Multi-Stage Facility Location Problems with Transient Agents

Xuezhen Wang[1], Vincent Chau[2*], Hau Chan[3], Ken C.K. Fong[4*], Minming Li[1]

[1] Department of Computer Science, City University of Hong Kong, HKSAR China
[2] School of Computer Science and Engineering, Southeast University, China
[3] Department of Computer Science and Engineering, University of Nebraska-Lincoln, USA
[4] Department of Computing and Decision Sciences, Lingnan University, HKSAR China
xuezhenwang2@gmail.com, vincentchau@seu.edu.cn, hchan3@unl.edu, kenfong@ln.edu.hk, minming.li@cityu.edu.hk

Abstract

We study various models for the one-dimensional multi-stage facility location problems with transient agents, where a transient agent arrives in some stage and stays for a number of consecutive stages. In the problems, we need to serve each agent in one of their stages by determining the location of the facility at each stage. In the first model, we assume there is no cost for moving the facility across the stages. We focus on optimal algorithms to minimize both the social cost objective, defined as the total distance of all agents to the facility over all stages, and the maximum cost objective, defined as the max distance of any agent to the facility over all stages. For each objective, we give a slice-wise polynomial (XP) algorithm (i.e., solvable in $m^{f(k)}$ for some fixed parameter k and computable function f, where m is the input size) and show that there is a polynomial-time algorithm when a natural first-come-first-serve (FCFS) order of agent serving is enforced. We then consider the mechanism design problem, when the agents' locations and arrival stages are private, and design a group strategy-proof mechanism that achieves good approximation ratios for both objectives and settings with and without FCFS ordering. In the second model, we consider the facility's moving cost between adjacent stages under the social cost objective, which accounts for the total moving distance of the facility. Correspondingly, we design XP (and polynomial time) algorithms and a group strategy-proof mechanism for settings with or without the FCFS ordering.

1 Introduction

In this paper, we consider the *Multi-Stage Facility Location Problems with Transient Agents (MSFLP-TA)*. In the basic setting, we have a set of agents arriving at different stages with a tolerance rate that indicates at most how many stages an agent can stay for before departure. A facility is reallocated at every stage to serve some of the (existing) agents such that all agents are served eventually, i.e., no agent will leave without being served. We consider two different models of the *MSFLP-TA* based on whether there is a moving cost of the facility between adjacent stages. As agents appear at different positions across stages, we also study the variants where agents should be served according to the natural *first-come-first-serve (FCFS) ordering*, i.e., an agent can be served after all agents arriving in earlier stages are served.

Many real-world scenarios can be modeled using the *MSFLP-TA*. For instance, consider the motivating example for the no-moving-cost (**NMC**) model where we have a residential street and the government aims to find a place each stage (e.g., few hours or days) to set up a COVID-19 testing site for residents who are at high risk of being infected (Peto 2020). Naturally, residents want to visit the testing site at different time from different locations. Because the government wants to test all the individuals timely (i.e., within some period or tolerance rate) without overcrowding the testing site, it needs to decide when/where the agents should visit the testing site given the resident preference while optimizing some social objectives, e.g., minimizing the social cost and maximum cost, associated with the travel distances of individuals to the testing site across the stages. In this scenario, the government may choose to serve the individuals in an FCFS fashion to ensure some form of fairness.

Alternatively, consider another motivating example for the with-moving-cost (**WMC**) model (De Keijzer and Wojtczak 2018) where some customers on a line are trying to buy snacks from a vendor. In their scenario, the vendor and the same set of customers move around in each stage and the vendor needs to serve all customers at every stage. Differently, in our model, the vendor can move to a different location at each stage to serve the customers, whereas customers stay at the same place for a certain period (i.e., tolerance rate) upon arrival and different customers dynamically enter the system. The goal for the vendor is to serve all the customers once (possibly in FCFS ordering) before they leave and minimize some social objectives.

We study each model and its FCFS variant from the algorithmic and mechanism design perspectives. From the algorithmic perspective, all of the information about agents are publicly known. We are interested in minimizing the social cost and the maximum cost, which are the total travel cost (plus the facility's moving costs in the **WMC** model) and the maximum cost of the agents to the facilities across stages, respectively. From the mechanism design perspective, agents' locations and arrival stages are private. Our goal is to design group strategy-proof mechanisms such that there is no joint deviation of agents to jointly misreport their private information such that all of them can gain.

*Corresponding author

Result Summary	NMC[1]		WMC[2]
	Social	Max.	Social
Complexity without FCFS	XP	XP	XP
Complexity with FCFS	P	P	P
Group SP[3] Mechanism	FC	FC	GS
Approx. ***(w/, w/o FCFS)***	n	2	$\frac{n}{2} + O(T)$

[1]NMC = No Moving Cost, [2]WMC = With Moving Cost
[3]SP = Strategy-Proof

Table 1: Algorithmic & Mechanism Results for Different Models.

Contribution. Our main modeling and technical contributions are summarized as follows in Table 1.

We introduce the **NMC** model. From the algorithmic perspective, we give slice-wise polynomial (XP) algorithms for the optimization problems of minimizing the social cost and maximum cost objectives, where the tolerance rate and the total number of stages are regarded as fixed parameters. Surprisingly, we show that under the FCFS ordering, the optimization problems can be solved in polynomial time. From the mechanism design perspective, we introduce the **Full-Coverage (FC)** randomized group strategy-proof mechanism, which is similar to the *Equal-Cost* mechanism proposed in (Fotakis and Tzamos 2013) that guarantees strategy-proofness in the k-facility location problems. The first difference between the two mechanisms is the computation of the *SR profile*, which requires us to develop new algorithms for computing the *SR profile* (see Section 3). Moreover, in the *Equal-Cost* mechanism, it requires all output facility locations to share one random variable to guarantee strategy-proofness. This is no longer required in the **FC** mechanism and the facility location at each stage can be independent from each other since the mechanism outputs the assignment of agents to facilities such that agents can only be served by the designated facilities.

