

Approximation Algorithms for Clustering with Dynamic Points

Shichuan Deng^{*1}Jian Li^{†1}Yuval Rabani^{‡2}¹Tsinghua University, China²The Hebrew University of Jerusalem, Israel

Abstract

In many classic clustering problems, we seek to sketch a massive data set of n points (a.k.a. *clients*) in a metric space, by segmenting them into k categories or clusters, each cluster represented concisely by a single point in the metric space (a.k.a. the cluster's *center* or its *facility*). The goal is to find such a sketch that minimizes some objective that depends on the distances between the clients and their respective facilities (the objective is a.k.a. the *service cost*). Two notable examples are the k -center/ k -supplier problem where the objective is to minimize the maximum distance from any client to its facility, and the k -median problem where the objective is to minimize the sum over all clients of the distance from the client to its facility.

In practical applications of clustering, the data set may evolve over time, reflecting an evolution of the underlying clustering model. Thus, in such applications, a good clustering must simultaneously represent the temporal data set well, but also not change too drastically between time steps. In this paper, we initiate the study of a dynamic version of clustering problems that aims to capture these considerations. In this version there are T time steps, and in each time step $t \in \{1, 2, \dots, T\}$, the set of clients needed to be clustered may change, and we can move the k facilities between time steps. The general goal is to minimize certain combinations of the service cost and the facility movement cost, or minimize one subject to some constraints on the other. More specifically, we study two concrete problems in this framework: the Dynamic Ordered k -Median and the Dynamic k -Supplier problem. Our technical contributions are as follows:

- We consider the Dynamic Ordered k -Median problem, where the objective is to minimize the weighted sum of ordered distances over all time steps, plus the total cost of moving the facilities between time steps. We present one constant-factor approximation algorithm for $T = 2$ and another approximation algorithm for fixed $T \geq 3$.
- We consider the Dynamic k -Supplier problem, where the objective is to minimize the maximum distance from any client to its facility, subject to the constraint that between time steps the maximum distance moved by any facility is no more than a given threshold. When the number of time steps T is 2, we present a simple constant factor approximation algorithm and a bi-criteria constant factor approximation algorithm for the outlier version, where some of the clients can be discarded. We also show that it is NP-hard to approximate the problem with any factor for $T \geq 3$.

1 Introduction

Clustering a data set of points in a metric space is a fundamental abstraction of many practical problems of interest and has been subject to extensive study as a fundamental problem of both machine learning and combinatorial optimization. In particular, cluster analysis is one of the main methods of unsupervised

*E-mail: dsc15@mails.tsinghua.edu.cn.

†E-mail: lijian83@mail.tsinghua.edu.cn.

‡E-mail: ywabani@cs.huji.ac.il.

learning, and clustering often models facility location problems.¹ More specifically, some of the most well-studied clustering problems involve the following generic setting. We are given a set C of points in a metric space, and our goal is to compute a set of k centers that optimizes a certain objective function which involves the distances between the points in C and the computed centers. Two prominent examples are the k -median problem and the k -center problem. They are formally defined as follows. Let S denote the computed set of k cluster centers, let $d(j, S) = \min_{i \in S} d(i, j)$ be the minimum distance from a point $j \in C$ to S , and let $D = (d(j, S))_{j \in C}$ be called the *service cost vector*. The k -median problem aims to minimize the L_1 objective $\|D\|_1 = \sum_{j \in C} d(j, S)$ over the choices of S , and the k -center aims to minimize the L_∞ objective $\|D\|_\infty = \max_{j \in C} d(j, S)$. In general metric spaces and when k is not a fixed constant, both problems are APX-hard and exhibit constant factor approximation algorithms [4, 6, 12, 23, 32]. An important generalization is the *ordered k -median* problem. Here, in addition to C and k , we are given also a non-increasing weight vector $w \in \mathbb{R}_{\geq 0}^{|C|}$. Letting D^\downarrow denote the sorted version of D in non-increasing order, the objective of ordered k -median is to minimize $w \cdot D^\downarrow$. This problem generalizes both k -center and k -median and has attracted significant attention recently and several constant factor approximation algorithms have been developed [3, 7, 9, 10].

In this paper, we study several dynamic versions of the classical clustering problems, in which the points that need to be clustered may change for each time step, and we are allowed to move the cluster centers in each time step, either subject to a constraint on the distance moved, or by incurring a cost proportional to that distance. These versions are motivated in general by practical applications of clustering, where the data set evolves over time, reflecting an evolution of the underlying clustering model. Consider, for instance, a data set representing the active users of a web service, and a clustering representing some meaningful segmentation of the user base. The segmentation should be allowed to change over time, but if it is changed drastically between time steps, then it is probably meaningless. For a more concrete example, consider the following application scenario. There is a giant construction company with several construction teams working in a city. The company has k movable wireless base stations for their private radio communication, and each construction team also has a terminal device. The teams need to put their devices at a certain energy level, in order to maintain the communication channel between the device and the nearest base station. Some construction team may finish some project and move to another place at some time. Note that the wireless base stations are also movable at a certain expense. Our high level objective is to have all teams covered by the base stations at all times, meanwhile minimizing the energy cost of all teams plus the cost of moving these base stations.

We study two problems of this flavor. The first problem, a dynamic version of the *ordered k -median* problem, is a very general model that captures a wide range of dynamic clustering problems where the objective is to minimize the sum of service cost and movement cost. In particular, it generalizes dynamic versions of k -center and k -median. The problem is defined as follows. We are given a metric space and there are T time steps. In each time step t , there is a set C_t of clients that we need to serve. In each time step, we can also choose the locations for k movable facilities to serve the clients (each client is served by its closest facility). Our goal is to minimize the total ordered service distance (i.e., the ordered k -median objective), summed over all times steps, plus the total distances traveled by the k movable facilities. We define the problem formally as follows.

Definition 1. (Dynamic Ordered k -Median) We are given a metric space (X, d) . An instance of Dynamic Ordered k -Median is specified by $(\{C_t\}_{t=1}^T, \{F_t\}_{t=1}^T, \{w_t \in \mathbb{R}_{\geq 0}^{|C_t|}\}_{t=1}^T, \gamma > 0, k \in \mathbb{N}_+)$, where $T \geq 2$ is a constant integer, $C_t \subset X$ is the set of clients for time t , $F_t \subset X$ is the set of candidate locations where we can place facilities. For a vector v , denote by v^\downarrow the vector derived from v by sorting its entries in non-increasing order. Also denote by $m(X, Y) = \min_{M_0 \in M(X, Y)} \sum_{(i, i') \in M_0} d(i, i')$ the total weight of minimum-weight perfect matching between two equal-sized multi-sets X, Y . We are required to compute a sequence of multi-sets of facilities $\{A_t\}_{t=1}^T$ with $A_t \subset F_t, |A_t| = k$, so that the following sum of ordered service cost

¹In the facility location literature, points are called clients and centers are called facilities, and we will use these terms interchangeably.

and movement cost is minimized:

$$\sum_{t=1}^T w_t \cdot (d(j, A_t))_{j \in C_t}^\downarrow + \gamma \cdot \sum_{t=1}^{T-1} m(A_t, A_{t+1}). \quad (1)$$

It is also natural to formulate dynamic clustering problems where the objective is to minimize just the service cost, subject to a constraint on the movement cost. This turns out to be technically very different from Dynamic Ordered k -Median. Our second problem, which we call Dynamic k -Supplier, is such a concrete problem, motivated by the above-mentioned construction company application. In this problem the service cost is the k -supplier objective, i.e. the maximum client service distance over all time steps, and the constraints are that any facility cannot be moved further than a fixed value $B > 0$ between any two consecutive time steps. More formally:

Definition 2. (Dynamic k -Supplier) We are given a metric space (X, d) . An instance of Dynamic k -Supplier is specified by $(\{C_t\}_{t=1}^T, \{F_t\}_{t=1}^T, B > 0, k \in \mathbb{N}_+)$, where $T \geq 2$ is the number of time steps, $C_t \subset X$ is the set of clients for time t , $F_t \subset X$ is the set of candidate locations where we can place facilities. We are required to compute a sequence of multi-sets of facilities $\{A_t\}_{t=1}^T$, with $A_t \subset F_t$, $|A_t| = k$, minimizing the maximum service cost of any client $\max_t \max_{j \in C_t} d(j, A_t)$, subject to the constraint that there must exist a one-to-one matching between A_t and A_{t+1} for any t , such that the distance between each matched pair is at most B .

In the outlier version, we are additionally given the outlier constraints $\{l_t \in \mathbb{N}\}_{t=1}^T$. We are asked to identify a sequence of multi-sets of facilities $\{A_t\}_{t=1}^T$, as well as a sequence of subsets of served clients $\{S_t \subset C_t\}_{t=1}^T$. The goal is to minimize the maximum service cost of any served client $\max_t \max_{j \in S_t} d(j, A_t)$, with the constraints that $A_t \subset F_t$, $|A_t| = k$, $|S_t| \geq l_t$, and there must exist a one-to-one matching between A_t and A_{t+1} for any t , such that the distance between each matched pair is at most B .

Note: The solutions for both Dynamic Ordered k -Median and Dynamic k -Supplier may be multi-sets, since we allow multiple centers to travel to the same location.

1.1 Our Results

We first study Dynamic Ordered k -Median. We assume the number of time steps T is a constant and all entries of the weight vector are larger than some small constant $\epsilon > 0$. We present a polynomial-time approximation on general metrics. Moreover, if $T = 2$ we present a constant-factor approximation algorithm without the assumption on the entries of the weight vectors.

Theorem 3. 1. If $T = 2$, for any constant $\delta > 0$ there exists a polynomial-time $(48 + 20\sqrt{3} + \delta)$ -approximation for Dynamic Ordered k -Median.

2. If $T \geq 3$ is a constant and all entries in $\{w_t\}_{t=1}^T$ are at least $\epsilon > 0$, for any constant $\delta > 0$ there exists a polynomial-time $(48 + 20\sqrt{3} + \delta + 6\gamma/\epsilon)$ -approximation algorithm for Dynamic Ordered k -Median.

Our techniques: The key idea in our algorithm is to design a surrogate relaxed LP as in [7] and embed the fractional LP solution in a network flow instance. We then proceed to round the fractional flow to an integral flow, thus obtaining the integral solution. The network is constructed based on a filtering process introduced by Charikar and Li [12]. We also adapt the oblivious clustering arguments by Byrka *et al.* [7], but with a slight increase in approximation factors due to the structure of our network flow.

Our approach can also give a constant approximation to the facility-weighted minimum total movement mobile facility location problem (facility-weighted TM-MFL), with a simpler analysis than the previously known local-search based algorithm [1], which achieves an approximation factor of $3 + O(\sqrt{\log \log p / \log p})$ for a p -Swap algorithm. Swamy [33] also points out an approximation-preserving reduction of TM-MFL to matroid median, and gives an 8-approximation algorithm. The approximation ratio for matroid median is further improved to 7.081 by Krishnaswamy *et al.* [30]. We note that the same reduction does not work for Dynamic Ordered k -Median even when the service costs are not ordered, as it either causes the triangle

inequality to fail, or turns out to require the intersection of at least 3 laminar families, and neither of them is feasible for approximation.

Theorem 4. *There exists a polynomial-time 10-approximation algorithm for facility-weighted minimum total movement mobile facility location problem.*

As a second result, we consider Dynamic k -Supplier and its outlier version. We show that if $T \geq 3$, it is not possible to obtain efficient approximation algorithms for Dynamic k -Supplier with *any* approximation factor, unless $P = NP$, via a simple reduction from perfect 3D matching [27]. However, for the case of $T = 2$, we present a flow-based 3-approximation, which is the best possible factor since vanilla k -supplier is NP-hard to approximate within a factor of $(3 - \epsilon)$ for any constant $\epsilon > 0$ [24].

Theorem 5. 1. *There exists a 3-approximation for Dynamic k -Supplier when $T = 2$.*

2. *There is no polynomial time algorithm for solving Dynamic k -Supplier with any approximation factor if $T \geq 3$, unless $P = NP$.*

We also study the outlier version of the problem for $T = 2$. In the outlier version, we can exclude a certain fraction of the clients as outliers in each time step. We obtain a bi-criteria approximation for the problem.

Theorem 6. *For any constant $\epsilon > 0$, there exists a bi-criteria 3-approximation algorithm for Dynamic k -Supplier with outliers when $T = 2$, that outputs a solution which covers at least $(1 - \epsilon)l_t$ clients within radius $3R^*$ at time t , where $t = 1, 2$ and R^* is the optimal radius.*

Our techniques: We first guess a constant-size portion of facilities in the optimal solution, remove these facilities and solve the LP relaxation of the remaining problem. This guessing step is standard as in multi-objective optimizations in [20]. From the LP solution, we form clusters as in Harris *et al.* [22], cast the outlier constraints as budget constraints over the LP solution, and finally round the fractional LP solution to an integral solution using the budgeted optimization methods by Grandoni *et al.* [20]. Note that since our outlier constraints translate naturally to budget lower bounds, and our optimization goal is minimization, we are only able to achieve bi-criteria approximations instead of pure approximations. For more details, please see Appendix C.

1.2 Other Related Work

The ordered k -median problem generalizes a number of classic clustering problems like k -center, k -median, k -facility l -centrum, and has been studied extensively in the literature. There are numerous approximation algorithms known for its special cases. We survey here only the results most relevant to our work (ignoring, for instance, results regarding restricted metric spaces or fixed k). Constant approximations for k -median can be obtained via local search, Lagrangian relaxations and the primal-dual schema, or LP-rounding [4, 6, 25, 26]. Constant approximations for k -center are obtained via greedy algorithms [23]. Aouad and Segev [3] employ the idea of surrogate models and give the first $O(\log n)$ -approximation for ordered k -median. Later, Byrka *et al.* [7], Chakrabarty and Swamy [9] both successfully design constant-factor approximations for k -facility l -centrum and ordered k -median. Chakrabarty and Swamy [10] subsequently improve the approximation factor for ordered k -median to $(5 + \epsilon)$, using deterministic rounding in a unified framework.

The outlier setting of clustering problems, specifically for center-type clustering problems, was introduced by Charikar *et al.* [11] and later further studied by Chakrabarty *et al.* [8]. Many other variants of different clustering constraints are also extensively studied, including matroid and knapsack center with outliers [13], and fair center-type problems with outliers [22].

Our problems are closely related to the mobile facility location problems (MFL), introduced by Demaine *et al.* [16]. In these problems, a static set of clients has to be served by a set of facilities that are given initial locations and can be moved to improve the service cost at the expense of incurring a facility movement cost. For the minimum total movement MFL problem (TM-MFL), Friggstad and Salavatipour [18] give

an 8-approximation using LP-rounding, where all facilities have unit weights. Ahmadian *et al.* [1] give a local search algorithm for TM-MFL with weighted facilities using p -swaps with an approximation ratio of $3 + O(\sqrt{\log \log p / \log p})$, and specifically show that the approximation ratio is at most 499 for $p = 1$. Swamy [33] shows an approximation-preserving reduction from TM-MFL to the matroid median problem, and gives an 8-approximation algorithm. Krishnaswamy *et al.* [30] later improve the approximation ratio for matroid median to 7.081, which consequently improves the approximation ratio for TM-MFL as well.

The dynamic formulations of our problems are closely related to the facility location problem with evolving metrics, proposed by Eisenstat *et al.* [17]. In this problem, there are also T time steps, while the facilities and clients are fixed, and the underlying metric is changing. The total cost is the sum of facility-opening cost, client-serving cost and additional switching costs for each client. The switching cost is paid whenever a client switches facility between adjacent time steps. In comparison, our problem **Dynamic k -Supplier** considers the cost of moving facilities instead of opening costs, and allows the number of clients to change over time. Eisenstat *et al.* [17] consider the problem when the open facility set A is fixed, and give a $O(\log(nT))$ -approximation, where n is the number of clients. They also show a hardness result on $o(\log T)$ -approximations. An *et al.* [2] consider the case when the open facilities are allowed to evolve as well, and give a 14-approximation.

