Characterization of Truthful Mechanisms for One-Dimensional Single Facility Location Game with Payments

Pinyan Lu[1] and Lan Yu[2,⋆]

[1] Microsoft Research, China
pinyanl@microsoft.com
[2] Division of Mathematical Sciences, School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore
YULA0001@ntu.edu.sg

Abstract. In a one-dimensional single facility location game, each player resides at a point on a straight line (his location); the task is to decide the location of a single public facility on the line. Each player derives a nonnegative cost, which is a monotonically increasing function of the distance between the location of the facility and himself, so he may misreport his location to minimize his cost. It is desirable to design an incentive compatible allocation mechanism, in which no player has an incentive to misreport.

Offering/Charging payments to players is a usual tool for a mechanism to adjust incentives. Our game setting without payment is equivalent to the voting setting where voters have single-peaked preferences. A complete parametric characterization of incentive compatible allocation mechanisms in this setting was given by [17], while the problem for games with payments is left open. We give a characterization for the case of linear and strictly convex cost functions by showing the sufficiency of weak-monotonicity, which, more importantly, implies an interesting monotone triangular structure on every single-player subfunction.

1 Introduction

People live in communities, where public facilities need to be built to serve the residents. The location of a public facility is one of the most important decisions to make since the convenience of accessing a facility is mainly affected by the distance to the facility. Although the social goal is to provide convenient service to the whole community, tradeoffs have to be made as people reside at different locations. Hence the well-known public facility location problem has been a long-lasting attraction to researchers.

Conventionally, the convenience of accessing a facility is quantified by the negation of a cost, indicating the effort needed to reach the facility. The cost can be calculated through a cost function C on the distance d to the facility, which grows with the distance. For different people or circumstances, the growth

⋆ Supported by National Research Foundation (Singapore) under grant RF2009-08.

Y. Chen and N. Immorlica (Eds.): WINE 2013, LNCS 8289, pp. 333–346, 2013.

rate may vary; the cost function may be linear, convex, concave, or a combination of the three. A linear cost function $C(d) = \alpha d$ is the simplest case, which has a stable growth rate α. A convex cost function has higher growth rate for longer distances (such as the taxi fare), which captures the nature that people become exhausted after long distances. In contrast, a concave cost function has lower growth rate for longer distances (such as the subway fare). This represents situations where people become adapted after long distances.

In algorithm design, people solve optimization problems, such as to minimize the total/maximum cost of the people served in the community. There are also interesting variations in the number of facilities, the cost function, and the location space (discrete or continuous). In our paper, we consider one of the simplest variations: one-dimensional single facility location, where the location space is the one-dimensional real line $\mathbb{R}$ and only a single facility needs to be built. All three types of cost functions are investigated.

However, optimization is not the focus of our paper: We investigate the game-theoretic setting where each resident is modeled as a player in the facility location game intending to maximize his *utility*, i.e., his overall benefit. Each player i's true location r_i is his private information, and the location of the facility is chosen based on the reported locations from all players $\mathbf{x} = (x_1, \ldots, x_n)$. Hence a player may have an incentive to misreport his location to make the facility closer to him. Naturally, to achieve good public service, we would like a solution where no player has an incentive to misreport. This property is called *incentive compatibility*, or simply *truthfulness*, which is the main solution concept of the field of *mechanism design* [21,20].

Offering/Charging payments to players is a usual tool in mechanism design to adjust incentives. Our work allows solutions with payments. Thus, an *allocation mechanism*, i.e., a solution to our facility location game, is composed of an *allocation function* and a *payment function* vector. The allocation function f takes the reported locations of all players $\mathbf{x}$ as input and outputs a location y of the public facility; The payment function vector $\mathbf{p}$ contains a payment function p_i as its ith component for each player i. Function p_i takes the same input as f, and assigns a (positive or negative) payment to player i. The setting without payments restricts $\mathbf{p} \equiv \mathbf{0}$.

Under a mechanism $(f, \mathbf{p})$, the utility u_i of player i is his payment under reported locations $p_i(\mathbf{x})$ minus his cost under true location $C(|f(\mathbf{x}) - r_i|)$. The mechanism is public knowledge, and a truthful mechanism ensures truth-telling in the following sense: No matter what other players may report, for each player, given the mechanism and other players' reports, reporting his true location always maximizes his utility.

The goal of our work is to characterize the set of truthful allocation functions, i.e., allocation functions f for which there exists a payment function vector $\mathbf{p}$ such that $(f, \mathbf{p})$ is truthful. Characterization of truthful functions is meaningful in mechanism design, since it allows mechanism designers to focus on the function and not to worry about payments, whose existence is already guaranteed by the characterization. Furthermore, in most applications, there are other desirable properties the allocation function should also satisfy, such as fairness or efficiency. A good characterization provides a useful description of the set of truthful functions for a designer to start with to work further on the other properties, or to prove the impossibility to satisfy other properties simultaneously. A great number of results in mechanism design follow this path [3,10,22].