We introduce the **WMC** model, where the total travel distance of the facility is added to the social cost. From the algorithmic perspective, we give an XP algorithm for the general setting without FCFS ordering, while under the FCFS ordering, it can be solved in polynomial time. From the mechanism design perspective, we consider the **Global-Swinging (GS)** group strategy-proof mechanism which works differently depending on whether the FCFS ordering is enforced. Specifically, **GS** works as a deterministic mechanism when there is no FCFS constraint and a randomized mechanism otherwise (see Section 4.2). As the maximum cost for the moving cost is not well-defined across stages, we do not consider this objective in the **WMC** model.

1.1 Related Work

Algorithmic Perspective. The simplest problem of locating a single facility on a line to minimize social cost and maximum cost can be solved in polynomial time (Konforty and Tamir 1997). For locating k facilities, the problem of minimizing social cost or maximum cost is already NP-hard in the 2-dimensional setting (Megiddo and Supowit 1984). However, it can be solved in polynomial time in the one-dimensional setting (Love 1976; Tamir 1994; Megiddo, Zemel, and Hakimi 1983). Notice that locating k facilities on a line is a special case for our **NMC** model when the agents' tolerance rate is set to k and they all arrive at the initial stage. Our models are more general and can be harder to analyze as they are dynamic and highly stage-dependent. In particular, our models have multiple stages where new agents arrive at each stage and agents that arrive earlier may or may not stay in the system depending on whether the facility has served them or not. De Keijzer and Wojtczak (2018) studied the dynamic facility relocation problem where a facility can be moved with a cost in each stage to serve the same set of agents whose locations change over time. Their goal is to minimize the total travel distance of agents and the facility. Later, Fotakis et al. (2021) generalized the problem to relocate k facilities where the optimal solution can be computed in polynomial time. In their settings, the facility serves all agents at each stage, and the served agents stay in the system with different locations in each stage. In contrast, our models assume the agents will leave the system after they get the service, though they can wait for it at the same location for a period of time (given by the tolerance rate), and the facility is relocated at each stage to serve some agents such that all of them get the service before leaving.

Mechanism Design Perspective. Procaccia and Tennenholtz (2013) initiated the approximate mechanism design without money paradigm using a simple facility location problem, which inspires works related to our models. Wada et al. (2018) studied the dynamic facility location problem from the mechanism design perspective where the same set of agents, though at the same position, can choose to participate in the game or not at different stages. However, in our model, each agent is served only once in the duration of his stay, and different agents dynamically enter and leave the system over time. Fotakis and Tzamos (2014) showed that there is no deterministic anonymous strategy-proof mechanisms with a bounded approximation ratio for the k-facility location problem on the line. Later, they proposed a randomized group strategy-proof mechanism to overcome the impossibility in the k-facility location problems with approximation ratio $O(n)$ for social cost and 2 for maximum cost (Fotakis and Tzamos 2013). Interestingly, our **WMC** model, which is a generalization of k-facility location problems, shows that incorporating the moving cost circumvents the inapproximability result for deterministic mechanisms. De Keijzer and Wojtczak (2018) also proposed a strategy-proof mechanism for k-facility relocation problem with social cost objective. A recent survey on facility location problems can be found in (Chan et al. 2021).

2 Preliminaries

For simplicity, $\forall a \in \mathbb{N}$, we use $[a]$ to denote the set $\{1, \ldots, a\}$ and all sets are regarded as multisets in this paper. An *instance* or input of the *Multi-Stage Facility Location Problems with Transient Agents (MSFLP-TA)* is a tuple

$\mathcal{I} = (T, r, X, N)$. The element $T \in \mathbb{N}$ is the total number of stages during which the agents might arrive. The element $r \in \mathbb{N}$ is the tolerance rate of the agents, indicating the maximum number of stages the agents are willing to stay. The element $N = (N_1, \dots, N_T) \in \mathbb{N}^T$ is a vector containing the number of agents arriving at each stage. We use $n = \sum_{i=1}^{T} N_i$ to denote the total number of agents and $n_{max} = \max(N)$ to denote the maximum number of agents among all N_t. The element $X = (X^1, \dots, X^T)$ is the *agent profile* where $X^t \in \mathbb{R}^{N_t}$ is the location of agents arriving at stage $t \in [T]$, referred to as *type t agents*. Without loss of generality, we assume $\forall t$, X^t is an ordered set, which indicates that $X_i^t \leq X_j^t$, i.e., the i^{th} agent is on the left of the j^{th} agent in X^t, if $i < j$. We use tuple $x \in \mathbb{R}^n$ to represent all agents' locations and similarly where x is ordered. The *output* is a tuple $\mathcal{O} = (y, p)$. The *facility profile* $y = (y_0, \dots, y_{T+r-1}) \in \mathbb{R}^{T+r}$ is the locations to place the facilities, where y_0 is the initial position of the facility. We call $p = (p^1, \dots, p^T)$ *service profile* where $p^t \in \mathbb{N}^{N_t}$ is a vector recording the stage where the agents in X^t are served. We say a model is *with FCFS ordering* if $\forall t < t'$, $\forall i, j$, we have $p_i^t \leq p_j^{t'}$. For all models in this paper, we call a *service profile valid* when every agent is served after its arrival and before its departure, i.e., if $\forall t \in T, i \in N_t$, $t \leq p_i^t \leq t + r - 1$.

