

Improved Approximation Ratio for Strategyproof Facility Location on a Cycle*

Krzysztof Rogowski, Marcin Dziubiński

University of Warsaw, Institute of Informatics
{kr418382, m.dziubinski}@mimuw.edu.pl,

Abstract

We study the problem of design of strategyproof in expectation (SP) mechanisms for facility location on a cycle, with the objective of minimizing the sum of costs of n agents. We show that there exists an SP mechanism that attains an approximation ratio of $7/4$ with respect to the sum of costs of the agents, thus improving the best known upper bound of $2 - 2/n$ in the cases of $n \geq 5$. The mechanism obtaining the bound randomizes between two mechanisms known in the literature: the Random Dictator (RD) and the Proportional Circle Distance (PCD) mechanism of [Meir, 2019]. To prove the result, we propose a cycle-cutting technique that allows for estimating the problem on a cycle by a problem on a line.

1 Introduction

The facility location problem involves a group of n agents, each having a preference over a set of locations within a metric space. The agents have their own ideal location in the space and they prefer locations that are closer to their ideal location. A central authority (the social planner), not knowing the ideal points of the agents, has to choose one of the locations, aiming to minimize the sum of distances between the locations of the agents and the chosen location (the so-called utilitarian welfare objective). An additional problem, beyond the choice of optimal location, arises due to the lack of knowledge of the ideal points of the agents. To address this problem the social planner asks the agents to report their location and uses a mechanism which, given the reports of the agents, determines the location to be chosen. Since the agents will report their location in order to minimize their own distance to the location chosen by the mechanism, there is no guarantee that their reports will be truthful. To resolve this problem, the social planner is restricted to choosing a strategyproof mechanism, under which reporting true ideal points is individually weakly optimal, regardless of the reports of other agents. The outcomes of the mechanisms may be either deterministic or randomized.

*Extended version of the paper is available at <https://arxiv.org/abs/2505.12943>

The strategyproof facility location problem attracted interest from researchers for nearly 50 years now and it actively researched to this day. This is due to natural applications, like locating public facilities (schools, healthcare facilities, etc.) in towns where each of the citizens has own preference regarding the best location. These natural, physical, applications extend to virtual ones, like choosing the best time for a meeting.

In this paper we are interested in strategyproof facility location problems on graphs, specifically on cycles. In the seminal work, [Moulin, 1980] obtained a complete characterization of strategyproof mechanisms when the space of possible locations is a line segment. He showed, in particular, that the mechanism choosing the median of the reported points is not only strategyproof but also efficient, in the sense that it minimizes the sum of distances to the ideal points of the agents. [Schummer and Vohra, 2002] extended the characterization of strategyproof mechanisms to graphs. It follows from their characterization that when the graph is a tree then there exists a strategyproof mechanism which is also efficient. If a graph contains a cycle, however, no deterministic strategyproof mechanism is efficient. In a later work, [Meir, 2019] showed that this is also true for randomized mechanisms. This raises a question: how close to efficiency can we get when cycles are present?

To address possible inefficiency of strategy proof mechanisms, [Procaccia and Tennenholtz, 2013] proposed the idea of approximate mechanism design (without money). They introduced the approximation ratio which measures how good the best outcome of a strategy proof mechanism is, related to the optimal outcome. The best known upper bound on the approximation ratio of strategyproof mechanisms for facility location on a graph (under utilitarian welfare objective), $2 - n/2$, was obtained by [Alon *et al.*, 2009] with use of a random dictator (RD) mechanism. [Meir, 2019] improved this bound for the case of a cycle with $n = 3$ agents to $7/6$. In this paper we are interested in the approximation ratio of the strategyproof facility location on a cycle.

1.1 Related Work

The literature on facility location, in general, and on strategyproof facility location, in particular, is vast. In the interest of space, in this review we restrict attention to the paper that are closest related to our work. For an excellent recent litera-

ture review of this field see [Chan *et al.*, 2021].

Our work falls into the area of mechanism design without money as applied to facility location. Fundamental papers in this area, [Moulin, 1980], [Schummer and Vohra, 2002], and [Procaccia and Tennenholtz, 2013] were already discussed above.

We are specifically interested in the approximation ratio of strategy proof facility location mechanisms on a cycle, with a single location being selected and under utilitarian welfare objective. As already mentioned, [Alon *et al.*, 2009] obtained the upper bound of $2 - 2/n$ on this approximation ratio, by the RD mechanism. This result was improved by [Meir, 2019] for the case of $n = 3$ agents. He defines new strategy proof mechanisms for a cycle, the Proportional Circle Distance (PCD) mechanism and the q -Quadratic Circle Distance (q -QCD) mechanism. The $1/4$ -QCD is strategyproof in the case of $n = 3$ agents and obtains the tight bound of $7/6$. The PCD mechanism is strategyproof for any odd number of agents. However, when the number of agents grows, its approximation ratio approaches 2. [Dokow *et al.*, 2012] studied the strategyproof facility location on graphs where only vertices of the graph can be chosen. In the case of cycles they showed that when the number of vertices is sufficiently large, the strategyproof mechanisms must be close to dictatorial. In the case of small numbers of vertices, non-dictatorial, anonymous, strategyproof mechanisms exist. Other notable works on strategyproof facility locations on cycles, with objective other than utilitarian social welfare, include [Alon *et al.*, 2010] and [Cai *et al.*, 2016].

