

On Tight FPT Time Approximation Algorithms for k -Clustering Problems

Han Dai ✉ 

School of Computer Science, Nanjing University, China

Shi Li ✉ 

School of Computer Science, Nanjing University, China

Sijin Peng ✉ 

CSAIL, Massachusetts Institute of Technology, Cambridge, MA, USA

Abstract

Following recent advances in combining approximation algorithms with fixed-parameter tractability (FPT), we study FPT-time approximation algorithms for minimum-norm k -clustering problems, parameterized by the number k of open facilities.

For the capacitated setting, we give a tight $(3+\epsilon)$ -approximation for the general-norm capacitated k -clustering problem in FPT-time parameterized by k and ϵ . Prior to our work, such a result was only known for the capacitated k -median problem [25]. As a special case, our result yields an FPT-time 3-approximation for capacitated k -center. The problem has not been studied in the FPT-time setting, with the previous best known polynomial-time approximation ratio being 9 [6].

In the uncapacitated setting, we consider the **top- cn** norm k -clustering problem, where the goal of the problem is to minimize the **top- cn** norm of the connection distance vector. Our main result is a tight $(1 + \frac{2}{ec} + \epsilon)$ -approximation algorithm for the problem with $c \in (\frac{1}{e}, 1]$. (For the case $c \leq \frac{1}{e}$, there is a simple tight $(3 + \epsilon)$ -approximation.) Our framework can be easily extended to give a tight $(3, 1 + \frac{2}{e} + \epsilon)$ -bi-criteria approximation for the (k -center, k -median) problem in FPT time, improving the previous best polynomial-time $(4, 8)$ guarantee [5].

All results are based on a unified framework: computing a $(1 + \epsilon)$ -approximate solution using $O(\frac{k \log n}{\epsilon})$ facilities S via LP rounding, sampling a few client representatives R based on the solution S , guessing a few pivots from $S \cup R$ and some radius information on the pivots, and solving the problem using the guesses. We believe this framework can lead to further results on k -clustering problems.

2012 ACM Subject Classification Theory of computation $\rightarrow$ Design and analysis of algorithms; Theory of computation $\rightarrow$ Fixed parameter tractability

Keywords and phrases Approximation algorithms, Monotone symmetric norms, Clustering, Fixed parameter tractability

Digital Object Identifier 10.4230/LIPIcs.ICALP.2026.72

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: <https://arxiv.org/abs/2512.04614>

Funding *Han Dai*: Supported by State Key Laboratory for Novel Software Technology, and New Cornerstone Science Foundation.

Shi Li: Supported by State Key Laboratory for Novel Software Technology, and New Cornerstone Science Foundation.

1 Introduction

In the supplier setting of clustering problems, we are given a set F of facilities, a set C of n clients, a metric d over $F \cup C$, and a non-negative integer k . The objective is to open a set $S \subseteq F$ of k facilities and assign each client to an open facility, so as to minimize



© Han Dai, Shi Li, and Sijin Peng;

licensed under Creative Commons License CC-BY 4.0

53rd International Colloquium on Automata, Languages, and Programming (ICALP 2026).

Editors: Sayan Bhattacharya, Danupon Nanongkai, Michael Benedikt, and Gabriele Puppis;

Article No. 72; pp. 72:1–72:23



Leibniz International Proceedings in Informatics

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



some function on the connection distances. Chakrabarty and Swamy [15] introduced the *minimum-norm k -clustering problem*, where the goal is to minimize a monotone symmetric norm $f : \mathbb{R}_{\geq 0}^C \rightarrow \mathbb{R}_{\geq 0}$ applied to the connection distances. Specifically, the objective is to find a set $S \subseteq F$ of k facilities so as to minimize $f((d(j, S))_{j \in C})$, where $d(j, S) := \min_{i \in S} d(j, i)$.

When the norm f is L_∞ , L_1 or L_2 norm, the problem respectively corresponds to the well-known *k -supplier-center*, *k -median* and *metric k -means*¹ problems. These problems have been extensively studied in the literature [4, 17, 12, 37, 34, 22, 30, 28, 40, 21, 31, 33, 42]. The best known approximation ratios for the three problems are respectively 3 [47], $2 + \epsilon$ [23] and 6.357 [4], while the known hardness of approximation bounds are 3 , $1 + \frac{2}{e}$ and $1 + \frac{8}{e}$ [38, 36].

Two other important classes of monotone symmetric norms studied in the literature are the *top- ℓ* norms and *ordered* norms. For any $\ell \in [n]$, the *top- ℓ* norm of $v \in \mathbb{R}_{\geq 0}^n$ is the sum of ℓ largest coordinates of v . An ordered norm is a non-negative linear combination of several *top- ℓ* norms. The clustering problem with this objective is called the *ordered k -median* problem. Independently, Byrka, Sornat, and Spoerhase [14] and Chakrabarty and Swamy [15] gave the first constant-factor approximation algorithms for ordered k -median. The approximation ratio was subsequently improved by Chakrabarty and Swamy [16] to $5 + \epsilon$.

Another aspect that naturally arises in clustering problems is the presence of facility capacities. In the supplier setting, a facility may not be able to serve too many clients. In clustering applications, there may be upper bounds on the sizes of clusters. In the *minimum-norm capacitated k -clustering* problem, each facility $i \in F$ is additionally given a capacity $u_i \in \mathbb{Z}_{>0}$. In a feasible solution, every facility $i \in S$ can serve at most u_i clients. Formally, we need to output a set $S \subseteq F$ of k facilities and an assignment vector $\sigma \in S^C$ such that $|\sigma^{-1}(i)| \leq u_i$ for every $i \in S$, so as to minimize $f(d(j, \sigma_j)_{j \in C})$.

The main problems that have been studied in the capacitated setting are the capacitated k -center and capacitated k -median problems. For capacitated k -center, the current best approximation ratio is 9, due to [6], which improves upon the earlier $O(1)$ -approximation of Cygan, Hajiaghayi and Khuller [26]. When capacities are uniform, a better approximation ratio of 6 is known [8, 41]. In contrast, the approximation status of the capacitated k -median problem is less satisfactory. Many $O(1)$ bi-criteria approximation algorithms are known, which either violate the limit k on the number of open facilities [1] or the capacity constraints [20, 43, 11]. In both settings, the violation factor can be reduced to $1 + \epsilon$ [45, 44, 13, 27]. For true approximation algorithms, the folklore result of $O(\log n)$ -distortion embedding of any n -point metric into a distribution of HST metrics [29] leads to an $O(\log n)$ -approximation for the problem. This ratio was improved to $O(\log k)$ by Adamczyk, Byrka, Marcinkowski, Meesum and Włodarczyk [3], who applied the HST embedding technique, but on a k -point metric obtained by considering the k -center objective. This remains the current best approximation ratio for the capacitated k -median problem, and obtaining a true $O(1)$ -approximation for this problem – without any violation – remains a notorious open question.

Recently, combining paradigms of approximation algorithms and *Fixed-Parameter-Tractability (FPT)* has led to exciting new results, especially for problems where the improvement on polynomial time approximation algorithms has stalled, including many clustering problems. In many clustering applications, k is a small number, allowing us to have a running time of the form $g(k) \cdot n^{O(1)}$, where g is a function only of k . Such a running time is called FPT time, parameterized by k .

¹ The objective of the metric k -means problem is to minimize the sum of squared distances, which is equivalent to minimize the L_2 norm.

Cohen-Addad, Gupta, Kumar, Lee and Li [24] studied the k -median and metric k -means problems in FPT time, obtaining approximation factors of $1 + \frac{2}{\epsilon} + \epsilon$ and $1 + \frac{8}{\epsilon} + \epsilon$ respectively. On the negative side, these approximation factors are tight in FPT time under the assumption $\text{FPT} \neq W[1]$ [24]. For the capacitated k -median problem, Cohen-Addad and Li [25] developed a tight 3-approximation in FPT time. More recently, Abbasi, Banerjee, Byrka, Chalmersook, Gaddekar, Khodamoradi, Marx, Sharma, and Spoerhase [2] proposed an Efficient Parameterized Approximation Scheme (EPAS) for the minimum-norm k -clustering problem on metrics with bounded ϵ -scatter dimension, which include Euclidean metrics, metrics of bounded doubling dimension and planar metrics. This scheme achieves a $(1 + \epsilon)$ -approximation in $g(k, \epsilon) \cdot \text{poly}(n)$ time, for some function g on k and ϵ .

Goyal and Jaiswal [35] studied FPT time approximation algorithms for a family of constrained clustering problems, with the objective of minimizing the z -th power of the maximum connection distance for any $z > 0$. A main constraint they consider is the *cluster-size constraint*: we are given k size bounds $r_1, r_2, \dots, r_k \geq 0$, and the partition $(O_1, O_2, \dots, O_k)$ of clients must satisfy $|O_1| \leq r_1, |O_2| \leq r_2, \dots, |O_k| \leq r_k$ (the partition can be arbitrarily permuted). They also considered many other constraints such as lower bounds on cluster sizes, color constraints, and fault-tolerant requirements. They developed a general FPT-time framework for these problems that achieve an approximation ratio of 2^z for k -center-type problems (where facilities can be put anywhere in the metric), and 3^z for k -supplier-type problems (where facilities can only be put at specified locations). Both ratios were shown to be tight under GAP-ETH.

In particular, the problem most closely related to the capacitated k -center problem we study is the r -capacity k -supplier problem. The goal is to find a clustering of clients satisfying the cluster-size constraint and choose a facility for each cluster, so as to minimize the maximum connection distance. When all size bounds $r_1, r_2, \dots, r_k$ are the same, the problem reduces to the soft and uniform capacitated k -center problem where each facility i has capacity $u_i = r_1$. However, as noted in [35], for general (non-uniform) size bounds, there is no easy reduction between the two problems, since one imposes capacities on clusters whereas the other imposes capacities on facilities.

1.1 Our Results

Following this line of research, we study approximation algorithms for k -clustering problems in FPT time, with k being the parameter.

Our first result is a $(3 + \epsilon)$ -approximation algorithm for the minimum-norm capacitated k -clustering problem in FPT time parameterized by k and ϵ .

► **Theorem 1.** *For any $\epsilon > 0$, there is a $(3 + \epsilon)$ -approximation algorithm for the minimum-norm capacitated k -clustering problem, that runs in time $g(k, \epsilon) \cdot \text{poly}(n)$, where g is a computable function depending on k and ϵ .*

In particular, this implies

► **Corollary 2.** *There is a 3-approximation algorithm for the capacitated k -center problem, that runs in time $g(k) \cdot \text{poly}(n)$, where g is a computable function depending on k .²*

² For the k -center objective, $3 + \epsilon$ can be reduced to 3. By guessing and scaling distances, we can assume the optimum value is 1. Changing the distance $d(i, j)$ to $\lceil d(i, j) \rceil$ does not change the optimum solution of the instance. For such an instance, a $(3 + \epsilon)$ -approximation is a 3-approximation if $\epsilon < 1$.

To the best of our knowledge, the result was previously only known for capacitated k -median [25]. In particular, it was not known even for the capacitated k -center problem with general capacities. The best known approximation ratio for the problem remains 9, achieved by a polynomial time algorithm due to [6]. Our result gives a 3-approximation in FPT time. This is tight under the assumption that $\text{FPT} \neq W[1]$, even for the (uncapacitated) k -supplier center problem.