For the game setting without payment, a complete parametric characterization of truthful allocation mechanisms was given by Moulin [17]. (One unnecessary assumption in the proof is dropped by Barberà and Jackson [5], and Sprumont [26].) Observe that, without payment, a player's utility is simply the negation of his cost, and the definition of truthfulness is only concerned with the comparison of the cost of two locations. Moreover, the cost is *single-peaked*: it reaches its minimum 0 at the player's true location and increases monotonically on both sides; the formula of the cost function becomes irrelevant. In fact, this setting is essentially equivalent to the voting setting where voters have single-peaked preferences. Moulin considered the voting setting, and hence his characterization, a parametric representation of truthful allocations is called a *generalized median voter scheme*. It is an extension of the function that selects the median voter's preference peak, which, interpreted into our setting, is the median location out of the locations of all players.

The characterization for games with payments is left open, which is what we studied in this paper. This question is interesting in its own right: In real life, some facility builders are willing to provide payments. For example, when a company chooses the location of its office, employees living far away from the office are subsidized. On the other hand, the generalized median voter scheme is very restricted, and does not satisfy certain other desirable properties, such as fairness or cost minimization. For example, the average function of all agents' locations minimizes the sum of squares of agents' cost, which is a widely used objective function in operational research to balance the social welfare and fairness. This nice function is not truthfully implementable without payment. However, by our characterization, it can be made truthful with payment, so designers may want to consider investing some money to realize this allocation function.

1.1 Our Contributions

It turns out that the set of truthful allocation functions with payments is a much wider class than the generalized median voter scheme. We show that weak monotonicity, an easily proven necessary condition for truthfulness [20], is also sufficient in this setting for linear and strictly convex cost functions. There has been a series of works on characterization of truthfulness for various settings [23,19,3,15,24,27,18,16,11,16,14,9,7,1,4,2], and most of them involve weak monotonicity, or some other kind of monotonicity properties. It turns out that the characterization results are closely related to the *domain* of the problem setting, i.e., the set of all possible *valuation functions* on the set of outcomes. In our setting, the domain of player i is $\{-C(|y - r_i|) : r_i \in \mathbb{R}\}$, where each element $-C(|y - r_i|)$ is a single-peaked function mapping each location of facility $y \in \mathbb{R}$ to the valuation (convenience) player i derives when his true location is r_i.

Evidently, the domain of our setting is restricted, so Roberts' Theorem that every truthful function in an unrestricted domain is an affine maximizer [23] does not apply. Furthermore, our domain is not a special case of the convex domain or single-parameter domain for which the sufficiency of weak monotonicity is proved [19,3,24]. Clearly, our domain is not convex, and an easy way to see this is that the average of two valuation functions no longer has peak of value 0. On the other hand, though the domain of each player i is associated with a single parameter r_i, the single-peaked function does not conform to the function in the definition of single-parameter domain. It is interesting that none of the previous characterization results covers our setting although it is very simple and realistic. In particular, most of the previous result involves some kind of convexity: either the valuation is convex or the type space is convex. The fact that there are infinite (uncountable) many alternatives also makes the result interesting since most of previous results assume a finite set of alternatives.

On the other hand, from the mechanism design point of view, weak monotonicity, as a condition on any two locations, is not directly applicable; it is more desirable to derive its equivalent properties that describe global features of the allocation function (usually on every single player subfunction), from which truthful payments can also be described. The characterization of the single-parameter domain in [19,3] is successful in this aspect: Various mechanisms for specific settings with different objectives are derived based on this characterization [19,3,12,10,25,6]. For our problem, we also succeed in providing a characterization of this kind. In fact, the sufficiency of weak monotonicity is shown indirectly through the correctness of this characterization.

More specifically, our characterization results are presented in three steps: In Section 3, we derive some properties from weak-monotonicity on every single player subfunction:[1] For strictly convex cost functions, the allocation function is simply monotonically non-decreasing in the usual sense; Linear cost functions imply a weaker condition, which we call *partially monotonically non-decreasing*. As shown in Section 4, this condition implies a *monotone triangular partition*, which graphically divides the allocation function into pieces each within a triangle, and the set of triangles obeys some "monotone" property. For strictly convex cost functions, this part is evident as the allocation function is monotone. Finally we provide a payment function with respect to a monotone triangular partition and prove truthfulness in Section 5 for linear and strictly convex cost functions respectively.

In summary, here are our main characterization results for one-dimensional single facility location game with payments (which also apply to the setting where the location space is a closed interval):

Theorem 1. *For linear and strictly convex cost functions, an allocation function is truthful if and only if it satisfies weak-monotonicity.*

[1] A single player subfunction on player i is the allocation function restricted to some fixed reported locations of players other than i. See Section 2 for a formal definition.

Theorem 2. *For linear cost functions, an allocation function is truthful if and only if each of its single player subfunctions is partially monotonically non-decreasing.*

Theorem 3. *For strictly convex cost functions, an allocation function is truthful if and only if each of its single player subfunctions is monotonically non-decreasing.*

Although Theorem 1 has its own theoretical significance (there are domains for which weak monotonicity is not sufficient [15]), Theorems 2 and 3 are more informative: they provide a global monotone structure on every single player subfunction, which is more intuitive and easier to verify for practical mechanism design. Strictly convex cost functions enforce a simple monotone structure; the linear cost function case is more intriguing: here monotonicity is required in a hidden (partial) way, captured in our notion of monotone triangular partition.

