Mechanism Design for Building Optimal Bridges Between Regions

Zining Qin[2], Hau Chan[3], Chenhao Wang[1,2(✉)], and Ying Zhang[4]

[1] Beijing Normal University, Zhuhai 519087, China
chenhwang@bnu.edu.cn
[2] BNU-HKBU United International College, Zhuhai 519087, China
[3] University of Nebraska-Lincoln, Lincoln, NE 68508, USA
[4] Beijing Electronic Science and Technology Institute, Beijing 100070, China

Abstract. We study the bridge-building problem from the mechanism design perspective. In this problem, a social planner is tasked with building a bridge to connect two regions separated by an obstacle (e.g., a river or valley). Each agent in a region has a private location and is interested in traveling to a facility (e.g., a transportation hub) in the other region. The cost of an agent with respect to a bridge is the distance from their location to the facility of interest via the bridge. Our goal is to design strategy-proof mechanisms that elicit truthful locations from the agents and approximately optimize an objective by determining a location for building a bridge. We consider the social cost and maximum cost objectives, which are the total cost and maximum cost of agents, respectively. For the social cost objective, we characterize an optimal solution and show that it is strategy-proof. For the maximum cost objective, any optimal solution is no longer strategy-proof. We present deterministic $\frac{5}{3}$-approximation and randomized $\frac{3}{2}$-approximation strategy-proof mechanisms. We complement the results by providing tight lower bounds.

1 Introduction

In many urban planning infrastructure projects, a social planner is often tasked with building a bridge to connect two different regions that are in between an obstacle (e.g., a body of water/river, a road, or a valley) [11]. Not surprisingly, building a bridge can lead to many societal benefits. For example, building a bridge over a river can help agents in a region to cross the river to reach the other region more directly. Building a bridge between two mountains can shorten the travel distances of agents compared to using spiral roads on either mountain. Building a viaduct can help to carry a road or railway and reduce the travel distances of agents between different regions. As a result, a bridge would enable agents, at their starting locations, to travel from one region to a point of interest (e.g., an access point, a transit station, or a region center) in another region more efficiently and safely without going through the obstacle directly.

In Fig. 1, we provide an example in which agents are in two regions (modeled simply as a line) separated by a river, which divides the whole region into A

X. Chen and B. Li (Eds.): TAMC 2024, LNCS 14637, pp. 332–343, 2024.
https://doi.org/10.1007/978-981-97-2340-9_28

and B. The agents from Region A need to be connected to a fixed access point in Region B, and agents in Region B need to access the fixed access point in Region A. The access points (referred to as facilities) can be viewed as a hub, such as a center of the community or transit station (determined by the social planner), which the agents can use to reach other places in different regions. The agents' starting locations are alongside the regions (denoted as dots), and the access points are denoted in squares. The social planner would like to improve the distances of the agents from one region to the access point of another region (e.g., from Region A to the access point in Region B) by building a bridge (denoted as a green line in Fig. 1).

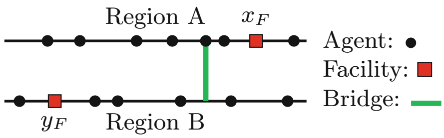

Fig. 1. An example for bridging two regions. The agents are represented as dots. The two access points (facilities) are at x_F and y_F (represented as squares). (Color figure online)

Existing studies have examined building optimal bridges between two different regions as an optimization problem, aiming to minimize the maximum distance between any two points from the two regions (see, e.g., [2,9,10,16,17]). These existing studies have designed polynomial algorithms for building (approximately) optimal bridges between different types of convex polygons (see related work for more details).

Unfortunately, there are two main assumptions that make existing optimization literature not ideal for capturing real-world situations. First, existing literature assumes that all of the points in the regions are the starting locations of some agents. However, in many real-world situations, agents' starting positions consist only of a subset of locations in the regions. Second, existing literature assumes that each agent is required to connect to or access other agents' starting locations in another region. However, agents may not necessarily be concerned about other agents' starting locations and would only be interested in connecting or accessing a given access point (e.g., a shopping mall, region center, or station) Therefore, our goal is to build optimal bridges to account for agents' starting locations and the access points in regions.

1.1 Our Contribution

In this paper, we initiate the mechanism design study of building (approximately) optimal bridges between two regions to connect agents (at their starting locations in their regions) to the corresponding facilities in other regions.

In such a setting, the agent's starting locations in the regions might not be publicly known or visible to the social planner. Therefore, our main goal is to design strategy-proof mechanisms that elicit agents' starting locations truthfully and output bridges that approximately minimize objectives based on the agents' starting locations in their regions to the facilities in other regions.

More specifically, we focus on a basic setting where the two regions are represented by two separate parallel real lines $\mathbb{R}_1$ and $\mathbb{R}_2$ divided by an obstacle (see Fig. 1). Agents' starting locations are points in the real lines (i.e., $x_i \in \mathbb{R}_1$ or $y_j \in \mathbb{R}_2$ for any agent i in $\mathbb{R}_1$ and any agent j in $\mathbb{R}_2$), and the facilities correspond to fixed points (i.e., $x_F \in \mathbb{R}_1$ and $y_F \in \mathbb{R}_2$) in the regions.