We then turn to the uncapacitated setting. Our main focus is the $\text{top-}cn$ norm for a constant $c \in (0, 1]$. (The problem can be referred to as the cn -centrum problem in the literature.) On the negative side, the problem is hard to approximate within a factor better than $\min\{3, 1 + \frac{2}{ec}\}$ in FPT time, assuming $\text{FPT} \neq W[1]$ [46]. We complement this with a matching positive result:

► **Theorem 3.** *For any $\epsilon > 0$, there is a $(\min\{3, 1 + \frac{2}{ec}\} + \epsilon)$ -approximation algorithm for the $\text{top-}cn$ norm k -clustering problem with running time $(\frac{k}{\epsilon})^{O(k)} \cdot n^{O(1)}$.*

The theorem suggests that considering the $\text{top-}\ell$ norm for $\ell = cn$ is appropriate, as the approximation ratio is a function of $\frac{\ell}{n}$. The negative result implies that if $\frac{\ell}{n} < \frac{1}{e} + \epsilon$, then the problem is hard to approximate within a factor better than 3 even in FPT time, assuming $\text{FPT} \neq W[1]$. On the other hand, achieving a $(3 + \epsilon)$ -approximation for minimum-norm k -clustering in FPT time is easy (See Corollary 14).

Finally, we show that the framework can be easily extended to give a $(3, 1 + \frac{2}{e} + \epsilon)$ -bi-criteria approximation for the clustering problem with both k -center and k -median objectives. This was introduced by Alamdari and Shmoys [5], who gave a polynomial-time $(4, 8)$ -approximation. We improve the bi-criteria approximation factors to $(3, 1 + \frac{2}{e} + \epsilon)$, albeit with FPT time. Indeed, our result works for the more general $\text{top-}cn$ norm. See Section 5 for formal definitions used in the theorem:

► **Theorem 4.** *For a constant $\epsilon > 0$, there is a $(3, \min\{3, 1 + \frac{2}{ec}\} + \epsilon)$ -bi-criteria approximation algorithm for the $(L_\infty, \text{top-}cn)$ -norms k -clustering problem, with running time $(\frac{k}{\epsilon})^{O(k)} \cdot \text{poly}(n)$.*

1.2 Our Techniques

Overview of FPT Time $(3 + \epsilon)$ -Approximation for Minimum-Norm Capacitated k -Clustering. In the overview, we mainly focus on the $\text{top-}\ell$ norm capacitated k -clustering problem, as it already captures the essence of our algorithm. The problem generalizes both capacitated k -center (with $\ell = 1$) and capacitated k -median (with $\ell = n$). Letting ℓ go from 1 to n gives a smooth transition from the k -center to the k -median objective. Therefore, we need to achieve $(3 + \epsilon)$ -approximation algorithms for both extreme cases while unifying their key ideas to handle the general ℓ case.

Our FPT time 3-approximation for capacitated k -center is new, and we sketch the intuition. For simplicity, we first focus on the soft-capacitated setting, where facilities may be opened multiple times. (The total number of open facilities, counting multiplicities, remains k .)

By guessing the optimum value, scaling and rounding distances, we can assume all distances are integers and the optimum value is 1. We construct a solution S with cost 1 and $O(k \log n)$ open facilities that respect the capacity constraints; this can be easily obtained via LP rounding. For each facility in S , we sample $O(\log n)$ clients connected to it.

For an optimum cluster (i^*, J) , i^* is the open facility and J is the assigned clients. If every $j \in J$ is connected to a facility i with $u_i \leq |J|$ in the solution S , then with high probability, some client in J will be sampled. This sampled client j can then serve as a *pivot* for the cluster: we open the facility with the largest capacity within distance at most 1 of j .

On the other hand, if some $j \in J$ is connected to a facility $i \in S$ with capacity $u_i > |J|$, then we can directly use i as a replacement for i^* . In either case, by guessing, we can identify a facility $i \in S$ with $u_i \geq |J|$ and $d(i, i^*) \leq 2$, which leads to a 3-approximation.

The hard-capacitated case can be addressed using the coloring idea [25], along with additional care. By guessing a coloring function, we can ensure with reasonable probability that the optimum facilities have distinct colors, which can be used to guide our selection of facilities. We have a stricter condition for the second case above: We require that $i \in S$ has capacity $u_i \geq k|J|$ so that i have enough capacity to serve k clusters in the optimum solution. We also need to take care of the case where an optimum facility also appears in S .

Our algorithm for the $\text{top-}\ell$ norm needs to capture an FPT time $(3 + \epsilon)$ -approximation for capacitated k -median. Such an algorithm was known [25]. It relies on the coresets technique, which has been extensively studied in the literature [10, 9, 39, 18, 19, 32]. However, the technique is based on sampling and only works well when $\ell = \Omega(n)$. This makes the algorithm hard to combine with the k -center objective. Instead, we use an alternative $(3 + \epsilon)$ -approximation for capacitated k -median that avoids coresets.

We obtain a $(1 + \epsilon)$ -approximation for the problem using $O(\frac{k \log n}{\epsilon})$ facilities S . We sample a few clients from each cluster in solution S . Additionally, we sample a few clients in C , with probabilities proportional to their costs in S . If for an optimum cluster (i^*, J) , the total cost of J in the solution S is large, we likely sampled a “good” client in J , which can serve as a pivot. Otherwise, the total cost of J in the solution S can be essentially ignored, and some facility with large enough capacity in S can be used to replace i . Again, extra care is needed to handle hard capacities.

To unify the approaches for both the capacitated k -center and k -median problems, we guess the ℓ -th largest connection distance t in the optimum solution. The $\text{top-}\ell$ norm of the optimum distance vector $v \in \mathbb{R}_{\geq 0}^C$ is $\ell t + \sum_{j \in C} (v_j - t)^+$. At a very high level, we treat the ℓt term as the k -center component, and the $\sum_{j \in C} (v_j - t)^+$ term as the k -median component. Our algorithm combines the techniques to handle these two parts effectively.

To extend the idea to a general monotone symmetric norm f , we apply the above idea for all distances t that is an integer power of $1 + \epsilon$. This will bound the $\text{top-}\ell$ norm cost of the solution, for every $\ell \in [0, n]$. Every symmetric norm is the maximum of many *order norms*, which in turn is a non-negative linear combination of $\text{top-}\ell$ norms [15]. Therefore, a bound on the $\text{top-}\ell$ norms leads to a bound on the f -norm. This proves Theorem 1. This technique has been used in the literature for many other problems [48, 7].

Overview of $(1 + \frac{2}{\epsilon} + \epsilon)$ -Approximation for $\text{Top-}cn$ Norm (Uncapacitated) k -Clustering Problem for $c \in (\frac{1}{e}, 1]$. To describe our algorithm, we first give an overview of the $(1 + \frac{2}{\epsilon})$ -approximation for k -median in FPT time by [24]. By using the coresets technique, one can assume that the set of clients is located in only $O(\frac{k \log n}{\epsilon})$ positions. Given a cluster (i^*, J) in the optimum clustering $\mathcal{C}$, a *leader* of the cluster is defined as the client in J that is nearest to i^* .

The algorithm proceeds by guessing the k leaders and the approximate distances of the leaders to their respective centers. With this information, a $(1 + \frac{2}{\epsilon})$ -approximation can be obtained via submodular maximization. Roughly speaking, we consider a ball centered at each leader, with the radius equal to the guessed distance to its center. Ensuring that

every ball contains one facility yields a cost of $3 \cdot \text{opt}$. Choosing the optimum facilities within these balls improves the cost by $2 \cdot \text{opt}$. The goal then becomes maximizing this improvement, which can be cast as a submodular maximization problem under a partition matroid constraint. The standard $(1 - \frac{1}{e})$ -approximation ratio for submodular maximization translates into an approximation ratio of $3 - 2(1 - \frac{1}{e}) = 1 + \frac{2}{e}$ for k -median.

To generalize this approach to the **top- cn** setting, we modify the algorithm in two ways. First, one would need to construct a coreset for the **top- cn** objective. We believe this is possible for any constant $c \in (0, 1)$; however, we instead adopt an alternative approach that is similar to our algorithm in the capacitated setting and avoids the use of coresets. We obtain a bi-criteria $(1 + \epsilon)$ -approximation for the problem by opening $O(\frac{k \log n}{\epsilon})$ facilities S . We sample a few clients R with probability proportional to their costs in S , and then guess a set of k pivots from $R \cup S$. Pivots play a role similar to leaders but can be any points in the metric space, not necessarily the clients in the respective optimum clusters. With reasonable probability, when the guesses are correct, the pivots have the desired property that leaders have.

The second modification is essential. With the **top- cn** cost, the improvement function is no longer submodular. Instead, we use an LP-based approach to construct the set of centers. A useful notion here is the *occurrence time vector* of a cost vector: in such a vector δ , δ_t represents the number of times distance t occurs in the cost vector. We naturally allow occurrence times to be fractional. We then formulate an LP relaxation that uses the pivot and radius information, with the objective of minimizing the **top- cn** cost of the occurrence time vector δ .

We design a randomized rounding algorithm such that, in expectation, the occurrence time vector of the integral solution is no worse than $(1 - \frac{1}{e})\delta + \frac{1}{e}(\delta \otimes 3)$, where $\delta \otimes 3$ denotes the vector obtained from δ by scaling all distances by a factor of 3. We bound the ratio of the **top- cn** cost of the vector to that of δ . The worst-case scenario arises when the cost vector is the all-one vector, i.e., when the occurrence time vector has distance 1 appearing n times and all other distances appearing zero times. This yields a bound of $1 + \frac{2}{ec}$. Finally, using the occurrence time representation, the **top- cn** function is concave, leading to an overall approximation ratio of $1 + \frac{2}{ec} + \epsilon$. This proves Theorem 3.

A simple modification to the algorithm can handle the L_∞ norm, and simultaneously provide a 3-approximation for it, proving Theorem 4.

Organization. In Section 2, we give some preliminaries. In Section 3, we give the $(3 + \epsilon)$ -approximation for the minimum-norm capacitated k -clustering problem in FPT time, proving Theorem 1. In Section 4, we give the tight $(\min\{1 + \frac{2}{ec}, 3\} + \epsilon)$ -approximation algorithm for the **top- cn** norm k -clustering problem, proving Theorem 3. All missing proofs can be found in the full version of the paper; in particular, we defer the proof of Theorem 4 to Section 5.

2 Preliminaries

For any $a \in \mathbb{R}$, we let $(a)^+ = \max\{a, 0\}$. For every real-valued vector y over some domain U and any $S \subseteq U$, we let $y(S) := \sum_{i \in S} y_i$, unless otherwise defined.

Given a metric space (V, d) , $v \in V$ and $S \subseteq V$, we define $d(v, S) := \min_{i \in S} d(v, i)$ to be the distance from v to the set S . For any subset $U \subseteq V$, $v \in V$ and $r \geq 0$, we define $\text{ball}_U(v, r) := \{u \in U : d(u, v) \leq r\}$ to be the set of points in U with distance at most r to v .

2.1 Norms

► **Definition 5** (Norms). A function $f : \mathbb{R}_{\geq 0}^n \rightarrow \mathbb{R}_{\geq 0}$ is a norm, if it satisfies the following 3 properties:

(5a) $f(x) = 0$ iff $x = 0$ (non-negativity).

(5b) For any $x \in \mathbb{R}_{\geq 0}^n, \lambda \in \mathbb{R}_{\geq 0}$, we have $f(\lambda x) = \lambda f(x)$ (homogeneity).

(5c) For any $x, y \in \mathbb{R}_{\geq 0}^n$, we have $f(x + y) \leq f(x) + f(y)$ (subadditivity).

For a norm f , we shall also use $|x|_f$ to denote $f(x)$, the f -norm of the vector x .

► **Definition 6** (Monotone and Symmetric Norms). A norm $f : \mathbb{R}_{\geq 0}^n \rightarrow \mathbb{R}_{\geq 0}$ is said to be monotone, if for any $x, y \in \mathbb{R}_{\geq 0}^n$ satisfying $x \leq y$, we have $f(x) \leq f(y)$. It is said to be symmetric, if for any $x \in \mathbb{R}_{\geq 0}^n$, and any permutation matrix $A \in \{0, 1\}^{n \times n}$, we have $f(x) = f(Ax)$.