1.2 Our contribution

We obtain a new upper bound on the approximation ratio for the strategyproof facility location on a cycle. We show that this ratio is bounded from above by $7/4$ for any odd $n \geq 3$. This improves the previously known upper bound of $2 - 2/n$ for the case of $n \geq 5$. To obtain this result, we propose a mixed randomized mechanism, RD+PCD, which mixes between using the RD mechanism and the PCD mechanism. The key challenge in proving the bound are non-linear distances on the cycle. To overcome this issue we propose a cut technique, which “cuts” the cycle in a properly chosen point and allows for reducing the analysis of the bounds on a cycle to a line segment. We complement the theoretical analysis of the bound with computational analysis of the performance of the RD+PCD mechanism. The analysis suggests that the upper bound could be improved even further, to $3/2$. It is important to note that in the paper we consider the notion of strategyproofness in expectation, weaker to the notion of universal strategyproofness [Mu’alem and Schapira, 2018].

2 Preliminaries

A facility location problem consists of a set of agents N , where $|N| \geq 3$ is assumed to be an odd number throughout this paper, and a domain (a metric space). Typically, the domain is considered to be a graph equipped with a metric. In this work, we restrict attention to a continuous cycle of length 1. Introducing that object as subsets of \mathbb{R} simplifies the notation in subsequent sections of this paper.

A *cycle* of length 1, obtained from the line segment $[-0.5, 0.5]$ by joining its endpoints, is denoted by G . For the purposes like comparisons or arithmetic operations, the joined endpoint of the cycle is -0.5 . The distance between any two points $v_1, v_2 \in G$ on the cycle is given by $d(v_1, v_2) = \min(|v_2 - v_1|, 1 - |v_2 - v_1|)$. Let us orientate the cycle in such a way that clockwise movement corresponds to increasing point values with one exception of passing through the joined endpoint -0.5 . An *arc* between two points $v_1, v_2 \in G$ on the cycle is defined as the set of points traversed on the cycle when moving clockwise from v_1 to v_2 .

Let $k = (|N| - 1)/2$. For convenience, we assume that the set of agents is given by $N = \{-k, -k+1, \dots, -1, 0, 1, \dots, k\}$.

Every agent $i \in N$ has preferences over the points of the graph, determined by their ideal point, b_i . These preferences are determined by the *cost function*, defined as the distance between the agent’s ideal point b_i and a chosen point $v \in G$. When comparing two points on the graph, agent i prefers the point with lower cost.

The collection of agents’ ideal points constitutes a *profile* and is represented by a function $b : N \rightarrow G$, where b_i is the ideal point of agent i and $b^{-1}(v)$ are the agents reporting v .

For any set A , a *lottery* over A is defined as a discrete probability distribution on A . We represent a lottery as a function $l : A \rightarrow [0, 1]$ satisfying the condition that the sum of values in its image equals 1. The set of all lotteries over A is denoted by $\Delta(A)$.

Agents’ preferences are extended to lotteries over the points of the graph by using the expected values. The *cost* for an agent with ideal point v under a lottery $l \in \Delta(G)$ is a linear extension of the deterministic cost, defined as the expected distance between v and the outcome of the lottery:

$$c_v(l) = \mathbb{E}_{X \sim l}[d(v, X)].$$

An agent prefers lotteries with lower costs.

In this work, the quality of any lottery $l \in \Delta(G)$ is evaluated using the *social cost*, defined as the sum of costs of all agents. Formally, for any profile b , let

$$sc_b(l) = \sum_{i \in N} c_{b_i}(l).$$

For simplicity, we sometimes abuse the notation and compute the social cost for a single point $v \in G$, treating it as a degenerate lottery where the given point is always selected. The minimal social cost is referred to as the *optimal cost* and denoted by opt_b . Formally,

$$opt_b = \inf_{v \in G} sc_b(v).$$

It is straightforward to see that, due to the definition of agent costs via expected values, the optimal cost is always achieved by some point.

In this paper, we consider the problem faced by a social planner who, without knowledge of the true profile, aims to choose a lottery over the points of the graph to minimize social cost. To achieve this, the social planner employs a mechanism design approach.