In the **NMC** model, the *social cost* function is defined as

$$SC_{\text{NMC}}(y, p, X) = \sum_{t=1}^{T} \sum_{i=1}^{N_t} |y_{p_i^t} - X_i^t|$$

Furthermore, the *maximum cost* function is given by

$$MC_{\text{NMC}}(y, p, X) = \max_{t \in [T], i \in N^t} |y_{p_i^t} - X_i^t|$$

On the other hand, in the **WMC** model, we define the *social cost* function as

$$SC_{\text{WMC}}(y, p, X) = \sum_{t=1}^{T} \sum_{i=1}^{N_t} |y_{p_i^t} - X_i^t| + \sum_{t=1}^{T+r-1} |y_t - y_{t-1}|$$

A *mechanism* F is a function mapping an *instance* of the problem to an *output*. We allow agents to have private information about their locations and arrival stages. For the convenience of defining *group strategy-proofness*, we further define X^* and p^* as one-to-one mappings from X and p, respectively. Specifically, $X_i^* \in \mathbb{R} \times [T]$ is a tuple containing the i^{th} agent's information about the location and the arrival stage and p_i^* is the stage that i^{th} agent is served, where $i \in [n]$. In other words, X^* includes all agents' location and arrival stage information and p^* records in which stage they are served. Based on that, we define the cost of agent $i \in [n]$ as $c_i(y, p^*, X^*) = |y_{p_i^*} - X_i^*|$. Therefore, we also have $F(T, r, X^*, N) = (y, p^*)$. A mechanism F is called *group strategy-proof* if for any coalition $s \in [n]$, there is no joint deviation of location and arrival stage by agents in s, such that they all gain, i.e., $\forall s, \widetilde{X_s^*} \neq X_s^*$, $\exists i \in s$, $c_i(F(T, r, (\widetilde{X_s^*}, X_{-s}^*), N), X^*) \leq c_i(F(T, r, X^*, n), X^*)$. The *approximation ratio* is the canonical metric for assessing a *group strategy-proof mechanism*'s performance (expected performance when the mechanism is randomized). Over all potential inputs, it compares the performance of the *group strategy-proof mechanism* to that of an optimal solution, which might not induce *group strategy-proofness*.

We further introduce some auxiliary functions which will be used in later sections. We define $med(W)$ as the median point of agents in W. Based on that, function $g(W) = \sum_{w \in W} |w - med(W)|$ is the cost of serving agents in W when we put the facility at $med(W)$. Particularly, $g(W) = 0$ when $W = \emptyset$.

Due to the space constraint, some proofs are either replaced by proof sketches or omitted.

3 No Moving Cost Model

3.1 Algorithm for without FCFS Ordering

We first discuss the most general model, i.e., the one without moving cost and without *FCFS ordering* (**NMC-NFCFS**). In terms of social cost (resp. maximum cost), we give an XP algorithm for fixed T and r (resp. r).

Social Cost. The hardness of the problem comes from the inter-dependency of agents in different stages. Namely, the facility is allowed to choose arbitrary agents at any stage provided that the *service profile* is *valid*. As we show next, there exists a *consecutive* property in this model which allows the optimal social cost to be computed in polynomial time for fixed T and r. We say an agent set W is *consecutive* on t if $W = \{X_i^t, X_{i+1}^t, \dots, X_j^t\}$ for some i and $j \in [r]$. We call S a *splitting of* X, obtained by splitting each X^t for all $t \in [T]$ into r agent sets $(S_{t,1}, \dots, S_{t,r})$, where $\forall i \in [r]$, $S_{t,i} = \bigcup_{j \in [N^t], p_j^t = t+i-1} (X_j^t)$ and may be empty. The following lemma states that all *type t* agents served at the same stage are *consecutive* on t.

Lemma 1. *$\forall t \in [T], i \in [r]$, $S_{t,i}$ is* consecutive *on t.*

Proof. Assume the optimal *facility profile* is y^*. We first show that $\forall t \in [T], i \in [r]$, agents in $S_{t,i}$'s nearest facility is y_{t+i-1}^* among all y_{t+j-1}^* where $j \in [r]$, in the optimal *service profile* p. Suppose there exists some agent $X_k^t \in S_{t,i}$, such that $|y_{t+i-1}^* - X_k^t| > |y_{t+j-1}^* - X_k^t|$, where $i \neq j$. If we let y_{t+j-1}^* serve X_k^t, the social cost will decrease, which contradicts the optimality where $p_k^t = t + i - 1$. Furthermore, suppose for contradiction there is some S_i^t that is not consecutive on t. It suggests that there exists an agent in X^t, that is between the leftmost and rightmost agents of S_i^t but served by another facility other than y_{t+i-1}^*. However, its nearest facility should be the same as either the leftmost or the rightmost agent in S_i^t, which are both y_{t+i-1}^*, a contradiction. □

We enumerate all the splittings, and the splitting in an optimal *output* must be one of them by Lemma 1. For each splitting, we use dynamic programming to find the optimal assignment of each agent set $S_{t,i}$ to some stage $t + j - 1$, where $i, j \in [r]$. The facility is put at the median point of all agents assigned to it in each stage. We call $M \in (0, 1)^{T \times r}$ agent state such that $M_{t',i}$ denotes whether the set $S_{t',i}$ is served or not. Let $OPT_S^{\text{NMC}}(t, M)$ be the minimum social cost for the **NMC-NFCFS** model to serve the remaining agents indicated by M starting from stage t to $T+r-1$ given a fixed S. $OPT_S^{\mathrm{NMC}}(T+r, M)$ is therefore 0 as there is no agent left, and for $t < T+r$, we have $OPT_S^{\mathrm{NMC}}(t, M) =$

$$\min_{\substack{M' \in (0,1)^{T\times r}\\ M' \leq M}} \left\{ \; OPT_S^{\mathrm{NMC}}(t+1, M') + g_{t,r}(M, M') \; \right\}$$

where M' is changed from M by changing one entry for each row among $\{\max\{1, t-r+1\}, \dots, t\}$ from 0 to 1 and $g_{t,r}(M, M') = g(\bigcup_{\substack{t' \in \{\max\{1,t-r+1\},\dots,t\}\\ i \mid (M'-M)_{t',i}=1}} S_{t',i})$, i.e., it serves one agent in each stage from $\max\{1, t-r+1\}$ to t given by the difference between M and M'. The subscript $M' \leq M$ indicates that $\forall t', i, M'_{t',i} \leq M_{t,i}$.