Our problem is also related to stochastic k -server [15] and the page migration problem [5, 35]. Dehghani *et al.* [15] first study the *stochastic k -server* problem. In this problem, we also have T time steps, and the distributions $\{P_t\}_{t \in [T]}$ are given *in advance*. The t -th client is drawn from P_t , and we can use k movable servers. One variant they consider is that, after a client shows up, its closest server goes to the client's location and comes back, and the optimization objective is the total distance travelled by all servers. They provide an $O(\log n)$ -approximation for general metrics, where n is the size of the distribution support. In expectation, their objective is the same as in **Dynamic Ordered k -Median**, if we consider non-ordered weighted clients and total weights sum up to 1 for each time slot. However, we note that our result does not imply a constant approximation for their problem. The difficulty is that if one maps the stochastic k -server problem to our problem, the corresponding weight coefficient γ is not necessarily a constant and our approximation ratio is proportional to γ . Obtaining a constant factor approximation algorithm for stochastic k -server is still an interesting open problem.

2 A Constant Approximation for Dynamic Ordered k -Median

We devise an LP-based algorithm, which generalizes the oblivious-clustering argument by Byrka *et al.* [7]. At the center of our algorithm, a network flow method is used, where an integral flow is used to represent our solution.

2.1 Flow-based Rounding of LP Solution

We first formulate the LP relaxation. By adding a superscript to every variable to indicate the time step, we denote $x_{ij}^{(t)} \in [0, 1]$ the partial assignment of client j to facility i and $y_i^{(t)} \in [0, 1]$ the extent of opening facility location i at time step t . Moreover, denote $z_{ii'}^{(t)}$ the fractional movement from facility i to facility i' , between neighboring time steps t and $t + 1$.

The following surrogate LP is designed using the cost reduction trick by Byrka *et al.* [7]. When the reduced cost functions are exactly guessed, the LP relaxation has an objective value at most the total cost of the optimal solution, denoted by OPT. Call $d' : X \times X \rightarrow R_{\geq 0}$ a reduced cost function (not necessarily a metric) of distance function d , if for any $x, y \in X$, $d'(x, y) \geq 0$, $d'(x, y) = d'(y, x)$, and $d(x_1, y_1) \leq d(x_2, y_2) \Rightarrow d'(x_1, y_1) \leq d'(x_2, y_2)$. For a sequence of reduced cost functions $\mathfrak{D} = \{d^t\}_{t=1}^T$ of d ,

the modified LP relaxation is defined as follows.

$$\text{minimize : } \sum_{t=1}^T \sum_{j \in C_t} \sum_{i \in F_t} d^t(i, j) x_{ij}^{(t)} + \gamma \sum_{t=1}^{T-1} \sum_{i \in F_t} \sum_{i' \in F_{t+1}} d(i, i') z_{ii'}^{(t)} \quad \text{LP}(\mathfrak{D})$$

$$\text{subject to : } \sum_{i \in F_t} x_{ij}^{(t)} = 1, \quad \forall j \in C_t, t \in [T] \quad (2)$$

$$\sum_{i \in F_t} y_i^{(t)} = k, \quad \forall t \in [T] \quad (3)$$

$$0 \leq x_{ij}^{(t)} \leq y_i^{(t)}, \quad \forall i \in F_t, j \in C_t, t \in [T] \quad (4)$$

$$\sum_{i' \in F_{t+1}} z_{ii'}^{(t)} = y_i^{(t)}, \quad \forall i \in F_t, t \in [T-1] \quad (5)$$

$$\sum_{i \in F_t} z_{ii'}^{(t)} = y_{i'}^{(t+1)}, \quad \forall i' \in F_{t+1}, t \in [T-1] \quad (6)$$

Suppose we have solved the corresponding surrogate $\text{LP}(\mathfrak{D})$. In the optimal solution (x, y, z) , we assume that whenever $x_{ij}^{(t)} > 0$, we have $x_{ij}^{(t)} = y_i^{(t)}$, via the standard duplication technique of facility locations (for example, see [12]). Denote $\text{Ball}^o(j, R) = \{x \in X : d(x, j) < R\}$ the open ball centered at j with radius R , and $E_j^{(t)} = \{i \in F_t : x_{ij}^{(t)} > 0\}$ the relevant facilities for client j . For any specific time step t , denote $d_{\text{av}}^{(t)}(j) = \sum_{i \in E_j^{(t)}} d(i, j) x_{ij}^{(t)}$ the average *unweighted* service cost of client j and $y^{(t)}(S) = \sum_{i \in S} y_i^{(t)}$ the amount of fractional facilities in $S \subset F_t$. We perform a filtering-and-matching algorithm (see Appendix A.1) to obtain a subset $C'_t \subset C_t$ for each t , a bundle $\mathcal{U}_j^{(t)} \subset F_t$ for each $j \in C'_t$, as well as P_t a partition of C'_t , where

1. C'_t is a subset of “well-separated” clients of C_t , such that for any client in $C_t \setminus C'_t$, there exists another relatively close client in C'_t . To be more precise, for any $j \neq j'$ in C'_t , $d(j, j') \geq 4 \max\{d_{\text{av}}^{(t)}(j), d_{\text{av}}^{(t)}(j')\}$, and for any $j'' \in C_t \setminus C'_t$, there exists $j''' \in C'_t$ such that $d_{\text{av}}^{(t)}(j''') \leq d_{\text{av}}^{(t)}(j'')$, $d(j'', j''') \leq 4 \max\{d_{\text{av}}^{(t)}(j''), d_{\text{av}}^{(t)}(j''')\}$;
2. $\mathcal{U}_j^{(t)}$ is a subset of fractionally open facility locations that are relatively close to client j ;
3. P_t is a judiciously created partition of C'_t , where every subset contains either a pair of clients, or a single client. Each pair $\{j, j'\}$ in P_t is chosen such that either j or j' is the closest neighbor of the other, and we guarantee to open a facility location in $\mathcal{U}_j^{(t)}$ or $\mathcal{U}_{j'}^{(t)}$.

The filtering-and-matching algorithm (see Appendix A.1) is fairly standard in several LP-based methods for median-type problems (see e.g. [7, 10, 12]). It is worth noting that, while we define the objective value of $\text{LP}(\mathfrak{D})$ using reduced cost functions \mathfrak{D} with respect to the weights, the filtering algorithm is completely oblivious of the weights and only uses the underlying metric d .

Network construction: We construct an instance of network flow \mathcal{N} , and embed the LP solution as a fractional flow \tilde{f} . The network \mathcal{N} consists of a source \mathbf{s} , a sink \mathbf{t} and $6T$ intermediate layers L_1, L_2, \dots, L_{6T} arranged in a linear fashion.

For each time step $t \in [T]$, we create two nodes for every pair $p \in P_t$, every bundle $\mathcal{U}_j^{(t)}$ and every candidate facility location $i \in F_t$. All these nodes are contained in the layers L_{6t-5}, \dots, L_{6t} . To distinguish between the two mirror nodes, we use $\mathcal{L}(\cdot)$ and $\mathcal{R}(\cdot)$ to represent the nodes in $\{L_{6t-5}, L_{6t-4}, L_{6t-3}\}$ and the nodes in $\{L_{6t-2}, L_{6t-1}, L_{6t}\}$, respectively. The network is constructed as follows. An example figure is shown in Figure 1.

1. For all $t \in [T]$, add $\mathcal{L}(i)$ to L_{6t-5} and $\mathcal{R}(i)$ to L_{6t} for each $i \in F_t$.
2. For all $t \in [T]$, add $\mathcal{L}(\mathcal{U}_j^{(t)})$ to L_{6t-4} and $\mathcal{R}(\mathcal{U}_j^{(t)})$ to L_{6t-1} for each $\mathcal{U}_j^{(t)}$.

3. For all $t \in [T]$, add $\mathcal{L}(p)$ to L_{6t-3} and $\mathcal{R}(p)$ to L_{6t-2} for each $p \in P_t$.
4. For all $t \in [T]$, $j \in C'_t$, $p \in P_t$ such that $j \in p$, connect $(\mathcal{L}(U_j^{(t)}), \mathcal{L}(p)), (\mathcal{R}(p), \mathcal{R}(U_j^{(t)}))$ in neighboring layers with an edge of capacity $\lceil [y^{(t)}(U_j^{(t)})], \lceil y^{(t)}(U_j^{(t)}) \rceil \rceil$. Let their initial fractional flow values be $\tilde{f}(\mathcal{L}(U_j^{(t)}), \mathcal{L}(p)) = \tilde{f}(\mathcal{R}(p), \mathcal{R}(U_j^{(t)})) = y^{(t)}(U_j^{(t)})$. The capacity is either $[0, 1]$ or $\{1\}$.
5. For all $t \in [T]$, $p \in P_t$, connect $(\mathcal{L}(p), \mathcal{R}(p))$ with an edge of capacity $\lceil [y^{(t)}(p)], \lceil y^{(t)}(p) \rceil \rceil$, and define $\tilde{f}(\mathcal{L}(p), \mathcal{R}(p)) = y^{(t)}(p) = \sum_{j \in p} y^{(t)}(U_j^{(t)})$. If p is a normal pair, the capacity is either $[1, 2]$ or $\{1\}$ or $\{2\}$; if p is a singleton pair, the capacity is either $[0, 1]$ or $\{1\}$.
6. For all $t \in [T]$, $j \in C'_t$ and $i \in U_j^{(t)}$, connect $(\mathcal{L}(i), \mathcal{L}(U_j^{(t)})), (\mathcal{R}(U_j^{(t)}), \mathcal{R}(i))$ in neighboring layers with an edge of unit capacity. Let the initial fractional flows be $\tilde{f}(\mathcal{L}(i), \mathcal{L}(U_j^{(t)})) = \tilde{f}(\mathcal{R}(U_j^{(t)}), \mathcal{R}(i)) = y_i^{(t)}$.
7. For all $t \in [T]$ but $i \in F_t - \bigcup_{j \in C'_t} U_j^{(t)}$, connect $(\mathcal{L}(i), \mathcal{R}(i))$ with an edge of unit capacity (across intermediate layers $L_{6t-4}, \dots, L_{6t-1}$). Let its initial fractional flow be $\tilde{f}(\mathcal{L}(i), \mathcal{R}(i)) = y_i^{(t)}$.
8. For all $z_{ii'}^{(t)}$, $i \in F_t, i' \in F_{t+1}$, connect $(\mathcal{R}(i), \mathcal{L}(i'))$ with an edge of unit capacity. Let its initial fractional flow be $\tilde{f}(\mathcal{R}(i), \mathcal{L}(i')) = z_{ii'}^{(t)}$.

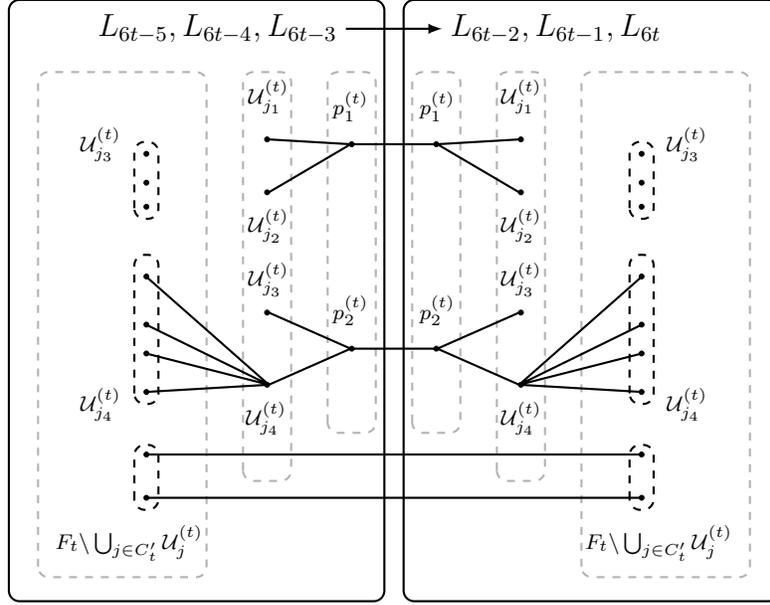


Figure 1: Some intermediate layers of \mathcal{N} representing a single time step t .

Notice \tilde{f} is naturally a flow with value k . Since the flow polytope is defined by a totally unimodular matrix, and our capacity constraints are all integers, it is a well-known result (see e.g. [19]) that we can efficiently and stochastically round \tilde{f} to an integral flow \bar{f} , such that \bar{f} is guaranteed to have value k , and $\mathbb{E}[\bar{f}] = \tilde{f}$. Next, given the integral flow \bar{f} , we deterministically construct the facilities to open $\{A_t\}_{t \in [T]}$ as follows.

- If $T = 2$, there are 12 layers L_1, L_2, \dots, L_{12} in the network. For each link $e = (\mathcal{R}(i_1), \mathcal{L}(i_2))$ between L_6 and L_7 such that $\bar{f}(e) = 1$, we add the original facility corresponding to i_1 to A_1 , and the original facility of i_2 to A_2 .

- If $T \geq 3$, the integral flow \bar{f} may enter L_{6t-5} and exit from L_{6t} at sets of different facility locations. For illustration, denote $A_{t,1}$ the set in L_{6t-5} and $A_{t,2}$ the set in L_{6t} . Notice it may happens that $|A_{t,1} \cup A_{t,2}| > k$ and we cannot open them both, so we design an algorithm to find $A_t \subseteq A_{t,1} \cup A_{t,2}$ and $|A_t| = k$, and open the facilities in A_t for time t .

The algorithm looks at each pair $(j_1, j_2) = p \in P_t$, and consider the 1 or 2 units of flow \bar{f} on the link $(\mathcal{L}(p), \mathcal{R}(p))$. For a facility i , if there is one unit of flow through $\mathcal{L}(i)$ or $\mathcal{R}(i)$, we call the facility i *activated*. But if $\mathcal{L}(i_1)$ and $\mathcal{R}(i_2)$ are activated and $i_1, i_2 \in \mathcal{U}_{j_1}, i_1 \neq i_2$, we only open one of them. The same is true when $i_1 \in \mathcal{U}_{j_1}, i_2 \in \mathcal{U}_{j_2}, i_1 \neq i_2$.

For each unit flow, our algorithm either always choose i_1 to open where $\mathcal{L}(i_1)$ is activated, or always choose the facility in \mathcal{U}_{j_2} if j_1 is not the closest neighbor of j_2 . The algorithm is deferred to Appendix A.2. As a result, we give the following lemma estimating the movement cost and defer the proof to Appendix A.3.

Lemma 7. *Let $d(A, A')$ denote the cost of minimum weight matching between A, A' . If $T = 2$, the expected movement cost of solution $\{A_1, A_2\}$ satisfies*

$$\mathbb{E}[d(A_1, A_2)] = \sum_{i \in F_1} \sum_{i' \in F_2} d(i, i') z_{ii'}^{(1)}.$$

If $T \geq 3$, the expected movement cost of solution $\{A_t\}_{t=1}^T$ after rerouting satisfies

$$\mathbb{E} \left[\sum_{t \in [T-1]} d(A_t, A_{t+1}) \right] \leq \sum_{t \in [T-1]} \sum_{i \in F_t} \sum_{i' \in F_{t+1}} d(i, i') z_{ii'}^{(t)} + 6 \sum_{t=1}^T \sum_{j \in C_t} d_{\text{av}}^{(t)}(j).$$

2.2 From Rectangular to General Cases

We first provide a lemma to bound the stochastic k -facility l -centrum cost of A_t for any fixed time t . Consequently, the ordered cost can be nicely bounded as well. We defer the proof of the following lemma to Appendix A.4.

Lemma 8. *(adapted from [7]) Fix $t \in [T]$ and let $m \in \mathbb{N}_+, h > 0$. Define $\text{rect}(a, b)$ the rectangular vector of length b , where the first a elements are 1s and the rest are 0s. For A_t as the (random) set of activated locations returned by our algorithm, and $d(C_t, A_t) = (d(j, A_t))_{j \in C_t}$ as the service cost vector, we have*

$$\mathbb{E}_{A_t}[\text{rect}(m, |C_t|) \cdot d(C_t, A_t)^\downarrow] \leq (24 + 10\sqrt{3})m \cdot h + (24 + 10\sqrt{3}) \sum_{j \in C_t} d_{\text{av}}^{-h(t)}(j),$$

where $d^{-h}(j, j') = 0$ if $d(j, j') < h$ and $d^{-h}(j, j') = d(j, j')$ otherwise. Similar to $d_{\text{av}}(j)$, the average clipped service cost $d_{\text{av}}^{-h(t)}(j)$ is defined as $d_{\text{av}}^{-h(t)}(j) = \sum_{i \in F_t} d^{-h}(i, j) x_{ij}^{(t)}$.