Definition 1. Let B be an arbitrary set of profiles. A function $M : B \rightarrow \Delta(G)$, which maps profiles to lotteries over the points of the graph, is called a *mechanism*.

Given a mechanism, agents report their ideal points, not necessarily truthfully. Based on these reports, the mechanism determines a lottery, which is then used to select a point of the graph.

2.1 Common Mechanism Properties

Let b denote an arbitrary profile.

Definition 2. A mechanism M is *strategyproof (SP)* if, for any agent $i \in N$, the cost incurred by i when truthfully reporting their ideal location b_i is no greater than the cost incurred when reporting any other point $v \in G$. Formally:

$$c_{b_i}(M(b)) \leq c_{b_i}(M(b[i \rightarrow v])),$$

where $b[i \rightarrow v]$ denotes the profile in which agent i reports point v instead of their true ideal location b_i .

Strategyproofness implies that no agent can gain by misreporting their ideal location to the mechanism. Consequently, when considering SP mechanisms, and under the assumption that agents act rationally, they will truthfully report their ideal points.

In this work, we focus exclusively on SP mechanisms. Based on the above reasoning, we make little distinction between the true profiles of agents' ideal points and the profile reported by the agents, assuming they are equal.

In addition, mechanism M is said to be:

- *Anonymous*, if it does not distinguish between agents, that is for any permutation of agents $\pi : N \rightarrow N$, the outcome of the mechanism M remains unchanged. Formally, $M(b) = M(b \circ \pi)$.
- *Neutral*, if it does not distinguish between similar graph's locations. Formally, for any automorphism f of the graph G , the outcome of the mechanism M is transformed according to f i.e., $M(f \circ b) \circ f = M(b)$.
- *Peaks-only*, if the support of the resulting lottery $M(b)$ is contained within the points reported by the agents, i.e., in the image $\text{Im } b$.
- α -*approximation* if the social cost of the mechanism's outcome is no greater than α times the optimal cost. Formally, $\text{sc}_b(M(b)) \leq \alpha \cdot \text{opt}_b$.

The *approximation ratio* of a mechanism M for a set of profiles B is defined as the smallest α such that M is an α -approximation for all profiles from B . We denote this value by $\text{apx}_M(B)$. We often abuse the notation and write $\text{apx}_M(b)$, where b is a single profile, to denote the approximation ratio of the mechanism M for the set $\{b\}$. In this paper we are interested in mechanisms with the smallest approximation ratio for all profiles on the cycle. This is equivalent to finding the upper bound of apx_M for singletons that holds uniformly across all profiles.

2.2 Known SP mechanisms

We now present two important mechanisms from the literature that are central to our work.

Definition 3. *RD* (Random Dictator) is a mechanism that returns a lottery $l \in \Delta(G)$, where the probability of selecting any point $v \in G$ is proportional to the number of agents choosing it in profile b . Formally, $l(v) = |b^{-1}(v)|/|N|$.

It is known [Alon *et al.*, 2009] that for any graph and any profile, the approximation ratio of this mechanism does not exceed $2 - 2/|N|$.

Definition 4. The *Proportional Circle Distance (PCD)* mechanism [Meir, 2019] is defined for cycles and an odd number of agents. The mechanism operates according to following steps:

1. Fix a linear order \prec of agents corresponding to the clockwise arrangement of their reports on the cycle, with ties broken arbitrarily (it can be verified that this does not affect the outcome of the mechanism).
2. For each agent i , define the *arc opposing* agent i as the arc between the report of the agent who is k positions after i in the order \prec and the report of the agent who is k positions before i in the order \prec . The terms “after” and “before” are understood with respect to the clockwise traversal of the cycle; for instance, the agent immediately after the last agent in \prec is the first agent in \prec .
3. Select the report of each agent i with a probability proportional to the length of the arc opposing agent i .

The approximation ratio of the PCD mechanism is upper-bounded by 2. This bound is supported by the existence of a sequence of profiles (for varying sets of agents) in which the approximation ratio of the PCD mechanism approaches 2 [Meir, 2019]. The *RD* and *PCD* mechanisms are strategyproof, anonymous, neutral, and peaks-only.

3 Analysis

In this section, we prove the main result of the paper:

Theorem 1. *For any set of agents with an odd cardinality and a cyclic graph G of length 1, there exists a strategyproof mechanism M whose approximation ratio is bounded from above by $7/4$.*

Remark 1. *The above theorem can be easily generalized to a cycle G_z of arbitrary length $z > 0$. Let f be a mapping that uniformly scales the cycle G_z to G . The mechanism M from Theorem 1 can be adapted to G_z by combining it with the mapping f i.e., considering mechanism M' defined for every profile $b \in G_z^N$ as $M'(b) = M(f \circ b) \circ f$. This operation preserves both the strategyproofness and the approximation ratio of the mechanism.*

The proof of Theorem 1 is constructive: we present a mechanism that satisfies conditions of the theorem, i.e., it is SP and achieves an approximation ratio bounded by $7/4$.