Here are some common monotone symmetric norms studied in the literature:

- L_p norms, $p \geq 1$: $L_p(x) := \left(\sum_{j=1}^n x_j^p\right)^{1/p}$, $\forall x \in \mathbb{R}_{\geq 0}^n$. We also use $|x|_p$ for $L_p(x)$.
- L_∞ norm: $L_\infty(x) := \max_{j \in [n]} x_j$, $\forall x \in \mathbb{R}_{\geq 0}^n$. We also use $|x|_\infty$ for $L_\infty(x)$.
- Top- ℓ norm, $\ell \in [0, n]$ is a real: $|x|_{\text{top-}\ell}$ for any $x \in \mathbb{R}_{\geq 0}^n$ is the maximum of $\alpha^\top x$ over all $\alpha \in [0, 1]^n$ satisfying $|\alpha|_1 = \ell$. When ℓ is an integer, $|x|_{\text{top-}\ell}$ is the sum of the ℓ largest coordinates in x . In this paper, we shall allow ℓ to take fractional values. Notice that $\text{top-}1 \equiv L_\infty$ and $\text{top-}n \equiv L_1$.
- Ordered norms: This is a non-negative linear combination of top- ℓ norms, for different values of ℓ . See the following definition.

► **Definition 7** (Ordered norms). Given a non-negative and non-increasing weight vector $w = (w_1, \dots, w_n)$. The w -ordered norm of any vector $x \in \mathbb{R}_{\geq 0}^n$, denoted as $|x|_{w\text{-ordered}}$, is defined as follows:

$$|x|_{w\text{-ordered}} := \sum_{i=1}^n w_i x_i^\downarrow,$$

where $x_i^\downarrow$ is the i -th largest value in $\{x_1, x_2, \dots, x_n\}$ (counting multiplicities).

The following folklore lemma is useful when analyzing top- ℓ norms:

► **Lemma 8.** For any $x \in \mathbb{R}_{\geq 0}^n$, we have $|x|_{\text{top-}\ell} = \min_{t \geq 0} \left(\sum_{j=1}^n (x_j - t)^+ + \ell t \right)$. Moreover, the minimum is achieved when t is the ℓ -th largest coordinate in x .

The following observation is easy to see:

► **Observation 9.** Given a non-negative and non-increasing weight vector $w = (w_1, \dots, w_n)$. Let $w_{n+1} = 0$. Then for any $x \in \mathbb{R}_{\geq 0}^n$, we have $|x|_{w\text{-ordered}} = \sum_{\ell=1}^n (w_\ell - w_{\ell+1}) \cdot |x|_{\text{top-}\ell}$.

Any monotone symmetric norm can be expressed as the maximum of ordered norms [15]:

► **Lemma 10** ([15]). Let $f : \mathbb{R}_{\geq 0}^n \rightarrow \mathbb{R}_{\geq 0}$ be a monotone symmetric norm. There exists a closed set $\mathcal{W}$ of non-increasing vectors in $\mathbb{R}_{\geq 0}^n$, such that for any $x \in \mathbb{R}_{\geq 0}^n$, we have

$$|x|_f = \max_{w \in \mathcal{W}} |x|_{w\text{-ordered}} = \max_{w \in \mathcal{W}} \sum_{\ell=1}^n (w_\ell - w_{\ell+1}) |x|_{\text{top-}\ell}.$$

2.2 Definitions of Problems

► **Definition 11** (Minimum-Norm k -Clustering). *Let $f : \mathbb{R}_{\geq 0}^n \rightarrow \mathbb{R}_{\geq 0}$ be a monotone symmetric norm. In the Minimum-Norm k -Clustering problem under f -norm, we are given a set F of facilities, a set C of n clients, a positive integer $k \leq |F|$, and a metric space $(F \uplus C, d)$ over $F \uplus C$. The goal of the problem is to find a set $S \subseteq F$ with $|S| = k$ so as to minimize $f((d(j, S))_{j \in C})$.*

When f are the L_∞, L_1 and L_2 norms, the problem becomes k -supplier-center, k -median and metric k -means problems respectively.

► **Definition 12** (Minimum-Norm Capacitated k -Clustering). *Let $f : \mathbb{R}_{\geq 0}^n \rightarrow \mathbb{R}_{\geq 0}$ be a monotone symmetric norm. In the Minimum-Norm Capacitated k -Clustering problem under f -norm, we are given F, C, n, k and d as in Definition 11. Additionally, we are given a capacity $u_i \in \mathbb{Z}_{\geq 0}$ for every $i \in F$. The goal of the problem is to find a set $S \subseteq F$ with $|S| = k$, and an assignment vector $\sigma \in S^C$ such that $|\sigma^{-1}(i)| \leq u_i$ for every $i \in S$, so as to minimize $f((d(j, \sigma_j))_{j \in C})$.*

When f are the L_∞ and L_1 norms, the problem becomes capacitated k -center³ and capacitated k -median respectively.

Throughout the paper, we assume we are given an efficient oracle for the norm f . We can assume distances are either ∞ or integers in $[0, \text{poly}(n)]$, by losing a factor of $1 + \epsilon$ in the approximation ratio. See the full version of the paper. As is typical, we do sampling and make guesses during our algorithms. When analyzing the algorithm, we define the conditions for the sampling steps being successful, and the correct answers for all the guesses. In the actual algorithm, we have to enumerate all possibilities for the guesses. It suffices for our algorithm to succeed with $\frac{1}{g(k, \epsilon) \text{poly}(n)}$ probability, as we can repeat it many times and output the best solution generated, to increase the probability to $1 - 1/\text{poly}(n)$. For the guessing and repetition to be feasible, it is required that we can compute the cost of the constructed solutions. For most natural norms, the optimal assignment σ , and thus the minimum cost, can be computed efficiently given the set S of open facilities. For general norms f , the cost can be approximated within a factor of $1 + \epsilon$.

2.3 $(1 + \epsilon)$ -Approximation for Capacitated k -Clustering with $O(\frac{k \log n}{\epsilon})$ Facilities

It is useful to run an LP rounding algorithm to obtain a $(1 + \epsilon)$ -approximate solution for capacitated k -clustering with $O(\frac{k \log n}{\epsilon})$ facilities S . With FPT time allowed, we can afford to guess $\text{poly}(k, \frac{1}{\epsilon})$ facilities in S .

We are given F, C, n, d, k and capacities $(u_i)_{i \in F}$ as in a capacitated k -clustering problem. Instead of a monotone symmetric norm f , we are given a convex monotone function $h : \mathbb{R}_{\geq 0}^C \rightarrow \mathbb{R}_{\geq 0}$ (which is not necessarily symmetric or a norm). The goal is the same as that of the minimum-norm capacitated k -clustering problem, except that we are minimizing the h function of the connection distance vector. Let opt' be the value of the instance.

For some $i \in F$ and $J \subseteq C$, we say (i, J) is valid star if $|J| \leq u_i$. The following theorem summarizes the $(1 + \epsilon)$ -approximation.

³ In the literature, capacitated k -center refers to the problem where F and C can be unrelated, instead of the problem where $F = C$.

► **Theorem 13.** *Let $\epsilon > 0$ be a constant. We can efficiently find a set $\mathcal{S}$ of $O\left(\frac{k \ln n}{\epsilon}\right)$ valid stars such that*

- $\bigcup_{(i,J) \in \mathcal{S}} J = C$.
 - *Let $b \in \mathbb{R}_{\geq 0}^C$ be the connection distance vector induced by $\mathcal{S}$: For every $j \in C$, we have $b_j = d(i, j)$ for the unique star (i, J) with $j \in J$. Then $h((1 - \epsilon)b) \leq \text{opt}'$.*
- The algorithm succeeds with high probability.*

The following corollary is an immediate consequence of our analysis, and we were unable to find an explicit statement of the result in the literature. As we discussed, it was implied by [35] when all facilities have uniform capacities. However, extending their result to the general capacitated setting seems non-trivial. For the **top- c** norm k -clustering problem, we need to use the corollary to handle the case $c \leq \frac{1}{\epsilon}$.

► **Corollary 14.** *Consider the minimum-norm k -clustering problem under a monotone norm f (which is not necessarily symmetric). There is a $(3 + \epsilon)$ -approximation algorithm for the problem in FPT time with parameters k and ϵ .*

The proofs of Theorem 13 and Corollary 14 are deferred to the full version of the paper.

2.4 Finding $(1 + \epsilon)$ -Approximate Assignment for Minimum-Norm Capacitated k -Clustering with Open Facilities

Consider an instance of the Minimum-Norm Capacitated k -Clustering problem. For many natural norms such as L_p and **top- ℓ** norms, when given the set S of k open facilities, the optimum assignment of C to S can be found easily. However, this is not trivial for general monotone symmetric norms f . Instead, we present an FPT-time algorithm to find a $(1 + \epsilon)$ -approximate assignment. The proof of the following theorem is described in the full version of the paper.

► **Theorem 15.** *Consider a minimum-norm capacitated k -clustering instance defined by F, C, n, d, u and k , under a symmetric monotone norm f . Assume $|F| = k$ and $\epsilon > 0$ is a constant. We can find a $(1 + \epsilon)$ -approximate assignment $\sigma \in F^C$ for the instance in time $g(k, \epsilon) \cdot \text{poly}(n)$, for a computable function g depending on k and ϵ .*

3 FPT Time $(3 + \epsilon)$ -Approximation for Minimum-Norm Capacitated k -Clustering

In this section, we give the FPT time $(3 + \epsilon)$ -approximation algorithm for the minimum-norm capacitated k -clustering problem, proving Theorem 1. Recall that we say $(i \in F, J \subseteq C)$ is a valid star if $|J| \leq u_i$. We fix an optimum solution $\mathcal{S}^*$ of k valid stars that is unknown to the algorithm. $\mathcal{S}^*$ cover all clients in C , and all facilities are distinct. Let S^* be the set of facilities in $\mathcal{S}^*$. So, we have $|S^*| = |\mathcal{S}^*| = k$. Let opt be the f -norm cost of $\mathcal{S}^*$.

The algorithm is stated in Algorithm 1. In Step 1, we define a set $\mathbf{T}$ of distances depending on the problem, as follows:

- If the norm to minimize is the **top- ℓ** norm, then let $\mathbf{T} = \{t\}$, with t being the ℓ -th largest connection distance in the optimum solution $\mathcal{S}^*$. There are $n|F|$ possibilities for t , so it is affordable to guess its correct value.
- If the norm is a general monotone symmetric norm, then $\mathbf{T} = \left\{ (1 + \epsilon)^{\lceil \log_{1+\epsilon} d(i,j) \rceil} : i \in F, j \in C \right\}$. Notice that $|\mathbf{T}| = O\left(\frac{\log n}{\epsilon}\right)$ as we assumed that distances that are not ∞ are integers bounded by $\text{poly}(n)$.

■ **Algorithm 1** $(3 + \epsilon)$ -Approximation for Minimum-Norm Capacitated k -Clustering.

