be a colored group and $h \in |G| \cdot \mathcal{B}$ be a feasible capacity constraint of G . The following holds:

$$|\text{cost}_z(G, C, \mathcal{B}, h) - \text{cost}_z(D_G, C, \mathcal{B}, h)| \leq \varepsilon \cdot \text{cost}_z(G, C, \mathcal{B}, h) + \frac{\varepsilon}{k \log(z\varepsilon^{-1})} \cdot \text{cost}_z(P_i, c_i^*).$$

Proof. By Theorem 3.2, we know that $\text{cost}_z(G, c_i^*) \leq (\frac{\varepsilon}{6z})^z \cdot \frac{\text{cost}_z(P_i, c_i^*)}{k \cdot \log(48z/\varepsilon)}$ since G is a group. Then by Observation 2, we have

$$\begin{aligned} & |\text{cost}_z(G, C, \mathcal{B}, h) - \text{cost}_z(D_G, C, \mathcal{B}, h)| \\ & \leq \varepsilon \cdot \text{cost}_z(G, C, \mathcal{B}, h) + \left(\frac{6z}{\varepsilon}\right)^{z-1} \cdot (\text{cost}(G, c_i^*) + \text{cost}(D_G, c_i^*)) && \text{(Ob. 2)} \\ & = \varepsilon \cdot \text{cost}_z(G, C, \mathcal{B}, h) + 2 \cdot \left(\frac{6z}{\varepsilon}\right)^{z-1} \cdot \text{cost}(G, c_i^*) && \text{(Defn. 3.5)} \\ & \leq \varepsilon \cdot \text{cost}_z(G, C, \mathcal{B}, h) + \frac{\varepsilon}{k \log(z\varepsilon^{-1})} \cdot \text{cost}_z(P_i, c_i^*), && \text{(Thm. 3.2)} \end{aligned}$$

which completes the proof. \square

Since there are at most $O(k \log(z\varepsilon^{-1}))$ colored groups, the above lemma provides an upper bound of the error induced by all colored groups $G \in \mathcal{G}_i$. It remains to upper bound the error of uncolored groups. To this end, we have the following lemma that further classifies the uncolored groups according to their C_{far}^G .

LEMMA 3.14. (Equivalent classes of uncolored groups w.r.t. C_{far}^G) There exists a partition $\mathcal{U}_1, \dots, \mathcal{U}_k$ of \mathcal{G}_i such that for every $j \in [k]$, it holds that $\forall G, G' \in \mathcal{U}_j, C_{\text{far}}^G = C_{\text{far}}^{G'}$.

Proof. To prove the desired result, it is sufficient to demonstrate that there are at most k distinct C_{far}^G for $G \in \mathcal{G}_i$. In other words, we need to show that $|\{C_{\text{far}}^G\}_{G \in \mathcal{G}_i}| \leq k$.

For any two groups $G, G' \in \mathcal{G}_i$, we can assume, without loss of generality, that

$$\forall p \in G, p' \in G', \quad d(p, c_i^*) \leq d(p', c_i^*).$$

According to the definition of C_{far}^G and $C_{\text{far}}^{G'}$ (see Definition 3.8), this implies that $C_{\text{far}}^{G'} \subseteq C_{\text{far}}^G$. Consequently, if we let C_1, \dots, C_l (with $|C_1| \leq |C_2| \leq \dots \leq |C_l|$) represent different center sets in $\{C_{\text{far}}^G\}_{G \in \mathcal{G}_i}$, it follows that $C_1 \subsetneq C_2 \subsetneq \dots \subsetneq C_l \subseteq C$. Therefore, we have $|C_1| < |C_2| < \dots < |C_l| \leq k$, which implies $|\{C_{\text{far}}^G\}_{G \in \mathcal{G}_i}| \leq k$. This completes the proof. \square

Our plan is to merge uncolored groups in each \mathcal{U}_j and analyze their error as a whole. Specifically, for every $j \in [k]$, we define $U_j := \bigcup_{G \in \mathcal{U}_j} G$ as the union of all groups in \mathcal{U}_j , and define $D_{U_j} = \bigcup_{G \in \mathcal{U}_j} D_G$ as the union of all two-point coresets of groups in \mathcal{U}_j . The following lemma provides an upper bound on the error for each U_j .

LEMMA 3.15. (Error analysis of D_{U_j} 's) Fix $j \in [k]$ and let $h \in |U_j| \cdot \text{conv}(\mathcal{B}^o)$ be a feasible capacity constraint of U_j , the following holds:

$$|\text{cost}_z(U_j, C, \mathcal{B}, h) - \text{cost}_z(D_{U_j}, C, \mathcal{B}, h)| \leq O(\varepsilon) \cdot \text{cost}_z(U_j, C, \mathcal{B}, h) + O\left(\frac{\varepsilon}{k \log(z\varepsilon^{-1})}\right) \cdot \text{cost}_z(P_i, c_i^*).$$

Lemma 3.15 provides a guarantee for each U_j similar to that of colored groups in Lemma 3.13. Since there are at most k such groups, we can use Lemma 3.15 to upper bound the error induced by all uncolored groups in \mathcal{G}_i . We defer the proof of Lemma 3.15 later.

Proof. [Proof of Lemma 3.4] Fix a center set $C \in \mathcal{X}^k$ and a capacity constraint $h \in |G[i]| \cdot \text{conv}(\mathcal{B}^o)$. Let $\mathcal{C}_i := \{G \in \mathcal{G}_i : G \text{ is colored}\}$ be the collection of all colored groups in \mathcal{G}_i . Suppose a collection $\{h^G \in |G| \cdot \text{conv}(\mathcal{B}^o) : G \in \mathcal{C}_i\} \cup \{h^{(j)} \in |U_j| \cdot \text{conv}(\mathcal{B}^o) : j \in [k]\}$ of capacity constraints satisfy that

$$\sum_{G \in \mathcal{C}_i} h^G + \sum_{j=1}^k h^{(j)} = h$$

and

$$(3.27) \quad \text{cost}_z(G[i], C, \mathcal{B}, h) = \sum_{G \in \mathcal{C}_i} \text{cost}_z(G, C, \mathcal{B}, h^G) + \sum_{j=1}^k \text{cost}_z(U_j, C, \mathcal{B}, h^{(j)}).$$

We have

$$\begin{aligned} & \text{cost}_z(D[i], C, \mathcal{B}, h) \\ & \leq \sum_{G \in \mathcal{C}_i} \text{cost}_z(D_G, C, \mathcal{B}, h^G) + \sum_{j=1}^k \text{cost}_z(D_{U_j}, C, \mathcal{B}, h^{(j)}) && \text{(by optimality)} \\ & \leq (1 + O(\varepsilon)) \sum_{G \in \mathcal{C}_i} \text{cost}_z(G, C, \mathcal{B}, h^G) + k \log(z\varepsilon^{-1}) \cdot O\left(\frac{\varepsilon \cdot \text{cost}_z(P_i, c_i^*)}{k \log(z\varepsilon^{-1})}\right) && \text{(Lemma 3.13)} \\ & \quad + (1 + O(\varepsilon)) \sum_{j=1}^k \text{cost}_z(U_j, C, \mathcal{B}, h^{(j)}) + k \cdot O\left(\frac{\varepsilon \cdot \text{cost}_z(P_i, c_i^*)}{k \log(z\varepsilon^{-1})}\right) && \text{(Lemma 3.15)} \\ & \leq (1 + O(\varepsilon)) \text{cost}_z(G[i], C, \mathcal{B}, h) + O(\varepsilon) \cdot \text{cost}_z(P_i, c_i^*) && \text{(Eq. (3.27))} \end{aligned}$$

Similarly, we can obtain $\text{cost}_z(G[i], C, \mathcal{B}, h) \leq (1 + O(\varepsilon)) \cdot \text{cost}_z(D[i], C, \mathcal{B}, h) + O(\varepsilon) \cdot \text{cost}_z(P_i, c_i^*)$ and thus complete the proof. \square

It remains to prove Lemma 3.15. In the following discussion, we fix $j \in [k]$. For the sake of simplicity, we slightly abuse the notation by denoting U_j as U , D_{U_j} as D , and the common C_{far}^G among all groups $G \in \mathcal{U}_j$ as C_{far} . Our strategy is similar to that in the ring case, which first shifts the focus from C to $\nu(C)$ based on Lemma 3.6. Note that for uncolored groups, $\nu(C) \equiv c_i^*$ since all centers in C are either in C_{close} or C_{far} , we have the following inequality:

$$\text{cost}_z(U, C, \mathcal{B}, h) \in \left(1 + \frac{\varepsilon}{4}\right) \cdot (\text{cost}_z(U, c_i^*, \mathcal{B}, h) + \phi(C, h)),$$

where $\phi(C, h) = \sum_{c \in C_{\text{far}}} \|h(\cdot, c)\|_1 \cdot d(c, c_i^*)$. Notice that

$$\text{cost}_z(U, c_i^*, \mathcal{B}, h) = \min_{\sigma \sim (\mathcal{B}, h)} \sum_{p \in U} \|\sigma(p, \cdot)\|_1 \cdot d(p, c_i^*),$$

where the constraint h no longer constrains the capacities of the centers but still indicates the number (fraction) of points in U that should be discarded as outliers. We formalize this result in the following lemma, which is a direct corollary of Lemma 3.6 combined with above observations.

LEMMA 3.16. (Reduce the capacity constraint of h on centers) *For a feasible capacity constraint $h \in |U| \cdot \text{conv}(\mathcal{B}^o)$ of U , let $m := w_U(U) - \|h\|_1$ denote the number of outliers, then it holds that*

$$\text{cost}_z(U, C, \mathcal{B}, h) \in \left(1 + \frac{\varepsilon}{4}\right) \cdot \left(\text{cost}_z^{(m)}(U, c_i^*) + \phi(C, h)\right),$$

Here $\text{cost}_z^{(m)}(U, c_i^*)$ is defined as

$$(3.28) \quad \text{cost}_z^{(m)}(U, c_i^*) := \min_{w: U \rightarrow \mathbb{R}_{\geq 0}: \|w\|_1 = m, w \leq w_U} \sum_{p \in U} (w_U(p) - w(p)) d^z(p, c_i^*),$$

where $w \leq w_U$ denotes that $\forall p \in U, w(p) \leq w_U(p)$.

Clearly, Lemma 3.16 works for D as well, i.e.,

$$\text{cost}_z(D, C, \mathcal{B}, h) \in \left(1 + \frac{\varepsilon}{4}\right) \cdot \left(\text{cost}_z^{(m)}(D, c_i^*) + \phi(C, h)\right).$$

We remark that $\text{cost}_z^{(m)}(U, c_i^*)$ defined in (3.28) is the same as the objective of (k, z) -CLUSTERING with m outliers, which has been investigated in prior studies, such as [HJLW23]. The following lemma upper bounds the error between $\text{cost}_z^{(m)}(U, c_i^*)$ and $\text{cost}_z^{(m)}(D, c_i^*)$.

LEMMA 3.17. (Error analysis for $\text{cost}_z^{(m)}(U, c_i^*)$) *For real number $0 \leq m \leq w_U(U)$, it holds that*

$$\left| \text{cost}_z^{(m)}(U, c_i^*) - \text{cost}_z^{(m)}(D, c_i^*) \right| \leq \varepsilon \cdot \text{cost}_z^{(m)}(U, c_i^*) + O\left(\frac{\varepsilon}{k \log(z\varepsilon^{-1})}\right) \cdot \text{cost}_z(P_i, c_i^*).$$

Proof. The lemma is implied by the proof of [HJLW23, Lemma 3.8]. For completeness, we present the proof here.

We separately prove the following two directions.

$$(3.29) \quad \text{cost}_z^{(m)}(D, c_i^*) \leq (1 + \varepsilon) \text{cost}_z^{(m)}(U, c_i^*) + O\left(\frac{\varepsilon}{k \log(z\varepsilon^{-1})}\right) \cdot \text{cost}_z(P_i, c_i^*),$$

$$(3.30) \quad \text{cost}_z^{(m)}(U, c_i^*) \leq (1 + \varepsilon) \text{cost}_z^{(m)}(D, c_i^*) + O\left(\frac{\varepsilon}{k \log(z\varepsilon^{-1})}\right) \cdot \text{cost}_z(P_i, c_i^*).$$

Proof of (3.29) Let $w^* : U \rightarrow \mathbb{R}_{\geq 0}$ be the solution of the optimization problem (3.28). Namely, it holds that $\|w^*\|_1 = m$, $w^*(p) \leq w_U(p), \forall p \in U$ and $\text{cost}_z^{(m)}(U, c_i^*) = \sum_{p \in U} (w_U(p) - w^*(p)) d^z(p, c_i^*)$. Recall that for every $G \in \mathcal{U}_j$ and $p \in G$, there exists a unique $\lambda_p \in [0, 1]$ such that $d^z(p, c_i^*) = \lambda_p \cdot d^z(p_{\text{close}}^G, c_i^*) + (1 - \lambda_p) \cdot d^z(p_{\text{far}}^G, c_i^*)$, then we construct $w' : D \rightarrow \mathbb{R}_{\geq 0}$ as follows: for every $G \in \mathcal{U}_j$ and $c \in C$,

$$w'(p_{\text{close}}^G) = \sum_{p \in G} \lambda_p \cdot w^*(p), \quad \text{and} \quad w'(p_{\text{far}}^G) = \sum_{p \in G} (1 - \lambda_p) \cdot w^*(p).$$

Note that $\|w'\|_1 = \|w^*\|_1 = m$, and $w^*(p) \leq w_U(p), \forall p \in U$ implies that

$$w'(p_{\text{close}}^G) \leq \sum_{p \in G} \lambda_p \cdot w_U(p) = w_D(p_{\text{close}}^G), \quad \text{and} \quad w'(p_{\text{far}}^G) \leq \sum_{p \in G} (1 - \lambda_p) \cdot w_U(p) = w_D(p_{\text{far}}^G).$$

Hence, w' is a feasible solution of the optimization problem $\text{cost}_z^{(m)}(D, c_i^*)$. We have

$$\begin{aligned} & \text{cost}_z^{(m)}(D, c_i^*) \\ & \leq \sum_{p \in D} (w_D(p) - w'(p)) d^z(p, c_i^*) \\ & = \sum_{G \in \mathcal{U}_j} ((w_D(p_{\text{close}}^G) - w'(p_{\text{close}}^G)) d^z(p_{\text{close}}^G, c_i^*) + (w_D(p_{\text{far}}^G) - w'(p_{\text{far}}^G)) d^z(p_{\text{far}}^G, c_i^*)) \\ & = \sum_{G \in \mathcal{U}_j} \left(\sum_{p \in G} (w_U(p) - w^*(p)) \lambda_p \cdot d^z(p_{\text{close}}^G, c_i^*) + \sum_{p \in G} (w_U(p) - w^*(p)) (1 - \lambda_p) \cdot d^z(p_{\text{far}}^G, c_i^*) \right) \\ & = \sum_{G \in \mathcal{U}_j} \sum_{p \in G} (w_U(p) - w^*(p)) \cdot d^z(p, c_i^*) \\ & = \text{cost}_z^{(m)}(U, c_i^*), \end{aligned}$$

which completes the proof of (3.29).