Let $M_1 + M_2$ denote a mechanism that is a *mixture* of mechanisms M_1 and M_2 . The mechanism $M_1 + M_2$, for every profile b , returns the outcome of M_1 on b with probability $1/2$ and the outcome of M_2 on b with probability $1/2$. Formally, for any point $v \in G$, the probability of selecting v in the resulting lottery $l = (M_1 + M_2)(b)$ of the mixed

mechanism is the average of its probabilities in the lotteries $l_1 = M_1(b)$ and $l_2 = M_2(b)$:

$$l(v) = \frac{l_1(v) + l_2(v)}{2}.$$

Mixing mechanisms has several desirable properties, making it a promising tool for constructing SP mechanisms with low approximation ratio:

Observation 1. *For any mechanisms M_1 and M_2 that are strategyproof, anonymous, neutral, or peaks-only, their mixture $M_1 + M_2$ also satisfies these properties.*

Additionally, the approximation ratio of the mixed mechanism behaves as follows:

Observation 2. *For a profile b , the approximation ratio of the mixed mechanism $M_1 + M_2$ is the arithmetic mean of the approximation ratios of M_1 and M_2 :*

$$\text{apx}_{M_1+M_2}(b) = \frac{\text{apx}_{M_1}(b) + \text{apx}_{M_2}(b)}{2}.$$

This result follows directly from the definition of the approximation ratio and the linearity of the social cost.

Consider the mechanism $RD + PCD$. By Observation 1, it is strategyproof. Therefore, to conclude the proof of Theorem 1, it suffices to show that the approximation ratio of the $RD + PCD$ mechanism is bounded above by $7/4$.

Let us begin with some trivial bounds.

Remark 2. *The approximation ratio of the $RD + PCD$ mechanism equals $(\text{apx}_{RD}(b) + \text{apx}_{PCD}(b)) / 2$ (by Observation 2), which:*

- cannot be bounded from above by any value smaller than $3/2$, since there exist profiles for which approximation ratio of RD is arbitrarily close to 2, while approximation ratio of PCD is at least 1,
- is at most 2, under the conjecture that approximation ratio of PCD mechanism equals 2 [Meir, 2019].

The bound of $7/4$, stated in Theorem 1, lies in the middle between these two values and we will improve this result even further in Section 4.

Approximation ratio of the $RD + PCD$ mechanism, by definition, is the smallest α such that inequality:

$$\text{sc}_b((RD + PCD)(b)) \leq \alpha \cdot \text{opt}_b$$

holds for every profile b . The above formula is not the best to work with and can be simplified if the optimal cost is greater than 0. Let us start with considering the border case of $\text{opt}_b = 0$, to exclude it from further analysis:

Remark 3. *Consider a profile b , for which $\text{opt}_b = 0$. Then the approximation ratio of the $RD + PCD$ mechanism is 1.*

Proof. If $\text{opt}_b = 0$ then all the agents report the same point under b . Let us denote this point by v . By Observation 1, the $RD + PCD$ mechanism is peaks-only and so its outcome on b is v with probability 1. Hence, the social cost of the mechanism output is 0 and the approximation ratio is 1. \square

If $\text{opt}_b > 0$, which we assume for the rest of the paper, the following holds:

$$\text{apx}_{RD+PCD}(b) = \frac{\text{sc}_b((RD + PCD)(b))}{\text{opt}_b}.$$

Estimating the above expression for a mechanism operating on a cycle presents new challenges as compared to the case of a line segment:

1. For the line segment, it is known that the optimal cost opt_b corresponds to the social cost for the point preferred by the median agent. For the cycle, we are not aware of any compact description of opt_b .
2. The social cost depends on the distances between the lottery points chosen by the mechanism and the agents' ideal points. Distances on a line segment have significantly simpler form in comparison to those on a cycle.

We address the first of these problems in Section 3.1, by selecting a subset of normalized profiles $B' \subsetneq B$, for which the point generating optimal cost is fixed. In addition, we require that agents are ordered according to the clockwise traversal of the cycle, which simplifies further analysis. We show that due to the anonymity and neutrality of the $RD + PCD$ mechanism, set B' is representative in terms of the approximation ratio values achieved by $RD + PCD$, i.e., for every $b \in B$, there exists $b' \in B'$ such that $\text{apx}_{RD+PCD}(b) = \text{apx}_{RD+PCD}(b')$. This allows us to narrow down further considerations to the set of normalized profiles B' .

