



To Close Is Easier Than To Open: Dual Parameterization To k -Median

Jarosław Byrka¹, Szymon Dudycz¹, Pasin Manurangsi², Jan Marcinkowski¹,
and Michał Włodarczyk³(✉)

¹ University of Wrocław, Wrocław, Poland

{jby, szymon.dudycz, jan.marcinkowski}@cs.uni.wroc.pl

² Google Research, Mountain View, USA

pasin@google.com

³ Eindhoven University of Technology, Eindhoven, The Netherlands

m.wlodarczyk@tue.nl

Abstract. The k -MEDIAN problem is one of the well-known optimization problems that formalize the task of data clustering. Here, we are given sets of facilities F and clients C , and the goal is to open k facilities from the set F , which provides the best division into clusters, that is, the sum of distances from each client to the closest open facility is minimized. In the CAPACITATED k -MEDIAN, the facilities are also assigned capacities specifying how many clients can be served by each facility.

Both problems have been extensively studied from the perspective of approximation algorithms. Recently, several surprising results have come from the area of parameterized complexity, which provided better approximation factors via algorithms with running times of the form $f(k) \cdot \text{poly}(n)$. In this work, we extend this line of research by studying a different choice of parameterization. We consider the parameter $\ell = |F| - k$, that is, the number of facilities that remain closed. It turns out that such a parameterization reveals yet another behavior of k -MEDIAN. We observe that the problem is W[1]-hard but it admits a parameterized approximation scheme. Namely, we present an algorithm with running time $2^{\mathcal{O}(\ell \log(\ell/\varepsilon))} \cdot \text{poly}(n)$ that achieves a $(1 + \varepsilon)$ -approximation. On the other hand, we show that under the assumption of Gap Exponential Time Hypothesis, one cannot extend this result to the capacitated version of the problem.

1 Introduction

Recent years have brought many surprising algorithmic results originating from the intersection of the areas of approximation algorithms and parameterized complexity. It turns out that the combination of techniques from these theories can be very fruitful and a new research area has emerged, devoted to studying *parameterized approximation algorithms*. The main goal in this area it to design

Part of this work was done while the third and the fifth author were visiting University of Wrocław. The fifth author was supported by the Foundation for Polish Science (FNP).

© Springer Nature Switzerland AG 2021

C. Kallamanis and A. Levin (Eds.): WAOA 2020, LNCS 12806, pp. 113–126, 2021.

https://doi.org/10.1007/978-3-030-80879-2_8

an algorithm processing an instance (I, k) in time $f(k) \cdot |I|^{\mathcal{O}(1)}$, where f is some computable function, and producing an approximate solution to the optimization problem in question. Such algorithms, called *FPT approximations*, are particularly interesting in the case of problems for which (1) we fail to make progress on improving the approximation factors in polynomial time, and (2) there are significant obstacles for obtaining exact parameterized algorithms. Some results of this kind are FPT approximations for k -CUT [22], DIRECTED ODD CYCLE TRANSVERSAL [27], and PLANAR STEINER NETWORK [10]. A good introduction to this area can be found in the survey [20].

One problem that has recently enjoyed a significant progress in this direction is the famous k -MEDIAN problem. Here, we are given a set F of facilities, a set C of clients, a metric d over $F \cup C$ and an upper bound k on the number of facilities we can open. A solution is a set $S \subseteq F$ of at most k open facilities and a connection assignment $\phi : C \rightarrow S$ of clients to the open facilities. The goal is to find a solution that minimizes the connection cost $\sum_{c \in C} d(c, \phi(c))$. The problem can be approximated in polynomial time up to a constant factor [3, 9] with the currently best approximation factor being $(2.675 + \varepsilon)$ [5]. On the other hand, we cannot hope for a polynomial-time $(1 + 2/e - \varepsilon)$ -approximation, since it would entail $P = NP$ [21]. Therefore, there is a gap in our understanding of the optimal approximability of k -MEDIAN.

Surprisingly, the situation becomes simpler if we consider parameterized algorithms, with k as the natural choice of parameterization. Such a parameterized problem is $W[2]$ -hard [1] so it is unlikely to admit an exact algorithm with running time of the form $f(k) \cdot n^{\mathcal{O}(1)}$, where n is the size of an instance. However, Cohen-Addad et al. [12] have obtained an algorithm with approximation factor $(1 + 2/e + \varepsilon)$ and running time¹ $2^{\mathcal{O}(k \log k)} \cdot n^{\mathcal{O}(1)}$. This result is essentially tight, as the existence of an FPT-approximation with factor $(1 + 2/e - \varepsilon)$ would contradict the *Gap Exponential Time Hypothesis*² (Gap-ETH) [12]. The mentioned hardness result has also ruled out running time of the form $f(k) \cdot n^{g(k)}$, where $g = k^{\text{poly}(1/\varepsilon)}$. This lower bound has been later strengthened: under Gap-ETH no algorithm with running time $f(k) \cdot n^{o(k)}$ can achieve approximation factor $(1 + 2/e - \varepsilon)$ [28].

The parameterized approach brought also a breakthrough to the understanding of CAPACITATED k -MEDIAN. In this setting, each facility f is associated with a capacity $u_f \in \mathbb{Z}_{\geq 0}$ and the connection assignment ϕ must satisfy $|\phi^{-1}(f)| \leq u_f$ for every facility $f \in S$. The best known polynomial-time approximation for CAPACITATED k -MEDIAN is burdened with a factor $\mathcal{O}(\log k)$ [1, 8] and relies on the generic technique of metric tree embeddings with expected logarithmic distortion [19]. All the known constant-factor approximations violate either the number of facilities or the capacities. Li has provided such an algorithm by

¹ We omit the dependency on ε in the running time except for approximation schemes.