Finally we turn to the generally-weighted case, where the weight vectors $w_t, t \in [T]$ are not necessarily rectangular ones like $\text{rect}(m, |C_t|)$. The guessing of underlying reduced cost functions \mathfrak{D} is exactly the same as in Byrka *et al.* [7], thus omitted here. We solve $\text{LP}(\mathfrak{D})$ using these induced reduced cost functions and proceed accordingly. The following lemma is similar to that of Lemma 5.1 in Byrka *et al.* [7], and we defer the proof to Appendix A.5 due to the space limit.

Lemma 9. *When $T = 2$, the procedure described above is a $(48 + 20\sqrt{3})$ -approximation for Dynamic Ordered k -Median.*

If $T \geq 3$ is a constant and the smallest entry in $\{w_t\}_{t=1}^T$ is at least some constant $\epsilon > 0$, the above-described procedure is a $(48 + 20\sqrt{3} + 6\gamma/\epsilon)$ -approximation for Dynamic Ordered k -Median.

In both cases, the procedure makes $O\left(\prod_{t=1}^T (|F_t| \cdot |C_t|)^{N_t}\right)$ calls to its subroutines, where N_t is the number of distinct entries in the weight vector $w_t, t \in [T]$.

Fix some positive parameter $\delta > 0$ and recall the distance bucketing trick by Aouad and Segev [3]. When T is a constant, it is possible to guess the largest service distance for each time step by paying a polynomial factor in the running time. Then we make logarithmically many buckets for each time step to hold the service cost values of clients. For each bucket, its average weight is also guessed up to a small multiplicative error $(1 + \delta)$. Since there are at most $O(\log_{1+\delta}(\frac{n}{\delta})) = O(\frac{1}{\delta} \log(\frac{n}{\delta}))$ buckets for each time step, where $n = |F| + |C|$, guessing a non-increasing sequence of the average weights only causes another polynomial factor $\exp(O(\frac{1}{\delta} \log(\frac{n}{\delta}))) = n^{O(1/\delta)}$. Finally, because T is a constant, the overall number of guesses is still bounded by a polynomial. For more details, see [3, 7].

Theorem 10. *When $T = 2$, for any $\delta > 0$ there exists a $(48 + 20\sqrt{3})(1 + \delta)$ -approximation algorithm for Dynamic Ordered k -Median, with running time $(|F_1| + |C_1|)^{O(1/\delta)} \cdot (|F_2| + |C_2|)^{O(1/\delta)}$.*

When $T \geq 3$ is a constant, and the smallest entry in $\{w_t\}_{t=1}^T$ is at least some constant $\epsilon > 0$, for any $\delta > 0$ there exists a $(48 + 20\sqrt{3} + 6\gamma/\epsilon)(1 + \delta)$ -approximation algorithm for Dynamic Ordered k -Median, with running time $\prod_{t \in [T]} (|F_t| + |C_t|)^{O(1/\delta)}$.

Proof. This is almost a direct consequence of Theorem 5.2 in [7], with the constant factor replaced by our $\mu = 24 + 10\sqrt{3}$. Notice that we need to slightly modify the way of constructing rounded weights $\{w_t^*\}_{t=1}^T$ in the following way,

$$\forall t \in [T], r \in [|C_t|], w_{tr}^* = \begin{cases} w_{t1} & r = 1, \\ \min \left\{ (1 + \delta)^{\lceil \log_{1+\delta} w_{tr} \rceil}, w_{t1} \right\} & w_{tr} \geq \epsilon w_{t1} / |C_t|, r \neq 1, \\ \epsilon w_{t1} / |C_t| & w_{tr} < \epsilon w_{t1} / |C_t|, \end{cases}$$

so that the perturbed weight vectors are rounded larger, but at most $(1 + \delta)$ times larger in terms of the overall objective, and there are $O(\log_{1+\delta}(|C_t|/\delta))$ different values in w_t^* .

Plugging in the approximations of individual time steps does not affect the analysis of movement costs in the proof of Lemma 9, hence the approximation factor follows. We omit the technical details here due to space limit. They can be found in Appendix D of [7]. \square

3 Approximating Dynamic k -Supplier

We present a flow-based algorithm that gives a 3-approximation for Dynamic k -Supplier when $T = 2$, and show it is NP-hard to obtain polynomial-time approximation algorithms for Dynamic k -Supplier with any approximation factor when $T \geq 3$. We also briefly introduce our bi-criteria approximation algorithm for Dynamic k -Supplier with outliers and $T = 2$, while the full version and analysis can be found in Appendix C.

3.1 A 3-Approximation for Dynamic k -Supplier, $T = 2$

In contrast to the NP-hardness of approximating Dynamic k -Supplier for $T \geq 3$, we consider Dynamic k -Supplier when $T = 2$ on general metrics and present a simple flow-based constant approximation. Suppose we are given the client sets C_1, C_2 and F_1, F_2 as candidate facility locations and the movement constraint is $B > 0$.

First, since the optimal radius R^* is obviously the distance between some client and some facility location, we assume we have successfully guessed the optimal radius R^* (using binary search). Next, we construct the following network flow instance $\mathcal{G}(\mathcal{V}, \mathcal{E})$. \mathcal{V} consists of 4 layers of vertices (two layers $\mathcal{L}^{11}, \mathcal{L}^{12}$ for $t = 1$, two layers $\mathcal{L}^{21}, \mathcal{L}^{22}$ for $t = 2$), a source s and sink t . We define the layers and links in \mathcal{G} as follows:

- For each $i \in F_1$, add a vertex in \mathcal{L}^{12} . For $i' \in F_2$, add a vertex in \mathcal{L}^{21} .
- Repeatedly pick an arbitrary client $j \in C_1$ and remove from C_1 every client within distance $2R^*$ from j . Denote these clients a new cluster corresponding to j . Since we have guessed the optimal radius R^* , it is easy to see we can get at most k such clusters. And if there are less than k clusters, we create some extra dummy clusters to obtain exactly k clusters, while dummy clusters do not correspond to any client. For each cluster, add a vertex to \mathcal{L}^{11} . Repeat this for C_2 and form \mathcal{L}^{22} .

- The four layers are arranged in order as $\mathcal{L}^{11}, \mathcal{L}^{12}, \mathcal{L}^{21}, \mathcal{L}^{22}$. With a slight abuse of notation, for $u \in \mathcal{L}^{11}, v \in \mathcal{L}^{12}$, connect them using a link with unit capacity if $d(u, v) \leq R^*$; for $w \in \mathcal{L}^{21}, z \in \mathcal{L}^{22}$, connect them using a link with unit capacity if $d(w, z) \leq R^*$. For $v \in \mathcal{L}^{12}, w \in \mathcal{L}^{21}$, connect them using a link with *unbounded* capacity if $d(v, w) \leq B$.
- Connect every dummy cluster in \mathcal{L}^{11} with every facility location vertex in \mathcal{L}^{12} . Connect every dummy cluster in \mathcal{L}^{22} with every facility location vertex in \mathcal{L}^{21} . Every such link has unit capacity.
- Finally, the source s is connected to every vertex in \mathcal{L}^{11} and the sink t is connected to every vertex in \mathcal{L}^{22} , with every edge having unit capacity.

Lemma 11. $\mathcal{G}(\mathcal{V}, \mathcal{E})$ admits a flow of value k . Moreover, we can obtain a feasible solution of cost at most $3R^*$ from an integral flow of value k in $\mathcal{G}(\mathcal{V}, \mathcal{E})$.

Proof. As an optimal solution with objective R^* , there exist two multi-sets $A_1 \subset F_1, A_2 \subset F_2$ such that $|A_1| = |A_2| = k$ and there exists a perfect matching between them. For any $i \in F_1, i' \in F_2$, if the pair (i, i') appears m times in the perfect matching, define a flow value $f(i, i') = m$ over link (i, i') .

Consider the first time step. For any facility location i and clusters j, j' , either $d(i, j)$ or $d(i, j')$ is larger than R^* , otherwise $d(j, j') \leq 2R^*$, contradicting with our construction. Because A_1 also covers all $j \in C_1$ with radius R^* , for every $j \in \mathcal{L}^{11}$, we can always find a different element $i \in A_1$ such that $d(i, j) \leq R^*$, and we add a unit flow as $f(j, i) = 1$. The same process is repeated for \mathcal{L}^{22} and A_2 .

The total flow between \mathcal{L}^{12} and \mathcal{L}^{21} is now obviously k , since the perfect matching between A_1 and A_2 has size k . After the construction of unit flows for non-dummy clusters, we arbitrarily direct the remaining flows from facility locations to the dummy clusters, one unit each time. Finally, for any cluster with unit flow, define the flow between it and the source/sink to be 1. This completes an integral flow of value k on \mathcal{G} .

For the second part, suppose we have an integral flow \bar{f} of value k on \mathcal{G} . For any facility location $i \in F_1$, denote $g(i)$ the total flow through i . We place $g(i)$ facilities at location i , and repeat the same procedures for $i' \in F_2$. If $\bar{f}(i, i') = m$ for $i \in F_1, i' \in F_2$, move m facilities from i to i' in the transition between 2 time steps.

For any $j' \in C_1$, if j is the cluster center it belongs to, there exists a facility at most $d(j', i) \leq d(j', j) + d(j, i) \leq 3R^*$ away. \square

Theorem 12. *There exists a 3-approximation for Dynamic k -Supplier when $T = 2$.*

Proof. Consider the network flow instance we construct. It only has integer constraints and the coefficient matrix is totally unimodular. Moreover, there exists a flow of value k due to Lemma 11, hence we can efficiently compute an integral solution \bar{f} of value k , thus obtaining a 3-approximation solution. \square

3.2 The Hardness of Approximating Dynamic k -Supplier, $T \geq 3$

We show it is NP-hard to design approximation algorithms for Dynamic k -Supplier with any approximation factor when $T \geq 3$. The proof is via reduction from the perfect 3D matching problem, which is known to be NP-Complete [27].

Theorem 13. *There is no polynomial time algorithm for solving Dynamic k -Supplier with any approximation factor if $T \geq 3$, unless $P = NP$.*

Proof. We reduce an arbitrary instance of perfect 3D-matching to Dynamic k -Supplier to show the NP-hardness. Recall for an instance of perfect 3D-matching, we are given three finite sets A, B, C with $|A| = |B| = |C|$, and a triplet set $\mathcal{T} \subset A \times B \times C$. Suppose $|A| = n$ and $|\mathcal{T}| = m$, and we are asked to decide whether there exists a subset $\mathcal{S} \subset \mathcal{T}$, such that $|\mathcal{S}| = n$, and each element in A, B, C appears exactly once in some triplet in \mathcal{S} . We construct the following graph $G = (V, E)$, where V, E are initially empty.

- For each triplet $g = (a, b, c) \in \mathcal{T}$, add three new vertices a_g, b_g, c_g to V correspondingly. Connect a_g, b_g with an edge of length α . Connect b_g, c_g with an edge of length α .

- Denote V_A all the vertices that correspond to vertices in A . Similarly for V_B and V_C .
- For any two vertices in V_A corresponding to the same element $a \in A$, connect them with an edge of length 1. Repeat the same procedure for V_B, V_C .

Assume we are able to solve Dynamic k -Supplier for $T = 3$ with an approximation factor α . We solve Dynamic k -Supplier for G on its graph metric d_G , with $k = n$ and the movement constraint $B = \alpha$, where the client sets are $\{V_A, V_B, V_C\}$ and facility sets are $\{V_A, V_B, V_C\}$ for the three time steps, respectively.

It is easy to see that the reduced Dynamic k -Supplier instance has covering radius $R^* = 1$ if and only if there exists a perfect 3D-matching, otherwise the covering radius is at least $2\alpha + 1$. Since our approximation factor is α , this concludes the NP-hardness of approximation algorithms with any factor for Dynamic k -Supplier when $T \geq 3$. \square

3.3 A Bi-criteria Approximation for Dynamic k -Supplier with Outliers

Lastly, we present our bi-criteria approximation algorithm that solves Dynamic k -Supplier, when $T = 2$ and outliers are allowed. As a useful ingredient, let us first briefly review the *m-budgeted bipartite matching* problem. The input consists of a bipartite graph $G = (V, E)$, and each edge $e \in E$ is associated with a weight $w(e) \geq 0$ and m types of lengths $f_i(e) \geq 0, i = 1, \dots, m$. The problem asks for a maximum weight matching M with m budget constraints, where the i th constraint is that the sum of all f_i lengths in M is no more than L_i , i.e. $\sum_{e \in M} f_i(e) \leq L_i$. When the number of constraints m is a constant, a pure $(1 - \epsilon)$ -approximation algorithm for any constant $\epsilon > 0$ is devised by Grandoni *et al.* [20].

Sketch: Due to space limit, we provide a sketch here and defer the full details to Appendix C. Consider Dynamic k -Supplier with outliers and $T = 2$. In the solution, we place k facilities for time $t = 1$, serving in total at least l_1 clients in C_1 , then move each of these facilities for a distance at most B to serve at least l_2 clients in C_2 , and the maximum service distance is our minimization goal. Clearly, the optimal solution R^* only has a polynomial number of possible values and can be guessed efficiently, so we assume that R^* is known to us in the following analysis.

For a fixed R^* , denote c_i the number of clients that facility location i can serve within distance R^* . We assign two lengths $f_1(e) = c_i, f_2(e) = c_{i'}$ and weight $w(e) = 1$ for every candidate edge $e = (i, i')$, where $i \in F_1, i' \in F_2$. By duplicating each possible facility location in F_1 and F_2 and only allowing vertices within distance B to be matched, the required solution can be fully represented by a k -cardinality matching M between F_1 and F_2 . Let us temporarily assume that any client miraculously contribute only once to the total number of clients served. Then the problem naturally translates to deciding whether there exists a bipartite matching M between F_1 and F_2 (with candidate facility locations duplicated) with weight k , such that the sum of all f_1 lengths in M is at least l_1 , and the sum of all f_2 lengths in M is at least l_2 .

This new problem is very similar to *2-budgeted bipartite matching*, but there are still some major differences. In the approximation algorithm in [20], every integral matching M is obtained by first finding a *feasible* fractional matching M' , which has at most $2m$ edges being fractional, and then dropping these fractionally-matched edges completely. Back to our problem where $m = 2$. If we obtain such a fractional solution M' which satisfies the constraints and only has at most 4 fractional edges, we would like to find an integral matching M in a way that uses more "budget" instead of using less, so as to cover at least as many clients as M' does and not violate any budget constraint (in other words, outlier constraints), and we have to drop these fractional edges again from M' .

Contrary to *2-budgeted bipartite matching*, we want to control the portion of budget dropped in this case. We achieve this by guessing a constant number of edges, which has either the top- θ f_1 lengths or top- θ f_2 lengths in the optimal solution, using a suitably chosen constant $\theta > 0$. We are able to devise a bi-criteria approximation algorithm that violates both budget constraints by any small constant ϵ -portion. The bi-criteria method is developed in line with the multi-criteria approximation schemes in [20].

To fully avoid counting any served client multiple times, whenever we duplicate a facility location, we make sure that only one copy induces non-zero lengths on edges that reside on it. We also use a greedy

algorithm to remove some facility locations in F_1, F_2 and form client clusters around the remaining ones. Now, instead of defining c_i as the number of clients that facility location i can serve within distance R^* , we change c_i to the number of clients that are gathered around i . More specifically, for client set C_t and facility location set F_t , we find a subset $F'_t \subset F_t$ and a corresponding sub-partition $\{K_i\}_{i \in F'_t}$ of C_t (i.e., K_i s are pair-wise disjoint and their union is a subset of C_t), such that $\forall j \in K_i, d(i, j) \leq 3R^*$ and we define $c_i = |K_i|$ for $i \in F'_t$, $c_i = 0$ for $i \in F_t \setminus F'_t$. Using this method, every client is counted at most once in all c_i s, hence its contribution to the total number is always at most 1. The same filtering process can be found in [22]. See Appendix C for the details.

4 Future Directions

We list some interesting future directions and open problems.