The second problem, of the more complex nature of distances between points on a cycle, is addressed in Section 3.2. We estimate these distances through their counterparts on a line segment obtained by cutting the cycle before some point. Such estimation increases the social costs associated with certain points of the graph, which could potentially increase the optimal cost (and thus decrease the estimated value yielding invalid bound). Fortunately, due to the earlier restriction to the set of normalized profiles B' , the point minimizing the social cost is fixed. This enables us to select the cutting point so that the minimal social cost is preserved, thereby obtaining a valid estimate.

In Section 3.3, we provide concise formulas for parts of the approximation ratio after the cut of the $RD + PCD$ mechanism corresponding to the social costs of the RD and PCD mechanisms. We substitute these expressions into the formula for the approximation ratio after the cut of the $RD + PCD$ mechanism, obtaining an estimate ϕ . Finally, in Section 3.4, we upper-bound the value of ϕ by $7/4$ on a set of normalized profiles via a series of technical lemmas.

The formal proof of Theorem 1, based on the developed lemmas, is presented in Section 3.5.

3.1 Normalization of profiles

We start with the definition of a subset of profiles for which the formula for the approximation ratio of the $RD + PCD$ mechanism is simplified, but still can obtain the same values.

Definition 5. A profile b is *normalized* if:

- point 0 minimizes the social cost in b over all the points of the cycle,
- b is a non-decreasing function with respect to the natural order of agents and the clockwise order of the points of the cycle,
- agent 0 reports the point 0.

The following lemma allows us to restrict attention to the normalized profiles.

Lemma 1. *Let M be any mechanism that is anonymous and neutral. Let b be any profile. There exists a normalized profile, b' , such that $\text{apx}_M(b) = \text{apx}_M(b')$.*

Proof. Let b be any profile. A normalized profile b' with the same approximation ratio is constructed as follows:

1. By the neutrality of the M mechanism, we cyclically shift the naming of the points of the graph, so that point 0 generates the minimal social cost.
2. By the anonymity property, we can freely rename agents without affecting the result of the mechanism. We traverse the cycle G clockwise starting from the joined endpoint -0.5 , numbering agents in the order their reports are encountered (breaking ties arbitrarily).

Due to the properties of the point generating the minimal social cost on the cycle, the above two steps ensure that the third condition of normalization is satisfied, namely, that agent 0 reports point 0. Assume, to the contrary, that this is not true, i.e., $b'_0 \neq 0$. Firstly consider the case $b'_0 > 0$. By step 2, b' is non-decreasing, so any agent $i > 0$ reports a point greater than or equal to b'_0 , i.e., $b'_i \geq b'_0$. At the same time, $b'_i \leq 0.5$ since agents selecting negative points are numbered before agent 0. It follows that moving from point 0 to point b'_0 brings us closer to the reports of strictly more than half of the agents (all agents with indices $i \geq 0$). Thus, the social cost of point b'_0 is less than the social cost of point 0, a contradiction to the assumption that point 0 generates the minimal social cost. The case of $b'_0 < 0$ is handled analogously. \square

3.2 Cut of the profile

After introducing the concept of normalization of profiles, let us proceed to analyzing the approximation ratio of the $RD + PCD$ mechanism. This quantity, by definition, depends on the social cost, which, in turn, depends on the distances between the agents' ideal points and the points in the support of the output lottery of the mechanism. The definition of distance between the points on a cycle involves a minimum. This non-linearity makes a direct analysis of the approximation ratio of the $RD + PCD$ mechanism challenging. We address this issue by replacing the distance function used to compute the social cost with a simpler one. Let d' denote a line-segment-like distance function on cycle G , defined for any two points $v_1, v_2 \in G$ as $d'(v_1, v_2) = |v_2 - v_1|$. Let us denote quantities such as cost, social cost, and the approximation ratio computed with respect the distance function d' , as c' , sc' , apx' , respectively.

For any pair of points $v_1, v_2 \in G$, the metric d is defined as $d(v_1, v_2) = \min(|v_2 - v_1|, 1 - |v_2 - v_1|)$. This corresponds to the general definition of distance on graphs as the length

of the shortest path between two points. The metric d' is derived from d by removing one of the two components of the minimum. Conceptually, this is equivalent to excluding paths that traverse the joined endpoint of the cycle when defining the distance. Intuitively, we can imagine that for the purpose of measuring distances, the cycle is “cut” just before point -0.5 . Due to this intuition, the transition from the cycle with the metric d to the cycle with the metric d' will henceforth be referred to as the *cut* of the cycle.

The following inequality connects the value of the social cost for different points of the cycle G computed before and after the cut (with respect to the original distance function, d , and the new distance function, d'):

Lemma 2. *Let b be any normalized profile. For any lottery $l \in \Delta(G)$, it holds that:*

$$sc'_b(l) \geq sc_b(l).$$

Moreover, the social cost associated with the point 0 is preserved by the cut, i.e., $sc'_b(0) = sc_b(0)$.