² The Gap Exponential Time Hypothesis [18, 29] states that, for some constant $\gamma > 0$, there is no $2^{o(n)}$ -time algorithm that can, given a 3SAT instance, distinguish between (1) the instance is fully satisfiable or (2) any assignment to the instance violates at least γ fraction of the clauses.

opening $(1 + \varepsilon) \cdot k$ facilities [24, 25]. Afterwards analogous results, but violating the capacities by a factor of $(1 + \varepsilon)$ were also obtained [6, 17]. This is in contrast with other capacitated clustering problems such as FACILITY LOCATION or k -CENTER, for which constant factor approximation algorithms have been constructed [15, 23]. However, no superconstant lower bound for CAPACITATED k -MEDIAN is known.

When it comes to parameterized algorithms, Adamczyk et al. [1] have presented a $(7 + \varepsilon)$ -approximation algorithm with running time $2^{\mathcal{O}(k \log k)} \cdot n^{\mathcal{O}(1)}$ for CAPACITATED k -MEDIAN. Xu et al. [31] proposed a similar algorithm for the related CAPACITATED k -MEANS problem, where one minimizes the sum of squares of distances. These results have been improved by Cohen-Addad and Li [13], who obtained factor $(3 + \varepsilon)$ for CAPACITATED k -MEDIAN and $(9 + \varepsilon)$ for CAPACITATED k -MEANS, within the same running time.

Our Contribution. In this work, we study a different choice of parameterization for k -MEDIAN. Whereas k is the number of facilities to open, we consider the dual parameter $\ell = |F| - k$: the number of facilities to be closed. We refer to this problem as CO- ℓ -MEDIAN in order to avoid ambiguity. Note that even though this is the same task from the perspective of polynomial-time algorithms, it is a different problem when seen through the lens of parameterized complexity. First, we observe that CO- ℓ -MEDIAN is W[1]-hard (Theorem 3), which motivates the study of approximation algorithms also for this choice of parameterization. It turns out that switching to the dual parameterization changes the approximability status dramatically and we can obtain an arbitrarily good approximation factor. More precisely, we present an efficient parameterized approximation scheme (EPAS), i.e., $(1 + \varepsilon)$ -approximation with running time of the form $f(\ell, \varepsilon) \cdot n^{\mathcal{O}(1)}$. This constitutes our main result.

Theorem 1. *The CO- ℓ -MEDIAN problem admits a deterministic $(1 + \varepsilon)$ -approximation algorithm running in time $2^{\mathcal{O}(\ell \log(\ell/\varepsilon))} \cdot n^{\mathcal{O}(1)}$ for any constant $\varepsilon > 0$.*

We obtain this result by combining the technique of color-coding from the FPT theory with a greedy approach common in the design of approximation algorithms. The running time becomes polynomial whenever we want to open all but $\mathcal{O}\left(\frac{\log n}{\log \log n}\right)$ facilities. To the best of our knowledge, this is the first non-trivial setting with general metric space which admits an approximation scheme.

A natural question arises about the behavior of the capacitated version of the problem in this setting, referred to as CAPACITATED CO- ℓ -MEDIAN. Both in polynomial-time regime or when parameterized by k , there is no evidence that the capacitated problem is any harder and the gap between the approximation factors might be just a result of our lack of understanding. Somehow surprisingly, for the dual parameterization ℓ we are able to show a clear separation between the capacitated and uncapacitated case. Namely, we present a reduction from the MAX k -COVERAGE problem which entails the same approximation lower bound as for the uncapacitated problem parameterized by k .

Theorem 2. *Assuming Gap-ETH, there is no $f(\ell) \cdot n^{o(\ell)}$ -time algorithm that can approximate CAPACITATED CO- ℓ -MEDIAN to within a factor of $(1 + 2/e - \epsilon)$ for any function f and any constant $\epsilon > 0$.*

Related Work. A simple example of dual parameterization is given by k -INDEPENDENT SET and ℓ -VERTEX COVER. From the perspective of polynomial-time algorithms, these problems are equivalent (by setting $\ell = |V(G)| - k$), but they differ greatly when analyzed as parameterized problems: the first one is $W[1]$ -hard while the latter is FPT and admits a polynomial kernel [14]. Another example originates from the early work on k -DOMINATING SET, which is a basic $W[2]$ -complete problem. When parameterized by $\ell = |V(G)| - k$, the problem is known as ℓ -NONBLOCKER. This name can be interpreted as a task of choosing ℓ vertices so that none is blocked by the others, i.e., each chosen vertex has a neighbor which has not been chosen. Under this parameterization, the problem is FPT and admits a linear kernel [16]. The best known running time for ℓ -NONBLOCKER is $1.96^\ell \cdot n^{O(1)}$ [30]. It is worth noting that ℓ -NONBLOCKER is a special case of CO- ℓ -MEDIAN with a graph metric and $F = C = V(G)$, however this analogy works only in a non-approximate setting.

The Gap Exponential Time Hypothesis was employed for proving parameterized inapproximability by Chalermsook et al. [7], who presented hardness results for k -CLIQUE, k -DOMINATING SET, and DENSEST k -SUBGRAPH. It was later used to obtain lower bounds for DIRECTED ODD CYCLE TRANSVERSAL [27], DIRECTED STEINER NETWORK [10], PLANAR STEINER ORIENTATION [11], and UNIQUE SET COVER [28], among others. Moreover, Gap-ETH turned out to be a sufficient assumption to rule out the existence of an FPT algorithm for k -EVEN SET [4].

2 Preliminaries

Parameterized Complexity and Reductions. A parameterized problem instance is created by associating an input instance with an integer parameter k . We say that a problem is *fixed parameter tractable* (FPT) if it admits an algorithm solving an instance (I, k) in time $f(k) \cdot |I|^{O(1)}$, where f is some computable function. Such an algorithm we shall call an *FPT algorithm*.

To show that a problem is unlikely to be FPT, we use *parameterized reductions* analogous to those employed in the classical complexity theory (see [14]). Here, the concept of W -hardness replaces the one of NP-hardness, and we need not only to construct an equivalent instance in time $f(k) \cdot |I|^{O(1)}$, but also to ensure that the value of the parameter in the new instance depends only on the value of the parameter in the original instance. In contrast to the NP-hardness theory, there is a hierarchy of classes $FPT = W[0] \subseteq W[1] \subseteq W[2] \subseteq \dots$ and these containments are believed to be strict. If there exists a parameterized reduction transforming a $W[t]$ -hard problem to another problem Π , then the problem Π is $W[t]$ -hard as well. If a parameterized reduction transforms parameter linearly,

i.e., maps an instance (I_1, k) to $(I_2, \mathcal{O}(k))$, then it also preserves running time of the form $f(k) \cdot |I|^{\mathcal{O}(k)}$.