Proof of (3.30) Let $\{m_G\}_{G \in \mathcal{U}_j}$ be a sequence of positive real numbers such that $m = \sum_{G \in \mathcal{U}_j} m_G$ and $\text{cost}_z^{(m)}(D, c_i^*) = \sum_{G \in \mathcal{U}_j} \text{cost}_z^{(m_G)}(D_G, c_i^*)$. We then de a case study for every group $G \in \mathcal{U}_j$.

- If $m_G = 0$, then $\text{cost}_z^{(m_G)}(G, c_i^*) = \text{cost}_z(G, c_i^*)$ and $\text{cost}_z^{(m_G)}(D_G, c_i^*) = \text{cost}_z(D_G, c_i^*)$. By definition of two-point coresets (Definition 3.5), we have $\text{cost}_z(D_G, c_i^*) = \text{cost}_z(G, c_i^*)$ and hence $\text{cost}_z^{(m_G)}(G, c_i^*) = \text{cost}_z^{(m_G)}(D_G, c_i^*)$.
- If $m_G = w_D(D_G) = w_U(G)$, then $\text{cost}_z^{(m_G)}(G, c_i^*) = \text{cost}_z^{(m_G)}(D_G, c_i^*) = 0$.
- If $0 < m_G < w_D(D_G)$, we call such group a *special* group. By [HJLW23, Lemma 3.17], there are at most $O(1)$ special groups. By a similar argument as in the proof of Lemma 3.13, we can obtain the following inequality.

$$\text{cost}_z^{(m_G)}(G, c_i^*) \leq (1 + \varepsilon) \text{cost}_z^{(m_G)}(D_G, c_i^*) + \frac{\varepsilon \cdot \text{cost}_z(P_i, c_i^*)}{k \log(z\varepsilon^{-1})}.$$

Putting everything together, we have

$$\begin{aligned} & \text{cost}_z^{(m)}(U, c_i^*) \\ & \leq \sum_{G \in \mathcal{U}_j} \text{cost}_z^{(m_G)}(G, c_i^*) \\ & \leq (1 + \varepsilon) \sum_{G \in \mathcal{U}_j} \text{cost}_z^{(m_G)}(D_G, c_i^*) + O(1) \cdot \frac{\varepsilon \cdot \text{cost}_z(P_i, c_i^*)}{k \log(z\varepsilon^{-1})} \\ & = (1 + \varepsilon) \text{cost}_z^{(m)}(D, c_i^*) + O\left(\frac{\varepsilon}{k \log(z\varepsilon^{-1})}\right) \cdot \text{cost}_z(P_i, c_i^*) \end{aligned}$$

which completes the proof. \square

Proof. [Proof of Lemma 3.15] Lemma 3.15 is a direct corollary of Lemma 3.16 and 3.17. \square

4 Bounding the Lipschitz Constant $\text{Lip}(\mathcal{B})$

In this section, we give upper and lower bounds for the Lipschitz constant $\text{Lip}(\mathcal{B})$, for various notable cases of \mathcal{B} . In particular, we provide an upper bound $\text{Lip}(\mathcal{B})$ for an important class of assignment structure constraints \mathcal{B} , called matroid basis polytopes (Theorems 4.1 and 4.2). We also show that $\text{Lip}(\mathcal{B})$ may be unbounded with knapsack constraints (Theorem 4.3). We review the definitions and some useful properties as follows.

Firstly, we introduce a general assignment structure constraint captured by the so-called matroid basis polytope. For more information about matroid and matroid basis polytope, see the classic reference [Sch03].

DEFINITION 4.1. (Matroid) *Given a ground set E , a family \mathcal{M} of subsets of E is a matroid if*

- $\emptyset \in \mathcal{M}$;
- If $I \in \mathcal{M}$ and $I' \subset I$, then $I' \in \mathcal{M}$;
- If $I, I' \in \mathcal{M}$ and $|I| < |I'|$, then there must exist an element $a \in I' \setminus I$ such that $I \cup \{a\} \in \mathcal{M}$.

Each $I \in \mathcal{M}$ is called an independent set. The maximum size of an independent set is called the rank of \mathcal{M} , denoted by $\text{rank}(\mathcal{M})$. Each set $I \in \mathcal{M}$ of size equal to the rank is called a basis of \mathcal{M} .

Matroid is a very general combinatorial structure that generalizes many set systems including uniform matroid, partition matroid, laminar matroid, regular matroid, graphic matroid, transversal matroid and so on. Now we define the matroid basis polytope.

DEFINITION 4.2. (Matroid basis polytope) *Let E be a ground set and let \mathcal{M} be a matroid on E . For each basis $I \in \mathcal{M}$, let $e_I := \sum_{i \in I} e_i$ denote the indicator vector of I , where $e_i \in \mathbb{R}^{|E|}$ is the standard i -th unit vector. The matroid basis polytope $P_{\mathcal{M}}$ is the convex hull of the set*

$$P_{\mathcal{M}} := \{e_I : I \text{ is a basis of } \mathcal{M}\}.$$

The following definition provides another description of $P_{\mathcal{M}}$ by rank functions.

DEFINITION 4.3. (Rank function and $P_{\mathcal{M}}$) *Let E be a ground set of size $n \geq 1$ and let \mathcal{M} be a matroid on E . The rank function $\text{rank} : 2^E \rightarrow \mathbb{Z}_{\geq 0}$ of \mathcal{M} is defined as follows: for every $A \subseteq E$, $\text{rank}(A) = \max_{I \in \mathcal{M}} |I \cap A|$. Moreover, the matroid basis polytope can be equivalently defined by the following linear program (see e.g., [Sch03]):*

$$P_{\mathcal{M}} = \left\{ x \in \mathbb{R}_{\geq 0}^n : \sum_{i \in A} x_i \leq \text{rank}(A), \forall A \in \mathcal{A}; \|x\|_1 = \text{rank}(\mathcal{M}) \right\}.$$

Let $h \in P_{\mathcal{M}}$ be a point inside the matroid basis polytope. We say a subset $A \subseteq E$ is tight on h if $\sum_{i \in A} h_i = \text{rank}(A)$.

We also consider a specific type of matroid, called laminar matroid.

DEFINITION 4.4. (Laminar matroid) *Given a ground set E , a family of \mathcal{A} of subsets of E is called laminar if for every two subsets $A, B \in \mathcal{A}$ with $A \cap B \neq \emptyset$, either $A \subseteq B$ or $B \subseteq A$. Assume $E \in \mathcal{A}$ and we define the depth of a laminar \mathcal{A} to be the largest integer $\ell \geq 1$ such that there exists a sequence of subsets $A_1, \dots, A_\ell \in \mathcal{A}$ with $A_1 \subsetneq A_2 \subsetneq \dots \subsetneq A_\ell = E$.*

We say a family \mathcal{M} of subsets of E is a laminar matroid if there exists a laminar \mathcal{A} and a capacity function $u : \mathcal{A} \rightarrow \mathbb{Z}_{\geq 0}$ such that $\mathcal{M} = \{I \subseteq E : |I \cap A| \leq u(A), \forall A \in \mathcal{A}\}$.

By Definition 4.3, we know that $\text{rank}(A) \leq u(A)$ holds for every $A \in \mathcal{A}$. Specifically, we can see that a *uniform matroid* is a laminar matroid of depth 1 (i.e., there is a single cardinality constraint over the entire set E) and a *partition matroid* is a laminar matroid of depth 2 (i.e., there is a partition of E and each partition has a cardinality constraint).

Technical Overview For $\mathcal{B} = \Delta_k$ which is the unconstrained case, we have $\text{Lip}(\mathcal{B}) = 1$. However, the geometric structure of \mathcal{B} can indeed result in a significantly larger, even unbound, $\text{Lip}(\mathcal{B})$. In particular, we show that the value $\text{Lip}(\mathcal{B})$ is unbounded, even for a very simple \mathcal{B} defined by two knapsack constraints (see Theorem 4.3). On the other hand, we do manage to show $\text{Lip}(\mathcal{B}) \leq k - 1$ for matroid basis polytopes (Theorem 4.1). To analyze $\text{Lip}(\mathcal{B})$ for the matroid case, we first show in Lemma 4.1 that it suffices to restrict our attention to the case where the initial assignment σ corresponds to vertices of the polytope (or basis) and so does terminal assignment σ' (which we need to find). Our argument is combinatorial and heavily relies on constructing augmenting paths (Definition 4.7). Roughly speaking, an augmenting path is defined by a sequence of basis (I_1, \dots, I_m) and a sequence of elements (a_0, a_1, \dots, a_m) where each I_i is an initial basis and the following exchange property holds: $I_i \cup \{a_{i-1}\} \setminus \{a_i\} \in \mathcal{M}$ for all i . The exchange property allows us to perform a sequence of basis exchanges to transport the mass from a_0 to a_m . Hence, the key is to find such an augmenting path in which a_0 is an element with $h_{a_0} < h'_{a_0}$ and a_m is one with $h_{a_m} > h'_{a_m}$, for some $h, h' \in \mathcal{B}$. Then performing the exchange operation reduces the value of $\|h - h'\|_1$ by 2τ with total transportation cost $2m\tau$ for some small $\tau > 0$. Thus, $\text{Lip}(\mathcal{B})$ is bounded by the path length m . We show that $m \leq k - 1$ when \mathcal{B} is a matroid basis polytope (Claim 4.4). To show the existence of such an augmenting path, we leverage several combinatorial properties of a matroid, as well as the submodularity of the rank function, and the non-crossing property of tight subsets. We further improve the bound to $m \leq \ell + 1$ for laminar matroid basis polytopes of depth at most $\ell \geq 1$ (Claim 4.5), which results in bounds $\text{Lip}(\mathcal{B}) \leq 2$ for uniform matroids and $\text{Lip}(\mathcal{B}) \leq 3$ for partition matroids.

4.1 Lipschitz Constant for (Laminar) Matroid Basis Polytopes For ease of analysis, we first propose another Lipschitz constant on \mathcal{B} for general polytopes (Definition 4.6). For preparation, we need the following notion of coupling of distributions (see e.g., [MU17]).

DEFINITION 4.5. (Coupling of distributions) Given two distributions $(\mu, \mu') \in \Delta_m$ ($m \geq 1$), $\kappa : [m] \times [m] \rightarrow \mathbb{R}_{\geq 0}$ is called a coupling of (μ, μ') , denoted by $\kappa \vdash (\mu, \mu')$, if

- For any $i \in [m]$, $\sum_{j \in [m]} \kappa(i, j) = \mu_i$;
- For any $j \in [m]$, $\sum_{i \in [m]} \kappa(i, j) = \mu'_j$.

Now we define the Lipschitz constant for a restricted assignment transportation problem, in which the initial and terminal assignments are restricted to the vertices of polytope \mathcal{B} .

DEFINITION 4.6. (Lipschitz constant for restricted OAT) Let $\mathcal{B} \subseteq c \cdot \Delta_k$ for some $c > 0$ be a polytope. Let V denote the collection of vertices of \mathcal{B} . Let $h, h' \in \mathcal{B}$ be two points inside \mathcal{B} . Let $\mu \in \Delta_{|V|}$ denote a distribution on V satisfying that

$$\sum_{v \in V} \mu_v \cdot v = h.$$

We define

$$\widetilde{\text{OAT}}(\mathcal{B}, h, h', \mu) := \min_{\substack{\mu' \in \Delta_{|V|} : \sum_{v \in V} \mu'_v \cdot v = h' \\ \kappa \vdash (\mu, \mu')}} \sum_{v, v' \in V} \kappa(v, v') \cdot \|v - v'\|_1$$

to be the optimal assignment transportation cost constrained to vertices of \mathcal{B} . Define the Lipschitz constant for the restricted OAT on \mathcal{B} to be

$$\widetilde{\text{Lip}}(\mathcal{B}) := \max_{\substack{h, h' \in \mathcal{B} \\ \mu \in \Delta_{|V|} : \sum_{v \in V} \mu_v \cdot v = h}} \frac{\widetilde{\text{OAT}}(\mathcal{B}, h, h', \mu)}{\|h - h'\|_1}.$$

By definition, we know that $\widetilde{\text{Lip}}(\mathcal{B}) = \widetilde{\text{Lip}}(c \cdot \mathcal{B})$ for any $c > 0$, i.e., $\widetilde{\text{Lip}}$ is scale-invariant on \mathcal{B} . We will analyze $\widetilde{\text{Lip}}(P_{\mathcal{M}})$ for matroid basis polytopes in the later sections. Now, we show the following lemma that connects two Lipschitz constants.

LEMMA 4.1. (Relation between $\text{Lip}(\mathcal{B})$ and $\widetilde{\text{Lip}}(\mathcal{B})$) Let $\mathcal{B} \subseteq \Delta_k$ be a polytope. We have $\text{Lip}(\mathcal{B}) \leq \widetilde{\text{Lip}}(\mathcal{B})$.

Proof. Let V denote the collection of vertices of \mathcal{B} . We define a collection \mathcal{H} of assignment functions σ such that for every $\sigma \in \mathcal{H}$: for every $p \in [n]$, there exists a vertex $v \in V$ such that $\frac{\sigma(p, \cdot)}{\|\sigma(p, \cdot)\|_1} = v$, i.e., the assignment vector of each $p \in [n]$ is equal to a scale of some vertex of \mathcal{B} . We have the following claim.

CLAIM 9. (**An equivalent formulation of $\text{Lip}(\mathcal{B})$)** $\text{Lip}(\mathcal{B}) = \max_{\substack{h, h' \in \mathcal{B} \\ \sigma \in \mathcal{H}}} \frac{\text{OAT}(\mathcal{B}, h, h', \sigma)}{\|h - h'\|_1}$.

Proof. Fix $h, h' \in \mathcal{B}$ and an assignment function $\sigma \sim (\mathcal{B}, h)$. It suffices to show the existence of another assignment function $\pi \in \mathcal{H}$ such that

$$(4.31) \quad \text{OAT}(\mathcal{B}, h, h', \sigma) \leq \text{OAT}(\mathcal{B}, h, h', \pi).$$

For every $p \in [n]$, since $\frac{\sigma(p, \cdot)}{\|\sigma(p, \cdot)\|_1} \in \mathcal{B}$, we can rewrite

$$\sigma(p, \cdot) = \|\sigma(p, \cdot)\|_1 \cdot \sum_{v \in V} \alpha_v^p \cdot v$$

for some distribution $\alpha^p \in \Delta_{|V|}$. Now we consider a weighted set V together with weights $w_V(v) = \sum_{p \in [n]} \|\sigma(p, \cdot)\|_1 \cdot \alpha_v^p$, and we construct $\pi : V \times C \rightarrow \mathbb{R}_{\geq 0}$ as follows: for every $v \in V$, let $\pi(v, \cdot) = w_V(v) \cdot v$. Since

$$\begin{aligned} \sum_{v \in V} \pi(v, \cdot) &= \sum_{v \in V, p \in [n]} \|\sigma(p, \cdot)\|_1 \cdot \alpha_v^p \cdot v && \text{(Defn. of } \pi) \\ &= \sum_{p \in [n]} \sigma(p, \cdot) && \text{(Defn. of } \alpha^p) \\ &= h, && (\sigma \sim h) \end{aligned}$$

we know that $\pi \sim (\mathcal{B}, h)$, which implies that $\pi \in \mathcal{H}$.