1. It would be very interesting to remove the dependency of γ (the coefficient of movement cost) and ϵ (the lower bound of the weight) from the approximation factor for Dynamic Ordered k -Median in Theorem 10, or show such dependency is inevitable. We leave it as an important open problem. We note that a constant approximation factor for Dynamic Ordered k -Median without depending on γ would imply a constant approximation for stochastic k -server, for which only a logarithmic-factor approximation algorithm is known [15].
2. Our approximation algorithm for Dynamic Ordered k -Median is based on the technique developed in Byrka *et al.* [7]. The original ordered k -median problem has subsequently seen improved approximation results in [9, 10]. We did not try hard to optimize the constant factors. Nevertheless, it is an interesting future direction to further improve the constant approximation factors by leveraging the techniques from [9, 10] or other ideas.
3. From Theorem 13, we can see that Dynamic k -Supplier is hard to approximate when $T \geq 3$. However, it makes sense to relax the hard constraint B (we allow the distance a facility can move be at most αB for some constant α).

It is possible to formulate other concrete problems that naturally fit into the dynamic clustering theme and are well motivated by realistic applications, but not yet considered in the paper. For example, one can use the k -median objective for the service cost and the maximum distance of any facility movement as the movement cost. One can also consider combining the cost in more general fashion like in [10], or extending the problems to the fault-tolerant version [21, 28, 34] or the capacitated version [14, 31].

5 Acknowledgements

Shichuan Deng and *Jian Li* are supported in part by the National Natural Science Foundation of China Grant 61822203, 61772297, 61632016, 61761146003, the Zhongguancun Haihua Institute for Frontier Information Technology, Turing AI Institute of Nanjing, and Xi'an Institute for Interdisciplinary Information Core Technology. *Yuval Rabani* is supported in part by ISF grant number 2553-17.

We would like to thank Chaitanya Swamy for kindly pointing out studies relevant to our results. We also want to thank the reviewers for their insightful and constructive comments.

References

- [1] Sara Ahmadian, Zachary Friggstad, and Chaitanya Swamy. Local-search based approximation algorithms for mobile facility location problems. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, pages 1607–1621, 2013. doi:10.1137/1.9781611973105.115.

- [2] Hyung-Chan An, Ashkan Norouzi-Fard, and Ola Svensson. Dynamic facility location via exponential clocks. *ACM Transactions on Algorithms*, 13(2):1–20, 2017. doi:10.1145/2928272.
- [3] Ali Aouad and Danny Segev. The ordered k -median problem: surrogate models and approximation algorithms. *Mathematical Programming*, 177(1-2):55–83, 2019. doi:10.1007/s10107-018-1259-3.
- [4] Vijay Arya, Naveen Garg, Rohit Khandekar, Adam Meyerson, Kamesh Munagala, and Vinayaka Pandit. Local search heuristics for k -median and facility location problems. *SIAM Journal on Computing*, 33(3):544–562, 2004. doi:10.1137/S0097539702416402.
- [5] David L Black, Anoop Gupta, and Wolf-Dietrich Weber. Competitive management of distributed shared memory. *COMPCON Spring 89*, pages 184–185, 1989. doi:10.1109/COMPCON.1989.301925.
- [6] Jarosław Byrka, Thomas Pensyl, Bartosz Rybicki, Aravind Srinivasan, and Khoa Trinh. An improved approximation for k -median and positive correlation in budgeted optimization. *ACM Transactions on Algorithms*, 13(2):1–31, 2017. doi:10.1145/2981561.
- [7] Jarosław Byrka, Krzysztof Sornat, and Joachim Spoerhase. Constant-factor approximation for ordered k -median. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 620–631, 2018. doi:10.1145/3188745.3188930.
- [8] Deeparnab Chakrabarty, Prachi Goyal, and Ravishankar Krishnaswamy. The non-uniform k -center problem. In *43rd International Colloquium on Automata, Languages, and Programming*, volume 55, page 67, 2016. doi:10.4230/LIPIcs.ICALP.2016.67.
- [9] Deeparnab Chakrabarty and Chaitanya Swamy. Interpolating between k -median and k -center: Approximation algorithms for ordered k -median. In *45th International Colloquium on Automata, Languages, and Programming*, volume 107, page 29, 2018. doi:10.4230/LIPIcs.ICALP.2018.29.
- [10] Deeparnab Chakrabarty and Chaitanya Swamy. Approximation algorithms for minimum norm and ordered optimization problems. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 126–137, 2019. doi:10.1145/3313276.3316322.
- [11] Moses Charikar, Samir Khuller, David M Mount, and Giri Narasimhan. Algorithms for facility location problems with outliers. In *Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*, pages 642–651, 2001. URL: <http://dl.acm.org/citation.cfm?id=365411.365555>.
- [12] Moses Charikar and Shi Li. A dependent LP-rounding approach for the k -median problem. In *Proceedings of the 39th International Colloquium on Automata, Languages, and Programming-Volume Part I*, pages 194–205, 2012. doi:10.1007/978-3-642-31594-7_17.
- [13] Danny Z Chen, Jian Li, Hongyu Liang, and Haitao Wang. Matroid and knapsack center problems. *Algorithmica*, 75(1):27–52, 2016. doi:10.1007/s00453-015-0010-1.
- [14] Julia Chuzhoy and Yuval Rabani. Approximating k -median with non-uniform capacities. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 952–958, 2005. URL: <http://dl.acm.org/citation.cfm?id=1070432.1070569>.
- [15] Sina Dehghani, Soheil Ehsani, MohammadTaghi Hajiaghayi, Vahid Liaghat, and Saeed Seddighin. Stochastic k -server: How should Uber work? In *44th International Colloquium on Automata, Languages, and Programming*, volume 80, page 126, 2017. doi:10.4230/LIPIcs.ICALP.2017.126.
- [16] Erik D Demaine, MohammadTaghi Hajiaghayi, Hamid Mahini, Amin S Sayedi-Roshkhar, Shayan Oveisgharan, and Morteza Zadimoghaddam. Minimizing movement. *ACM Transactions on Algorithms*, 5(3):30, 2009. doi:10.1145/1541885.1541891.

- [17] David Eisenstat, Claire Mathieu, and Nicolas Schabanel. Facility location in evolving metrics. In *International Colloquium on Automata, Languages, and Programming*, pages 459–470, 2014. doi:[10.1007/978-3-662-43951-7_39](https://doi.org/10.1007/978-3-662-43951-7_39).
- [18] Zachary Friggstad and Mohammad R Salavatipour. Minimizing movement in mobile facility location problems. *ACM Transactions on Algorithms*, 7(3):28, 2011. doi:[10.1145/1978782.1978783](https://doi.org/10.1145/1978782.1978783).
- [19] Rajiv Gandhi, Samir Khuller, Srinivasan Parthasarathy, and Aravind Srinivasan. Dependent rounding and its applications to approximation algorithms. *Journal of the ACM*, 53(3):324–360, 2006. doi:[10.1145/1147954.1147956](https://doi.org/10.1145/1147954.1147956).
- [20] Fabrizio Grandoni, R Ravi, Mohit Singh, and Rico Zenklusen. New approaches to multi-objective optimization. *Mathematical Programming*, 146(1-2):525–554, 2014. doi:[10.1007/s10107-013-0703-7](https://doi.org/10.1007/s10107-013-0703-7).
- [21] Mohammadtaghi Hajiaghayi, Wei Hu, Jian Li, Shi Li, and Barna Saha. A constant factor approximation algorithm for fault-tolerant k -median. *ACM Transactions on Algorithms*, 12(3):36, 2016. doi:[10.1145/2854153](https://doi.org/10.1145/2854153).
- [22] David G Harris, Thomas Pensyl, Aravind Srinivasan, and Khoa Trinh. A lottery model for center-type problems with outliers. *ACM Transactions on Algorithms*, 15(3):1–25, 2019. doi:[10.1145/3311953](https://doi.org/10.1145/3311953).
- [23] Dorit S Hochbaum and David B Shmoys. A best possible heuristic for the k -center problem. *Mathematics of Operations Research*, 10(2):180–184, 1985. doi:[10.1287/moor.10.2.180](https://doi.org/10.1287/moor.10.2.180).
- [24] Dorit S Hochbaum and David B Shmoys. A unified approach to approximation algorithms for bottleneck problems. *Journal of the ACM*, 33(3):533–550, 1986. doi:[10.1145/5925.5933](https://doi.org/10.1145/5925.5933).
- [25] Kamal Jain and Vijay V Vazirani. Primal-dual approximation algorithms for metric facility location and k -median problems. In *40th Annual Symposium on Foundations of Computer Science*, pages 2–13, 1999. doi:[10.1109/SFFCS.1999.814571](https://doi.org/10.1109/SFFCS.1999.814571).
- [26] Kamal Jain and Vijay 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. doi:[10.1145/375827.375845](https://doi.org/10.1145/375827.375845).
- [27] Richard M Karp. Reducibility among combinatorial problems. *50 Years of Integer Programming 1958-2008*, pages 219–241, 2010. doi:[10.1007/978-3-540-68279-0_8](https://doi.org/10.1007/978-3-540-68279-0_8).
- [28] Samir Khuller, Robert Pless, and Yoram J Sussmann. Fault tolerant k -center problems. *Theoretical Computer Science*, 242(1-2):237–245, 2000. doi:[10.1016/S0304-3975\(98\)00222-9](https://doi.org/10.1016/S0304-3975(98)00222-9).
- [29] Ravishankar Krishnaswamy, Amit Kumar, Viswanath Nagarajan, Yogish Sabharwal, and Barna Saha. The matroid median problem. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete algorithms*, pages 1117–1130, 2011. doi:[10.1137/1.9781611973082.84](https://doi.org/10.1137/1.9781611973082.84).
- [30] Ravishankar Krishnaswamy, Shi Li, and Sai Sandeep. Constant approximation for k -median and k -means with outliers via iterative rounding. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 646–659, 2018. doi:[10.1145/3188745.3188882](https://doi.org/10.1145/3188745.3188882).
- [31] Shi Li. Approximating capacitated k -median with $(1 + \epsilon)k$ open facilities. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 786–796, 2016. doi:[10.1137/1.9781611974331.ch56](https://doi.org/10.1137/1.9781611974331.ch56).
- [32] Shi Li and Ola Svensson. Approximating k -median via pseudo-approximation. In *Proceedings of the forty-fifth annual ACM symposium on Theory of Computing*, pages 901–910, 2013. doi:[10.1145/2488608.2488723](https://doi.org/10.1145/2488608.2488723).

- [33] Chaitanya Swamy. Improved approximation algorithms for matroid and knapsack median problems and applications. *ACM Transactions on Algorithms*, 12(4):49:1–49:22, 2016. doi:10.1145/2963170.
- [34] Chaitanya Swamy and David B Shmoys. Fault-tolerant facility location. *ACM Transactions on Algorithms*, 4(4):1–27, 2008. doi:10.1145/1383369.1383382.
- [35] Jeffery Westbrook. Randomized algorithms for multiprocessor page migration. *SIAM Journal on Computing*, 23(5):951–965, 1994. doi:10.1137/S0097539791199796.

A Missing Proofs and Algorithms in Section 2

A.1 Filtering and Matching

Algorithm 1 OBLIVIOUSFILTER(x, y)

```

1: for all  $t \in [T]$  do
2:    $C'_t \leftarrow \emptyset, C''_t \leftarrow C_t$ 
3:   while  $C''_t$  is nonempty do
4:     Choose  $j \in C''_t$  with the smallest  $d_{\text{av}}^{(t)}(j)$ 
5:      $C'_t \leftarrow C'_t \cup \{j\}, C''_t \leftarrow C''_t \setminus \{j\}$ . Delete each  $j' \in C''_t$  that has  $d(j, j') \leq 4d_{\text{av}}^{(t)}(j)$ 
6:   end while
7:   for all  $j \in C'_t$  do
8:      $n_j^{(t)} \leftarrow \arg \min_{j' \in C'_t, j' \neq j} d(j, j')$ ;  $R_j^{(t)} \leftarrow \frac{1}{2}d(j, n_j^{(t)})$ ;  $\mathcal{U}_j^{(t)} \leftarrow E_j^{(t)} \cap \text{Ball}^o(j, R_j)$ 
9:   end for
10:   $P_t \leftarrow \emptyset, C''_t \leftarrow C'_t$ 
11:  while  $\exists j \in C''_t$  such that  $n_j^{(t)} \in C''_t$  do
12:    Choose  $j \in C''_t$  such that  $d(j, n_j^{(t)})$  is minimized
13:     $P_t \leftarrow P_t \cup \{(j, n_j^{(t)})\}$ . Delete from  $C''_t$   $j$  and  $n_j^{(t)}$ 
14:  end while
15:  while  $C''_t$  is nonempty do
16:    Choose arbitrary  $j \in C''_t$ .  $P_t \leftarrow P_t \cup \{(j)\}$ . Delete  $j$  from  $C''_t$ 
17:  end while
18: end for
19: return  $\{C'_t\}_{t \in [T]}, \{\mathcal{U}_j^{(t)}\}_{j \in C'_t}, \{P_t\}_{t \in [T]}$ 

```

It is naturally desired that we discuss the differences between our algorithm and the celebrated filtering and matching algorithm used in Charikar and Li [12]. Fix any time step t , our algorithm has the same filtering process as in [12], while the matching process is different. For the filtered client set C'_t , Charikar and Li use a simple greedy algorithm, matching the closest unmatched pair in C'_t whenever possible, so their matching phase leaves at most one client in C'_t unmatched.

In our algorithm, given C'_t , we first compute for every $j \in C'_t$ its closest neighbor $n_j^{(t)} \in C'_t$. Whenever there exists unmatched j such that its closest neighbor $n_j^{(t)}$ is also unmatched, we choose such a j that minimizes $d(j, n_j^{(t)})$, and match $(j, n_j^{(t)})$. Notice that this process may leave an arbitrary number of clients in C'_t unmatched, or in other words, in singleton pairs.

However, this potential problem is easily addressed by the careful design of the network flow instance. If j ends up in a singleton pair, it is easy to see that $n_j^{(t)}$ has to be matched with $j' \neq j$ with $d(n_j^{(t)}, j') \leq d(n_j^{(t)}, j)$. Our stochastically-rounded integral solution guarantees that the probability of opening an facility in the singleton pair (j) is exactly $y^{(t)}(\mathcal{U}_j^{(t)})$, as well as the existence of an open facility in $\mathcal{U}_{n_j^{(t)}}^{(t)} \cup \mathcal{U}_{j'}^{(t)}$. Therefore,

since the marginal distribution on $\mathcal{U}_j^{(t)}$ is preserved, we are able to formulate an argument that is similar to Lemma 3 in Charikar and Li [12], but serves the ordered objectives better. For more details, please see Appendix A.4.