In order to prove hardness of parameterized approximation, we use parameterized reductions between *promise problems*. Suppose we are given an instance (I_1, k_1) of a minimization problem with a promise that the answer is at most D_1 and we want to find a solution of value at most $\alpha \cdot D_1$. Then a reduction should map (I_1, k_1) to such an instance (I_2, k_2) so that the answer to it is at most D_2 and any solution to (I_2, k_2) of value at most $\alpha \cdot D_2$ can be transformed in time $f(k_1) \cdot |I_1|^{\mathcal{O}(1)}$ to a solution to (I_1, k_1) of value at most $\alpha \cdot D_1$. If an FPT α -approximation exists for the latter problem, then it exists also for the first one. Again, if we have $k_2 = \mathcal{O}(k_1)$, then this relation holds also for algorithms with running time of the form $f(k) \cdot |I|^{\mathcal{O}(k)}$.

Problem Definitions. Below we formally introduce the main studied problem and the problems employed in reductions.

(CAPACITATED) CO- ℓ -MEDIAN

Parameter: ℓ

Input: set of facilities F , set of clients C , metric d over $F \cup C$, sequence of capacities $u_f \in \mathbb{Z}_{\geq 0}$, integer ℓ

Task: find a set $S \subseteq F$ of at most $|F| - \ell$ facilities and a connection assignment $\phi : C \rightarrow F \setminus S$ that satisfies $|\phi^{-1}(f)| \leq u_f$ for all $f \in F \setminus S$, and minimizes $\sum_{c \in C} d(c, \phi(c))$

A metric $d : (F \cup C) \times (F \cup C) \rightarrow \mathbb{R}_{\geq 0}$ is a symmetric function that obeys the triangle inequality $d(x, y) + d(y, z) \geq d(x, z)$ and satisfies $d(x, x) = 0$. In the uncapacitated version we assume that all capacities are equal $|C|$, so any assignment $\phi : C \rightarrow F \setminus S$ is valid. In the approximate version of CAPACITATED CO- ℓ -MEDIAN we treat the capacity condition $|\phi^{-1}(f)| \leq u_f$ as a hard constraint and we allow only the connection cost $\sum_{c \in C} d(c, \phi(c))$ to be larger than the optimum.

k -INDEPENDENT SET

Parameter: k

Input: graph $G = (V, E)$, integer k

Task: decide whether there exists a set $S \subseteq V(G)$ of size k such that for all pairs $u, v \in S$ we have $uv \notin E(G)$

MAX k -COVERAGE

Parameter: k

Input: universe U , family of subsets $T_1, \dots, T_n \subseteq U$, integer k

Task: find k subsets T_{i_1}, \dots, T_{i_k} that maximizes $|T_{i_1} \cup \dots \cup T_{i_k}|$

3 Uncapacitated CO- ℓ -MEDIAN

We begin with a simple reduction, showing that the exact problem remains hard under the dual parameterization.

Theorem 3. *The CO- ℓ -MEDIAN problem is $W[1]$ -hard.*

Proof. We reduce from ℓ -INDEPENDENT SET, which is W[1]-hard. We transform a given graph G into a CO- ℓ -MEDIAN instance by setting $F = V(G)$, and placing a client in the middle of each edge. The distance from a client to both endpoints of its edge is 1 and the shortest paths of such subdivided graph induce the metric d .

If we did not close any facilities, the cost of serving all clients would equal $|E(G)|$. The same holds if each client has an open facility within distance 1, so the set of closed facilities forms an independent set of vertices in G . On the other hand, if we close a set of facilities containing two endpoints of a single edge then the cost increases. Therefore the answer to the created instance is $|E(G)|$ if and only if G contains an independent set of size ℓ .

We move on to designing a parameterized approximation scheme for CO- ℓ -MEDIAN. We use notation $d(c, S)$ for the minimum distance between c and any element of the set S . In the uncapacitated setting the connection assignment ϕ_S is unique for a given set of closed facilities S : each client is assigned to the closest facility outside S . Whenever we consider a solution set $S \subseteq F$, we mean that this is the set of closed facilities and denote $cost(S) = \sum_{c \in C} d(c, F \setminus S)$. We define $V(f)$ to be the *Voronoi cell* of facility f , i.e., the set of clients for which f is the closest facility. We can break ties arbitrarily and for the sake of disambiguation we assume an ordering on F and whenever two distances are equal we choose the facility that comes first in the ordering.

Let $C(f)$ denote the cost of the cell $V(f)$, i.e., $\sum_{c \in V(f)} d(c, f)$. For a solution S and $f \in S, g \notin S$, we define $C(S, f, g) = \sum d(c, g)$ over $\{c \in V(f) \mid \phi_S(c) = g\}$, that is, the sum of connections of clients that switched from f to g . Note that as long as $f \in F$ remains open, there is no need to change connections of the clients in $V(f)$. We can express the difference of connection costs after closing S as

$$\Delta(S) = \sum_{f \in S} \sum_{c \in V(f)} d(c, F \setminus S) - \sum_{f \in S} C(f) = \sum_{f \in S} \sum_{g \in F \setminus S} C(S, f, g) - \sum_{f \in S} C(f).$$

We have $cost(S) = \sum_{f \in F} C(f) + \Delta(S)$, therefore the optimal solution closes set S of size ℓ minimizing $\Delta(S)$.

The crucial observation is that any small set of closed facilities S can be associated with a small set of open facilities that are relevant for serving the clients from $\bigcup_{f \in S} V(f)$. Intuitively, if $C(S, f, g) = \mathcal{O}(\frac{\varepsilon}{\ell^2}) \cdot cost(S)$ for all $f \in S$, then we can afford replacing ℓ such facilities g with others that are not too far away.

