

Truthful Interval Covering

Argyrios Deligkas¹, Aris Filos-Ratsikas² and Alexandros A. Voudouris³

¹Royal Holloway University of London, UK

²University of Edinburgh, UK

³University of Essex, UK

argyrios.deligkas@rhul.ac.uk, aris.filos-ratsikas@ed.ac.uk, alexandros.voudouris@essex.ac.uk

Abstract

We initiate the study of a novel problem in mechanism design without money, which we term *Truthful Interval Covering* (TIC). An instance of TIC consists of a set of agents each associated with an individual interval on a line, and the objective is to decide where to place a *covering interval* to minimize the total social or egalitarian cost of the agents, which is determined by the intersection of this interval with their individual ones. This fundamental problem can model situations of provisioning a public good, such as the use of power generators to prevent or mitigate load shedding in developing countries. In the strategic version of the problem, the agents wish to minimize their individual costs, and might misreport the position and/or length of their intervals to achieve that. Our goal is to design *truthful* mechanisms to prevent such strategic misreports and achieve good approximations to the best possible social or egalitarian cost. We consider the fundamental setting of known intervals with equal lengths and provide tight bounds on the approximation ratios achieved by truthful deterministic mechanisms. For the social cost, we also design a randomized truthful mechanism that outperforms all possible deterministic ones. Finally, we highlight a plethora of natural extensions of our model for future work, as well as some natural limitations of those settings.

1 Introduction

We introduce the *Truthful Interval Covering* (TIC) problem, a novel problem in the field of mechanism design without money [Procaccia and Tennenholtz, 2013]. In this problem, there is a set N of n agents, each of whom is associated with an interval I_i on the line of real numbers. There is also a *covering interval* C , which should be placed somewhere on the line. The cost of agent $i \in N$ is a function of the portion of I_i that C covers; in the simplest version of the problem, the cost is just the part of I_i that is not covered by C . The goal is to place the interval so as to minimize the *social cost* (total cost of the agents) or the *max cost* (maximum individual agent cost), while taking the incentives of the agents into account.

Indeed, agents might misreport information about their intervals (e.g., their position or length) if that would lead to an outcome that is preferable for them.

The TIC problem captures applications in which a public good is provisioned and shared among a set of participants. We provide a few indicative of many examples below.

- The covering interval could represent the time interval during which a power generator can be operated, and the individual intervals capture the times during which each citizen would like to have access to electricity. The minimum-social cost solution is one that covers as much demand for electricity as possible. This is particularly relevant in developing countries where electricity might be a scarce resource, and can be used to prevent or mitigate the effects of load shedding.¹
- The covering interval could correspond to the range of a public WiFi hotspot to be placed in an area with low broadband connectivity, when the agents' intervals are the signal ranges of their devices.
- The covering interval could capture the time in which to schedule a university open-day or a job fair, given the preferences of the potential attendees over the different time intervals in the day.
- The covering interval could be an express public transportation line connecting parts of a city or intercity network, and the agents express which parts of the route they would like this service to cover.

Despite its fundamental nature, and its resemblance to other classic algorithmic problems like the interval scheduling problem and its variants [Kolen *et al.*, 2007], the interval covering problem has seemingly not been studied systematically from a purely algorithmic point of view. This can likely be attributed to the fact that the optimal covering can be found in polynomial time via a rather simple algorithm (see Theorem 2 in Section 2). Once we move to a mechanism design regime however, where the incentives of the agents for misreporting come into effect, the problem becomes much more challenging. *Truthful* mechanisms, which eliminate those incentives, are necessarily suboptimal, and resort to approximations. Our goal is to design truthful mechanisms that achieve

¹E.g., see https://en.wikipedia.org/wiki/South_African_energy_crisis

approximations that are as small as possible, and identify the limitations of such mechanisms via appropriate inapproximability results.

1.1 Our Contribution

In this paper, we introduce the Truthful Interval Covering (TIC) problem as a novel and interesting problem in mechanism design without money. Our technical contribution is as follows.

- We provide upper and lower bounds on the approximation ratio of truthful mechanisms for the most fundamental version of the problem, where all of the interval lengths are *known* and *equal*, which already turns out to be quite challenging. We start with the social cost objective and deterministic truthful mechanisms, for which, in Section 3, we prove a *tight bound* of $2 - 2/n$ on the approximation ratio. In Section 4, we present a simple *randomized, universally truthful* mechanism that achieves an approximation ratio of $5/3$, thus outperforming all deterministic ones. In Section 5, we turn our attention to the max cost objective, for which we show a tight approximation ratio of 2 for deterministic mechanisms, and a lower bound of 2 for a natural class of randomized mechanisms, thus showing that randomization might not be able to lead to improvements for this objective.
- We also consider two natural extensions of the main model in Section 6. In the first one, the agent intervals are assumed to be unknown, and thus the agents can misreport their starting positions as well as the lengths of their intervals. We show that no truthful mechanism can achieve a meaningful approximation ratio in terms of both the social and the max cost. In the second extension, we consider the case where the interval lengths are known but might be unequal. We show that a simple mechanism, which places the covering interval at the starting position of the agent with the maximum-length interval is truthful and achieves a linear approximation ratio in terms of the social cost, and an approximation ratio of at most 2 for the max cost; the latter is best possible when the interval lengths are known (any might be equal or unequal).

