Interpolating between k-Median and k-Center: Approximation Algorithms for Ordered k-Median

Deeparnab Chakrabarty
Dept. of Computer Science, Dartmouth College, Hanover, NH 03755-3510, USA
deeparnab@dartmouth.edu

Chaitanya Swamy[1]
Dept. of Combinatorics and Optimization, Univ. Waterloo, Waterloo, ON N2L 3G1, Canada
cswamy@uwaterloo.ca

Abstract

We consider a generalization of k-median and k-center, called the *ordered k-median* problem. In this problem, we are given a metric space $(\mathcal{D}, \{c_{ij}\})$ with $n = |\mathcal{D}|$ points, and a non-increasing weight vector $w \in \mathbb{R}_+^n$, and the goal is to open k centers and assign each point $j \in \mathcal{D}$ to a center so as to minimize $w_1 \cdot$ (largest assignment cost) $+ w_2 \cdot$ (second-largest assignment cost) $+ \ldots + w_n \cdot$ (n-th largest assignment cost). We give an $(18 + \epsilon)$-approximation algorithm for this problem. Our algorithms utilize Lagrangian relaxation and the primal-dual schema, combined with an enumeration procedure of Aouad and Segev. For the special case of $\{0, 1\}$-weights, which models the problem of minimizing the ℓ largest assignment costs that is interesting in and of by itself, we provide a novel reduction to the (standard) k-median problem, showing that LP-relative guarantees for k-median translate to guarantees for the ordered k-median problem; this yields a nice and clean $(8.5 + \epsilon)$-approximation algorithm for $\{0, 1\}$ weights.

2012 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems, G.1.6 Optimization, G.2 Discrete Mathematics

Keywords and phrases Approximation algorithms, Clustering, Facility location, Primal-dual method

Digital Object Identifier 10.4230/LIPIcs.ICALP.2018.417

1 Introduction

Clustering is an ubiquitous problem that finds applications in various fields including data mining, machine learning, image processing, and bioinformatics. Many clustering problems involve finding a set F of at most k "centers" from an underlying set $\mathcal{D}$ of data points located in some metric space $\{c_{ij}\}_{i,j \in \mathcal{D}}$, and an assignment of data points to centers, so as to minimize some objective function of the assignment costs, i.e., the distances between data points and their assigned centers. These problems can typically also be stated as *facility-location* problems, wherein we seek a cost-effective way of opening facilities ($\equiv$ centers) and assigning clients ($\equiv$ data points) to open facilities. Given their widespread applicability, clustering and facility-location problems have been extensively studied in the Computer Science and Operations Research literature; see, e.g., [16, 22], as also the literature on the classical *k-median* (minimize *sum* of the assignment costs) [6, 13, 15, 4]), and *k-center* (minimize *maximum* assignment cost [10, 11]) problems.

We consider a common generalization of k-median and k-center, called the *ordered k-median problem* [17, 9]. As before, we are given a metric space $(\mathcal{D}, \{c_{ij}\}_{i,j \in \mathcal{D}})$, and an integer

[1] Supported in part by NSERC grant 327620-09 and an NSERC Discovery Accelerator Supplement Award.

45th International Colloquium on Automata, Languages, and Programming (ICALP 2018).
Editors: Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Don Sannella; Article No. 417; pp. 417:1–417:14
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

$k \geq 0$. We will often refer to points in $\mathcal{D}$ as clients. We are also given non-increasing, nonnegative weights $w_1 \geq w_2 \geq \ldots \geq w_n \geq 0$, where $n = |\mathcal{D}|$. For a vector $v \in \mathbb{R}^{\mathcal{D}}$, we use $v^{\downarrow}$ to denote the vector v with coordinates sorted in non-increasing order. That is, we have $v^{\downarrow}_i = v_{\sigma(i)}$, where σ is a permutation of $\mathcal{D}$ such that $v_{\sigma(1)} \geq v_{\sigma(2)} \geq \ldots v_{\sigma(n)}$. The goal in the ordered k-median problem is to choose a set F of k points from $\mathcal{D}$ as centers (or "facilities"), and assign each client $j \in \mathcal{D}$ to a center $i(j) \in F$, so as to minimize

$$cost\big(w; \vec{c} := \{c_{i(j)j}\}_{j \in \mathcal{D}}\big) := w^T \vec{c}^{\downarrow} = \textstyle\sum_{j=1}^{n} w_j \vec{c}^{\downarrow}_j.$$

Observe that when all the w_is are 1, we obtain the k-median problem; on the other hand, setting $w_1 = 1$, $w_2 = \ldots = w_n = 0$, yields the k-center problem. Indeed the special case with $\{0, 1\}$ weights is already interesting: that is, for some $\ell \in [n]$, we have $w_1 = \ldots = w_\ell = 1$ and all the remaining w_is are 0; this captures the problem of minimizing the ℓ largest assignment costs, which Tamir [23] calls the *ℓ-centrum problem*.

The ordered k-median problem can be motivated from various perspectives. The problem was proposed in network location theory as a convenient way of unifying the k-median and k-center objectives, as also some other objective functions considered in location theory (see, e.g., [17]). Such a versatile model is also useful in the context of clustering applications, wherein the clustering objective (e.g., k-median or k-center) is often a means to an end, namely, producing a "good" clustering. The ordered k-median problem yields a suite of clustering objectives, including those that interpolate between the k-median and k-center objectives, and thereby offers a useful means of obtaining a variety of clustering solutions (which motivates the question of developing efficient algorithms for (approximately) solving this problem). Another motivation for studying ordered k-median comes from a fairness perspective: if the weights decrease geometrically (at a sufficiently large rate), then an optimal ordered-k-median solution yields a *min-max fair* assignment-cost vector: that is, a solution that minimizes the maximum assignment cost, subject to which, it minimizes the second largest assignment cost, and so on. Finally, the ℓ-centrum problem can also be interpreted as the following *robust-optimization* version of k-median. Suppose there is some uncertainty in the client-set that needs to be clustered: in every scenario, some (at most) ℓ clients need to be clustered, and we need to determine the k centers and the assignment of clients to centers before knowing the scenario realization. Robust optimization seeks to minimize the maximum scenario cost, which leads to precisely the ℓ-centrum problem.

While the special cases of k-median and k-center have been considered extensively from the viewpoint of developing approximation algorithms, much less is known about the approximability of the ordered k-median problem, especially in general metrics. Aouad and Segev [2] obtained a logarithmic-approximation ratio for general metrics, and Alamdari and Shmoys [1] obtain a bicriteria approximation for the special case, where w is a convex combination of $(1, 0, \ldots, 0)$ and $\big(\frac{1}{n}, \frac{1}{n}, \ldots, \frac{1}{n}\big)$, which is called the *centridian* problem [12].

Our results. We obtain constant-factor approximation algorithms for the ordered k-median problem. Together with the concurrent work of [3], these constitute the first constant-factor approximation guarantees for ordered k-median. Our main result is an (deterministic) $(18+\epsilon)$-approximation algorithm for the ordered k-median problem (Theorem 7). Our algorithm utilizes the *primal-dual schema and Lagrangian relaxation*, and, hence, is *combinatorial*.

En route, in Section 2, we first develop constant-factor approximation algorithms for the case of $\{0, 1\}$-weights. This introduces many of the ideas needed to handle the general setting. We design two algorithms for this setting. Both algorithms are derived using a novel LP-relaxation that we propose for the problem, which leverages a key insight to circumvent the issue that the natural LP-relaxation has a large (non-constant) integrality gap.

Our first algorithm is a clean, combinatorial $(12+\epsilon)$-approximation algorithm that is based on the Jain-Vazirani primal-dual schema coupled with Lagrangian relaxation (Theorem 4). Both the algorithm and its analysis are versatile, and we show in Section 3 that the underlying ideas extend easily and, in combination with an enumeration procedure of [2], yield an $(18+\epsilon)$-approximation for the general setting. Our second algorithm for $\{0,1\}$-weights is based on LP-rounding, and yields an improved approximation factor via a novel *black-box reduction* to LP-relative algorithms for (standard) k-median. We show that an LP-relative α-approximation for k-median yields (essentially) a $2(\alpha+1)$-approximation; taking $\alpha = 3.25$ [7], we obtain an $(8.5+\epsilon)$-approximation for ordered k-median with $\{0,1\}$-weights (Theorem 5); we believe that this reduction is of independent interest.