Definition 1. *The ε -support of a solution $S \subseteq F, |S| = \ell$, referred to as ε -SUPP(S), is the set of all open facilities g (i.e., $g \notin S$) satisfying one of the following conditions:*

1. *there is $f \in S$ such that g minimizes distance $d(f, g)$ among all open facilities,*
2. *there is $f \in S$ such that $C(S, f, g) > \frac{\varepsilon}{6\ell^2} \cdot cost(S)$.*

We break ties in condition (1) according to the same rule as in the definition of $V(f)$, so there is a single g satisfying condition (1) for each f .

Lemma 1. For a solution S of size ℓ , we have $|\varepsilon\text{-SUPP}(S)| \leq 6 \cdot \ell^3 / \varepsilon + \ell$.

Proof. We get at most ℓ facilities from condition (1). Since the sets of clients being served by different $g \in F \setminus S$ are disjoint and $\sum_{g \in F} C(S, f, g) \leq \text{cost}(S)$, we obtain at most $6 \cdot \ell^3 / \varepsilon$ facilities from condition (2).

Even though we will not compute the set $\varepsilon\text{-SUPP}(\text{Opt})$ directly, we are going to work with partitions $F = A \uplus B$, such that $\text{Opt} \subseteq A$ and $\varepsilon\text{-SUPP}(\text{Opt}) \subseteq B$. Such a partition already gives us a valuable hint. By looking at each facility $f \in A$ separately, we can deduce that if $f \in \text{Opt}$ and some other facility g belongs to A (so it cannot belong to $\varepsilon\text{-SUPP}(\text{Opt})$) then in some cases g must also belong to Opt . More precisely, if $g \in A$ is closer to f than the closest facility in B , then g must be closed, as otherwise it would violate condition (1). Furthermore, suppose that $h \in A$ serves clients from $V(f)$ (assuming f is closed) of total cost at least $\frac{\varepsilon}{6\ell^2} \cdot \text{cost}(S)$. If we keep h open and close some other facilities, this relation is preserved and having h in A violates condition (2). We formalize this idea with the notion of *required sets*, given by the following procedure, supplied additionally with a real number D , which can be regarded as the guessed value of $\text{cost}(\text{Opt})$.

Algorithm 1. COMPUTE-REQUIRED-SET($A, B, f, \varepsilon, \ell, D$) (assume $f \in A$ and $A \cap B = \emptyset$)

```

1:  $s_f \leftarrow \min_{y \in B} d(f, y)$ 
2:  $R_f \leftarrow \{g \in A \mid d(f, g) < s_f\}$  (including  $f$ )
3: while  $\exists g \in A : C(R_f, f, g) > \frac{\varepsilon}{3\ell^2} \cdot D$  do
4:    $R_f \leftarrow R_f \cup \{g\}$ 
5: end while
6: return  $R_f$ 

```

Lemma 2. Let $\text{Opt} \subseteq F$ be the optimal solution. Suppose $F = A \uplus B$, $f \in \text{Opt} \subseteq A$, $\varepsilon\text{-SUPP}(\text{Opt}) \subseteq B$, and $\text{cost}(\text{Opt}) \leq 2D$. Then the set R_f returned by the routine COMPUTE-REQUIRED-SET($A, B, f, \varepsilon, \ell, D$) satisfies $R_f \subseteq \text{Opt}$.

Proof. Let y_f be the facility in B that is closest to f . Due to condition (1) in Definition 1, all facilities $g \in A$ satisfying $d(f, g) < d(f, y_f)$ must be closed in the optimal solution, so we initially add them to R_f . We keep invariant $R_f \subseteq \text{Opt}$, so for any $g \in F \setminus \text{Opt}$ it holds that $C(\text{Opt}, f, g) \geq C(R_f, f, g)$. Whenever there is $g \in A$ satisfying $C(R_f, f, g) > \frac{\varepsilon}{3\ell^2} \cdot D$, we get

$$C(\text{Opt}, f, g) \geq C(R_f, f, g) > \frac{\varepsilon}{3\ell^2} \cdot D \geq \frac{\varepsilon}{6\ell^2} \cdot \text{cost}(\text{Opt}).$$

Since g does not belong to $\varepsilon\text{-SUPP}(\text{Opt}) \subseteq B$, then by condition (2) it must be closed. Hence, adding g to R_f preserves the invariant.

Before proving the main technical lemma, we need one more simple observation, in which we exploit the fact that the function d is indeed a metric.

Lemma 3. *Suppose $c \in V(f_0)$ and $d(f_0, f_1) \leq d(f_0, f_2)$. Then $d(c, f_1) \leq 3 \cdot d(c, f_2)$.*

Proof. An illustration is given in Fig. 1. Since c belongs to the Voronoi cell of f_0 , we have $d(c, f_0) \leq d(c, f_2)$. By the triangle inequality

$$\begin{aligned} d(c, f_1) &\leq d(c, f_0) + d(f_0, f_1) \leq d(c, f_0) + d(f_0, f_2) \\ &\leq d(c, f_0) + d(c, f_0) + d(c, f_2) \leq 3 \cdot d(c, f_2). \end{aligned}$$




Fig. 1. An example of a Voronoi diagram with squares representing facilities and dots being clients. Lemma 3 states that even if $d(c, f_1) > d(c, f_2)$ for $c \in V(f_0)$ and $d(f_0, f_1) \leq d(f_0, f_2)$, then $d(c, f_1)$ cannot be larger than $3 \cdot d(c, f_2)$.

