On the Fine-Grained Complexity of Small-Size Geometric Set Cover and Discrete k-Center for Small k

Timothy M. Chan
Department of Computer Science, University of Illinois at Urbana-Champaign, IL, USA

Qizheng He
Department of Computer Science, University of Illinois at Urbana-Champaign, IL, USA

Yuancheng Yu
Department of Computer Science, University of Illinois at Urbana-Champaign, IL, USA

Abstract

We study the time complexity of the discrete k-center problem and related (exact) geometric set cover problems when k or the size of the cover is small. We obtain a plethora of new results:

- We give the first subquadratic algorithm for *rectilinear discrete 3-center* in 2D, running in $\widetilde{O}(n^{3/2})$ time.
- We prove a lower bound of $\Omega(n^{4/3-\delta})$ for rectilinear discrete 3-center in 4D, for any constant $\delta > 0$, under a standard hypothesis about triangle detection in sparse graphs.
- Given n points and n *weighted* axis-aligned unit squares in 2D, we give the first subquadratic algorithm for finding a minimum-weight cover of the points by 3 unit squares, running in $\widetilde{O}(n^{8/5})$ time. We also prove a lower bound of $\Omega(n^{3/2-\delta})$ for the same problem in 2D, under the well-known APSP Hypothesis. For arbitrary axis-aligned rectangles in 2D, our upper bound is $\widetilde{O}(n^{7/4})$.
- We prove a lower bound of $\Omega(n^{2-\delta})$ for Euclidean discrete 2-center in 13D, under the Hyperclique Hypothesis. This lower bound nearly matches the straightforward upper bound of $\widetilde{O}(n^{\omega})$, if the matrix multiplication exponent ω is equal to 2.
- We similarly prove an $\Omega(n^{k-\delta})$ lower bound for Euclidean discrete k-center in $O(k)$ dimensions for any constant $k \geq 3$, under the Hyperclique Hypothesis. This lower bound again nearly matches known upper bounds if $\omega = 2$.
- We also prove an $\Omega(n^{2-\delta})$ lower bound for the problem of finding 2 boxes to cover the largest number of points, given n points and n boxes in 12D. This matches the straightforward near-quadratic upper bound.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Geometric set cover, discrete k-center, conditional lower bounds

Digital Object Identifier 10.4230/LIPIcs.ICALP.2023.34

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: https://arxiv.org/abs/2305.01892 [19]

Funding *Timothy M. Chan*: Work supported by NSF Grant CCF-2224271.

Acknowledgements We thank Yinzhan Xu for helpful discussions.

1 Introduction

1.1 The discrete k-center problem for small k

The *Euclidean k-center* problem is well-known in computational geometry and has a long history: given a set P of n points in $\mathbb{R}^d$ and a number k, we want to find k congruent balls covering S, while minimizing the radius. Euclidean 1-center can be solved in linear time for any constant dimension d by standard techniques for low-dimensional linear programming or LP-type problems [21, 24, 27, 41, 50]. In a celebrated paper from SoCG'96, Sharir [44] gave the first $\widetilde{O}(n)$-time[1] algorithm for Euclidean 2-center in $\mathbb{R}^2$, which represented a significant improvement over previous near-quadratic algorithms (the hidden logarithmic factors have since been reduced in a series of subsequent works [29, 16, 49, 23]). The problem is more difficult in higher dimensions: the best time bound for Euclidean 2-center in $\mathbb{R}^d$ is about n^d (see [3, 2] for some results on the $\mathbb{R}^3$ case), and Cabello et al. [14] proved a conditional lower bound, ruling out $n^{o(d)}$-time algorithms, assuming the Exponential Time Hypothesis (ETH). We are not aware of any work specifically addressing the Euclidean 3-center problem.

The k-center problem has also been studied under different metrics. The most popular version after Euclidean is L_∞ or *rectilinear k-center*: here, we want to find k congruent hypercubes covering P, while minimizing the side length of the hypercubes.[2] As expected, the rectilinear version is a little easier than the Euclidean. Sharir and Welzl in SoCG'96 [45] showed that rectilinear 3-center problem in $\mathbb{R}^2$ can be solved in linear time, and that rectilinear 4-center and 5-center in $\mathbb{R}^2$ can be solved in $\widetilde{O}(n)$ time (the logarithmic factors have been subsequently improved by Nussbaum [42]). Katz and Nielsen's work in SoCG'96 [35] implied near-linear-time algorithms for rectilinear 2-center in any constant dimension d, while Cabello et al. in SODA'08 [14] gave an $O(n \log n)$-time algorithm for rectilinear 3-center in any constant dimension d. Cabello et al. also proved a conditional lower bound for rectilinear 4-center, ruling out $n^{o(\sqrt{d})}$-time algorithms under ETH.