We aim to build a bridge that connects the two real lines $\mathbb{R}_1$ and $\mathbb{R}_2$, where one endpoint is in $\mathbb{R}_1$ and the other is in $\mathbb{R}_2$. The bridge must be perpendicular to both $\mathbb{R}_1$ and $\mathbb{R}_2$, and we use a single point s to denote the bridge location. Given a bridge location s, the cost of an agent at $x_i \in \mathbb{R}_1$ is the distance from their starting location to the endpoint of the bridge in $\mathbb{R}_1$, plus the distance from the facility location to the endpoint of the bridge in $\mathbb{R}_2$, i.e., $|x_i - s| + |y_F - s|$.[3] The cost of an agent at $y_j \in \mathbb{R}_2$ is defined similarly. We study two objectives that aim to minimize the social cost and the maximum cost.

For the social cost objective, we characterize an optimal solution and show that an optimal solution is strategy-proof. For the maximum cost objective, we characterize an optimal solution and show that any optimal solution is not strategy-proof.

We provide a deterministic strategy-proof mechanism that has an approximation ratio of $\frac{5}{3}$. We complement this result by providing a tight matching lower bound. We also design a randomized strategy-proof mechanism that has an approximation ratio of $\frac{3}{2}$. We also provide a tight matching lower bound for any randomized strategy-proof mechanisms.

We note that our setting can be reduced to a special setting of Fukui et al. [8], who studied a variant of facility location problems. Fukui et al. [8] proposed a group strategy-proof mechanism that minimizes the social cost, but they did not consider the maximum cost objective. Our characterization of an optimal solution provides a more succinct mechanism result.

1.2 Related Work

The optimal bridge-building problem has been widely studied in the literature [1,2,9,10,16,17]. Cai et al. [2] introduced the problem of adding a line segment to connect two disjoint convex polygonal regions in a plane, such that the length of the longest path from a point in one polygon, passing through the bridge, to a point in another region is minimized. They proposed an $O(n^2 \log n)$-time algorithm, where n is the maximum number of extreme points of the polygons. Later, Bhattacharya and Benkoczi [1] proposed a linear-time algorithm that improves the $O(n^2 \log n)$-time algorithm in [2]. Tan [16] independently presented an alternate linear-time algorithm for the above setting and further generalized it to an $O(n^2)$-time algorithm for bridging two convex polyhedra in space. Kim and Shin

[3] We assume that the bridge has zero cost for the agent using the bridge because it is a constant term that each agent would incur. A positive bridge cost can only help improve our approximation ratios.

[10] provided algorithms to find an optimal bridge between two convex polygons, two simple non-convex polygons, and one convex and one simple non-convex polygons in $O(n)$, $O(n^2)$, and $O(n \log n)$, respectively. Later, Tan [17] provided an $O(n \log^3 n)$-time algorithm for the settings of two simple non-convex polygons. The most related setting is by Kim et al. [9], who proposed a linear-time algorithm to compute an optimal bridge between two parallel lines separated by an obstacle to minimize the length of the longest path connecting two points on the lines. On a related note, there are works that focus on minimizing the diameter or average shortest distances between pairs of nodes of a network using new edges (see, e.g., [5,14]). However, all of the works mentioned focusing on *all* points in the regions. Our work focuses on a subset of points, which are the agents' starting points, and aims to bridge agents to their corresponding facilities in other regions from the mechanism design perspective.

Our work is within the paradigm of approximate mechanism design without money. The paradigm of approximate mechanism design without money is initialized by Procaccia and Tennenholtz [15], who used facility location problems as case studies. This paradigm investigates strategy-proof mechanisms through the lens of the approximation ratio. In a typical setting of facility location problems, the agents report their private locations on the real line to a mechanism. The mechanism determines the locations for building facilities, where the cost of agents is the distance to the facilities. Following this work, variations of facility location problems have been introduced (see, e.g., [6,7,12,13]). We refer readers to a survey on models and results for mechanism design for facility location [3]. The most relevant mechanism design work to ours is by Fukui et al. [8], which considers a more general setting called *pit-stop facility game*, where all agents are in a real line, and each agent i reports an interval $[x_i, y_i]$. A mechanism determines a point $s \in \mathbb{R}$. The cost of agent i is $|s - x_i| + |s - y_i|$. It is easy to see that our setting is the special case when the agents in $\mathbb{R}_1$ has 0 (resp. the agents in $\mathbb{R}_2$ has 1) as an endpoint of their intervals. Fukui et al. [8] proposed a deterministic group strategy-proof mechanism that minimizes the social cost (called *lowest balanced mechanism*), but they did not consider the maximum cost objective. Another most relevant work to ours is the work of [4] in which they considered modifying the structure of regions by adding a shuttle or road to improve the distances of the agents to a prelocated facility in a real line. However, they do not consider two regions separated by an obstacle.