A.2 The Algorithm for Rerouting

Algorithm 2 REROUTE(\bar{f}, \mathcal{N})

```

1: for all  $t \in [T]$  do
2:   Initialize multi-set  $A_t \leftarrow \emptyset$ 
3:   for all  $p \in P_t$  do
4:     if  $\bar{f}(\mathcal{L}(p), \mathcal{R}(p)) = 2$  then
5:       Pick  $\mathcal{L}(i_1), \mathcal{L}(i_2)$  where the flow passes through,  $A_t \leftarrow A_t \cup \{i_1, i_2\}$ 
6:     else if  $\bar{f}(\mathcal{L}(p), \mathcal{R}(p)) = 1$  then
7:       if  $p = (j_1, j_2), n_{j_1}^{(t)} = j_2, n_{j_2}^{(t)} = j_1$  (i.e.  $j_1$  and  $j_2$  are closest to each other) then
8:         Pick  $\mathcal{L}(i)$  where the unit flow over  $(\mathcal{L}(p), \mathcal{R}(p))$  passes through,  $A_t \leftarrow A_t \cup \{i\}$ 
9:       else if  $p = (j_1, j_2), n_{j_1}^{(t)} = j_2, n_{j_2}^{(t)} \neq j_1$  then
10:        if  $\bar{f}$  passes through the same bundle in  $L_{6t-5}$  and  $L_{6t}$  then
11:          Pick  $\mathcal{L}(i)$  where the unit flow over  $(\mathcal{L}(p), \mathcal{R}(p))$  passes through,  $A_t \leftarrow A_t \cup \{i\}$ 
12:        else
13:          Pick  $i \in \mathcal{U}_{j_2}^{(t)}$  where the unit flow over  $(\mathcal{L}(p), \mathcal{R}(p))$  passes through,  $A_t \leftarrow A_t \cup \{i\}$ 
14:        end if
15:      else if  $p = (j)$  then
16:        Pick  $\mathcal{L}(i)$  where the unit flow over  $(\mathcal{L}(p), \mathcal{R}(p))$  passes through,  $A_t \leftarrow A_t \cup \{i\}$ 
17:      end if
18:    end if
19:  end for
20: end for
21: return  $\{A_t\}_{t=1}^T$ 

```

A.3 Missing Proof of Lemma 7

Proof. If $T = 2$, recall that we choose A_1, A_2 solely based on \bar{f} and the links between L_6 and L_7 . Each flow on link $e = (\mathcal{R}(i), \mathcal{L}(i'))$ between L_6 and L_7 is rounded to 1 with probability $z_{ii'}^{(1)}$. Recall that $d(A, A')$ is the cost of the minimum weight matching between A, A' , the expectation of total movement cost is exactly,

$$\mathbb{E}[d(A_1, A_2)] = \mathbb{E} \left[\sum_{i \in F_1, i' \in F_2} \Pr[\bar{f}(\mathcal{R}(i), \mathcal{L}(i')) = 1] \cdot d(i, i') \right] = \sum_{i \in F_1} \sum_{i' \in F_2} d(i, i') z_{ii'}^{(1)}.$$

From now on we assume $T \geq 3$. Denote $\bar{f}(L)$ the multi-set of facilities that \bar{f} passes through in layer L . According to Algorithm REROUTE, $|\bar{f}(L)| = k$ for each $L \in \{L_{6t-5}, L_{6t}\}$, and using triangle inequality, the expected total movement cost is at most

$$\mathbb{E} \left[\sum_{t \in [T-1]} d(A_t, A_{t+1}) \right] \leq \sum_{t \in [T]} \mathbb{E} [d(\bar{f}(L_{6t-5}), \bar{f}(L_{6t}))] + \sum_{t \in [T-1]} \mathbb{E} [d(\bar{f}(L_{6t}), \bar{f}(L_{6t+1}))].$$

Denote the random variable $\Delta_t \triangleq d(\bar{f}(L_{6t-5}), \bar{f}(L_{6t}))$. Similarly, we have

$$\mathbb{E} [d(\bar{f}(L_{6t}), \bar{f}(L_{6t+1}))] \leq \sum_{i \in F_t} \sum_{i' \in F_{t+1}} d(i, i') z_{ii'}^{(t)},$$

and we only need to bound the expectation of Δ_t . We fix t and leave out the superscripts in the following. The minimum-weight matching between $\bar{f}(L_{6t-5})$ and $\bar{f}(L_{6t})$ can be upper bounded by considering the set of matchings that only match facilities in the same pair. Consider any pair $p \in P_t$ and define $\Delta_{t,p}$ the additional cost of moving facilities within p . Therefore, we always obtain $\Delta_t \leq \sum_{p \in P_t} \Delta_{t,p}$, since for every $i \in F_t - \bigcup_{j \in C'_t} \mathcal{U}_j^{(t)}$ that are matched, the design of \mathcal{N} makes sure that no additional movement is required for i .

To match the facilities and bound $\Delta_{t,p}$, we prioritize the facilities in the same bundle, and only make a cross-bundle matching if we have to. For example for $p = (j_1, j_2)$, if $\bar{f}(\mathcal{L}(p), \mathcal{R}(p)) = 2$, $\mathcal{L}(i_1), \mathcal{L}(i_2)$ with $i_1 \in \mathcal{U}_{j_1}, i_2 \in \mathcal{U}_{j_2}$ are picked, and $\mathcal{R}(i'_1), \mathcal{R}(i'_2)$ with $i'_1 \in \mathcal{U}_{j_1}, i'_2 \in \mathcal{U}_{j_2}$ are picked, we match i_1 with i'_1 and i_2 with i'_2 .

Using triangle inequality, we have $d(i_1, i'_1) \leq d(i_1, j_1) + d(i'_1, j_1)$, then it is not hard to obtain the following for $p = (j_1, j_2)$, i.e. when p is not singleton,

$$\begin{aligned} \mathbb{E}[\Delta_{t,p}] &\leq \sum_{i \in \mathcal{U}_{j_1}} d(i, j_1) \cdot \Pr[\bar{f}(\mathcal{L}(i), \mathcal{L}(j_1)) = 1] + d(i, j_1) \cdot \Pr[\bar{f}(\mathcal{R}(j_1), \mathcal{R}(i)) = 1] \\ &\quad + \sum_{i \in \mathcal{U}_{j_2}} d(i, j_2) \cdot \Pr[\bar{f}(\mathcal{L}(i), \mathcal{L}(j_2)) = 1] + d(i, j_2) \cdot \Pr[\bar{f}(\mathcal{R}(j_2), \mathcal{R}(i)) = 1] \\ &\quad + \Pr[\bar{f}(\mathcal{L}(p), \mathcal{R}(p)) = 1, L_{6t-5}, L_{6t} \text{ disagree on } \mathcal{U}_{j_1}, \mathcal{U}_{j_2}] \cdot d(j_1, j_2) \\ &\leq 2d_{\text{av}}(j_1) + 2d_{\text{av}}(j_2) + \Pr[\bar{f}(\mathcal{L}(p), \mathcal{R}(p)) = 1, L_{6t-5}, L_{6t} \text{ disagree on } \mathcal{U}_{j_1}, \mathcal{U}_{j_2}] \cdot d(j_1, j_2), \end{aligned}$$

where $\bar{f}(\mathcal{L}(p), \mathcal{R}(p)) = 1, L_{6t-5}, L_{6t}$ disagree on $\mathcal{U}_{j_1}, \mathcal{U}_{j_2}$ means the unit flow on the link $(\mathcal{L}(p), \mathcal{R}(p))$ passes through different bundles in L_{6t-5} and L_{6t} .

Since (j_1, j_2) is a pair, w.l.o.g. we have j_2 is the closest neighbor of j_1 among C'_t , $R_{j_1} = 0.5d(j_1, j_2)$ and $d_{\text{av}}(j_1) \geq (1 - y(\mathcal{U}_{j_1})) \cdot R_{j_1} = 0.5(1 - y(\mathcal{U}_{j_1})) \cdot d(j_1, j_2)$, hence $(1 - y(\mathcal{U}_{j_1}))d(j_1, j_2) \leq 2d_{\text{av}}(j_1)$. Furthermore, we can obtain a simple bound on the probability

$$\begin{aligned} &\Pr[\bar{f}(\mathcal{L}(p), \mathcal{R}(p)) = 1, L_{6t-5}, L_{6t} \text{ disagree on } \mathcal{U}_{j_1}, \mathcal{U}_{j_2}] \\ &= \Pr[\bar{f}(\mathcal{L}(\mathcal{U}_{j_1}), \mathcal{L}(p)) = 1, \bar{f}(\mathcal{L}(p), \mathcal{R}(p)) = 1, \bar{f}(\mathcal{R}(p), \mathcal{R}(\mathcal{U}_{j_2})) = 1] \\ &\quad + \Pr[\bar{f}(\mathcal{L}(\mathcal{U}_{j_2}), \mathcal{L}(p)) = 1, \bar{f}(\mathcal{L}(p), \mathcal{R}(p)) = 1, \bar{f}(\mathcal{R}(p), \mathcal{R}(\mathcal{U}_{j_1})) = 1] \\ &\leq \Pr[\bar{f}(\mathcal{R}(p), \mathcal{R}(\mathcal{U}_{j_1})) = 0] + \Pr[\bar{f}(\mathcal{L}(\mathcal{U}_{j_1}), \mathcal{L}(p)) = 0] \\ &\leq 2(1 - y(\mathcal{U}_{j_1})). \end{aligned}$$

Therefore, the expectation above can be further bounded as

$$\mathbb{E}[\Delta_{t,p}] \leq 2d_{\text{av}}(j_1) + 2d_{\text{av}}(j_2) + 2d(j_1, j_2) \cdot (1 - y(\mathcal{U}_{j_1})) \leq 6d_{\text{av}}(j_1) + 2d_{\text{av}}(j_2).$$

By summing over all pairs and all time steps, we have

$$\sum_{t \in [T]} \mathbb{E}[\Delta_t] \leq 6 \sum_{t \in [T]} \sum_{j \in C'_t} d_{\text{av}}^{(t)}(j) \leq 6 \sum_{t \in [T]} \sum_{j \in C_t} d_{\text{av}}^{(t)}(j).$$

□

A.4 Missing Proof of Lemma 8

We first need the following lemma for filtered clients and bundles.

Lemma 14. *For any $j \in C'_t$, we have $y^{(t)}(\mathcal{U}_j^{(t)}) = \sum_{i \in \mathcal{U}_j^{(t)}} y_i^{(t)} \in [1/2, 1]$. For any $j, j' \in C'_t$, $j \neq j'$, we have $d(j, j') \geq 4 \max\{d_{\text{av}}^{(t)}(j), d_{\text{av}}^{(t)}(j')\}$, and $\forall i \in \mathcal{U}_j^{(t)}, i' \in \mathcal{U}_{j'}^{(t)}, d(i, i') > 0$.*

This is a fairly standard result used in several approximation algorithms for k -median (see e.g. [7, 12]). Thus, we omit the proof. We are ready to present the proof of Lemma 8.

Proof. We fix and leave out the superscript $t \in [T]$ in the following. The proof follows the proof of Lemma 4.5 by Byrka *et al.* in [7], by upper-bounding the service cost of every client $j \in C_t$ using a deterministic budget D_j and another random budget variable X_j so that $\mathbb{E}_{A_t}[d(j, A_t)] \leq D_j + \mathbb{E}_{A_t}[X_j]$, and aim to prove $D_j \leq a_1 \cdot h, \mathbb{E}[X_j] \leq a_2 \cdot d^{-h}_{\text{av}}(j)$ for some constant a_1, a_2 . In this case, for the expected ordered sum $\mathbb{E}_{A_t}[\text{rect}(m, |C_t|) \cdot d(C_t, A_t)^\downarrow]$, since we are only paying for at most m of the deterministic budgets, it follows that,

$$\mathbb{E}_{A_t}[\text{rect}(m, |C_t|) \cdot d(C_t, A_t)^\downarrow] \leq m \cdot \max_j D_j + \sum_{j \in C_t} \mathbb{E}[X_j] \leq a_1 \cdot mh + a_2 \cdot \sum_{j \in C_t} d^{-h}_{\text{av}}(j).$$

In the following, for any fixed client j , we progressively charge parts of $d(j, A_t)$ to either D_j or X_j , where we initially charge 0 to D_j and X_j for any $j \in C_t$.

Consider a client j , there exists $j' \in C'_t$ (probably $j = j'$ when $j \in C'_t$) such that $d(j, j') \leq 4d_{\text{av}}(j)$ and $d_{\text{av}}(j') \leq d_{\text{av}}(j)$. Notice by definition of d^{-h} , it is obvious that $d_{\text{av}}(j) \leq d^{-h}_{\text{av}}(j) + h$, hence if facility i serves j' , using triangle inequality, $\mathbb{E}_i[d(i, j)] \leq \mathbb{E}_i[d(i, j')] + d(j, j') \leq \mathbb{E}[d(i, j')] + 4d_{\text{av}}(j) \leq \mathbb{E}[d(i, j')] + 4d^{-h}_{\text{av}}(j) + 4h$. We charge $4h$ to D_j and $4d^{-h}_{\text{av}}$ with probability 1 to X_j , so $D_j = 4h$ and $\mathbb{E}[X_j] = 4d^{-h}_{\text{av}}(j)$ for now.

Next, we charge the stochastic service cost $d(j', A_t)$ to D_j or X_j , and we only consider the case where j' and its nearest neighbor $j'' \in C'_t$ are *not* matched (the case where they are matched is simpler but somewhat different, which will be explained later).

Fix some constant $\beta > 5$ to be determined later. First, if any facility in the open ball $\text{Ball}^o(j', \beta h)$ is opened, we charge βh to D_j ; then if a facility in $\mathcal{U}_{j'} \setminus \text{Ball}^o(j', \beta h)$ is opened instead, we charge the stochastic cost to X_j . In Algorithm **REROUTE**, if j' is in a singleton pair, the marginal distribution over bundle $\mathcal{U}_{j'}$ is not changed, hence the expectation of this stochastic cost is at most

$$\begin{aligned} \sum_{i \in \mathcal{U}_{j'} \setminus \text{Ball}^o(j', \beta h)} x_{ij'} d(i, j') &= \sum_{i \in \mathcal{U}_{j'} \setminus \text{Ball}^o(j', \beta h)} x_{ij'} d^{-h}(i, j') \\ &\leq \sum_{i \in E_{j'} \setminus \text{Ball}^o(j', \beta h)} x_{ij'} d^{-h}(i, j') \triangleq d^{-h}_{\text{far}}(j') \leq d^{-h}_{\text{av}}(j'), \end{aligned}$$

where the last inequality is due to the definition of $d^{-h}_{\text{av}}(j')$; when j' is not in a singleton pair, since j' is not matched to its closest neighbor j'' , it has to be matched to \tilde{j} , and j' is the closest neighbor of \tilde{j} . The marginal distribution over $\mathcal{U}_{j'}$ is modified in this case, because whenever the rounded \tilde{f} disagrees on the bundles $\mathcal{U}_{j'}$ versus $\mathcal{U}_{\tilde{j}}$ in layers L_{6t-5} and L_{6t} , the algorithm always chooses the facility that is in $\mathcal{U}_{j'}$. Considering that $\mathcal{U}_{j'}$ has probability mass at least 1/2 itself, the probabilities of each $i \in \mathcal{U}_{j'}$ are perturbed by the same multiplicative factor at most 2, we can still bound the expectation of the stochastic service cost outside $\text{Ball}^o(j', \beta h)$ as,

$$\sum_{i \in \mathcal{U}_{j'} \setminus \text{Ball}^o(j', \beta h)} 2x_{ij'} d(i, j') \leq 2d^{-h}_{\text{far}}(j').$$

Then, if none of the facilities in $\mathcal{U}_{j'}$ is open, because j' is not matched to its nearest neighbor j'' , j'' has to be matched with another j''' such that $d(j'', j''') \leq d(j', j'')$. Denote $R = R_{j'} = \frac{1}{2}d(j', j'')$, since our algorithm guarantees that there exists an open facility in $\mathcal{U}_{j''} \cup \mathcal{U}_{j'''}$, j' could be served by such a facility at a distance of at most

$$d(j', j'') + d(j'', j''') + \max\{R_{j''}, R_{j'''}\} \leq 2R + 2R_{j''} + \max\{R_{j''}, R_{j'''}\} \leq 5R.$$

But R could be unbounded with respect to our goal, $a_1 h$ or $a_2 d^{-h}_{\text{av}}(j)$, so we need to bound the probability that this case happens. Nonetheless, if $R \leq \beta h$, we simply charge $5R \leq 5\beta h$ to D_j , hence we assume $R > \beta T$ in the following.

When j' is in a singleton pair, because $R > \beta h \Rightarrow E_{j'} \cap \text{Ball}^o(j', \beta h) \subseteq \mathcal{U}_{j'}$,

$$d^{-h}_{\text{far}}(j') = \sum_{i \in E_{j'} \setminus \text{Ball}^o(j', \beta h)} x_{ij'} d^{-h}(i, j') \geq \sum_{i \in E_{j'} \setminus \mathcal{U}_{j'}} x_{ij'} d^{-h}(i, j') \geq R \sum_{i \in E_{j'} \setminus \mathcal{U}_{j'}} x_{ij'} = R(1 - y(\mathcal{U}_{j'})),$$

therefore according to our rounding algorithm, there is a probability exactly $1 - y(\mathcal{U}_{j'})$ that none of the facilities in $\mathcal{U}_{j'}$ is chosen, and we will charge $5R$ to X_j , resulting in an increase of at most $5R(1 - y(\mathcal{U}_{j'})) \leq 5d_{\text{far}}^{-h}(j')$ to the expectation $\mathbb{E}[X_j]$.

When j' is in a normal pair with \tilde{j} , we know j' has to be the nearest neighbor of \tilde{j} , the marginal distribution of not selecting any facility in $\mathcal{U}_{j'}$ is actually decreased, hence the argument above still holds true.