Consider a single player subfunction f. Since the distance to the facility $|f(x)-x|$ switches sign at $f(x) = x$, the sign of $f(x) - x$ (thus the line $y = x$) is important. Our monotone triangular partition is a partition of the real line into intervals such that, for each interval I, all $f(x)$ are within the closure $\tilde{I}$ of the interval and on the same side of line $y = x$. Hence f is monotonically non-decreasing between different intervals (i.e., f on a right interval is never below f on a left interval), but need not be monotone within an interval I. Graphically, each interval I corresponds to one of the two triangles generated by dividing $I \times \tilde{I}$ with line $y = x$. The sign of the interval, i.e., the uniform sign of $f(x) - x$, corresponds to which side of $y = x$ the triangle resides. Therefore, f is contained in these monotone triangles, and we call this nice interesting structure a monotone triangular partition. For intervals where $f(x) \equiv x$, we allow the corresponding triangle to degenerate into the line segment on $y = x$ intersecting $I \times \tilde{I}$.

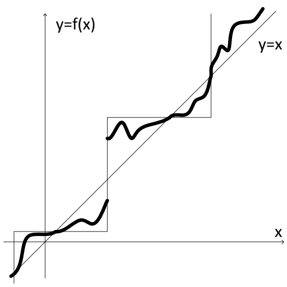

Fig. 1. A monotone triangular partition when C is linear

Unsurprisingly, the payment function in Section 5 is closely related to the triangular partition. Since truthfulness is unaffected by shifting the payment function by any arbitrary constant, we pick an arbitrary reference point and set its payment 0. Then to find out the payment for any point x, imagine taking a walk from the reference point towards the allocated facility location $f(x)$ and counting throughout the way. For cost function $C(d) = d$, we simply count the distance we have walked, but with a sign according to the sign of the interval we are walking at (We take negation of it if $f(x)$ is to the left of the reference point). Hence in the formula, the payment is a directed summation of lengths of intervals corresponding to a monotone triangular partition. For linear cost functions with slope $\alpha \neq 1$ or strictly convex cost functions, we need to adjust the quantity (not as easy as distance here) counted into the payment, but the idea is the same.

2 Preliminaries and Notation

Now we formally define a one-dimensional single facility location game. Suppose there are n players and player i's location is represented by a real number $x_i \in \mathbb{R}$. Given a location vector $\mathbf{x} = (x_1, \ldots, x_n)$ of n players, an allocation mechanism chooses a location $y = f(\mathbf{x}) \in \mathbb{R}$ for the single facility and assigns payments $\mathbf{p}(\mathbf{x}) = (p_1(\mathbf{x}), \ldots, p_n(\mathbf{x}))$ to players where player i gets payment $p_i(\mathbf{x})$. The setting without payments restricts $\mathbf{p}(\mathbf{x}) \equiv \mathbf{0}$.

Let $C(d)$ denote the cost function of all players, which is a smooth monotonically-increasing function on nonnegative distances, and can always be normalized to satisfy $C(0) = 0$. Let $\mathbf{r} = (r_1, \ldots, r_n)$ be the true location vector of the n players, in which r_i is player i's private information. Then the utility of player i is $u_i(\mathbf{x}) = -C(|f(\mathbf{x}) - r_i|) + p_i(\mathbf{x})$.

In the game-theoretic model, each player intends to maximize his utility. An allocation mechanism $(f, \mathbf{p})$ is incentive compatible, or truthful, if for each player i, reporting his true location r_i always maximizes his utility. Formally, it requires that, for each player i, for each fixed reported locations of players other than i, written as $\mathbf{x}_{-i} = (x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n)$, and for any r_i and x_i, we have $u_i((r_i, \mathbf{x}_{-i})) \geq u_i((x_i, \mathbf{x}_{-i}))$. We call an allocation function f truthful if there exists a payment function vector $\mathbf{p}$ such that $(f, \mathbf{p})$ is truthful. Our goal is to characterize the set of truthful allocation functions.

For a player i, each fixed reported locations of other players $\mathbf{x}_{-i}$ induces a subfunction of the allocation function f on player i's location: $f^i_{\mathbf{x}_{-i}}(x_i) = f((x_i, \mathbf{x}_{-i}))$, which can be viewed as an allocation function for a game of a single player i. Thus the notion of truthfulness also applies to such single player subfunctions of f. The following easily proved fact is used extensively in the literature:

Proposition 4. *The allocation function f is truthful if and only if every single player subfunction of f is truthful.*

By Proposition 4, it is meaningful to characterize the set of truthful allocation functions of one player: now an allocation function $f : \mathbb{R} \to \mathbb{R}$ maps a location x to location y of the facility. We want to know for which f there exists payment function $p : \mathbb{R} \to \mathbb{R}$ such that (f, p) is truthful.

Weak monotonicity is a well-known necessary condition for a truthful function. In our setting, it translates to the following: for any $x, x' \in \mathbb{R}$, $C(|f(x) - x|) - C(|f(x) - x'|) \leq C(|f(x') - x|) - C(|f(x') - x'|)$. We obtain a nice characterization by showing weak monotonicity is also sufficient and providing more illustrative conditions equivalent to weak monotonicity on the allocation function f.