Relationship with the work of [3]. Recently, we learnt that Byrka et al. [3] have also obtained a (randomized) $O(1)$-approximation guarantee (equal to $38+\epsilon$) for the ordered k-median problem. Our work was done independently and concurrently; a manuscript with the same approximation guarantees was posted on the arXiv in November 2017 [5]. In particular, our results for $\{0,1\}$ weights were obtained without knowledge of the work of [3]. But it was after we learnt of the results in [3] that we realized that our results can be extended to the general weighted setting.

While we use similar LP relaxations, our techniques are different. Whereas [3] crucially exploit properties of the Charikar-Li [7] LP-rounding algorithm, we leverage the (primal-dual + Lagrangian relaxation) methodology for k-median due to Jain and Vazirani [13]. Our algorithms are thus combinatorial. Our approximation factors improve upon those obtained in [3], *both* for $\{0,1\}$ weights and general weights; we believe that our algorithms and analyses are also simpler. Finally, our reduction to LP-relative algorithms for k-median shows that we do not need to rely on a *specific* k-median LP-rounding algorithm in order to tackle ordered k-median with $\{0,1\}$ weights, and suggests that the same might be true for general weights.

Our techniques. It is instructive to first discuss the $\{0,1\}$-weighted case. One of the main challenges is in coming up with a good LP-relaxation for this ℓ-centrum problem. The natural LP-relaxation augments the natural LP for k-median by imposing constraints encoding that the total assignment cost of any set of ℓ clients is at most B, where B is a new variable that we seek to minimize. It is well known that, even for (standard) k-median, one cannot hope to round an LP solution while approximately preserving the assignment cost of *each* client [6].[2] More significantly, whereas we can round and approximately preserve the sum of *all* assignment costs (as shown by k-median rounding), it turns out that we cannot preserve the sum of the ℓ largest assignment costs: the natural LP has a large (non-constant) integrality gap. This integrality gap is robust and cannot be alleviated by guessing the maximum assignment cost and incorporating this in the LP and the lower bound.[3] In essence, the cause for this disparity (between k-median and ℓ-centrum) is that the k-median objective crucially also includes the contribution from clients with small assignment costs.

The key insight that allows us to circumvent this difficulty is the following. Suppose we aim to find a solution of objective value $O(B)$. Then, it suffices to find a solution where the total assignment cost of clients having assignment cost at least B/ℓ is $O(B)$: the remaining clients can contribute at most additional B towards the ℓ-centrum objective, since we consider at most ℓ clients in the ℓ-centrum objective value. Moreover, if there is a solution of ℓ-centrum

[2] This is possible if we open $O(k)$ centers, using, e.g., the filtering-based algorithm of [21] for facility location.

[3] This is in contrast with k-center, where such preprocessing does mitigate the bad integrality gap of the natural LP and reduces it to a constant.

objective value at most B, then the total assignment cost of clients with assignment cost at least B/ℓ is at most B. Thus, given a "guess" B of the optimal value, *our new LP* (P_B) *seeks to minimize the total assignment cost of clients having assignment cost larger than* B/ℓ.

The LP (P_B) corresponds to the LP-relaxation for k-median with *non-metric distances* given by $\{f_B(c_{ij})\}_{i,j\in\mathcal{D}}$, where $f_B(d) = d$ if $d \geq B/\ell$, and is 0 otherwise. Despite this complication, we devise two ways of leveraging (P_B) to obtain a solution of ℓ-centrum cost $O(OPT_B + B)$ (which yields an $O(1)$-approximation for the correct choice of B), both of which involve simple procedures with a clean analysis; here, OPT_B denotes the optimal value of (P_B). Our first algorithm is based on the Jain-Vazirani (JV) template [13]. This is our main result for $\{0, 1\}$ weights (see Section 2.1), and this algorithm extends easily to the setting with general weights. We Lagrangify the cardinality constraint and move to the facility-location (FL) version where we may choose any number of centers but incur a fixed cost of (say) λ for each center we choose. We adapt the JV primal-dual algorithm and its analysis to obtain a so-called Lagrangian-multiplier-preserving guarantee for this FL version. By fine-tuning λ, we can then find two solutions, one with less than k centers and the other with more than k centers, whose convex combination has low cost; rounding this *bipoint solution* yields the final solution. This yields our 12-approximation algorithm.

The second algorithm utilizes LP-rounding. We show that after a clustering step, where we merge clients that are distance at most $\frac{B}{\ell}$-apart, the problem of rounding a solution to (P_B) reduces to that of rounding a fractional k-median solution on the cluster centers. Thus, any LP-relative α-approximation algorithm for k-median can be used to obtain a solution of cost at most $2(\alpha + 1)\overline{B}$.

For general weights, the key again is to consider k-median with suitable (non-metric) proxy distances analogous to the $f_B(c_{ij})$s. We utilize a clever enumeration idea due to [2] to obtain these proxy distances. Whereas with $\{0, 1\}$ weights, we created two distance buckets ($c_{ij} \geq B/\ell$ and $c_{ij} < B/\ell$) with weight multipliers 1 and 0, we now create $O\big(\log_{1+\epsilon}(\frac{n}{\epsilon})\big)$ buckets by grouping distances in powers of $(1 + \epsilon)$. We guess the average weight (roughly speaking) incurred for a bucket by an optimal solution, and use this as the weight multiplier for the bucket. As argued in [2]: (a) if we enumerate average weights in powers of $(1+\epsilon)$ then there are only polynomially many choices; and (b) the resulting proxy distances provide a good approximation for the actual $cost(w; .)$-cost. Finally, we show that the primal-dual algorithm and its analysis developed in Section 2.1 extends to solve the k-median problem with these new proxy distances. Combining these ingredients, we obtain an $(18 + \epsilon)$-approximation.

Other related work. While the ordered k-median problem, and its special cases, have been well studied in the Operations Research literature (see, e.g., [18, 14]), much of this work has focused either on modeling issues and formulations, or on solving the problem exactly in special cases, or via (non-polynomial time) heuristics. There is little prior work (i.e., discounting [3]) on the design of approximation algorithms for this problem, in general metrics. As mentioned earlier, for general metrics, we are only aware of the work of [2], who obtain a logarithmic-approximation ratio, and [1], who obtain a bicriteria approximation for the special case of the centridian problem.

A significant amount of research has taken place for special cases of the problem, e.g., the $k = 1$ setting [17], and the "continuous" version of the problem where centers can also be opened "in the middle of an edge" [19]. For these settings, fast exact algorithms have been developed in many interesting cases; see, e.g., [8, 23, 20] and the references therein. There is also a large body of work looking at compact integer-programming formulations, branch and bound methods etc.; for a detailed account of this and other work related to location theory and ordered-median models, we refer the reader to the books [18, 14].

2 The setting with $\{0,1\}$-weights

We first consider the setting with $\{0,1\}$ weights. Let $\ell \in [n]$ be such $w_1 = \ldots = w_\ell = 1$, $w_{\ell+1} = 0 = \ldots = w_n$. We abbreviate $cost(w; \vec{c})$ to $cost(\ell; \vec{c})$, or simply $cost(\vec{c})$. The $\{0,1\}$-weight setting serves as a natural starting point for two reasons. First, the problem of minimizing the ℓ most expensive assignment costs is a natural, well-motivated problem that is interesting in its own right. Second, the study of the $\{0,1\}$-case serves to introduce some of the key underlying ideas that are also used to handle the general setting. Notice also that a non-decreasing weight vector w can be written as a nonnegative linear-combination of such $\{0,1\}$ weight vectors.