Let $\pi' : V \times C \rightarrow \mathbb{R}_{\geq 0}$ with $\pi' \sim (\mathcal{B}, h')$ be an assignment function such that

$$\text{OAT}(\mathcal{B}, h, h', \pi) = \sum_{v \in V} \|\pi(v, \cdot) - \pi'(v, \cdot)\|_1.$$

We construct an assignment function $\sigma' : [n] \times [k] \rightarrow \mathbb{R}_{\geq 0}$ as follows: for every $p \in [n]$, let

$$\sigma'(p, \cdot) = \sum_{v \in V} \frac{\|\sigma(p, \cdot)\|_1 \cdot \alpha_v^p}{w_V(v)} \cdot \pi'(v, \cdot).$$

We first verify that $\sigma' \sim (\mathcal{B}, h')$. Note that for every $p \in [n]$

$$\begin{aligned} \|\sigma'(p, \cdot)\|_1 &= \sum_{v \in V} \frac{\|\sigma(p, \cdot)\|_1 \cdot \alpha_v^p}{w_V(v)} \cdot \|\pi'(v, \cdot)\|_1 && \text{(Defn. of } \sigma') \\ &= \|\sigma(p, \cdot)\|_1 \cdot \sum_{v \in V} \alpha_v^p && \text{(Defn. of } \pi') \\ &= \|\sigma(p, \cdot)\|_1, && (\alpha^p \in \Delta_{|V|}) \end{aligned}$$

which implies that $\sigma'(p, \cdot) \in \|\sigma(p, \cdot)\|_1 \cdot \mathcal{B}$. Also, we have

$$\begin{aligned} \sum_{p \in [n]} \sigma'(p, \cdot) &= \sum_{p \in [n], v \in V} \frac{\|\sigma(p, \cdot)\|_1 \cdot \alpha_v^p}{w_V(v)} \cdot \pi'(v, \cdot) && \text{(Defn. of } \pi') \\ &= \sum_{v \in V} \pi'(v, \cdot) && \text{(Defn. of } w_V(v)) \\ &= h', && (\pi' \sim h') \end{aligned}$$

which implies that $\sigma' \sim h'$. Thus, $\sigma' \sim (\mathcal{B}, h')$ holds. Finally, we have

$$\begin{aligned}
 & \sum_{p \in [n]} \|\sigma(p, \cdot) - \sigma'(p, \cdot)\|_1 \\
 = & \sum_{p \in [n]} \|\sigma(p, \cdot)\|_1 \cdot \left\| \sum_{v \in V} \alpha_v^p \cdot v - \frac{\alpha_v^p}{w_V(v)} \cdot \pi'(v, \cdot) \right\|_1 && \text{(Defns. of } \alpha^p \text{ and } \sigma') \\
 \leq & \sum_{p \in [n]} \sum_{v \in V} \frac{\|\sigma(p, \cdot)\|_1 \cdot \alpha_v^p}{w_V(v)} \cdot \|w_V(v) \cdot v - \pi'(v, \cdot)\|_1 \\
 \leq & \sum_{v \in V} \sum_{p \in [n]} \frac{\|\sigma(p, \cdot)\|_1 \cdot \alpha_v^p}{w_V(v)} \cdot \|\pi(v, \cdot) - \pi'(v, \cdot)\|_1 && \text{(Defn. of } \pi) \\
 = & \sum_{v \in V} \|\pi(v, \cdot) - \pi'(v, \cdot)\|_1, && \text{(Defn. of } w_V(v))
 \end{aligned}$$

which completes the proof of Claim 9. \square

Fix $h, h' \in \mathcal{B}$ and $\sigma : [n] \times [k] \rightarrow \mathbb{R}_{\geq 0} \in \mathcal{H}$. We construct a distribution $\mu \in \Delta_{|V|}$ as follows: for every $v \in V$, let

$$\mu_v = \sum_{p \in [n]: \sigma(p, \cdot) = \|\sigma(p, \cdot)\|_1 \cdot v} \|\sigma(p, \cdot)\|_1.$$

By construction, we know that

$$\begin{aligned}
 \sum_{v \in V} \mu_v \cdot v &= \sum_{v \in V} \sum_{p \in [n]: \sigma(p, \cdot) = \|\sigma(p, \cdot)\|_1 \cdot v} \|\sigma(p, \cdot)\|_1 \cdot v && \text{(Defn. of } \mu) \\
 &= \sum_{p \in [n]} \sigma(p, \cdot) \\
 &= h, && (\sigma \sim h).
 \end{aligned}$$

Let $\mu' \in \Delta_{|V|}$ with $\sum_{v \in V} \mu'_v \cdot v = h'$ and $\kappa \vdash (\mu, \mu')$ satisfy that

$$\widetilde{\text{OAT}}(\mathcal{B}, h, h', \mu) = \sum_{v, v' \in V} \kappa(v, v') \cdot \|v - v'\|_1.$$

Next, we construct another assignment function $\sigma' \sim (\mathcal{B}, h')$ as follows: for every $p \in [n]$ with $\sigma(p, \cdot) = \|\sigma(p, \cdot)\|_1 \cdot v$, let

$$\sigma'(p, \cdot) = \frac{\|\sigma(p, \cdot)\|_1}{\mu_v} \cdot \sum_{v' \in V} \kappa(v, v') \cdot v'.$$

Note that for every $p \in [n]$, $\|\sigma'(p, \cdot)\|_1 = \frac{\|\sigma(p, \cdot)\|_1}{\mu_v} \cdot \sum_{v' \in V} \kappa(v, v') = \|\sigma(p, \cdot)\|_1$, which implies that $\sigma'(p, \cdot) \in \|\sigma(p, \cdot)\|_1 \cdot \mathcal{B}$. Also, we have

$$\begin{aligned}
 \sum_{p \in [n]} \sigma'(p, \cdot) &= \sum_{v \in V} \sum_{p \in [n]: \sigma(p, \cdot) = \|\sigma(p, \cdot)\|_1 \cdot v} \frac{\|\sigma(p, \cdot)\|_1}{\mu_v} \cdot \sum_{v' \in V} \kappa(v, v') \cdot v' && \text{(Defn. of } \sigma') \\
 &= \sum_{v, v' \in V} \kappa(v, v') \cdot v' \\
 &= \sum_{v' \in V} \mu_{v'} \cdot v' && (\kappa \vdash (\mu, \mu')) \\
 &= h',
 \end{aligned}$$

which implies that $\sigma' \sim (\mathcal{B}, h')$. Moreover,

$$\begin{aligned}
 & \text{OAT}(\mathcal{B}, h, h', \sigma) \\
 & \leq \sum_{p \in [n]} \|\sigma(p, \cdot) - \sigma'(p, \cdot)\|_1 \quad (\text{by optimality}) \\
 & = \sum_{v \in V} \sum_{p \in [n]: \sigma(p, \cdot) = \|\sigma(p, \cdot)\|_1 \cdot v} \|\sigma(p, \cdot) - \sigma'(p, \cdot)\|_1 \\
 (4.32) \quad & = \sum_{v \in V} \sum_{p \in [n]: \sigma(p, \cdot) = \|\sigma(p, \cdot)\|_1 \cdot v} \frac{\|\sigma(p, \cdot)\|_1}{\mu_v} \cdot \|\mu_v \cdot v - \sum_{v' \in V} \kappa(v, v') \cdot v'\|_1 \quad (\text{Defns. of } \sigma \text{ and } \sigma') \\
 & = \sum_{v \in V} \|\mu_v \cdot v - \sum_{v' \in V} \kappa(v, v') \cdot v'\|_1 \\
 & \leq \sum_{v, v' \in V} \kappa(v, v') \cdot \|v - v'\|_1 \\
 & = \widetilde{\text{OAT}}(\mathcal{B}, h, h', \mu), \quad (\text{Defns. of } \mu' \text{ and } \kappa)
 \end{aligned}$$

which implies that

$$\begin{aligned}
 \text{Lip}(\mathcal{B}) & = \max_{\substack{h, h' \in \mathcal{B} \\ \sigma \in \mathcal{H}: \|\sigma\|_1 = 1}} \frac{\text{OAT}(\mathcal{B}, h, h', \sigma)}{\|h - h'\|_1} \quad (\text{Claim 9}) \\
 & \leq \max_{\substack{h, h' \in \mathcal{B} \\ \mu \in \Delta_{|V|}: \sum_{v \in V} \mu_v \cdot v = h}} \frac{\widetilde{\text{OAT}}(\mathcal{B}, h, h', \mu)}{\|h - h'\|_1} \quad (\text{Ineq. (4.32)}) \\
 & = \widetilde{\text{Lip}}(\mathcal{B}).
 \end{aligned}$$

Thus, we complete the proof. \square

4.1.1 Lipschitz Constant for Matroid Basis Polytopes For preparation, we need the following lemmas that provide well-known properties of matroids; see [Sch03] for more details. The first is an easy consequence of the submodularity of the rank function and the second easily follows from the exchange property of matroid.

LEMMA 4.2. (Properties of rank function) *Let E be a ground set and let \mathcal{M} be a matroid on E with a rank function $\text{rank} : 2^E \rightarrow \mathcal{Z}_{\geq 0}$. Let $h \in P_{\mathcal{M}}$ be a point inside the matroid basis polytope. Recall that we say a subset $A \subseteq E$ is tight on h if $\sum_{i \in A} h_i = \text{rank}(A)$. If two subsets $A, B \subseteq E$ are tight on h , then both $A \cup B$ and $A \cap B$ are tight on h .*

LEMMA 4.3. (Circuit) *Let E be a ground set of size $n \geq 1$ and let \mathcal{M} be a matroid on E . Let $I \in \mathcal{M}$ be a basis and $a \in E \setminus I$ be an element. Let circuit $C(I, a)$ be the smallest tight set on e_I that contains a . We have that $C(I, a) \subseteq I$ and for every element $b \in C(I, a)$, $I \cup \{a\} \setminus \{b\} \in \mathcal{M}$.*

We show that $\text{Lip}(P_{\mathcal{M}}) \leq |E| - 1$ for matroid basis polytopes by the following theorem. This theorem is useful since $|E| = k$ for every assignment structure constraint \mathcal{B} , and hence, $\text{Lip}(\mathcal{B}) \leq k - 1$ when \mathcal{B} is a scaled matroid basis polytope.

THEOREM 4.1. (Lipschitz constant for matroid basis polytopes) *Let E be a ground set and let \mathcal{M} be a matroid on E . We have $\text{Lip}(P_{\mathcal{M}}) \leq \widetilde{\text{Lip}}(P_{\mathcal{M}}) \leq |E| - 1$.*

Proof. Let $V = \{e_I : I \text{ is a basis of } \mathcal{M}\}$ be the vertex set of $P_{\mathcal{M}}$. Fix $h, h' \in P_{\mathcal{M}}$ and a distribution $\mu \in \Delta_{|V|}$ with $\sum_{v \in V} \mu_v \cdot v = h$. We aim to show that

$$\widetilde{\text{OAT}}(\mathcal{B}, h, h', \mu) = \min_{\substack{\mu' \in \Delta_{|V|}: \sum_{v \in V} \mu'_v \cdot v = h' \\ \kappa(\mu, \mu')}} \sum_{v, v' \in V} \kappa(v, v') \cdot \|v - v'\|_1 \leq (|E| - 1) \cdot \|h - h'\|_1.$$

We first have the following claim.

CLAIM 10. (**Path decomposition from h to h' .**) *There exists a sequence $h^{(0)} = h, h^{(1)}, \dots, h^{(m)} = h' \in P_{\mathcal{M}}$ satisfying*

1. For every $i \in [m]$, $\|h^{(i-1)} - h^{(i)}\|_0 = 2$;
2. $\sum_{i \in [m]} \|h^{(i-1)} - h^{(i)}\|_1 = \|h - h'\|_1$.

Proof. We first show how to construct $h^{(1)}$ such that $\|h^{(0)} - h^{(1)}\|_0 = 2$ and $\|h^{(0)} - h^{(1)}\|_1 + \|h^{(1)} - h'\|_1 = \|h - h'\|_1$.

Let $H^+ = \{i \in E : h_i > h'_i\}$ and $H^- = \{i \in E : h_i < h'_i\}$. Consider for every element $j \in H^-$, the smallest tight set S_j on h that contains j (inclusion-wise). By Lemma 4.2, we know that such S_j is unique. If $S_j \cap H^+ \neq \emptyset$ for some $j^* \in H^-$, we can choose an arbitrary element $i^* \in S_{j^*} \cap H^+$. Let $\tau = \min_{A \subseteq E: \sum_{i \in A} h_i < \text{rank}(A)} \{\text{rank}(A) - \sum_{i \in A} h_i\}$. We construct $h^{(1)} = h + \tau \cdot (e_{j^*} - e_{i^*})$. It is easy to see that $h^{(1)} \in P_{\mathcal{M}}$, since for every $S \subseteq E$,

1. If S is not tight on h , we have $\sum_{i \in S} h_i^{(1)} \leq \sum_{i \in S} h_i^{(0)} + \tau \leq \text{rank}(S)$ by the definition of τ .
2. If S does not contain j^* , then $\sum_{i \in S} h_i^{(1)} \leq \sum_{i \in S} h_i^{(0)} \leq \text{rank}(S)$ by the construction of $h^{(1)}$.
3. If S is tight on h and $j^* \in S$, then S must contain S_{j^*} by Lemma 4.2, which makes $\sum_{i \in S} h_i^{(1)} = \sum_{i \in S} h_i^{(0)} \leq \text{rank}(S)$.

Now, suppose for every $j \in H^-$ we have $S_j \cap H^+ = \emptyset$. Consider the union U of all such S_j s. We can easily see that $H^- \subseteq U$ and $U \cap H^+ = \emptyset$. By Lemma 4.2, U is also tight on h , say $\sum_{i \in U} h_i = \text{rank}(U)$. Since $h' \in P_{\mathcal{M}}$, $\sum_{i \in U} h'_i \leq \text{rank}(U)$. However, by the definition of H^- , we have

$$\sum_{i \in U} h_i < \sum_{i \in U} h'_i \leq \text{rank}(U),$$

which is a contradiction. Hence, this case is impossible and we can always find $j \in H^-$ with $S_j \cap H^+ \neq \emptyset$ and construct $h^{(1)} = h + \tau \cdot (e_{j^*} - e_{i^*})$.

We then can repeat the above procedure and construct a sequence $h^{(1)}, h^{(2)}, \dots$. By the previous argument, at each iteration t we increase the number of tight sets on $h^{(t)}$ by at least one from $h^{(t-1)}$. Since there is a finite number of tight sets (at most $2^{|E|}$), the above process terminates in a finite number of times and arrives h' . Hence, we complete the proof of Claim 10. \square

Suppose $\widetilde{\text{OAT}}(\mathcal{B}, h, h', \mu) \leq (|E| - 1) \cdot \|h - h'\|_1$ holds for the case of $\|h - h'\|_0 = 2$. Let $\mu^{(0)} = \mu$. For every $i \in [m]$, we consecutively construct a distribution $\mu^{(i)} \in \Delta_{|V|}$ with $\sum_{v \in V} \mu_v^{(i)} \cdot v = h^{(i)}$ and a coupling $\kappa^{(i)} \vdash (\mu^{(i-1)}, \mu^{(i)})$ such that

$$(4.33) \quad \widetilde{\text{OAT}}(\mathcal{B}, h^{(i-1)}, h^{(i)}, \mu^{(i-1)}) = \sum_{v, v' \in V} \kappa^{(i)}(v, v') \cdot \|v - v'\|_1 \leq (|E| - 1) \cdot \|h^{(i-1)} - h^{(i)}\|_1.$$

Then we have

$$\begin{aligned} \widetilde{\text{OAT}}(P_{\mathcal{M}}, h, h', \mu) &= \min_{\substack{\mu' \in \Delta_{|V|}: \sum_{v \in V} \mu'_v \cdot v = h' \\ \kappa \vdash (\mu, \mu')}} \sum_{v, v' \in V} \kappa(v, v') \cdot \|v - v'\|_1 \\ &\leq \sum_{i \in [m]} \sum_{v, v' \in V} \kappa^{(i)}(v, v') \cdot \|v - v'\|_1 && \text{(Defn. of } \kappa^{(i)}) \\ &\leq (|E| - 1) \cdot \sum_{i \in [m]} \|h^{(i-1)} - h^{(i)}\|_1 && \text{(Ineq. (4.33))} \\ &= (|E| - 1) \cdot \|h - h'\|_1. && \text{(Claim 10)} \end{aligned}$$

Thus, we only prove for the case that h and h' differ in exactly two entries $\|h - h'\|_0 = 2$. We define the following augmenting path.