2 Model

There are two parallel real lines, denoted by $\mathbb{R}_1$ and $\mathbb{R}_2$. Assume that $\mathbb{R}_1$ is above $\mathbb{R}_2$, where $\mathbb{R}_1$ and $\mathbb{R}_2$ are regions A and B, respectively, as shown in Fig. 1. A set of $M = \{1, \ldots, m\}$ agents is located in $\mathbb{R}_1$, and a set of $N = \{1, \ldots, n\}$ agents is located in $\mathbb{R}_2$. Each agent $i \in M$ has a location $x_i \in \mathbb{R}_1$, and each agent $j \in N$ has a location $y_j \in \mathbb{R}_2$. Let $\mathbf{x} = (x_1, \ldots, x_m) \in \mathbb{R}_1^m$ be the location profile of the agents in M, and $\mathbf{y} = (y_1, \ldots, y_n) \in \mathbb{R}_2^n$ be the location profile of the agents in N. There are two facilities F_1 and F_2 located at $\mathbb{R}_1$ and $\mathbb{R}_2$, respectively. Both facilities have fixed locations. Let $x_F \in \mathbb{R}_1$ be the location of facility F_1, and $y_F \in \mathbb{R}_2$ be the location of facility F_2.

The agents in M want to access facility F_2, and the agents in N want to access facility F_1. To this end, we aim to build a *bridge* to connect the two real lines $\mathbb{R}_1$ and $\mathbb{R}_2$, where one endpoint of this bridge is in $\mathbb{R}_1$ and the other endpoint is in $\mathbb{R}_2$. The bridge is perpendicular to both $\mathbb{R}_1$ and $\mathbb{R}_2$. Hence, we can use a single point s to denote the location of a bridge, where the two endpoints are $s \in \mathbb{R}_1$ and $s \in \mathbb{R}_2$, respectively. Given a bridge s, the cost of each agent $i \in M$ is the length of the shortest path from this agent to facility F_2,

$$c_1(x_i, s) := |x_i - s| + |s - y_F|.$$

Similarly, the cost of each agent $j \in N$ is defined as

$$c_2(y_j, s) := |y_j - s| + |s - x_F|.$$

A (deterministic) mechanism $f : \mathbb{R}_1^m \times \mathbb{R}_2^n \to \mathbb{R}$ is a function that maps the location profiles $\mathbf{x}, \mathbf{y}$ of agents into a real value as the bridge location. A randomized mechanism is a function f from $R_1^m \times \mathbb{R}_2^n$ to probability distributions over $\mathbb{R}$. If $f(\mathbf{x}, \mathbf{y}) = P$ is a probability distribution, the cost of agent $i \in M$ is defined as the expectation $c_1(x_i, P) = \mathbb{E}_{s \sim P}[c_1(x_i, s)]$, and the cost of $j \in N$ is defined similarly.

A mechanism is *strategy-proof* if no agent can decrease their cost by misreporting their location, where we restrict the misreporting on their own side, that is, an agent in $\mathbb{R}_1$ cannot report a location in $\mathbb{R}_2$, and vice versa.

Definition 1. *A mechanism f is* strategy-proof *if for any location profiles* $\mathbf{x}, \mathbf{y}$*, it satisfies two conditions: (1) $c_1(x_i, f(\mathbf{x}, \mathbf{y})) \leq c_1(x_i, f(x_i', \mathbf{x}_{-i}, \mathbf{y}))$ for any agent $i \in M$ and $x_i' \in \mathbb{R}_1$, where $\mathbf{x}_{-i}$ is the location profile of the agents in $M \setminus \{i\}$, and (2) $c_2(y_j, f(\mathbf{x}, \mathbf{y})) \leq c_2(y_j, f(\mathbf{x}, y_j', \mathbf{y}_{-j}))$ for any agent $j \in N$ and $y_j' \in \mathbb{R}_2$, where $\mathbf{y}_{-j}$ is the location profile of the agents in $N \setminus \{j\}$.*

We study two objective functions, minimizing the social cost, and minimizing the maximum cost. Given profiles $\mathbf{x}, \mathbf{y}$ and bridge location s, the *social cost* is the total cost of all agents $SC(s, \mathbf{x}, \mathbf{y}) = \sum_{i \in M} c_1(x_i, s) + \sum_{j \in N} c_2(y_j, s)$, and the *maximum cost* is the maximum value among the cost of all agents $MC(s, \mathbf{x}, \mathbf{y}) = \max\{\max_{i \in M} c_1(x_i, s), \max_{j \in N} c_2(y_j, s)\}$. We say that a mechanism f has an approximation ratio α or is α-approximation for the objective $\mathcal{S} \in \{SC, MC\}$, if for any instance $(\mathbf{x}, \mathbf{y}, x_F, y_F)$, we have $\frac{\mathcal{S}(f(\mathbf{x},\mathbf{y}),\mathbf{x},\mathbf{y})}{\min_{s \in \mathbb{R}} \mathcal{S}(s,\mathbf{x},\mathbf{y})} \leq \alpha$.

Within the agenda of approximate mechanism design, the goal is to design strategy-proof mechanisms with good approximation ratios. Note that when $x_F = y_F$, a trivial mechanism that always returns x_F as bridge location is clearly strategy-proof and optimal for both objectives (as all agents attain their best possible cost). Hence, we focus on the situation $x_F \neq y_F$, and assume that $x_F = 1$, $y_F = 0$ throughout the remainder of this paper. This assumption is without loss of generality because we can scale the locations in real lines.