Theorem 1. $\min_{S \in \{\text{splittings of } X\}} \{OPT_S^{\mathrm{NMC}}(1, \{0\}^{T\times r})\}$ *computes the optimal social cost for the* ***NMC-NFCFS*** *model in* $O(T \times r^{r+1} \times (2(n_{max} + r))^{T\times r})$ *using dynamic programming.*

Maximum Cost. Our algorithm is inspired by the idea in *Equal-Cost* mechanism given by Fotakis and Tzamos (2013), which studies the group strategy-proof mechanism for k-facility location problems. We define the *serving range (SR)* of the facility at some stage such that the facility serves all existing agents within that range. From this perspective, if we place an *SR* at each stage and these *SRs* cover all agents at least once within their duration of stay, the corresponding *service profile* is *valid*.

Lemma 2. *The minimum length l for* SRs*, which guarantees a* valid service profile *must be a distance between two agents* X_i^{t1} *and* X_j^{t2}*, where* $t1, t2 \in [T], i \in [N_{t1}]$ *and* $j \in [N_{t2}]$.

Lemma 2 shows that the number of possible lengths of l is bounded by $O(n^2)$, through which we reduce an optimization problem to a feasibility problem. We may also observe a simple yet useful property which serves as a building block of the algorithm. Namely, the *SRs* with the minimum length l can be placed in such a way that their leftmost points are located exactly at some agents' locations. This property is true because we can always move such *SR* to the right until the left boundary touches the leftmost agent it covers, without serving fewer agents.

We denote $C = (L, l)$ as an *SR profile*. *SR anchor* is defined as $L = (L_1, \dots, L_{T+r-1}) \in n^{T+r}$, where the leftmost point of the *SR* in stage t is at x_{L_t} and l is the length of those *SRs*. From the above analysis, the optimal *output* $\mathcal{O}$ can be easily obtained from *SR profile* C, as the facility is located at the middle point of each *SR* and each agent is served by the first facility that covers him. If an *SR profile* is obtained, the corresponding *facility profile* is given by $y = \{x_{L_1} + \frac{l}{2}, \dots, x_{L_{T+r-1}} + \frac{l}{2}\}$. In terms of the *service profile* p, p_i^t is given by the minimum $j \in \{t, \dots, t+r-1\}$ such that $X_i^t \leq x_{L_j} + l$ and $X_i^t \geq x_{L_j}$. We call an *SR profile* feasible if and only if its corresponding *service profile* is *valid*. FEASIBLE1 (Algorithm 1) computes whether a feasible *SR profile* starting from stage t can be obtained. Notice that when this function is called, we assume that agents leaving before stage t are already served. Moreover, the parameter a_i is set to -1 when $i \leq 0$, representing that there is no facility placed at stage i.

Algorithm 1: FEASIBLE1$(a_{t-(r-1)}, \dots, a_{t-1}, t)$

Parameter: *SR* length l

for $1 \leq i \leq r-1$ **do**
  if $a_{t-i} \neq -1$ **then**
    mark agents with locations among $[x_{a_{t-i}}, x_{a_{t-i}}+l]$
  end if
end for
$res \leftarrow$ FALSE, $k \leftarrow 1$
while $res =$ FALSE and $1 \leq k \leq n$ **do**
  mark agents whose locations are within $[x_k, x_k + l]$
  if all agents leaving at stage t are marked **then**
    if $t = T+r-1$ **then**
      return TRUE
    else
      $res \leftarrow$ FEASIBLE1 $(a_{t-r+2}, \dots, k, t+1)$
    end if
  end if
  $k \leftarrow k+1$
end while
return res

Theorem 2. *The optimal maximum cost for the* ***NMC-NFCFS*** *model can be computed by choosing the minimum parameter l among $O(n^2)$ possible agents' distance*[1] *such that Algorithm 1 returns* TRUE *in* $O(T \times n^{r+2} \times (n+r))$ *using dynamic programming.*

3.2 Algorithm for with FCFS Ordering

In this subsection, we discuss the **NMC** model *with FCFS ordering* (**NMC-WFCFS**). We show that by enforcing this reasonable assumption of serving order based on arrival time, we can solve the optimization problem for both social cost and maximum cost in polynomial time.

Social Cost. *The FCFS ordering* restricts *type* t agents to be served after all *type* $(t-1)$ agents (if any). This implies that if agents served at stage t contains those of different $types$, these $type$ numbers must be contiguous. Moreover, suppose these types are $\{i, i+1, \dots, j\}$, then all agents in $\{X^{i+1}, \dots, X^{j-1}\}$ plus part of (or possibly all) agents in X^i and X^j are served at t. This significantly reduces the constitution of *types* of agents served at some stage t compared with that in the **NMC-NFCFS** model, where some other *types* of agents might be served at t. Inspired by the above observation on the *types* of agents served in each stage, we partition each X^i into three subsets, namely, *previous batch*, *individual batch* and *next batch*. With the *FCFS ordering*, agents are in *previous batch* (resp. *next batch*) of X^i if they are served together with some agents in X^j, where $i > j$ (resp. $i < j$). An agent is in the *individual batch* of X^i if all agents served together with him are in X^i.