In this paper, we focus on a natural variant of the problem called *discrete k-center*, which has also received considerable attention: here, given a set P of n points in $\mathbb{R}^d$ and a number k, we want to find k congruent balls covering P, while minimizing the radius, with the extra constraint that the centers of the chosen balls are from P.[3] The Euclidean discrete 1-center problem can be solved in $O(n \log n)$ time in $\mathbb{R}^2$ by a straightforward application of farthest-point Voronoi diagrams; it can also be solved in $O(n \log n)$ (randomized) time in $\mathbb{R}^3$ with more effort [15], and in subquadratic $\widetilde{O}(n^{2-2/(\lceil d/2 \rceil+1)})$ time for $d \geq 4$ by standard range searching techniques [4, 40]. Agarwal, Sharir, and Welzl in SoCG'97 [6] gave the first subquadratic algorithm for Euclidean discrete 2-center in $\mathbb{R}^2$, running in $\widetilde{O}(n^{4/3})$ time.

One may wonder whether Euclidean discrete 2-center in higher constant dimensions could also be solved in subquadratic time via range searching techniques. No results have been reported, but an $\widetilde{O}(n^\omega)$-time algorithm is not difficult to obtain, where $\omega < 2.373$ denotes the matrix multiplication exponent: by binary search, the problem reduces to finding two balls of a given radius r with centers in S covering S, which is equivalent to finding a pair $p, q \in S$ such that $c_{pq} = \bigvee_{z \in S}(a_{pz} \wedge a_{zq})$ is false, where a_{pz} is true iff p and z has distance more than r – this computation reduces to a Boolean matrix product. This approach works for arbitrary

---

[1] The $\widetilde{O}$ notation hides polylogarithmic factors.

[2] All squares, rectangles, hypercubes, and boxes are axis-aligned in this paper.

[3] Some authors define the problem slightly more generally, where the constraint is that the centers are from a second input set; in other words, the input consists of two sets of points ("demand points" and "supply points"). The results of this paper will apply to both versions of the problem.

Table 1 Summary of results on k-center for small k in $\mathbb{R}^2$.

k	Euclidean	rectilinear	Euclidean discrete	rectilinear discrete
1	$O(n)$	$O(n)$	$O(n \log n)$	$O(n)$
2	$\widetilde{O}(n)$ [44]	$O(n)$ [45]	$\widetilde{O}(n^{4/3})$ [6]	$\widetilde{O}(n)$
3		$\widetilde{O}(n)$ [45]		$\widetilde{O}(n^{3/2})$ **(new)**

(not necessarily geometric) distance functions. The main question is whether geometry could help in obtaining faster algorithms in the higher-dimensional Euclidean setting, as Agarwal, Sharir, and Welzl were able to exploit successfully in $\mathbb{R}^2$:

▶ **Question 1.** *Is there an algorithm running in faster than n^ω time for the Euclidean discrete 2-center problem in higher constant dimensions?*

We can similarly investigate the rectilinear version of the discrete k-center problem, which is potentially easier. For example, the rectilinear discrete 2-center problem can be solved in $\widetilde{O}(n)$ time in any constant dimension d, by a straightforward application of orthogonal range searching, as reported in several papers [10, 11, 34]. The approach does not seem to work for the rectilinear discrete 3-center problem. Naively, rectilinear discrete 3-center can be reduced to n instances of (some version of) rectilinear discrete 2-center, and solved in $\widetilde{O}(n^2)$ time. However, no better results have been published, leading to the following questions:

▶ **Question 2.** *Is there a subquadratic-time algorithm for the rectilinear discrete 3-center problem?*

▶ **Question 3.** *Are there lower bounds to show that the rectilinear discrete 3-center problem does not have near-linear-time algorithm (and is thus strictly harder than rectilinear discrete 2-center, or rectilinear continuous 3-center)?*