As a preliminary result, the following lemma says that an optimal bridge location is between the two facilities (public locations 0 and 1).

Lemma 1. *For both social cost and maximum cost objectives, there exists an optimal solution* $s^* \in [0,1]$.

Proof. Let s be an optimal solution. If $s \in [0,1]$, the lemma is proved. Assume without loss of generality that $s < 0$. Compared with the solution $s^* = 0$, we show that the cost of every agent under s^* is no more than the cost under s. For every agent $i \in M$, we have

$$c_1(x_i, s) = |x_i - s| + |s - 0| \geq |x_i| = c_1(x_i, 0).$$

For every agent $j \in N$, we have

$$c_2(y_j, s) = |y_j - s| + (1 - s) = |y_j - s| + |s| + 1 \geq |y_j| + 1 = c_2(y_j, 0).$$

Therefore, $s^* = 0$ is also an optimal solution, establishing the proof. □

3 Social Cost

In this section, we study the objective of minimizing the social cost $SC(s, \mathbf{x}, \mathbf{y}) = \sum_{i \in M} c_1(x_i, s) + \sum_{j \in N} c_2(y_j, s)$. We first present an intuitive algorithm for computing an optimal solution and then show that it is strategy-proof.

Given location profiles $\mathbf{x}, \mathbf{y}$, for any $s \in [0,1]$, define $M_l(s) = \{i \in M | x_i \leq s\}$ to be the set of agents in M whose locations are on the left of s, and $M_r(s) = \{i \in M | x_i > s\}$ to be the set of agents in M on the right of s. Similarly, define $N_l(s) = \{j \in N | y_j \leq s\}$ and $N_r(s) = \{j \in N | y_j > s\}$. The algorithm moves the bridge location s from 0 to 1 continuously until it cannot decrease the social cost or it reaches 1. During the process of moving s to $s + \epsilon$ for any $\epsilon > 0$, if there is no agent located at interval $(s, s + \epsilon)$, then the cost of any agents in $M_r(s) \cup N_l(s)$ does not change, the cost of the agents in $M_l(s)$ increases by 2ϵ, and the cost of agents in $N_r(s)$ decreases by 2ϵ. The algorithm stops when the number of agents in $M_l(s)$ is no less than that in $N_r(s)$. Formally, the algorithm returns a bridge location

$$b(\mathbf{x}, \mathbf{y}) := \begin{cases} 1, & \text{if } |M_l(1)| < |N_r(1)| \\ \min\{s \in [0,1] \mid |M_l(s)| \geq |N_r(s)|\}, & \text{otherwise.} \end{cases}$$

Theorem 1. *For the social cost, bridge location* $b(\mathbf{x}, \mathbf{y})$ *is optimal. The mechanism that returns* $b(\mathbf{x}, \mathbf{y})$ *is strategy-proof.*

Proof. By Lemma 1, there is an optimal solution $s^* \in [0,1]$. We prove the optimality of $b(\mathbf{x}, \mathbf{y})$ by discussing two cases $s^* < b$ and $s^* > b$ (we write $b(\mathbf{x}, \mathbf{y})$ as b when no confusion arises). If $s^* < b$, by the definition of b, it must be $|M_l(s^*)| < |N_r(s^*)|$. When moving the bridge location from s^* to b, the cost of each agent in $M_r(b) \cup N_l(s^*)$ does not change, the cost of each agent in $M_l(b)$ increases, and the cost of each agent in $N_r(s^*)$ decreases. Let $\epsilon > 0$ be a sufficiently small value so that $s < b - \epsilon$ and no agent lies in interval $(b - \epsilon, b)$. Then, in particular, the cost of each agent in $M_l(b - \epsilon)$ increases by at most

$2(b - s^*)$, and the cost of each agent in $N_r(b-\epsilon) \subseteq N_r(s^*)$ decreases by exactly $2(b - s^*)$. Hence,

$$\begin{aligned} SC(b, \mathbf{x}, \mathbf{y}) &\leq SC(s^*, \mathbf{x}, \mathbf{y}) - 2(b - s^*)|N_r(b-\epsilon)| + 2(b - s^*)|M_l(b-\epsilon)| \\ &< SC(s^*, \mathbf{x}, \mathbf{y}), \end{aligned}$$

where the last inequality comes from the fact that $|M_l(b-\epsilon)| < |N_r(b-\epsilon)|$ by the definition of b.

If $s^* > b$, it must be $|M_l(s^*)| \geq |N_r(s^*)|$. When moving the bridge location from b to s^*, the cost of each agent in $M_r(s^*) \cup N_l(b)$ does not change, the cost of each agent in $M_l(s^*)$ is non-decreasing, and the cost of each agent in $N_r(b)$ decreases. In particular, the cost of each agent in $M_l(b) \subseteq M_l(s^*)$ increases by exactly $2(s^* - b)$, and the cost of each agent in $N_r(b)$ decreases by at most $2(s^* - b)$. Hence, we have

$$\begin{aligned} SC(s^*, \mathbf{x}, \mathbf{y}) &\geq SC(b, \mathbf{x}, \mathbf{y}) - 2(s^* - b)|N_r(b)| + 2(s^* - b)|M_l(b)| \\ &\geq SC(b, \mathbf{x}, \mathbf{y}), \end{aligned}$$

where the last inequality comes from the fact that $|M_l(b)| \geq |N_r(b)|$ by the definition of b. Therefore, b must be an optimal solution for the social cost.