Lemma 4. *Suppose we are given a partition $F = A \uplus B$, such that $Opt \subseteq A$, $\varepsilon\text{-SUPP}(Opt) \subseteq B$, and a number $D \in \mathbb{R}_{>0}$, such that $\text{cost}(Opt) \in [D, 2D]$. Then we can find a solution $S \subseteq A$, such that $\text{cost}(S) \leq (1 + \varepsilon) \cdot \text{cost}(Opt)$, in polynomial time.*

Proof. We compute the set $R_f = \text{COMPUTE-REQUIRED-SET}(A, B, f, \varepsilon, \ell, D)$ for each facility $f \in A$. The subroutine from Algorithm 1 clearly runs in polynomial time. Furthermore, for each $f \in A$ we compute its *marginal cost* of closing

$$m_f = \sum_{c \in V(f)} d(c, F \setminus R_f) - C(f).$$

If $|R_f| > \ell$ then f cannot belong to any solution consistent with the partition (A, B) and in this case we set $m_f = \infty$. Since the marginal cost depends only on f , we can greedily choose ℓ facilities from A that minimize m_f – we refer to this set as S .

We first argue that $\sum_{f \in F} C(f) + \sum_{f \in S} m_f$ is at most the cost of the optimal solution. By greedy choice we have that $\sum_{f \in S} m_f \leq \sum_{f \in Opt} m_f$. We

have assumed $cost(Opt) \leq 2D$ so by Lemma 2 we get that if $f \in Opt$, then $R_f \subseteq Opt$. The set of facilities $F \setminus Opt$ that can serve clients from $V(f)$ is a subset of $F \setminus R_f$ and the distances can only increase, thus for $f \in Opt$ we have $m_f \leq \sum_{c \in V(f)} d(c, F \setminus Opt) - C(f)$. We conclude that $\sum_{f \in F} C(f) + \sum_{f \in S} m_f$ is upper bounded by

$$\sum_{f \in F} C(f) + \sum_{f \in Opt} \sum_{c \in V(f)} d(c, F \setminus Opt) - \sum_{f \in Opt} C(f) = cost(Opt). \tag{1}$$

The second argument is that after switching benchmark from the marginal cost to the true cost of closing S , we will additionally pay at most εD . These quantities differ when for a facility $f \in S$ we have ‘connected’ some clients from $V(f)$ to $g \in S \setminus R_f$ when computing m_f . More precisely, we want to show that for each $f \in S$ we have

$$\sum_{c \in V(f)} d(c, F \setminus S) \leq \sum_{c \in V(f)} d(c, F \setminus R_f) + \frac{\varepsilon D}{\ell}. \tag{2}$$

By the construction of R_f , whenever $g \in S \setminus R_f$ we are guaranteed that there exists a facility $y \in B$ such that $d(f, g) \geq d(f, y)$ and, moreover, $C(R_f, f, g) \leq \frac{\varepsilon}{3\ell^2} \cdot D$. We can reroute all such clients c to the closest open facility and we know it is not further than $d(c, y)$. By Lemma 3 we know that $d(c, y) \leq 3 \cdot d(c, g)$ so rerouting those clients costs at most $\frac{\varepsilon}{\ell^2} \cdot D$. Since there are at most ℓ such facilities $g \in S \setminus R_f$, we have proved Formula (2). Combining this with bound from (1) implies that $cost(S) \leq cost(Opt) + \varepsilon D$. As we have assumed $D \leq cost(Opt)$, the claim follows.

In order to apply Lemma 4, we need to find a partition $F = A \uplus B$ satisfying $Opt \subseteq A$ and $\varepsilon\text{-SUPP}(Opt) \subseteq B$. Since $\varepsilon\text{-SUPP}(Opt) = \mathcal{O}(\ell^3/\varepsilon)$, we can do this via randomization. Consider tossing a biased coin for each facility independently: with probability $\frac{\varepsilon}{\ell^3}$ we place it in A , and with remaining probability in B . The probability of obtaining a partitioning satisfying $Opt \subseteq A$ and $\varepsilon\text{-SUPP}(Opt) \subseteq S$ equals $(\frac{\varepsilon}{\ell^3})^\ell$ times $(1 - \frac{\varepsilon}{\ell^3})^{\mathcal{O}(\frac{\ell^3}{\varepsilon})} = \Omega(1)$. Therefore $2^{\mathcal{O}(\ell \log(\ell/\varepsilon))}$ trials give a constant probability of sampling a correct partitioning. In order to derandomize this process, we take advantage of the following construction which is a folklore corollary from the framework of color-coding [2]. As we are not aware of any self-contained proof of this claim in the literature, we provide it for completeness.

Lemma 5. *For a set U of size n , there exists a family \mathcal{H} of partitions $U = A \uplus B$ such that $|\mathcal{H}| = 2^{\mathcal{O}(\ell \log(\ell+r))} \log n$ and for every pair of disjoint sets $A_0, B_0 \subseteq U$ with $|A_0| \leq \ell$, $|B_0| \leq r$, there is $(A, B) \in \mathcal{F}$ satisfying $A_0 \subseteq A, B_0 \subseteq B$. The family \mathcal{H} can be constructed in time $2^{\mathcal{O}(\ell \log(\ell+r))} n \log n$.*

Proof. Let us denote $[n] = \{1, 2, \dots, n\}$ and identify $U = [n]$. We rely on the following theorem: for any integers n, k there exists a family \mathcal{F} of functions $f : [n] \rightarrow [k^2]$, such that $|\mathcal{F}| = k^{\mathcal{O}(1)} \log n$ and for each $X \subseteq [n]$ of size k there is a function $f \in \mathcal{F}$ which is injective on X ; moreover, \mathcal{F} can be constructed in time $k^{\mathcal{O}(1)} n \log n$ [14, Theorem 5.16].

We use this construction for $k = \ell + r$. Next, consider the family \mathcal{G} of all functions $g : [(\ell + r)^2] \rightarrow \{0, 1\}$ such that $|g^{-1}(0)| \leq \ell$. Clearly, $|\mathcal{G}| \leq (\ell + r)^{2\ell}$. The family \mathcal{H} is given by taking all compositions $\{h = g \circ f \mid g \in \mathcal{G}, f \in \mathcal{F}\}$ and setting $(A_h, B_h) = (h^{-1}(0), h^{-1}(1))$. We have $|\mathcal{H}| \leq |\mathcal{G}| \cdot |\mathcal{F}| = 2^{\mathcal{O}(\ell \log(\ell+r))} \log n$. Let us consider any pair of disjoint subsets $A_0, B_0 \subseteq [n]$ with $|A_0| \leq \ell$, $|B_0| \leq r$. There exists $f \in \mathcal{F}$ injective on $A_0 \cup B_0$ and $g \in \mathcal{G}$ that maps $f(A_0)$ to 0 and $f(B_0)$ to 1, so $A_0 \subseteq A_{g \circ f}, B_0 \subseteq B_{g \circ f}$.