DEFINITION 4.7. (Augmenting path from h to h') Let h and h' be two vectors in \mathcal{B} that differ in exactly two entries, $\|h - h'\|_0 = 2$. W.l.o.g., we assume the corresponding two elements are $s, t \in E$ with $h_s < h'_s$ and $h_t > h'_t$ respectively. Let $\mu \in \Delta_{|V|}$ be the distribution associated with h (i.e., $\sum_{v \in V} \mu_v \cdot v = h$). We say a sequence of indicator vector of basis $(v_1 = e_{I_1}, \dots, v_m = e_{I_m})$ together with a sequence of indices of elements $(a_0 = s, a_1, \dots, a_{m-1}, a_m = t)$ ($a_i \in E$ for $0 \leq i \leq m$) form a weak augmenting path of length $m \geq 1$ from h to h' if

1. For every $i \in [m]$, $\mu_{v_i} > 0$ (i.e., every v_i appears in the support of μ);
2. For every $i \in [m]$, $I_i \cup \{a_{i-1}\} \setminus \{a_i\} \in \mathcal{M}$ (weak exchange property).

Here, we do not require that I_i s and a_i s be distinct.

Moreover, we say $(v_1 = e_{I_1}, \dots, v_m = e_{I_m})$ and $(a_0 = s, a_1, \dots, a_{m-1}, a_m = t)$ form a strong augmenting path if it is a weak augmenting path with the following additional properties:

1. All elements a_i are distinct;
2. For every basis $I \in \mathcal{M}$, letting $A_I = \{i \in [m] : I_i = I\}$, we have

$$\widehat{I} = I \cup \{a_{i-1} : i \in A_I\} \setminus \{a_i : i \in A_I\} \in \mathcal{M}. \quad (\text{strong exchange property})$$

Intuitively, given a weak augmenting path, we can perform a sequence of exchange of basis according to the weak exchange property. As a result, we get a sequence of new basis so that the mass in $a_0 = s$ is transported to $a_m = t$. However, there is a subtle technical problem that one basis may be used several times and we need to guarantee it is still a basis after all exchange operations, which motivates the notion of strong exchange property. The following crucial lemma shows that a strong augmenting path from h to h' exists and its length can be bounded.

LEMMA 4.4. (Existence of a strong augmenting path of length $\leq |E| - 1$) Given h and h' in \mathcal{B} that differ in exactly two entries, for any $\mu \in \Delta_{|V|}$ with $\sum_{v \in V} \mu_v \cdot v = h$, there exists a strong augmenting path of length at most $|E| - 1$ from h to h' .

Proof. For preparation, we construct a (multi-edge) directed graph G_{ex} , called the *exchange graph*. G_{ex} has vertex set E and the following set of edges: for every basis $I \in \mathcal{M}$ with $\mu_{e_I} > 0$, every element $a \in E \setminus I$ and every element $b \in C(I, a)$, add a directed edge (a, b) with a certificate $C(I, a)$ to the graph G_{ex} .

Existence of A Weak Augmenting Path We first prove the existence of a weak augmenting path from h to h' , which is equivalent to proving that there exists a path on G_{ex} from vertex s to vertex t . By contradiction assume that there does not exist such a path. Let E^- be the collection of elements $i \in E$ such that there exists a path from s to i , and let $E^+ = E \setminus E^-$. We have that $t \in E^+$. Since $\sum_{i \in E^-} h_i < \sum_{i \in E^-} h'_i \leq \text{rank}(E^-)$, we know that E^- is not a tight set on h . Consequently, there must exist a basis $I \in \mathcal{M}$ with $\mu_{e_I} > 0$ such that $|I \cap E^-| < \text{rank}(E^-)$. Let $a \in E^- \setminus I$ be an element such that $(I \cap E^-) \cup \{a\} \in \mathcal{M}$. If $C(I, a) \subseteq E^-$, we have $C(I, a) \cup \{a\} \notin \mathcal{M}$. However, $C(I, a) \cup \{a\} \subseteq (I \cap E^-) \cup \{a\} \in \mathcal{M}$, which is a contradiction. Hence, we have that $C(I, a) \cap E^+ \neq \emptyset$ and there exists an edge from $a \in E^-$ to some element in E^+ , which contradicts the definition of E^- .

Existence of A Strong Augmenting Path of Length $\leq |E| - 1$ Among all weak augmenting path from h to h' , we select a shortest one, say $(v_1 = e_{I_1}, \dots, v_m = e_{I_m})$ and $(a_0 = s, a_1, \dots, a_{m-1}, a_m = t)$. We claim that $a_0 \neq a_1 \neq \dots \neq a_m$. Assume that $a_i = a_j$ for some $0 \leq i < j \leq m$, we can see that $(e_{I_1}, \dots, e_{I_i}, e_{I_{j+1}}, \dots, e_{I_m})$ and $(a_0 = s, a_1, \dots, a_i, a_{j+1}, \dots, a_{m-1}, a_m = t)$ form a shorter weak augmenting path, which is a contradiction. Since there are at most $|E|$ different elements, we have that $m \leq |E| - 1$.

Now we prove that $(v_1 = e_{I_1}, \dots, v_m = e_{I_m})$ and $(a_0 = s, a_1, \dots, a_{m-1}, a_m = t)$ already form a strong augmenting path. Note that for every $i \in [m]$, (a_{i-1}, a_i) is a direct edge on G_{ex} with a certificate $C(I_i, a)$. It is easy to verify this when all v_1, \dots, v_m are distinct (i.e., each edge (a_{i-1}, a_i) is defined by a distinct basis).

Now, we consider the more difficult case where v_i s are not all distinct. Observe $v_i \neq v_{i+1}$ (i.e., two adjacent edges correspond to two different bases), since $a_i \in I_i$ and $a_i \notin I_{i+1}$ which implies $I_i \neq I_{i+1}$. Fix a basis $I \in \mathcal{M}$ with $A_I = \{i \in [m] : t \in [T], I_i = I\}$ and $|A_I| \geq 2$. W.l.o.g., assume $0 \leq i_1 < i_2 < \dots < i_T \leq m$ and we have $i_{l+1} - i_l \geq 2$ for $l \in [T - 1]$. We first note that for any $i < j \in A_I$, $a_j \notin C(I, a_{i-1})$. Suppose, for contradiction that $a_j \in C(I, a_{i-1})$. Based on the construction of graph G_{ex} , (a_{i-1}, a_j) is also an edge with certificate $C(I, a_j)$.

Consequently, $(e_{I_1}, \dots, e_{I_i}, e_{I_{j+1}}, \dots, e_{I_m})$ and $(a_0 = s, a_1, \dots, a_{i-1}, a_j, a_{j+1}, \dots, a_{m-1}, a_m = t)$ form a shorter weak augmenting path, which is a contradiction.

Hence, we have that $a_j \notin C(I, a_{i-1})$ holds for any $i < j \in A_I$. We first swap $a_{i_{T-1}}$ and a_{i_T} and obtain a basis $I^{(T)} = I \cup \{a_{i_{T-1}}\} \setminus \{a_{i_T}\}$. The circuit $C(I, a_{i_{T-1}-1})$ is completely contained in $I^{(T)}$ since $a_{i_T} \notin C(I, a_{i_{T-1}-1})$, and hence we can continuously perform the swap between $a_{i_{T-1}-1}$ and $a_{i_{T-1}}$. By reduction, we have that $I^{(l)} = I \cup \{a_{i_{l-1}}, \dots, a_{i_{T-1}}\} \setminus \{a_{i_l}, \dots, a_{i_T}\}$ for every $l \in [T]$ is still a basis. Hence, $\widehat{I} = I^{(1)}$ is a basis, which completes the proof. \square

We consider the following procedure:

1. Find a strong augmenting path of length $m \leq |E| - 1$ from h to h' , say $(v_1 = e_{I_1}, \dots, v_m = e_{I_m})$ and $(a_0 = s, a_1, \dots, a_{m-1}, a_m = t)$.
2. Let $\tau = \min \{h'_1 - h_1, \min_{i \in [m]} \mu_{v_i}\}$. For every basis $I \in \mathcal{M}$, let $A_I := \{i \in [m] : I_i = I\}$ and let $\widehat{I} = I \cup \{a_i : i \in A_I\} \setminus \{a_{i-1} : i \in A_I\}$.
3. Construct $\mu'' \in \Delta_{|V|}$ as the resulting distribution of the following procedure: for every $I \in \mathcal{M}$ with $A_I \neq \emptyset$, reduce μ_{e_I} by τ and increase $\mu_{e_{\widehat{I}}}$ by τ .
4. Let $h'' = \sum_{v \in V} \mu''_v \cdot v$. If $h'' = h'$, we are done. Otherwise, let $h \leftarrow h''$ and $\mu \leftarrow \mu''$, and iteratively run the above steps for tuple (h, h', μ) .

After running an iteration, we can easily verify the following

1. $\|h'' - h'\|_0 \leq 2$, $h''_1 - h_1 = \tau$ and $h_2 - h''_2 = \tau$.
2. $\|h'' - h'\|_1 = \|h - h'\|_1 - 2\tau$.
3. The total assignment transportation from μ to μ'' , say $\min_{\kappa \vdash (\mu, \mu'')} \sum_{v, v' \in V} \kappa(v, v') \cdot \|v - v'\|_1$, is at most $\sum_{I \in \mathcal{M}: |I| = \text{rank}(\mathcal{M})} \tau \cdot \|e_I - e_{\widehat{I}}\|_1 = 2m\tau \leq 2(|E| - 1)\tau$.

It means that we can reduce the value $\|h - h'\|_1$ by 2τ , by introducing at most $2(|E| - 1)\tau$ assignment transportation for $\widehat{\text{OAT}}(\mathcal{B}, h, h', \mu)$. Hence, the required assignment transportation from h to h' is at most $(|E| - 1) \cdot \|h - h'\|_1$, which implies that

$$\widehat{\text{OAT}}(P_{\mathcal{M}}, h, h', \mu) \leq (|E| - 1) \cdot \|h - h'\|_1.$$

Due to the arbitrary selection of h, h' and μ , we complete the proof of Theorem 4.2. \square

4.1.2 Lipschitz Constant for Laminar Matroid Basis Polytopes

Our main theorem for laminar matroid basis polytopes is as follows.

THEOREM 4.2. *Let E be a ground set and let \mathcal{M} be a laminar matroid on E of depth $\ell \geq 1$. We have $\text{Lip}(P_{\mathcal{M}}) \leq \widehat{\text{Lip}}(P_{\mathcal{M}}) \leq \ell + 1$.*

Proof. Let $V = \{e_I : I \text{ is a basis of } \mathcal{M}\}$ be the vertex set of $P_{\mathcal{M}}$. Let \mathcal{A} be the corresponding laminar and u be the corresponding capacity function of \mathcal{M} . W.l.o.g., we assume $u(A) = \text{rank}(A)$ for every $A \in \mathcal{M}$. Let $h, h' \in P_{\mathcal{M}}$ and $\mu \in \Delta_{|V|}$ be a distribution on V satisfying that

$$\sum_{v \in V} \mu_v \cdot v = h.$$

By the proof of Theorem 4.1, we only need to prove for the case that $\|h - h'\|_0 = 2$. W.l.o.g., we assume $h_1 < h'_1$ and $h_2 > h'_2$. Again, by the same argument as in Theorem 4.1, it suffices to prove the following claim.

LEMMA 4.5. *There exists a strong augmenting path of length at most $\ell + 1$ from h to h' .*

Proof. We again construct the exchange graph G_{ex} as in the proof of Claim 4.4, and iteratively construct a strong augmenting path from h to h' . Also recall that we assume $h_s < h'_s$ and $h_t > h'_t$ for some $s, t \in E$. Let $A^* \in \mathcal{A}$ be the minimal set that contains $\{s, t\}$. Firstly, we find a vector $v_1 = e_{I_1}$ with $\mu_{v_1} > 0$ and $a_0 = s \notin I_1$. Let $A_1 \in \mathcal{A}$ be the minimal set with $a_0 \in A_1$ and $|I_1 \cap A_1| = u(A_1)$. Note that such A_1 must exist, since $I_1 \in P_{\mathcal{M}}$ and we have $|I_1 \cap E| = \text{rank}(\mathcal{M}) = u(E)$. We discuss the following cases.

Case 1: $A_1 \supseteq A^*$ Since $h_t > h'_t \geq 0$, there must exist a vertex $v_2 = e_{I_2}$ with $\mu_{v_2} > 0$ and $a_2 = t \in I_2$. Note that $|(I_2 \setminus \{a_2\}) \cap A_1| < |I_1 \cap A_1| = u(A_1)$. There must exist an element $a_1 \in I_1 \cap A_1$ such that $(I_2 \cap A_1) \cup \{a_1\} \setminus \{a_2\} \in \mathcal{M}$. Due to the laminar structure, we know that $I_2 \cup \{a_1\} \setminus \{a_2\} \in \mathcal{M}$ also holds. Also, since $(I_1 \cap A) \setminus \cup \{a_0\} \setminus \{a_1\} \in \mathcal{M}$, we conclude that $I_1 \cup \{a_0\} \setminus \{a_1\} \in \mathcal{M}$. Consequently, (v_1, v_2) and (a_0, a_1, a_2) form an augmenting path of length 2 from h to h' .

Case 2: $A_1 \subsetneq A^*$ We know that $\sum_{i \in A_1} h'_i - h_i = h'_s - h_s > 0$. Thus, we must have

$$\begin{aligned} \sum_{e_I \in V} \mu_{e_I} \cdot |I \cap A_1| &= \sum_{i \in A_1} h_i && \text{(Defn. of } \mu) \\ &< \sum_{i \in A_1} h'_i \\ &\leq u(A_1), && (h' \in P_{\mathcal{M}}) \end{aligned}$$

which implies the existence of a vertex $v_2 = e_{I_2} \in V$ with $\mu_{v_2} > 0$ and $|I_2 \cap A_1| < u(A_1)$. Since $u(A) = |I_1 \cap A_1| > |I_2 \cap A_1|$, there must exist an element $a_1 \in I_1 \cap A_1$ such that $(I_2 \cap A_1) \cup \{a_1\} \in \mathcal{M}$. Also note that since $I_1 \cup \{a_0\}$ only violates the capacity constraint on sets $A \supseteq A_1$ ($A \in \mathcal{A}$) by the fact that \mathcal{M} is a laminar matroid, we have that $I_1 \cup \{a_0\} \setminus \{a_1\} \in \mathcal{M}$.

