

Facility Location Problem with Capacity Constraints: Algorithmic and Mechanism Design Perspectives

Haris Aziz,¹ Hau Chan,² Barton E. Lee,¹ Bo Li,³ Toby Walsh⁴

¹UNSW Sydney and Data61 CSIRO

²Department of Computer Science and Engineering, University of Nebraska-Lincoln

³Department of Computer Science, University of Oxford

⁴TU Berlin, UNSW Sydney and Data61 CSIRO

haris.aziz@unsw.edu.au, hchan3@unl.edu, barton.e.lee@gmail.com, boli@cs.ox.ac.uk, tw@cse.unsw.edu.au

Abstract

We consider the facility location problem in the one-dimensional setting where each facility can serve a limited number of agents from the algorithmic and mechanism design perspectives. From the algorithmic perspective, we prove that the corresponding optimization problem, where the goal is to locate facilities to minimize either the total cost to all agents or the maximum cost of any agent is NP-hard. However, we show that the problem is fixed-parameter tractable, and the optimal solution can be computed in polynomial time whenever the number of facilities is bounded, or when all facilities have identical capacities. We then consider the problem from a mechanism design perspective where the agents are strategic and need not reveal their true locations. We show that several natural mechanisms studied in the uncapacitated setting either lose strategyproofness or a bound on the solution quality for the total or maximum cost objective. We then propose new mechanisms that are strategyproof and achieve approximation guarantees that almost match the lower bounds.

1 Introduction

In this paper, we study the *facility location problem with capacity constraints (FLP-CC)* from the *algorithmic* and *mechanism design* perspectives. In this version of the facility location problem, we have a set of agents and a set of facilities, where each agent is located somewhere on a line, and each facility has a capacity limiting the number of agents it can serve. From the *algorithmic perspective*, the locations of the agents are *publicly* known, and we are interested in the question of how to best locate the facilities to minimize the total travel distance/cost or the maximum cost of the agents to the located facilities. On the other hand, in our *mechanism design* setting, the locations of the agents are *privately* known to the agents themselves, and our goal is to design mechanisms that elicit the true locations of the agents and locate the facilities to minimize the total or maximum cost of the agents subject to the agents' reported locations.

Our FLP-CC models many real-world problems. The problems include locating schools, hospitals, warehouses, and libraries, all of which actually face capacity constraints.

In the one-dimensional setting, our models could be used to describe the setting of locating wastewater plants along a river or distribution centres along a highway. There are also various non-geographical settings that can be viewed as one-dimensional facility location problems (e.g. choosing the temperature for a classroom, or selecting a committee to represent people with different political views). In addition, there are settings where we can use the one-dimensional problem to solve more complex problems (e.g. decomposing the 2-d rectilinear problem into a pair of 1-d problems). The one-dimensional problem is also the starting point to consider more complex metrics (e.g., trees and networks).

In FLP-CC, we have n agents located on the real line, and we need to locate m facilities on the line to serve all the agents¹. The i th facility can serve up to c_i agents. We assume that $n \leq \sum_{i=1}^m c_i$ so that every agent can be served. Each agent j is at location x_j , and we suppose that the agents are ordered so that $x_1 \leq \dots \leq x_n$. Given an agent j , let $a_j \in \{1, \dots, m\}$ denote the facility that agent j is assigned and let N_i denote the set of agents assigned to facility i , i.e., $N_i := \{j \mid a_j = i\}$. A solution is a location y_i for each facility i , and an assignment of agents to facilities such that the capacity constraint c_i of each facility is not exceeded, i.e., $|N_i| \leq c_i$ for all $i \in \{1, \dots, m\}$. Accordingly, a solution is denoted by $\{(y_i, N_i)\}_{i=1}^m$. We consider a utilitarian measure: the *total cost*, $\sum_{j=1}^n |x_j - y_{a_j}|$; and an egalitarian measure: the *maximum cost*, $\max_{j \in \{1, \dots, n\}} |x_j - y_{a_j}|$. Our goal is to locate facilities on the line and assign agents to these facilities to minimize the total or maximum cost.

Contribution. Our main contributions are as follows. Firstly, for FLP-CC, we provide algorithmic results identifying the complexity of the corresponding optimization problem (NP-hard) but also provide tractability results, via dynamic programs (DPs), for various restricted settings. Secondly, we explore the mechanism design challenges introduced by this new setting. We show that many mechanisms which are considered desirable in uncapacitated FLPs become undesirable (with respect to strategyproofness and/or approximation ratio bounds) in FLP-CC. We introduce the *in-*

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Our problem naturally generalizes to higher dimensions as well as to non-Euclidean distance metrics.

nerpoint mechanism which performs relatively well in special cases, and characterize this mechanism as the only such strategyproof mechanism within a larger class of mechanisms, which we call *rank* mechanisms. Finally, we introduce a new strategyproof mechanism, *extended endpoint* mechanism (EEM), which achieves approximation guarantees that almost match the lower bounds. We summarize our contributions in Table 1. We leave open the question of

Table 1: Algorithmic & Mechanism Approximation Results

Algorithmic Results	Total Cost	Max. Cost
NP-hard	Exact	Exact
DP $O(2^m mn^2)$	Exact	Exact
SP Mechanisms*	Total Cost	Max. Cost
Median	Unbounded	Unbounded
Endpoint	Unbounded	Unbounded
Innerpoint ¹	$n/2 - 1$	2
EMM ²	$3n/2$	4

SP = Strategyproof, *two facilities, ¹when $c_1 = c_2 = n/2$,
²when $c_1 + c_2 \geq n$

upper bounds for three facilities which has remained open since the seminal work (Procaccia and Tennenholtz 2009; 2013), even when all facilities have infinite capacities.