Similar questions may be asked about rectilinear discrete k-center for $k \geq 4$. Here, the complexity of the problem is upper-bounded by $\widetilde{O}(n^{\omega(\lfloor k/2\rfloor,1,\lceil k/2\rceil)})$, where $\omega(a,b,c)$ denotes the exponent for multiplying an $n^a \times n^b$ and an $n^b \times n^c$ matrix: by binary search, the problem reduces to finding k hypercubes of a given edge length r with centers in S covering S, which is equivalent to finding a dominating set of size k in the graph with vertex set S where an edge pz exists iff the distance of p and z is more than r – the dominating set problem reduces to rectangular matrix multiplication with the time bound stated, as observed by Eisenbrand and Grandoni [28]. Note that the difference $\omega(\lfloor k/2\rfloor, 1, \lceil k/2\rceil) - k$ converges to 0 as $k \to \infty$ by known matrix multiplication bounds [25] (and is exactly 0 if $\omega = 2$).

As k gets larger compared to d, a better upper bound of $n^{O(dk^{1-1/d})}$ is known for both the continuous and discrete k-center problem under the Euclidean and rectilinear metric [5, 31, 32]. Recently, in SoCG'22, Chitnis and Saurabh [22] (extending earlier work by Marx [38] in the $\mathbb{R}^2$ case) proved a nearly matching conditional lower bound for discrete k-center in $\mathbb{R}^d$, ruling out $n^{o(k^{1-1/d})}$-time algorithms under ETH. However, these bounds do not answer our questions concerning very small k's. In contrast, the conditional lower bounds by Cabello et al. [14] that we have mentioned earlier are about very small k and so are more relevant, but are only for the continuous version of the k-center problem. (The continuous version behaves differently from the discrete version; see Tables 1–2.)

Table 2 Summary of results on k-center for small k in $\mathbb{R}^d$ for an arbitrary constant d. (CLB stands for "conditional lower bound".)

k	Euclidean	rectilinear	Euclidean discrete	rectilinear discrete
1	$O(n)$	$O(n)$	$\widetilde{O}(n^{2-2/(\lceil d/2\rceil+1)})$	$O(n)$
2	$n^{O(d)}$ CLB: $n^{\Omega(d)}$ [14]	$\widetilde{O}(n)$ [35]	$\widetilde{O}(n^{\omega})$ CLB: $\Omega(n^{2-\delta})$ **(new)**	$\widetilde{O}(n)$
3		$\widetilde{O}(n)$ [14]	$\widetilde{O}(n^{\omega(1,1,2)})$ CLB: $\Omega(n^{3-\delta})$ **(new)**	$\widetilde{O}(n^2)$ CLB: $\Omega(n^{4/3-\delta})$ **(new)**
4		$n^{O(d)}$ CLB: $n^{\Omega(\sqrt{d})}$ [14]	$\widetilde{O}(n^{\omega(2,1,2)})$ CLB: $\Omega(n^{4-\delta})$ **(new)**	$\widetilde{O}(n^3)$

1.2 The geometric set cover problem with small size k

The decision version of the discrete k-center problem (deciding whether the minimum radius is at most a given value) reduces to a *geometric set cover* problem: given a set P of n points and a set R of n objects, find the smallest subset of objects in R that cover all points of P. Geometric set cover has been extensively studied in the literature, particularly from the perspective of approximation algorithms (since for most types of geometric objects, set cover remains NP-hard); for example, see the references in [18]. Here, we are interested in *exact* algorithms for the case when the optimal size k is a small constant.

For the application to Euclidean/rectilinear k-center, the objects are congruent balls/hypercubes, or by rescaling, unit balls/hypercubes, but other types of objects may be considered, such as arbitrary rectangles or boxes.

We can also consider the *weighted* version of the problem: here, given a set P of n points, a set R of n weighted objects, and a small constant k, we want to find a subset of k objects in R that cover all points of P, while minimizing the total weight of the chosen objects.

A "dual" problem is *geometric hitting set*, which in the weighted case is the following: given a set P of n weighted points, a set R of n objects, and a small constant k, find a subset of k points in P that hit all objects of R, while minimizing the total weight of the chosen points. (The continuous unweighted version, where the chosen points may be anywhere, is often called the *piercing* problem.) In the case of unit balls/hypercubes, hitting set is equivalent to set cover due to self-duality.