Note that $I_2 \cup \{a_1\}$ can only violates the capacity constraint on sets $A \supseteq A_1$ ($A \in \mathcal{A}$). We can again find $A_2 \in \mathcal{A}$ be the minimal set with $a_1 \in A_2$ and $|I_2 \cap A_2| = u(A_2)$, and recursively apply the above argument to tuple (I_2, a_1, A_2) until arriving $A_{m-1} \supseteq A^*$. Since the operation on I_1 , say $\tilde{I}_1 = I_1 \cup \{a_0\} \setminus \{a_1\} \in \mathcal{M}$, maintains the number $|\tilde{I}_1 \cap A_2| = |I_1 \cap A_2|$, we have that if $A_2 \subsetneq A^*$, the following inequality holds

$$\sum_{e_I \in V} \mu_{e_I} \cdot |I \cap A_1| + \mu_{e_{\tilde{I}_1}} (|\tilde{I}_1 \cap A_2| - |I_1 \cap A_2|) < u(A_2).$$

Hence, the above argument can be applied to tuple (I_2, a_1, A_2) .

Overall, we can get an augmenting path $(v_1 = e_{I_1}, \dots, v_m = e_{I_m})$ and $(a_0 = s, a_1, \dots, a_{m-1}, a_m = t)$, together with a certificate set sequence $A_1 \subsetneq A_2 \subsetneq \dots \subsetneq A_{m-1}$ where $A_{m-1} \supseteq A^*$. Since the depth of \mathcal{M} is ℓ , we have $m - 1 \leq \ell$. Thus, we complete the proof of Claim 4.5. \square

Overall, we complete the proof of Theorem 4.2. \square

4.2 Lipschitz Constant $\text{Lip}(\mathcal{B})$ can be Unbounded We show that $\text{Lip}(\mathcal{B})$ can be extremely large by considering the following assignment structure constraints \mathcal{B} , called knapsack polytopes.

DEFINITION 4.8. (Knapsack polytope) Let $A \in \mathbb{R}_{\geq 0}^{m \times k}$ be a non-negative matrix. We say an assignment structure constraint $\mathcal{B} \in \Delta_k$ is A -knapsack polytope if

$$\mathcal{B} = \{x \in \Delta_k : Ax \leq 1\},$$

where $Ax \leq 1$ is called knapsack constraints.

Our result is as follows.

THEOREM 4.3. (Lipschitz constant may be unbounded for knapsack polytope) Let $m = 2$ and $k = 3$. For any $U > 0$, there exists a matrix $A \in \mathbb{R}_{\geq 0}^{m \times k}$ such that the Lipschitz constant $\text{Lip}(\mathcal{B})$ of the A -knapsack polytope \mathcal{B} satisfies $\text{Lip}(\mathcal{B}) \geq U$.

Proof. We construct A by $A_1 = (\frac{10U+2}{5U+3}, \frac{4}{5U+3}, 0)$ and $A_2 = (\frac{10U+2}{5U+3}, 0, \frac{4}{5U+3})$. Let $h = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$ and $h' = (\frac{1}{2} + \frac{1}{10U}, \frac{1}{4} - \frac{1}{20U}, \frac{1}{4} - \frac{1}{20U})$. We can check that $h, h' \in \mathcal{B}$ and specifically, h' is a vertex of \mathcal{B} . Let $\sigma : [2] \times [3] \rightarrow \mathbb{R}_{\geq 0}$ be defined as follows:

$$\sigma(1, \cdot) = (\frac{1}{4}, \frac{1}{4}, 0), \text{ and } \sigma(2, \cdot) = (\frac{1}{4}, 0, \frac{1}{4}).$$

We know that $\sigma \sim (\mathcal{B}, h)$. Since h' is a vertex of \mathcal{B} , there is only one $\sigma' : [2] \times [3] \rightarrow \mathbb{R}_{\geq 0}$ with $\sigma' \sim (\mathcal{B}, h')$ and $\|\sigma'(p, \cdot)\|_1 = \|\sigma(p, \cdot)\|_1$ for $p \in [2]$, say

$$\sigma'(1, \cdot) = \frac{1}{2} \cdot h', \text{ and } \sigma'(2, \cdot) = \frac{1}{2} \cdot h'$$

for some $\alpha \in [0, 1]$. Thus, we have

$$\sum_{p \in [2]} \|\sigma(p, \cdot) - \sigma'(p, \cdot)\|_1 = \frac{1}{2} + \frac{3}{20U}.$$

However, $\|h - h'\|_1 = \frac{1}{5U}$, which implies that

$$\text{Lip}(\mathcal{B}) \geq \text{OAT}(\mathcal{B}, h, h', \sigma) \cdot \|h - h'\|_1 = \left(\frac{1}{2} + \frac{3}{20U}\right) / \left(\frac{1}{5U}\right) > U.$$

This completes the proof. \square

5 Bounding the Covering Exponent $\Lambda_\varepsilon(\mathcal{X})$

As we show in Theorem 3.1, the covering exponent $\Lambda_\varepsilon(\mathcal{X})$ is a key parameter for the complexity of the metric space. In this section, we give several upper bounds of $\Lambda_\varepsilon(\mathcal{X})$, against several other well-known notions of “dimension” measure of the metric space (\mathcal{X}, d) , which enables us to obtain coresets for constrained clustering in various metric spaces (see Remark 1 for a summary of the concrete metric families that we can handle). We first introduce the shattering dimension (Definition 5.2), and show its relation to the covering exponent in Lemma 5.1. This relation helps to translate the upper bounds for the shattering dimension (which is rich in the literature) to small-size coresets (Corollary 5.1). Then we consider the doubling dimension (Definition 5.3), and show the covering exponent is bounded by the doubling dimension (Lemma 5.2).

Shattering Dimension As was also considered in recent papers [BBH⁺20, BJKW21a, BCJ⁺22], we employ the following notion of shattering dimension of (the ball range space of) metric space (see also e.g. [Hp11]).

DEFINITION 5.1. (Ball range space) A ball $\text{Ball}(x, r) = \{y \in \mathcal{X} \mid d(x, y) \leq r\}$ is determined by the center $c \in \mathcal{X}$ and the radius $r > 0$. Let $\text{Balls}(\mathcal{X}) = \{\text{Ball}(x, r) \mid x \in \mathcal{X}, r > 0\}$ denote the collection of all balls in \mathcal{X} . $(\mathcal{X}, \text{Balls}(\mathcal{X}))$ is called the ball range space of \mathcal{X} . For $P \subset \mathcal{X}$, let

$$P \cap \text{Balls}(\mathcal{X}) := \{P \cap \text{Ball}(x, r) \mid \text{Ball}(x, r) \in \text{Balls}(\mathcal{X})\}.$$

We are ready to define the shattering dimension.

DEFINITION 5.2. (Shattering dimension of ball range spaces) The shattering dimension of $(\mathcal{X}, \text{Balls}(\mathcal{X}))$, denoted by $\text{dim}(\mathcal{X})$ is the minimum positive integer t such that for every $P \subset \mathcal{X}$ with $|P| \geq 2$,

$$|P \cap \text{Balls}(\mathcal{X})| \leq |P|^t$$

A closely related notion to shattering dimension is the well-known Vapnik-Chervonenkis dimension [VC71]. Technically, they are equivalent up to only a logarithmic factor (see e.g., [Hp11]). The applications of shattering dimension in designing small-sized coresets for unconstrained (k, z) -CLUSTERING stem from the work [FL11], which are followed by [HJLW18, BBH⁺20, BJKW21a]. However, all these works require to bound the shattering dimension of a more complicated weighted range space. Only recently, this complication of weight in the range space is removed in the uniform sampling framework by [BCJ⁺22], which only requires to consider the unweighted ball range space (Definition 5.1).

We have the following lemma that upper bounds the covering number via the shattering dimension.

LEMMA 5.1. For every $\alpha > 0$, we have $\Lambda_\alpha(\mathcal{X}) \leq O(\text{dim}(\mathcal{X}) \cdot z^2 \alpha^{-1} \varepsilon^{-1})$.

Proof. We fix an unweighted dataset $P \subseteq \text{Ball}(a, r_{\max})$ of $n \geq 2$ points. For ease of statement, we use r to represent r_{\max} in the following. We upper bound $N_{\mathcal{X}}(P, \alpha)$ by constructing an α -covering \mathcal{C} . For every $x \in \text{Ball}(a, 24zr/\varepsilon)$,

let $f_x(y) = d(x, y), y \in P$ denote its distance function. We round $f_x(y)$ to obtain an approximation \tilde{f}_x such that $\forall y \in P, \tilde{f}_x(y) = \lfloor \frac{24z \cdot d(x, y)}{\alpha r} \rfloor \cdot \frac{\alpha r}{24z}$. Let $\mathcal{F}_\alpha = \{\tilde{f}_x \mid x \in \text{Ball}(a, 24zr/\varepsilon)\}$ and for each $\tilde{f} \in \mathcal{F}_\alpha$, we add exactly one x_0 such that $\tilde{f}_{x_0} = \tilde{f}$ into \mathcal{C} (notice that there can be multiple x 's that have the same approximation \tilde{f}_x , we include only one of them).

We claim that \mathcal{C} is already an α -covering of P . To see this, note that for every $x \in \text{Ball}(a, 24zr/\varepsilon)$, by construction, $\forall y \in P, |\tilde{f}_x(y) - f_x(y)| \leq \frac{\alpha r}{24z}$ and also by construction, there exists $c \in \mathcal{C}$ such that $\tilde{f}_c = \tilde{f}_x$. So we conclude that for every

$$\max_{p \in P} |d(x, y) - d(c, y)| \leq \frac{\alpha r}{24z} + \frac{\alpha r}{24z} = \frac{\alpha r}{12z}.$$

It remains to upper bound the size of \mathcal{C} . To achieve this goal, we first construct a family of distance functions G_α . Let $T = O(\frac{z^2}{\alpha \varepsilon})$. We start with enumerating all possible collections of subsets $\{H_i \mid i = 0, 1, \dots, T\}$ satisfying the followings

1. $H_i \in P \cap \text{Balls}(\mathcal{X})$ (Definition 5.1),
2. $H_T \subseteq H_{T-1} \subseteq \dots \subseteq H_0 \subseteq P$.

By the definition of the shattering dimension, we know that $|P \cap \text{Balls}(\mathcal{X})| \leq |P|^{\dim(\mathcal{X})}$. So there are at most $(|P|^{\dim(\mathcal{X})})^{O(z^2 \alpha^{-1} \varepsilon^{-1})} = |P|^{O(\dim(\mathcal{X}) \cdot z^2 \alpha^{-1} \varepsilon^{-1})}$ collections of such $\{H_i \mid i = 0, \dots, T\}$.

Fix a collection $\mathcal{H} = \{H_i \mid i = 0, 1, \dots, T\}$, we construct a corresponding distance function $g_{\mathcal{H}}$ as the following. For every $p \in P$, let i_p denote the maximum integer $i \in \{0, \dots, T\}$ such that $p \in H_i$. Let $g_{\mathcal{H}}(p) = i_p \cdot \frac{\alpha r}{24z}$. Let G_α denote the subset of all possible $g_{\mathcal{H}}$'s constructed as above.

Now we claim that $|\mathcal{C}| \leq |G_\alpha|$ since $|G_\alpha| \leq n^{O(z^2 \alpha^{-1} \varepsilon^{-1} \cdot \dim(\mathcal{M}))}$, which completes the proof. It suffices to show that $\mathcal{F}_\alpha \subseteq G_\alpha$. To prove $\mathcal{F}_\alpha \subseteq G_\alpha$, we fix $\tilde{f}_x \in \mathcal{F}_\alpha$ and show that there must exist a realization $\mathcal{H} = \{H_i \mid i = 0, 1, \dots, T\}$ such that $g_{\mathcal{H}} = \tilde{f}_x$. To see this, we simply let $H_i = \{y \in P \mid \tilde{f}_x(y) \leq i \cdot \frac{\alpha r}{24z}\}$. It is obvious that $H_i \in P \cap \text{Balls}(\mathcal{X})$ and $H_T \subseteq \dots \subseteq H_0$, thus \mathcal{H} is a valid realization of the enumeration. Moreover, by construction, we know that $g_{\mathcal{H}} = \tilde{f}_x$. \square

Applying $\Lambda_\varepsilon(\mathcal{X}) = O_z(\dim(\mathcal{X}) \cdot \varepsilon^{-2})$ to Theorem 3.1, we have the following corollary that bounds the coresets size via $\dim(\mathcal{X})$.

COROLLARY 5.1. (Relating coresets size to shattering dimension) *Let (\mathcal{X}, d) be a metric space, $k \geq 1, m \geq 0$ be integers, and $z \geq 1$ be a constant. Let $\varepsilon, \delta \in (0, 1)$ and $\mathcal{B} \subseteq \Delta_k$ be a convex body specifying the assignment structure constraint. There exists a randomized algorithm that given a dataset $P \subseteq \mathcal{X}$ of size $n \geq 1$ and an $(2^{O(z)}, O(1), O(1))$ -approximation $C^* \in \mathcal{X}^k$ of P for (k, z) -CLUSTERING with m outliers, constructs an $(\varepsilon, \mathcal{B}, m)$ -coresets for (k, z) -CLUSTERING with general assignment constraints of size*

$$O(m) + 2^{O(z \log z)} \cdot \tilde{O}(\text{Lip}(\mathcal{B})^2 \cdot (\dim(\mathcal{X}) \cdot \varepsilon^{-2} + k) \cdot k^2 \varepsilon^{-2z}) \cdot \log \delta^{-1},$$

in $O(nk)$ time with probability at least $1 - \delta$. Moreover, when $\mathcal{B} = \Delta_k$, the coresets size can be further improved to

$$O(m) + 2^{O(z \log z)} \cdot \tilde{O}(\dim(\mathcal{X}) \cdot k^2 \varepsilon^{-2z-2}) \cdot \log \delta^{-1}.$$

Doubling Dimension Doubling dimension is an important generalization of Euclidean and more generally ℓ_p spaces with the motivation to capture the intrinsic complexity of a metric space [Ass83, GKL03]. Metric spaces with bounded doubling dimensions are known as doubling metrics. For unconstrained clustering, small-sized coresets were found in doubling metrics [HJLW18, CSS21].

DEFINITION 5.3. (Doubling dimension [Ass83, GKL03]) *The doubling dimension of a metric space (\mathcal{X}, d) is the least integer t , such that every ball can be covered by at most 2^t balls of half the radius.*

LEMMA 5.2. *For every $\alpha > 0$, we have $\Lambda_\alpha(\mathcal{X}) \leq O(\text{ddim}(\mathcal{X}) \cdot \log(z\alpha^{-1}\varepsilon^{-1}))$.*

Proof. The proof is standard and we provide it for completeness. Fix an unweighted dataset $P \subseteq \text{Ball}(a, r_{\max})$ of $n \geq 2$ points. For ease of statement, we use r to represent r_{\max} in the following. By the definition of doubling

dimension, $\text{Ball}(a, 24zr/\varepsilon)$ has an $\frac{\alpha r}{12z}$ -net of size $(\frac{z}{\alpha \cdot \varepsilon})^{O(\text{ddim}(\mathcal{X}))}$. By the triangle inequality, we can see that this net is an α -covering of \mathcal{X} . So we have

$$N_{\mathcal{X}}(P, \alpha) \leq \left(\frac{z}{\alpha \cdot \varepsilon}\right)^{O(\text{ddim}(\mathcal{X}))} \leq n^{O(\text{ddim}(\mathcal{X}) \cdot \log(z\alpha^{-1}\varepsilon^{-1}))}.$$

Due to the arbitrary selection of P and $n \geq 2$, we have $\Lambda_{\alpha}(\mathcal{X}) \leq O(\text{ddim}(\mathcal{X}) \cdot \log(z\alpha^{-1}\varepsilon^{-1}))$, which completes the proof. \square

Similarly, applying $\Lambda_{\varepsilon}(\mathcal{X}) = \tilde{O}(\text{ddim}(\mathcal{X}))$ to Theorem 3.1, we have the following corollary that bounds the coresets size via $\text{ddim}(\mathcal{X})$.