1.1 Background

Algorithmic Perspective. The classic FLP with a single uncapacitated facility can be solved optimally in polynomial time when the objective is to minimize the total or maximum cost (e.g., Procaccia and Tennenholtz (2013)). Similarly, when extending to the setting with multiple uncapacitated facilities, the optimization problem remains tractable admitting a polynomial time solution for either objective function (Megiddo et al. 1981; Megiddo, Zemel, and Hakimi 1983). If there are multiple facilities with identical capacity constraints, the problem is tractable if the objective is to minimize the total cost but becomes intractable (NP-hard) for more general objective functions (Brimberg et al. 2001). Complementing these results, we show in the present paper that for our setting, i.e., the multiple capacitated FLP with non-identical capacity constraints, that minimizing either objective function is NP-hard. The intractability result in Brimberg et al. (2001) does not imply intractability in our setting for our objective functions. Moreover, we provide an alternate dynamic program to minimize the total cost for the identical capacity constraint setting studied by Brimberg et al. (2001). Finally, we note that our setting is different from the “capacitated k -facility location problem” (see e.g., (Pal, Tardos, and Wexler 2001; Levi, Shmoys, and Swamy 2012; Aardal et al. 2015)) where the set of potential facility locations is countable and bounded while, in ours, the set of facility locations is infinite (i.e., the real line). Thus, we cannot use previous results in the literature directly.

Mechanism Design Perspective. The classic FLP focuses on the setting where a single facility with unlimited capacity, i.e., an uncapacitated facility, must be located along a line with the goal of satisfying certain prop-

erties (most importantly strategyproofness) and/or optimizing some objective function. This setting escapes the famous impossibility result of Gibbard-Satterthwaite (1973; 1975) and inherits many well-known characterization results from social choice, such as those from Moulin (1980).

Two extensions of the classic FLP have attracted attention in the literature: multiple but uncapacitated facilities (see e.g., Miyagawa (2001), Heo (2013), Fotakis and Tzamos (2013), and Golowich, Narasimhan, and Parkes (2018)), and, to a lesser extent, a single capacitated facility. Considering the mechanism design problem for facilities with capacity constraints has only recently been considered by Aziz, Chan, Lee, and Parkes (2019) for a single facility. In their setting, additional strategic complications arise since some agents will not be served. The set of agents served by the facility is determined via an equilibrium outcome arising from the induced subgame.² We do not have such an issue since all our agents can be serviced. Another closely related work is by Procaccia and Tennenholtz (2013) which considers the problem of designing strategyproof mechanisms with approximation guarantees on an objective function (such as total cost and egalitarian welfare) for one or two facilities without capacity constraints. The works of (Fotakis and Tzamos 2013; Lu et al. 2010) prove that any deterministic strategyproof mechanism has an approximation ratio of $\Omega(n)$ for the total cost. Our EEM mechanism has approximation guarantees that almost match the lower bound of the uncapacitated setting.

2 An Algorithmic Perspective

We first show that it is intractable to find a solution that minimizes either the total or maximum cost. This result complements the result in Brimberg et al. (2001, Theorem 2) where intractability is proven for facilities with identical capacity constraints and a total cost objective function that can be non-monotonic in the distance between agents and facilities.

Theorem 1. *Computing a solution that minimizes the total or maximum cost is NP-hard even when there is no spare capacity in the FLP-CC.*

Proof. We show that the problem is NP-hard by reducing from the 3-partition problem, which is known to be strongly NP-hard (Garey and Johnson 1979). In a 3-partition problem, we are given a multiset $T = \{t_1, \dots, t_{3n}\}$ of positive integers of size $3n$, and we want to know if T can be partitioned into n subsets T_1, T_2, \dots, T_n , each of size three, such that the sum of the numbers in each subset is the same. Let $B = \frac{\sum_{t \in T} t}{n}$. In particular, we consider the 3-partition instances in which B and integers are polynomially bounded in n and $\frac{B}{4} < t_i < \frac{B}{2}$ for all i such that B is a positive integer and the integers are at least 1. Take such instance of 3-partition, we reduce it to the decision version of FLP-CC where we want to place the facilities to achieve the total cost and maximum cost of 0. We let $m = |S| = 3n$ to be the number of facilities, and, for each facility s_j , $j = 1, \dots, 3n$, set capacity $k_j = t_j$. With a slight abuse of notation, we let

²They show an impossibility result that ‘reasonable’ mechanisms which dictate the agents served can never be strategyproof.

$nB = |N|$ to be the number of agents such that there are n groups of B agents which are equally spaced apart and the agents in the same group are located at a single location.

Suppose that the instance of the 3-partition problems has a solution. It follows that there are T_1, T_2, \dots, T_n such that each sum up to B . For each group i of B agents, we place three facilities of capacities $t_j \in T_i$ at the location of group i . Clearly, we achieve the total cost or maximum cost of 0.

Suppose that we have a solution to the constructed problem. We have a location a_j for each facility s_j such that the total cost or maximum cost is 0. In fact, this is the best cost we can obtain where the agents do not have to travel at all. Moreover, each facility must be located on the location of one of the groups and all of the B agents within each group can use one of the facilities at its location. Since the capacity of facility s_j is $\frac{B}{4} < k_j = t_j < \frac{B}{2}$, each location of the group has at least three facilities. Also since the total number of facilities is $3n$ and there are n groups, the number of facilities at each group is exactly 3. Finally, the sum of the capacities of the three facilities is B . Thus, for each group i , we can construct T_i by take elements with the same values as capacities of the facilities located at group i .

To conclude, the problem instance above is of polynomial size bounded by n and we obtain our claimed result. \square

We observe that the computational complexity of finding an optimal solution does not come from having to decide where to distribute any spare capacity. The reduction uses a FLP-CC instance where the facilities have a total capacity exactly equal to the number of agents to be served. The computational complexity comes from having to find the best of the $m!$ possible left to right orderings of these facilities.