-
- 1: define set $\mathbf{T}$ of non-negative reals with $|\mathbf{T}| \leq O\left(\frac{\log n}{\epsilon}\right)$, as described in text
 - 2: **for** every $t \in \mathbf{T}$ **do**
 - 3: with objective $h(v) := \sum_{j \in C} (v_j - t)^+$ for every $v \in \mathbb{R}_{\geq 0}^C$, use Theorem 13 to obtain a set $\mathcal{S}^t$ of $O\left(\frac{k \ln n}{\epsilon}\right)$ valid stars
 - 4: $S^t \leftarrow \{i \in F : \exists J, (i, J) \in \mathcal{S}^t\}$
 - 5: randomly choose a color function $\text{color} : F \rightarrow [k]$
 - 6: guess the types of each color $c \in [k]$ ▷ each color is of type-1, 2 or 3
 - 7: **for** every $t \in \mathbf{T}$ **do** $R^t \leftarrow \text{MNCKC-choose-R}(t)$ ▷ clients in R^t are called *representatives*
 - 8: $S \leftarrow \bigcup_{t \in \mathbf{T}} S^t, R \leftarrow \bigcup_{t \in \mathbf{T}} R^t$
 - 9: guess a pivot p_c for each type-1 or 2 color $c \in [k]$, and i_c^* for each type-3 color $c \in [k]$ such that
 - ▷ p_c 's for type-1 and 2 colors c , i_c^* 's for all colors $c \in [k]$ are defined in text
 - $p_c \in R$ if c is of type-1, $p_c \in S$ if c is of type-2, and
 - $i_c^* \in S$ has color c for a type-3 color c .
 - 10: guess a $(1 + \epsilon)$ -approximate overestimation r_c for $d(i_c^*, p_c)$ for every type-1 or 2 color c
 - 11: **return** $\text{MNCKC-clustering-with-pivots}()$ ▷ See Algorithm 3 for its definition
-

Readers seeking for a more efficient understanding of the core ideas can focus on the **top- ℓ** norm case, where we have only one t in $\mathbf{T}$.

For every $t \in \mathbf{T}$, in Step 3, we obtain a set $\mathcal{S}^t$ of $O\left(\frac{k \ln n}{\epsilon}\right)$ valid stars covering C using Theorem 13, and define S^t to be the set of facilities used in $\mathcal{S}^t$ in Step 4. We let i_j^t be the center of the star in $\mathcal{S}^t$ containing j for every $j \in C$, and let $b_j^t = d(j, i_j^t)$ be the connection distance of j in $\mathcal{S}^t$. The conditions of the Theorem 13 hold with high probability. We assume they are satisfied:

(P) For every $t \in \mathbf{T}$, we have $\sum_{j \in C} ((1 - \epsilon)b_j^t - t)^+ \leq \sum_{(i^*, J) \in \mathcal{S}^*, j \in J} (d(i^*, j) - t)^+$.

We describe the remaining steps in more detail.

3.1 Steps 5 and 6 of Algorithm 1: Guessing Colors and Types

In Step 5 of Algorithm 1, we randomly choose a function $\text{color} : F \rightarrow [k]$. With probability $\frac{k!}{k^k}$, the k facilities in S^* have distinct colors. We assume this happens. For every color $c \in [k]$, let (i_c^*, J_c^*) be the star in $\mathcal{S}^*$ such that i_c^* is of color c . We define $\bar{J}_c^*$ to be the $\lceil \epsilon |J_c^*| \rceil$ clients in J_c^* closest to i_c^* . Let $\tilde{J}_{t,c}^*$ be the $\left\lceil \frac{|\bar{J}_c^*|}{2} \right\rceil$ clients in $\bar{J}_c^*$ with smallest b_j^t values for every $t \in \mathbf{T}$. Notice that the i_c^* 's, J_c^* 's, $\bar{J}_c^*$'s and $\tilde{J}_{t,c}^*$'s are not known to our algorithm after the steps.

We define a *type* for each color $c \in [k]$ as follows:

- If $i_c^* \notin S$, then
 - if $\exists t \in \mathbf{T}, \forall j \in \tilde{J}_{t,c}^*, u_{i_j^t} < k |J_c^*|$, then c is of *type-1a*, otherwise,
 - if $\exists t \in \mathbf{T}, \sum_{j \in \bar{J}_c^*} ((1 - \epsilon)b_j^t - t)^+ \geq \frac{\epsilon^2}{k} \cdot \sum_{j \in C} ((1 - \epsilon)b_j^t - t)^+$, then c is of *type-1b*, and otherwise,
 - c is of *type-2*.
- If $i_c^* \in S$, then c is of *type-3*.

Notice that for the type-1a colors, we consider the set $\tilde{J}_{t,c}^*$, but for type-1b colors, we consider the set $\bar{J}_c^*$. We say c is of *type-1* if it is of type-1a or type-1b. Notice the types 1, 2 and 3 partition the set $[k]$ of colors. We guess the types in Step 6 of Algorithm 1. Again, we assume our guesses are correct.

3.2 Step 7 of Algorithm 1: Constructing Representatives R^t

In Step 7 of Algorithm 1, for every $t \in \mathbf{T}$, we construct a *representative set* R^t through the process `MNCKC-choose-R( $t$ )`, described in Algorithm 2.

■ **Algorithm 2** `MNCKC-choose-R( $t$ )`.

-
- 1: for every $(i, J) \in \mathcal{S}^t$, uniformly sample $\min(|J|, \lceil \frac{2k}{\epsilon} \ln(kn) \rceil)$ clients from J , without replacement; they constitute set R_1^t
 - 2: sample $O\left(\frac{k \ln n}{\epsilon^2}\right)$ clients from C , for a sufficiently large hidden constant, each client j being chosen from C with probabilities proportional to $((1 - \epsilon)b_j^t - t)^+$, with replacement; they constitute set R_2^t
 - 3: **return** $R^t = R_1^t \cup R_2^t$
-

As $|\mathcal{S}^t| = O\left(\frac{k \ln n}{\epsilon^2}\right)$, we have $|R^t| = O\left(\frac{k^2 \ln^2 n}{\epsilon^2}\right)$ for every $t \in \mathbf{T}$.

In Step 8, we merge all sets $\mathcal{S}^t$ into S and R^t into R . Therefore, $|S| = O\left(\frac{k \log^2 n}{\epsilon^2}\right)$ and $|R| = O\left(\frac{k^2 \log^3 n}{\epsilon^3}\right)$.

► **Lemma 16.** *With probability at least $1 - \frac{1}{n}$, the following event happens: for every color $c \in [k]$ of type-1, we have $\bar{J}_c^* \cap R \neq \emptyset$.*

The proof of Lemma 16 is deferred to the full version of the paper.

3.3 Step 9 and 10 of Algorithm 1: Guessing Pivots, Optimum Facilities and Radius

In Step 9 of Algorithm 1, we guess the *pivot* p_c for type-1 or 2 colors c , which is defined as follows.

- If c is of type-1, then p_c is defined as any client in $\bar{J}_c^* \cap R$. $p_c \in R \subseteq C$ in this case.
- If c is of type-2, then p_c is defined as the i_j^t for some $t \in \mathbf{T}$ and $j \in \bar{J}_{t,c}^*$ satisfying $u_{i_j^t} \geq k|J_c^*|$ with the smallest b_j^t value. Then $p_c \in S \subseteq F$. Such a (t, j) pair exists, since otherwise c would be of type-1a.

We also guess i_c^* for each type-3 color c . By definition, we have $i_c^* \in S$ and thus we can afford this.

In Step 10, we guess $d(i_c^*, p_c)$ for every type-1 or type-2 color c . We can afford to guess a $(1 + \epsilon)$ -approximation of $d(i_c^*, p_c)$: r_c is the smallest $(1 + \epsilon)^z$ that is at least $d(i_c^*, p_c)$ for an integer z .

So, after the guessing, we know the p_c 's for type-1 and 2 colors c , and i_c^* 's for type-3 colors c . However, we do not know i_c^* 's for type-1 and 2 colors c .

3.4 Step 11 of Algorithm 1: Finding Clusters Using Pivots and Radius

With all the guessed information, we can construct a feasible clustering. This is done in Step 11 of Algorithm 1, which calls `MNCKC-clustering-with-pivots()` as described in Algorithm 3.

Notice that for a type-2 color c , p_c is not always of color c , so Step 8 in Algorithm 3 is not redundant. Further, we treat T as a set, instead of a multi-set, so $|T| \leq k$ and it may happen that $|T| < k$.

Algorithm 3 MNCKC-clustering-with-pivots().

```

1:  $T \leftarrow \emptyset$ 
2: for every color  $c \in [k]$  of type-1 do
3:    $q_c \leftarrow \arg \max \{u_i : i \in \text{ball}_{F \setminus S}(p_c, r_c) \cap \text{color}^{-1}(c)\}; T \leftarrow T \cup \{q_c\}$ 
4: for every color  $c \in [k]$  of type-3 do
5:    $T \leftarrow T \cup \{i_c^*\}$ 
6: for every color  $c \in [k]$  of type-2 do
7:   if  $p_c \neq i_{c'}^*$  for every type-3 color  $c'$  then  $T \leftarrow T \cup \{p_c\}$ 
8:   else  $g_c \leftarrow \arg \max \{u_i : i \in \text{ball}_{F \setminus S}(p_c, r_c) \cap \text{color}^{-1}(c)\}; T \leftarrow T \cup \{g_c\}$ 
9: return  $T$  and the feasible assignment  $\sigma \in T^C$  respecting the capacity constraints, either
    directly (if  $f$  is the top- $\ell$  norm) or using Theorem 15 (if  $f$  is general)
  
```

3.5 Analysis of Cost

To analyze the cost of the solution, we explicitly construct an assignment for T with small cost. We build a bipartite graph $H = ([k], T, E_H)$ as follows; it is instructive to correlate the construction with Algorithm 3.

- For every c of type-1, we add (c, q_c) to E_H .
- For every c' of type-3 (which implies $i_{c'}^* \in T$), we add $(c', i_{c'}^*)$ to E_H .
- For every c of type-2, we add (c, p_c) to E_H . If additionally $p_c = i_{c'}^*$ for some type-3 color c' , then we add (c', g_c) to E_H .


We construct a solution in the following way: Initially, all clients are moved to their corresponding facilities in S^* , with moving cost precisely opt. After that, i_c^* contains $|J_c^*| \leq u_{i_c^*}$ clients for any color c . Then, we move the clients from S^* to T according to a “transportation” function $\phi \in \mathbb{Z}_{\geq 0}^{E_H}$: $\phi_{c,i}$ clients will be moved from i_c^* to i . In order for ϕ to be a solution, ϕ must satisfy the following properties:

- $\forall c \in [k], \sum_{(c,i) \in E_H} \phi_{c,i} = |J_c^*|$;
- $\forall i \in T, \sum_{(c,i) \in E_H} \phi_{c,i} \leq u_i$.

In the following, we construct ϕ . Connected components in H are only of the following three possibilities, and we construct ϕ for each of them (See Figure 1 for an illustration):

- (1) Some type-1 color c connected to $q_c \in T$. Then q_c has degree 1 and $\{c, q_c\}$ is maximally connected. We set $\phi_{c,q_c} = |J_c^*|$. We have $u_{q_c} \geq u_{i_c^*} \geq |J_c^*|$.
- (2) Several type-2 colors and no type-3 color connect to facility $i \in T$. Let $D := \{c : p_c = i\}$, then $D \cup \{i\}$ is maximally connected. Set $\phi_{c,i} = |J_c^*|$ for every $c \in D$. We have $\sum_{c \in D} |J_c^*| \leq \sum_{c \in D} \frac{u_i}{k} \leq u_i$ as all colors in D are of type-2.
- (3) Facility $i = i_{c'}^* \in T \cap S^*$ is connected to one type-3 color c' . Some type-2 colors may also connect to i . Let $D := \{c : p_c = i, c \text{ is of type-2}\}$; D may be empty. $\{c', i\} \cup D \cup \{g_c \mid c \in D\}$ is one connected component in H . Similar to (ii), $\sum_{c \in D} |J_c^*| \leq u_i$. We define $\phi_{c,i} = |J_c^*|$ for every $c \in D$. We then define $\phi_{c',i} = \min\{|J_{c'}^*|, u_i - \sum_{c \in D} |J_c^*|\}$. If $u_i - \sum_{c \in D} |J_c^*| < |J_{c'}^*|$, then we define ϕ_{c',g_c} values for all $c \in D$ so that $\phi_{c',g_c} \leq |J_c^*|$ for every $c \in D$, and $\sum_{c \in D} \phi_{c',g_c} = |J_{c'}^*| - \phi_{c',i}$. This is feasible since $\sum_{c \in D} |J_c^*| = u_i - \phi_{c',i} \geq |J_{c'}^*| - \phi_{c',i}$. As for $c \in D$, i_c^* is a candidate for g_c , $u_{g_c} \geq u_{i_c^*} \geq |J_c^*|$, so ϕ does not violate the capacity constraint for g_c .