Next, we prove the strategy-proofness. Note that the agents in $M_r(b) \cup N_l(b)$ have no incentive to lie because they attain the best possible cost. The agent $i \in M_l(b)$ with $x_i = b$ also has the best possible cost and, thus, will not misreport. For each agent $i \in M_l(b)$ with $x_i < b$, the only way to change the solution is to misreport a location $x'_i > b$, which leads to a solution $b(x'_i, \mathbf{x}_{-i}, \mathbf{y}) \geq b$; however, this can only increase the cost of agent i. For each agent $j \in N_r(b)$, the only way to change the solution is to misreport a location $y'_j < b$, which leads to a solution $b(\mathbf{x}, y'_j, \mathbf{y}_{-j}) \leq b$ and cannot bring any benefit. Therefore, no agent has an incentive to misreport. □

4 Maximum Cost

In this section, we consider the objective of minimizing the maximum cost. Before presenting the mechanisms, we first characterize the optimal maximum cost. Given location profiles $\mathbf{x}, \mathbf{y}$, define $x_l = \min_{i \in M} x_i$, $x_r = \max_{i \in M} x_i$, $y_l = \min_{j \in N} y_j$, and $y_r = \max_{j \in N} y_j$. The maximum cost in any solution must be attained by one of the four extreme agents. Intuitively, a good bridge location should balance the cost of these extreme agents. Recall from Lemma 1 that there is an optimal solution in $[0, 1]$. Indeed, the bridge location in $[0, 1]$ that minimizes the difference between the cost of x_l and y_r is optimal, that is,

$$\arg\min_{w \in [0,1]} |c_1(x_l, w) - c_2(y_r, w)| = ||x_l - w| - |y_r - w| - 1 + 2w|.$$

Proposition 1. *For the maximum cost, $s^* \in \arg\min_{w \in [0,1]} |c_1(x_l, w) - c_2(y_r, w)|$ is an optimal solution.*

Proof. Let $s \in [0,1]$ be an optimal solution. If $c_1(x_l, s) = c_2(y_r, s)$, the proof is done. Assume without loss of generality that $c_1(x_l, s) > c_2(y_r, s)$. Note that for any solution $s' > s$, the difference $|c_1(x_l, s') - c_2(y_r, s')|$ is no less than $c_1(x_l, s) - c_2(y_r, s)$. If the optimal solution s does not minimize the difference between the cost of x_l and y_r, then there is a solution s' with $0 \le s' < s$ that minimizes the difference. Indeed, compared with s, the cost of the agent at x_l under s' is non-increasing, i.e., $c_1(x_l, s') \le c_1(x_l, s)$, and the cost of the agent at y_r is non-decreasing, i.e., $c_2(y_r, s') \le c_2(y_r, s)$. It follows that the difference between these two agents is non-increasing. Further, while the cost of the agent at x_r is non-increasing, the cost of the agent at y_l may be increasing. If it increases, then it must be $0 \le y_l \le y_r$, and y_l cannot be responsible for the maximum cost. Hence, the maximum cost is non-increasing, and the solution s' is optimal. □

Define a rounding function $r : \mathbb{R} \to [0,1]$ that maps any real number $x \in \mathbb{R}$ to the nearest point in $[0,1]$ from it, i.e.,

$$r(x) = \arg \min_{w \in [0,1]} |x - w|.$$

For example, $r(1.5) = 1, r(-2) = 0$ and $r(0.3) = 0.3$. We call $r(x_l)$ and $r(y_r)$ the *rounding extremes*, and as a corollary of Proposition 1, there exists an optimal solution between them.

Corollary 1. *For the maximum cost, there exists an optimal solution that lies in* $[r(x_l), r(y_r)]$ *or* $[r(y_r), r(x_l)]$.

Proof. If an optimal solution $s^* \in [0,1]$ is larger than both $r(x_l)$ and $r(y_r)$, then we can move it to the location $\max\{r(x_l), r(y_r)\}$, in which no agent would increase their cost. Thus, $\max\{r(x_l), r(y_r)\}$ is also optimal. Similarly, if s^* is smaller than $r(x_l)$ and $r(y_r)$, the location $\min\{r(x_l), r(y_r)\}$ is optimal. □

In contrast to the social cost objective that admits an optimal strategy-proof mechanism, the mechanism that returns an optimal solution for maximum cost is not strategy-proof. Consider the instance with $\mathbf{x} = (0)$ and $\mathbf{y} = (1)$. The unique optimal solution is $\frac{1}{2}$, where both agents have a cost equal to 1. Now, suppose that the agent in N misreports the location as $y_1' = 3$. Then the mechanism takes $\mathbf{x} = (0)$ and $\mathbf{y}' = (3)$ as input, and outputs the unique optimal solution 1. After misreporting, this agent with true location $y_1 = 1$ decreases the cost to 0.