As we show next, to minimize the maximum cost (total cost), we locate each facility in the midpoint (median) of the continuous region of agents it serves. Thus, for a fixed number of facilities or facilities of identical capacity, we can use dynamic programming to compute the optimal (total or maximum cost) solution in polynomial time.

Order the set of agents N such that $x_1 \leq x_2 \leq \dots \leq x_{n-1} \leq x_n$. We say that a set of agents N_j is *continuous* if $N_j = \{x_{j_1}, x_{j_1+1}, \dots, x_{j_2}\}$ for some $j_1 \leq j_2$.

Lemma 1. *For either objective function (total cost or maximum cost), there exists an optimal solution $\{(y_j, N_j)\}_{j=1}^m$ such that each N_j is continuous.*

Proof. Suppose that the solution $\{(y_j, N_j)\}_{j=1}^m$ provides minimal total cost but is such that there exists j for which N_j is not continuous. That is, there exists some pair of agents $p, s \in N_j$ such that (1) $y_j \leq x_p < x_s, p \notin N_j$ and $s \in N_j$, or (2) $x_s < x_p \leq y_j, p \notin N_j$ and $s \in N_j$. Both cases are dealt with similarly and so we assume the first case.

Let $\ell : p \in N_\ell$. For any location of x_s , if $y_j \leq x_p < y_\ell$, then swapping p from N_ℓ and s from N_j leads to strictly lower total cost. On the other hand, if $y_\ell \leq y_j$, then swapping p from N_ℓ and s from N_j leads to no change in total cost. Repeating this method until each set N_j is continuous leads to a solution with at least as low total cost. Since the original solution was optimal, we conclude that there exists an optimal solution with each N_j continuous.

Now consider the maximum cost objective. A similar argument applies. The only difference is that in the case where $y_j \leq x_p < y_\ell$ we can only guarantee weakly lower maximum cost. Nonetheless, the same conclusion is reached — there exists an optimal maximum cost solution such that each N_j is continuous. \square

The above lemma shows that partitions, or sets N_j , must be continuous but it does not require that facilities be located ‘inside’ such a partition. However, it is well-known that the median of a set of points minimizes the sum of absolute deviations and the midpoint minimizes the maximum deviation. This leads to the following corollary.

Corollary 1. *An optimal solution $\{(y_j, N_j)\}_{j=1}^m$ is given by by continuous partitions of the agents which correspond to facilities such that each facility s_j locates at the median and midpoint of its associated partition N_j for total cost and maximum cost, respectively.*

The above results show that to solve FLP-CC for total and maximum costs, it suffices to look for continuous sets which match the capacity of each facility and locate the facility at the median or midpoint of each set.

2.1 Bounded Number of Facilities

Total Cost. We compute the optimal way (for minimizing the total cost) of locating m facilities and partition agents into m parts using dynamic program. We will use $[j, j']$ to denote the set of agents $\{j, \dots, j'\}$ for any $j \leq j'$. We use $OPT(j, j', k, S)$ to represent the optimal total cost when agents $[1, j']$ are partitioned into k parts, and $j' \geq j$ are in the k -th partition, and S is the subset of k facilities used to serve agents $[1, j']$. We will make use of a $n \times n \times m \times 2^m$ array M , whose entries are initially set to *empty*. We will use $M[j, j', k, S]$ to store $OPT(j, j', k, S)$. We use $v([j, j'], s)$ to denote the optimal total cost when agents $[j, j']$ are served by facility s (with capacity c). In order to compute $v([j, j'], s)$ we can simply try to locate s at each of the agent locations and serve c nearest agents.³

We invoke $OPT^T(n, n, m, S)$ (Algorithm 1) to compute the optimal total cost and recover the solution from the values stored in M .

Theorem 2. *Algorithm 1 computes the optimal total cost solution in $O(2^m mn^2)$ time. Hence, it is fixed parameter tractable in m .*

Proof. It is easy to see that the running time of Algorithm 1 is $O(2^m mn^2)$. We argue for the proof of correctness. The base case is clear. The total cost is zero if there are no agents or no facilities. Suppose we want to compute $OPT(j, j', k, S)$ and we have already computed values for $OPT(h, h', \hat{k}, \hat{S})$ where $h < j$ or $h' < j'$ or $\hat{k} < k$ or $\hat{S} \subset S$. We distinguish between two cases depending on whether $j - 1$ is covered by the last facility in S .

In the first case, $OPT(j, j', k, S)$ is equivalent to $OPT(j - 1, j', k, S)$. In the second case we know that $j - 1$

³The optimal point will be the median of those subset of agents who are actually served.

Algorithm 1 $OPT^T(j, j', k, S)$

```
1: if  $j = 0$  or  $k = 0$  or  $S = \emptyset$  then
2:   return 0
3: else if  $M[j, j', k, S]$  not empty then
4:   return  $M[j, j', k, S]$ 
5: else
6:    $s^* \leftarrow \arg \min_{s \in S} \{OPT^T(j-1, j-1, k-1, S \setminus \{s\}) + v([j, j'], s)\}$ 
    $M[j, j', k, S] \leftarrow \min\{OPT^T(j-1, j', k, S),$ 
    $OPT^T(j-1, j-1, k-1, S \setminus \{s^*\}) + v([j, j'], s^*)\}$ 
7:   return  $M[j, j', k, S]$ 
8: end if
```

is not served by the last facility in S but j is. Hence, if j is served by facility $s \in S$, then $OPT(j, j', k, S')$ is equivalent to $OPT(j-1, j-1, k-1, S \setminus \{s\}) + v([j, j'], s)$. \square