It turns out that the shape of the cost function C plays an important role. In our work, we investigate linear, strictly convex and strictly concave cost functions. $C(d) = \alpha d$ where $\alpha > 0$ is a linear cost function; C is strictly convex if for any two points $d_1 \neq d_2$ and $t \in (0, 1)$, it holds that $C(td_1 + (1 - t)d_2) < tC(d_1) + (1 - t)C(d_2)$. Symmetrically C is strictly concave if for any two points $d_1 \neq d_2$ and $t \in (0, 1)$, it holds that $C(td_1 + (1 - t)d_2) > tC(d_1) + (1 - t)C(d_2)$.

3 Implication of Weak-Monotonicity

In this section, for linear and strictly convex cost functions respectively, derive from weak monotonicity an equivalent condition on every single player subfunction.

3.1 Convex Cost Functions

Lemma 5. *If the cost function is strictly convex, a single player allocation function f satisfies weak monotonicity if and only if it is monotonically non-decreasing, i.e., $f(x_1) \leq f(x_2)$ for any $x_1 < x_2$.*

Proof. We use the following property of strictly convex functions:

Proposition 6. *If function C is strictly convex, $C(d_1)+C(d_4) > C(d_2)+C(d_3)$ holds for any $d_1 < d_2 \leq d_3 < d_4$ satisfying $d_1 + d_4 = d_2 + d_3$.*

Now given a strictly convex cost function C, for any $x_1 < x_2$, we claim that function $\Delta(z) = C(|z - x_1|) - C(|z - x_2|)$ is a monotonically increasing function: For any $z_1 < z_2 \leq x_1$, set $d_1 = x_1 - z_2$, $d_4 = x_2 - z_1$, $d_2 = \min(x_1 - z_1, x_2 - z_2)$ and $d_3 = \max(x_1 - z_1, x_2 - z_2)$. We can easily check $d_1 < d_2 \leq d_3 < d_4$ and $d_1 + d_4 = d_2 + d_3$. By Proposition 6, $C(x_1 - z_2) + C(x_2 - z_1) > C(x_1 - z_1) + C(x_2 - z_2)$. We rearrange the terms and change the distances to the form of absolute values to get $C(|z_1 - x_1|) - C(|z_1 - x_2|) < C(|z_2 - x_1|) - C(|z_2 - x_2|)$, i.e., $\Delta(z_1) < \Delta(z_2)$.

The case $x_2 \leq z_1 < z_2$ is symmetric. For $x_1 \leq z_1 < z_2 \leq x_2$, we have $C(|z_1 - x_1|) = C(z_1 - x_1) < C(z_2 - x_1) = C(|z_2 - x_1|)$ and $C(|z_1 - x_2|) = C(x_2 - z_1) > C(x_2 - z_2) = C(|z_2 - x_1|)$. Taking the difference of the two inequalities gives $C(|z_1 - x_1|) - C(|z_1 - x_2|) < C(|z_2 - x_1|) - C(|z_2 - x_2|)$, i.e., $\Delta(z_1) < \Delta(z_2)$.

The monotonicity of the entire function Δ can be easily derived by its monotonicity on the three closed intervals $z \leq x_1$, $x_1 \leq z \leq x_2$, $z \geq x_2$ derived above. The condition of weak monotonicity can be rewritten as $\Delta(f(x_1)) \leq \Delta(f(x_2))$, which holds if and only if $f(x_1) \leq f(x_2)$ since function Δ is strictly monotonically increasing.

3.2 Linear Cost Functions

Lemma 7. *If the cost function is linear, a single player allocation function f satisfies weak monotonicity if and only if for any $x_1 < x_2$, $f(x_1) > x_1$ implies $f(x_2) \geq \min(x_2, f(x_1))$.*

This property is weaker than being monotonically non-decreasing, which we call *partially monotonically non-decreasing.*

Proof. If the cost function $C(d) = \alpha d$ $(\alpha > 0)$, for any $x_1 < x_2$, function $\Delta(z) = C(|z - x_1|) - C(|z - x_2|)$ is a continuous non-decreasing piecewise linear function: For $z \leq x_1$, $\Delta(z)$ is a negative constant $\alpha(x_1 - x_2)$, whereas it constantly equals its negation $\alpha(x_2 - x_1)$ for $z \geq x_2$. Between $z = x_1$ and $z = x_2$ is a linear piece of slope $2\alpha > 0$.

The condition of weak monotonicity can be rewritten as $\Delta(f(x_1)) \leq \Delta(f(x_2))$. This always holds for $f(x_1) \leq x_1$ since $\Delta(f(x_1))$ reaches the minimum. If $f(x_1) > x_1$, there are two cases: for $f(x_1) < x_2$, $f(x_1)$ belongs to the linearly increasing piece, so $\Delta(f(x_1)) \leq \Delta(f(x_2))$ if and only if $f(x_1) \leq f(x_2)$; otherwise, $f(x_1) \geq x_2$, $\Delta(f(x_1))$ reaches the maximum, thus $\Delta(f(x_2))$ is also the maximum, i.e., $f(x_2) \geq x_2$. The summary of the two cases is exactly $f(x_2) \geq \min(x_2, f(x_1))$.