The natural LP-relaxation for this ℓ-centrum problem has an $\Omega(\ell)$ integrality gap, and, as noted earlier, the integrality gap does not decrease even if we guess the maximum assignment cost and incorporate this in our LP and lower bound. Our constant-factor approximation algorithms are based on an alternate novel LP-relaxation, where, given a "guess" B of the optimal value, we seek to minimize the total assignment cost of clients having assignment cost at least B/ℓ. The rationale is that assignment costs that are smaller than B/ℓ can contribute at most B to the ℓ-centrum cost, and can hence be ignored when searching for a solution of ℓ-centrum cost $O(B)$. For $d \geq 0$, define $f_B(d) = d$ if $d \geq B/\ell$, and 0 otherwise. Throughout, i and j index points of $\mathcal{D}$. We consider the following LP.

$$\min \quad \sum_j \sum_i f_B(c_{ij}) x_{ij} \tag{P$_B$}$$

$$\text{s.t.} \quad \sum_i x_{ij} \geq 1 \qquad \text{for all } j \tag{1}$$

$$0 \leq x_{ij} \leq y_i \qquad \text{for all } i, j \tag{2}$$

$$\sum_i y_i \leq k. \tag{3}$$

Variable y_i indicates if facility i is open (i.e., i is chosen as a center), and x_{ij} indicates if client j is assigned to facility i. The first two constraints say that each client must be assigned to an open facility, and the third constraint encodes that at most k centers may be chosen.

An atypical aspect of our relaxation is that, while an integer solution corresponds to a solution to our problem, its objective value under (P$_B$) may *underestimate* the actual objective value; however, as alluded to above, the objective value of (P$_B$) is within an additive B of the actual objective value. Let OPT_B denote the optimal value of (P$_B$), and opt denote the optimal value of the ℓ-centrum problem.

▶ Claim 1. If $B \geq opt$, then $OPT_B \leq opt \leq B$.

Proof. Let $(\tilde{x}, \tilde{y})$ be the integer point corresponding to an optimal solution. Clearly, $(\tilde{x}, \tilde{y})$ is feasible to (P$_B$). There are at most ℓ assignment costs that are at least opt/ℓ (and hence at least B/ℓ). Therefore, the objective value of $(\tilde{x}, \tilde{y})$ is at most opt. ◀

▶ Claim 2. Let $\vec{c}$ be an assignment-cost vector (where $\vec{c}_j$ is the assignment cost of j). Then, $cost(\ell; \vec{c}) \leq \sum_j f_B(\vec{c}_j) + B$.

▶ Claim 3. For any $B \geq 0$, we have: (i) $f_B(x) \leq f_B(y)$ if $x \leq y$; (ii) $\max\{f_B(x), f_B(y), f_B(z)\} \geq f_B\left(\frac{x+y+z}{3}\right)$ for any $x, y, z \geq 0$; and (iii) $3 f_B(x/3) = f_{3B}(x)$ for any $x \geq 0$.

We may assume that we have $\overline{B} \leq (1+\epsilon) opt$ (e.g., by enumerating all possible choices for opt in powers of $(1+\epsilon)$, or using binary search to find, within a $(1+\epsilon)$-factor, the smallest B such that $OPT_B \leq B$). While (P$_B$) closely resembles the LP-relaxation for k-median, notice that the assignment costs $\{f_B(c_{ij})\}$ used in the objective of (P_B) *do not form a metric*. Despite this complication, we show that ($\mathrm{P}_{\overline{B}}$) can be leveraged to obtain a solution of $cost(\ell;.)$-cost $O(\overline{B})$. We devise two algorithms for obtaining such a guarantee. The first algorithm is based on the primal-dual method and the Jain-Vazirani (JV) template [13]; this yields a 12-approximation algorithm. The second algorithm is based on LP-rounding, and shows that any LP-relative α-approximation algorithm for k-median can be used to obtain a solution of $cost(\ell.)$-cost at most $2(\alpha+1)\overline{B}$.

▶ **Theorem 4.** *We can obtain a solution to the ℓ-centrum problem of cost at most $\big(12+O(\epsilon)\big)\cdot\overline{B}\leq\big(12+O(\epsilon)\big)\,opt$.*

▶ **Theorem 5.** *Let (kmed-P) denote the k-median LP:* $\min\left\{\sum_{j,i}c_{ij}x_{ij}\,:\,(1)\text{–}(3)\right\}$. *Let $\mathcal{A}$ be an α-approximation algorithm for k-median whose approximation guarantee is proved relative to (kmed-P). We can obtain a solution to the ℓ-centrum problem of cost at most $2(\alpha+1)\overline{B}$. Thus, taking $\mathcal{A}$ to be the 3.25-approximation algorithm in [7], we obtain an $(8.5+\epsilon)$-approximation algorithm for the ℓ-centrum problem.*

Although Theorem 4 yields a worse approximation factor, the underlying primal-dual algorithm and analysis are quite versatile and extend easily to the setting with general weights. We prove Theorem 4 in this extended abstract. The proof of Theorem 5 can be found in Appendix A of the arXiv version [5] of the paper.

2.1 Proof of Theorem 4

As noted earlier, the proof relies on the primal-dual method. The dual of ($\mathrm{P}_{\overline{B}}$) is as follows.

$$\max \quad \sum_j \alpha_j - k\cdot\lambda \tag{$\mathrm{D}_{\overline{B}}$}$$

$$\text{s.t.} \quad \alpha_j \leq f_{\overline{B}}(c_{ij})+\beta_{ij} \qquad \forall i,j \tag{4}$$

$$\sum_j \beta_{ij}\leq\lambda \qquad \forall i \tag{5}$$

$$\alpha,\lambda\geq 0.$$

Let $OPT := OPT_{\overline{B}}$ denote the optimal value of ($\mathrm{P}_{\overline{B}}$). We first fix λ and construct a solution that may open more than k centers but will have some near-optimality properties (see Theorem 6).

P1. **Dual-ascent.** Initialize $\mathcal{D}'=\mathcal{D}$, $\alpha_j=\beta_{ij}=0$ for all $i,j\in\mathcal{D}$, $F=\emptyset$. The clients in $\mathcal{D}'$ are called *active clients*. If $\alpha_j\geq f_{\overline{B}}(c_{ij})$, we say that j reaches i. (So if $c_{ij}\leq\overline{B}/\ell$, then j reaches i from the very beginning.)
We repeat the following until all clients become inactive. Uniformly raise the α_js of all active clients, and the β_{ij}s for (i,j) such that $i\notin F$, j is active, and can reach i until one of the following events happen.
- Some client $j\in\mathcal{D}$ reaches some i (and previously could not reach i): if $i\in F$, we *freeze* j, and remove j from $\mathcal{D}'$.
- Constraint (5) becomes tight for some $i\notin F$: we add i to F; for every $j\in\mathcal{D}'$ that can reach i, we freeze j and remove j from $\mathcal{D}'$.

P2. **Pruning.** Pick a maximal subset T of F with the following property: for every $j\in\mathcal{D}$, there is at most one $i\in T$ such that $\beta_{ij}>0$. Let $P=\{j:\exists i\in T \text{ s.t. } \beta_{ij}>0\}$.

P3. Return T as the set of centers, and assign every j to the nearest point (in terms of c_{ij}) in T, which we denote by $i(j)$.

▶ **Theorem 6.** *The solution satisfies* $3\lambda|T| + \sum_{j\in P} f_{\overline{B}}(c_{i(j)j}) + \sum_{j\notin P} f_{3\overline{B}}(c_{i(j)j}) \leq 3\sum_j \alpha_j$.

Proof. The proof resembles the analysis of the JV primal-dual algorithm for facility location, but the subtlety is that we need to deal with the complication that the $\{f_{\overline{B}}(c_{ij})\}_{i,j\in\mathcal{D}}$ "distances" do not form a metric.