Maximum Cost. In a similar way, we use $OPT^M(i, i', j, S')$ to represent the optimal maximum cost when agents $[1, i']$ are partitioned into j parts, and $i' \geq i$ are in the j -th partition, and S' is the subset of j facilities. We use $OPT^M(i, i', j, S')$ to denote optimal maximum cost when agents $[1, i']$ are partitioned into j parts, and i and $i' \geq i$ are in the j -th partition, and S' is the subset of j facilities. We will make use of a $n \times n \times m \times 2^m$ array E , whose entries are initially set to *empty*. We will use $E[i, i', j, S']$ to store $OPT^M(i, i', j, S')$. We use $v^M([i, i'], s_j)$ to denote the optimal maximum cost when agents $[i, i']$ are covered by s_j (with capacity k_j). We invoke $OPT^M(n, n, m, S)$ to compute the optimal maximum cost and recover the solution from the values stored in E . The subroutine to solve OPT is summarized as Algorithm 2.

Algorithm 2 $OPT^M(i, i', j, S')$

```
1: if  $i = 0$  or  $j = 0$  or  $S' = \emptyset$  then
2:   return 0
3: else if  $E[i, i', j, S']$  not empty then
4:   return  $E[i, i', j, S']$ 
5: else
6:    $E[i, i', j, S'] \leftarrow \min(\max\{OPT^M(i-1, i', j, S'), \min_{s \in S'} \{\max(OPT^M(i-1, i-1, j-1, S \setminus \{s\}), v^M([i, i'], s))\}\}$ 
7:   return  $E[i, i', j, S']$ 
8: end if
```

Theorem 3. *Algorithm 2 computes the optimal total cost solution in $O(2^m mn^2)$ time. Hence, it is fixed parameter tractable in m .*

Identical capacities When facilities have the same capacity, we can use a different dynamic program to compute an optimal solution with minimum total cost. We exploit the property that an optimal solution consists of m continuous regions served by a facility located at the median agent of

the region. We construct an array $C(i, j)$ which is the optimal total cost supposing the leftmost i agents are served by the first j facilities. Let c be the capacity of the facilities. We initialize $C(i, j)$ with ∞ . In the first round, we set $C(i, 1)$ for i in 1 to c with the optimal total cost for the first i agents served by a single facility. In the j th round, for each k from j to $\min(n, jc)$ and i from 1 to $\min(c, k-j+1)$ we update $C(k, j)$ with the minimum of $C(k, j)$ and $C(k-i, j-1)$ plus the total cost of serving the $k-i$ th to the k th agents from left by the j th facility. In the final round, the optimal total cost will be computed at $C(n, m)$. The runtime to update each of the $O(nm)$ entries is $O(c^2)$ giving a total cost of $O(nmc^2)$. Note that we can suppose $m < n$ otherwise the optimal solution will have zero total cost. Similarly, we can suppose $c < n$. Hence, the total cost is polynomial in n . We can construct a similar dynamic program to compute the optimal maximum cost, exploiting the fact that an optimal solution consists in this case of m continuous regions served by a facility located at the midpoint between the leftmost and rightmost agents served by the facility. The total run time, in this case, is $O(nmc)$. These results complement those in Brimberg et al. (2001) where an alternate dynamic program with run time $O(m^3 n^2)$ is provided for minimizing the total cost when facilities have identical capacities.

Theorem 4. *There exists an $O(nmc^2)$ time algorithm to compute the optimal total cost and an $O(nmc)$ time algorithm to compute the optimal maximum cost with n agents when all m facilities have the same capacity c .*

3 A Mechanism Design Perspective

In our mechanism design setting, each agent j 's position x_j is her private information and a mechanism M operates on the reported locations (x'_1, \dots, x'_n) of all agents, which may be different from their actual locations. Based on the reported locations, the mechanism $M(x'_1, \dots, x'_n)$ locates m facilities, $y_1, \dots, y_m \in \mathbb{R}$, and allocates disjoint sets of agents to each facility while respecting the capacity constraints (i.e., assign $a_j \in \{1, \dots, m\}$ for each agent j where $|\{j \mid a_j = i\}| \leq c_i$ for each facility i). The total and maximum costs of a mechanism are defined in the same way. Accordingly, every agent j 's cost is $u_j(x_j, M(x'_1, \dots, x'_j, \dots, x'_n)) = |x_j - y_{a_j}|$ and they prefer mechanism outcomes that allocate them to a facility closer to their true location x_j . A mechanism is said to be *strategyproof* if for every agent j and for any profile of other agent reports, agent j weakly prefers the mechanism outcome achieved by reporting her true location $x'_j = x_j$ to the outcome achieved by any other report x''_j . Our goal in this section is to find strategyproof mechanisms.

When $m = 1$ and $c_1 \geq n$, locating the facility at the median location is strategyproof and optimal for minimizing the total cost. As a result, we focus on the case of two facilities. When $m = 2$ and $c_1, c_2 \geq n$, Procaccia and Tennenholtz (2013) first argue that no deterministic strategyproof mechanism is guaranteed to return the optimal solution for either the total or the maximum cost.

In general, if we set the capacities of the locations high enough, we can inherit the lower bounds results from exist-

ing lower bound results. Particularly, (Lu et al. 2010) proved that any deterministic strategyproof mechanism has an approximation ratio for the total cost of at least $\frac{n}{2} - 1$, which is further improved to be $n - 2$ by (Fotakis and Tzamos 2013). Procaccia and Tennenholtz (2013) proved that any deterministic strategyproof mechanism has an approximation ratio for the maximum cost of at least 2. An approximation ratio of α (w.r.t. some objective function to be minimized) implies that, across all instances of agent locations, the mechanism outcome evaluated according to the objective function is no more than a factor of α larger than the optimal outcome.