In summary, the two properties needed for the transportation function ϕ are satisfied, which implies the assignment of C to T respects all capacity constraints.



■ **Figure 1** 3 types of connected components of H . In the first type, a type-1 color c is connected to q_c . Notice that q_c is defined via a representative client $p_c \in R \cap \bar{J}_c^*$. In the second type, we have many type-2 colors $c \in D$ connected to their common p_c . p_c for a type-2 color c is defined via a client j . In the third type, we have many type-2 colors $c \in D$ (D' in the figure to avoid confusion) connected to their common p_c , which is $i_{c'}^*$ for a type-3 color c' . Then we have edges $(c', i_{c'}^*)$ and (c', g_c) for every $c \in D$ in the component. Solid squares in T are in S , and empty squares in T are not in S .

We then bound the f -norm cost of moving clients from S^* to T using the transportation function ϕ . For each client j , let β_j be the moving distance of this step, and b_j^* be the connection distance of j in the optimum solution $\mathcal{S}^*$. The main lemma we prove is the following:

► **Lemma 17.** *For every $t \in \mathbf{T}$, we have $\sum_{j \in C} (\beta_j - (2 + O(\epsilon))t)^+ \leq (2 + O(\epsilon)) \sum_{j \in C} (b_j^* - t)^+$.*

The proofs of Lemma 17 are deferred to the full version of the paper.

Recall that in our constructed solution, we moved clients from C to T in two steps. We first moved C to S^* according to $\mathcal{S}^*$, which incur a cost of opt . Then we moved clients from S^* to T , with moving distance vector β .

Now we describe the $f \equiv \text{top-}\ell$ and general f case separately. For the case $f \equiv \text{top-}\ell$, Lemma 17 holds for the ℓ -th largest coordinate of b^* . By Lemma 8, the moving cost of the second step is

$$\begin{aligned} |\beta|_{\text{top-}\ell} &\leq \ell \cdot (2 + O(\epsilon))t + \sum_{j \in C} (\beta_j - (2 + O(\epsilon))t)^+ \\ &\leq (2 + O(\epsilon))\ell t + (2 + O(\epsilon)) \sum_{j \in C} (b_j^* - t)^+ \\ &\leq (2 + O(\epsilon)) \left(\ell t + \sum_{j \in C} (b_j^* - t)^+ \right) = (2 + O(\epsilon)) \cdot |b^*|_{\text{top-}\ell} = (2 + O(\epsilon))\text{opt}. \end{aligned}$$

This implies that $\text{top-}\ell$ cost of our constructed solution is at most $(3 + O(\epsilon)) \cdot \text{opt}$.

Now consider a general monotone symmetric norm f . We show that for every integer $\ell \in [n]$, we have $|\beta|_{\text{top-}\ell} \leq (2 + O(\epsilon))|b^*|_{\text{top-}\ell}$. By Lemma 10, this implies $|\beta|_f \leq (2 + O(\epsilon))|b^*|_f = (2 + O(\epsilon))\text{opt}$. Focus on any $\ell \in [n]$ and let t be the ℓ -th largest coordinate in b^* . Then, we have $t' = (1 + \epsilon)^{\lceil \log_{1+\epsilon} t \rceil} \in \mathbf{T}$. Applying Lemma 17 with t being this t' will prove $|\beta|_{\text{top-}\ell} \leq (2 + O(\epsilon))|b^*|_{\text{top-}\ell}$. Therefore, the f -norm cost our constructed solution is at most $(3 + O(\epsilon)) \cdot \text{opt}$. In both cases, scaling ϵ at the beginning gives us a $(3 + \epsilon)$ -approximation.

3.6 Analysis of Runtime

Now we analyze the running time of our algorithm. When all guesses are correct, the coloring satisfies our requirement, and the events mentioned in Lemma 16 happen, the algorithm returns a solution with value at most $(3 + O(\epsilon))\text{opt}$.

For the runtime, all steps except coloring and guessing are polynomial. With probability $\frac{k!}{k^k}$, the color function in Step 5 satisfies our requirement, so we need to run the algorithm $O(k^k \ln n)$ times to boost the success probability to $1 - \frac{1}{n}$. For the guessing part, Step 6 has 3^k possible choices for the types, Step 9 has $(|R| + |S|)^k = (\frac{k \ln n}{\epsilon})^{O(k)}$ choices for pivots, and finally, as we can always assume $d(c, f) = \text{poly}(n)$ for all $c \in C$ and $f \in F$, Step 10 has $(\log_{1+\epsilon} \text{poly}(n))^k = (\frac{\ln n}{\epsilon})^{O(k)}$ choices for the radius. In total, the running time is $(\frac{k \ln n}{\epsilon})^{O(k)} n^{O(1)}$, which can be bounded by $(\frac{k}{\epsilon})^{O(k)} n^{O(1)}$.⁴

4 FPT Time $(1 + \frac{2}{\epsilon c} + \epsilon)$ -Approximation for Top- cn Norm k -Clustering

In this section, we give the tight $(1 + \frac{2}{\epsilon c} + \epsilon)$ -approximation algorithm for the **top- cn** norm k -clustering problem for $c \in (\frac{1}{e}, 1]$, proving Theorem 3. Notice that the case $c \leq \frac{1}{e}$ has an approximation ratio of $3 + \epsilon$, as stated in Corollary 14.

4.1 Useful Tools for Top- ℓ Norms

In this section, we describe some useful tools. We assume the distances in the metric d are non-negative integers. It will be convenient to use the *occurrence-times vectors* to represent distance vectors: for each $a \in \mathbb{Z}_{\geq 0}$, we have a coordinate indicating the number of times a appear in the distance vector; that is, how many clients have distance a to their nearest facility. With this in mind, we define the set of occurrence-times vectors to be

$$\mathcal{S}_n := \{\delta \in \mathbb{R}_{\geq 0}^{\mathbb{Z}_{\geq 0}} : |\delta|_1 = n \text{ and } \delta \text{ has a finite support}\}.$$

For a $\delta \in \mathcal{S}_n$ and $a \in \mathbb{Z}_{\geq 0}$, we can think of δ_a as the number of times a appears in the distance vector. It is convenient to allow fractional occurrence times.

For any vector $\phi \in \mathbb{R}_{\geq 0}^{\mathbb{Z}_{\geq 0}}$ with finite support, we define $\bar{L}_1(\phi) = \sum_{a=0}^{\infty} \phi_a \cdot a$. (Throughout, we shall typically use $\bar{f}$ to denote a norm function f using the occurrence-times vector representation: for an occurrence-time vector $\delta \in \mathcal{S}_n$, $\bar{f}(\delta)$ is the value of f on the corresponding multiset of distances.)

► **Definition 18.** For every $\delta \in \mathcal{S}_n$ and real $\ell \in [0, n]$, we define $\overline{\text{top-}\ell}(\delta)$ to be the value of the following linear program with variables $\alpha \in \mathbb{R}_{\geq 0}^{\mathbb{Z}_{\geq 0}}$: maximize $\bar{L}_1(\alpha)$ subject to $0 \leq \alpha \leq \delta$ and $|\alpha|_1 = \ell$.

To get some intuition about the definition, consider the case where d is a distance vector of dimension n and δ is its correspondent occurrence-times vector. Then $\overline{\text{top-}\ell}(\delta) = |d|_{\text{top-}\ell}$. We extended the definition to real vectors δ .

► **Lemma 19.** For every real $\ell \in [0, n]$, we have that $\overline{\text{top-}\ell}(\cdot)$ is concave on $\mathcal{S}_n$.

This is in contrast to the **top- ℓ** norm function using the normal representation of vectors, which is convex.

⁴ We need to bound $(\ln n)^{O(k)}$. If $k \leq \sqrt{\log n}$, then this is upper bounded by $n^{O(1)}$. Otherwise, it is upper bounded by $k^{O(k)}$.

► **Definition 20.** We say a vector $\delta \in \mathcal{S}_n$ dominates a vector $\delta' \in \mathcal{S}_n$ with a factor of γ for some real $\gamma \geq 0$, denoted as $\delta' \preceq_\gamma \delta$, if $\overline{\text{top-}\ell}(\delta') \leq \gamma \cdot \overline{\text{top-}\ell}(\delta)$ for every real $\ell \in [0, n]$. We use $\preceq$ for $\preceq_1$.

► **Lemma 21.** Assume $\delta^1, \delta^2, \dots, \delta^H, \delta'^1, \delta'^2, \dots, \delta'^H$ are $2H$ vectors in $\mathcal{S}_n$, $\gamma \in \mathbb{R}_{>0}$, and $\delta'^h \preceq_\gamma \delta^h$ for every $h \in [H]$. Let $\beta_1, \beta_2, \dots, \beta_H \in [0, 1]$ satisfy $\sum_{h=1}^H \beta_h = 1$. Then

$$\sum_{h=1}^H \beta_h \delta'^h \preceq_\gamma \sum_{h=1}^H \beta_h \delta^h.$$

► **Corollary 22.** Let $\delta \in \mathcal{S}_n$, and $\delta = \sum_{h=0}^H \theta^h$ where $H \in \mathbb{Z}_{\geq 0}$ and $\theta^0, \theta^1, \dots, \theta^H \geq 0$. Let $\gamma \geq 1$. For every $h \in [H]$, let $b_h \in \mathbb{Z}_{\geq 0}$ satisfy $b_h \leq \gamma \cdot \frac{L_1(\theta^h)}{|\theta^h|_1}$. Let $\delta' = \theta^0 + \sum_{h=1}^H |\theta^h|_1 \cdot \mathbf{e}_{b_h}$, where $\mathbf{e}_b$ for any $b \in \mathbb{Z}_{\geq 0}$ is the vector in $\mathbb{R}_{\geq 0}^{\mathbb{Z}_{\geq 0}}$ with $(\mathbf{e}_b)_b = 1$ and $(\mathbf{e}_b)_{b'} = 0$ if $b' \neq b$. Then $\delta' \preceq_\gamma \delta$.

We treat δ as n fractional values. Each θ^h contains a disjoint portion of δ . For every $h \in H$, we replace θ^h with $|\theta^h|_1$ fractional values equaling to b_h . The lemma says that the resulting occurrence-times vector is dominated by the original one with a factor of γ .

► **Definition 23.** For any two vectors $\delta, \delta' \in \mathcal{S}_n$, we say some $\phi \in \mathcal{S}_n$ can be obtained by adding δ and δ' if there exists some $z \in \mathbb{R}_{\geq 0}^{\mathbb{Z}_{\geq 0} \times \mathbb{Z}_{\geq 0}}$ such that

- $\sum_{a'} z_{a,a'} = \delta_a$ for every $a \in \mathbb{Z}_{\geq 0}$,
- $\sum_a z_{a,a'} = \delta'_{a'}$ for every $a' \in \mathbb{Z}_{\geq 0}$, and
- $\phi_t = \sum_{a+a'=t} z_{a,a'}$ for every $t \in \mathbb{Z}_{\geq 0}$.

z gives a matching between the vectors δ and δ' . Then the occurrence-time vector ϕ is obtained by adding the distance vectors for δ and δ' , using the matching z .

► **Definition 24.** Let $\delta, \delta' \in \mathcal{S}_n$, and $z \in \mathbb{R}_{\geq 0}^{\mathbb{Z}_{\geq 0} \times \mathbb{Z}_{\geq 0}}$ be the unique vector satisfying the first two conditions in Definition 23 and the following condition:

- There are no integers $a < a', b > b'$ such that $z_{a,b} > 0$ and $z_{a',b'} > 0$.