In Section 7, we present and discuss several other interesting variants of the main model, which capture a wealth of different possible application domains. We believe that there is great potential for follow-up work, and the problem could enjoy similar success as other problems within the research agenda of mechanism design without money, such as truthful facility location [Chan *et al.*, 2021; Procaccia and Tennenholtz, 2013], truthful resource allocation [Krysta *et al.*, 2014; Filos-Ratsikas *et al.*, 2014; Abebe *et al.*, 2020], or impartial selection [Alon *et al.*, 2011; Fischer and Klimm, 2014; Bjelde *et al.*, 2017]. Due to space constraints, the proofs of some statement are omitted.

1.2 Related Work and Discussion

The research agenda of *approximate mechanism design without money* was put forward by Procaccia and Tennenholtz [2013] and aims to capture settings involving selfish

participants, in which truthful mechanisms are used to optimize a social objective. These mechanisms are compared, via their approximation ratio, against the performance of the best-possible outcome, which would be achievable if the participants were not selfish. The prototypical problem in this field is that of truthful facility location, which has flourished into an extremely fruitful research area, giving rise to a plethora of works on several different variants; see the survey of Chan *et al.* [2021] for details, and the related work discussion in [Procaccia and Tennenholtz, 2013] for earlier references.

Our setting is markedly different from facility location, where the cost of an agent depends on the *distance* from the location of the facility. In contrast, in our case, the cost of an agent is a function of how much her associated interval is covered. Still, there are some conceptual similarities between the two problems, namely in terms of the truthful mechanisms employed to achieve the approximation guarantees. In particular, similarly to the literature of facility location, we also employ mechanisms that are based around k -th ordered statistics (e.g., the median) of the agents' reports. These are in fact not particular to facility location, but more generally centered around the concept of single-peaked preferences [Black, 1948; Moulin, 1980]. Despite this superficial connection, the proofs for the performance of these mechanisms are very much different in the TIC problem; it is worth mentioning that, contrary to the setting of [Procaccia and Tennenholtz, 2013], in TIC these mechanisms provably do not admit social-cost minimizing outcomes.

Another related problem is that of strategyproof activity scheduling studied by Xu *et al.* [2020], in which an activity (represented by an interval) is to be placed on a line based on the preferences of self-interested agents. Despite the superficial similarity, this setting is again notably different from ours; in activity scheduling, each agent reports a single point and her cost is the distance from the closest endpoint of the activity interval. This makes the problem much closer to facility location rather than our covering problem. The work of Bei *et al.* [2022] on truthful cake sharing is also related to our paper. In contrast to our model, where the agents report their intervals on the line and have costs for the chosen covering interval, in the cake cutting model of Bei *et al.*, the agents report piecewise uniform utilities over a cake (represented as a fixed-size interval) and the objective is to choose an interval of certain length that they will all share.

From a purely algorithmic point of view (without incentives), problems related to intervals (like scheduling or coloring) are rather fundamental and included in most textbooks on algorithms, e.g., see [Kleinberg and Tardos, 2006; Roughgarden, 2022]. As we mentioned earlier, the algorithmic variant of TIC admits an easy polynomial-time algorithm, and the problem becomes challenging once studied under the mechanism design regime.

Finally, we remark that the term “truthful interval cover” has been used before in the literature for mechanisms with money for solving a crowdsourcing problem, where agents bid for intervals of tasks that they are willing to get; see [Dayama *et al.*, 2015; Markakis *et al.*, 2022]. This model is completely different compared to the one we study here, and, thus, we do not expect any ambiguity to arise.

2 The Setting

In the Truthful Interval Covering (TIC) problem, there is a set N of n agents, each associated with an interval $I_i = [s_i, t_i]$ on the line of real numbers. There is also a covering interval $C = [s, t]$ whose position needs to be determined. We focus on the most fundamental version of the problem, in which the interval lengths are all known and equal, i.e., $|I_i| = |I_j| = |C|$ for any agents $i, j \in N$. Given this, we can assume without loss of generality that $|I_i| = 1$ for any $i \in N$ and $|C| = 1$. Let $\mathcal{I} = (I_1, I_2, \dots, I_n)$ be the vector of the agents' intervals, to which we refer to as an *instance*. Without loss of generality, we assume that for any two agents $i, j \in N$ with $i < j$, $s_i \leq s_j$, i.e., the intervals in \mathcal{I} appear in non-decreasing order of left endpoints. Using this, we may refer to an agent i being before or after another agent if $i < j$ or $i > j$, respectively. We will also say that an agent i is before or after a point x if $s_i \leq x$ or $s_i > x$.

Given a position for the covering interval C , the cost of an agent $i \in N$ is the part of her interval I_i that does not overlap with C , i.e., $\text{cost}_i(C) = |I_i \setminus (I_i \cap C)| = 1 - |I_i \cap C|$. The *social cost* of C is the total cost of all agents:

$$\text{SC}(C) = \sum_i \text{cost}_i(C).$$

The *max cost* of C is the maximum cost over all agents:

$$\text{MC}(C) = \max_i \text{cost}_i(C).$$

When the objective (social or max cost) is clear from context, for any instance \mathcal{I} , we will use $O(\mathcal{I})$ to denote the covering interval that minimizes the objective for instance for \mathcal{I} ; when \mathcal{I} is also clear from context, we will simply write O .