3.1 Limitations of Existing Approaches

The first contender is the **median mechanism** which has been studied for the case of one facility with unlimited capacity. We can extend this mechanism to two or even more facilities with capacity constraints as follows: locate all facilities at the median agent. The mechanism remains strategyproof. However, we can no longer bound the approximation ratio for either the total or maximum cost. Suppose, for instance, most of the facilities are concentrated at the end points in the optimal solution. We next explore the idea of placing facilities at the endpoints with the **endpoint mechanism** which locates the facilities at the reported endpoints (Procaccia and Tennenholtz 2009). With two uncapacitated facilities, the endpoint mechanism has several nice properties. It is, for instance, strategyproof and returns a 2-approximation of the maximum cost, which is proved to be optimal by (Procaccia and Tennenholtz 2009). Unfortunately, it is not hard to verify that, when we add (even uniform) capacities, the endpoint mechanism is not strategyproof (i.e., agents can strategically exploit capacity limits to ensure a more favourable outcome). The main difficulty here is how to allocate the agents to the facilities.

3.2 Mechanisms for Two Facilities with Equal Capacity for Total and Maximum Costs

We now present another mechanism for two capacitated facilities. Suppose we have two facilities of capacities c_1 and c_2 where $n = c_1 + c_2$. If we order agents from left to right, the **innerpoint mechanism** places one facility at the c_1 th agent from left, serving the leftmost c_1 agents, and the other facility at the $c_1 + 1$ th agent from the left, serving the rightmost c_2 agents. If $c_1 \neq c_2$, we suppose as before that the order of facilities is fixed in advance of the agents reporting.

Theorem 5. *When n is even and 2 facilities have equal capacity $n/2$, the innerpoint mechanism is strategyproof and has a bounded approximation ratio of 2 for maximum cost and a ratio of $\frac{n}{2} - 1$ for total cost.*

Proof. To show strategyproofness, we consider three cases. In the first, consider one of the $c_1 - 1$ leftmost agents. The only way they can change the outcome is if they report a location to the right of the c_1 th agent. But this will give them a worst outcome. In the second case, consider the c_1 th agent from the left. They have zero cost. Reporting any other location leaves this the same or worse. In the third case, consider one of the c_2 rightmost agents. This case is analogous to the first two cases. In all cases, misreporting is not beneficial.

To determine the approximation ratio for the maximum cost, we rescale the problem so leftmost agent is at 0, and rightmost agent at 1. This does not affect the approximation ratio. By assumption, the two facilities have equal capacity c . Suppose the c th agent from left is at x , and the $c + 1$ th agent is at $1 - y$. The maximum cost is then $\max(x, y)$. An optimal location of facilities would put one facility at $\frac{x}{2}$, and the other at $1 - \frac{y}{2}$. This has a maximum cost half as much of $\max(\frac{x}{2}, \frac{y}{2})$. Hence, the approximation ratio is 2.

To determine the approximation ratio for the total cost, we again rescale the problem so the leftmost agent is at 0, and the rightmost agent at 1. Suppose the c th agent from left is at x , and the $c + 1$ th agent is at $x + u$. The worst case for the total cost has $c - 1$ agents at 0, the c th agent at x , the $c + 1$ th agent at $x + u$ and the $c - 1$ remaining agents at 1. This gives a total cost of $(c - 1)x + (c - 1)(1 - x - u)$. That is, $(c - 1)(1 - u)$. The optimal solution, in this case, puts one facility at 0 and the other at 1, giving a total cost of $x + (1 - x - u)$. That is, $1 - u$. Hence, the approximation ratio for the total cost is at least $c - 1$ (or $\frac{n}{2} - 1$). \square

Unfortunately, innerpoint no longer has a bounded approximation ratio for the total or maximum cost when facilities have spare capacity i.e., $c_1 + c_2 > n$. The following theorem can be proved by constructing a suitable example.

Theorem 6. *With 4 agents, and 2 capacitated facilities of size 3 or larger, the innerpoint mechanism has an unbounded approximation ratio for the total or maximum cost.*

Proof. Suppose two agents are at 0, one at ϵ and the fourth agent at 1. The innerpoint mechanism will locate one facility at 0, serving the two leftmost agents, and the other at ϵ , serving the two rightmost agents. The total and maximum costs are both $1 - \epsilon$ while the optimal total and maximum costs for two facilities with capacity 3 or greater are both ϵ . The approximation ratio for total or maximum costs is therefore $\frac{1-\epsilon}{\epsilon}$ which tends to infinity as ϵ tends to zero. \square

Inspired by the median, endpoint and innerpoint mechanisms, we introduce a family of mechanisms that generalize all three. Given $m \geq 2$ facilities and n agents, we let the capacity of the i th facility be c_i and agent j report location x_j where $x_1 \leq \dots \leq x_n$. We suppose $n = \sum_{i=1}^m c_i$. The **rank mechanism** has $m + 1$ parameters, t_1 to t_m with $t_1 \leq \dots \leq t_m$ and a permutation π of 1 to m . The mechanism locates facility $\pi(i)$ at x_{t_i} . Agents are then allocated to the facilities from left to right (while also respecting each facility's capacity constraint). If facilities have identical capacities, then we can ignore the permutation. We also note that our rank mechanism is similar to the percentile mechanisms introduced in (Sui, Boutilier, and Sandholm 2013), but the later only works for uncapacitated facilities.

We now provide a simple characterization of those rank mechanisms which are strategyproof.