It is worth pointing out that, when j' is actually matched to its nearest neighbor j'' , but j' is not the nearest neighbor of j'' , some additional arguments need to be made clear. Due to our algorithm **REROUTE**, the probability that none of the facilities in $\mathcal{U}_{j'}$ is picked is increased, whenever the integral flow \bar{f} disagrees on $\mathcal{U}_{j'}$ and $\mathcal{U}_{j''}$ in layers L_{6t-5} and L_{6t} , which occurs with a probability of at most $2(1 - y(\mathcal{U}_{j'}))$. Because at least one facility is opened in the bundle pair $(\mathcal{U}_{j'}, \mathcal{U}_{j''})$, at most $d(j', j'') + R_{j''} \leq 3R$ away from j' , we charge $3R$ more to X_j with probability at most $2(1 - y(\mathcal{U}_{j'}))$, hence an increase of $6d_{\text{far}}^{-h}(j')$ in expectation. Moreover, an additional of at most $3d_{\text{far}}^{-h}(j')$ needs to be charged when \bar{f} picks $\mathcal{U}_{j''}$ consistently in L_{6t-5} and L_{6t} , so the total increase in expectation is $9d_{\text{far}}^{-h}(j')$.

When j' and j'' are matched and they are closest neighbors of each other, the analysis and upper bounds are subsumed by the above. Up till now, we have charge incremental costs to D_j and X_j in several different cases, such that $D_j = (4 + 5\beta)h$ and the most amount we have charged for X_j in expectation is when j' is matched to $j'' = n_{j'}$ and $n_{j''} \neq j'$, where

$$D_j = (4 + 5\beta)h, \mathbb{E}[X_j] \leq 4d_{\text{av}}^{-h}(j) + d_{\text{far}}^{-h}(j') + 9d_{\text{far}}^{-h}(j') = 4d_{\text{av}}^{-h}(j) + 10d_{\text{far}}^{-h}(j'),$$

where the second term $d_{\text{far}}^{-h}(j')$ is because the marginal distribution on $\mathcal{U}_{j'} \setminus \text{Ball}^o(j', \beta h)$ cannot increase in this case. Now we only need to relate $d_{\text{far}}^{-h}(j')$ to $d_{\text{av}}^{-h}(j)$.

First, assume that $d(j, j') > \alpha h$ for some parameter $\alpha \in (4, \beta - 1]$ to be determined later (recall that $\beta > 5$). Then from $\alpha h < d(j, j') \leq 4d_{\text{av}}(j)$ and $d_{\text{far}}^{-h}(j') \leq d_{\text{av}}^{-h}(j') \leq d_{\text{av}}(j') \leq d_{\text{av}}(j) \leq d_{\text{av}}^{-h}(j) + h$, it is easy to see that

$$h < \frac{4}{\alpha - 4} \cdot d_{\text{av}}^{-h}(j) \Rightarrow d_{\text{far}}^{-h}(j') \leq \frac{\alpha}{\alpha - 4} \cdot d_{\text{av}}^{-h}(j),$$

so in this case, we can upper bound $\mathbb{E}[X_j]$ by

$$\mathbb{E}[X_j] \leq \left(4 + \frac{10\alpha}{\alpha - 4}\right) d_{\text{av}}^{-h}(j) = \frac{14\alpha - 16}{\alpha - 4} \cdot d_{\text{av}}^{-h}(j).$$

Second, assume that $d(j, j') \leq \alpha h$. We claim that in the fractional assignment x , j is served by facilities in $F_t \setminus \text{Ball}^o(j', \beta h)$ no less than j' does, otherwise we could shift a positive amount of assignment from $F_t \setminus \text{Ball}^o(j', \beta h)$ into $\text{Ball}^o(j', \beta h)$ for j' , resulting in a better assignment of j' , which is a contradiction. Therefore, we know for sure that

$$\sum_{i \in F_t \setminus \text{Ball}^o(j', \beta h)} x_{ij'} \leq \sum_{i \in F_t \setminus \text{Ball}^o(j', \beta h)} x_{ij}.$$

Then we try to modify the assignment of j' and get a sub-optimal solution, which helps us relate $d_{\text{far}}^{-h}(j')$ to $d_{\text{av}}^{-h}(j)$. In the altered assignment x' , only assignments of j' are changed, where $x'_{ij'} = x_{ij}$ for every facility $i \in F_t \setminus \text{Ball}^o(j', \beta h)$. It can be shown easily that this is possible, simply by allocating some fractional assignment of j' in $\text{Ball}^o(j', \beta h)$ to the outside, and we do not care about $x'_{ij'}, i \in \text{Ball}^o(j', \beta h)$ in this case.

Using triangle inequality, for any $i \in F_t \setminus \text{Ball}^o(j', \beta h)$, $d(i, j) \geq d(i, j') - d(j, j') \geq (\beta - \alpha)h \geq h$, so $d^{-h}(i, j) = d(i, j)$ and

$$\frac{d^{-h}(i, j')}{d^{-h}(i, j)} = \frac{d(i, j')}{d(i, j)} \leq \frac{d(i, j')}{d(i, j') - \alpha h} \leq \frac{\beta h}{\beta h - \alpha h} = \frac{\beta}{\beta - \alpha}.$$

On the other hand, x' may not be the optimal assignment for j' , therefore we know

$$d_{\text{far}}^{-h}(j') \leq \sum_{i \in F_t \setminus \text{Ball}^o(j', \beta h)} x'_{i,j'} d^{-h}(i, j') \leq \frac{\beta}{\beta - \alpha} \sum_{i \in F_t \setminus \text{Ball}^o(j', \beta h)} x_{ij} d^{-h}(i, j) \leq \frac{\beta}{\beta - \alpha} \cdot d_{\text{av}}^{-h}(j).$$

So to summarize, we can bound the two connection budgets by

$$D_j \leq (4 + 5\beta)h, \mathbb{E}[X_j] \leq \max \left\{ \frac{14\alpha - 16}{\alpha - 4}, \frac{14\beta - 4\alpha}{\beta - \alpha} \right\} d^{-h}_{\text{av}}(j), \text{ where } \beta > 5, 4 < \alpha \leq \beta - 1.$$

Plugging in $\alpha = 2 + 2\sqrt{3}, \beta = 4 + 2\sqrt{3}$, we have the desired result

$$D_j \leq (24 + 10\sqrt{3})h, \mathbb{E}[X_j] \leq (24 + 10\sqrt{3})d^{-h}_{\text{av}}(j).$$

□

A.5 Missing Proof of Lemma 9

Proof. There are $|F_t| \cdot |C_t|$ possible distinct distances at time t , hence the number of guesses is obvious, and each guess is associated with a family of reduced cost functions \mathfrak{D} , which fully dictates the running of our algorithm.

Assume our guessed thresholds are exactly those in the optimal solution in the following. Let $\{A_t\}_{t=1}^T$ be the stochastic output facility locations, and $\{O_t\}_{t=1}^T$ be the optimal solution, with a total cost $\text{OPT} = \text{OPT}_{\text{service}} + \text{OPT}_{\text{move}}$, and LP our corresponding LP value. We have $\text{OPT} \geq \text{LP}$. We aim to show an upper-bound on the overall expectation of $\text{SOL} = \text{SOL}_{\text{move}} + \text{SOL}_{\text{service}}$ of our algorithm's solution output, where the two parts represent the facility movement cost (γ -scaled) and client service cost.

For each $t \in [T], r \in [N_t]$, let $I_r^{(t)}$ be the largest index in w_t that has value $w_{tI_r^{(t)}} = \bar{w}_{tr}$, and denote $\mu = 24 + 10\sqrt{3}$. Using Lemma 8, for any $t \in [T], r \in [N_t]$ we have

$$\mathbb{E}[\text{rect}(I_r^{(t)}, |C_t|) \cdot d(C_t, A_t)^\downarrow] \leq \mu I_r^{(t)} T_r^{(t)} + \mu \sum_{j \in C_t} d^{-T_r^{(t)}(t)}(j),$$

and we decompose the true cost $w_t \cdot d(C_t, A_t)^\downarrow$ as rectangular pieces,

$$\begin{aligned} \mathbb{E}[w_t \cdot d(C_t, A_t)^\downarrow] &= \sum_{r=1}^{N_t} (\bar{w}_{tr} - \bar{w}_{t(r+1)}) \cdot \mathbb{E} \left[\text{rect}(I_r^{(t)}, |C_t|) \cdot d(C_t, A_t)^\downarrow \right] \\ &\leq \sum_{r=1}^{N_t} (\bar{w}_{tr} - \bar{w}_{t(r+1)}) \left(\mu I_r^{(t)} T_r^{(t)} + \mu \sum_{j \in C_t} d^{-T_r^{(t)}(t)}(j) \right) \\ &= \mu \cdot \sum_{r=1}^{N_t} (\bar{w}_{tr} - \bar{w}_{t(r+1)}) I_r^{(t)} T_r^{(t)} + \mu \cdot \sum_{r=1}^{N_t} \sum_{j \in C_t} (\bar{w}_{tr} - \bar{w}_{t(r+1)}) d^{-T_r^{(t)}(t)}(j). \end{aligned}$$

We aim to bound each term above using OPT. Since we've made an assumption that our guessed thresholds are correct, in the inner product $w_t \cdot d(A_t, C_t)^\downarrow$, every weight that is equal to \bar{w}_{tr} should be multiplied by a distance at least $T_r^{(t)}$, thus the optimal solution OPT has service cost OPT_t for time step t at least

$$\begin{aligned} \text{OPT}_t &\geq \sum_{r=1}^{N_t} \bar{w}_{tr} (I_r^{(t)} - I_{r-1}^{(t)}) T_r^{(t)} \geq \sum_{r=1}^{N_t} \left(\bar{w}_{tr} I_r^{(t)} T_r^{(t)} - \bar{w}_{tr} I_{r-1}^{(t)} T_{r-1}^{(t)} \right) \\ &\geq \sum_{r=1}^{N_t} \left(\bar{w}_{tr} I_r^{(t)} T_r^{(t)} - \bar{w}_{t(r+1)} I_r^{(t)} T_r^{(t)} \right) = \sum_{r=1}^{N_t} (\bar{w}_{tr} - \bar{w}_{t(r+1)}) I_r^{(t)} T_r^{(t)}, \end{aligned}$$

where OPT_t denotes the service cost of all clients in C_t , and $\sum_{t \in [T]} \text{OPT}_t = \text{OPT}_{\text{service}}$.

Further, for the second term above, we have

$$\begin{aligned}
\sum_{r=1}^{N_t} \sum_{j \in C_t} (\bar{w}_{tr} - \bar{w}_{t(r+1)}) d^{-T_r^{(t)}}(j) &= \sum_{r=1}^{N_t} \sum_{j \in C_t} \sum_{i \in F_t} (\bar{w}_{tr} - \bar{w}_{t(r+1)}) x_{ij}^{(t)} \cdot d^{-T_r^{(t)}}(i, j) \\
&= \sum_{j \in C_t} \sum_{i \in F_t} x_{ij}^{(t)} \cdot \sum_{r=1}^{N_t} (\bar{w}_{tr} - \bar{w}_{t(r+1)}) d^{-T_r^{(t)}}(i, j) \\
&= \sum_{j \in C_t} \sum_{i \in F_t} x_{ij}^{(t)} \cdot \sum_{r: T_r^{(t)} \leq d(i, j)} (\bar{w}_{tr} - \bar{w}_{t(r+1)}) d(i, j) \\
&= \sum_{j \in C_t} \sum_{i \in F_t} x_{ij}^{(t)} \cdot w^t(i, j) \cdot d(i, j) \\
&= \sum_{j \in C_t} \sum_{i \in F_t} x_{ij}^{(t)} \cdot d^t(i, j).
\end{aligned}$$

To summarize, our algorithm outputs a stochastic solution $\{A_t\}_{t \in [T]}$ such that the expected service cost of all clients is at most

$$\mathbb{E}[\text{SOL}_{\text{service}}] = \mathbb{E} \left[\sum_{t=1}^T w_t \cdot d(C_t, A_t)^\downarrow \right] \leq \mu \cdot \sum_{t \in [T]} \text{OPT}_t + \mu \cdot \sum_{t=1}^T \sum_{j \in C_t} \sum_{i \in F_t} x_{ij}^{(t)} \cdot d^t(i, j).$$

Meanwhile by Lemma 7, the movement cost of all facilities in $\{A_t\}_{t \in [T]}$ has an expectation of at most

$$\mathbb{E}[\text{SOL}_{\text{move}}] \leq \gamma \sum_{t=1}^{T-1} \sum_{i \in F_t} \sum_{i' \in F_{t+1}} d(i, i') z_{ii'}^{(t)} + 6\gamma \cdot \mathbf{1}_{T \geq 3} \cdot \sum_{t \in [T]} \sum_{j \in C_t} d_{\text{av}}^{(t)}(j).$$

Back to the lemma, if $T = 2$, our overall solution has total expected value at most

$$\begin{aligned}
\mathbb{E}[\text{SOL}] &\leq \mu \cdot \text{OPT}_{\text{service}} + \mu \cdot \sum_{t=1}^T \sum_{j \in C_t} \sum_{i \in F_t} x_{ij}^{(t)} \cdot d^t(i, j) + \gamma \sum_{t=1}^{T-1} \sum_{i \in F_t} \sum_{i' \in F_{t+1}} d(i, i') z_{ii'}^{(t)} \\
&\leq \mu \cdot \text{OPT} + \mu \cdot \sum_{t=1}^T \sum_{j \in C_t} \sum_{i \in F_t} x_{ij}^{(t)} \cdot d^t(i, j) + \gamma \mu \cdot \sum_{t=1}^{T-1} \sum_{i \in F_t} \sum_{i' \in F_{t+1}} d(i, i') z_{ii'}^{(t)} \\
&\leq \mu \cdot \text{OPT} + \mu \cdot \text{LP} \leq (48 + 20\sqrt{3}) \cdot \text{OPT}.
\end{aligned}$$

For the case when $T \geq 3$, recall that we assume the smallest entry in $\{w_t\}_{t \in [T]}$ is at least $\epsilon > 0$, using Lemma 7, we have the following regarding our overall expected cost

$$\begin{aligned}
\mathbb{E}[\text{SOL}] &\leq \mu \cdot \text{OPT}_{\text{service}} + \mu \cdot \sum_{t=1}^T \sum_{j \in C_t} \sum_{i \in F_t} x_{ij}^{(t)} \cdot d^t(i, j) + \gamma \sum_{t=1}^{T-1} \sum_{i \in F_t} \sum_{i' \in F_{t+1}} d(i, i') z_{ii'}^{(t)} \\
&\quad + 6\gamma \sum_{t \in [T]} \sum_{j \in C_t} \sum_{i \in F_t} x_{ij}^{(t)} \cdot d(i, j) \\
&\leq \mu \cdot \text{OPT} + \mu \cdot \sum_{t=1}^T \sum_{j \in C_t} \sum_{i \in F_t} x_{ij}^{(t)} \cdot d^t(i, j) + \gamma \cdot \sum_{t=1}^{T-1} \sum_{i \in F_t} \sum_{i' \in F_{t+1}} d(i, i') z_{ii'}^{(t)} \\
&\quad + \frac{6\gamma}{\epsilon} \cdot \sum_{t \in [T]} \sum_{j \in C_t} \sum_{i \in F_t} x_{ij}^{(t)} \cdot d^t(i, j) \\
&\leq \mu \cdot \text{OPT} + \left(\mu + \frac{6\gamma}{\epsilon} \right) \cdot \text{LP} = \left(2\mu + \frac{6\gamma}{\epsilon} \right) \cdot \text{OPT} = \left(48 + 20\sqrt{3} + \frac{6\gamma}{\epsilon} \right) \cdot \text{OPT},
\end{aligned}$$

where the second to last inequality is because the smallest entry in $\{w_t\}_{t \in [T]}$ is at least ϵ , thus our reduced metric d^t always satisfies $d^t \geq \epsilon d$ for any t , since $d^t(i, j) = w^t(i, j)d(i, j) \geq \epsilon d(i, j)$ by definition. \square

B An Application to Facility-Weighted TM-MFL

The minimum total movement mobile facility location (TM-MFL) problem is studied in [1, 16, 18], where the facility-weighted version has the best known approximation ratio of $3 + O(\sqrt{\log \log p / \log p})$ using p -Swap based local search algorithms in [1], specifically with an approximation ratio at most 499 when $p = 1$. Swamy [33] shows that there exists an approximation-preserving reduction from TM-MFL to the matroid median problem, which is introduced by Krishnaswamy *et al.* [29], and gives an 8-approximation thereof. The current best approximation ratio for matroid median is later improved to 7.081 by Krishnaswamy *et al.* [30]. In light of our algorithm for Dynamic Ordered k -Median, we present an LP-rounding algorithm with a much simpler analysis.