A deterministic *mechanism* \mathcal{M} takes as input an instance \mathcal{I} and outputs the position of a covering interval $\mathcal{M}(\mathcal{I})$, i.e., the position $s \in \mathbb{R}$ of the left endpoint of the covering interval. We also consider *randomized* mechanisms, which, instead of a single position, output a probability distribution $D_{\mathcal{M}}(\mathcal{I})$ over possible positions of the covering interval.

The *approximation ratio* of \mathcal{M} in terms of an objective $f \in \{\text{SC}, \text{MC}\}$ is the worst-case ratio of the objective value of the covering interval computed by \mathcal{M} over the smallest possible objective value achieved by any covering interval, over all instances of the problem:

$$\sup_{\mathcal{I}} \frac{f(\mathcal{M}(\mathcal{I}))}{\min_C f(C)}.$$

For randomized mechanisms, the definition is very similar, with the only difference that the *expected objective value* $\mathbb{E}_{\mathcal{M}(\mathcal{I}) \sim D_{\mathcal{M}}(\mathcal{I})}[f(\mathcal{M}(\mathcal{I}))]$ appears in the numerator.

The term ‘‘Truthful’’ in the name of the TIC problem comes from the fact that the information about the intervals is not public knowledge, but has to be elicited from the agents. The agents are self-interested entities who might misreport this information if that results in them achieving a smaller cost. In the setting we consider, since the interval lengths are all 1, the elicited information is the position of the interval of each agent $i \in N$, i.e., the left endpoint $s_i \in \mathbb{R}$ of I_i . For simplicity, we say that each agent reports her interval I_i rather than s_i .

A mechanism \mathcal{M} is said to be *truthful* if does not incentivize the agents to misreport their intervals, that is, for every agent i and every possible interval I'_i that the agent could report,

$$\text{cost}_i(\mathcal{M}(\mathcal{I})) \leq \text{cost}_i(\mathcal{M}(I'_i, \mathcal{I}_{-i})) \quad (1)$$

where \mathcal{I}_{-i} is the vector \mathcal{I} without the i -th coordinate.

For randomized mechanisms, the definition of truthfulness extends to *truthfulness in expectation*, which stipulates that no agent can decrease her expected cost by deviating. In our positive results, we will actually use a stronger truthfulness guarantee called *universal truthfulness*. A mechanism is universally truthful if it is truthful for any realization of truthfulness, i.e., Inequality (1) holds for any $\mathcal{M}(\mathcal{I}) \sim D_{\mathcal{M}}(\mathcal{I})$.

Our goal is to design truthful mechanisms (either deterministic truthful or universally truthful) with approximation ratios as close to 1 as possible. To achieve this, we focus on a class of mechanisms called *k-ordered statistics*.

Definition 1 (*k-ordered statistic*). For $k \in [n]$, the *k-ordered statistic mechanism* \mathcal{M} outputs the interval reported by the k -th ordered agent i in instance \mathcal{I} , i.e., $\mathcal{M}(\mathcal{I}) = I_i$.

For example, for $k = \lfloor n/2 \rfloor$, the *k-ordered statistic mechanism* outputs exactly the interval reported by the median agent. *k-ordered statistic mechanisms* (as well as their convex combinations) are well-known to be truthful in other contexts, e.g., see [Dummett and Farquharson, 1961; Moulin, 1980]. For similar reasons, any *k-ordered statistic mechanism* is truthful in our setting.

Theorem 1. *For any $k \in [n]$, the k -ordered statistic mechanism is truthful. Furthermore, any convex combination over k -ordered statistic mechanisms is universally truthful.*

Before we proceed with the design of truthful mechanisms, we state and prove the following statement, which establishes that the purely algorithmic version of the problem, without any regard to agent incentives, can be solved in polynomial time with respect to the social cost and the max cost. We refer to this problem as the INTERVAL COVERING PROBLEM.

Theorem 2. *The social cost-minimizing and the max cost-minimizing positions for the covering interval in the INTERVAL COVERING PROBLEM can be computed in linear time.*

3 Social Cost: Deterministic Mechanisms

We start by showing bounds on the approximation ratio of deterministic truthful mechanisms for the social cost. As mentioned in Section 1.1, for this case we obtain a tight bound of $2 - 2/n$ on the approximation ratio achievable by any such mechanism. The mechanism that achieves this bound is the MEDIAN mechanism, the *k-ordered statistic mechanism* with $k = \lfloor n/2 \rfloor$ (see Definition 1). Similar mechanisms (that choose the reported action of the median agents) have played a prominent role in other domains in mechanism design without money [Chan *et al.*, 2021]. However, as we mentioned earlier, the nature of our problem is different from those, and hence the proof is also rather different.

Theorem 3. *The MEDIAN mechanism achieves an approximation ratio of $2 - \frac{2}{n}$ for the TIC problem.*

Proof. Consider an arbitrary instance. Let m be the median agent so that $I_m = [s_m, t_m]$ is the unit-size interval that is chosen by the mechanism. Let $O = [s_o, t_o]$ be the optimal unit-size interval. Without loss of generality, we can assume that $s_o \geq s_m$ and that n is even. Let $x \in [0, 1]$ be the length of the intersection $I_m \cap O = [s_o, t_m]$ between the interval chosen by the mechanism and the optimal interval. Let L be the set of $n/2 - 1$ agent at the left of m , M the set of agents between m and s_o , and R the set of agents at the right of s_o ; note that $|M \cup R| = n/2$. Clearly, the maximum cost of the mechanism is n and the minimum optimal cost is 0. We make the following observations:

- The median agent decreases the cost of the mechanism by 1 and increases the optimal cost by $1 - x$ (the part of the median interval that the optimal solution does not cover).
- Any agent $i \in L$ such that $|I_i \cap [s_o, t_m]| = 0$ increases the optimal cost by 1; let A be the set of such agents.
- Any agent $i \in L$ such that $|I_i \cap [s_o, t_m]| = x_i > 0$ decreases the cost of the mechanism by at least $1 - x + x_i$ and increases the optimal cost by at least $1 - x_i$ (since the interval of i starts before s_m but reaches s_o); let B be the set of such agents.
- Any agent $i \in M$ decreases the cost of the mechanism by at least x .
- Any agent $i \in R$ such that $|I_i \cap [s_o, t_m]| = x_i > 0$ decreases the cost of the mechanism by some length $x_i \in [0, x]$ and increases the optimal cost by $1 - x_i - (1 - x) = x - x_i$; let Γ be the set of such agents.
- Any agent $i \in R$ such that $|I_i \cap [s_o, t_m]| = 0$ increases the optimal cost by at least x ; let Δ be the set of such agents.

Hence, we have

$$\begin{aligned} \text{SC}(I_m) &\leq n - 1 - |B|(1 - x) - \sum_{i \in B} x_i - |M|x - \sum_{i \in \Gamma} x_i \\ \text{SC}(O) &\geq 1 - x + |A| + |B| - \sum_{i \in B} x_i + |\Gamma|x - \sum_{i \in \Gamma} x_i + |\Delta|x. \end{aligned}$$

So, the approximation ratio is at most

$$\frac{n - 1 - |B|(1 - x) - \sum_{i \in B} x_i - |M|x - \sum_{i \in \Gamma} x_i}{1 - x + |A| + |B| - \sum_{i \in B} x_i + |\Gamma|x - \sum_{i \in \Gamma} x_i + |\Delta|x}.$$

Since the ratio is at least 1 (by definition), it is an increasing function in terms of $\sum_{i \in B} x_i \leq |B|x$ and it terms of $\sum_{i \in \Gamma} x_i \leq |\Gamma|x$, and is thus at most

$$\frac{n - 1 - |B| - (|M| + |\Gamma|)x}{1 - x + |A| + |B|(1 - x) + |\Delta|x}.$$

Since $|A| + |B| + 1 = n/2$ and $|M| + |\Gamma| + |\Delta| = n/2$, we further obtain

$$\frac{n - 1 - |B| - (n/2) \cdot x + |\Delta|x}{n/2 - x - |B|x + |\Delta|x}.$$

This is a decreasing function in terms of $|\Delta| \geq 0$, and is thus at most

$$\frac{n - 1 - |B| - (n/2) \cdot x}{n/2 - x - |B|x}.$$

This is a decreasing function in terms of $|B| \geq 0$, and is thus at most

$$\frac{n - 1 - (n/2) \cdot x}{n/2 - x}.$$

Finally, this is a decreasing function in terms of x and thus attains its maximum value of $2 - 2/n$ when $x = 0$. \square

Next, we present a lower bound for deterministic truthful mechanisms that matches the upper bound of Theorem 3. Before we do so though, we will provide a structural property of any deterministic truthful mechanism. This property will be repeatedly used in order to prove the lower bound.

Lemma 1. *Consider a deterministic truthful mechanism \mathcal{M} , an instance \mathcal{I} , and an agent i such that $I_i \cap \mathcal{M}(\mathcal{I}) \neq \emptyset$. In addition, consider the instance $\mathcal{I}' = (I'_i, \mathcal{I}_{-i})$, where $I'_i \cap I_i \cap \mathcal{M}(\mathcal{I}) \neq \emptyset$, and let $\mathcal{M}(\mathcal{I}')$ be the location of the covering interval in \mathcal{I}' under mechanism \mathcal{M} . Then, it must hold that $I'_i \cap \mathcal{M}(\mathcal{I}') \neq \emptyset$.*

Theorem 4. *Let \mathcal{M} be any deterministic truthful mechanism. Then the approximation ratio of \mathcal{M} is at least $2 - \frac{2}{n}$.*

Proof. Let \mathcal{M} be any deterministic truthful mechanism. At a high level, the proof will construct a series of instances and will use the truthfulness of \mathcal{M} to argue about the possible positions of the covering interval on each one of those instances.

The starting point is the following instance \mathcal{I}^0 , with two groups of agents: group \mathcal{G}_0 contains $\frac{n}{2}$ agents with $I_i = [0, 1]$ for all $i \in \mathcal{G}_0$ and group \mathcal{G}_1 contains $\frac{n}{2}$ agents with $I_i = [n, n+1]$ for all $i \in \mathcal{G}_1$. Without loss of generality, we will assume that on instance \mathcal{I}^0 , mechanism \mathcal{M} locates the covering interval $[a_0, b_0]$ such that it covers some part of the intervals of the agents from cluster \mathcal{G}_0 ; the other case is symmetric. Observe here that throughout the proof it is without loss of generality to assume that the covering interval *always* covers a strictly positive part of an agent; if this was not the case, then the optimal cost would be $n/2$ while the mechanism would achieve cost of n and thus it would be 2-approximate. In addition, again without loss of generality, we will assume that $0 \leq a_0 \leq 1$. Observe that since the covering interval has length 1, then it cannot cover any agent from cluster \mathcal{G}_1 .