Observe that for every $i \in T$, every client $j \in P$ for which $\beta_{ij} > 0$ satisfies $i(j) = i$. So

$$\sum_{j\in P} 3\alpha_j \geq \sum_{j\in P}\Big(3\beta_{i(j)j} + f_{\overline{B}}(c_{i(j)j})\Big) = 3\lambda|T| + \sum_{j\in P} f_{\overline{B}}(c_{i(j)j}).$$

We show that for each client $j \notin P$, there is some $i'' \in T$ such that $f_{3\overline{B}}(c_{i''j}) \leq 3\alpha_j$, which will complete the proof. Let $i \in F$ be the facility that caused j to freeze, so $f_{\overline{B}}(c_{ij}) \leq \alpha_j$. If $i \in T$, then we are done. Otherwise, since T is maximal, there is some $i' \in T$ and some client $k \in P$ such that $\beta_{i'k}, \beta_{ik} > 0$. Notice that $\alpha_j \geq \alpha_k$, since α_j grows at least until the time point when i joins F, and α_k grows until at most this time point. Therefore, $f_{\overline{B}}(c_{ik}), f_{\overline{B}}(c_{i'k}) \leq \alpha_k \leq \alpha_j$. We have $c_{i'j} \leq c_{i'k} + c_{ik} + c_{ij}$. Now, by Claim 3, we have $f_{3\overline{B}}(c_{i'j}) \leq f_{3\overline{B}}(c_{i'k} + c_{ik} + c_{ij}) = 3f_{\overline{B}}((c_{i'k} + c_{ik} + c_{ij})/3) \leq 3\max(f_{\overline{B}}(c_{ik}), f_{\overline{B}}(c_{i'k}), f_{\overline{B}}(c_{ij})) \leq 3\alpha_j$. ◀

Using standard arguments, by performing binary search on λ, we can achieve one of the following two outcomes.

(a) Obtain some λ such that the above algorithm returns a solution T with $|T| = k$: in this case, Theorem 6 implies that $\sum_j f_{3\overline{B}}(c_{i(j)j}) \leq 3OPT$, and Claim 2 then implies that the $cost(\ell; .)$-cost of our solution is at most $3OPT + 3\overline{B} \leq 6\overline{B}$.

(b) Obtain $\lambda_1 < \lambda_2$ with $\lambda_2 - \lambda_1 < \frac{\epsilon\overline{B}}{n}$ such that letting T_1 and T_2 be the solutions returned for λ_1 and λ_2, we have $k_1 := |T_1| > k > k_2 := |T_2|$. We describe below the procedure for extracting a low-cost feasible solution from T_1 and T_2, and analyze it, which will complete the proof of Theorem 4.

Extracting a feasible solution from T_1 and T_2 in outcome (b). Let $a, b \geq 0$ be such that $ak_1 + bk_2 = k$, $a + b = 1$. Thus, a convex combination of T_1 and T_2, called a *bipoint solution*, yields a feasible fractional solution and our task is to round this into a feasible solution. Let (α_1, β_1), (α_2, β_2) denote the dual solutions obtained for λ_1 and λ_2 respectively. Let $i_1(j)$ and $i_2(j)$ denote the centers to which j is assigned in T_1 and T_2 respectively. Let $d_{1,j} = f_{3\overline{B}}(c_{i_1(j)j})$ and $d_{2,j} = f_{3\overline{B}}(c_{i_2(j)j})$. Let $C_1 := \sum_j d_{1,j}$ and $C_2 := \sum_j d_{2,j}$. Then,

$$\begin{aligned} aC_1 + bC_2 &\leq 3a\Big(\sum_j \alpha_{1,j} - k_1\lambda_1\Big) + 3b\Big(\sum_j \alpha_{2,j} - k_2\lambda_2\Big) \\ &\leq 3a\Big(\sum_j \alpha_{1,j} - k\lambda_2\Big) + 3b\Big(\sum_j \alpha_{2,j} - k\lambda_2\Big) + 3ak_1(\lambda_2 - \lambda_1) \leq 3OPT + 3\epsilon\overline{B} \end{aligned}$$

where the last inequality follows since $(\alpha_1, \beta_1, \lambda_2)$, $(\alpha_2, \beta_2, \lambda_2)$ are feasible solutions to $(\mathrm{D}_{\overline{B}})$. If $b \geq 0.5$, then T_2 yields a feasible solution of $cost(\ell; .)$-cost at most $C_2 + 3\overline{B} \leq 6OPT + (3+\epsilon)\overline{B}$. So suppose $a \geq 0.5$. The procedure for rounding the bipoint solution is as follows.

B1. **Clustering.** We first match facilities in T_2 with a subset of facilities in T_1 as follows. Initialize $\mathcal{D}' \leftarrow \mathcal{D}$, $A \leftarrow \emptyset$, and $M \leftarrow \emptyset$. While $\mathcal{D}' \neq \emptyset$, we repeatedly pick the client $j \in \mathcal{D}'$ with minimum $d_{1,j} + d_{2,j}$ value, and add j to A. We add the tuple $(i_1(j), i_2(j))$ to M, remove from $\mathcal{D}'$ all clients k (including j) such that $i_1(k) = i_1(j)$ or $i_2(k) = i_2(j)$, and set $\sigma(k) = j$ for all such clients. Let $M_1 = M$ denote the matching when $\mathcal{D}' = \emptyset$. Next, for each unmatched $i \in T_2$, we pick an arbitrary unmatched facility $i' \in T_1$, and add (i', i) to M. Let F_1 be the set of T_1-facilities that are matched, and $S := \{j \in \mathcal{D} : i_1(j) \in F_1\}$. Note that $|F_1| = |M| = k_2$.

B2. **Opening facilities.** We will open k_2 facilities at locations in $A \cup M$, and $k - k_2$ facilities from $T_1 \setminus F_1$. We solve the following LP to determine how to do this. Variables z_i for every $i \in T_1 \setminus F_1$ indicate if we open facility i; variable θ indicates if we give preference to F_1 (i.e., the T_1-facilities in M), or the facilities in T_2 (which are always matched).

$$\min \quad \sum_{j \in S}\big(\theta d_{1,j} + (1-\theta) d_{2,j}\big) + \sum_{k \notin S}\big(z_{i_1(k)} d_{1,k} + (1 - z_{i_1(k)})(d_{2,k} + d_{1,\sigma(k)} + d_{2,\sigma(k)})\big) \quad \text{(R-P)}$$

$$\text{s.t.} \quad \sum_{i \in T_1 \setminus F_1} z_i \le k - k_2, \qquad \theta \in [0,1], \quad z_i \in [0,1] \ \ \forall i \in T_1 \setminus F_1.$$

The above LP is integral. Given an integral optimal solution $(\tilde{\theta}, \tilde{z})$ to (R-P), we open facilities as follows. We open the facilities in $T_1 \setminus F_1$ specified by the $\tilde{z}_i$ variables that are 1. If $\tilde{\theta} = 1$, we open all the T_1-facilities in $M \setminus M_1$, and if $\tilde{\theta} = 0$, we open all the T_2-facilities in $M \setminus M_1$. For some clients $j \in A$, we may open a facility at j (instead of at $i_1(j)$ or $i_2(j)$). For every $j \in A$, if $\tilde{\theta} d_{1,j} + (1 - \tilde{\theta}) d_{2,j} = 0$, then we open a facility at j; otherwise, we open a facility at $i_1(j)$ if $\tilde{\theta} = 1$ and at $i_2(j)$ if $\tilde{\theta} = 0$.

Analysis. It suffices to show that (R-P) has a fractional solution of small objective value, and that the integral optimal solution $(\tilde{\theta}, \tilde{z})$ yields a feasible solution to our problem whose $cost(\ell; .)$-cost is comparable to the objective value of $(\tilde{\theta}, \tilde{z})$ in (R-P).

For the former, we argue that setting $\theta = a$, $z_i = a$ for all $i \in T_1 \setminus F_1$ yields a feasible solution of objective value at most $2(aC_1 + bC_2)$. We have $\sum_{i \in T_1 \setminus F_1} z_i = a(k_1 - k_2) = k - k_2$. Every $j \in S$ contributes $ad_{1,j} + bd_{2,j}$ to the objective value of (R-P), which is also its contribution to $aC_1 + bC_2$. Consider $k \notin S$ with $\sigma(k) = j$, so $d_{1,j} + d_{2,j} \le d_{1,k} + d_{2,k}$. Its contribution to the objective value of (R-P) is $ad_{1,k} + b(d_{2,k} + d_{1,j} + d_{2,j}) \le (a+b)d_{1,k} + 2bd_{2,k}$, which is at most twice its contribution to $aC_1 + bC_2$.