4 Monotone Triangular Partition

In this section, we show that weak monotonicity implies a monotone triangular partition. We start with the following key separation theorem:

Theorem 8. *If f is partially monotonically non-decreasing, then for any $x_1 < x_2$ satisfying $(f(x_1) - x_1)(f(x_2) - x_2) < 0$, there exists $x^* \in [x_1, x_2]$ such that $f(x) \leq x^*$ for $x < x^*$ and $f(x) \geq x^*$ for $x > x^*$. In particular, $f(x^*) = x^*$ for the case $f(x_1) > x_1$ and $f(x_2) < x_2$.*

Proof. $(f(x_1) - x_1)(f(x_2) - x_2) < 0$ implies that $f(x_1) - x_1$ and $f(x_2) - x_2$ have different signs. There are two cases:

If $f(x_1) < x_1$ and $f(x_2) > x_2$, we take $x^* = \inf\{x : f(x) > x, x \geq x_1\}$, where the infimum exists since the set is non-empty (contains x_2) and bounded below by x_1. Clearly $x^* \in [x_1, x_2]$.

First, we show $f(x) \leq x^*$ for $x < x^*$ in this case. This is immediate for $x \geq x_1$ by the definition of x^*. For $x < x_1$, suppose $f(x) > x^*$ for contradiction. Then since f is partially monotonically non-decreasing, $x < x_1$ and $f(x) > x$ implies $f(x_1) \geq \min(x_1, f(x)) \geq \min(x_1, x^*) = x_1$, contradicting that $f(x_1) < x_1$.

Next, for $x > x^*$, we want to show $f(x) \geq x^*$. By the definition of x^*, there exists $x' \in [x^*, x)$ satisfying $f(x') > x'$. Now we apply the partial monotonicity condition again with $x' < x$: $f(x') > x'$ implies $f(x) \geq \min(x, f(x')) > x^*$.

Interestingly, the second case $f(x_1) > x_1$ and $f(x_2) < x_2$ is not symmetric. Here we take $x^* = \inf\{f(x) : f(x) < x, x \geq x_1\}$. Again the set is non-empty since it contains $f(x_2)$. It is bounded below by x_1 since we can apply the partial monotonicity condition with $x_1 < x$, $f(x_1) > x_1$ and get $f(x) \geq \min(x, f(x_1)) > x_1$.

Hence x^* is also well-defined in this case. Moreover, the above argument plus $f(x_2) < x_2$ guarantees $x^* \in [x_1, x_2]$. For this case, we need to show a slightly stronger statement: $f(x) \le x^*$ for $x \le x^*$ and $f(x) \ge x^*$ for $x \ge x^*$, which at $x = x^*$ implies $f(x^*) = x^*$.

First we prove $f(x) \le x^*$ for $x \le x^*$. For contradiction, suppose $f(x) > x^*$. By the definition of x^*, there exists $x'(\ge x_1)$ such that $x^* \le f(x') < f(x)$ and $f(x') < x'$: This immediately implies $f(x') < \min(x', f(x))$; On the other hand, $x \le x^* \le f(x') < x'$ and $f(x) > x^* \ge x$ allows us to apply the partial monotonicity condition and get $f(x') \ge \min(x', f(x))$, which is a contradiction.

Now $f(x) \ge x^*$ for $x \ge x^*$. For those x satisfying $f(x) \ge x$, $x \ge x^*$ immediately gives $f(x) \ge x^*$; otherwise, $f(x) < x$, then the definition of x^* implies that $f(x)$ is at least the infimum x^*.

Theorem 8 enables us to repeatedly partition the real line into intervals: as long as there exist two points $x_1 < x_2$ within the same interval I whose signs of $f(x) - x$ are different, we disect the interval at x^*. Point x^* belongs to the left subinterval I_1 if $f(x^*) < x^*$, to the right subinterval I_2 if $f(x^*) > x^*$, and to either one of the two if $f(x^*) = x^*$ (Note that I_1 and I_2 are both nonempty but may only contain a single point). This disection at the same time disects the allocation function by line $y = x^*$: for x in I_1 and all intervals left to I_1, $f(x) \le x^*$, the allocation function does not exceed this line; symmetrically for x in I_2 and all intervals right to I_2, $f(x) \ge x^*$, the allocation function never goes below the line. Graphically, f appears within the region $x \le x^*, y \le x^*$ and $x \ge x^*, y \ge x^*$.

Eventually, we get a partition $\mathcal{P} = \{I\}$ of $\mathbb{R}$ satisfying the following:

- Within each interval I, the sign of $f(x) - x$ is uniformly $\delta_I \in \{-1, 1\}$, i.e., $\delta_I(f(x) - x) \ge 0$ for all $x \in I$.
- For each interval I, $f(x) \in \tilde{I}$ for all $x \in I$, where $\tilde{I}$ is the closure of I.
- Between different intervals I_1 and I_2, if I_1 is to the left of I_2, $f(x_1) \le f(x_2)$ for any $x_1 \in I_1, x_2 \in I_2$.

The second property is immediate from the disecting argument in the description of our partition process; and the last property immediately follows from the second.