Similarly, we exploit the consecutive property. Informally, for *type* i agents, its *previous batch*, *next batch*, and all agent sets that are the subsets of *individual batch* and served at the same stage, are *consecutive on* i. Notice that this property

[1]For Theorem 2 and Theorem 4, doing binary search for parameter l over $O(n^2)$ possibilities can reduce the complexity.

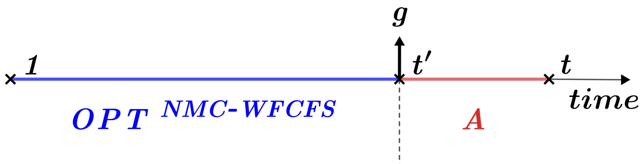

Figure 1: Decomposition of $OPT^{\text{NMC-WFCFS}}(t,e,i,j)$

follows directly from Lemma 1. We first introduce an auxiliary algorithm $A(W,i,j)$, which returns the optimal cost of serving agents in an agent set W ordered from left to right (a subset of agents in X^t for some $t \in [T]$), where $\{W_1,\dots,W_i\}$ need to be served in remaining j stages.

$$A(W,i,j) = \min_{1\le i'\le i}\left\{ g\left(\bigcup_{k=i'}^{i} W_k\right) + A(W,i'-1,j-1)\right\}$$

Since all agents must be served and the algorithm terminates when there is no agent to serve, the initial states are $A(W,i,1) = g\left(\bigcup_{k=1}^{i} W_k\right)$ and $A(W,0,j) = 0$, respectively.

Next, we define $OPT^{\text{NMC-WFCFS}}(t,e,i,j)$ to be the minimum social cost for the **NMC-WFCFS** model to serve agents of *types* $[e]$ except for those in $\bigcup_{k=i}^{j} X_k^e$ during the first t stages. Based on that, we have $OPT^{\text{NMC-WFCFS}}(t,1,i,j) = A(X^1 \setminus \bigcup_{k=i}^{j} X_k^1, N_1-(j-i-1), t)$ if only *type* 1 agents need to be served. Meanwhile, $OPT^{\text{NMC-WFCFS}}(1,e,i,j) = g\left(\bigcup_{k=1}^{e} X^k \setminus \bigcup_{k=i}^{j} X_k^e\right)$ if all remaining agents need to be served at stage 1. Finally, for other cases, we have

$$OPT^{\text{NMC-WFCFS}}(t,e,i,j) = \min_{\substack{0<i'\le j'+1\\ 0<e'<e\\ 0<a\le b+1\\ [a,b]\cap[i,j]=\emptyset\\ I_2<t'\le t-I_1}} \left\{ \begin{array}{l} OPT^{\text{NMC-WFCFS}}(t'-I_2,e',i',j') \\ +g\left(\bigcup_{k=e'+1}^{e-1} X^k \cup \bigcup_{k=i'}^{j'} X_k^{e'} \cup \bigcup_{k=a}^{b} X_k^e\right) \\ +A\left(\begin{array}{l} X^e \setminus \left\{\bigcup_{k=a}^{b} X_k^e \cup \bigcup_{k=i}^{j} X_k^e\right\}, \\ N_e-(j-i+1)-(b-a+1), \\ t-t' \end{array}\right) \end{array} \right\}$$

where $i' = j'+1, a = b+1$ suggests that the corresponding union of agents is empty. With a slight abuse of notation, $I_z \in \{0,1\}$ indicates whether $X^e \setminus \{\bigcup_{k=a}^{b} X_k^e \cup \bigcup_{k=i}^{j} X_k^e\}$, $\{\bigcup_{k=e'+1}^{e-1} X^k \cup \bigcup_{k=i'}^{j'} X_k^{e'} \cup \bigcup_{k=a}^{b} X_k^e\}$, $\bigcup_{k=i'}^{j'} X_k^{e'}$ and $\bigcup_{k=a}^{b} X_k^e$ are empty or not for $z = 1,2,3,4$, respectively. In each recursion, the selection of e' and t' should also satisfy the following restriction: $\max\{e-1+I_4, e'+I_2\} \le t' \le e'+r-1+I_2-I_3$, which guarantees all agents are served before they leave and after their arrival. See Figure 1 for an illustration of the dynamic programming.

Theorem 3. $OPT^{\text{NMC-WFCFS}}(T+r-1,T,0,0)$ *computes the optimal social cost for the* ***NMC-WFCFS*** *model in* $O((T+r)^2 \times T^2 \times n_{max}^6 + T \times n_{max}^7 \times r)$ *using dynamic programming.*

Algorithm 2: FEASIBLE2(a,t,s)

Parameter: *SR* length l

```
if t = T + 1 then
    return TRUE
end if
if s − 1 ≥ t then
    mark agents in X^t that are covered by [x_a, x_a + l]
end if
if all agents in X^t are marked then
    res ← FEASIBLE2(a, t + 1, s)
else
    res ← FALSE
    i ← 1
    while res = FALSE and 1 ≤ i ≤ n do
        num ← number of stages needed to locate SRs
        so as to mark those unmarked agents in X^t whose
        locations are outside [x_i, x_i + l].
        if s + num ≤ t + r − 1 then
            res ← FEASIBLE2 (i, t + 1, s + num + 1)
        end if
        i ← i + 1
    end while
    return res
end if
```

Maximum Cost. *FCFS ordering* significantly reduces the number of states in the dynamic programming for maximum cost. Moreover, with the *FCFS ordering*, we modify the transition from *SR profile* C to *service profile* p. p_i^t is now the minimum j such that $x_{L_j} \le X_i^t \le x_{L_j}+l$, and $\forall k < t$, agents in X^k are served.