Theorem 7. *The rank mechanism with parameters t_1 to t_m and π is strategyproof if and only if either $t_i = t_j$ for any i and j , or there exists $1 \leq k < m$ with $t_1 = t_k = \sum_{i=1}^k c_{\pi(i)}$ and $t_{k+1} = t_m = t_1 + 1$.*

Proof. (Necessity) There are two cases. In the first case, $t_i = t_j$ for any i and j . All facilities are located at the same location. An agent to the left of this can only mis-report and move the location of the facilities to the right which is not in their interest. Similarly, an agent to the right can only mis-report and move the location of the facilities to the left which is not in their interest. In the second case, there exists $1 \leq k < m$ with $t_1 = t_k = \sum_{i=1}^k c_{\pi(i)}$ and $t_{k+1} = t_m = t_1 + 1$. We can view this as equivalent to lumping the facilities into two “super” facilities, setting $m = 2$, $t_1 = c_{\pi(1)}$ and $t_2 = t_1 + 1$. This is the innerpoint mechanism which we have shown previously to be strategyproof.

(Sufficiency). There are again two cases. In the first case, $m > 2$. We give a counter-example for strategyproofness for $m = 3$. For greater m , we pad the counter-example with additional facilities and agents. We give a counter-example for $t_1 = t_2 = 1$ and $t_3 = 2$. Similar counter-examples can be constructed for other values of t_i . We suppose facilities have capacity 2, and $x_1 = 0, x_2 = 3, x_3 = 4, x_4 = 5, x_5 = 6, x_6 = 7$. Agent 2 has an incentive to mis-report their position as location 7. In the second case, $m = 2$ and $t_1 \neq c_{\pi(i)}$. We give a counter-example for $t_1 = 1$ and $t_2 = 2$. Similar counter-examples can be constructed for other values of t_1 and t_2 . Suppose the two facilities have capacity 2, and $x_1 = 0, x_2 = 3, x_3 = 4, x_4 = 5$. Agent 2 has an incentive to mis-report their location to be position 5. \square

As an example, consider the **quartile mechanism** that places one facility at the first quartile, and the second at the third quartile. It follows immediately from the last theorem that the quartile mechanism is not strategyproof.

We now provide a simple necessary condition for the rank mechanism to have bounded approximation ratio.

Theorem 8. *With 2 (or more) facilities with capacity constraints, if the rank mechanism has bounded approximation ratio for the total and maximum cost it must be that $t_1 \leq c_1$ and $t_m \geq c_1 + 1$ (recall that $t_1 \leq t_m$), where c_1 is the capacity of facility $\pi(1)$.*

Proof. Consider an instance where c_1 agents located at 0 and the remaining $n - c_1 = \sum_{j=2}^m c_j$ agents are located at 1. The optimal total and maximum costs are zero. If either $t_1 > c_1$ or $t_m < c_1 + 1$ holds, since $t_j \leq t_k$ for all $j \leq k$, all facilities will be located at the same location (either 0 or 1). This necessarily gives non-zero total and maximum costs and hence unbounded approximation ratios. \square

Theorem 5, Theorem 7, and Theorem 8 collectively lead to the following theorem which characterizes the innerpoint mechanism as the only strategyproof rank mechanism with bounded approximation ratio (for either total or maximum cost) when there are two facilities of equal capacities.

Theorem 9. *With 2 capacitated facilities of equal capacity, the only rank mechanism which is both strategyproof and provides a bounded approximation ratio for both the maximum and total cost is the innerpoint mechanism, i.e., $t_1 = c_1$ and $t_2 = c_1 + 1$.*

In summary, the mechanisms in this section are either unbounded or not strategyproof or only satisfy the desirable

properties if $c_1 = c_2 = n/2$. Next we present a new mechanism that is strategyproof and provides bounded approximations even for two capacitated facilities and the case where there may be spare capacities.

3.3 Mechanisms for Two Facilities with Arbitrary Capacities for Total and Maximum Costs

In this section, we present a more general strategyproof mechanism that almost matches the existing lower bounds of (Lu et al. 2010) and (Fotakis and Tzamos 2013). Let $x = (x_i)_{i \in N}$ be the location profile such that $x_1 \leq \dots \leq x_n$ and $c_1 \geq c_2 \geq 1$ be the capacities such that $c_1 + c_2 \geq n$. Let f_1 and f_2 be the output locations of the two facilities.

Extended Endpoint Mechanism (EEM) Let $X_1 = \{i | x_i - x_1 \leq \frac{1}{2}(x_n - x_1)\}$ and $X_2 = \{i | x_n - x_i < \frac{1}{2}(x_n - x_1)\}$. If $|X_1| \geq |X_2|$, execute one of the following three cases.

Case 1. If $|X_1| \leq c_1$ and $|X_2| \leq c_2$, $f_1 = x_1$ and $f_2 = x_n$. Agents in X_1 are allocated to f_1 and the others are allocated to f_2 .

Case 2. If $|X_1| > c_1$ and $|X_2| \leq c_2$, $f_1 = 2x_{c_1+1} - x_n$ and $f_2 = x_n$. Agents $\{1, \dots, c_1\}$ are allocated to f_1 and the others are allocated to f_2 .

Case 3. If $|X_1| \leq c_1$ and $|X_2| > c_2$, $f_1 = x_1$ and $f_2 = 2x_{n-c_2} - x_1$. Agents in $\{1, \dots, n - c_2\}$ are allocated to f_1 and the others are allocated to f_2 .

If $|X_1| < |X_2|$, switch the roles of the two facilities in above cases and execute one of them.

EEM is essentially an endpoint mechanism, but to restore strategyproofness, it locates one facility outside of $[x_1, x_n]$. Although EEM has slightly poorer performance than the classical endpoint mechanism, with respect to the lower bound results, EEM is actually optimal up to a constant.

Theorem 10. *EEM is strategyproof.*

Proof. Without loss of generality, assume $|X_1| \geq |X_2|$ as the other case is symmetric. Note that $f_k \leq x_1$ and $f_{2-k} \geq x_n$ always hold for some $k \in \{1, 2\}$.