4.1 Deterministic Strategy-Proof Mechanisms

In this subsection, we provide upper and lower bounds on the approximation ratio of deterministic strategy-proof mechanisms for the maximum cost objective. Intuitively, based on Proposition 1, a good solution should balance the cost of the two extremes x_l and y_r. However, to guarantee the strategy-proofness, a mechanism cannot always achieve such a balance perfectly. Instead, we consider specific locations between $r(x_l)$ and $r(y_r)$ by Corollary 1. The following deterministic mechanism returns the bridge location $\frac{1}{2}$ if it lies between the two rounding extremes $r(x_l)$ and $r(y_r)$ and returns the rounding extreme that is closer to $\frac{1}{2}$ otherwise.

Mechanism 1. *Given location profiles* $\mathbf{x}, \mathbf{y}$*, if* $r(x_l) \le \frac{1}{2} \le r(y_r)$ *or* $r(y_r) \le \frac{1}{2} \le r(x_l)$*, then return* $\frac{1}{2}$*. Otherwise, return* $\arg\min_{s \in \{r(x_l), r(y_r)\}} |s - \frac{1}{2}|$.

Theorem 2. *Mechanism 1 is strategy-proof and* $\frac{5}{3}$*-approximation for the maximum cost.*

Proof. We first prove the strategy-proofness. If $r(x_l) \le \frac{1}{2} \le r(y_r)$ and the mechanism returns $\frac{1}{2}$, the misreports from the agents in M can only lead to a situation where both rounding extremes are on the right of $\frac{1}{2}$, and the outcome solution is larger than $\frac{1}{2}$. Thus, the agents in M cannot decrease their cost. Similarly, the misreports from the agents in N can only lead to an outcome smaller than $\frac{1}{2}$, and these agents cannot decrease their cost. If $r(x_l) \ge \frac{1}{2} \ge r(y_r)$, all agents have their best possible cost under the solution $\frac{1}{2}$, and thus they have no incentive to misreport. When both rounding extremes $r(x_l)$ and $r(y_r)$ are larger than $\frac{1}{2}$, if $r(x_l) \ge r(y_r) \ge \frac{1}{2}$, all agents have their best possible cost. If $r(y_r) \ge r(x_l) \ge \frac{1}{2}$, at the solution $r(x_l)$, all agents in M and those agents in N located on the left of $r(x_l)$ already achieve their best possible cost, and only the agents in N located on the right of $r(x_l)$ have the potential incentive to misreport. However, their misreporting cannot benefit them. Finally, the symmetric argument/analysis works for the case when both $r(x_l)$ and $r(y_r)$ are smaller than $\frac{1}{2}$.

Next, we prove the approximation ratio. When $r(y_r) \le \frac{1}{2} \le r(x_l)$, it must be that $y_r \le \frac{1}{2} \le x_l$, and the mechanism outputs $\frac{1}{2}$. In this case, all agents achieve their best possible cost, and thus the solution $\frac{1}{2}$ is optimal.

When $r(x_l) \le \frac{1}{2} \le r(y_r)$, it must be that $x_l \le \frac{1}{2} \le y_r$, and the mechanism outputs $\frac{1}{2}$. If the maximum cost $MC(\frac{1}{2}, \mathbf{x}, \mathbf{y})$ is attained by the agents at x_r or y_l, they already achieve the best possible cost, and an optimal solution should have a maximum cost equal to $MC(\frac{1}{2}, \mathbf{x}, \mathbf{y})$. Hence, we only need to focus on the case when the maximum cost is attained by the agents at x_l or y_r, which is

$$\begin{aligned} MC(\frac{1}{2}, \mathbf{x}, \mathbf{y}) &= \max\{c_1(x_l, \frac{1}{2}), c_2(y_r, \frac{1}{2})\} = \max\{\frac{1}{2} - x_l + \frac{1}{2}, y_r - \frac{1}{2} + (1 - \frac{1}{2})\} \\ &= \max\{1 - x_l, y_r\}. \end{aligned}$$

Since an optimal solution s^* lies in interval $[r(x_l), r(y_r)]$ by Corollary 1, the optimal maximum cost is

$$\begin{aligned} MC(s^*, \mathbf{x}, \mathbf{y}) &\ge \max\{c_1(x_l, s^*), c_2(y_r, s^*), \frac{c_1(x_l, s^*) + c_2(y_r, s^*)}{2}\} \\ &= \max\{c_1(x_l, s^*), c_2(y_r, s^*), \frac{(2s^* - x_l) + (y_r - s^* + 1 - s^*)}{2}\} \\ &\ge \max\{-x_l, y_r - 1, \frac{y_r + 1 - x_l}{2}\}. \end{aligned}$$

Therefore, the approximation ratio is

$$\frac{MC(\frac{1}{2}, \mathbf{x}, \mathbf{y})}{MC(s^*, \mathbf{x}, \mathbf{y})} \le \frac{\max\{1 - x_l, y_r\}}{\max\{-x_l, y_r - 1, \frac{y_r + 1 - x_l}{2}\}}.$$

Taking into account the constraint that $x_l \leq \frac{1}{2} \leq y_r$, it is easy to verify that this ratio is no more than $\frac{5}{3}$, as desired. It equals $\frac{5}{3}$ when $(x_l, y_r) = (-\frac{3}{2}, \frac{1}{2})$ or $(x_l, y_r) = (\frac{1}{2}, \frac{5}{2})$.