Algorithm 2 decides whether we can obtain a feasible *SR profile* starting from stage s, given that agents of *type* $[t-1]$ are served and a is the *SR anchor* at stage $s-1$, i.e., the last *SR* whose corresponding facility serves agents of *type* $t-1$.

Notice that under the *FCFS ordering*, among all *SRs* which have some agents in X^t served in the corresponding *service profile*, only one of them can also serve agents of *type(s)* greater than t. This indicates that when the algorithm deals with agents in X^t, it only needs the *SR anchor* a of the last *SR* covering agents in X^{t-1}. However, the *SR anchor* a can cover agents in X^t only if $s-1 \ge t$, i.e., agents can only be served after their arrival. When $s-1 \ge t$, in case $[x_a, x_a+l]$ cannot cover all agents in X^t, it also needs to track the *SR anchor* i of the last *SR* covering agents in X^t, and pass it as an argument to FEASIBLE2 $(i, t+1, s+num+1)$ which is called recursively. Once i is fixed, num, the number of intermediate *SRs* needed to cover remaining agents in X^t whose locations are outside $[x_a, x_a+l] \cup [x_i, x_i+l]$, can be greedily computed by assigning corresponding *SR anchors* from left to right. When $s-1 < t$, it simply ignores the *SR anchor* a.

Theorem 4. *Algorithm 2 computes the optimal maximum cost for the* ***NMC-WFCFS*** *model in* $O(n^4 \times n_{max} \times T \times r)$ *using dynamic programming.*

3.3 Strategy-Proof Mechanism Design

In this subsection, we consider the **Full-Coverage (FC)** randomized mechanism which can be used for our model with or without *FCFS ordering*. This mechanism extends the result of *Equal-Cost* mechanism proposed by (Fotakis and Tzamos 2013). They consider the setting of serving agents using k facilities in one stage with any concave cost function. Their setting can be regarded as a degeneration of our model where all agents arrive at the first stage with *tolerance rate* $r = k$. Interestingly, as discussed in Section 1, a similar mechanism can be used in our model, and it satisfies our desired properties.

Mechanism 1. *The mechanism has* 3 *steps.*

1. *Compute the feasible* SR profile $C = (L, l)$ *with minimum* l*.*
2. *For each stage* t*, place the facility at* x_{L_t} *with probability* $1/2$ *and* $x_{L_t} + l$ *with probability* $1/2$*.*
3. *Output* $\mathcal{O} = (y^{random}, p)$ *is converted from* C*, where* y^{random} *contains the probability distribution of the facility location at each stage.*

Notice that in step 1 of Mechanism 1, the resulting *SR profile* depends on whether the *FCFS ordering* is enforced. Therefore, we can use different optimal algorithms to compute the maximum cost.

To show that Mechanism 1 is group strategy-proof, we need to prove the following lemma which follows from (Fotakis and Tzamos 2013). It states that there exists a probability distribution of facility locations such that all agents covered by *SRs* have the same expected cost and cannot benefit from misreporting.

Lemma 3. *All the agents within* SR $[start, end]$ *have the same expected cost, specifically* $\frac{end-start}{2}$*, if we place the facility at* $start$ *with probability* $1/2$ *and at* end *with probability* $1/2$*. Those agents who are outside the* SR *have their expected cost equal to the distance to the center of the* SR.

Theorem 5. *Mechanism 1 is group strategy-proof for the* ***NMC*** *model.*

Proof Sketch. A part of the proof is similar to the one for *Equal-Cost* mechanism proposed in (Fotakis and Tzamos 2013). However, we need to show that the mechanism remains strategy-proof where the agents' locations and arrival stages are both private in our setting. The key idea is that the misreporting agent is either served outside of or within the agent's true duration of stay after misreporting. In the first case, the agent will not misreport since it prefers to be served with any cost than get no service at all. In the second case, they can only gain by making the length of *SRs* smaller by Lemma 3. However, if all misreporting agents are served during their true duration of stay, it indicates that the resulting *SR profile* must be a subset of all *SR profile* considered by the mechanism when agents report truthfully. Therefore, they can only make the length of *SRs* larger by step 1.

Notice that if the *SR profile* computed from step 1 in Mechanism 1 has $l = 0$, the mechanism returns the optimal solution and guarantees group strategy-proofness at the same time. The following theorem gives the upper bound for **FC**'s performance in terms of social cost and maximum cost when $l > 0$.

Theorem 6. *When the* SR profile *has* $l > 0$*, Mechanism 1 achieves a tight approximation ratio of* 2 *for maximum cost and* n *for social cost in the* ***NMC*** *model.*

4 With Moving Cost Model

We investigate in this section the model with moving cost (**WMC**), i.e., the moving cost of the facility between adjacent stages. The target objective function is the social cost as discussed before.

Our algorithm relies on the following lemma, which extends the result of De Keijzer and Wojtczak (2018) to our model. Its key idea is that it suffices to consider the facility profile y where in each stage, its location is either one of the agents' locations or the initial facility's location.

Lemma 4. *The optimal* facility profile y *satisfies the property that for all* $t \in [T + r - 1]$, $y_t \in \{y_0\} \cup x$*.*

Since we incorporate the moving cost in $SC_{\text{WMC}}(y, p)$, putting the facility at the median of agents to be served at some stage, i.e., using function g, is no longer optimal. In the following two subsections, we will modify the algorithms in Section 3 to accommodate the introduced moving cost using Lemma 4. The key point is to explicitly examine all possible $O(n)$ optimal placement of the facility, i.e., $x \cup \{y_0\}$ in each stage instead of applying the median rule when the agents to be served at some stage are determined.