The proof will construct a sequence of *families* of instances $\mathcal{J}^0, \mathcal{J}^1, \mathcal{J}^2, \dots, \mathcal{J}^k$, where $k \leq n/2 - 1$, which will guarantee that, on any instance in any of these families:

1. mechanism \mathcal{M} *cannot* place the covering interval and cover (part of) the cluster of agents located at $[n, n+1]$;
2. mechanism \mathcal{M} can move the covering interval only to the right of its position in the previous instance and never to the left;
3. the maximum approximation ratio that mechanism \mathcal{M} can achieve will strictly decrease, compared to the previous family.

Before we present the formal argument, we define some notation that will make the exposition more clear. For every instance \mathcal{I} , let $X(\mathcal{I}) = \{i \in [n] : I_i \cap \mathcal{M}(\mathcal{I}) \neq \emptyset\}$, i.e., the set $X(\mathcal{I})$ contains the agents that have a non-empty intersection with the covering interval $\mathcal{M}(\mathcal{I})$. We will prove by

induction that for each \mathcal{J}^k with $k \in \{0, \dots, \frac{n}{2} - 1\}$ and every instance $\mathcal{I} \in \mathcal{J}^k$ the following two conditions are satisfied.

1. The left endpoint of $\mathcal{M}(\mathcal{I})$ is in $[k, k + 1]$.
2. On instance $\mathcal{I} \in \mathcal{J}^k$, due to truthfulness, mechanism \mathcal{M} is able to cover a total mass of at most $n/2 - k$, for a social cost of at least $n - (n/2 - k) = n/2 + k$, while the optimal social cost will remain the same, equal to $n/2$. Hence, when the argument reaches family $\mathcal{J}^{\frac{n}{2}-1}$, i.e., when $k = n/2 - 1$, the social cost of \mathcal{M} becomes at least $n - 1$, leading to an approximation ratio of $2 - 2/n$. For a visual representation of instance $\mathcal{I} \in \mathcal{J}^{\frac{n}{2}-1}$, see the right-hand side of Figure 1.

Now we are ready to complete the proof. Observe that by assumption Conditions 1 and 2 above hold for \mathcal{I}^0 , so they hold for the base case, \mathcal{J}^0 , of our induction. For the induction step, assume that Conditions 1 and 2 hold for instance $\mathcal{I} \in \mathcal{J}^k$, for some $k \leq n/2 - 1$. Hence, we have that the left endpoint of $\mathcal{M}(\mathcal{I})$ is in $[k, k + 1]$ and that $\mathcal{M}(\mathcal{I})$ intersects with at most $n/2 - k$ agents, formally, $|X(\mathcal{I})| \leq n/2 - k$.

In what follows, without loss of generality we will assume that $|X(\mathcal{I})| > 1$; if $|X(\mathcal{I})| \leq 1$, then $\mathcal{M}(\mathcal{I})$ has zero intersection with the intervals of at least $n - 1$ agents, and hence has a social cost of at least $n - 1$. The optimal cost is $n/2$, and hence \mathcal{M} has an approximation ratio of at least $2 - 2/n$ and we are done.

Given the above, we can pick a ‘‘rightmost’’ agent $i \in X(\mathcal{I})$ with interval $I_i = [s_i, t_i]$, i.e., an agent i such that $t_i \geq t_{i'}$ for all $i' \in X(\mathcal{I})$. We define instance $\mathcal{I}' = (I'_i, \mathcal{I}_{-i})$ as follows, where $\text{left}(I')$ and $\text{right}(I')$ denote the left and right endpoints of interval I' respectively.

- If $\text{left}(\mathcal{M}(\mathcal{I})) < t_i < \text{right}(\mathcal{M}(\mathcal{I}))$, then $I'_i = [t_i, t_i + 1]$.
- If $t_i \geq \text{right}(\mathcal{M}(\mathcal{I}))$, then $I'_i = [\text{right}(\mathcal{M}(\mathcal{I})) - \delta, \text{right}(\mathcal{M}(\mathcal{I})) - \delta + 1]$, where $\delta > 0$ is an arbitrarily small quantity.

Observe that in both cases we have that $I_i \cap I'_i \cap \mathcal{M}(\mathcal{I}) \neq \emptyset$. Thus, from Lemma 1 and due to the truthfulness of mechanism \mathcal{M} , it must hold that $I'_i \cap \mathcal{M}(\mathcal{I}') \neq \emptyset$. We distinguish between three cases.

- $\text{left}(\mathcal{M}(\mathcal{I}')) < k + 1$ and $|X(\mathcal{I}')| > 1$. Then, we create a new instance \mathcal{I}'' as before; formally, we set $\mathcal{I} = \mathcal{I}'$ and $X(\mathcal{I}) = X(\mathcal{I}')$ and we choose an agent from $X(\mathcal{I})$ to move.
- $\text{left}(\mathcal{M}(\mathcal{I}')) < k + 1$ and $|X(\mathcal{I}')| = 1$. Then, as we have argued above, the approximation ratio of mechanism \mathcal{M} is $2 - 2/n$.
- $\text{left}(\mathcal{M}(\mathcal{I}')) \in [k + 1, k + 2]$. In this case, it holds that $\mathcal{I}' \in \mathcal{J}^{k+1}$. Observe that since we have assumed that $|X(\mathcal{I})| > 1$ and that since we have created instance \mathcal{I}' by moving the ‘‘rightmost’’ interval of $X(\mathcal{I})$, it should hold that $|X(\mathcal{I}')| \leq |X(\mathcal{I})| - 1 \leq \frac{n}{2} - k - 1$, where for the last inequality we have used the induction hypothesis.