When $r(x_l), r(y_r) \leq \frac{1}{2}$ or $r(x_l), r(y_r) \geq \frac{1}{2}$, by symmetry we can only consider the case when $r(x_l), r(y_r) \leq \frac{1}{2}$. If $r(x_l) \geq r(y_r)$, then the mechanism will output $r(x_l)$, in which all agents achieve their best possible cost. If $r(x_l) \leq r(y_r)$, the mechanism returns $r(y_r)$. All agents in N, and the agents in M who are located on the right of $r(y_r)$, already achieve their best possible cost. Hence, we only need to consider the case when the induced maximum cost of $r(y_r)$ is attained by x_l, that is, $MC(r(y_r), \mathbf{x}, \mathbf{y}) = c_1(x_l, y_r) = 2y_r - x_l$. The optimal solution s^* has a maximum cost

$$\begin{aligned} MC(s^*, \mathbf{x}, \mathbf{y}) &\geq \max\{c_1(x_l, s^*), \frac{c_1(x_l, s^*) + c_2(y_r, s^*)}{2}\} \\ &= \max\{c_1(x_l, s^*), \frac{|s^* - x_l| + |s^*| + |y_r - s^*| + |1 - s^*|}{2}\} \\ &\geq \max\{|x_l|, \frac{y_r - x_l + 1}{2}\}. \end{aligned}$$

Therefore, the approximation ratio is

$$\frac{MC(\frac{1}{2}, \mathbf{x}, \mathbf{y})}{MC(s^*, \mathbf{x}, \mathbf{y})} \leq \frac{2y_r - x_l}{\max\{|x_l|, \frac{y_r - x_l + 1}{2}\}}.$$

Taking into account the constraint that $x_l \leq y_r \leq \frac{1}{2}$, it is easy to verify that this ratio is no more than $\frac{5}{3}$. It equals $\frac{5}{3}$ when $(x_l, y_r) = (-\frac{3}{2}, \frac{1}{2})$. □

The following theorem provides a matching lower bound $\frac{5}{3}$ for deterministic mechanisms, indicating that Mechanism 1 has tight approximation ratio.

Theorem 3. *For the maximum cost objective, no deterministic strategy-proof mechanism has an approximation ratio less than* $\frac{5}{3}$.

4.2 Randomized Strategy-Proof Mechanisms

While the bounds in Sect. 4.1 for deterministic mechanisms are tight, in this section, we consider randomized mechanisms and also derive tight bounds. Inspired by Corollary 1 that an optimal solution lies between $r(x_l)$ and $r(y_r)$, the following randomized mechanism returns a random point between them. Precisely, it outputs $r(x_l)$, $r(y_r)$, and the point that is between them and is closest to $\frac{1}{2}$, with specified probabilities.

Mechanism 2. *Given location profiles* $\mathbf{x}, \mathbf{y}$, *with probability* $\frac{1}{4}$, *return* $r(x_l)$; *with probability* $\frac{1}{4}$, *return* $r(y_r)$; *with probability* $\frac{1}{2}$, *return*

$$mid(x_l, y_r) = \begin{cases} \frac{1}{2}, & if x_l \leq \frac{1}{2} \leq y_r or x_l \geq \frac{1}{2} \geq y_r \\ \arg\min_{s \in \{r(x_l), r(y_r)\}} |s - \frac{1}{2}|, & otherwise \end{cases}$$

The point $mid(x_l, y_r)$ coincides with $r(x_l)$ or $r(y_r)$ when both x_l and y_r are larger than $\frac{1}{2}$ or both are smaller than $\frac{1}{2}$. For example, when $x_l = 1.2$ and $y_r = 0.6$, the mechanism returns $r(x_l) = 1$ with probability $\frac{1}{4}$, and returns $r(y_r) = mid(x_l, y_r) = 0.6$ with probability $\frac{1}{4} + \frac{1}{2} = \frac{3}{4}$.

Lemma 2. *Mechanism 2 is strategy-proof.*

Proof. We show that any agent i in M cannot decrease the cost by misreporting. The analysis for the agents in N is the same. Clearly, agent i cannot change y_r. When the location of agent i is $x_i > x_l$, the only way to change the solution is to misreport a location $x_i' < x_l$. Then the realizations of the random point returned by the mechanism become $r(x_i')$, $r(y_r)$ and $mid(x_i', y_r)$. Since $x_i > x_l$ and $0 \leq r(x_i') \leq r(x_l)$, we have $c_1(x_i, r(x_i')) = c_1(x_i, r(x_l)) = x_i$. On the other hand, if $x_i \geq mid(x_l, y_r)$, then we have $cost(x_i, mid(x_l, y_r)) = cost(x_i, mid(x_i', y_r)) = x_i$. If $x_i < mid(x_l, y_r)$, then it must be $x_l < x_i < mid(x_l, y_r) \leq r(y_r)$, and the misreport of agent i would not change $mid(x_l, y_r)$, that is, $mid(x_i', y_r) = mid(x_l, y_r) = \min\{\frac{1}{2}, r(y_r)\}$. Hence, agent i cannot benefit from misreporting.