Graphically, the partition $\mathcal{P}$ defines a triangular structure: each interval $I \in \mathcal{P}$ corresponds to a triangle T_I : $T_I = \{(x, y) : x \in I, y \in \tilde{I}, y \le x\}$ for $\delta_I = -1$ and $T_I = \{(x, y) : x \in I, y \in \tilde{I}, y \ge x\}$ for $\delta_I = 1$. And the allocation function f only appears within the set of triangles. Moreover, the triangular structure is "monotonic" in the sense that "a triangle to the right is always above". Therefore, we call such a partition $\mathcal{P}$ a *monotone triangular partition.*

Combining Theorem 8 with Lemma 7 in Subsection 3.2, we derive that weak monotonicity guarantees the existence of such a partition for linear cost functions. For convex cost functions, Lemma 5 says that weak monotonicity implies that the allocation function f is monotonically non-decreasing, which is stronger than the condition of partially monotonically non-decreasing required in Theorem 8. Thus a monotone triangular partition exists for convex cost functions as well. This can also be derived directly from the monotonicity of the allocation function.

5 Incentive Compatible Payments

In this section, for any allocation function f that admits a monotone triangular partition, we would like to provide a payment function p such that (f, p) is truthful. We have an explicit formula of p for partitions where any finite range $[a, b]$ only intersects finitely many intervals, i.e., $\mathcal{P}$ can be written as $\{I_i : b_\ell \leq i \leq b_r\}$, where the I_i's are ordered from left to right (possibly $b_\ell = -\infty$, or $b_r = +\infty$, or both). For general f that does not admit a partition of this form, the same idea works; yet it involves infinite summations and makes our argument notationally much more complicated. Handling such technical details is not the focus of our paper here.

Now given a monotone triangular partition $\mathcal{P} = \{I_i : b_\ell \leq i \leq b_r\}$, let $\{a_i : b_\ell \leq i \leq b_r + 1\}$ be the set of boundary points, where $a_i \leq a_{i+1}$ and the left/right endpoint of I_i is a_i/a_{i+1}. $\mathcal{P}$ may contain only finitely many intervals, including the very special case $|\mathcal{P}| = 1$ where $b_\ell = b_r$; otherwise, there is an infinite sequence of intervals to the left end of the real line ($b_\ell = -\infty$), or to the right end ($b_r = +\infty$), or both. If b_ℓ is finite, $a_{b_\ell} = -\infty$; If b_r is finite, $a_{b_r+1} = +\infty$. Other than these two, all a_i's are finite.

A monotone triangular partition $\mathcal{P} = \{I_i : b_\ell \leq i \leq b_r\}$ of $\mathbb{R}$ satisfies the following three properties:

- Each interval I_i is associated with $\delta_i \in \{-1, 0, 1\}$, which denotes the uniform sign of $f(x) - x$. We have $\delta_i(f(x) - x) \geq 0$ for all $x \in I_i$, and in particular, $\delta_i = 0$ requires $f(x) \equiv x$ for all $x \in I_i$.
- For each i, $f(x) \in \tilde{I}_i$ for all $x \in I_i$, where $\tilde{I}_i$ is the closure of I_i.
- For any $i < j$ and $x \in I_i, x' \in I_j$, we have $f(x) \leq f(x')$.

Here we allow $\delta_i = 0$ for an interval I_i where $f(x) \equiv x$, while for such an interval, the other two choices -1 and 1 are also allowed. Graphically $\delta_i = 0$ indicates that the corresponding triangle of I_i shrinks to the line segment $\{(x, y) : x \in I_i, y = x\}$. This extra freedom does not add any difficulty to our proofs, but as shown in Subsection 5.2, now our payment function includes the no-payment case, i.e., for an allocation function that is truthful without payments, there is a monotone triangular partition with associated δ's under which our payment function is exactly $p(x) \equiv 0$.

For linear and strictly convex cost functions respectively, we present a formula of the payment function p and show its incentive compatibility based on the above properties. This, combined with Section 4, and the necessity of weak monotonicity, completes the proof of Theorem 1-3. Due to the space limit, we defer the convex cost function part to the full paper.

5.1 Linear Cost Functions

Given a monotone triangular partition $\mathcal{P} = \{I_i : b_\ell \leq i \leq b_r\}$, we define a function $q : \mathbb{R} \to \mathbb{R}$ on the location $y \in \mathbb{R}$ of the public facility as follows:

$q(y) = \delta_0 y$ if $|\mathcal{P}| = 1$; otherwise, choose a reference boundary point a_{b_0}, where $b_\ell < b_0 \le b_r$.

$$q(y) = \begin{cases} \delta_k(y - a_k) + \sum_{i=b_0}^{k-1} \delta_i(a_{i+1} - a_i), & y \in I_k, b_0 \le k \le b_r \\ -\delta_k(a_{k+1} - y) - \sum_{i=k+1}^{b_0-1} \delta_i(a_{i+1} - a_i), & y \in I_k, b_\ell \le k < b_0 \end{cases}$$

Our definition of q only involves a_i with $b_l + 1 \le i \le b_r$, which are all finite, thus function q is well-defined. Moreover, observe that for any finite boundary point a_k, the value of $q(a_k)$ is the same no matter whether a_k belongs to interval I_{k-1} or I_k. Hence the above formula holds for any $y \in \tilde{I}_k$ as well.