Then, we use $\delta \oplus \delta'$ to denote the vector $\phi \in \mathcal{S}_n$ satisfying the third condition of Definition 23. We use $\delta \otimes 2$ to denote $\delta \oplus \delta$.

► **Lemma 25.** Let $\delta, \delta' \in \mathcal{S}_n$, $\phi = \delta \oplus \delta'$, $\phi' \in \mathcal{S}_n$ be obtained by adding δ and δ' . Then $\phi' \preceq \phi$.

► **Theorem 26.** Let $\delta, \delta' \in \mathcal{S}_n$ such that $\delta' \preceq_\gamma \delta$ for some $\gamma > 0$. Let $c \in (\frac{1}{e}, 1]$ and $\alpha \in [0, c]$. Then we have

$$\overline{\text{top-cn}}((1 - \alpha)\delta + \alpha \cdot (\delta \oplus \delta' \otimes 2)) \leq \left(1 + \frac{2\alpha\gamma}{c}\right) \overline{\text{top-cn}}(\delta).$$

We only apply the theorem for $\alpha = \frac{1}{e}$ and $\gamma = 1 + O(\epsilon)$.

The proofs of Lemma 19, Lemma 21, Corollary 22, and Theorem 26 are deferred to the full version of the paper.

4.2 The Algorithm

Throughout this and the next section, given any set $S \subseteq F$ (it is possible that $|S| \neq k$), we shall use $\vec{d}(S)$ to denote the vector $(d(j, S))_{j \in C}$ of n distances; we call it the distance vector for S . So the goal of the **top-cn** norm k -clustering problem is to minimize $|\vec{d}(S)|_{\text{top-cn}}$ subject to $S \subseteq F, |S| = k$.

Let S^* be the unknown set of k facilities in the optimum solution. We let $\text{opt} = |\vec{d}(S^*)|_{\text{top-cn}}$. For every $i^* \in S^*$, and the set $J \subseteq C$ of clients connected to i^* in the optimum solution, we refer to (i^*, J) as an *optimum cluster*.

► **Definition 27** (Per-Client Costs and Cores). *For every optimum cluster (i^*, J) , we define its per-client cost to be $\frac{1}{|J|} \sum_{j \in J} d(i^*, j)$. We define its core to be the set of the $\lceil \epsilon |J| \rceil$ clients in J with the smallest $d(i^*, j)$ values.*

■ **Algorithm 4** Top-cn Norm k -Clustering.

-
- 1: obtain a set $\mathcal{S}$ of valid stars using Theorem 13 with $h \equiv \text{top-cn}$ and $u_i = \infty, \forall i \in F$
 - 2: $S \leftarrow \{i \in F : \exists J, (i, J) \in \mathcal{S}\}$
 - 3: $R \leftarrow \text{TpcnC-choose-R}()$ ▷ Clients in R are called *representatives*
 - 4: guess a multi-set $P \subseteq S \cup R$ of size k ▷ Points in P are called *pivots*
 - 5: guess a vector $(r_p)_{p \in P}$ where each r_p is 0 or an integer power of $(1 + \epsilon)$ in $[1, \Delta]$ ⁵
 - 6: **return** $\text{TpcnC-clustering-with-pivots}()$
-

The pseudo-code for the algorithm is given in Algorithm 4. In Step 2, we apply Theorem 13 to obtain a solution $\mathcal{S}$ and let S be the set of open facilities. As there are no capacities, we simply use S to denote the solution and discard the notion $\mathcal{S}$. We have $|S| \leq \frac{O(k \log n)}{\epsilon}$. Moreover, $|\vec{d}(S)|_{\text{top-cn}} \leq \frac{\text{opt}}{1 - \epsilon}$.

We describe the remaining steps of the algorithm in more detail.

4.3 Step 3 of Algorithm 4: Choosing Representatives

■ **Algorithm 5** $\text{TpcnC-choose-R}()$.

-
- 1: $R \leftarrow \emptyset$
 - 2: **repeat** k times:
 - 3: randomly choose a client $j \in C$, with probabilities proportional to $d(j, S)$
 - 4: $R \leftarrow R \cup \{j\}$
 - 5: **return** R
-

In this step, we call the procedure $\text{TpcnC-choose-R}()$, described in Algorithm 5. Before analyzing the properties of the representative set R , we make some definitions and partition the optimum clusters into 3 types: Given an optimum cluster (i^*, J) with core J' , we say (i^*, J) is of

- *type-1* if $\sum_{j \in J'} d(j, S) \geq \frac{\epsilon^3}{k} \cdot |\vec{d}(S)|_1$,
- *type-2* if $\sum_{j \in J'} d(j, S) < \frac{\epsilon^3}{k} \cdot |\vec{d}(S)|_1$ and $\sum_{j \in J'} d(i^*, j) \geq \frac{\epsilon^2}{k} \cdot |\vec{d}(S)|_1$, and
- *type-3* if $\sum_{j \in J'} d(j, S) < \frac{\epsilon^3}{k} \cdot |\vec{d}(S)|_1$ and $\sum_{j \in J'} d(i^*, j) < \frac{\epsilon^2}{k} \cdot |\vec{d}(S)|_1$.

⁵ Let $\Delta := \max\{d(u, v) : u, v \in F \cup C, d(u, v) < \infty\}$. Recall that after the standard discretization, all finite distances are integers in $[0, \Delta]$, with $\Delta = \text{poly}(n)$.

Notice that the three types partition all the optimum clusters.

► **Lemma 28.** *With probability at least $\frac{\epsilon^{3k}}{k^k}$, the following event happens: For every optimum cluster (i^*, J) of type-1 with core J' , we have $J' \cap R \neq \emptyset$. In other words, R intersects the core of every type-1 optimum cluster.*

From now on, we assume Step 3 of Algorithm 4 is successful, which means the event in Lemma 28 occurs.

► **Lemma 29.** *For every type-1 or 2 optimum cluster (i^*, J) with per-client cost $\bar{d}$, we have $d(i^*, R \cup S) \leq \frac{\bar{d}}{1-\epsilon}$.*

The proofs of Lemma 28 and Lemma 29 are deferred to the full version of the paper.

4.4 Steps 4 and 5 of Algorithm 4: Guessing Pivots and Radius Vector

In Steps 4 and 5 of Algorithm 4, we guess a multi-set P of k pivots from $R \cup S$, and a radius r_p for every $p \in P$. There is a *desired* pivot p for every optimum cluster (i^*, J) , and the pivot p has a *desired* radius. Step 4 is successful if P is the set of k desired pivots for the k optimum clusters, and Step 5 is successful if r_p for each $p \in P$ is the desired radius for p . For every optimum cluster (i^*, J) , the desired pivot p for the cluster and its desired radius are defined as follows:

- If (i^*, J) is of type-1 or 2, then $d(i^*, S \cup R) \leq \frac{\bar{d}}{1-\epsilon}$ by Lemma 29. The desired pivot p is the closest point in $S \cup R$ to i^* . Thus, we have $d(i^*, p) \leq \frac{\bar{d}}{1-\epsilon}$. The desired radius is $d(i^*, p)$ rounded up to the nearest integer power of $1 + \epsilon$. Therefore, if r_p is the desired radius, we have $r_p \leq \frac{1+\epsilon}{1-\epsilon} \bar{d}$.
- If (i^*, J) is of type-3, then the desired pivot p is the closest facility in S to i^* , and its desired radius is 0.

Again, we assume the two steps are successful from now on. Therefore, P is the set of desired pivots for the optimum clusters and r_p for each $p \in P$ is the desired radius for p .

► **Lemma 30.** *There is a set S'^* of k facilities, one from $\text{ball}_F(p, r_p)$ for each $p \in P$, and an assignment vector $\sigma \in (S'^*)^C$ such that the following properties hold.*

$$(30a) \quad \left| (d(j, \sigma_j))_{j \in C} \right|_{\text{top-cn}} \leq (1 + O(\epsilon)) \cdot \text{opt}.$$

(30b) *For every $p \in P$, and the facility $i'^* \in S'^*$ in $\text{ball}_F(p, r_p)$, we have*

$$r_p \leq \frac{1+\epsilon}{1-\epsilon} \cdot \frac{\sum_{j \in \sigma^{-1}(i'^*)} d(i'^*, j)}{|\sigma^{-1}(i'^*)|}.$$

We remark that the balls $\text{ball}_F(p, r_p)$ may overlap with each other and thus one facility in P may be in two different balls. However, we guarantee that there is precisely one facility in S'^* that is designated to $\text{ball}_F(p, r_p)$ for each $p \in P$, and the facility is inside the ball.

Proof of Lemma 30. We construct the solution S'^* with the assignment σ as follows. For every optimum cluster (i^*, J) with per-client-cost $\bar{d}$, we include a facility i'^* in S'^* . Let p be the desired pivot for this optimum cluster. If (i^*, J) is of type-1 or 2, we let $i'^* = i^*$. In this type, we have $i'^* = i^* \in \text{ball}_F(p, r_p)$ as r_p is the desired radius for p . Otherwise, we are of type-3 and the pivot p is the nearest facility in S to i^* . We let $i'^* = p$ and thus $i'^* \in \text{ball}_F(p, r_p = 0)$. In any type, we let $\sigma_j = i'^*$ for every $j \in J$. b holds if (i^*, J) is of type-1 or 2 as $r_p \leq \frac{1+\epsilon}{1-\epsilon} \bar{d}$. of type-3, we have $r_p = 0$ and the property holds trivially.

It remains to show a. Notice that the difference between the left side of the inequality in a and opt come from type-3 optimum clusters. Therefore, for every type-1 or 2 optimum cluster (i^*, J) , we define $\rho_j = 0$ for every $j \in J$. Fix a type-3 optimum cluster (i^*, J) , and its correspondent desired pivot p , we define $\rho_j = (d(j, p) - d(j, i^*))^+$ for every $j \in J$.

$$\sum_{j \in J} \rho_j \leq |J| \cdot d(i^*, p) \leq |J| \cdot \frac{1}{|J'|} \left(\frac{\epsilon^3}{k} + \frac{\epsilon^2}{k} \right) |\vec{d}(S)|_1 \leq \frac{1}{\epsilon} \cdot \frac{2\epsilon^2}{k} |\vec{d}(S)|_1 = \frac{2\epsilon}{k} |\vec{d}(S)|_1.$$

The second inequality comes from triangle inequality and the definition of type-3.

Summing up over all type-3 optimum clusters, we have $|\rho|_1 \leq 2\epsilon |\vec{d}(S)|_1$. Therefore,

$$|\rho|_{\text{top-cn}} \leq |\rho|_1 \leq O(\epsilon) \cdot |\vec{d}(S)|_1 \leq O(\epsilon) \cdot |\vec{d}(S)|_{\text{top-cn}} \leq O(\epsilon) \cdot \text{opt}.$$

The third inequality used that $c \in (\frac{1}{\epsilon}, 1]$. ◀

4.5 Step 6 of Algorithm 4: Find Clustering using P and $(r_p)_{p \in P}$

In this final step, we try to find a clustering using the pivot set P and the radius vector $(r_p)_{p \in P}$ as a guide, by calling `TpcnC-clustering-with-pivots`($P, (r_p)_{p \in P}$) in Algorithm 6. The existence of a good clustering is guaranteed by Lemma 30.

■ **Algorithm 6** `TpcnC-clustering-with-pivots`().