For rectangles in $\mathbb{R}^2$ or boxes in $\mathbb{R}^d$, size-2 geometric set cover (unweighted or weighted) can be solved in $\widetilde{O}(n)$ time, like discrete rectilinear 2-center [10, 11, 34], by orthogonal range searching. Analogs to Questions 2–3 may be asked for size-3 geometric set cover for rectangles/boxes.

Surprisingly, the complexity of exact geometric set cover of small size k has not received as much attention, but very recently in SODA'23, Chan [17] initiated the study of similar questions for geometric independent set with small size k, for example, providing subquadratic algorithms and conditional lower bounds for size-4 independent set for boxes.

For larger k, hardness results by Marx and Pilipczuk [39] and Bringmann et al. [13] ruled out $n^{o(k)}$-time algorithms for size-k geometric set cover for rectangles in $\mathbb{R}^2$ and unit hypercubes (or orthants) in $\mathbb{R}^4$, and $n^{o(\sqrt{k})}$-time algorithms for unit cubes (or orthants) in $\mathbb{R}^3$ under ETH. But like the other fixed-parameter intractability results mentioned, these proofs do not appear to imply any nontrivial lower bound for very small k such as $k = 3$.

1.3 New results

New algorithms. In this paper, we answer Question 2 in the affirmative for dimension $d = 2$, by presenting the first subquadratic algorithms for rectilinear discrete 3-center in $\mathbb{R}^2$, and more generally, for (unweighted and weighted) geometric size-3 set cover for unit squares, as well as arbitrary rectangles in $\mathbb{R}^2$. More precisely, the time bounds of our algorithms are:

- $\widetilde{O}(n^{3/2})$ for rectilinear discrete 3-center in $\mathbb{R}^2$ and unweighted size-3 set cover for unit squares in $\mathbb{R}^2$;
- $\widetilde{O}(n^{8/5})$ for weighted size-3 set cover for unit squares in $\mathbb{R}^2$;
- $\widetilde{O}(n^{5/3})$ for unweighted size-3 set cover for rectangles in $\mathbb{R}^2$;
- $\widetilde{O}(n^{7/4})$ for weighted size-3 set cover for rectangles in $\mathbb{R}^2$.

New conditional lower bounds. We also prove the first nontrivial conditional lower bounds on the time complexity of rectilinear discrete 3-center and related size-3 geometric set cover problems. More precisely, our lower bounds are:[4]

- $\Omega(n^{3/2-\delta})$ for weighted size-3 set cover (or hitting set) for unit squares in $\mathbb{R}^2$, assuming the APSP Hypothesis;
- $\Omega(n^{4/3-\delta})$ for rectilinear discrete 3-center in $\mathbb{R}^4$ and unweighted size-3 set cover (or hitting set) for unit hypercubes in $\mathbb{R}^4$, assuming the Sparse Triangle Hypothesis;
- $\Omega(n^{4/3-\delta})$ for unweighted size-3 set cover for boxes in $\mathbb{R}^3$, assuming the Sparse Triangle Hypothesis.

The lower bound in the first bullet is particularly attractive, since it implies that conditionally, our $\widetilde{O}(n^{8/5})$-time algorithm for weighted size-3 set cover for unit squares in $\mathbb{R}^2$ is within a small factor (near $n^{0.1}$) from optimal, and that our $\widetilde{O}(n^{7/4})$-time algorithm for weighted size-3 set cover for rectangles in $\mathbb{R}^2$ is within a factor near $n^{0.25}$ from optimal. The second bullet answers Question 3, implying that rectilinear discrete 3-center is strictly harder than rectilinear discrete 2-center and rectilinear (continuous) 3-center, at least when the dimension is 4 or higher. (In contrast, rectilinear (continuous) 4-center is strictly harder than rectilinear discrete 4-center for sufficiently large constant dimensions [14]; see Table 2.)

In addition, we prove the following conditional lower bounds:

- $\Omega(n^{2-\delta})$ for Euclidean discrete 2-center in $\mathbb{R}^{13}$ and unweighted size-3 set cover (or hitting set) for unit balls in $\mathbb{R}^{13}$, assuming the Hyperclique Hypothesis;
- $\Omega(n^{k-\delta})$ for Euclidean discrete k-center in $\mathbb{R}^{7k}$ and unweighted size-k set cover for unit balls in $\mathbb{R}^{7k}$ for any constant $k \geq 3$, assuming the Hyperclique Hypothesis.