Proof. The social cost of a point $v \in G$ is defined as a linear combination of distances between v and the agents' ideal points. This definition is extended to lotteries by taking expected values. The first statement holds because d' never decreases the distances between the points, as compared to d , and therefore the values on which the social cost is based can only increase after the cut. The second statement follows from the fact that the cut does not alter the shortest path between 0 and any other point. \square

Based on Lemma 2 and Observation 2, we can establish the following upper bound on the approximation ratio of the $RD + PCD$ mechanism:

Corollary 1. *For any normalized profile b , the approximation ratio of the $RD + PCD$ mechanism is bounded by*

$$\phi(b) := \frac{sc'_b(RD(b)) + sc'_b(PCD(b))}{2sc'_b(0)}.$$

Proof.

$$\begin{aligned} \text{apx}_{RD+PCD}(b) &= \frac{\text{apx}_{RD}(b) + \text{apx}_{PCD}(b)}{2} = \\ &= \frac{sc_b(RD(b)) + sc_b(PCD(b))}{2\text{opt}_b} = \\ &= \frac{sc_b(RD(b)) + sc_b(PCD(b))}{2sc_b(0)} \leq \\ &= \frac{sc'_b(RD(b)) + sc'_b(PCD(b))}{2sc'_b(0)} = \phi(b). \quad \square \end{aligned}$$

3.3 Estimations of the social cost after the cut

In the previous subsection we bounded the value of the approximation ratio of the $RD + PCD$ mechanism by the expression $\phi(b)$, which depends on $sc'_{RD}(b)$, $sc'_{PCD}(b)$: the social costs of outputs of RD and PCD after the cut. In this subsection we provide concise forms for these quantities. Since deriving those values involves only algebraic manipulations, we postpone the detailed proofs to the appendix.

Lemma 3. For any normalized profile b , the social cost of the outcome of mechanism RD after the cut is equal to:

$$sc'_b(RD(b)) = \sum_i \frac{4|i|}{2k+1} |b_i|.$$

Lemma 4. For any normalized profile b the social cost of the outcome of mechanism PCD after the cut is equal to:

$$sc'_b(PCD(b)) = \sum_{j>0} |b_j|(2k+1-2j)(|b_{j-k-1}| - |b_{j-k}|) + \sum_j |b_j| + \sum_{j<0} |b_j|(2k+1+2j)(|b_{j+k+1}| - |b_{j+k}|).$$

3.4 Estimating $\phi(b)$

Substituting the results from Section 3.3 into the expression for $\phi(b)$, allows us to express it as a fraction with numerator:

$$\begin{aligned} & \sum_i \frac{4|i|}{2k+1} |b_i| + \sum_j |b_j| + \\ & \sum_{j>0} |b_j|(2k+1-2j)(|b_{j-k-1}| - |b_{j-k}|) + \\ & \sum_{j<0} |b_j|(2k+1+2j)(|b_{j+k+1}| - |b_{j+k}|) \end{aligned} \quad (1)$$

and denominator:

$$2sc'_b(0) = 2 \sum_i d(b_i, b_0) = 2 \sum_i |b_i - 0| = 2 \sum_i |b_i|. \quad (2)$$

In this section, we bound the value of $\phi(b)$ by $7/4$ for any normalized profile b . To achieve this, we break the proof into several steps, stated as lemmas. We provide some intuition behind each lemma and defer the detailed proofs to the appendix.

We begin by extending the set of considered profiles, as the current set (the normalized profiles) is not easily described due to the requirement that point 0 minimizes the social cost. To address this, we define a larger set, B'' , which contains all the non-decreasing functions from the set of agents to $[-0.5, 0.5]$, except the constant $\mathbf{0}$ function, such that $b''_0 = 0$. One can view the elements of B'' as profiles on a closed line segment of length 1, where agent 0 reports the point 0.

We extend the definition of ϕ to be valid on B'' and proceed to estimating the supremum of ϕ over B'' . This approach provides an upper bound for the supremum over the normalized profiles, as B'' is a larger set.

A challenging aspect of estimating the supremum of ϕ over B'' is that B'' is not a closed set, because it excludes $\mathbf{0}$. To address this, we begin by proving that ϕ attains small values in the neighborhood of $\mathbf{0}$.

We introduce the concept of a *dominated* function in B'' : a function $b \in B''$ is dominated if it equals zero for at least $k+1$ agents. Let $B''_0 \subseteq B''$ denote the subset of functions in B'' that are dominated.

Lemma 5. For any dominated $b \in B''_0$, it holds that $\phi(b) \leq 3/2 - 1/n$. Moreover, the equality can be achieved as witnessed by the function $b^\circ \in B''_0$ corresponds to tuple $(0, 0, \dots, 0.5)$ i.e., the function getting value 0 for every agent except agent k for which it gets value 0.5.