This completes the induction step and the proof. \square

4 Social Cost: Randomized Mechanisms

In this section, we turn our focus to randomized mechanism for the social cost. We present a simple randomized, universally truthful mechanism that achieves an approximation ratio of $5/3$, thus outperforming all deterministic truthful mechanisms. In particular, we consider the following mechanism, which we coin UNIFORM-STATISTIC:

Definition 2 (UNIFORM-STATISTIC). Let ℓ be the $(n/3)$ -th leftmost agent, m be the median agent, and r be the $(2n/3)$ -th leftmost agent. Place the covering interval at the starting position of each of $\{\ell, m, r\}$ with probability $1/3$.

The mechanism is a convex combination of k -ordered statistics, and hence by Theorem 1, it is universally truthful. What remains is to bound its approximation ratio, established by the following theorem.

Theorem 5. *The approximation of UNIFORM-STATISTIC is at most $5/3$.*

Proof idea. We show that any arbitrary instance can be transformed into one of the following two possible worst-cases instances (up to symmetries), by appropriately moving the agents so that the approximation ratio of the mechanism does not decrease.

- **WC11:** The first instance is such that there are (approximately) $2n/3$ agents grouped together, while the remaining $n/3$ agents are all singletons without any intersection with any other agent. The optimal interval completely covers the group of $2n/3$ agents for a social cost of $n/3$. The output of the mechanism coincides with the optimal with probability $2/3$ (due to agents ℓ and m , or agents m and r), and has social cost (approximately) n with probability $1/3$ when it chooses a singleton. So, the approximation ratio is $2/3 + \frac{1}{3} \cdot \frac{n}{n/3} = 5/3$. See Figure 1a.
- **WC12:** The second instance is such that there are $n/2$ singleton agents (including m) without any intersection with any other agent, while the remaining $n/2$ agents are grouped together. Here, the optimal interval completely covers $n/2$ agents for a social cost of $n/2$. The output of the mechanism coincides with the optimal with probability $1/3$ (due to agent r , or agent ℓ) and has cost (approximately) n with probability $2/3$ when it chooses a singleton. So, the approximation ratio is $1/3 + \frac{2}{3} \cdot \frac{n}{n/2} = 5/3$. See Figure 1b. \square

We complement the aforementioned positive result with a lower bound of $3/2$ on the approximation ratio of any randomized truthful mechanism that is a convex combination of k -ordered statistic mechanisms.

Theorem 6. *The approximation ratio of any convex combination of k -ordered statistic mechanisms is at least $3/2$.*

Proof. Let p be the total probability with which any of the first $n/2$ agents is chosen; hence, $1 - p$ is the total probability with which any of the remaining $n/2$ is chosen. Without loss of generality, $p \geq 1/2$. Now, consider an instance in which the first $n/2$ agents are singletons, whereas the other $n/2$ agents are all grouped together. The optimal social cost




Figure 1: The two worst-case instances for the UNIFORM-STATISTIC mechanism. In each figure there are some singleton agents and a group of agents whose intervals overlap (these are depicted as a shaded rectangle).

is exactly $n/2$, while the expected social cost of the mechanism is $p \cdot n + (1 - p) \cdot n/2$, since with probability p we choose some of the first $n/2$ agents leading to social cost n and with probability $1 - p$ we choose some of the last $n/2$ agents leading to the optimal social cost of $n/2$. Therefore, the approximation ratio is $1 + p \geq 3/2$. \square

5 Max Cost

In this section, we focus on the max cost and show that the best possible approximation ratio of any deterministic mechanism is 2, and this is achieved by any k -ordered statistic mechanism.

Theorem 7. *For the max cost, the approximation ratio of any k -ordered statistic mechanism is at most 2.*

Proof. Consider an arbitrary instance and any k -statistic mechanism. The max cost of the solution computed by the mechanism as well that of the optimal solution depend on how close the intervals of the leftmost and rightmost agents are. We consider the following cases:

- If the intervals of the leftmost and the rightmost agents are disjoint and the distance between them is at least 1, then the max cost of the mechanism and the optimal max cost are both 1; hence, the mechanism is optimal.
- If the intervals of the leftmost and the rightmost agents are disjoint and the distance between them is equal to $1 - x$ for some $x \in (0, 1)$, then the optimal interval can cover $x/2$ of each of these two agents (and any other agent inbetween them), leading to an optimal max cost of $1 - x/2 \geq 1/2$. Since the max cost of the mechanism is again 1, the approximation ratio is at most 2.
- If the interval of the leftmost and the rightmost agents have an overlap of $x < 1$, then the max cost of the mechanism is $1 - x$. The optimal solution can cover $x + y$ from each of the two agents, where y is such that $x + 2y = 1 \Leftrightarrow y = \frac{1-x}{2}$, leading to an optimal max cost of $1 - x - y = \frac{1-x}{2}$. So, the approximation ratio is 2.

Overall, in any case, the approximation ratio is at most 2. \square

Next, we show that there is no better deterministic truthful mechanism.