In particular, this answers Question 1 in the negative if $\omega = 2$ (as conjectured by some researchers): geometry doesn't help for Euclidean discrete 2-center when the dimension is a sufficiently large constant. Similarly, the second bullet indicates that the upper bound $\widetilde{O}(n^{\omega(\lfloor k/2\rfloor,1,\lceil k/2\rceil)})$ for Euclidean discrete k-center is basically tight for any fixed $k \geq 3$ in a sufficiently large constant dimension, if $\omega = 2$. (See Tables 1–3.)

Lastly, we prove a lower bound for a standard variant of set cover known as *maximum coverage*: given a set P of n points, a set R of n objects, and a small constant k, find k objects in R that cover the largest number (rather than all) of points of P. Geometric versions of the maximum coverage problem have been studied before from the approximation algorithms perspective (e.g., see [9]). It is also related to "outliers" variants of k-center problems (where we allow a certain number of points to be uncovered), which have also been studied for small

[4] Throughout this paper, δ denotes an arbitrarily small positive constant.

Table 3 Summary of results on size-3 geometric set cover in $\mathbb{R}^2$.

objects	unweighted		weighted	
unit squares	$\widetilde{O}(n^{3/2})$	**(new)**	$\widetilde{O}(n^{8/5})$	**(new)**
			CLB: $\Omega(n^{3/2-\delta})$	**(new)**
rectangles	$\widetilde{O}(n^{5/3})$	**(new)**	$\widetilde{O}(n^{7/4})$	**(new)**
			CLB: $\Omega(n^{3/2-\delta})$	**(new)**

k (e.g., see [5]). Recall that the size-2 geometric set cover problem for boxes in $\mathbb{R}^d$ can be solved in $\widetilde{O}(n)$ time (which was why our attention was redirected to the size-3 case). In contrast, we show that maximum coverage for boxes cannot be solved in near-linear time even for size $k = 2$. More precisely, we obtain the following lower bound:

- $\Omega(n^{2-\delta})$ for size-2 maximum coverage for unit hypercubes in $\mathbb{R}^{12}$, assuming the Hyperclique Hypothesis.

What is notable is that this lower bound is tight (up to $n^{o(1)}$ factors), regardless of ω, since there is an obvious $\widetilde{O}(n^2)$-time algorithm for boxes in $\mathbb{R}^d$ by answering n^2 orthogonal range counting queries – our result implies that this obvious algorithm can't be improved!

On hypotheses from fine-grained complexity. Let us briefly state the hypotheses used.

- The *APSP Hypothesis* is among the three most popular hypotheses in fine-grained complexity [46] (the other two being the 3SUM Hypothesis and the Strong Exponential Time Hypothesis): it asserts that there is no $O(n^{3-\delta})$-time algorithm for the all-pairs shortest paths problem for an arbitrary weighted graph with n vertices (and $O(\log n)$-bit integer weights). This hypothesis has been used extensively in the algorithms literature (but less often in computational geometry).
- The *Sparse Triangle Hypothesis* asserts that there is no $O(m^{4/3-\delta})$-time algorithm for detecting a triangle (i.e., a 3-cycle) in a sparse unweighted graph with m edges. The current best upper bound for triangle detection, from a 3-decade-old paper by Alon, Yuster, and Zwick [8], is $\widetilde{O}(m^{2\omega/(\omega+1)})$, which is $\widetilde{O}(m^{4/3})$ if $\omega = 2$. (In fact, a stronger version of the hypothesis asserts that there is no $O(m^{2\omega/(\omega+1)-\delta})$-time algorithm.) As supporting evidence, it is known that certain "listing" or "all-edges" variants of the triangle detection problem have an $O(m^{4/3-\delta})$ lower bound, under the 3SUM Hypothesis or the APSP Hypothesis [43, 48, 20]. See [1, 33] for more discussion on the Sparse Triangle Hypothesis, and [17] for a recent application in computational geometry.
- The *Hyperclique Hypothesis* asserts that there is no $O(n^{k-\delta})$-time algorithm for detecting a size-k hyperclique in an ℓ-uniform hypergraph with n vertices, for any fixed $k > \ell \geq 3$. See [37] for discussion on this hypothesis, and [12, 17, 36] for some recent applications in computational geometry, including Künnemann's breakthrough result on conditional lower bounds for Klee's measure problem [36].