4.1 Algorithm

In this subsection, we first study the **WMC** model without *FCFS ordering*. $OPT_S^{\text{WMC}}(t, M, P)$ is defined by adding a new parameter in $OPT_S^{\text{NMC}}(t, M)$, where $P \in \mathbb{N}$ adds an additional constraint that we place the facility at the position of the P^{th} element in $x \cup \{y_0\}$ at stage t. By Lemma 4, once we find the optimal *service profile*, it must consist of locations in $x \cup \{y_0\}$. The recurrence works almost the same except that we examine all possible $P \in [|x \cup \{y_0\}|]$ after determining which agents to serve at t. This parameter enlarges the size of the dynamic programming table by a factor of $O(n)$, and it increases the running time by $O(n)$ during the computation for each state.

Corollary 1. *There exists an algorithm that computes the optimal social cost for the* ***WMC-NFCFS*** *model in* $O(T \times r^{r+1} \times (2(n_{max} + r))^{T \times r} \times n^2)$*.*

We then investigate the **WMC** model with the *FCFS ordering*. Similarly, we obtain new algorithms $B(x, i, j, P)$ and $OPT^{\text{WMC-WFCFS}}(t, e, i, j, P)$ by adding a new dimension P to record the placement of the facility in algorithms $A(x, i, j)$ and $OPT^{\text{NMC-WFCFS}}(t, e, i, j)$. Utilizing B and $OPT^{\text{WMC-WFCFS}}$, we obtain the following corollary.

Corollary 2. *There exists an algorithm that computes the optimal social cost for the* ***WMC-WFCFS*** *model in* $O(((T + r)^2 \times T^2 \times n_{max}^6 + T \times n_{max}^7 \times r) \times n^2)$*.*

4.2 Strategy-Proof Mechanism Design

We consider the **Global-Swinging (GS)** group strategy-proof mechanism for the **WMC** model in this subsection. We denote the leftmost (resp. the rightmost) agent among all stages as $lt(x)$ (resp. $rt(x)$).

Mechanism 2. *The mechanism works differently in terms of the* ***WMC-NFCFS*** *model and the* ***WMC-WFCFS*** *model. It outputs* $\mathcal{O} = (y, p)$ *given* $\mathcal{I} = (T, r, X, N)$.

For WMC-NFCFS: *Deterministic Mechanism*

- *When* $r = 1$, $\forall t \in [T]$, $y_t = med(X^t)$ *and* $\forall t \in [T], i \in [N_t], p_i^t = t$.
- *When* $r > 1$, $\forall t \in [T+r-1]$, $y_t = lt(x)$ *if* t *is odd and* $y_t = rt(x)$ *if* t *is even.* $\forall i \in [T], j \in [N_i]$, *if* i *is odd and* $X_j^i \in [lt(x), (lt(x) + rt(x))/2]$ *or* i *is even and* $X_j^i \in [(lt(x) + rt(x))/2, rt(x)]$, $p_j^i = i$. *Otherwise,* $p_j^i = i + 1$.

For WMC-WFCFS: *Randomized Mechanism*
$\forall t \in [T]$, $y_t = lt(x)$ *with probability* $1/2$ *and* $y_t = rt(x)$ *with probability* $1/2$. $\forall t \in [T], i \in [N_t]$, $p_i^t = t$.

The idea of the mechanism for **WMC-NFCFS** is that if $r = 1$, all agents are immediately served upon arrival by placing the facility at the median point of them. Notice that when $r = 1$, the model degenerates to the multi-stage dynamic facility location problem considered by De Keijzer and Wojtczak (2018), except that in their model, the number of agents in each stage is the same. Similarly, they also place the facility at the median of the arrived agents in each stage. When $r > 1$, the facility is put at the global leftmost agent's location and rightmost agent's location, i.e. $lt(x)$ and $rt(x)$ in a round-robin fashion and for each stage, it serves all existing agents with locations between $(lt(x) + rt(x))/2$ and itself. However, when $r > 1$, as agents of different *types* can be served at the same stage even though not all smaller *type* agents have been served, the deterministic version fails to obey the FCFS ordering in the **WMC-WFCFS** model. A randomized version is therefore proposed to fulfill the requirement by serving all agents immediately. Interestingly, both versions of **GS** have the same approximation ratio.

Theorem 7. *Mechanism 2 is group strategy-proof for the* ***WMC*** *model.*

Fotakis and Tzamos (2014) showed that there is no anonymous deterministic strategy-proof mechanism for the k-facility location problem, which is a degeneration of our **NMC** model. Their intuition of the unbounded approximation ratio comes from the fact that a strategy-proof mechanism needs to place a facility far from an agent which incurs a large cost, while the optimal solution has a negligible cost. However, in the **WMC** model, the following lemma shows the incorporated moving cost also needs to be considered in the optimal solution, making the cost of the optimal solution no longer negligible.

Lemma 5. *The social cost for any mechanism to serve all the agents (regardless of enforcing* FCFS ordering *or not) is at least* $rt(x) - lt(x) + \min(|y_0 - lt(x)|, |y_0 - rt(x)|)$.

Proof Sketch. It suffices to consider the cost of serving $lt(x)$ and $rt(x)$ plus the moving cost of the facility. The sum of these two costs must be at least $rt(x) - lt(x) + \min(|y_0 - lt(x)|, |y_0 - rt(x)|)$ by applying the triangle inequality.