COROLLARY 5.2. (Relating coresets size to doubling dimension) *Let (\mathcal{X}, d) be a metric space, $k \geq 1, m \geq 0$ be integers, and $z \geq 1$ be a constant. Let $\varepsilon, \delta \in (0, 1)$ and $\mathcal{B} \subseteq \Delta_k$ be a convex body specifying the assignment structure constraint. There exists a randomized algorithm that given a dataset $P \subseteq \mathcal{X}$ of size $n \geq 1$ and an $(2^{O(z)}, O(1), O(1))$ -approximation $C^* \in \mathcal{X}^k$ of P for (k, z) -CLUSTERING with m outliers, constructs an $(\varepsilon, \mathcal{B}, m)$ -coresets for (k, z) -CLUSTERING with general assignment constraints of size*

$$O(m) + 2^{O(z \log z)} \cdot \tilde{O}(\text{Lip}(\mathcal{B})^2 \cdot (\text{ddim}(\mathcal{X}) + k + \varepsilon^{-1}) \cdot k^2 \varepsilon^{-2z}) \cdot \log \delta^{-1},$$

in $O(nk)$ time with probability at least $1 - \delta$. Moreover, when $\mathcal{B} = \Delta_k$, the coresets size can be further improved to

$$O(m) + 2^{O(z \log z)} \cdot \tilde{O}(\text{ddim}(\mathcal{X}) \cdot k^2 \varepsilon^{-2z}) \cdot \log \delta^{-1}.$$

REMARK 1. (Covering exponents for special metrics) *We list below several examples of metric spaces that have bounded covering exponent, which is obtained by using Corollary 5.1 or Corollary 5.2.*

- *Let (\mathcal{X}, d) be an Euclidean metric \mathbb{R}^d . We have $\text{ddim}(\mathcal{X}) \leq d + 1$ and hence, $\Lambda_{\varepsilon}(\mathcal{X}) \leq \tilde{O}(d)$ by Lemma 5.2. Once we have the bound $\Lambda_{\varepsilon}(\mathcal{X}) \leq \tilde{O}(d)$, we can further assume $d = \tilde{O}(\varepsilon^{-2} \log k)$ by applying a standard iterative size reduction technique introduced in [BJKW21a], which has been applied in other coresets works [CLSS22, BCJ⁺22]. This idea yields a coresets of size $2^{O(z \log z)} \cdot \tilde{O}(\text{Lip}(\mathcal{B})^2 \cdot k^2 \varepsilon^{-2z-2}) \cdot \log \delta^{-1}$, which removes the dependence of d .*
- *Let (\mathcal{X}, d) be a doubling metric with bounded doubling dimension $\text{ddim}(\mathcal{X})$. We directly have $\Lambda_{\varepsilon}(\mathcal{X}) \leq \tilde{O}(\text{ddim}(\mathcal{X}))$ by Lemma 5.2.*
- *Let (\mathcal{X}, d) be a general discrete metric. Note that $\text{ddim}(\mathcal{X}) \leq \log |\mathcal{X}|$, and hence, we have $\Lambda_{\varepsilon}(\mathcal{X}) \leq \tilde{O}(\log |\mathcal{X}|)$ by Lemma 5.2.*
- *Let (\mathcal{X}, d) be a shortest-path metric of a graph with bounded treewidth t . By [BT15, BBH⁺20], we know that $\dim(\mathcal{M}) \leq O(t)$, which implies that $\Lambda_{\varepsilon}(\mathcal{X}) \leq \tilde{O}(t\varepsilon^{-2})$ by Lemma 5.1.*
- *Let (\mathcal{X}, d) be a shortest-path metric of a graph that excludes a fixed minor H . By [BT15], we know that $\dim(\mathcal{M}) \leq O(|H|)$, which implies that $\Lambda_{\varepsilon}(\mathcal{X}) \leq \tilde{O}(|H| \cdot \varepsilon^{-2})$ by Lemma 5.1.*

Acknowledgement

Jian Li was supported in part by the National Natural Science Foundation of China Grant 62161146004. Xuan Wu was supported by the Singapore Ministry of Education (AcRF) Tier 2 grant MOE-T2EP20122-0001.

References

- [ABM⁺19] Marek Adamczyk, Jaroslaw Byrka, Jan Marcinkowski, Syed Mohammad Meesum, and Michal Włodarczyk. Constant-factor FPT approximation for capacitated k -median. In *ESA*, volume 144 of *LIPIcs*, pages 1:1–1:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [AMT13] Zeinab Abbassi, Vahab S. Mirrokni, and Mayur Thakur. Diversity maximization under matroid constraints. In *KDD*, pages 32–40. ACM, 2013.

- [Ass83] Patrice Assouad. Plongements lipschitziens dans \mathbb{R}^n . *Bulletin de la Société Mathématique de France*, 111:429–448, 1983.
- [BBH⁺20] Daniel N. Baker, Vladimir Braverman, Lingxiao Huang, Shaofeng H.-C. Jiang, Robert Krauthgamer, and Xuan Wu. Coresets for clustering in graphs of bounded treewidth. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, pages 569–579. PMLR, 2020.
- [BCFN19] Suman Kalyan Bera, Deeparnab Chakrabarty, Nicolas Flores, and Maryam Negahbani. Fair algorithms for clustering. In *NeurIPS*, pages 4955–4966, 2019.
- [BCJ⁺22] Vladimir Braverman, Vincent Cohen-Addad, Shaofeng H.-C. Jiang, Robert Krauthgamer, Chris Schwiegelshohn, Mads Bech Tofttrup, and Xuan Wu. The power of uniform sampling for coresets. In *FOCS*, pages 462–473. IEEE, 2022.
- [BCP⁺24] Nikhil Bansal, Vincent Cohen-Addad, Milind Prabh, David Saulpic, and Chris Schwiegelshohn. Sensitivity sampling for k-means: Worst case and stability optimal coreset bounds. 2024.
- [BDHR05] Jonathan Bredin, Erik D. Demaine, Mohammad Taghi Hajiaghayi, and Daniela Rus. Deploying sensor networks with guaranteed capacity and fault tolerance. In *MobiHoc*, pages 309–319. ACM, 2005.
- [BEL13] Maria-Florina Balcan, Steven Ehrlich, and Yingyu Liang. Distributed k -means and k -median clustering on general communication topologies. In *NIPS*, pages 1995–2003, 2013.
- [BFS21] Sayan Bandyapadhyay, Fedor V. Fomin, and Kirill Simonov. On coresets for fair clustering in metric and Euclidean spaces and their applications. In Nikhil Bansal, Emanuela Merelli, and James Worrell, editors, *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12-16, 2021, Glasgow, Scotland (Virtual Conference)*, volume 198 of *LIPICs*, pages 23:1–23:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [BIK07] Moshe Babaioff, Nicole Immorlica, and Robert Kleinberg. Matroids, secretary problems, and online mechanisms. In *SODA*, pages 434–443. SIAM, 2007.
- [BIKK18] Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. Matroid secretary problems. *J. ACM*, 65(6):35:1–35:26, 2018.
- [BIO⁺19] Arturs Backurs, Piotr Indyk, Krzysztof Onak, Baruch Schieber, Ali Vakilian, and Tal Wagner. Scalable fair clustering. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pages 405–413. PMLR, 2019.
- [BJKW19] Vladimir Braverman, Shaofeng H.-C. Jiang, Robert Krauthgamer, and Xuan Wu. Coresets for ordered weighted clustering. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pages 744–753. PMLR, 2019.
- [BJKW21a] Vladimir Braverman, Shaofeng H.-C. Jiang, Robert Krauthgamer, and Xuan Wu. Coresets for clustering in excluded-minor graphs and beyond. In *SODA*, pages 2679–2696. SIAM, 2021.
- [BJKW21b] Vladimir Braverman, Shaofeng H.-C. Jiang, Robert Krauthgamer, and Xuan Wu. Coresets for clustering with missing values. In *NeurIPS*, pages 17360–17372, 2021.
- [BK18] Aditya Bhaskara and Srivatsan Kumar. Low rank approximation in the presence of outliers. In Eric Blais, Klaus Jansen, José D. P. Rolim, and David Steurer, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2018, August 20-22, 2018 - Princeton, NJ, USA*, volume 116 of *LIPICs*, pages 4:1–4:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
- [BLL18] Olivier Bachem, Mario Lucic, and Silvio Lattanzi. One-shot coresets: The case of k -clustering. In *AISTATS*, volume 84 of *Proceedings of Machine Learning Research*, pages 784–792. PMLR, 2018.
- [BRS08] Ivan Baev, Rajmohan Rajaraman, and Chaitanya Swamy. Approximation algorithms for data placement problems. *SIAM Journal on Computing*, 38(4):1411–1429, 2008.
- [BRU16] Jaroslav Byrka, Bartosz Rybicki, and Sumedha Uniyal. An approximation algorithm for uniform capacitated k -median problem with $1 + \epsilon$ capacity violation. In *IPCO*, volume 9682 of *Lecture Notes in Computer Science*, pages 262–274. Springer, 2016.
- [BT15] Nicolas Bousquet and Stéphan Thomassé. Vc-dimension and erdős-pósa property. *Discrete Mathematics*, 338(12):2302–2317, 2015.
- [BVX19] Aditya Bhaskara, Sharvaree Vadgama, and Hong Xu. Greedy sampling for approximate clustering in the presence of outliers. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 11146–11155, 2019.
- [CAL⁺22] Vincent Cohen-Addad, Kasper Green Larsen, David Saulpic, Chris Schwiegelshohn, and Omar Ali Sheikh-Omar. Improved coresets for Euclidean k -means. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022*, 2022.
- [CGTS02] Moses Charikar, Sudipto Guha, Éva Tardos, and David B. Shmoys. A constant-factor approximation algorithm for the k -median problem. *J. Comput. Syst. Sci.*, 65(1):129–149, 2002.
- [Che08] Ke Chen. A constant factor approximation algorithm for k -median clustering with outliers. In Shang-Hua Teng, editor, *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2008, San Francisco, California, USA, January 20-22, 2008*, pages 826–835. SIAM, 2008.

- [Che09] Ke Chen. On coresets for k -median and k -means clustering in metric and Euclidean spaces and their applications. *SIAM Journal on Computing*, 39(3):923–947, 2009.
- [CHK12] Marek Cygan, MohammadTaghi Hajiaghayi, and Samir Khuller. LP rounding for k -centers with non-uniform hard capacities. In *FOCS*, pages 273–282. IEEE Computer Society, 2012.
- [CKLV17] Flavio Chierichetti, Ravi Kumar, Silvio Lattanzi, and Sergei Vassilvitskii. Fair clustering through fairlets. In *NIPS*, pages 5029–5037, 2017.
- [CKMN01] Moses Charikar, Samir Khuller, David M. Mount, and Giri Narasimhan. Algorithms for facility location problems with outliers. In S. Rao Kosaraju, editor, *Proceedings of the Twelfth Annual Symposium on Discrete Algorithms, January 7-9, 2001, Washington, DC, USA*, pages 642–651. ACM/SIAM, 2001.
- [CL19] Vincent Cohen-Addad and Jason Li. On the fixed-parameter tractability of capacitated clustering. In *ICALP*, volume 132 of *LIPICs*, pages 41:1–41:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [CLLW16] Danny Z Chen, Jian Li, Hongyu Liang, and Haitao Wang. Matroid and knapsack center problems. *Algorithmica*, 75(1):27–52, 2016.
- [CLSS22] Vincent Cohen-Addad, Kasper Green Larsen, David Saulpic, and Chris Schwiegelshohn. Towards optimal lower bounds for k -median and k -means coresets. In *STOC*, pages 1038–1051. ACM, 2022.
- [CSS21] Vincent Cohen-Addad, David Saulpic, and Chris Schwiegelshohn. A new coreset framework for clustering. In *STOC*, pages 169–182. ACM, 2021.
- [DL16] H. Gökalp Demirci and Shi Li. Constant approximation for capacitated k -median with $(1+\epsilon)$ -capacity violation. In *ICALP*, volume 55 of *LIPICs*, pages 73:1–73:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.
- [FL11] Dan Feldman and Michael Langberg. A unified framework for approximating and clustering data. In *STOC*, pages 569–578. ACM, 2011.
- [FSS20] Dan Feldman, Melanie Schmidt, and Christian Sohler. Turning big data into tiny data: Constant-size coresets for k -means, pca, and projective clustering. *SIAM J. Comput.*, 49(3):601–657, 2020.
- [FZH⁺20] Qilong Feng, Zhen Zhang, Ziyun Huang, Jinhui Xu, and Jianxin Wang. A unified framework of FPT approximation algorithms for clustering problems. In *ISAAC*, volume 181 of *LIPICs*, pages 5:1–5:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [GKL03] Anupam Gupta, Robert Krauthgamer, and James R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *FOCS*, pages 534–543. IEEE Computer Society, 2003.
- [GKL⁺17] Shalmoli Gupta, Ravi Kumar, Kefu Lu, Benjamin Moseley, and Sergei Vassilvitskii. Local search methods for k -means with outliers. *Proc. VLDB Endow.*, 10(7):757–768, 2017.
- [HHHW23] Lingxiao Huang, Ruiyuan Huang, Zengfeng Huang, and Xuan Wu. On coresets for clustering in small dimensional euclidean spaces. In *International Conference on Machine Learning, 2023*. URL: <https://api.semanticscholar.org/CorpusID:257219804>.
- [HHL⁺16] Mohammadtaghi Hajiaghayi, Wei Hu, Jian Li, Shi Li, and Barna Saha. A constant factor approximation algorithm for fault-tolerant k -median. *ACM Transactions on Algorithms (TALG)*, 12(3):1–19, 2016.
- [HJLW18] Lingxiao Huang, Shaofeng H.-C. Jiang, Jian Li, and Xuan Wu. ϵ -coresets for clustering (with outliers) in doubling metrics. In *FOCS*, pages 814–825. IEEE Computer Society, 2018.
- [HJLW23] Lingxiao Huang, Shaofeng H.-C. Jiang, Jianing Lou, and Xuan Wu. Near-optimal coresets for robust clustering. In *ICLR*. OpenReview.net, 2023.
- [HJV19] Lingxiao Huang, Shaofeng H.-C. Jiang, and Nisheeth K. Vishnoi. Coresets for clustering with fairness constraints. In *NeurIPS*, pages 7587–7598, 2019.
- [HK07] Sarel Har-Peled and Akash Kushal. Smaller coresets for k -median and k -means clustering. *Discret. Comput. Geom.*, 37(1):3–19, 2007.
- [HK20] Monika Henzinger and Sagar Kale. Fully-dynamic coresets. In *ESA*, volume 173 of *LIPICs*, pages 57:1–57:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [HLW24] Lingxiao Huang, Jian Li, and Xuan Wu. On optimal coreset construction for euclidean (k, z) -clustering. *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, 2024*. URL: <https://api.semanticscholar.org/CorpusID:258426772>.
- [HM04] Sarel Har-Peled and Soham Mazumdar. On coresets for k -means and k -median clustering. In *STOC*, pages 291–300. ACM, 2004.
- [Hp11] Sarel Har-peled. *Geometric Approximation Algorithms*. American Mathematical Society, USA, 2011.
- [HSV21] Lingxiao Huang, K. Sudhir, and Nisheeth K. Vishnoi. Coresets for time series clustering. In *NeurIPS*, pages 22849–22862, 2021.
- [HV20] Lingxiao Huang and Nisheeth K. Vishnoi. Coresets for clustering in Euclidean spaces: Importance sampling is nearly optimal. In *STOC*, pages 1416–1429. ACM, 2020.
- [KKN⁺11] Ravishankar Krishnaswamy, Amit Kumar, Viswanath Nagarajan, Yogish Sabharwal, and Barna Saha. The matroid median problem. In *SODA*, pages 1117–1130. SIAM, 2011.