In particular, the value of q at the reference boundary point a_{b_0} is set to 0. Each interval I_i, or part of an interval, contributes to the payment if and only if it is between a_{b_0} and y. The contribution equals its sign δ_i times the length of the interval if it is to the right of a_{b_0}, and its negation if it is to the left of a_{b_0}. Under this summarization, the difference of the function value of any two points y and y' is irrelevant to the choice of the reference point a_{b_0}. The following lemma can be easily proven:

Lemma 9. *Suppose $y \in \tilde{I}_k$ and $y' \in \tilde{I}_{k'}$. For $k < k'$,*

$$q(y') - q(y) = \delta_{k'}(y' - a_{k'}) + \sum_{i=k+1}^{k'-1} \delta_i(a_{i+1} - a_i) + \delta_k(a_{k+1} - y);$$

For $k = k'$, $q(y') - q(y) = \delta_k(y' - y)$.

Theorem 10. *Let f be an allocation function that admits a monotone triangular partition $\mathcal{P} = \{I_i : b_\ell \le i \le b_r\}$, and $C(d) = \alpha d$ ($\alpha > 0$) is the cost function. Then mechanism (f, p) is truthful where the payment function is defined as $p(x) = \alpha q(f(x))$.*

Proof. To prove (f, p) is truthful, we need to show that for any true location x and $x' \ne x$,

$$-C(|f(x) - x|) + p(x) = u(x) \ge u(x') = -C(|f(x') - x|) + p(x'),$$

i.e., reporting true location x always maximizes the player's utility. Substituting $C(d) = \alpha d$, $p(x) = \alpha q(f(x))$ and $f(x) = y$, $f(x') = y'$, the inequality simplifies to

$$q(y) - |y - x| \ge q(y') - |y' - x|.$$

Now we verify this inequality in three cases as follows. Throughout our proof, we repeatedly use the simple fact that, for $x \in I_k$ and $y = f(x)$, $|y - x| = \delta_k(y - x)$. This is immediate from the first property of the partition.

Case 1: x and x' are in the same interval I_k.

In this case, $y, y' \in \tilde{I_k}$ from the second property of the partition. By Lemma 9, we have $q(y')-q(y) = \delta_k(y'-y)$. We substitute this and $|y-x| = \delta_k(y-x)$ and the inequality simplifies to $|y'-x| \geq \delta_k(y'-y) + \delta_k(y-x) = \delta_k(y'-x)$, which always holds given $\delta_k \in \{-1, 0, 1\}$.

Case 2 : $x \in I_k$, $x' \in I_{k'}$ and $k < k'$.

In this case, $y \in \tilde{I_k}$ and $y' \in \tilde{I_{k'}}$. Applying Lemma 9 gives

$$q(y') - q(y) = \delta_{k'}(y' - a_{k'}) + \sum_{i=k+1}^{k'-1} \delta_i(a_{i+1} - a_i) + \delta_k(a_{k+1} - y).$$

On the other hand, $|y' - x| = y' - x = (y' - a_{k'}) + \sum_{i=k+1}^{k'-1}(a_{i+1} - a_i) + (a_{k+1} - x)$.

Putting all equalities together, we get

$$\begin{aligned} q(y') - q(y) - |y' - x| + |y - x| &= (\delta_{k'} - 1)(y' - a_{k'}) + \sum_{i=k+1}^{k'-1} (\delta_i - 1)(a_{i+1} - a_i) \\ &\quad + \delta_k(a_{k+1} - y) - (a_{k+1} - x) + \delta_k(y - x) \\ &\leq \delta_k(a_{k+1} - y) - (a_{k+1} - x) + \delta_k(y - x) \\ &= (\delta_k - 1)(a_{k+1} - x) \leq 0, \end{aligned}$$

given $\delta_k, \delta_{k'} \in \{-1, 0, 1\}$. Rearranging the terms gives exactly the inequality we want to prove.

Case 3 : $x \in I_k$, $x' \in I_{k'}$ and $k > k'$. This case is symmetric to Case 2.

5.2 Generality and Non-uniqueness of Our Payment

As mentioned before, by allowing degenerated triangles (allowing $\delta_I = 0$ for interval I where $f(x) \equiv x$) in our monotone triangular partition, we make our payment formula include the all-zero payment function for truthful allocation functions in the no payment setting.

For games without payment, every single player subfunction behaves as follows: as player's location x grows, the facility location $y = f(x)$ either remains the same, or jumps to a symmetric (higher point) with respect to x, or continues to equal x. Formally, for any single player subfunction of a truthful allocation function, there exists a monotone triangular partition satisfying the following properties:

- For any I with $\delta_I = 0$, $f(x) \equiv x$. This is always required by a monotone triangular partition. We state it here for completeness.
- For any I with $\delta_I = 1$, $f(x)$ always equals to the right endpoint of I.
- For any I with $\delta_I = -1$, $f(x)$ always equals to the left endpoint of I.
- For any I_1 adjacent to I_2 and to the left of I_2, $\delta_{I_1} = -1$ implies $\delta_{I_2} = 1$ and the lengths of the two intervals are equal.

It can be verified that our payment function p in Subsection 5.1 based on the above monotone triangular partition is constant, thus can be made all-zero by a constant shift.

On the other hand, for an interval I where $f(x) \equiv x$, we can still set $\delta_I = -1$ or 1, or even divide it into more intervals and set different δ's. This freedom results in different monotone triangular partitions, which, plugged into our payment formula, results in payment functions that differ more than a constant shift. Therefore, the payment function for a truthful allocation function may not be unique. In contrast, the classic unique-payment theorem [20] states that the payment function is unique for a truthful mechanism when the domain is connected; and the domain of our setting is connected. The inconsistency comes from the fact that our outcome set (the set of possible facility locations) is uncountable, while the theorem assumes the outcome set to be finite. This is called the revenue equivalence in economics literature [13,8].