In this facility weighted TM-MFL problem, we are given a finite ground set X equipped with a metric d , a client set $C \subseteq X$ with each client $j \in C$ having a demand D_j , a facility set F with each facility $i \in F$ having a weight w_i . We are asked to assign a destination $v_i \in X$ for every $i \in F$, move i to v_i , and minimize the following objective,

$$\sum_{i \in F} w_i \cdot d(i, v_i) + \sum_{j \in C} D_j \cdot d(j, u_j),$$

where u_j denotes the closest facility to j after moving the facilities.

Consider the instance of facility-weighted TM-MFL as an instance of Dynamic Ordered k -Median with only 2 time steps, where in the first time step, the facilities in F are fixed but serve no clients, i.e. $C_1 = \emptyset$, and in the second time step, the facilities are moved in order to serve clients in $C_2 = C$. We use a new LP relaxation similar to LP(\mathfrak{D}). Denote x_{ij} the assignment from client j to location i , after moving the facilities, y_i the probability of some facility being moved to i , and $z_{ii'}$ the fractional movement from i to i' ,

$$\text{minimize : } \sum_{i \in F} \sum_{i' \in X} w_i \cdot z_{ii'} \cdot d(i, i') + \sum_{j \in C} \sum_{i \in X} D_j \cdot x_{ij} \cdot d(i, j) \quad \text{LP}_{\text{tm-mfl}}$$

$$\text{subject to : } \sum_{i \in X} x_{ij} = 1, \quad \forall j \in C \quad (7)$$

$$\sum_{i' \in X} z_{ii'} = 1, \quad \forall i \in F \quad (8)$$

$$\sum_{i \in F} z_{ii'} = y_{i'}, \quad \forall i' \in X \quad (9)$$

$$0 \leq x_{ij} \leq y_i \leq 1, \quad \forall i \in X, j \in C \quad (10)$$

Lemma 15. *The above LP_{tm-mfl} is feasible with optimum at most OPT, the optimal cost of the original TM-MFL instance.*

Proof. Let (x^*, y^*, z^*) be the integral solution constructed from the optimal solution. More precisely, $x_{ij}^* = 1$ iff client j is assigned to location i , $y_i = 1$ iff there exists a facility at location i after the movements, and $z_{ii'} = 1$ iff the facility originally at i is moved to i' . Obviously (x^*, y^*, z^*) satisfies the constraints, and the objective has value exactly OPT in this case. \square

We solve LP_{tm-mfl} for a fractional optimum solution (x, y, z) . Notice that all movements and demands are individually weighted. However, we ignore the demands during a filtering process similar to Algorithm 1 and still only use $d_{\text{av}}(j) = \sum_{i \in X} x_{ij}d(i, j)$. For each client $j \in C$, let $j' \in C'$ be the filtered client such that $d_{\text{av}}(j') \leq d_{\text{av}}(j)$, $d(j, j') \leq 4 \max\{d_{\text{av}}(j), d_{\text{av}}(j')\}$. If we can round (x, y, z) to a random integral solution, such that the expected distance from j' to the nearest facility is at most $\theta d_{\text{av}}(j')$, we have

$$\mathbb{E}[D_j \cdot d(j, u_j)] \leq D_j \cdot \mathbb{E}[d(j, j') + d(j', u_{j'})] \leq 4D_j d_{\text{av}}(j) + \theta D_j d_{\text{av}}(j') \leq (4 + \theta)D_j d_{\text{av}}(j).$$

To secure the expected service cost for $j' \in C'$, we construct a similar network \mathcal{N} as before. In fact, \mathcal{N} can be substantially simplified, as there are only 2 time steps, and the facilities for $t = 1$ are fixed. Hence we only keep the source, the sink, L_7 through L_{12} and L_0 which contains a node for each $i \in F$, serving in place for $t = 1$. There is, again, a natural fractional flow \bar{f} with value $|F|$ constructed from (x, y, z) , and we can round it to \tilde{f} so that it has value $|F|$ and $\mathbb{E}[\tilde{f}] = \bar{f}$. Recall that $\bar{f}(L_7) \subseteq X$ the (multi)-set of facility nodes in layer L_7 that \bar{f} passes through.

Indeed, we can further simplify the multi-set $\bar{f}(L_7)$ into a proper subset S_7 of X , so that the support of S_7 contains that of $\bar{f}(L_7)$, the overall service cost stays the same, while the movement cost decreases. This is because, whenever there exists a facility location i that appears multiple times in $\bar{f}(S_7)$, it means that at least some facility that was in location $i' \neq i$ moved into i . If we remove this movement, i still holds at least one facility after the arrangements, and the total movement cost is reduced. We repeat this process until there are no repeated locations and obtain S_7 . Clearly $|S_7| = |F|$.

Lemma 16. For $j' \in C'$, $\mathbb{E}[d(j', S_7)] \leq 6d_{\text{av}}(j')$.

Proof. The proof falls in line with the standard analysis of Charikar and Li [12]. For each filtered client j' , there is a probability $y(\mathcal{U}_{j'})$ that one facility is present in $\mathcal{U}_{j'}$, contributing at most $d_{\text{av}}(j')$ to $\mathbb{E}[d(j', S_7)]$.

On the other hand, there is a $1 - y(\mathcal{U}_{j'})$ probability that there exists no facility in $\mathcal{U}_{j'}$. Consider the worst case where j'' , the nearest neighbor of j' is matched with $j''' \neq j'$, hence $d(j', j'') \geq d(j'', j''')$. There is guaranteed to be a facility in the pair (j'', j''') , and this facility is at most $d(j', j'') + d(j'', j''') + 0.5d(j'', j''')$ away from j' . We also have $d_{\text{av}}(j') \geq 0.5d(j', j'')(1 - y(\mathcal{U}_{j'}))$, thus the expected service cost of j' is contributed to by at most $(d(j', j'') + d(j'', j''') + 0.5d(j'', j'''))(1 - y(\mathcal{U}_{j'})) \leq 5d_{\text{av}}(j')$.

In total, the expectation of $d(j', S_7)$ is at most $6d_{\text{av}}(j')$. \square

Meanwhile, because there are only 2 time steps, if we construct the network flow accordingly, the expected movement cost is unchanged, as shown in Lemma 7. Combined with the approximation of service cost above, we obtain a solution with expected total cost

$$\mathbb{E}[\text{SOL}] \leq 10\text{LP}_{\text{service}} + \text{LP}_{\text{move}} \leq 10\text{OPT},$$

where $\text{LP} = \text{LP}_{\text{service}} + \text{LP}_{\text{move}}$ is the optimal objective value of the LP relaxation, $\text{LP}_{\text{service}}$ is the service cost part and LP_{move} is the movement cost part. The last inequality is due to Lemma 15.

Theorem 17. *There exists a 10-approximation algorithm for the facility-weighted minimum total movement mobile facility location problem.*

C The Approximation Algorithm for Dynamic k -Supplier with Outliers and $T = 2$

C.1 Solving Dynamic k -Supplier with outliers, $T = 2$: Guessing and Modifying LP

We consider the outlier version of Dynamic k -Supplier when $T = 2$ from now on. A useful ingredient as shown in [20], is the multi-budget approximation algorithm for the general matching problem, when the number of budgets is constant and *the matching is a maximum version with budget upper bounds*. In the algorithm in [20], first the edges with the highest $O(1/\epsilon)$ weights in the optimal solution are guessed, then every other edge with potentially higher weights are removed from further consideration, which helps improves the worst-case guarantee in the remaining problem.

In our adaptation of the algorithm, we slightly modify the method of guessing the top $O(1/\epsilon)$ facilities we choose. Denote $B_1(i, R) = \{j \in C_1 : d(i, j) \leq R\}$, $B_2(i, R) = \{j \in C_2 : d(i, j) \leq R\}$. W.l.o.g. we assume the optimal R^* is guessed correctly, we do the following to modify the original instance,

1. Enumerate all possible choices of γ/ϵ distinct facilities in F_1 and γ/ϵ distinct facilities in F_2 , where $\gamma > 0$ is a constant. Denote the two chosen sets of facilities T_1, T_2 . Additionally, enumerate all possible movements associated with T_1 and T_2 , i.e. $g : T_1 \rightarrow F_2$ and $h : T_2 \rightarrow F_1$ such that $\forall i \in T_1, d(i, g(i)) \leq B, \forall i' \in T_2, d(i', h(i')) \leq B$;
2. Recursively sort T_1 , each time by choosing the unchosen $i \in T_1$ such that $B_1(i, 3R^*)$ covers the largest number of uncovered clients in C_1 , and letting this number be $u_i^{(1)}$. Denote $u_0^{(1)}$ the number of remaining clients that are covered by $\bigcup_{i \in h(T_2)} B_1(i, 3R^*)$. Same for T_2 and C_2 , and we have $u_0^{(2)}$ and $u_i^{(2)}$ for $i \in T_2$. Denote $C'_1 = C_1 - \bigcup_{i \in T_1} B_1(i, 3R^*) - \bigcup_{i \in h(T_2)} B_1(i, 3R^*)$ and $C'_2 = C_2 - \bigcup_{i \in T_2} B_2(i, 3R^*) - \bigcup_{i \in g(T_1)} B_2(i, 3R^*)$;
3. Remove from F_1 any $i' \in F_1 - T_1$ such that $B_1(i', 3R^*) \cap C'_1 > \min_{i \in T_1} \{u_i^{(1)}\}$. Remove from F_2 any $i' \in F_2 - T_2$ such that $B_2(i', 3R^*) \cap C'_2 > \min_{i \in T_2} \{u_i^{(2)}\}$. Denote the remaining two facility sets F'_1 and F'_2 . Note that we still have $T_1 \subseteq F'_1, T_2 \subseteq F'_2$;
4. Denote the reduced problem \mathcal{P}' , where the client sets are C'_1, C'_2 , the facility sets are F'_1, F'_2 and the outlier constraints become $l'_1 = \max \left\{ l_1 - u_0^{(1)} - \sum_{i \in T_1} u_i^{(1)}, 0 \right\}$ and $l'_2 = \max \left\{ l_2 - u_0^{(2)} - \sum_{i \in T_2} u_i^{(2)}, 0 \right\}$, respectively. Since we guess T_1 and T_2 to be chosen with certainty, the corresponding variable y_i s should be explicitly or implicitly set to ≥ 1 in the relaxed LP, through a set of additional constraints, similarly for $g(T_1)$ and $h(T_2)$, as well as the movement variables representing $T_1 \rightarrow g(T_1), h(T_2) \rightarrow T_2$.

Lemma 18. *Assume we have guessed the optimal R^* , then there exists a guess (T_1, T_2, g, h) such that T_1, T_2 both correspond to the top γ/ϵ facilities in the optimal solution, $g(T_1) \subseteq F'_2, h(T_2) \subseteq F'_1$, and the rest of the facilities in the optimal solution all belong to F'_1, F'_2 , respectively. Moreover, there are at most $O\left((|F_1| \cdot |F_2|)^{2\gamma/\epsilon}\right)$ such different guesses.*

Proof. The number of possible guesses is easy to see from the definition. Consider the optimal solution $(U_1 \subseteq F_1, U_2 \subseteq F_2)$. Choose γ/ϵ facilities in $U'_1 \subseteq U_1$, such that the total number of clients they cover in radius $3R^*$ is maximized. Same goes for $U'_2 \subseteq U_2$, and suppose we have made the correct guess $T_1 = U'_1, T_2 = U'_2$.

Sort U'_1 in the same manners as above, it is easy to see that any $i' \in U_1 - U'_1$ can only cover at most $\min_{i \in U'_1} \{u_i^{(1)}\}$ new clients, otherwise it is chosen as a member of U'_1 . Therefore, $U_1 - U'_1 \subseteq F'_1$. The same argument is valid for $U_2 - U'_2 \subseteq F'_2$. \square

From now on, we further assume we have made the correct guess about (T_1, T_2, g, h) , and successfully reduced the problem to \mathcal{P}' , where the outlier constraints, client sets and facilities sets are modified corre-

spondingly. We restate the modified $\text{LP}(\mathcal{P}')$ here.

$$\sum_{j \in C'_t} \sum_{i \in F'_t, d(i,j) \leq R^*} x_{ij}^{(t)} \geq l'_t, \quad \forall t = 1, 2 \quad (11)$$

$$\sum_{i \in F'_t, d(i,j) \leq R^*} x_{ij}^{(t)} \leq 1, \quad \forall j \in C'_t, t = 1, 2 \quad (12)$$

$$\sum_{i \in F'_t} y_i^{(t)} = k, \quad \forall t = 1, 2 \quad (13)$$

$$0 \leq x_{ij}^{(t)} \leq y_i^{(t)}, \quad \forall i \in F'_t, j \in C_t, t = 1, 2 \quad (14)$$

$$\sum_{i \in \text{Ball}(i', B) \cap F_1} z_{ii'}^{(1)} = y_{i'}^{(2)}, \quad \forall i' \in F'_2 \quad (15)$$

$$\sum_{i' \in \text{Ball}(i, B) \cap F_2} z_{ii'}^{(1)} = y_i^{(1)}, \quad \forall i \in F'_1 \quad (16)$$

$$z_{ii'}^{(1)} \geq 0, \quad \forall i \in F'_1, i' \in F'_2 \quad (17)$$

$$z_{ig(i)} = 1, \quad \forall i \in T_1 \quad (18)$$

$$z_{h(i')i'} = 1, \quad \forall i' \in T_2 \quad (19)$$

Lemma 19. *If (T_1, T_2, g, h) and R^* are guessed correctly, then $\text{LP}(\mathcal{P}')$ is feasible.*

Proof. Consider the optimal solution $U_1 \subseteq F_1, U_2 \subseteq F_2$, and define the variables $x_{ij}^{*(t)}, y_i^{*(t)}$ and $z_{ii'}^{*(1)}$ accordingly. Using Lemma 18, if we compute F'_1, F'_2 according to T_1 and T_2 , we have $U_1 \subseteq F'_1, U_2 \subseteq F'_2$, and it is easy to check that all but the first constraint of $\text{LP}(\mathcal{P}')$ are satisfied by (x^*, y^*, z^*) .

Now consider the first constraint. In the optimal solution, U_1 covers at least l_1 clients in C_1 with radius R^* and we have $T_1 \subseteq U_1, h(T_2) \subseteq U_1$. Among these l_1 clients, the facilities in $T_1 \cup h(T_2)$ cover exactly $u_0^{(1)} + \sum_{i \in T_1} u_i^{(1)}$ clients of them with radius $3R^* > R^*$, hence the remaining facilities in $U_1 - T_1 - h(T_2)$ cover at least $\max\{l_1 - u_0^{(1)} - \sum_{i \in T_1} u_i^{(1)}, 0\} = l'_1$ clients with radius R^* , and finally we notice $U_1 - T_1 - h(T_2) \subseteq F'_1$. Same goes for C'_2 and l'_2 , thus the first constraint of the LP is also satisfied by (x^*, y^*, z^*) , and $\text{LP}(\mathcal{P}')$ is feasible. \square

Corollary 20. *In the remaining problem \mathcal{P}' , any facility $i \in F'_1 - T_1$ can cover at most $(|C_1 - C'_1| - u_0^{(1)}) \epsilon / \gamma$ clients within distance of $3R^*$; any facility $i \in F'_2 - T_2$ can cover at most $(|C_2 - C'_2| - u_0^{(2)}) \epsilon / \gamma$ clients within distance of $3R^*$.*

Proof. According to our construction of F'_1 , any facility $i \in F'_1 - T_1$ must cover at most $\min_{i \in T_1} u_i^{(1)}$ remaining clients in C'_1 , with radius $3R^*$, otherwise it would have been removed during constructing \mathcal{P}' . Of course, this number is also smaller than the average of all $u_i^{(1)}$ s, which is exactly

$$\frac{\epsilon}{\gamma} \sum_{i \in T_1} u_i^{(1)} = \frac{\epsilon}{\gamma} \left(|C_1 - C'_1| - u_0^{(1)} \right),$$

where we recall $C_1 - C'_1$ contains all clients covered by $T_1 \cup h(T_2)$ and $|T_1| = \gamma / \epsilon$. The case with $F'_2 - T_2$ is identical. \square

C.2 Solving Dynamic k -Supplier with outliers, $T = 2$: Rounding the LP Solution

Given a fractional solution (x, y, z) to a Dynamic k -Supplier instance, specifically for an solution to $\text{LP}(\mathcal{P}')$, define $E_j^{(t)} = \{i \in \text{Ball}(j, R^*) \cap F_t : x_{ij}^{(t)} > 0\}$ the relevant facilities with respect to j and $s_j^{(t)} = \sum_{i \in E_j^{(t)}} x_{ij}^{(t)}$ the fraction demand assigned to j . We conduct the filtering algorithm in [22] to filter the clients, defined

as in Algorithm 3. For any $t \in [T]$ and $j \in C_t$, there exist some $j' \in C'_t$ such that $E_j^{(t)} \cap E_{j'}^{(t)} \neq \emptyset$, hence $d(j, j') \leq 2R^*$. Therefore, if we can cover all the clients in C'_t with a radius of at most κR^* , using triangle inequality, every client in C_t can be covered within distance $(\kappa + 2)R^*$.