- [KKN⁺15] Ravishankar Krishnaswamy, Amit Kumar, Viswanath Nagarajan, Yogish Sabharwal, and Barna Saha. Facility location with matroid or knapsack constraints. *Math. Oper. Res.*, 40(2):446–459, 2015.
- [KM13] Jonathan Korman and Robert J McCann. Insights into capacity-constrained optimal transport. *Proceedings of the National Academy of Sciences*, 110(25):10064–10067, 2013.
- [KM15] Jonathan Korman and Robert McCann. Optimal transportation with capacity constraints. *Transactions of the American Mathematical Society*, 367(3):1501–1521, 2015.
- [KMS15] Jonathan Korman, Robert J McCann, and Christian Seis. Dual potentials for capacity constrained optimal transport. *Calculus of Variations and Partial Differential Equations*, 54(1):573–584, 2015.
- [KPS00] Samir Khuller, Robert Pless, and Yoram J Sussmann. Fault tolerant k -center problems. *Theoretical Computer Science*, 242(1-2):237–245, 2000.
- [Law01] Eugene L Lawler. *Combinatorial optimization: networks and matroids*. Courier Corporation, 2001.
- [Li16] Shi Li. Approximating capacitated k -median with $(1 + \epsilon)k$ open facilities. In *SODA*, pages 786–796. SIAM, 2016.
- [Li17] Shi Li. On uniform capacitated k -median beyond the natural LP relaxation. *ACM Trans. Algorithms*, 13(2):22:1–22:18, 2017.
- [LS10] Michael Langberg and Leonard J. Schulman. Universal ϵ -approximators for integrals. In *SODA*, pages 598–607. SIAM, 2010.
- [LSS12] Retsef Levi, David B. Shmoys, and Chaitanya Swamy. LP-based approximation algorithms for capacitated facility location. *Math. Program.*, 131(1-2):365–379, 2012.
- [LSV13] Jon Lee, Maxim Sviridenko, and Jan Vondrák. Matroid matching: The power of local search. *SIAM J. Comput.*, 42(1):357–379, 2013.
- [MNP⁺14] David M. Mount, Nathan S. Netanyahu, Christine D. Piatko, Ruth Silverman, and Angela Y. Wu. On the least trimmed squares estimator. *Algorithmica*, 69(1):148–183, 2014.
- [MSSW18] Alexander Munteanu, Chris Schwiegelshohn, Christian Sohler, and David P. Woodruff. On coresets for logistic regression. In *NeurIPS*, pages 6562–6571, 2018.
- [MU17] Michael Mitzenmacher and Eli Upfal. *Probability and computing: Randomization and probabilistic techniques in algorithms and data analysis*. Cambridge university press, 2017.
- [PC⁺19] Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- [Sch03] Lex Schrijver. *Combinatorial optimization - polyhedra and efficiency*. *Algorithms and Combinatorics*, 2003.
- [SRF20] Adiel Statman, Liat Rozenberg, and Dan Feldman. k -means: Outliers-resistant clustering+++. *Algorithms*, 13(12):311, 2020.
- [SS08] Chaitanya Swamy and David B. Shmoys. Fault-tolerant facility location. *ACM Trans. Algorithms*, 4(4):51:1–51:27, 2008.
- [SSS19] Melanie Schmidt, Chris Schwiegelshohn, and Christian Sohler. Fair coresets and streaming algorithms for fair k -means. In *WAOA*, volume 11926 of *Lecture Notes in Computer Science*, pages 232–251. Springer, 2019.
- [SW18] Christian Sohler and David P. Woodruff. Strong coresets for k -median and subspace approximation: Goodbye dimension. In *FOCS*, pages 802–813. IEEE Computer Society, 2018.
- [Swa16] Chaitanya Swamy. Improved approximation algorithms for matroid and knapsack median problems and applications. *ACM Trans. Algorithms*, 12(4):49:1–49:22, 2016.
- [TT13] Xiaolu Tan and Nizar Touzi. Optimal transportation under controlled stochastic dynamics. *The annals of probability*, 41(5):3201–3240, 2013.
- [TWZ⁺22] Murad Tukan, Xuan Wu, Samson Zhou, Vladimir Braverman, and Dan Feldman. New coresets for projective clustering and applications. In *AISTATS*, volume 151 of *Proceedings of Machine Learning Research*, pages 5391–5415. PMLR, 2022.
- [VC71] VN Vapnik and A Ya Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264, 1971.
- [vH14] Ramon van Handel. *Probability in high dimension*. 2014.
- [ZG03] Shi Zhong and Joydeep Ghosh. A unified framework for model-based clustering. *J. Mach. Learn. Res.*, 4:1001–1037, 2003.

A Application of Theorem 3.1: Simultaneous Coresets for Multiple \mathcal{B} 's

Since our coreset can handle all capacity constraints simultaneously, one may be interested in whether our coresets can also handle multiple assignment structure constraints \mathcal{B} 's simultaneously as well. We show that this is indeed possible for several families of \mathcal{B} . A similar idea, called simultaneous or one-shot coresets, has also been considered in the literature [BLL18, BJKW19]. This type of coresets is particularly useful when some key hyper-parameters, in our case \mathcal{B} , are not given/known but need to be figured out by experiments, since a single coreset can be reused to support all the experiments. Fortunately, our coreset in Theorem 1.1 has a powerful feature that it does not

use any specific structure of the given \mathcal{B} except an upper bound of $\text{Lip}(\mathcal{B})$ in the algorithm. Hence, for a family \mathcal{F} of $m \geq 1$ different \mathcal{B} 's, one can apply Theorem 1.1 with a union bound, so that with an additional $\text{poly}(\log(m))$ factor in the coreset size, the returned coreset is a coreset for all \mathcal{B} 's in \mathcal{F} simultaneously. Below, we discuss the size bounds of the simultaneous coresets for two useful families \mathcal{F} and their potential applications.

1. Let \mathcal{F} be the collection of all (scaled) uniform matroid basis polytopes ($\text{Lip}(\mathcal{B}) \leq 2$). Since the ground set is $[k]$, we have $|\mathcal{F}| = k$ (each corresponding to a cardinality constraint). Then we can achieve a simultaneous coreset S for every $\mathcal{B} \in \mathcal{F}$ by increasing the coreset size in Theorem 1.1 by a multiplicative $\log k$ factor, say $\tilde{O}_z(\Lambda_\varepsilon(\mathcal{X}) \cdot k^2 \varepsilon^{-2z})$. Consequently, S is a simultaneous coreset for all fault-tolerant clusterings.
2. Let \mathcal{F} be the collection of all (scaled) partition matroid basis polytopes ($\text{Lip}(\mathcal{B}) \leq 3$). Since the ground set is $[k]$, there are at most k^k partition ways. For each partition, there are at most k^k different selections of rank functions (see Definition 4.3). Thus, we have $|\mathcal{F}| \leq k^{2k}$ and we can achieve a simultaneous coreset for every $\mathcal{B} \in \mathcal{F}$ by increasing the coreset size in Theorem 1.1 by a multiplicative $k \log(2k)$ factor, say $\tilde{O}_z(\Lambda_\varepsilon(\mathcal{X}) \cdot k^3 \varepsilon^{-2z})$.

B Missing Proofs in Section 2

B.1 Proof of Claim 1: Capacitated Clustering

Proof. Fix a center set $C \in \mathcal{X}^k$. Let σ^* be an assignment that achieves the optimal capacitated clustering cost on P with respect to C , i.e.,

$$\min_{\substack{\sigma: \|\sigma\|_1 = n-m, \sigma \sim \mathcal{B} \\ \ell_c \leq \|\sigma(\cdot, c)\|_1 \leq u_c, \forall c \in C}} \text{cost}_z^\sigma(P, C) = \text{cost}_z^{\sigma^*}(P, C).$$

Let $h^* \in (n-m) \cdot \mathcal{B}$ such that $\sigma^* \sim h^*$. Then we have

$$\text{cost}_z^{\sigma^*}(P, C) = \text{cost}_z(P, C, \mathcal{B}, h^*).$$

Since $\ell_c \leq \|\sigma(\cdot, c)\|_1 \leq u_c$ for $\forall c \in C$, we know that $\ell_c \leq h_c \leq u_c$. Thus, we have

$$\text{cost}_z^{\sigma^*}(P, C) = \min_{\substack{h: h^* \in (n-m) \cdot \mathcal{B} \\ \ell_c \leq h_c \leq u_c, \forall c \in C}} \text{cost}_z^\sigma(P, C).$$

Similarly, we have

$$\min_{\substack{\sigma: \|\sigma\|_1 = n-m, \sigma \sim \mathcal{B} \\ \ell_c \leq \|\sigma(\cdot, c)\|_1 \leq u_c, \forall c \in C}} \text{cost}_z^\sigma(S, C) = \min_{\substack{h: h^* \in (n-m) \cdot \mathcal{B} \\ \ell_c \leq h_c \leq u_c, \forall c \in C}} \text{cost}_z^\sigma(S, C).$$

By Definition 2.4, $\text{cost}_z(S, C, \mathcal{B}, h^*) \leq (1 + \varepsilon) \text{cost}_z(P, C, \mathcal{B}, h^*)$. We also have

$$\min_{\substack{\sigma: \|\sigma\|_1 = n-m, \sigma \sim \mathcal{B} \\ \ell_c \leq \|\sigma(\cdot, c)\|_1 \leq u_c, \forall c \in C}} \text{cost}_z^\sigma(S, C) \leq \text{cost}_z(S, C, \mathcal{B}, h^*)$$

Combining the above two inequalities, we conclude that

$$\min_{\substack{\sigma: \|\sigma\|_1 = n-m, \sigma \sim \mathcal{B} \\ \ell_c \leq \|\sigma(\cdot, c)\|_1 \leq u_c, \forall c \in C}} \text{cost}_z^\sigma(S, C) \leq (1 + \varepsilon) \cdot \min_{\substack{\sigma: \|\sigma\|_1 = n-m, \sigma \sim \mathcal{B} \\ \ell_c \leq \|\sigma(\cdot, c)\|_1 \leq u_c, \forall c \in C}} \text{cost}_z^\sigma(P, C).$$

The other direction can be proved in the same way, which completes the proof. \square

B.2 Proof of Claim 2: Fair Clustering

Proof. The idea is similar to that of [HJV19, Theorem 4.2], while we additionally consider a total capacity constraint. Fix a center set $C \in \mathcal{X}^k$.

We first consider the simple case that $\Gamma = 1$ in which all \mathcal{G}_p 's are the same. Fix a center c . For each $i \in \mathcal{G}_p$, we have that $\frac{|\sigma^{-1}(c) \cap G_i|}{|\sigma^{-1}(c)|} = 1$; and for each $i \notin \mathcal{G}_p$, we have that $\frac{|\sigma^{-1}(c) \cap G_i|}{|\sigma^{-1}(c)|} = 0$. Then if there exists some group G_i with $(i \in \mathcal{G}_p) \wedge (u_i < 1)$, or with $(i \notin \mathcal{G}_p) \wedge (u_i > 0)$ holds, there is no feasible solution for the fair clustering problem with outliers. Otherwise, the fair clustering problem with outliers is reduced to a clustering problem with outliers. Thus, by the assumption of Claim 2, an ε -coreset exists for the fair clustering problem with outliers of size at most A .

For the general case that $\Gamma \geq 1$, we partition P into Γ disjoint subsets P_1, \dots, P_Γ in which all points $p \in P_j$ correspond to the same collection \mathcal{G}_p for each $j \in [\Gamma]$. For each $j \in [\Gamma]$, we let \mathcal{G}_j denote the corresponding collection \mathcal{G}_p of $p \in P_j$. Suppose $S_j \subseteq P_j$ is an $(\varepsilon, \Delta_k, m')$ -coreset of P_j of size at most A for all $0 \leq m' \leq m$. Then it suffices to prove that $S := \bigcup_{j \in [\Gamma]} S_j$ is an ε -coreset for the fair clustering problem with outliers.