Theorem 8. *For the max cost, for any $k > 0$ the approximation ratio of any deterministic truthful mechanism is at least $2 - \frac{1}{k}$.*

Proof idea. In order to prove our theorem we produce a sequence of instances with two agents – Left agent and Right agent – such that the approximation ratio of any deterministic truthful mechanism will be monotonically increasing and tend to 2. The high level idea is that at every iteration of the sequence, either Left agent moves ε to the left, or Right agent moves ε to the right; at the same time though, due to truthfulness the mechanism can either (a) follow the agent that moves and lose a “large” fraction of the other agent that optimal solution covers, or (b) do not follow the agent that moves and thus lose a fraction of the agent that is covered by the optimal solution. Using this idea iteratively, we prove that the approximation ratio of any deterministic truthful mechanism tends to 2. \square

We finally show that no randomized mechanism that is a convex combination of k -ordered mechanism can achieve an approximation ratio better than 2.

Theorem 9. *For the max cost, the approximation ratio of any mechanism that is a convex combination of k -ordered mechanisms is at least 2.*

6 Two Natural Extensions

Our results so far concerned the case of known intervals of equal length and settled the problem for deterministic truthful mechanisms, while the best possible approximation ratio for randomized mechanisms is still to be determined. We now present two very natural extensions of the main model which could be better fitted to several of the potential applications of the problem. In particular, we first consider the setting where the lengths of the agent intervals are private information; for this model, it turns out that meaningful approximations and truthfulness are incompatible. Second, we consider the case where the lengths of the agent interval are known but unequal; for this, we show that a finite approximation ratio is possible for the social cost, and we also identify the best possible mechanism for the max cost.

6.1 Unknown Interval Lengths

In general, it seems natural to assume that the length of the agent intervals, as well as their positions, could constitute reported information. In this case, however, we prove the following impossibilities for both the social and the max cost.

Theorem 10. *For the social cost, when the lengths of the intervals are unknown, the approximation ratio of any randomized truthful in expectation mechanism is $\Omega(1/\varepsilon)$, for any $\varepsilon \in (0, 1)$.*

Theorem 11. *For the max cost, when the lengths of the intervals are unknown, the approximation ratio of randomized truthful in expectation mechanism is $\Omega(1/\varepsilon)$, for any $\varepsilon \in (0, 1)$.*

6.2 Known, Unequal Interval Lengths

Another interesting variant that directly generalizes our main setting is that in which the interval lengths are known, but they are not necessarily equal. In the case of electricity supply for example, it is reasonable to assume that the government has good estimates of how much time each household requires to complete essential chores, based on verifiable information (e.g., the size of their property or the number of family members), not about their preferences on the different times of the day.

For the social cost, it is not hard to see that the vanilla median mechanism, and in fact any unweighted k -th ordered statistic, has an infinite approximation ratio for this case. However, we can show a linear approximation ratio by considering the MAX-LENGTH mechanism, which places the interval at the starting position of the agent with the maximum-length interval among all agents. This mechanism is clearly truthful since the lengths are known. Without loss of generality, we will assume that the covering interval length is 1.

Theorem 12. *Let $\ell = \max_{i \in N} |I_i|$. For the social cost, the approximation ratio of the MAX-LENGTH mechanism is at most $n - 1$ when $\ell \leq 1$, and at most n when $\ell > 1$.*

For the max cost, we show that MAX-LENGTH achieves an approximation ratio of at most 2. In combination to Theorem 8, this show that MAX-LENGTH is essentially the best possible among all deterministic mechanisms when the lengths of the intervals are known (equal or unequal).

Theorem 13. *For the max cost, the approximation ratio of the MAX-LENGTH mechanism is at most 2.*

7 Other Open Problems and Directions

We envision that the model we have introduced in this paper can serve as a basis for a plethora of further extensions motivated from real life scenarios. Having completely resolved the foundational version of the model, at least with respect to deterministic mechanisms, below we highlight what we consider to be some of the most prominent avenues for future work.

Multiple Intervals. A very meaningful extension is the one where each agent is associated with multiple intervals (say, k_i intervals for agent i), and there are k_c covering intervals to be placed (think of the choice of several different open days at a university); those intervals could be of equal or unequal (known) length. A similar setting is one in which there is a covering budget (a total covering length ℓ_c) which can be partitioned into intervals freely over the line. Similarly, the agents themselves could also have such interval budgets ℓ_i ; one could even impose some restrictions on the number of intervals that can be used by each agent, or by the covering budget.

Different Cost Functions and Objectives. In this paper, we have considered perhaps the simplest and most intuitive cost function for the agents, namely the part of their intervals that is not covered by C . One could consider more complicated cost functions, e.g., functions where the cost is a convex or concave function of the proportion of the agents' interval(s) that are covered, or some most specific functions like a piecewise linear function (e.g., capturing cases where a certain amount of the interval *has* to be covered for the agent to have any reduction in cost).