Theorem 4. *The CO- ℓ -MEDIAN problem admits a deterministic $(1+\varepsilon)$ -approximation algorithm running in time $2^{\mathcal{O}(\ell \log(\ell/\varepsilon))} \cdot n^{\mathcal{O}(1)}$ for any constant $\varepsilon > 0$.*

Proof. We apply Lemma 5 for $U = F$, ℓ being the parameter, and $r = 6 \cdot \ell^3/\varepsilon + \ell$, which upper bounds the size of $\varepsilon\text{-SUPP}(Opt)$ (Lemma 1). The family \mathcal{H} contains a partition $F = A \uplus B$ satisfying $Opt \subseteq A$ and $\varepsilon\text{-SUPP}(Opt) \subseteq B$. Next, we need to find D , such that $cost(Opt) \in [D, 2D]$. We begin with any polynomial-time α -approximation algorithm for k -MEDIAN ($\alpha = \mathcal{O}(1)$) to get an interval $[X, \alpha X]$, which contains $cost(Opt)$. We cover this interval with a constant number of intervals of the form $[X, 2X]$ and one of these provides a valid value of D . We invoke the algorithm from Lemma 4 for each such triple (A, B, D) and return a solution with the smallest cost.

4 Hardness of Capacitated co- ℓ -Median

In this section we show that, unlike CO- ℓ -MEDIAN, its capacitated counterpart does not admit a parameterized approximation scheme.

We shall reduce from the MAX k -COVERAGE problem, which was also the source of lower bounds for k -MEDIAN in the polynomial-time regime [21] and when parameterized by k [12]. However, the latter reduction is not longer valid when we consider a different parameterization for k -MEDIAN, as otherwise we could not obtain Theorem 1. Therefore, we need to design a new reduction, that exploits the capacity constraints and translates the parameter k of an instance of MAX k -COVERAGE into the parameter ℓ of an instance of CAPACITATED CO- ℓ -MEDIAN. To the best of our knowledge, this is the first hardness result in which the capacities play a role and allow us to obtain a better lower bound.

We rely on the following strong hardness result. Note that this result is a strengthening of [12], which only rules out $f(k) \cdot n^{k^{poly(1/\delta)}}$ -time algorithm. This suffices to rule out a parameterized approximation scheme for CAPACITATED CO- ℓ -MEDIAN, but not for a strong running time lower bound of the form $f(\ell) \cdot n^{\mathcal{O}(\ell)}$.

Theorem 5 ([28]). *Assuming Gap-ETH, there is no $f(k) \cdot n^{\mathcal{O}(k)}$ -time algorithm that can approximate MAX k -COVERAGE to within a factor of $(1 - 1/e + \delta)$ for any function f and any constant $\delta > 0$. Furthermore, this holds even when every input subset is of the same size and with a promise that there exists k subsets that covers each element exactly once.*

We can now prove our hardness result for CAPACITATED CO- ℓ -MEDIAN.

Theorem 6. *Assuming Gap-ETH, there is no $f(\ell) \cdot n^{o(\ell)}$ -time algorithm that can approximate CAPACITATED CO- ℓ -MEDIAN to within a factor of $(1 + 2/e - \epsilon)$ for any function f and any constant $\epsilon > 0$.*

Proof. Let U, T_1, \dots, T_n be an instance of MAX k -COVERAGE. We create an instance (F, C) of CAPACITATED CO- ℓ -MEDIAN as follows.

- For each subset T_i with $i \in [n]$, create a facility f_i^{set} with capacity $|T_i|$. For each element $u \in U$, create a facility $f_u^{element}$ with capacity $|U| + 2$.
- For every $i \in [n]$, create $|T_i|$ clients $c_{i,1}^{set}, \dots, c_{i,|T_i|}^{set}$. For each $j \in [|T_i|]$, we define the distance from $c_{i,j}^{set}$ to the facilities by

$$\begin{aligned} d(c_{i,j}^{set}, f_i^{set}) &= 0, \\ d(c_{i,j}^{set}, f_u^{element}) &= 1 && \forall u \in T_i, \\ d(c_{i,j}^{set}, f_{i'}^{set}) &= 2 && \forall i' \neq i, \\ d(c_{i,j}^{set}, f_u^{element}) &= 3 && \forall u \notin T_i. \end{aligned}$$

- For every element $u \in U$, create $|U| + 1$ clients $c_{u,1}^{element}, \dots, c_{u,|U|+1}^{element}$ and, for each $j \in [|U| + 1]$, define the distance from $c_{u,j}^{element}$ to the facilities by

$$\begin{aligned} d(c_{u,j}^{element}, f_u^{element}) &= 0, \\ d(c_{u,j}^{element}, f_i^{set}) &= 1 && \forall T_i \ni u, \\ d(c_{u,j}^{element}, f_{u'}^{element}) &= 2 && \forall u' \neq u, \\ d(c_{u,j}^{element}, f_i^{set}) &= 3 && \forall T_i \not\ni u. \end{aligned}$$

- Let $\ell = k$.

Suppose that we have an $f(\ell) \cdot n^{o(\ell)}$ -time $(1 + 2/e - \epsilon)$ -approximation algorithm for CAPACITATED CO- ℓ -MEDIAN. We will use it to approximate MAX k -COVERAGE instance with $|T_1| = \dots = |T_n| = |U|/k$ with a promise that there exists k subsets that covers each element exactly once, as follows. We run the above reduction to produce an instance (F, C) and run the approximation algorithm for CAPACITATED CO- ℓ -MEDIAN; let $S \subseteq F$ be the produced solution. Notice that S may not contain any element-facility, as otherwise there would not even be enough capacity left to serve all clients. Hence, $S = \{f_{i_1}^{set}, \dots, f_{i_k}^{set}\}$. We claim that T_{i_1}, \dots, T_{i_k} is an $(1 - 1/e + \epsilon/2)$ -approximate solution for MAX k -COVERAGE.