Theorem 8. *Mechanism 2 achieves an approximation ratio of* $\frac{n}{2} + O(T)$ *for social cost in the* ***WMC*** *model.*

Proof. We first consider the mechanism for the **WMC-NFCFS** model. When $r = 1$, denote the optimal cost as OPT and the cost returned by the mechanism as $COST = C_{agent} + C_{mov}$, where C_{agent} is the cost of serving the agents and C_{mov} is the cost induced by moving the facility. Since we put the facility at the median of agents in each stage, C_{agent} is optimal if the moving cost is not considered. Hence, $C_{agent} \leq OPT$. On the other hand, the moving cost of the facility, i.e., C_{mov} is at most $T \times (rt(x) - lt(x)) + dis \leq (T + 1) \times OPT$, where $dis = \min(|y_0 - lt(x)|, |y_0 - rt(x)|)$. Therefore, the approximation ratio for deterministic version of Mechanism 2 when $r = 1$ is $\frac{COST}{OPT} = \frac{C_{agent}+C_{mov}}{OPT} \leq \frac{(T+2)\times OPT}{OPT} = T + 2 = O(T)$. When $r > 1$, the cost of serving any agent is at most $(rt(x) - lt(x))/2$ and the moving cost is bounded by $T \times (rt(x) - lt(x)) + (rt(x) - lt(x) + dis)$. Therefore, the social cost is upper bounded by $(\frac{n}{2} + T + 1) \times (rt(x) - lt(x)) + dis$. Moreover, Lemma 5 provides a lower-bound $rt(x) - lt(x) + dis$ for social cost given any mechanism. Correspondingly, the approximation ratio for deterministic version of Mechanism 2 when $r > 1$ is $\frac{(\frac{n}{2}+T+1)\times(rt(x)-lt(x))+dis}{rt(x)-lt(x)+dis} \leq \frac{(\frac{n}{2}+T+1)\times(rt(x)-lt(x))}{rt(x)-lt(x)} = \frac{n}{2} + O(T)$.

We then consider the mechanism for the **WMC-WFCFS** model. By Lemma 3, the expected cost for all agents is $\frac{rt(x)-lt(x)}{2}$. On the other hand, the expected moving cost of the facility is bounded by $O(T) \times (rt(x) - lt(x))$. Therefore, the approximation ratio for randomized version of Mechanism 2 is $\frac{n}{2} + O(T)$. □

5 Conclusion and Discussion

We study the multi-stage facility location problems with transient agents in different settings from both algorithmic and mechanism design perspectives. Collectively, we establish several non-trivial dynamic programming algorithms which require key insights into the properties of the models. Moreover, we consider these models under the private information setting and design group strategy-proof mechanisms with approximation guarantees.

There are many interesting directions for future work. The most important one is whether it is possible to improve the approximation ratio for group strategy-proof mechanisms. On the other hand, the model can be extended where agents' tolerance rates are different. All the algorithmic results and strategy-proof mechanisms in this paper can be modified to accommodate the setting when agents of different types can have different tolerance rates. Nevertheless, the fundamental building blocks of the positive results collapse when the tolerance rates for agents of the same type can be different. Additionally, we conjecture that the optimization problems without FCFS ordering are NP-hard.

Acknowledgments

We gratefully acknowledge funding from Research Grants Council of the Hong Kong Special Administrative Region, China (Project No. UGC/FDS13/E01/20), and by NSFC No.62202100. Hau Chan is supported by the National Institute of General Medical Sciences of the National Institutes of Health [P20GM130461] and the Rural Drug Addiction Research Center at the University of Nebraska-Lincoln. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health or the University of Nebraska.

References

Chan, H.; Filos-Ratsikas, A.; Li, B.; Li, M.; and Wang, C. 2021. Mechanism Design for Facility Location Problem: A Survey. In *The 30th International Joint Conference on Artificial Intelligence (IJCAI 2021)*, 1–17.

De Keijzer, B.; and Wojtczak, D. 2018. Facility Reallocation on the Line. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, 188–194.

Fotakis, D.; Kavouras, L.; Kostopanagiotis, P.; Lazos, P.; Skoulakis, S.; and Zarifis, N. 2021. Reallocating multiple facilities on the line. *Theoretical Computer Science*, 858: 13–34.

Fotakis, D.; and Tzamos, C. 2013. Strategyproof Facility Location for Concave Cost Functions. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, 435–452. ISBN 9781450319621.

Fotakis, D.; and Tzamos, C. 2014. On the Power of Deterministic Mechanisms for Facility Location Games. *ACM Transactions on Economics and Computation*, 2(4).

Konforty, Y.; and Tamir, A. 1997. The single facility location problem with minimum distance constraints. *Location Science*, 5(3): 147–163.

Love, R. F. 1976. One-Dimensional Facility Location-Allocation using Dynamic Programming. *Management science*, 22: 614–617.

Megiddo, N.; and Supowit, K. J. 1984. On the Complexity of Some Common Geometric Location Problems. *SIAM journal on computing*, 13(1): 182–196.

Megiddo, N.; Zemel, E.; and Hakimi, S. L. 1983. The Maximum Coverage Location Problem. *SIAM Journal on Algebraic Discrete Methods*, 4(2): 253–261.

Peto, J. 2020. Covid-19 mass testing facilities could end the epidemic rapidly. *Bmj*, 368.

Procaccia, A. D.; and Tennenholtz, M. 2013. Approximate mechanism design without money. *ACM Transactions on Economics and Computation*, 1(4): 1–26.

Tamir, A. 1994. A distance constrained p-facility location problem on the real line. *Mathematical Programming*, 66(1-3): 201–204.

Wada, Y.; Ono, T.; Todo, T.; and Yokoo, M. 2018. Facility Location with Variable and Dynamic Populations. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 336–344.