-
- 1: solve program (1) to obtain solution $((x_{ij}^{(p)})_{p,i,j}, (y_i^{(p)})_{p,i}, \delta \in \mathcal{S}_n)$
 - 2: **for** every $p \in P$ **do** randomly open a facility i , with probabilities $y_i^{(p)}$
 - 3: **return** the set of open facilities
-

In Step 1 of Algorithm 6, we solve the following program:

$$\min \quad \overline{\text{top-cn}}(\delta) \tag{1}$$

$$\sum_{i \in \text{ball}_F(p, r_p)} y_i^{(p)} = 1 \quad \forall p \in P \tag{2}$$

$$x_{ij}^{(p)} \leq y_i^{(p)} \quad \forall p \in P, i \in \text{ball}_F(p, r_p), j \in C \tag{3}$$

$$\sum_{p \in P, i \in \text{ball}(p, r_p)} x_{ij}^{(p)} = 1 \quad \forall j \in C \tag{4}$$

$$\frac{1+\epsilon}{1-\epsilon} \sum_{j \in C} d(i, j) x_{ij}^{(p)} - r_p \sum_{j \in C} x_{ij}^{(p)} \geq 0 \quad \forall p \in P, i \in \text{ball}_F(p, r_p) \tag{5}$$

$$\delta_a - \sum_{p, i, j: d(i, j) = a} x_{ij}^{(p)} = 0 \quad \forall a \in [0, \Delta] \tag{6}$$

$$x_{ij}^{(p)}, y_i^{(p)} \geq 0 \quad \forall p \in P, i \in F, j \in C \tag{7}$$

We now focus on the correspondent integer program to program (1), whose goal is to find the set S'^* and the assignment σ in Lemma 30. $y_i^{(p)}$ indicates whether i is the open facility in $\text{ball}_F(p, r_p)$, and $x_{ij}^{(p)}$ indicates whether the client j is connected to facility i in $\text{ball}_F(p, r_p)$. (2) indicates we open exactly one facility in $\text{ball}_F(p, r_p)$, (3) indicates that a client can only be connected to an open facility, and (4) requires every client to be connected to a facility in all balls. (5) is due to b, and (6) gives the definition of $\delta \in \mathcal{S}_n$.

We then describe the objective (1), which is not a linear function of δ . By Lemma 8, we know that $\overline{\text{top-cn}}(\delta) = \min_{t \in [0, \Delta]} (\sum_a (a - t)^+ \delta_a + cnt)$, we can solve the program by enumerating integers $t \in [0, \Delta]$.⁶

After solving the CP(1), we obtain the solution $((x_{ij}^{(p)})_{p,i,j}, (y_i^{(p)})_{p,i}, \delta \in \mathcal{S}_n)$. Let $\mathbf{1p}'$ be the value of the solution. By Property a, we have $\mathbf{1p}' \leq (1 + O(\epsilon))\text{opt}$.

Let ψ be integral occurrence-times vector of the solution given by the algorithm. That is, for every $a \in \mathbb{Z}_{\geq 0}$, ψ_a is the number of clients in C with connection distance a . So our goal is to upper bound $\mathbb{E}[\overline{\text{top-cn}}(\psi)]$ in terms of $\mathbf{1p}'$.

We define a vector $\delta' \in \mathcal{S}_n$ as follows: $\delta'_{a'} = \sum_{p,i,j: \lfloor r_p \rfloor = a'} x_{ij}^{(p)}$ for every $a' \in \mathbb{Z}_{\geq 0}$. That is, for every p, i, j , we include $x_{ij}^{(p)}$ fractional connection of distance $\lfloor r_p \rfloor$ in δ' . By Corollary 22 and (5), we have

$$\delta' \preceq_{\frac{1+\epsilon}{1-\epsilon}} \delta.$$

To see the inequality, we consider how δ and δ' is constructed: For every $p \in P$ and $i \in \text{ball}_F(p, r_p)$, we include $x_{ij}^{(p)}$ fraction of distance $d(i, j)$ in δ for every j , and we include $\sum_{j \in C} x_{ij}^{(p)}$ fraction of distance $\lfloor r_p \rfloor \leq \frac{1+\epsilon}{1-\epsilon} \sum_{j \in C} d(i, j) x_{ij}^{(p)} / \sum_{j \in C} x_{ij}^{(p)}$ in δ' . Thus Corollary 22 can be applied.

The main lemma we prove is the following:

► **Lemma 31.** $\mathbb{E}[\psi]$ is dominated by $(1 - \frac{1}{e})\delta + \frac{1}{e} \cdot (\delta \oplus \delta' \otimes 2)$.

We show why the lemma implies the desired approximation ratio. By the concavity of the $\overline{\text{top-cn}}(\cdot)$ function in Lemma 19, we have

$$\begin{aligned} \mathbb{E}[\overline{\text{top-cn}}(\psi)] &\leq \overline{\text{top-cn}}(\mathbb{E}[\psi]) \leq \overline{\text{top-cn}}\left(\left(1 - \frac{1}{e}\right)\delta + \frac{1}{e} \cdot (\delta \oplus \delta' \otimes 2)\right) \\ &\leq \left(1 + \frac{2}{ec} \cdot \frac{1+\epsilon}{1-\epsilon}\right) \cdot \overline{\text{top-cn}}(\delta) \leq \left(1 + \frac{2}{ec} + O(\epsilon)\right) \text{opt}. \end{aligned}$$

Proof of Lemma 31. We focus on a single client $j \in C$. It is optimal to connect j to the nearest open facility. But for the sake of analysis, it is convenient to connect j to a random and possibly sub-optimal facility. Also, the connection cost we impose on j could be larger than its actual cost.

Abusing notations slightly, we use a pair (p, i) to denote the copy of facility i dedicated to $\text{ball}_F(p, r_p)$. The probability that one facility in $\{(p, i) : x_{ij}^{(p)} > 0\}$ is open is at least $1 - \frac{1}{e}$. Using a contention resolution scheme, we can connect j to 0 or 1 open facility in the set, such that

$$\Pr[j \text{ connected to } (p, i)] = \left(1 - \frac{1}{e}\right) x_{ij}^{(p)}, \forall (p, i).$$

When j is connected to (p, i) , we impose a connection cost of $d(i, j)$ for j . In this type, we say j is *directly* connected.

With the remaining probability of $1 - (1 - \frac{1}{e}) \sum_{i,p} x_{ij}^{(p)} = 1 - (1 - \frac{1}{e}) = \frac{1}{e}$, we say j is *indirectly* connected. We then specify the connection cost in this type. For every facility (p, i) with $x_{ij}^{(p)} > 0$, we know that some facility in $\text{ball}_F(p, r_p)$ must be open. Then we can

⁶ By allowing a $1 + \epsilon$ loss, we can assume t is an integer power of $1 + \epsilon$, and we only need to enumerate $O(\frac{\log n}{\epsilon})$ different values of t .

impose a connection cost of $d(i, j) + 2 \lfloor r_p \rfloor$ on j by connecting it to the facility in the ball. We make the indirect connection randomly using $x_{ij}^{(p)}$ values: conditioned on that we make an indirect connection for j , we randomly choose a pair (p, i) with probabilities $x_{ij}^{(p)}$, and we impose a connection cost of $d(i, j) + 2 \lfloor r_p \rfloor$ on j .

Therefore, the connection cost of j is distributed as follows:

- For every facility (p, i) , j is directly connected to (p, i) (and thus incurs a cost of $d(i, j)$) with probability $(1 - \frac{1}{e}) x_{ij}^{(p)}$.
- For every facility (p, i) , j is indirectly connected via the facility (p, i) (and thus incurs a cost of $d(i, j) + 2 \lfloor r_p \rfloor$) with probability $\frac{1}{e} \cdot x_{ij}^{(p)}$.

Now we consider all clients $j \in C$. We have that $\mathbb{E}[\psi]$ is dominated by $(1 - \frac{1}{e})\delta + \frac{1}{e} \cdot \phi$ for some ϕ obtained by adding δ and $\delta' \otimes 2$, which is dominated by $\delta \oplus \delta' \otimes 2$ by Lemma 25. Therefore, by Lemma 21, we have $\mathbb{E}[\psi]$ is dominated by $(1 - \frac{1}{e})\delta + \frac{1}{e} \cdot (\delta \oplus \delta' \otimes 2)$. ◀

4.6 Wrapping up

Therefore, when all steps are successful, and all our guesses are correct, the algorithm returns a solution with expected top-cn norm being at most $(1 + \frac{2}{ec} + O(\epsilon)) \text{opt}$.

There are $|S \cup R|^k$ different choices for P in Step 4 of Algorithm 4, and $O\left(\frac{\log n}{\epsilon}\right)^k$ different choices for $(r_p)_{p \in P}$ in Step 5. The success probability of Steps 2 and 3 in Algorithm 5 are respectively $1 - \frac{1}{n^2}$ and $\frac{\epsilon^{3k}}{k^k}$. Therefore, if we enumerate all choices of P and $(r_p)_{p \in P}$, and repeat the algorithm $O\left(\frac{k^k}{\epsilon^{3k}} \cdot \log n\right)$ times, the success probability can be increased to $1 - 1/\text{poly}(n)$. Overall, we obtain a final algorithm with running time

$$O\left(\frac{k \log n}{\epsilon^4}\right)^k \cdot \text{poly}(n) \leq O\left(\frac{k^2 \log k}{\epsilon^4}\right)^k \cdot \text{poly}(n) \leq \text{poly}\left(k, \frac{1}{\epsilon}\right)^k \cdot \text{poly}(n),$$

that with high probability outputs a solution whose top-cn cost is at most $(1 + \frac{2}{ec} + O(\epsilon)) \cdot \text{opt}$. To see the first inequality, notice that either $\log n \leq k \log k$, or $(\log n)^k \leq \text{poly}(n)$; thus $(\log n)^k \leq \max\{(k \log k)^k, \text{poly}(n)\}$.

5 $(3, 1 + \frac{2}{ec} + \epsilon)$ -Bi-Criteria Approximation for $(L_\infty, \text{top-cn})$ -Norms k -Clustering Problem

In this section, we prove Theorem 4 by giving the $(3, 1 + \frac{2}{ec})$ -bi-criteria approximation for the $(L_\infty, \text{top-cn})$ -norms k -clustering problem, for $c \in (\frac{1}{e}, 1]$.

► **Definition 32.** In the $(L_\infty, \text{top-cn})$ -Norms k -Clustering problem, we are given F, m, C, n, k and d as in Definition 11, and a number $L \in \mathbb{R}_{\geq 0}$, the goal of the problem is to find a set $S \subseteq F$ of k facilities so as minimize $|\vec{d}(S)|_{\text{top-cn}}$, subject to $|\vec{d}(S)|_\infty \leq L$.

Let opt be the optimum value of the instance. A solution $S \subseteq F, |S| = k$ is called an (α, β) -bi-criteria approximation for the instance, for some $\alpha, \beta \geq 1$, if $|\vec{d}(S)|_\infty \leq \alpha L$ and $|\vec{d}(S)|_{\text{top-cn}} \leq \beta \cdot \text{opt}$.

To prove Theorem 4, the main modification to Algorithm 5 is that in Step 2, after applying Theorem 13 for the top-cn norm to obtain $\mathcal{S}$ and S , we add $O(k \log n)$ facilities that form a 1-approximation (this can be achieved using LP rounding or Theorem 13) for the L_∞ norm to S . Now S is a $(1 + O(\epsilon))$ -approximation for top-cn norm and has L_∞ cost L . Then, when guessing r_p 's in Step 5, we guarantee $r_p \leq L$. In the end, we return the solution with the smallest top-cn cost, whose L_∞ cost at most $3L$.

Clearly, if our guesses are correct, then the solution constructed have L_∞ cost at most $3L$. Consider an optimum cluster (i^*, J) , where every $j \in J$ has $d(i^*, j) \leq L$. If it is of type-1 or 2, the pivot p is in J and thus has $d(i^*, p) \leq L$. We are guaranteed to open a facility with distance $r_p \leq L$ to p . Therefore, the distance of the facility to i^* is at most $2L$, implying that the distance from all clients J to i^* is at most $3L$. When (i^*, J) is of type-3, then the nearest facility to i^* in S has distance at most $2L$ to i^* , and thus distance at most $3L$ to all clients in J . Finally, for the case $c \in (0, \frac{1}{e}]$, a $(3, 3 + \epsilon)$ -approximation can be achieved easily by modifying Algorithm in the proof of Theorem 13, referring to the full version of the paper.