For the latter, we first show that every $k \in S$ has assignment cost at most $\tilde{\theta} d_{1,k} + (1 - \tilde{\theta}) d_{2,k} + 6\overline{B}/\ell$. If a facility is opened in $\{k, i_1(k), i_2(k)\}$, then this clearly holds. Otherwise, it must be that $k \notin A$. Let $i = i_1(k)$ if $\tilde{\theta} = 1$, and $i_2(k)$ if $\tilde{\theta} = 0$. Since i is not open, it must be that i belongs to a tuple $(i_1(j), i_2(j))$ of M. Then, $j \in A$, and a facility is opened at j. we have that $c_{i,k} \le \tilde{\theta} d_{1,k} + (1 - \tilde{\theta}) d_{2,k} + 3\overline{B}/\ell$ and $c_{i,j} \le 3\overline{B}/\ell$. The last inequality follows since the fact that none of $i_1(j)$, $i_2(j)$ is open implies that $\tilde{\theta} d_{1,j} + (1 - \tilde{\theta}) d_{2,j} = 0$.

Now consider $k \notin S$ with $\sigma(k) = j$. If $\tilde{z}_{i_1(k)} = 1$, it's assignment cost is at most $d_{1,k} + 3\overline{B}/\ell$. Otherwise, a facility is opened in $\{j, i_1(j), i_2(j)\}$. If a facility is opened in $\{j, i_2(j)\}$, then k's assignment cost is at most $c_{i_2(k)k} + c_{i_2(j)j} \le d_{2,k} + d_{1,j} + d_{2,j} + 6\overline{B}/\ell$. Otherwise, it must be that $\tilde{\theta} = 1$ and $d_{1,j} = c_{i_1(j)j} > 3\overline{B}/\ell$; in this case, k' assignment cost is at most $c_{i_2(k)k} + c_{i_2(j)j} + c_{i_1(j)j} \le (d_{2,k} + 3\overline{B}/\ell) + (d_{2,j} + 3\overline{B}/\ell) + d_{1,j}$. Thus, the $cost(\ell; .)$-cost of our solution is at most the objective value of $(\tilde{\theta}, \tilde{z}) + 6\overline{B}$, which is at most $2(aC_1 + bC_2) + 6\overline{B} \le 6\,OPT + (6 + 3\epsilon)\overline{B} \le \big(12 + O(\epsilon)\big)\overline{B}$. This completes the proof.

3 The general weighted case

We now consider the general setting, where we have $n = |\mathcal{D}|$ non-increasing nonnegative weights $w_1 \ge \ldots \ge w_n \ge 0$, and the goal is to open k centers from $\mathcal{D}$ and assign each client $j \in \mathcal{D}$ to a center $i(j) \in F$, so as to minimize $cost\big(w; \vec{c} := \{c_{i(j)j}\}_{j \in \mathcal{D}}\big) := w^T \vec{c}^{\downarrow} = \sum_{j=1}^{n} w_j \vec{c}_j^{\downarrow}$.

By combining the ideas in Section 2 with an enumeration procedure due to Aouad and Segev in [2], we obtain the following result.

▶ **Theorem 7.** *We can obtain an $\big(18 + O(\epsilon)\big)$-approximation algorithm for ordered k-median that runs in time* $\mathrm{poly}\big((\frac{n}{\epsilon})^{1/\epsilon}\big)$.

As before, we define suitable proxy costs analogous to the $f_B(c_{ij})$s for the setting with general weights. By defining these appropriately, it will be easy to argue that the primal-dual algorithm and its analysis extend to the setting with general weights, since essentially the only property that we use about $\{f_B(c_{ij})\}$ costs in Section 2 is that they satisfy Claim 3. Instead of creating two distance buckets in the $\{0,1\}$ weighted case ($c_{ij} \geq B/\ell$ and $c_{ij} < B/\ell$), with weight multipliers 1 and 0, we now create $O\big(\log_{1+\epsilon}(\frac{n}{\epsilon})\big)$ buckets and utilize an enumeration idea due to Aouad and Segev [2]. In Section 3.1, we describe this enumeration procedure using our notation, and restate the main claims in [2] in a simplified form. Next, in Section 3.2, we discuss how to adapt the ideas in Section 2 to the k-median problem for the proxy costs (given by (7)) that we obtain from Section 3.1. At the end of this section, we combine this ingredients to prove Theorem 7.

3.1 Proxy costs and the enumeration idea of [2]

Throughout, let $\vec{o}^{\downarrow}$ denote the assignment-cost vector corresponding to an optimal solution, whose coordinates are sorted in non-increasing order. So the optimal cost opt is $\sum_{i=1}^{n} w_i \vec{o}_i^{\downarrow}$. By a standard argument, we can perturb w to eliminate very small weights w_i: for $i \in [n]$, set $\widetilde{w}_i = w_i$ if $w_i \geq \frac{\epsilon w_1}{n}$, and $\widetilde{w}_i = 0$ otherwise.

▶ Claim 8. For any vector $v \in \mathbb{R}_+^n$, we have $(1-\epsilon)cost(w;v) \leq cost(\widetilde{w};v) \leq cost(w;v)$.

Proof. Since $\widetilde{w}_i \leq w_i$ for all $i \in [n]$, the upper bound on $cost(\widetilde{w};v)$ is immediate. We have

$$cost(\widetilde{w};v) = \sum_{i=1}^{n} \widetilde{w}_i v_i^{\downarrow} = cost(w;v) - \sum_{i\in[n]:w_i<\epsilon w_1/n} w_i v_i^{\downarrow} \geq cost(w;v) - \tfrac{\epsilon w_1}{n} \cdot n v_1^{\downarrow}. \quad ◀$$

In the sequel, we always work with the $\widetilde{w}$-weights. We guess an estimate M of $\vec{o}_1^{\downarrow}$, and group distances in the range $\big[\frac{\epsilon M}{n}, M\big]$ (roughly speaking) by powers of $(1+\epsilon)$. Let T be the largest integer such that $\frac{\epsilon M}{n}(1+\epsilon)^T \leq M$. For $r = 0,\ldots,T$, we define the distance interval $I_r := \big[\frac{\epsilon M}{n}(1+\epsilon)^{T-r}, \frac{\epsilon M}{n}(1+\epsilon)^{T-r+1}\big)$. There are at most $1+\log_{1+\epsilon}\big(\frac{n}{\epsilon}\big) = O\big(\frac{1}{\epsilon}\log\frac{n}{\epsilon}\big)$ intervals.

Finally, we guess a non-increasing vector $w_0^{\mathsf{est}} \geq w_1^{\mathsf{est}} \geq \ldots \geq w_T^{\mathsf{est}}$, where the w_r^{est}s are powers of $(1+\epsilon)$ in the range $[\frac{\epsilon \widetilde{w}_1}{n}, \widetilde{w}_1(1+\epsilon))$. As argued in [2], there are only $\exp\big(O(\frac{1}{\epsilon}\log(\frac{n}{\epsilon}))\big) = O\big((\frac{n}{\epsilon})^{1/\epsilon}\big)$ choices for $w^{\mathsf{est}} := (w_0^{\mathsf{est}},\ldots,w_T^{\mathsf{est}})$. The intention is for w_r^{est} to represent (within a $(1+\epsilon)$-factor) the average $\widetilde{w}$-weight of the set $\{i \in [n] : \vec{o}_i^{\downarrow} \in I_r\}$. More precisely, we would like w_r^{est} to estimate the following quantity, for all $r \in \{0,\ldots,T\}$.

$$w_r^{\mathsf{avg}} := \begin{cases} \big(\sum_{i\in[n]:\vec{o}_i^{\downarrow}\in I_r} \widetilde{w}_i\big)/|\{i\in[n]:\vec{o}_i^{\downarrow}\in I_r\}| & \text{if } \{i\in[n]:\vec{o}_i^{\downarrow}\in I_r\}\neq\emptyset; \\ \min\{\widetilde{w}_i : \vec{o}_i^{\downarrow} \in \bigcup_{s<r} I_s\} & \text{if } \bigcup_{s<r} I_s \neq \emptyset; \\ \widetilde{w}_1 & \text{otherwise.} \end{cases} \tag{6}$$

The following claim will be useful.

▶ Claim 9. For any $r \in \{0,\ldots,T\}$, we have $w_r^{\mathsf{avg}} \geq \max\{\widetilde{w}_i : \vec{o}_i^{\downarrow} \notin \bigcup_{s\leq r} I_s\}$.