To see that T_{i_1}, \dots, T_{i_k} is an $(1 - 1/e + \epsilon/2)$ -approximate solution for MAX k -COVERAGE, notice that the cost of closing $\{f_{i_1}^{set}, \dots, f_{i_k}^{set}\}$ is exactly $|T_{i_1} \cup \dots \cup T_{i_k}| + 3 \cdot |U \setminus (T_{i_1} \cup \dots \cup T_{i_k})|$ because each element-facility $f_u^{element}$ can only serve one more client in addition to $c_{u,1}^{element}, \dots, c_{u,|U|+1}^{element}$. (Note that we may assume without loss of generality that $f_u^{element}$ serves $c_{u,1}^{element}, \dots, c_{u,|U|+1}^{element}$.) Moreover, there are exactly $|T_{i_1}| + \dots + |T_{i_k}| = |U|$ clients left to be served after the closure of $\{f_{i_1}^{set}, \dots, f_{i_k}^{set}\}$. Hence, each element-facility $f_u^{element}$ with $u \in T_{i_1} \cup \dots \cup T_{i_k}$

can serve a client of distance one from it. All other element-facilities will have to serve a client of distance three from it. This results in the cost of exactly $|T_{i_1} \cup \dots \cup T_{i_k}| + 3 \cdot |U \setminus (T_{i_1} \cup \dots \cup T_{i_k})|$. Now, since we are promised that there exists k subsets that uniquely covers the universe U , the optimum of CAPACITATED CO- ℓ -MEDIAN must be $|U|$. Since our (assumed) approximation algorithm for CAPACITATED CO- ℓ -MEDIAN has approximation factor $(1 + 2/e - \epsilon)$, we must have $|T_{i_1} \cup \dots \cup T_{i_k}| + 3 \cdot |U \setminus (T_{i_1} \cup \dots \cup T_{i_k})| \leq |U| \cdot (1 + 2/e - \epsilon)$, which implies that $|T_{i_1} \cup \dots \cup T_{i_k}| \geq |U| \cdot (1 - 1/e + \epsilon/2)$. Hence, the proposed algorithm is an $f(k) \cdot n^{o(k)}$ -time algorithm that approximates MAX k -COVERAGE to within a factor of $(1 - 1/e + \epsilon/2)$, which by Theorem 5 contradicts Gap-ETH.

5 Conclusions and Open Problems

We have presented a parameterized approximation scheme for CO- ℓ -MEDIAN and shown that its capacitated version does not admit such a scheme. It remains open whether CAPACITATED CO- ℓ -MEDIAN admits any constant-factor FPT approximation. Obtaining such a result might be an important step towards getting a constant-factor polynomial-time approximation, which is a major open problem.

Another interesting question concerns whether one can employ the framework of lossy kernelization [26] to get a polynomial size approximate kernelization scheme (PSAKS) for CO- ℓ -MEDIAN, which would be a strengthening of our main result. In other words, can we process an instance \mathcal{I} in polynomial time to produce an equivalent instance \mathcal{I}' of size $poly(\ell)$ so that solving \mathcal{I}' would provide a $(1 + \epsilon)$ -approximation for \mathcal{I} ?

References

1. Adamczyk, M., Byrka, J., Marcinkowski, J., Meesum, S.M., Włodarczyk, M.: Constant-factor FPT approximation for capacitated k -median. In: 27th Annual European Symposium on Algorithms, ESA 2019. Leibniz International Proceedings in Informatics (LIPIcs), vol. 144, pp. 1:1–1:14. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2019). <https://doi.org/10.4230/LIPIcs.ESA.2019.1>
2. Alon, N., Yuster, R., Zwick, U.: Color-coding. J. ACM **42**(4), 844–856 (1995). <https://doi.org/10.1145/210332.210337>
3. Arya, V., Garg, N., Khandekar, R., Meyerson, A., Munagala, K., Pandit, V.: Local search heuristics for k -median and facility location problems. SIAM J. Comput. **33**(3), 544–562 (2004). <https://doi.org/10.1137/S0097539702416402>
4. Bhattacharyya, A., Ghoshal, S., Karthik C.S., Manurangsi, P.: Parameterized intractability of even set and shortest vector problem from gap-ETH. In: 45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, Prague, Czech Republic, 9–13 July 2018, pp. 17:1–17:15 (2018). <https://doi.org/10.4230/LIPIcs.ICALP.2018.17>