The remaining proof is similar to that of Claim 1. Let $h^{(j)} \in \mathbb{R}_{\geq 0}^k$ be the capacity vector of group $j \in [\Gamma]$. Also, let $\sigma^{(j)} : P_j \times C \rightarrow \mathbb{R}_{\geq 0}$ be the assignment function of group j . Let $\sigma : P \times C \rightarrow \mathbb{R}_{\geq 0}$ be the union assignment function of $\sigma^{(1)}, \dots, \sigma^{(s)}$. We have for every $i \in [s]$ and $c \in C$,

$$|\sigma^{-1}(c)| = \sum_{j \in [\Gamma]} h_c^{(j)}, \text{ and } |\sigma^{-1}(c) \cap G_i| = \sum_{j \in [\Gamma]} \mathbb{I}[i \in \mathcal{G}_j] \cdot h_c^{(j)},$$

where $\mathbb{I}[\cdot]$ is the indicator function. Then a fairness constraint $\ell_i \leq \frac{|\sigma^{-1}(c) \cap G_i|}{|\sigma^{-1}(c)|} \leq u_i$ is equivalent to

$$\ell_i \leq \frac{\sum_{j \in [\Gamma]} \mathbb{I}[i \in \mathcal{G}_j] \cdot h_c^{(j)}}{\sum_{j \in [\Gamma]} h_c^{(j)}} \leq u_i.$$

Thus, to check whether σ satisfies all fairness constraints, we only need the information of vectors $h^{(1)}, \dots, h^{(\Gamma)}$. Consequently, there must exist vectors $h^{(1)}, \dots, h^{(\Gamma)}$ with $\sum_{j \in [\Gamma]} h^{(j)} = h$, $|P_j| - m \leq \|h^{(j)}\|_1 \leq |P_j|$ for all $j \in [\Gamma]$ and

$$\ell_i \leq \frac{\sum_{j \in [\Gamma]} \mathbb{I}[i \in \mathcal{G}_j] \cdot h_c^{(j)}}{\sum_{j \in [\Gamma]} h_c^{(j)}} \leq u_i$$

for all $i \in [s]$ and $c \in C$, such that

$$\min_{\substack{\sigma: \|\sigma\|_1 = n-m \\ \ell_i \leq \|\sigma(\cdot, c)\|_1 \leq u_i, \forall i \in [s], c \in C}} \text{cost}_z^\sigma(P, C) = \sum_{j \in [\Gamma]} \text{cost}_z(P_j, C, \Delta_k, h^{(j)}).$$

By the definition of S_j , we know that

$$\begin{aligned} \sum_{j \in [\Gamma]} \text{cost}_z(P_j, C, \Delta_k, h^{(j)}) &\geq (1 - \varepsilon) \cdot \sum_{j \in [\Gamma]} \text{cost}_z(S_j, C, \Delta_k, h^{(j)}) \\ &\geq (1 - \varepsilon) \cdot \min_{\substack{\sigma: \|\sigma\|_1 = n-m \\ \ell_i \leq \|\sigma(\cdot, c)\|_1 \leq u_i, \forall i \in [s], c \in C}} \text{cost}_z^\sigma(S, C). \end{aligned}$$

Thus, we conclude that

$$\min_{\substack{\sigma: \|\sigma\|_1 = n-m \\ \ell_i \leq \|\sigma(\cdot, c)\|_1 \leq u_i, \forall i \in [s], c \in C}} \text{cost}_z^\sigma(P, C) \geq (1 - \varepsilon) \cdot \min_{\substack{\sigma: \|\sigma\|_1 = n-m \\ \ell_i \leq \|\sigma(\cdot, c)\|_1 \leq u_i, \forall i \in [s], c \in C}} \text{cost}_z^\sigma(S, C).$$

Similarly, we can prove that

$$\min_{\substack{\sigma: \|\sigma\|_1 = n-m \\ \ell_i \leq \|\sigma(\cdot, c)\|_1 \leq u_i, \forall i \in [s], c \in C}} \text{cost}_z^\sigma(P, C) \leq (1 + \varepsilon) \cdot \min_{\substack{\sigma: \|\sigma\|_1 = n-m \\ \ell_i \leq \|\sigma(\cdot, c)\|_1 \leq u_i, \forall i \in [s], c \in C}} \text{cost}_z^\sigma(S, C).$$

This completes the proof. \square

C Proof of Lemma 3.7: Relation between Two Covering Notions

LEMMA C.1. (**Restatement of Lemma 3.7**) *Let \mathcal{B} be an assignment structure constraint. For every $\beta > 0$ and $t \in [k]$, we have*

$$\begin{aligned} \log N_{\mathcal{X}}(R, t, \beta, \mathcal{B}) &\leq O(t \cdot \log N_{\mathcal{X}}(n_R, \beta)) + zk \cdot \log(\text{Lip}(\mathcal{B}) \cdot z\varepsilon^{-1}\beta^{-1}) \\ &\leq O(\Lambda_{\beta}(\mathcal{X}) \cdot t \log n_R + zk \cdot \log(\text{Lip}(\mathcal{B}) \cdot z\varepsilon^{-1}\beta^{-1})). \end{aligned}$$

Moreover, when $\mathcal{B} = \Delta_k$, we have

$$\log N_{\mathcal{X}}(R, t, \beta, \mathcal{B}) \leq O(\Lambda_{\beta}(\mathcal{X}) \cdot t \log n_R + zk \cdot \log(zk\varepsilon^{-1}\beta^{-1})).$$

Proof. For preparation, we provide the following lemma.

LEMMA C.2. (**Extension of $\text{Lip}(\mathcal{B})$ to $\text{conv}(\mathcal{B}^o)$**) *Suppose there are real numbers $a \geq b > 0$, constraints $h \in a\mathcal{B}, h' \in b\mathcal{B}$, and $\sigma \sim (a\mathcal{B}, h)$ then there exists $\sigma' \sim (b\mathcal{B}, h')$ such that $\|\sigma - \sigma'\| \leq 3\text{Lip}(\mathcal{B}) \cdot \|h - h'\|_1$.*

Proof. By definition of $\text{Lip}(\mathcal{B})$, there exists $\pi \sim (a\mathcal{B}, \frac{a}{b}h')$ such that

$$\|\pi - \sigma\|_1 \leq \text{Lip}(\mathcal{B}) \cdot \|h - \frac{a}{b}h'\|_1.$$

Let $\sigma' = \frac{b}{a}\pi$, obviously $\sigma' \sim (b\mathcal{B}, h')$. Moreover,

$$\begin{aligned} \|\sigma' - \sigma\|_1 &\leq \|\sigma' - \pi\|_1 + \|\sigma - \pi\|_1 \\ &\leq \frac{|a-b|}{b} \cdot \|\sigma'\|_1 + \text{Lip}(\mathcal{B}) \cdot \|h - \frac{a}{b}h'\|_1 \\ &\leq |a-b| + \text{Lip}(\mathcal{B}) \cdot (\|h - h'\|_1 + \frac{|a-b|}{b} \cdot \|h'\|_1) \\ &= |a-b| + \text{Lip}(\mathcal{B}) \cdot (\|h - h'\|_1 + |a-b|) \\ &\leq 3\text{Lip}(\mathcal{B}) \cdot \|h - h'\|_1 \end{aligned}$$

where for the last inequality, we use the triangle inequality to obtain $\|h - h'\|_1 \geq \|h\|_1 - \|h'\|_1 = |a-b|$. □

We first show how to construct an β -covering $\mathcal{F} \subset \Phi_t$ w.r.t. (R, \mathcal{B}) with the desired size and covering property. For center sets, let \mathcal{C} denote the β -covering of $B(c_i^*, \frac{48zr}{\varepsilon})$ and for technical convenience we also add c_i^* into \mathcal{C} . Let $\mathcal{C}[t]$ denote the collection of k -tuples (c_1, c_2, \dots, c_k) satisfying that

- exactly $k-t$ centers $c = c_i^*$;
- the remaining t centers are selected from \mathcal{C} .

For capacity constraints, let \mathcal{N} denote an l_1 -distance $\frac{\beta \cdot \varepsilon^z}{12 \cdot (48z)^z \cdot \text{Lip}(\mathcal{B})}$ -net of $\text{conv}(\mathcal{B}^o)$, namely, for every $h \in \text{conv}(\mathcal{B}^o)$, there exists $h' \in \mathcal{N}$ such that $\|h - h'\|_1 \leq \frac{\beta \cdot \varepsilon^z}{12 \cdot (48z)^z \cdot \text{Lip}(\mathcal{B})}$. Since $\mathcal{B} \subseteq \Delta_k$, we know that $|\mathcal{N}| \leq (\frac{z \cdot \text{Lip}(\mathcal{B})}{\beta \cdot \varepsilon^z})^{O(k)}$. We let $\mathcal{F} := \mathcal{C}[t] \times n_R \mathcal{N}$ to be the Cartesian product of \mathcal{C} and $n_R \mathcal{N}$. By construction, we have

$$\log |\mathcal{F}| \leq t \log k + t \log |\mathcal{C}| + \log |\mathcal{N}| \leq O(t \cdot \log N_{\mathcal{X}}(n, \beta) + zk \cdot \log(\text{Lip}(\mathcal{B}) \cdot z\beta^{-1}\varepsilon^{-1})).$$

Then it remains to show that \mathcal{F} is an (t, β) -covering w.r.t. (R, \mathcal{B}) .

Now we fix a pair $(C, h) \in \Phi_t$. We show that there exists $(C', h') \in \mathcal{F}$ such that for every $Q \subseteq R, w_Q(Q) = n_R$,

$$\text{cost}_z(Q, C, \mathcal{B}, h) \in (1 \pm (\beta + \varepsilon)) \cdot \text{cost}_z(Q, C', \mathcal{B}, h') \pm \beta n_R r^z.$$

By Lemma 3.6, it suffices to prove that

$$(C.1) \quad \text{cost}_z(Q, \nu(C), \mathcal{B}, h) \in (1 \pm \beta) \cdot \text{cost}_z(Q, C', \mathcal{B}, h') \pm \beta n_R r^z.$$

For the construction of C' , we let $c' = c_i^*$ for every $c \in C$ with $\nu(c) = c_i^*$, and let c' be the closest point in C for the remaining $c \in C$. For the construction of h' , we know that there exists $\tilde{h} \in n_R \mathcal{N}$ such that $\|\tilde{h} - h\|_1 \leq \frac{\beta \cdot \varepsilon^z}{12 \cdot (48z)^z \cdot \text{Lip}(\mathcal{B})} \cdot n_R$. We define a capacity constraint h' on C' such that $\forall c \in C, h'(c') = \tilde{h}(c)$.

In the following, we fix a weighted $Q \subseteq \text{Ball}(a, r)$ with total weight n_R and prove Inequality (C.1). For ease of analysis, we slightly abuse the notation by using C to replace $\nu(C)$ in the following. Assume σ is the corresponding optimal assignment for $\text{cost}_z(Q, C, \mathcal{B}, h)$ and expand it as

$$\text{cost}_z(Q, C, \mathcal{B}, h) = \sum_{p \in Q} \sum_{c \in C} \sigma(p, c) \cdot d^z(p, c).$$

It suffices to prove the following inequality:

$$(C.2) \quad \text{cost}_z(Q, C', \mathcal{B}, h') \in (1 \pm \beta) \cdot \sum_{p \in Q} \sum_{c \in C} \sigma(p, c) \cdot d^z(p, c) \pm \beta n_R r^z.$$

We prove the two directions separately. Firstly, let σ' denote an optimal assignment of $\text{cost}_z(Q, C', \mathcal{B}, h')$, we prove that

$$(C.3) \quad \text{cost}_z(Q, C', \mathcal{B}, h') \geq (1 - \beta) \cdot \sum_{p \in Q} \sum_{c \in C} \sigma(p, c) \cdot d^z(p, c) - \beta n_R r^z.$$

For the sake of contradiction, suppose (C.3) does not hold. By Lemma C.2, we know that there exists an assignment $\sigma'' \sim (\mathcal{B}, h')$ such that

$$(C.4) \quad \sum_{p \in Q} \sum_{c \in C} |\sigma''(p, c') - \sigma'(p, c)| \leq \frac{\beta}{4} \cdot \left(\frac{\varepsilon}{48z}\right)^z \cdot n_R$$

It suffices to show that σ'' is actually a better assignment than σ to conclude a contradiction. We can see that

$$\begin{aligned} & \text{cost}_z^{\sigma''}(Q, C) \\ &= \sum_{p \in Q} \sum_{c \in C} \sigma''(p, c') \cdot d^z(p, c) \\ &= \sum_{p \in Q} \sum_{c \in C} \sigma'(p, c') \cdot d^z(p, c) + \sum_{p \in Q} \sum_{c \in C} |\sigma'(p, c') - \sigma''(p, c)| \cdot d^z(p, c)^z \\ (C.5) \quad & \leq \sum_{p \in Q} \sum_{c \in C} \sigma'(p, c') \cdot d^z(p, c) + \|\sigma'' - \sigma'\|_1 \cdot \left(\frac{48zr}{\varepsilon}\right)^z \quad (\text{Defn. of } \bar{C}_{\text{far}}) \\ & \leq \sum_{p \in Q} \sum_{c \in C} \sigma'(p, c') \cdot d^z(p, c) + \frac{\beta}{4} n_R r^z. \quad (\text{Ineq. (C.4)}) \end{aligned}$$

By Lemma 3.5, we know that

$$\begin{aligned} & \sum_{p \in Q} \sum_{c \in C} \sigma'(p, c') \cdot d^z(p, c) \\ &= \sum_{p \in Q} \sum_{c \in C} \sigma'(p, c') \cdot (d^z(p, c') + d^z(p, c) - d^z(p, c')) \\ (C.6) \quad & \leq (1 + \beta) \cdot \sum_{p \in Q} \sum_{c \in C} \sigma'(p, c') \cdot d^z(p, c') + \left(\frac{3z}{\beta}\right)^{z-1} \cdot n_R \cdot \left(\frac{\beta r}{12z}\right) \quad (\text{Lemma 3.5}) \\ & \leq (1 + \beta) \cdot \text{cost}_z(Q, C', \mathcal{B}, h') + \frac{\beta}{4} n_R r^z. \end{aligned}$$

Summing up the above three inequalities, we conclude that

$$\begin{aligned} \text{cost}_z^{\sigma''}(Q, C) &\leq (1 + \beta) \cdot \text{cost}_z(Q, C', \mathcal{B}, h') + \frac{\beta}{2} n_R r^z \quad (\text{Ineqs. (C.5) and (C.6)}) \\ &< \text{cost}_z^{\sigma}(Q, C), \quad (\text{by assumption}) \end{aligned}$$

which has been a contradiction to the fact that σ is an optimal assignment for $\text{cost}_z(Q, C, \mathcal{B}, h)$.

To prove the other direction of (C.2), it suffices to construct an assignment $\sigma' \sim (\mathcal{B}, h')$ such that

$$\text{cost}_z^{\sigma'}(Q, C) \leq (1 + \beta) \sum_{p \in Q} \sum_{c \in C} \sigma(p, c) \cdot d^z(p, c) + \beta n_R r^z.$$

To this end, we simply choose σ' to be an assignment consistent with (\mathcal{B}, h') and

$$\sum_{p \in Q} \sum_{c \in C} |\sigma'(p, c) - \sigma(p, c)| \leq \frac{\beta}{4} \cdot \left(\frac{\varepsilon}{48z}\right)^z \cdot n_R.$$

Again, by Lemma C.2, such σ' exists. Then by a similar argument as in the first direction, we can verify that

$$\sum_{p \in Q} \sum_{c \in C} \sigma'(p, c') \cdot d^z(p, c') \leq (1 + \beta) \sum_{p \in Q} \sum_{c \in C} \sigma(p, c) \cdot d^z(p, c) + \beta n_R r^z$$

which completes the proof of the first inequality of Lemma C.1.

When $\mathcal{B} = \Delta_k$, the only difference is the construction of \mathcal{N} . Note that for any $C \subset B(c_i^*, \frac{48zr}{\varepsilon})$ and any $h, h' \in n_R \text{conv}(\mathcal{B}^o)$ satisfying that 1) $h_c = h'_c$ for every $c \in C$ with $c \neq c_i^*$; and 2) $\sum_{c \in C: c=c_i^*} h_c = \sum_{c \in C: c=c_i^*} h'_c$, the following holds:

$$\text{cost}_z(Q, C, \Delta_k, h) = \text{cost}_z(Q, C, \Delta_k, h').$$

This observation enables us to construct \mathcal{N} as follows:

1. Enumerate all subsets $A \subseteq [k]$ of size $t - 1$ and let $\Delta_A = \{h \in \Delta_k : \forall i \in A, h_i = 0\}$.
2. Construct an l_1 -distance $\frac{\beta \cdot \varepsilon^z}{12 \cdot (48z)^z \cdot \text{Lip}(\mathcal{B})}$ -net for every Δ_A and let \mathcal{N} be their union.

We still let $\mathcal{F} := \mathcal{C}[t] \times n_R \mathcal{N}$, which can be shown an (t, β) -covering w.r.t. (R, Δ_k) . By construction, we have

$$\log |\mathcal{F}| \leq t \log k + t \log |\mathcal{C}| + \log |\mathcal{N}| \leq O\left(t \cdot \log N_{\mathcal{X}}(n, \beta) + zt \cdot \log(zk\beta^{-1}\varepsilon^{-1})\right),$$

which completes the proof of the second inequality of Lemma C.1. \square