Proof. If w_r^{avg} is defined by cases 1 or 2 of (6), then the inequality follows since for every $i' \in \bigcup_{s\leq r} I_r$ and $i \notin \bigcup_{s\leq r} I_s$, we have $\widetilde{w}_{i'} \geq \widetilde{w}_i$ (since $\vec{o}_{i'}^{\downarrow} \geq \vec{o}_i^{\downarrow}$). If w_r^{avg} is defined by case 3 of (6), then $w_r^{\mathsf{avg}} = \widetilde{w}_1$, and again, the inequality holds. ◀

Given M and the corresponding intervals $I_0, \ldots, I_T$, and the vector w^{est}, we can now finally define our proxy costs as follows. For $d \geq 0$ and $\gamma \geq 1$, define

$$g_{M,w^{\mathsf{est}}}(\gamma; d) = \begin{cases} \widetilde{w}_1(1+\epsilon)d & \text{if } d/\gamma \geq \frac{\epsilon M}{n}(1+\epsilon)^{T+1}; \\ w_r^{\mathsf{est}} d & \text{if } d/\gamma \in I_r \text{ (where } r \in \{0, \ldots, T\}) \\ 0 & \text{if } d/\gamma < \frac{\epsilon M}{n}. \end{cases} \tag{7}$$

The above definition is essentially the scaled surrogate function in [2]. We abbreviate $g_{M,w^{\mathsf{est}}}(1; d)$ to $g_{M,w^{\mathsf{est}}}(d)$. The following two key lemmas are analogous to Claims 1 and 2, and show that for the right choice of M and w^{est}, evaluating the above proxy costs on an assignment-cost vector $\vec{c}$ yields a good estimate of the actual $cost(\widetilde{w}; .)$-cost of $\vec{c}$. Similar statements, albeit stated somewhat differently, are proved in [2].

▶ **Lemma 10** (adapted from [2]). *Suppose $M \geq \vec{o}_1^{\downarrow}$ and the w^{est} satisfies $w_r^{\mathsf{est}} \leq (1+\epsilon)w_r^{\mathsf{avg}}$ for all $r \in \{0, \ldots, T\}$. Then, $\sum_{i=1}^n g_{M,w^{\mathsf{est}}}(\vec{o}_i^{\downarrow}) \leq (1+\epsilon)^2 cost(\widetilde{w}; \vec{o}^{\downarrow})$.*

Proof. Since $\frac{\epsilon M}{n}(1+\epsilon)^{T+1} > M \geq \vec{o}_1^{\downarrow}$, there is no i such that $\vec{o}_i^{\downarrow} \geq \frac{\epsilon M}{n}(1+\epsilon)^{T+1}$. Fix $r \in \{0, \ldots, T\}$, and consider all $i \in [n]$ such that $\vec{o}_i^{\downarrow} \in I_r$. We have

$$\begin{aligned} \sum_{i \in [n]: \vec{o}_i^{\downarrow} \in I_r} g_{M,w^{\mathsf{est}}}(\vec{o}_i^{\downarrow}) = w_r^{\mathsf{est}} \sum_{i \in [n]: \vec{o}_i^{\downarrow} \in I_r} \vec{o}_i^{\downarrow} &\leq \frac{\epsilon M}{n}(1+\epsilon)^{T-r+1} \cdot w_r^{\mathsf{est}} \cdot \left|\{i \in [n] : \vec{o}_i^{\downarrow} \in I_r\}\right| \\ &\leq (1+\epsilon) \cdot \frac{\epsilon M}{n}(1+\epsilon)^{T-r+1} \cdot w_r^{\mathsf{avg}} \cdot \left|\{i \in [n] : \vec{o}_i^{\downarrow} \in I_r\}\right| \\ &= (1+\epsilon) \cdot \frac{\epsilon M}{n}(1+\epsilon)^{T-r+1} \cdot \sum_{i \in [n]: \vec{o}_i^{\downarrow} \in I_r} \widetilde{w}_i \leq (1+\epsilon)^2 \sum_{i \in [n]: \vec{o}_i^{\downarrow} \in I_r} \widetilde{w}_i \vec{o}_i^{\downarrow}. \end{aligned}$$

It follows that $\sum_{i=1}^n g_{M,w^{\mathsf{est}}}(\vec{o}_i^{\downarrow}) \leq (1+\epsilon)^2 cost(\widetilde{w}; \vec{o}^{\downarrow})$. ◀

▶ **Lemma 11** (adapted from [2]). *Let $\gamma \geq 1$. Let $M \geq 0$, and suppose w^{est} satisfies $w_r^{\mathsf{avg}} \leq w_r^{\mathsf{est}}$ for all $r \in \{0, \ldots, T\}$. Let $\vec{c}$ be an assignment-cost vector. Then, we have the upper bound $cost(\widetilde{w}; \vec{c}) \leq \sum_{i=1}^n g_{M,w^{\mathsf{est}}}(\gamma; \vec{c}_i) + \gamma(1+\epsilon)cost(\widetilde{w}; \vec{o}^{\downarrow}) + \gamma\epsilon\widetilde{w}_1 M$.*

Proof. We have $cost(\widetilde{w}; \vec{c}) = \sum_{i=1}^n \widetilde{w}_i \vec{c}_i^{\downarrow} \leq \sum_{i=1}^n g_{M,w^{\mathsf{est}}}(\gamma; \vec{c}_i) + \sum_{i:\ \widetilde{w}_i \vec{c}_i^{\downarrow} > g_{M,w^{\mathsf{est}}}(\gamma; \vec{c}_i^{\downarrow})} \widetilde{w}_i \vec{c}_i^{\downarrow}$. Consider some $i \in [n]$ for which $\widetilde{w}_i \vec{c}_i^{\downarrow} > g_{M,w^{\mathsf{est}}}(\gamma; \vec{c}_i^{\downarrow})$. It must be that $\vec{c}_i^{\downarrow}/\gamma < \frac{\epsilon M}{n}(1+\epsilon)^{T+1}$ as otherwise (see (7)), we have $g_{M,w^{\mathsf{est}}}(\gamma; \vec{c}_i^{\downarrow}) = (1+\epsilon)\widetilde{w}_1 \vec{c}_i^{\downarrow} > \widetilde{w}_i \vec{c}_i^{\downarrow}$. If $g_{M,w^{\mathsf{est}}}(\gamma; \vec{c}_i^{\downarrow}) = 0$, then we have $\widetilde{w}_i \vec{c}_i^{\downarrow}/\gamma < \widetilde{w}_i \cdot \frac{\epsilon M}{n} \leq \widetilde{w}_1 \cdot \frac{\epsilon M}{n}$.

Otherwise, we claim that $\vec{c}_i^{\downarrow}/\gamma \leq (1+\epsilon)\vec{o}_i^{\downarrow}$. Suppose not. Suppose $\vec{c}_i^{\downarrow}/\gamma \in I_r$, where $r \in \{0, \ldots, T\}$. Since $\frac{\vec{c}_i^{\downarrow}/\gamma}{\vec{o}_i^{\downarrow}} > (1+\epsilon)$, we have that $\vec{o}_i^{\downarrow} \notin \bigcup_{s \leq r} I_s$. So by Claim 9, we have $w_r^{\mathsf{avg}} \geq \widetilde{w}_i$. Hence, $g_{M,w^{\mathsf{est}}}(\gamma; \vec{c}_i^{\downarrow}) = w_r^{\mathsf{est}} \vec{c}_i^{\downarrow} \geq w_r^{\mathsf{avg}} \vec{c}_i^{\downarrow} \geq \widetilde{w}_i \vec{c}_i^{\downarrow}$, which contradicts our assumption that $\widetilde{w}_i \vec{c}_i^{\downarrow} > g_{M,w^{\mathsf{est}}}(\gamma; \vec{c}_i^{\downarrow})$.