When agent i's location is $x_i = x_l$, if agent i misreports a location $x_i' < x_l$, the analysis is the same as above. If agent i misreports $x_i' > x_l$, as both functions $r(\cdot)$ and $mid(\cdot, \cdot)$ are non-decreasing with x_l, it cannot decrease the cost. □

We remark that the above proof for the strategy-proofness does not rely on the probabilities of the mechanism. Therefore, any constant probabilities assigned to the candidate locations $r(x_l), r(y_r), mid(x_l, y_r)$ can induce a strategy-proof mechanism. Nevertheless, we show that the probabilities specified in Mechanism 2 lead to the best possible approximation ratio.

Theorem 4. *Mechanism 2 is a randomized strategy-proof mechanism and is $\frac{3}{2}$-approximation for the maximum cost.*

Theorem 5. *For the maximum cost objective, no randomized strategy-proof mechanism has an approximation ratio less than $\frac{3}{2}$.*

5 Conclusion

We studied a novel mechanism design setting for building a bridge to connect two regions separated by an obstacle under the social and maximum cost objectives. For both objectives, we characterized their optimal solutions. While any optimal solution for the social cost objective is strategy-proof, it is not strategy-proof for the maximum cost objective. Therefore, for the maximum cost objective, we provided a deterministic $\frac{5}{3}$-approximation mechanism and a randomized $\frac{3}{2}$-approximation mechanism. Furthermore, we derived tight lower bounds, showing that no strategy-proof mechanisms can have better approximation ratios.

For the future directions, we note that our model is just a starting point of the mechanism design for bridge-building problems. While we model the two regions as two real lines, the regions could be a network, Euclidean plane, or other metric spaces. It is interesting to study how the existing and new methods can be applied to these regions. More generally, we can also further consider building multiple non-perpendicular bridges among multiple regions.

References

1. Bhattacharya, B., Benkoczi, R.: On computing the optimal bridge between two convex polygons. Inf. Process. Lett. **79**(5), 215–221 (2001)
2. Cai, L., Xu, Y., Zhu, B.: Computing the optimal bridge between two convex polygons. Inf. Process. Lett. **69**(3), 127–130 (1999)
3. Chan, H., Filos-Ratsikas, A., Li, B., Li, M., Wang, C.: Mechanism design for facility location problems: a survey. In: Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI), pp. 4356–4365 (2021)
4. Chan, H., Wang, C.: Mechanism design for improving accessibility to public facilities. In: Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS), pp. 2116–2124 (2023)
5. Demaine, E.D., Zadimoghaddam, M.: Minimizing the diameter of a network using shortcut edges. In: Kaplan, H. (ed.) SWAT 2010. LNCS, vol. 6139, pp. 420–431. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13731-0_39
6. Dokow, E., Feldman, M., Meir, R., Nehama, I.: Mechanism design on discrete lines and cycles. In: Proceedings of the 13th ACM Conference on Electronic Commerce (EC), pp. 423–440 (2012)
7. Filos-Ratsikas, A., Voudouris, A.A.: Approximate mechanism design for distributed facility location. In: Caragiannis, I., Hansen, K.A. (eds.) SAGT 2021. LNCS, vol. 12885, pp. 49–63. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-85947-3_4
8. Fukui, Y., Shurbevski, A., Nagamochi, H.: Group strategy-proof mechanisms for shuttle facility games. J. Inf. Process. **28**, 976–986 (2020)
9. Kim, B.J., Shin, C.S., Chwa, K.Y.: Linear algorithms for computing a variant segment center. In: Proceedings of Korean Information Science Society Conference (KISS), vol. A, pp. 708–710 (1998). (in Korean)
10. Kim, S.K., Shin, C.S.: Computing the optimal bridge between two polygons. Theory Comput. Syst. **34**(4), 337–352 (2001)
11. McCartney, G., Whyte, B., Livingston, M., Crawford, F.: Building a bridge, transport infrastructure and population characteristics: explaining active travel into glasgow. Transp. Policy **21**, 119–125 (2012)
12. Mei, L., Li, M., Ye, D., Zhang, G.: Facility location games with distinct desires. Discret. Appl. Math. **264**, 148–160 (2019)
13. Meir, R.: Strategyproof facility location for three agents on a circle. In: Fotakis, D., Markakis, E. (eds.) SAGT 2019. LNCS, vol. 11801, pp. 18–33. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-30473-7_2
14. Meyerson, A., Tagiku, B.: Minimizing average shortest path distances via shortcut edge addition. In: Proceedings of the 12th International Workshop and 13th International Workshop on Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM), vol. 5687, pp. 272–285 (2009)
15. Procaccia, A.D., Tennenholtz, M.: Approximate mechanism design without money. ACM Trans. Econ. Comput. (TEAC) **1**(4), 1–26 (2013)
16. Tan, X.: On optimal bridges between two convex regions. Inf. Process. Lett. **76**(4–6), 163–168 (2000)
17. Tan, X.: Finding an optimal bridge between two polygons. Int. J. Comput. Geom. Appl. **12**(03), 249–261 (2002)