Let us define function $w : B'' \rightarrow \mathbb{N}$ which, for any $b \in B''$, returns the number of distinct values of b that are not equal to $-0.5, 0$, or 0.5 . Formally,

$$w(b) = |\text{Im } b \setminus \{-0.5, 0, 0.5\}|.$$

We call a function $b \in B''$ *boundary* if $w(b) = 0$. Let B^* denote the set of all boundary functions in B'' .

We demonstrate that the supremum of ϕ over B'' is the same as its supremum over the smaller set B^* . To establish this, we start with a lemma that allows us to reduce the number of non-boundary values without decreasing the value of ϕ .

Lemma 6. Let $b \in B''$ be a function with non-boundary values, i.e., $w(b) > 0$. There exists $b' \in B''$ such that $w(b') < w(b)$ and $\phi(b') \geq \phi(b)$ holds.

The proof proceeds as follows: we consider a pair of functions $b_1, b_2 \in B''$, constructed by replacing one non-boundary value in b (for all agents reporting it) with either the value reported by other agents or with a boundary value. By construction, both b_1 and b_2 have smaller w -values than b . To conclude the proof, we show that ϕ for at least one of these modified functions is not smaller than $\phi(b)$.

By repeatedly applying the above lemma, we can iteratively reduce the number of non-boundary values in any $b \in B''$ until $w(b) = 0$, while ensuring that ϕ does not decrease. Thus, we have:

Corollary 2. For any element $b \in B''$, there exists a boundary element $b^* \in B^*$ such that $\phi(b) \leq \phi(b^*)$.

From Corollary 2, we immediately deduce:

$$\sup_{b \in B''} \phi(b) \leq \sup_{b \in B^*} \phi(b).$$

Finally, we bound the supremum of ϕ over B^* by $7/4$.

Lemma 7. For any $b^* \in B^*$, the value of ϕ for b^* is at most $7/4$.

We achieve this by simplifying the representation of boundary functions in B^* and directly estimating the value of ϕ for these functions.

3.5 Proof of the main theorem

We are now ready to prove Theorem 1.

Proof (of Theorem 1). Consider the mechanism $RD + PCD$. By Observation 1, this mechanism is strategyproof. Now, let us show that the approximation ratio of $RD + PCD$ mechanism on all profiles on cycle G is bounded by $7/4$. Let b be any profile on G . If all agents report the same point, then the approximation ratio of the $RD + PCD$ mechanism is equal to 1, as shown in Remark 3. Assume now that there are at least two distinct reports in the profile b . By Observation 1, the $RD + PCD$ mechanism is anonymous and neutral. Hence, by Lemma 1, there exists a normalized profile, \tilde{b} , on the cycle such that $\text{apx}_{RD+PCD}(b) = \text{apx}_{RD+PCD}(\tilde{b})$. By Corollary 1 we have that $\text{apx}_{RD+PCD}(\tilde{b}) \leq \phi(\tilde{b})$. By Corollary 2, there exists a boundary function $b^* \in B^*$ such that $\phi(\tilde{b}) \leq \phi(b^*)$. Lastly, by Lemma 7, we have that $\phi(b^*) \leq 7/4$. Combining the above inequalities we obtain the

desired bound for the approximation ratio of the $RD + PCD$ mechanism for the profile b . Since the profile b was arbitrary, this completes the proof. \square

4 Experiments

In the previous section, we proved an upper bound of $7/4$ on the approximation ratio of the $RD + PCD$ mechanism. On the other hand, in Remark 2, we established that no upper bound below $3/2$ is possible, leaving some uncertainty about the actual approximation ratio of the $RD + PCD$ mechanism. To complement these theoretical results, we conducted numerical experiments to investigate which values of the approximation ratio are actually attained by the $RD + PCD$ mechanism.

4.1 Setup

The idea of the experiments is to compute the approximation ratio of the $RD + PCD$ mechanism for a finite, computationally feasible, set of profiles, providing insight into the behavior of the mechanism for all possible profiles.

The first obvious parameter to consider is the number of agents n , as it directly affects the definition, and potentially the operation of, the $RD + PCD$ mechanism. During the experiments, we considered the number of agents n in the range $[2..60]$.

The second parameter to consider is the set of possible reports for each agent. In the theoretical analysis, no restrictions were imposed on the agents' reports. However, since cycle G is a continuous set with infinitely many points, we introduce such restrictions in the experiments in order to have a finite set of profiles to analyze. For $l \in \mathbb{N}$, let G_l denote the subset of l points on the cycle G (of length 1) that are equally spaced and include the point 0. During the experiments, we restricted agents' reports to the set G_l for $l \in [2..60]$. We hypothesize that profiles with agents' reports restricted to the set G_l provide increasingly better approximation of the behavior of the $RD + PCD$ mechanism for all possible profiles, as l grows (since the allowed reports cover the cycle more densely).