Putting everything together, we have that $\sum_{i: \widetilde{w}_i \vec{c}_i^{\downarrow} > g_{M,w^{\mathsf{est}}}(\gamma; \vec{c}_i^{\downarrow})} \widetilde{w}_i \vec{c}_i^{\downarrow} \leq n\gamma\widetilde{w}_1 \cdot \frac{\epsilon M}{n} + \gamma(1+\epsilon) \sum_{i \in [n]} \widetilde{w}_i \vec{o}_i^{\downarrow}$, which proves the lemma. ◀

Finally, we show that $g_{M,w^{\mathsf{est}}}$ satisfies the analogue of Claim 3, which will be crucial in arguing that our algorithms and analysis from Section 4 carry over and allow us to solve, in an approximate sense, the k-median problem with the $\{g_{M,w^{\mathsf{est}}}(c_{ij})\}$ proxy costs.

▶ **Lemma 12.** *For any $\gamma \geq 1$, $M \geq 0$, and w^{est}, we have: (i) $g_{M,w^{\mathsf{est}}}(\gamma; x) \leq g_{M,w^{\mathsf{est}}}(\gamma; y)$ if $x \leq y$; and (ii) $3\max\{g_{M,w^{\mathsf{est}}}(\gamma; x), g_{M,w^{\mathsf{est}}}(\gamma; y), g_{M,w^{\mathsf{est}}}(\gamma; z)\} \geq g_{M,w^{\mathsf{est}}}(3\gamma; x+y+z)$ for any $x, y, z \geq 0$.*

3.2 Solving the k-median problem with the $\left\{g_{M,w^{\text{est}}}(c_{ij})\right\}$ proxy costs

We now work with a fixed guess M, w^{est}, and give an algorithm for finding a near-optimal k-median solution with the $\{g_{M,w^{\text{est}}}(c_{ij})\}$ proxy costs. Our algorithm and analysis will be quite similar to the one in Section 4. The primal and dual LPs we consider are the same as (P_B) and ($\text{D}_{\overline{B}}$), except that we replace all occurrences of $f_B(c_{ij})$ and $f_{\overline{B}}(c_{ij})$ with $g_{M,w^{\text{est}}}(c_{ij})$. Let $OPT_{M,w^{\text{est}}}$ denote the optimal value of this LP.

The primal-dual algorithm for a given center-cost λ (steps P1–P3 in Section 4) is unchanged. The analysis also is essentially identical, since, previously, we only relied on the fact that the proxy costs satisfy an approximate triangle inequality, which is also true here (Lemma 12). We state below the guarantee from the primal-dual algorithm slightly differently, in the form suggested by part (ii) of Lemma 12; the proof mimics the proof of Theorem 6.

▶ **Theorem 13.** *For any $\lambda \geq 0$, the primal-dual algorithm (P1)–(P3) returns a set T of centers, an assignment $i(j) \in T$ for every $j \in \mathcal{D}$, and a dual feasible solution (α, β, λ) such that $3\lambda|T| + \sum_j g_{M,w^{\text{est}}}(3; c_{i(j)j}) \leq 3\sum_j \alpha_j$.*

Given Theorem 13, we can use binary search on λ, to either obtain: (a) some λ such for which we return a solution T with $|T| = k$; or (b) $\lambda_1 < \lambda_2$ with $\lambda_2 - \lambda_1 < \frac{\epsilon\widetilde{w}_1 M}{n}$ such that letting T_1 and T_2 be the solutions returned for λ_1 and λ_2, we have $k_1 := |T_1| > k > k_2 := |T_2|$. In case (a), Theorem 13 implies that $\sum_j g_{M,w^{\text{est}}}(3; c_{i(j)j}) \leq 3OPT_{M,w^{\text{est}}}$. In case (b), we again extract a low-cost feasible solution from T_1 and T_2 by rounding the bipoint solution given by their convex combination. As before, $a, b \geq 0$ be such that $ak_1 + bk_2 = k$, $a + b = 1$. Let (α_1, β_1), (α_2, β_2) denote the dual solutions obtained for λ_1 and λ_2 respectively. Let $i_1(j)$ and $i_2(j)$ denote the centers to which j is assigned in T_1 and T_2 respectively. Let $d_{1,j} = g_{M,w^{\text{est}}}(3; c_{i_1(j)j})$ and $d_{2,j} = g_{M,w^{\text{est}}}(3; c_{i_2(j)j})$. Let $C_1 := \sum_j d_{1,j}$ and $C_2 := \sum_j d_{2,j}$. Similar to before, we have $aC_1 + bC_2 \leq 3OPT_{M,w^{\text{est}}} + 3\epsilon\widetilde{w}_1 M$. The procedure for rounding this bipoint solution requires only minor changes to steps B1, B2 in Section 4.

Rounding the bipoint solution. If $b \geq 1/3$, then T_2 yields a feasible solution with $\sum_j g_{M,w^{\text{est}}}(3; c_{i_2(j)j}) = C_2 \leq 9OPT_{M,w^{\text{est}}} + 9\epsilon\widetilde{w}_1 M$. So suppose $a \geq 2/3$.

G1. **Clustering.** We match facilities in T_2 with a subset of facilities in T_1 as follows. Initialize $\mathcal{D}' \leftarrow \mathcal{D}$, $A \leftarrow \emptyset$, and $M \leftarrow \emptyset$. We repeatedly pick the client $j \in \mathcal{D}'$ with minimum $\max\{d_{1,j}, d_{2,j}\}$ value, and add j to A. (***This is the only change, compared to step B1.***) We add the tuple $(i_1(j), i_2(j))$ to M, remove from $\mathcal{D}'$ all clients k (including j) such that $i_1(k) = i_1(j)$ or $i_2(k) = i_2(j)$, and set $\sigma(k) = j$ for all such clients. Let $M_1 = M$ denote the matching when $\mathcal{D}' = \emptyset$. Next, for each unmatched $i \in T_2$, we pick an arbitrary unmatched facility $i' \in T_1$, and add (i', i) to M. Let F_1 be the set of T_1-facilities that are matched, and $S := \{j \in \mathcal{D} : i_1(j) \in F_1\}$. Note that $|F_1| = |M| = k_2$.

G2. **Opening facilities.** This is almost identical to step B2, except that we decide which facilities to open by now solving the following LP.

$$\min \sum_{j \in S}\big(\theta d_{1,j} + (1-\theta)d_{2,j}\big) + \sum_{k \notin S}\big(z_{i_1(k)}d_{1,k} + (1 - z_{i_1(k)}) \cdot 3\max\{d_{1,k}, d_{2,k}\}\big) \quad \text{(GR-P)}$$

$$\text{s.t.} \sum_{i \in T_1 \setminus F} z_i \leq k - k_2, \qquad \theta \in [0,1], \quad z_i \in [0,1] \ \ \forall i \in T_1 \setminus F.$$

Let $(\tilde{\theta}, \tilde{z})$ be an optimal integral solution to (GR-P). If $\tilde{\theta} = 1$, we open all facilities in F_1, and otherwise, all facilities in T_2. We also open the facilities from $T_1 \setminus F_1$ for which $\tilde{z}_i = 1$.

To analyze this, we first show that setting $\theta = a$, $z_i = a$ for all $i \in T_1 \setminus F_1$ yields a feasible solution to (GR-P) of objective value at most $3(aC_1 + bC_2)$. We have $\sum_{i \in T_1 \setminus F_1} z_i = a(k_1 - k_2) = k - k_2$. Every $j \in S$ contributes $ad_{1,j} + bd_{2,j}$ to the objective value of (GR-P). Consider $k \notin S$. Its contribution to the objective value of (GR-P) is

$$ad_{1,k} + 3b \max\{d_{1,k}, d_{2,k}\} = \max\{(a + 3b)d_{1,k}, ad_{1,k} + 3bd_{2,k}\} \leq 3(ad_{1,k} + bd_{2,k})$$

where the inequality follows since $a + 3b \leq 3a$ when $a \geq 2/3$. Thus, for every $j \in \mathcal{D}$, its contribution to the objective value of (GR-P) is at most thrice its contribution to $aC_1 + bC_2$.