Utilities and Social Welfare. We could even consider (positive) utilities rather than (negative) costs. For example, in the simplest case of known and equal length intervals that we studied here, the utility of an agent would be the part of her interval that is covered, and the approximations would be in terms of the *social welfare*, the total utility of the agents. It is not hard to observe that the social welfare and the social cost of a covering interval C are related in this case; it holds that $SW(C) = n - SC(C)$. Given this, if we have a mechanism \mathcal{M} with a provable approximation ratio of ρ in terms of the social cost, the approximation ratio of the same mechanism is at most $\frac{1}{\rho} + \frac{n(\rho-1)}{\rho \cdot SW(\mathcal{M})}$ in terms of the social welfare. This directly gives us that the approximation ratio of the MEDIAN mechanism is at most $n/2$, since $\rho = 2 - 2/n$ and $SW(\mathcal{M}) \geq 1$ (since at least one agent is completely covered). It is not hard to verify that the same arguments as in Theorem 4 can lead to an essentially matching lower bound for all deterministic truthful mechanisms, thus showing that the MEDIAN mechanism is best possible even in terms of utilities and the social welfare. For more general settings, such relations between the social cost and social welfare might not exist however, and studying both of them is interesting.

Obnoxious and Hybrid Models. There are also applications in which the agents might want to avoid any intersection with the covering interval. For example, when the interval corresponds to a public transportation line, the agents might want to not have any intersection with the interval since they have no interest in using the public transportation and want to avoid possible congestion or noise. On the other hand, some agents might want to have intersection with the interval in such a case as they want to use the public transportation, thus leading to interesting hybrid interval covering models. This sort of application draws parallels to similar models in the facility location literature [Chan *et al.*, 2021].

Higher-Dimensional Spaces. One does not have to restrict attention only to intervals; a very similar setting can be defined in which each agent is associated with one or more geometric shapes on a higher dimensional space (e.g., the plane) and there is also one or more covering geometric shapes to be placed, aiming to minimize the cost of the agents as a function of the intersection with their shapes. For example, think that the covering comes from cellular antennas and every agent wants to minimize the area that they are not covered by the radius of the antenna.

References

- [Abebe *et al.*, 2020] Rediet Abebe, Richard Cole, Vasilis Gkatzelis, and Jason D Hartline. A truthful cardinal mechanism for one-sided matching. In *Proceedings of the 14th annual ACM-SIAM symposium on discrete algorithms (SODA)*, pages 2096–2113, 2020.
- [Alon *et al.*, 2011] Noga Alon, Felix Fischer, Ariel Procaccia, and Moshe Tennenholtz. Sum of us: Strategyproof selection from the selectors. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 101–110, 2011.
- [Bei *et al.*, 2022] Xiaohui Bei, Xinhang Lu, and Warut Suksumpong. Truthful cake sharing. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, pages 4809–4817, 2022.
- [Bjelde *et al.*, 2017] Antje Bjelde, Felix Fischer, and Max Klimm. Impartial selection and the power of up to two choices. *ACM Transactions on Economics and Computation*, 5(4):1–20, 2017.
- [Black, 1948] Duncan Black. On the rationale of group decision-making. *Journal of political economy*, 56(1):23–34, 1948.
- [Chan *et al.*, 2021] Hau Chan, Aris Filos-Ratsikas, Bo Li, Minming Li, and Chenhao Wang. Mechanism design for facility location problem: A survey. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1–17, 2021.
- [Dayama *et al.*, 2015] Pankaj Dayama, Balakrishnan Narayanaswamy, Dinesh Garg, and Y Narahari. Truthful interval cover mechanisms for crowdsourcing applications. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1091–1099, 2015.
- [Dummett and Farquharson, 1961] Michael Dummett and Robin Farquharson. Stability in voting. *Econometrica: Journal of The Econometric Society*, pages 33–43, 1961.
- [Filos-Ratsikas *et al.*, 2014] Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, and Jie Zhang. Social welfare in one-sided matchings: Random priority and beyond. In *Proceedings of the 7th International Symposium on Algorithmic Game Theory (SAGT)*, pages 1–12, 2014.
- [Fischer and Klimm, 2014] Felix Fischer and Max Klimm. Optimal impartial selection. In *Proceedings of the 15th ACM conference on Economics and computation (EC)*, pages 803–820, 2014.
- [Kleinberg and Tardos, 2006] Jon Kleinberg and Eva Tardos. *Algorithm design*. Pearson Education India, 2006.
- [Kolen *et al.*, 2007] Antoon WJ Kolen, Jan Karel Lenstra, Christos H Papadimitriou, and Frits CR Spieksma. Interval scheduling: A survey. *Naval Research Logistics (NRL)*, 54(5):530–543, 2007.
- [Krysta *et al.*, 2014] Piotr Krysta, David Manlove, Baharak Rastegari, and Jinshan Zhang. Size versus truthfulness in the house allocation problem. In *Proceedings of the 15th ACM conference on Economics and computation (EC)*, pages 453–470, 2014.
- [Markakis *et al.*, 2022] Evangelos Markakis, Georgios Papatropoulos, and Artem Tsikiris. On improved interval cover mechanisms for crowdsourcing markets. In *Proceedings of the 15th International Symposium on Algorithmic Game Theory (SAGT)*, pages 94–112. Springer, 2022.
- [Moulin, 1980] Hervé Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.
- [Procaccia and Tennenholtz, 2013] Ariel D. Procaccia and Moshe Tennenholtz. Approximate mechanism design without money. *ACM Transactions on Economics and Computation*, 1(4):18:1–18:26, 2013.
- [Roughgarden, 2022] Tim Roughgarden. *Algorithms illuminated*. Soundlikeyourself publishing, 2022.
- [Xu *et al.*, 2020] Xinpeng Xu, Minming Li, and Lingjie Duan. Strategyproof mechanisms for activity scheduling. In *19th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1539–1547, 2020.