6 Conclusion and Open Questions

In this paper, we characterize the set of truthful allocation functions for one-dimensional single facility location game with payments: we show the sufficiency of weak monotonicity, and its equivalent global monotone structure on every single player subfunction for linear and strictly convex cost functions respectively.

When investigating concave cost functions, we observe certain anti-monotone feature implied by weak monotonicity, which makes this case greatly different from the cases we have solved. We would love to see characterization results of this case: it is not known yet whether weak monotonicity is sufficient or not. We note here that, when the cost function is concave, the global utility function is still not convex (it is convex in both sides of its true location but not if we view it globally).

Another direction is to consider the game for more facilities, say, two facilities. In this case, the valuation domain for each agent is more complicated. Even the characterization for truthful mechanisms without payment is still open.

References

1. André, B., Rudolf, M., Hossein, N.S.: Path-monotonicity and incentive compatibility. Technical report (2010)
2. Archer, A., Kleinberg, R.: Truthful germs are contagious: A local to global characterization of truthfulness. In: Proceedings of the 9th ACM Conference on Electronic Commerce, pp. 21–30 (2008)
3. Archer, A., Tardos, É.: Truthful mechanisms for one-parameter agents. In: FOCS, pp. 482–491. IEEE Computer Society (2001)
4. Ashlagi, I., Braverman, M., Hassidim, A., Monderer, D.: Monotonicity and implementability. Econometrica 78(5), 1749–1772 (2010)
5. Barberà, S., Jackson, M.: A characterization of strategy-proof social choice functions for economies with pure public goods. Social Choice and Welfare 11(3), 241–252 (1994)

6. Bei, X., Chen, N., Gravin, N., Lu, P.: Budget feasible mechanism design: from prior-free to bayesian. In: STOC, pp. 449–458 (2012)
7. Bikhchandani, S., Chatterji, S., Lavi, R., Mu'alem, A., Nisan, N., Sen, A.: Weak monotonicity characterizes deterministic dominant-strategy implementation. Econometrica 74(4), 1109–1132 (2006)
8. Carbajal, J.C., Ely, J.C.: Mechanism design without revenue equivalence. Journal of Economic Theory (2012)
9. Cuff, K., Hong, S., Schwartz, J.A., Wen, Q., Weymark, J.A.: Dominant strategy implementation with a convex product space of valuations. Social Choice and Welfare 39(2-3), 567–597 (2012)
10. Dhangwatnotai, P., Dobzinski, S., Dughmi, S., Roughgarden, T.: Truthful approximation schemes for single-parameter agents. In: FOCS, pp. 15–24 (2008)
11. Frongillo, R.M., Kash, I.A.: General truthfulness characterizations via convex analysis. arXiv preprint arXiv:1211.3043 (2012)
12. Goldberg, A.V., Hartline, J.D., Karlin, A.R., Saks, M., Wright, A.: Competitive auctions. Games and Economic Behavior 55(2), 242–269 (2006)
13. Heydenreich, B., Müller, R., Uetz, M., Vohra, R.V.: Characterization of revenue equivalence. Econometrica 77(1), 307–316 (2009)
14. Jehiel, P., Moldovanu, B., Stacchetti, E.: Multidimensional mechanism design for auctions with externalities. Journal of Economic Theory 85(2), 258–293 (1999)
15. Lavi, R., Mu'alem, A., Nisan, N.: Towards a characterization of truthful combinatorial auctions. In: FOCS, p. 574 (2003)
16. Mishra, D., Roy, S.: Implementation in multidimensional dichotomous domains. Theoretical Economics 8(2), 431–466 (2013)
17. Moulin, H.: On strategy-proofness and single peakedness. Public Choice 35(4), 437–455 (1980)
18. Müller, R., Perea, A., Wolf, S.: Weak monotonicity and bayes–nash incentive compatibility. Games and Economic Behavior 61(2), 344–358 (2007)
19. Myerson, R.B.: Optimal auction design. Mathematics of Operations Research 6(1), 58–73 (1981)
20. Nisan, N.: Introduction to mechanism design. In: Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V. (eds.) Algorithmic Game Theory, ch. 9, pp. 209–242 (2007)
21. Nisan, N., Ronen, A.: Algorithmic mechanism design. Games and Economic Behavior 35(1-2), 166–196 (2001)
22. Papadimitriou, C.H., Schapira, M., Singer, Y.: On the hardness of being truthful. In: FOCS, pp. 250–259 (2008)
23. Roberts, K.: The characterization of implementable choice rules. In: Aggregation and Revelation of Preferences, pp. 321–348 (1979)
24. Saks, M., Yu, L.: Weak monotonicity suffices for truthfulness on convex domains. In: Proceedings of ACM EC, pp. 286–293. ACM (2005)
25. Singer, Y.: Budget feasible mechanisms. In: FOCS, pp. 765–774 (2010)
26. Sprumont, Y.: The division problem with single-peaked preferences: A characterization of the uniform allocation rule. Econometrica, 509–519 (1991)
27. Vohra, R.V.: Mechanism Design: A linear programming approach, vol. 47. Cambridge Univ. Pr. (2011)