5. Byrka, J., Pensyl, T., Rybicki, B., Srinivasan, A., Trinh, K.: An improved approximation for k -median, and positive correlation in budgeted optimization. In: Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 737–756. SIAM (2015). <https://doi.org/10.1145/2981561>
6. Byrka, J., Rybicki, B., Uniyal, S.: An approximation algorithm for uniform capacitated k -median problem with $1 + \epsilon$ capacity violation. In: Proceedings of the 18th International Conference on Integer Programming and Combinatorial Optimization, IPCO 2016, Liège, Belgium, 1–3 June 2016, pp. 262–274 (2016). https://doi.org/10.1007/978-3-319-33461-5_22
7. Chalermsook, P., et al.: From gap-ETH to FPT-inapproximability: clique, dominating set, and more. In: 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS), pp. 743–754 (October 2017). <https://doi.org/10.1109/FOCS.2017.74>
8. Charikar, M., Chekuri, C., Goel, A., Guha, S.: Rounding via trees: deterministic approximation algorithms for group Steiner trees and k -median. In: Proceedings of the 30th Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, 23–26 May 1998, pp. 114–123 (1998). <https://doi.org/10.1145/276698.276719>
9. Charikar, M., Guha, S., Tardos, É., Shmoys, D.B.: A constant-factor approximation algorithm for the k -median problem. In: Proceedings of the 31st annual ACM Symposium on Theory of Computing, pp. 1–10. ACM (1999). <https://doi.org/10.1145/301250.301257>
10. Chitnis, R., Feldmann, A.E., Manurangsi, P.: Parameterized approximation algorithms for bidirected Steiner network problems. In: 26th Annual Symposium on Algorithms, ESA 2018. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2018). <https://doi.org/10.4230/LIPIcs.ESA.2018.20>
11. Chitnis, R., Feldmann, A.E., Suchý, O.: A tight lower bound for planar Steiner orientation. *Algorithmica* (2019). <https://doi.org/10.1007/s00453-019-00580-x>
12. Cohen-Addad, V., Gupta, A., Kumar, A., Lee, E., Li, J.: Tight FPT approximations for k -median and k -Means. In: 46th International Colloquium on Automata, Languages, and Programming, ICALP 2019. Leibniz International Proceedings in Informatics (LIPIcs), vol. 132, pp. 42:1–42:14. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2019). <https://doi.org/10.4230/LIPIcs.ICALP.2019.42>
13. Cohen-Addad, V., Li, J.: On the fixed-parameter tractability of capacitated clustering. In: Baier, C., Chatzigiannakis, I., Flocchini, P., Leonardi, S. (eds.) 46th International Colloquium on Automata, Languages, and Programming, ICALP 2019. Leibniz International Proceedings in Informatics (LIPIcs), vol. 132, pp. 41:1–41:14. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2019). <https://doi.org/10.4230/LIPIcs.ICALP.2019.41>
14. Cygan, M., et al.: *Parameterized Algorithms*. Springer, Cham (2015). <https://doi.org/10.1007/978-3-319-21275-3>
15. Cygan, M., Hajiaghayi, M., Khuller, S.: LP rounding for k -centers with non-uniform hard capacities. In: 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science (FOCS), pp. 273–282. IEEE (2012). <https://doi.org/10.1109/FOCS.2012.63>
16. Dehne, F., Fellows, M., Fernau, H., Prieto, E., Rosamond, F.: NONBLOCKER: parameterized algorithmics for MINIMUM DOMINATING SET. In: Wiedermann, J., Tel, G., Pokorný, J., Bieliková, M., Štuller, J. (eds.) SOFSEM 2006. LNCS, vol. 3831, pp. 237–245. Springer, Heidelberg (2006). https://doi.org/10.1007/11611257_21

17. Demirci, H.G., Li, S.: Constant approximation for capacitated k -median with $(1 + \epsilon)$ -capacity violation. In: 43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, Rome, Italy, 11–15 July 2016, pp. 73:1–73:14 (2016). <https://doi.org/10.4230/LIPIcs.ICALP.2016.73>
18. Dinur, I.: Mildly exponential reduction from gap 3SAT to polynomial-gap label-cover. *Electron. Colloquium Comput. Complex. (ECCC)* **23**, 128 (2016)
19. Fakcharoenphol, J., Rao, S., Talwar, K.: A tight bound on approximating arbitrary metrics by tree metrics. In: Proceedings of the 35th Annual ACM Symposium on Theory of Computing, San Diego, CA, USA, 9–11 June 2003, pp. 448–455 (2003). <https://doi.org/10.1145/780542.780608>
20. Feldmann, A.E., Karthik, C., Lee, E., Manurangsi, P.: A survey on approximation in parameterized complexity: hardness and algorithms. *Algorithms* **13**(6), 146 (2020). <https://doi.org/10.3390/a13060146>
21. Guha, S., Khuller, S.: Greedy strikes back: improved facility location algorithms. *J. Algorithms* **31**(1), 228–248 (1999). <https://doi.org/10.1006/jagm.1998.0993>
22. Gupta, A., Lee, E., Li, J.: An FPT algorithm beating 2-approximation for k -cut. In: Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 2821–2837. Society for Industrial and Applied Mathematics (2018). <https://doi.org/10.1137/1.9781611975031.179>
23. Korupolu, M.R., Plaxton, C.G., Rajaraman, R.: Analysis of a local search heuristic for facility location problems. *J. Algorithms* **37**(1), 146–188 (2000). <https://doi.org/10.1006/jagm.2000.1100>
24. Li, S.: On uniform capacitated k -median beyond the natural LP relaxation. In: Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 696–707. SIAM (2015). <https://doi.org/10.1145/2983633>
25. Li, S.: Approximating capacitated k -median with $(1 + \epsilon)k$ open facilities. In: Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 786–796. SIAM (2016). <https://doi.org/10.1137/1.9781611974331.ch56>
26. Lokshtanov, D., Panolan, F., Ramanujan, M.S., Saurabh, S.: Lossy kernelization. In: Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, pp. 224–237. Association for Computing Machinery, New York, NY, USA (2017). <https://doi.org/10.1145/3055399.3055456>
27. Lokshtanov, D., Ramanujan, M.S., Saurabh, S., Zehavi, M.: Parameterized complexity and approximability of directed odd cycle transversal. In: Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, 5–8 January 2020, pp. 2181–2200 (2020). <https://doi.org/10.1137/1.9781611975994.134>
28. Manurangsi, P.: Tight running time lower bounds for strong inapproximability of maximum k -coverage, unique set cover and related problems (via t -wise agreement testing theorem). In: Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 62–81. SIAM (2020). <https://doi.org/10.1137/1.9781611975994.5>
29. Manurangsi, P., Raghavendra, P.: A birthday repetition theorem and complexity of approximating dense CSPs. In: ICALP, pp. 78:1–78:15 (2017). <https://doi.org/10.4230/LIPIcs.ICALP.2017.78>
30. van Rooij, J.M.: Exact Exponential-Time Algorithms for Domination Problems in Graphs. BOXpress (2011)
31. Xu, Y., Möhring, R.H., Xu, D., Zhang, Y., Zou, Y.: A constant FPT approximation algorithm for hard-capacitated k -means. *Optim. Eng.* **21**(3), 709–722 (2020). <https://doi.org/10.1007/s11081-020-09503-0>