For Case 1 where $|X_1| \leq c_1$ and $|X_2| \leq c_2$, each agent is served by the facility located at the nearest endpoint to her, and the two endpoint agents get the best possible solution. If any agent except the two endpoint agents deviates and the resulting locations of the two facilities change, she can only be worse off as under no situation she could be served by a facility located in (x_1, x_n) . Thus our mechanism is strategyproof in this case.

Next we consider Case 2 where $|X_1| > c_1$ and $|X_2| \leq c_2$.

- For any agent $1 \leq i \leq c_1$, $i \in X_1$ and is served by the facility located at $f_1 = 2x_{c_1+1} - x_n$. Under EEM, any deviation by agent i to $x'_i < x_{c_1+1}$ will continue to allocate agent i to the left most facility whose location f'_1 either remains unchanged at $2x_{c_1+1} - x_n$, or, if Case 1 occurs, is moved to $f'_1 = x'_i < x_1$. It is straightforward to see that neither of these cases is profitable for the agent, and so we restrict our attention to deviations $x'_i \geq$

x_{c_1+1} . Let X'_1 and X'_2 be the new partitions of the EEM. If X'_1, X'_2 are such that Case 1 or Case 2 holds then agent i is now allocated to $f'_2 = \max\{x_n, x'_i\}$, this is never a profitable deviation. Finally, we note that Case 3 can never occur since this would require that $|X_1| = c_1 + 1$ and $|X_2| = c_2$ which contradicts $|X_1| + |X_2| = n$ and $c_1 + c_2 \geq n$. Thus i cannot be better off by deviating.

- Next, we consider agent $i \geq c_1 + 1$. If she reports some location on the left of x_{c_1} , then agent c_1 will replace the role of $c_1 + 1$ in the mechanism thus $f'_1 = 2x_{c_1} - x_n$ and i is connected to f'_1 . However, $x_i - f'_1 \geq x_{c_1} - f'_1 = x_n - x_{c_1} \geq x_n - x_i$, thus i does not profit from the misreporting their location. Note that i 's deviation cannot cause the number of agents in X_1 to be less than the number of agents in X_2 . Otherwise, $|X_1| = |X_2| + 1$, then c_1 must be at least as large as $|X_1|$ as $c_1 \geq c_2$ and $c_1 + c_2 \geq n$. Thus Case 1 must happen when i tells the truth. Accordingly, if i reports some location on the right of x_{c_1} , she is still connected to f_2 (or even the right of f_2), no matter which case happens after the deviation. Thus i never benefits from misreporting.

By a similar argument as for Case 2, EEM is strategyproof under Case 3. \square

Theorem 11. *For the objective of total cost, EEM has an approximation ratio of $\frac{3n}{2}$.*

Proof. Again, without loss of generality, assume $|X_1| \geq |X_2|$. We first note that the optimal solution has the following form: partition of the agents into left successive c_1 agents and right successive c_2 successive agents and locate each facility at the median point of each subset of the agents.

For Case 1, EEM is exactly the same as the classical Endpoint Mechanism (Procaccia and Tennenholtz 2009), thus has an approximation of $n - 2$.

For Case 2, $OPT \geq x_n - x_{c_1+1}$. This is because in the optimal solution, all the agents in X_1 cannot be served by a single facility as $|X_1| > c_1 \geq c_2$. Thus at least one of the agents in X_1 has to be grouped with x_n . Moreover,

$$\begin{aligned} ALG &\leq c_1(x_{c_1} - f_1) + (n - c_1)(f_2 - x_{c_1+1}) \\ &\leq c_1(x_{c_1+1} - f_1) + (n - c_1)(f_2 - x_{c_1+1}) \\ &= n(f_2 - x_{c_1+1}) \leq nOPT. \end{aligned}$$

For Case 3, $OPT \geq \frac{1}{2}(x_n - x_1)$. This is because in the optimal solution X_1 cannot be served by one facility and X_2 by another facility. Thus either at least one of the agents in X_1 has to be grouped with x_n or at least one of the agents in X_2 has to be grouped with x_1 . No matter which case happens, there is one group such that the longest distance between each pair of agents of this group is at least $\frac{1}{2}(x_n - x_1)$. Moreover,

$$\begin{aligned} ALG &\leq |X_1|(x_{|X_1|} - f_1) + (n - c_2 - |X_1|)(x_{n-c_2} - f_1) \\ &\quad + c_2(f_2 - x_{n-c_2+1}) \\ &\leq |X_1| \cdot \frac{1}{2}(x_n - x_1) + (n - |X_1|) \cdot (x_n - x_1) \\ &\leq \frac{3n}{4}(x_n - x_1) \leq \frac{3n}{2}OPT. \end{aligned}$$

The second inequality is because $f_2 - x_{n-c_2+1} \leq f_2 - x_{n-c_2} = x_{n-c_2} - x_1 \leq x_n - x_1$. The third inequality is because the term in the second line is maximized when $|X_1| = \frac{n}{2}$ as $|X_1| \geq |X_2|$. \square

Theorem 12. *For the objective of maximum cost, EEM has an approximation ratio of 4.*

Proof. For Case 1, EEM is exactly the same as the classical Endpoint Mechanism (Procaccia and Tennenholtz 2009), thus has an approximation of 2.