Algorithm 3 GREEDYFILTER(x, y)

```

1: for all  $t \in [T]$  do
2:    $C''_t \leftarrow \emptyset$ 
3:   for all unmarked cluster  $E_j^{(t)}$  in decreasing order of  $s_j^{(t)}$  do
4:      $C''_t \leftarrow C''_t \cup \{j\}$ 
5:     Set each unmarked  $E_{j'}^{(t)}$  that has non-empty intersection with  $E_j^{(t)}$  as marked
6:     Let  $c_j^{(t)}$  be the number of clusters marked in this iteration
7:   end for
8:    $c^{(t)} \leftarrow (c_j^{(t)} : j \in C''_t)$ 
9: end for
10: return  $((C''_t, c^{(t)}) : t \in [T])$ 

```

Lemma 21. For any $j \in C''_1$, $c_j^{(1)} \leq (|C_1 - C'_1| - u_0^{(1)}) \epsilon / \gamma$.

Proof. For any $j \in C''_1$, $c_j^{(1)}$ is at most the number of clients in C'_1 that are $\leq 2R^*$ away from j . Since we know there exists a facility i' in $E_j^{(1)} \subseteq F'_1$, these $c_j^{(1)}$ clients are at most $3R^*$ away from i' . Using Corollary 20, $c_j^{(1)}$ is at most the number of clients in C'_1 covered by i' , hence at most $(|C_1 - C'_1| - u_0^{(1)}) \epsilon / \gamma$. \square

Lemma 22. $\sum_{j \in C''_1} c_j^{(1)} s_j^{(1)} \geq l'_1$.

Proof. In Algorithm 3, the clusters are processed in descending order of $s_j^{(1)}$, and $c_j^{(1)}$ denotes the number of clusters that are removed in each step, hence we have

$$\sum_{j \in C''_1} c_j^{(1)} s_j^{(1)} \geq \sum_{j \in C'_1} s_j^{(1)} \geq l'_1,$$

where we recall that $s_j^{(1)} = \sum_{i \in F'_1, d(i, j) \leq R^*} x_{ij}^{(1)}$, and the last inequality is due to $x_{ij}^{(1)}$ being the LP solution. \square

We first present the following Algorithm 4 to modify the LP solution. Notice in line 22, we merge all the facilities in some filtered cluster $E_j^{(1)}$, and the remaining facility $v_j^{(1)}$ has y value guaranteed to be $s_j^{(1)}$. This is because after the previous splitting process, whenever a facility location i is in any cluster $E_j^{(1)}$ (meaning $x_{ij}^{(1)} > 0$), we have $x_{ij}^{(1)} = y_i^{(1)}$, and since $s_j^{(1)} = \sum_i x_{ij}^{(1)}$, our claim follows.

Lemma 23. In the modified LP solution output by Algorithm 4, if we create a bipartite graph $G = (V_1 \cup V_2, V_1 \times V_2)$, where V_1, V_2 represent the facilities in F''_1, F''_2 , respectively, and edge (i, i') has fractional matching value $z_{ii'}$, we have a fractional k -cardinality bipartite matching M_f over G .

Moreover, if we assign budget $l_1(e) = c_j^{(1)}$ for every edge e resting on merged facility $v_j^{(1)}$, $j \in C''_1$, and budget $l_2(e') = c_{j'}^{(2)}$ for every edge e' on merged facility $v_{j'}^{(2)}$, $j' \in C''_2$, then M_f satisfies the budget (coverage) constraints of \mathcal{P}^j , where

$$L_1(M_f) = \sum_e z_e l_1(e) \geq l'_1, \quad L_2(M_f) = \sum_e z_e l_2(e) \geq l'_2.$$

Algorithm 4 SPLIT $(x, y, z, C'_1, C'_2, c^{(1)}, c^{(2)}, T_1, T_2, g, h)$

- 1: $F''_1 \leftarrow F'_1, F''_2 \leftarrow F'_2, \hat{x} \leftarrow x, \hat{y} \leftarrow y, \hat{z} \leftarrow z$
 - 2: **for** each distinct edge $(i, i') \in \{(i, g(i)) : i \in T_1\} \cup \{(h(i'), i') : i' \in T_2\}$ **do**
 - 3: Split i and i' for two new nodes i_1, i'_1 . Set $F''_1 \leftarrow F'_1 \cup \{i_1\}, F''_2 \leftarrow F'_2 \cup \{i'_1\}$
 - 4: Set $\hat{y}_{i_1}^{(1)} = \hat{y}_{i'_1}^{(2)} = 1, \hat{z}_{i_1 i'_1}^{(1)} = 1$, reduce corresponding \hat{y} and \hat{z} values of i, i' and (i, i') accordingly
 - 5: **end for**
 - 6: **while** there exists $i \in E_j^{(1)}$ such that $\hat{y}_i^{(1)} > \hat{x}_{ij}^{(1)}$ **do**
 - 7: Fix such an i . Choose j such that $i \in E_j^{(1)}$ (so that $\hat{x}_{ij}^{(1)} > 0$) and $\hat{x}_{ij}^{(1)}$ is minimized
 - 8: Split i into i_1, i_2 with $\hat{y}_{i_1}^{(1)} = \hat{x}_{ij}^{(1)}, \hat{y}_{i_2}^{(1)} = \hat{y}_i^{(1)} - \hat{y}_{i_1}^{(1)}$. Split corresponding \hat{z} values accordingly
 - 9: **for all** $j' \in C'_1$ such that $i \in E_{j'}^{(1)}$ **do**
 - 10: **if** $\hat{x}_{ij'}^{(1)} = \hat{x}_{ij}^{(1)}$ **then**
 - 11: Let $\hat{x}_{i_1 j'}^{(1)} = \hat{y}_{i_1}^{(1)}, \hat{x}_{i_2 j'}^{(1)} = 0$. Delete i from $E_j^{(1)}$, add i_1 to $E_j^{(1)}$
 - 12: **else**
 - 13: Let $\hat{x}_{i_1 j'}^{(1)} = \hat{y}_{i_1}^{(1)}, \hat{x}_{i_2 j'}^{(1)} = \hat{x}_{ij'}^{(1)} - \hat{y}_{i_1}^{(1)}$. Delete i from $E_j^{(1)}$, add i_1, i_2 to $E_j^{(1)}$
 - 14: **end if**
 - 15: **end for**
 - 16: **end while**
 - 17: **while** there exists i such that $\hat{y}_i^{(1)} > 1$ **do**
 - 18: Split i into $\lceil \hat{y}_i^{(1)} \rceil$ copies, such that the first $\lfloor \hat{y}_i^{(1)} \rfloor$ of them all hold \hat{y} value exactly 1, and the last one (if any) holds \hat{y} value equal to $\hat{y}_i^{(1)} - \lfloor \hat{y}_i^{(1)} \rfloor$. Split corresponding \hat{z} values accordingly
 - 19: **end while**
 - 20: Repeat the two loops above for $t = 2$
 - 21: **for all** $j \in C''_1$ **do**
 - 22: Merge all locations in $E_j^{(1)}$ into a single one $v_j^{(1)}$, with $\hat{y}_{v_j^{(1)}}^{(1)} = s_j^{(1)}$. Merge corresponding \hat{z} values accordingly
 - 23: **end for**
 - 24: Repeat the loop above for $t = 2$
 - 25: **return** $(\hat{x}, \hat{y}, \hat{z}, F''_1, F''_2)$
-

Proof. In Algorithm 4, we have two phases, which first splits facility locations into copies so that each has y value at most 1, then merges every facility in some filtered cluster $E_j^{(t)}$ into a super facility, and every merged facility $v_j^{(t)}$ has $y_{v_j^{(t)}}^{(t)} = \sum_{i \in E_j^{(t)}} y_i^{(t)} = \sum_{i \in E_j^{(t)}} x_{ij}^{(t)} = s_j^{(t)} \leq 1$, hence the modified LP solution z is a fractional k -cardinality matching on G .

For the budget constraints, the budget values for both time steps can be written as,

$$L_1(M_f) = \sum_{e=(i,i')} z_e l_1(e), \quad L_2(M_f) = \sum_{e=(i,i')} z_e l_2(e).$$

To see that the two constraints of \mathcal{P}' are satisfied, recall that all the edges e residing on the same merged facility $v_j^{(1)} \in V_1$ has $l_1(e) = c_j^{(1)}$, similarly for $t = 2$, and all the other budget values are 0, so the budget values above can be rephrased as,

$$L_1(M_f) = \sum_{j \in C''_1} \sum_{e=(v_j^{(1)}, i'')} z_e c_j^{(1)} = \sum_{j \in C''_1} c_j^{(1)} \left(\sum_{e=(v_j^{(1)}, i'')} z_e \right) = \sum_{j \in C''_1} c_j^{(1)} y_{v_j^{(1)}}^{(1)} = \sum_{j \in C''_1} c_j^{(1)} s_j^{(1)} \geq l'_1,$$

where the second to last equality is due to our Algorithm 4, and the last inequality is due to Lemma 22. \square

With the feasibility lemma above, we present the following theorem.

Theorem 24. *For any $\epsilon > 0$, there exists a bi-criteria approximation scheme for Dynamic k -Supplier with outliers when $T = 2$, that computes a solution which covers at least $(1 - \epsilon)l_t$ clients within radius $3R^*$, $t = 1, 2$, where R^* is the radius of the optimal solution.*

Proof. The proof follows the procedures in Section 4 of [20], with a slight modification. After remove the $\kappa = \Theta(\gamma/\epsilon)$ edges reserved by Algorithm 4, consider z_0 as a basic solution to the following LP: $\max\{\min \mathbf{1}^T z \mid z \in P_{\mathcal{M}}, L_1(z) \geq l'_1, L_2(z) \geq l'_2\}$. Since the aforementioned fractional solution z satisfies the constraints and has $\mathbf{1}^T z = k - \kappa$, we must also have $\mathbf{1}^T z_0 \leq k - \kappa$.

Now that z_0 is a basic solution in the matching polytope with two additional linear constraints, it lies on a face of $P_{\mathcal{M}}$ of dimension at most 2, thus using Carathéodory theorem, it can be written as the convex combination of 3 (integral) matchings, say $z_0 = \alpha_1 z_1 + \alpha_2 z_2 + \alpha_3 z_3$, where $\alpha_i \in [0, 1], \alpha_1 + \alpha_2 + \alpha_3 = 1$.

Construction of matching z'_2 . First, we create an almost-matching m_2 that fractionally combines z_1, z_2 , as well as their values on $L_1(\cdot), L_2(\cdot)$ and $\langle \mathbf{1}, \cdot \rangle$. To be more precise, let $z_{12} = \frac{\alpha_1}{\alpha_1 + \alpha_2} z_1 + \frac{\alpha_2}{\alpha_1 + \alpha_2} z_2$, we want to find $m_2 \in [0, 1]^E$ such that

$$L_1(m_2) = L_1(z_{12}), L_2(m_2) = L_2(z_{12}), \mathbf{1}^T m_2 = \mathbf{1}^T z_{12},$$

and it is possible to set at most 4 edges of m_2 to 0 in order to obtain a matching.

Using Corollary 4.10 of [20], such an almost-matching m_2 exists and can be efficiently computed. We can further efficiently set at most 4 edges of m_2 to 0 and obtain a matching z'_2 . It is obvious that

$$L_1(z'_2) \geq L_1(m_2) - 4 \max_e l_1(e), L_2(z'_2) \geq L_2(m_2) - 4 \max_e l_2(e), \mathbf{1}^T z'_2 \leq \mathbf{1}^T m_2.$$

Construction of matching z'_3 . Let $z_{23} = (\alpha_1 + \alpha_2)z'_2 + \alpha_3 z_3$, and $m_3 \in [0, 1]^E$ be such that

$$L_1(m_3) = L_1(z_{23}), L_2(m_3) = L_2(z_{23}), \mathbf{1}^T m_3 = \mathbf{1}^T z_{23},$$

and it is possible to set at most 4 edges of m_3 to 0 in order to obtain a matching.

Using Corollary 4.10 of [20], such an almost-matching m_3 exists and can be efficiently computed. We can further efficiently set at most 4 edges of m_3 to 0 and obtain a matching z'_3 . It is obvious that

$$L_1(z'_3) \geq L_1(m_3) - 4 \max_e l_1(e), L_2(z'_3) \geq L_2(m_3) - 4 \max_e l_2(e), \mathbf{1}^T z'_3 \leq \mathbf{1}^T m_3.$$

Further from the construction of z'_2 , we know that

$$\begin{aligned} L_1(z'_3) &\geq L_1(m_3) - 4 \max_e l_1(e) = L_1(z_{23}) - 4 \max_e l_1(e) \\ &= (\alpha_1 + \alpha_2)L_1(z'_2) + \alpha_3 L_1(z_3) - 4 \max_e l_1(e) \\ &\geq (\alpha_1 + \alpha_2) \left(L_1(m_2) - 4 \max_e l_1(e) \right) + \alpha_3 L_1(z_3) - 4 \max_e l_1(e) \\ &\geq (\alpha_1 + \alpha_2)L_1(z_{12}) + \alpha_3 L_1(z_3) - 8 \max_e l_1(e) \\ &= L_1(\alpha_1 z_1 + \alpha_2 z_2 + \alpha_3 z_3) - 8 \max_e l_1(e) \geq l'_1 - \frac{8\epsilon}{\gamma} \left(|C_1 - C'_1| - u_0^{(1)} \right), \end{aligned}$$

and similarly $L_2(z'_3) \geq l'_2 - 8\epsilon \left(|C_2 - C'_2| - u_0^{(2)} \right) / \gamma$. Meanwhile, we can easily see $\mathbf{1}^T z'_3 \leq \mathbf{1}^T m_3 = \mathbf{1}^T z_{23} = (\alpha_1 + \alpha_2)\mathbf{1}^T z'_2 + \alpha_3 \mathbf{1}^T z_3 \leq (\alpha_1 + \alpha_2)\mathbf{1}^T z_{12} + \alpha_3 \mathbf{1}^T z_3 = \mathbf{1}^T z_0 \leq k - \kappa$.

Let $\gamma = 8$ and M be the set of edges matched in z'_3 plus the κ edges reserved in the beginning, and A_1, A_2 be two multi-sets that are initially empty. For each $e \in M$, choose $i_1 \in F'_1, i_2 \in F'_2$ such that $d(i_1, i_2) \leq B$, and either (i_1, i_2) is merged into e or e is split from (i_1, i_2) , then add i_1 to A_1 , i_2 to A_2 .

It is easy to see that $T_1 \cup h(T_2) \subseteq A_1, T_2 \cup g(T_1) \subseteq A_2$, since they are secured by the κ reserved edges. We already know $T_1 \cup h(T_2)$ covers all clients in $C_1 - C'_1$, and from $L_1(z'_3)$ we know additionally at least $l'_1 - \epsilon \left(|C_1 - C'_1| - u_0^{(1)} \right)$ clients are covered in C'_1 , so the total number of clients covered by A_1 is at least

$$L_1(z'_3) + |C_1 - C'_1| \geq l'_1 - \epsilon \left(|C_1 - C'_1| - u_0^{(1)} \right) + u_0 + \sum_{i \in T_1} u_i^{(1)} \geq l_1 - \epsilon l_1 = (1 - \epsilon)l_1,$$

where we assume that $|C_1 - C'_1| - u_0^{(1)} < l_1$, otherwise we would have already covered $\geq l_1$ clients using T_1 . The proof is the same for A_2 and the second time step. \square