References

- 1 Karen Aardal, Pieter L. van den Berg, Dion Gijswijt, and Shanfei Li. Approximation algorithms for hard capacitated k -facility location problems. *European Journal of Operational Research*, 242(2):358–368, 2015. doi:10.1016/J.EJOR.2014.10.011.
- 2 Fateme Abbasi, Sandip Banerjee, Jaroslaw Byrka, Parinya Chalermsook, Ameet Gadekar, Kamyar Khodamoradi, Dániel Marx, Roohani Sharma, and Joachim Spoerhase. Parameterized approximation schemes for clustering with general norm objectives. In *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023*, pages 1377–1399, 2023. doi:10.1109/FOCS57990.2023.00085.
- 3 Marek Adamczyk, Jaroslaw Byrka, Jan Marcinkowski, Syed Mohammad Meesum, and Michal Włodarczyk. Constant-factor FPT approximation for capacitated k -median. In *27th Annual European Symposium on Algorithms, ESA 2019*, volume 144, pages 1:1–1:14, 2019. doi:10.4230/LIPIcs.ESA.2019.1.
- 4 Sara Ahmadian, Ashkan Norouzi-Fard, Ola Svensson, and Justin Ward. Better guarantees for k -means and euclidean k -median by primal-dual algorithms. *SIAM Journal on Computing*, 49(4), 2020. doi:10.1137/18M1171321.
- 5 Soroush Alamdari and David B. Shmoys. A bicriteria approximation algorithm for the k -center and k -median problems. In *Workshop on Approximation and Online Algorithms, WAOA 2017*, volume 10787, pages 66–75, 2017. doi:10.1007/978-3-319-89441-6_6.
- 6 Hyung-Chan An, Aditya Bhaskara, Chandra Chekuri, Shalmoli Gupta, Vivek Madan, and Ola Svensson. Centrality of trees for capacitated k -center. *Mathematical Programming*, 154(1):29–53, 2015. doi:10.1007/S10107-014-0857-Y.
- 7 Ali Aouad and Danny Segev. The ordered k -median problem: surrogate models and approximation algorithms. *Mathematical Programming*, 177(1-2):55–83, 2019. doi:10.1007/S10107-018-1259-3.
- 8 Judit Barilan, Guy Kortsarz, and David Peleg. How to allocate network centers. *Journal of Algorithms*, 15(3):385–415, 1993. doi:10.1006/JAGM.1993.1047.
- 9 Vladimir Braverman, Vincent Cohen-Addad, Shaofeng H.-C. Jiang, Robert Krauthgamer, Chris Schlegelshohn, Mads Bech Tofttrup, and Xuan Wu. The power of uniform sampling for coresets. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022*, pages 462–473, 2022. doi:10.1109/FOCS54457.2022.00051.
- 10 Vladimir Braverman, Shaofeng H.-C. Jiang, Robert Krauthgamer, and Xuan Wu. Coresets for ordered weighted clustering. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019*, volume 97, pages 744–753, 2019. URL: <http://proceedings.mlr.press/v97/braverman19a.html>.
- 11 Jarosław Byrka, Krzysztof Fleszar, Bartosz Rybicki, and Joachim Spoerhase. Bi-factor approximation algorithms for hard capacitated k -median problems. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015*, pages 722–736, 2014. doi:10.1137/1.9781611973730.49.
- 12 Jaroslaw Byrka, Thomas W. Pensyl, Bartosz Rybicki, Aravind Srinivasan, and Khoa Trinh. An improved approximation for k -median and positive correlation in budgeted optimization. *ACM Transactions on Algorithms*, 13(2):23:1–23:31, 2017. doi:10.1145/2981561.

- 13 Jarosław Byrka, Bartosz Rybicki, and Sumedha Uniyal. An approximation algorithm for uniform capacitated k -median problem with $(1 + \epsilon)$ -capacity violation. In *International Conference on Integer Programming and Combinatorial Optimization, IPCO 2016*, volume 9682, pages 262–274, 2016. doi:10.1007/978-3-319-33461-5_22.
- 14 Jarosław Byrka, Krzysztof Sornat, and Joachim Spoerhase. Constant-factor approximation for ordered k -median. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018*, pages 620–631, 2018. doi:10.1145/3188745.3188930.
- 15 Deeparnab Chakrabarty and Chaitanya Swamy. Interpolating between k -median and k -center: Approximation algorithms for ordered k -median. In *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018*, volume 107, pages 29:1–29:14, 2018. doi:10.4230/LIPIcs.ICALP.2018.29.
- 16 Deeparnab Chakrabarty and Chaitanya Swamy. Approximation algorithms for minimum norm and ordered optimization problems. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019*, pages 126–137, 2019. doi:10.1145/3313276.3316322.
- 17 Moses Charikar and Sudipto Guha. Improved combinatorial algorithms for the facility location and k -median problems. In *40th Annual Symposium on Foundations of Computer Science, FOCS 1999*, pages 378–388, 1999. doi:10.1109/SFFCS.1999.814609.
- 18 Ke Chen. On k -median clustering in high dimensions. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006*, pages 1177–1185, 2006. URL: <http://dl.acm.org/citation.cfm?id=1109557.1109687>.
- 19 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. doi:10.1137/070699007.
- 20 Julia Chuzhoy and Yuval Rabani. Approximating k -median with non-uniform capacities. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2005*, volume 5, pages 952–958, 2005. URL: <http://dl.acm.org/citation.cfm?id=1070432.1070569>.
- 21 Vincent Cohen-Addad, Hossein Esfandiari, Vahab Mirrokni, and Shyam Narayanan. Improved approximations for euclidean k -means and k -median, via nested quasi-independent sets. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2022*, pages 1621–1628, 2022. doi:10.1145/3519935.3520011.
- 22 Vincent Cohen-Addad, Andreas Emil Feldmann, and David Saulpic. Near-linear time approximation schemes for clustering in doubling metrics. *Journal of the ACM*, 68(6):1–34, 2021. doi:10.1145/3477541.
- 23 Vincent Cohen-Addad, Fabrizio Grandoni, Euiwoong Lee, Chris Schwiegelshohn, and Ola Svensson. A $(2 + \epsilon)$ -approximation algorithm for metric k -median. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing, STOC 2025*, pages 615–624, 2025. doi:10.1145/3717823.3718299.
- 24 Vincent Cohen-Addad, Anupam Gupta, Amit Kumar, Euiwoong Lee, and Jason Li. Tight fpt approximations for k -median and k -means. In *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019*, volume 132, pages 42:1–42:14, 2019. doi:10.4230/LIPIcs.ICALP.2019.42.
- 25 Vincent Cohen-Addad and Jason Li. On the fixed-parameter tractability of capacitated clustering. In *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019*, volume 132, pages 41–1, 2019. doi:10.4230/LIPIcs.ICALP.2019.41.
- 26 Marek Cygan, MohammadTaghi Hajiaghayi, and Samir Khuller. Lp rounding for k -centers with non-uniform hard capacities. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012*, pages 273–282, 2012. doi:10.1109/FOCS.2012.63.
- 27 Gökalp Demirci and Shi Li. Constant approximation for capacitated k -median with $(1 + \epsilon)$ -capacity violation. In *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016*, volume 55, pages 73:1–73:14, 2016. doi:10.4230/LIPIcs.ICALP.2016.73.

- 28 David Eisenstat, Philip N. Klein, and Claire Mathieu. Approximating k -center in planar graphs. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014*, pages 617–627, 2014. doi:10.1137/1.9781611973402.47.
- 29 Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing, STOC 2003*, pages 448–455, 2003. doi:10.1145/780542.780608.
- 30 Tomás Feder and Daniel Greene. Optimal algorithms for approximate clustering. In *Proceedings of the 21th Annual ACM Symposium on Theory of Computing, STOC 1988*, pages 434–444, 1988. doi:10.1145/62212.62255.
- 31 Dan Feldman, Morteza Monemizadeh, and Christian Sohler. A ptas for k -means clustering based on weak coresets. In *Proceedings of the 23rd Annual Symposium on Computational Geometry, SoCG 2007*, pages 11–18, 2007. doi:10.1145/1247069.1247072.
- 32 Dan Feldman, Melanie Schmidt, and Christian Sohler. Turning big data into tiny data: Constant-size coresets for k -means, pca, and projective clustering. *SIAM Journal on Computing*, 49(3):601–657, 2020. doi:10.1137/18M1209854.
- 33 Zachary Friggstad, Mohsen Rezapour, and Mohammad R Salavatipour. Local search yields a ptas for k -means in doubling metrics. *SIAM Journal on Computing*, 48(2):452–480, 2019. doi:10.1137/17M1127181.
- 34 Teofilo F Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science*, 38:293–306, 1985. doi:10.1016/0304-3975(85)90224-5.
- 35 Dishant Goyal and Ragesh Jaiswal. Tight FPT approximation for constrained k -center and k -supplier. *Theoretical Computer Science*, 940:190–208, 2023. doi:10.1016/J.TCS.2022.11.001.
- 36 Sudipto Guha and Samir Khuller. Greedy strikes back: Improved facility location algorithms. *Journal of Algorithms*, 31(1):228–248, 1999. doi:10.1006/JAGM.1998.0993.
- 37 Dorit S Hochbaum and David B Shmoys. A best possible heuristic for the k -center problem. *Mathematics of Operations Research*, 10(2):180–184, 1985. doi:10.1287/MOOR.10.2.180.
- 38 Dorit S Hochbaum and David B Shmoys. A unified approach to approximation algorithms for bottleneck problems. *Journal of the ACM*, 33(3):533–550, 1986. doi:10.1145/5925.5933.
- 39 Lingxiao Huang, Jian Li, Pinyan Lu, and Xuan Wu. Coresets for constrained clustering: General assignment constraints and improved size bounds. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025*, pages 4732–4782, 2025. doi:10.1137/1.9781611978322.161.
- 40 Kamal Jain and Vijay V. Vazirani. Primal-dual approximation algorithms for metric facility location and k -median problems. In *40th Annual Symposium on Foundations of Computer Science, FOCS 1999*, pages 2–13, 1999. doi:10.1109/SFFCS.1999.814571.
- 41 Samir Khuller and Yoram J Sussmann. The capacitated k -center problem. *SIAM Journal on Discrete Mathematics*, 13(3):403–418, 2000. doi:10.1137/S0895480197329776.
- 42 Silvio Lattanzi and Christian Sohler. A better k -means++ algorithm via local search. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019*, volume 97, pages 3662–3671, 2019. URL: <http://proceedings.mlr.press/v97/lattanzi19a.html>.
- 43 Shanfei Li. An improved approximation algorithm for the hard uniform capacitated k -median problem. *arXiv preprint arXiv:1406.4454*, 2014. doi:10.4230/LIPIcs.APPROX-RANDOM.2014.325.
- 44 Shi Li. Approximating capacitated k -median with $(1 + \epsilon)k$ open facilities. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016*, pages 786–796, 2016. doi:10.1137/1.9781611974331.CH56.
- 45 Shi Li. On uniform capacitated k -median beyond the natural lp relaxation. *ACM Transactions on Algorithms*, 13(2):1–18, 2017. doi:10.1145/2983633.
- 46 Michal Osadnik. Fixed parameter tractable algorithm and coreset for the ordered k -median problem, master thesis, 2023. URL: <https://urn.fi/URN:NBN:fi:aalto-202405193486>.
- 47 Ján Plesník. A heuristic for the p -center problems in graphs. *Discrete Applied Mathematics*, 17(3):263–268, 1987. doi:10.1016/0166-218X(87)90029-1.
- 48 Arie Tamir. The k -centrum multi-facility location problem. *Discrete Applied Mathematics*, 109(3):293–307, 2001. doi:10.1016/S0166-218X(00)00253-5.