The total number of distinct profiles with n agents whose reports are restricted to G_l is l^n . To reduce the computational complexity, we leveraged the anonymity and neutrality of the $RD + PCD$ mechanism, grouping profiles that would yield the same social cost and avoiding redundant calculations. Despite this, for larger values of n and l , the number of profiles still exceeded 10^7 , making computations infeasible.

In such cases, we restricted our analysis to profiles in which agents reported no more than 3 distinct points. These profiles were considered good candidates for generating the highest approximation ratio for the $RD + PCD$ mechanism, as will be demonstrated later.

4.2 Results

The results of the experiments are presented in Fig. 1. The graph illustrates the maximum value of the approximation ratio of the $RD + PCD$ mechanism for profiles with n agents, whose reports are restricted to G_l . The dotted line represents results computed under the restriction that agents report no more than 3 distinct points.




Figure 1: The maximum value of the approximation ratio of the $RD + PCD$ mechanism in profiles with n agents, whose reports are restricted to G_k . The dotted line highlights the results calculated with the restriction that the number of different points reported by agents is limited to 3.

An immediate observation is that the approximation ratio of the $RD + PCD$ mechanism does not exceed $3/2$ for any profile considered in the experiments. This observation leads to the following hypothesis, which, if true, would further refine the theoretical result established in Theorem 1:

Hypothesis 1. *The approximation ratio of the $RD + PCD$ mechanism is strictly less than $3/2$ for any set of agents and any profile on the cycle G .*

The bound of $3/2$ is the lowest possible due to Observation 2 and would imply that the mechanisms RD and PCD perfectly complement each other, meaning the worst-case profile for one mechanism is the best-case profile for the other.

A deeper analysis of the experimental results provides further evidence supporting Hypothesis 1. Indeed, we observe that for a fixed number of agents, increasing the number of points on which agents can report (i.e., increasing l) beyond a certain threshold does not affect the highest approximation ratio of the $RD + PCD$ mechanism (excluding parity effect). This suggests that the approximation ratio of the $RD + PCD$ mechanism does not grow significantly if we were to extend the range of l .

5 Conclusions

We considered a strategyproof facility location on a cycle with the utilitarian welfare objectives. We showed that the approximation ratio of strategyproof mechanisms for this problem is bounded from above by $7/4$ and is guaranteed by the $RD + PCD$ mechanism. This bound improves the best previously known bound of $2 - 2/n$ in the cases of odd $n \geq 5$. Computational analysis of the approximation ratio of the $RD + PCD$ mechanism suggests that the bound could be further improved to $3/2$. Verifying this hypothesis is an interesting question for future research.

Acknowledgements

This work was supported by the Polish National Science Centre through grant 2021/42/E/HS4/00196.

References

- [Alon *et al.*, 2009] Noga Alon, Michal Feldman, Ariel D Procaccia, and Moshe Tennenholtz. Strategyproof approximation mechanisms for location on networks. *arXiv preprint arXiv:0907.2049*, 2009.
- [Alon *et al.*, 2010] Noga Alon, Michal Feldman, Ariel Procaccia, and Moshe Tennenholtz. Walking in circles. *Discrete Mathematics*, 310(23):3432–3435, 2010.
- [Cai *et al.*, 2016] Qingpeng Cai, Aris Filos-Ratsikas, and Pingzhong Tang. Facility location with minimax envy. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI’16*, pages 137–143. AAAI Press, 2016.
- [Chan *et al.*, 2021] Hau Chan, Aris Filos-Ratsikas, Bo Li, Minming Li, and Chenhao Wang. Mechanism design for facility location problems: A survey. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence*, pages 4356–4365, 2021.
- [Dokow *et al.*, 2012] Elad Dokow, Michal Feldman, Reshef Meir, and Ilan Nehama. Mechanism design on discrete lines and cycles. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 423–440, 2012.
- [Meir, 2019] Reshef Meir. Strategyproof facility location for three agents on a circle. In *International symposium on algorithmic game theory*, pages 18–33. Springer, 2019.
- [Moulin, 1980] Hervé Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.
- [Mu’alem and Schapira, 2018] Ahuva Mu’alem and Michael Schapira. Setting lower bounds on truthfulness. *Games and Economic Behavior*, 110:174–193, 2018.
- [Procaccia and Tennenholtz, 2013] Ariel D Procaccia and Moshe Tennenholtz. Approximate mechanism design without money. *ACM Transactions on Economics and Computation*, 1(4):1–26, 2013.
- [Schummer and Vohra, 2002] James Schummer and Rakesh V Vohra. Strategy-proof location on a network. *Journal of Economic Theory*, 104(2):405–428, 2002.