For Case 2, let $dis = x_n - x_{c_1+1}$. We claim that $OPT = \frac{1}{2}dis$. We first note that $(f_1^*, f_2^*) = (\frac{1}{2}(x_{c_1} - x_1), \frac{1}{2}(x_n + x_{c_1+1}))$ is a feasible solution with the announced cost, i.e. (f_1^*, f_2^*) an optimal solution. Let (l_1, l_2) be any solution. It is easy to see that to guarantee a cost that is not greater than $\frac{1}{2}dis$, $\max\{l_1, l_2\} \geq \frac{1}{2}(x_n + x_{c_1+1})$ as agent n has to be connected to one of them. Moreover, since the capacity of the leftmost facility is at most c_1 , at least one of agents $\{1, \dots, c_1, c_1 + 1\}$ has to be grouped with agent n , thus the optimal solution is to put $c_1 + 1$ into n 's group and its serving facility is located at $\frac{1}{2}(x_n + x_{c_1+1})$. We see that the cost of (l_1, l_2) cannot be smaller than $\frac{1}{2}dis$.

Let (f_1, f_2) be the output of EEM. Then $x_{c_1} - f_1 \leq x_{c_1+1} - f_1 = x_n - x_{c_1+1} = dis$. Thus, $ALG \leq dis = 2OPT$, which is a 2-approximation.

Case 3 is similar to Theorem 11. Since the optimal solution cannot serve X_1 by one facility and X_2 by another facility, either one of the agents in X_1 has to be grouped with x_n or one of the agents in X_2 has to be grouped with x_1 . Thus $OPT \geq \frac{1}{4}(x_n - x_1)$. Again, since $f_2 - x_{n-c_2+1} \leq f_2 - x_{n-c_2} = x_{n-c_2} - x_1 \leq x_n - x_1$. $ALG \leq 4OPT$, which is a 4-approximation. \square

4 Conclusion and Discussion

We considered the FLP-CC both algorithmically and from a mechanism design perspective. See Table 1 for a summary of our results. Collectively, our results show that the addition of capacity constraints to the FLP makes it more difficult compared to the uncapacitated case to solve optimally and to design mechanisms with desirable properties (e.g., strategyproofness and bounds on the approximation ratio). There are many directions for future work. The most important question is to understand whether any meaningful upper bound can be established for the case of three locations. Moreover, can we extend results beyond one dimension to trees, networks, or two-dimensional rectilinear and Euclidean metrics? Finally, it is interesting to consider if randomization can help to design more efficient and better algorithms or mechanisms for FLP-CC.

Acknowledgements

Aziz is supported by a UNSW Scientia Fellowship, and Defence Science and Technology (DST) under the project "Auctioning for distributed multi vehicle planning" (DST 9190). Lee is supported by a Data61 and a UNSW Scientia PhD fellowship. Li is supported by the ERC grant number 639945 (ACCORD). Walsh is funded by the ERC under Horizon 2020 via AMPLify 670077.

References

- Aardal, K.; van den Berg, P. L.; Gijswijt, D.; and Li, S. 2015. Approximation algorithms for hard capacitated k -facility location problems. *European Journal of Operational Research* 242(2):358–368.
- Aziz, H.; Chan, H.; Lee, B. E.; and Parkes, D. C. 2019. The capacity constrained facility location problem. In *Proceedings of the 15th Conference on Web and Internet Economics (WINE)*, To appear.
- Brimberg, J.; Korach, E.; Eben-Chaim, M.; and Mehrez, A. 2001. The capacitated p -facility location problem on the real line. *International Transactions in Operations Research* 8:727–738.
- Fotakis, D., and Tzamos, C. 2013. On the power of deterministic mechanisms for facility location games. In Fomin, F.; Freivalds, R.; Kwiatkowska, M.; and Peleg, D., eds., *Automata, Languages, and Programming*, 449–460.
- Garey, M., and Johnson, D. 1979. *Computers and intractability: a guide to the theory of NP-completeness*. W.H. Freeman.
- Gibbard, A. 1973. Manipulation of voting schemes: A general result. *Econometrica* 41:587–601.
- Golowich, N.; Narasimhan, H.; and Parkes, D. 2018. Deep learning for multi-facility location mechanism design. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, 261–267.
- Heo, E. 2013. Strategy-proof rules for two public goods: double median rules. *Social Choice and Welfare* 41(4):895–922.
- Levi, R.; Shmoys, D. B.; and Swamy, C. 2012. Lp-based approximation algorithms for capacitated facility location. *Mathematical programming* 131(1-2):365–379.
- Lu, P.; Sun, X.; Wang, Y.; and Zhu, Z. 2010. Asymptotically optimal strategy-proof mechanisms for two-facility games. In *Proceedings of the 11th ACM conference on Electronic commerce*, 315–324.
- Megiddo, N.; Tamir, A.; Zemel, E.; and Chandrasekaran, R. 1981. An $o(n \log^2 n)$ algorithm for the k -th longest path in a tree, with applications to location problems. *SIAM Journal on Computing* 10:328–337.
- Megiddo, N.; Zemel, E.; and Hakimi, S. 1983. The maximum coverage location problem. *SIAM Journal on Algebraic and Discrete Methods* 4:253–261.
- Miyagawa, E. 2001. Locating libraries on a street. *Social Choice and Welfare* 18(3):527–541.
- Moulin, H. 1980. On strategy-proofness and single peakedness. *Public Choice* 35(4):437–455.
- Pal, M.; Tardos, T.; and Wexler, T. 2001. Facility location with nonuniform hard capacities. In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, 329–338.
- Procaccia, A. D., and Tennenholtz, M. 2009. Approximate mechanism design without money. In *Proceedings 10th ACM Conference on Electronic Commerce*, 177–186.
- Procaccia, A., and Tennenholtz, M. 2013. Approximate mechanism design without money. *ACM Trans. Econ. Comput.* 1(4):18:1–18:26.
- Satterthwaite, M. 1975. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10:187–216.
- Sui, X.; Boutilier, C.; and Sandholm, T. 2013. Analysis and optimization of multi-dimensional percentile mechanisms. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI '13*, 367–374.