Suppose $\vec{c}$ is the assignment-cost vector resulting from $(\tilde{\theta}, \tilde{z})$. We show that $\sum_j g_{M,w^{\mathsf{est}}}(9; \vec{c}_j)$ is at most the objective value of $(\tilde{\theta}, \tilde{z})$ under (GR-P). For every $k \in S$, we have $g_{M,w^{\mathsf{est}}}(9; \vec{c}_k) \leq g_{M,w^{\mathsf{est}}}(3; \vec{c}_k) \leq \tilde{\theta} d_{1,k} + (1 - \tilde{\theta}) d_{2,k}$. Now consider $k \notin S$ with $\sigma(k) = j$, so $\max\{d_{1,j}, d_{2,j}\} \leq \max\{d_{1,k}, d_{2,k}\}$. If $\tilde{z}_{i_1(k)} = 1$, then $g_{M,w^{\mathsf{est}}}(9; \vec{c}_k) \leq g_{M,w^{\mathsf{est}}}(3; \vec{c}_k) \leq d_{1,k}$. Otherwise, $\vec{c}_k \leq c_{i_2(k)k} + c_{i_1(j)j} + c_{i_2(j)j}$, and so by Lemma 12, we have

$$\begin{aligned} g_{M,w^{\mathsf{est}}}(9; \vec{c}_k) &\leq g_{M,w^{\mathsf{est}}}(9; c_{i_2(k)k} + c_{i_1(j)j} + c_{i_2(j)j}) \\ &\leq 3 \max\{g_{M,w^{\mathsf{est}}}(3; c_{i_2(k)}), g_{M,w^{\mathsf{est}}}(3; c_{i_1(j)j}), g_{M,w^{\mathsf{est}}}(3; c_{i_2(j)j})\} \leq 3 \max\{d_{1,k}, d_{2,k}\}. \end{aligned}$$

So in every case, $g_{M,w^{\mathsf{est}}}(9; \vec{c}_k)$ is bounded by the contribution of k to the objective value of $(\tilde{\theta}, \tilde{z})$. Thus, we have proved the following theorem.

▶ **Theorem 14.** *For any $M \geq 0$, w^{est}, we can obtain a solution opening k centers whose assignment-cost vector $\vec{c}$ satisfies $\sum_j g_{M,w^{\mathsf{est}}}(9; \vec{c}_j) \leq 9 OPT_{M,w^{\mathsf{est}}} + 9\epsilon \widetilde{w}_1 M$.*

Proof of Theorem 7. The proof follows by combining Theorem 14, Lemmas 10 and 11, and Claim 8. Recall that $\vec{o}^{\downarrow}$ is the assignment-cost vector corresponding to an optimal solution with coordinates sorted in non-increasing order, and $opt = \sum_{i=1}^{n} w_i \vec{o}_i^{\downarrow}$ is the optimal cost.

There are only n^2 choices for M, and $O\big((\frac{n}{\epsilon})^{1/\epsilon}\big)$ choices for w^{est}, so we may assume that in polynomial time, we have obtained $M = \vec{o}_1^{\downarrow}$, and w_r^{est}s satisfying $w_r^{\mathsf{avg}} \leq w_r^{\mathsf{est}} \leq (1+\epsilon) w_r^{\mathsf{avg}}$ for all $r \in \{0, \ldots, T\}$. By Lemma 10, we know that $OPT_{M,w^{\mathsf{est}}} \leq (1+\epsilon)^2 cost(\widetilde{w}; \vec{o}^{\downarrow}) \leq (1+\epsilon)^2 opt$. Let $\vec{c}$ be the assignment-cost vector of the solution returned by Theorem 14 for this M, w^{est}. Combining Theorem 14, Lemma 11, and Claim 8, we obtain that

$$\begin{aligned} (1 - \epsilon) cost(w; \vec{c}) \leq cost(\widetilde{w}; \vec{c}) &\leq \big(9 OPT_{M,w^{\mathsf{est}}} + 9\epsilon \widetilde{w}_1 M\big) + 9(1 + \epsilon) cost(\widetilde{w}; \vec{o}^{\downarrow}) + 9\epsilon \widetilde{w}_1 M \\ &\leq 9(1 + \epsilon)^2 opt + 9 opt + O(\epsilon) opt = \big(18 + O(\epsilon)\big) opt. \end{aligned}$$

4 Conclusions and discussion

We have described algorithms achieving approximation guarantees of $12 + \epsilon$ and $18 + \epsilon$ for the ℓ-centrum and ordered k-median problems. Our algorithms are combinatorial, utilizing the primal-dual schema and Lagrangian relaxation, and improve upon the algorithms in [3], both in terms of approximation factors and simplicity of analysis.

One interesting research direction suggested by our work is to investigate the ordered-median and ℓ-centrum (i.e., ordered median with $\{0, 1\}$-weights) versions of other optimization problems. In further work, we have been able to develop a general framework for devising algorithms for ordered-median problems. Our framework also yields improved guarantees for the ℓ-centrum and ordered k-median problems studied here. We obtain analogous improvements for ordered k-median. We defer details to a forthcoming manuscript.

References

1. S. Alamdari and D. Shmoys. A bicriteria approximation algorithm for the k-center and k-median problems. In *Proceedings of WAOA*, 2017.
2. A. Aouad and D. Segev. The ordered k-median problem: surrogate models and approximation algorithms. In submission. Available at `https://pdfs.semanticscholar.org/39e8/391a9e918a2ba1eb21c1c2dbb52d3b1f0de1.pdf`, 2017.
3. J. Byrka, K. Sornat, and J. Spoerhase. Constant-factor approximation for ordered k-median. To appear in *Proceedings of STOC*, 2018. Available at `https://arXiv.org/abs/1711.01972`, Nov 6, 2017.
4. Jarosław Byrka, Thomas Pensyl, Bartosz Rybicki, Aravind Srinivasan, and Khoa Trinh. An improved approximation for k-median, and positive correlation in budgeted optimization. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 737–756. SIAM, 2015.
5. D. Chakrabarty and C. Swamy. Interpolating between k-median and k-center: Approximation algorithms for ordered k-median. `https://arxiv.org/abs/1711.08715`, Nov 23, 2017.
6. M. Charikar, S. Guha, É. Tardos, and D. B. Shmoys. A constant-factor approximation algorithm for the k-median problem. *Journal of Computer and System Sciences*, 65(1):129–149, 2002.
7. M. Charikar and S. Li. A dependent LP-rounding approach for the k-median problem. In *Proceedings of the 39th ICALP*, 2012.
8. Zvi Drezner and Stefan Nickel. Solving the ordered one-median problem in the plane. *European Journal of Operational Research*, 195(1):46–61, 2009.
9. R. L. Francis, T. J. Lowe, and A. Tamir. Aggregation error bounds for a class of location models,. *Operations Research*, 48:294–307, 2000.
10. 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.
11. 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.
12. J. N. Hooker, R. S. Garfinkel, and C. K. Chen. Finite dominating sets for network location problems. *Operations Research*, 39:100–118, 1991.
13. K. Jain and V. V. Vazirani. Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and Lagrangian relaxation. *Journal of the ACM*, 48(2):274–296, 2001.
14. G. Laporte, S. Nickel, and F. S. da Gama. *Location Science*. Springer, 2015.
15. Shi Li and Ola Svensson. Approximating k-median via pseudo-approximation. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing*, pages 901–910. ACM, 2013.
16. Pitu B Mirchandani and Richard L Francis. *Discrete location theory*. Wiley-Interscience, 1990.
17. S. Nickel and J. Puerto. A unified approach to network location problems. *Networks*, 34:283–290, 1999.
18. S. Nickel and J. Puerto. *Location Theory: A Unified Approach*. Springer Science & Business Media, 2005.
19. Justo Puerto and Antonio M. Rodríguez-Chía. On the exponential cardinality of FDS for the ordered p-median problem. *Oper. Res. Lett.*, 33(6):641–651, 2005.
20. Justo Puerto, Antonio M. Rodríguez-Chía, and Arie Tamir. Revisiting k-sum optimization. *Math. Program.*, 165(2):579–604, 2017.
21. D. B. Shmoys, É. Tardos, and K. Aardal. Approximation algorithms for facility location problems. In *Proceedings of the 29th STOC*, pages 265–274, 1997.

22 David B Shmoys. The design and analysis of approximation algorithms. *Trends in Optimization: American Mathematical Society Short Course, January 5-6, 2004, Phoeniz, Arizona*, 61:85, 2004.

23 A. Tamir. The k-centrum multi-facility location problem. *Discrete Applied Mathematics*, 109(3):293–307, 2001.