Techniques. Traditionally, in computational geometry, subquadratic algorithms with "intermediate" exponents between 1 and 2 tend to arise from the use of nonorthogonal range searching [4] (Agarwal, Sharir, and Welzl's $\widetilde{O}(n^{4/3})$-time algorithm for Euclidean discrete 2-center in $\mathbb{R}^2$ [6] being one such example). Our subquadratic algorithms for rectilinear discrete 3-center in $\mathbb{R}^2$ and related set-cover problems, which are about "orthogonal" or axis-aligned objects, are different. A natural first step is to use a $g \times g$ grid to divide into cases, for some carefully chosen parameter g. Indeed, a grid-based approach was used in some recent subquadratic algorithms by Chan [17] for size-4 independent set for boxes in any constant dimension, and size-5 independent set for rectangles in $\mathbb{R}^2$ (with running time $\widetilde{O}(n^{3/2})$ and $\widetilde{O}(n^{4/3})$ respectively). However, discrete 3-center or rectangle set cover is much more challenging than independent set (for one thing, the 3 rectangles in the solution may intersect each other). To make the grid approach work, we need new original ideas (notably, a sophisticated argument to assign grid cells to rectangles, which is tailored to the 2D case). Still, the entire algorithm description fits in under 3 pages.

Our conditional lower bounds for rectilinear discrete 3-center and the corresponding set cover problem for unit hypercubes are proved by reduction from unweighted or weighted triangle finding in graphs. It turns out there is a simple reduction in $\mathbb{R}^2$ by exploiting weights. However, lower bounds in the unweighted case (and thus the original rectilinear discrete 3-center problem) are much trickier. We are able to design a clever, simple reduction in $\mathbb{R}^6$ by hand, but reducing the dimension down to 4 is far from obvious and we end up employing a *computer-assisted search*, interestingly. The final construction is still simple, and so is easy to verify by hand.

Our conditional lower bound proofs for Euclidean discrete 2-center, and more generally discrete k-center, are inspired by a recent conditional hardness proof by Bringmann et al. [12] from SoCG'22 on a different problem. Specifically, they proved that deciding whether the intersection graph of n unit hypercubes in $\mathbb{R}^{12}$ has diameter 2 requires near-quadratic time under the Hyperclique Hypothesis. A priori, this diameter problem doesn't seem related to discrete k-center; moreover, it was a rectilinear problem, not Euclidean (and we know that in contrast, rectilinear discrete 2-center has a near-linear upper bound!). Our contribution is in realizing that Bringmann et al.'s approach is useful for Euclidean discrete 2-center and k-center, surprisingly. To make the proof work though, we need some new technical ideas (in particular, an extra dimension for the $k = 2$ case, and multiple extra dimensions for larger k, with carefully designed coordinate values). Still, the final proof is not complicated to follow.

Our conditional lower bound for size-2 maximum coverage for boxes is also proved using a similar technique, but again the adaptation is nontrivial, and we introduce some interesting counting arguments that proceed a bit differently from Bringmann et al.'s original proof for diameter (a problem that does not involve counting).

2 Subquadratic Algorithms for Size-3 Set Cover for Rectangles in $\mathbb{R}^2$

In this section, we describe the most basic version of our subquadratic algorithm to solve the size-3 geometric set cover problem for weighted rectangles in $\mathbb{R}^2$. The running time is $\widetilde{O}(n^{16/9})$. Refinements of the algorithm will be described in the full version of the paper, where we will improve the time bound further to $\widetilde{O}(n^{7/4})$, or even better for the unweighted case and unit square case. The rectilinear discrete 3-center problem in $\mathbb{R}^2$ reduces to the unweighted unit square case by standard techniques [15, 30].

We begin with a lemma giving a useful geometric data structure:

▶ **Lemma 1.** *For a set P of n points and a set R of n weighted rectangles in $\mathbb{R}^2$, we can build a data structure in $\widetilde{O}(n)$ time and space, to support the following kind of queries: given a pair of rectangles $r_1, r_2 \in R$, we can find a minimum-weight rectangle $r_3 \in R$ (if it exists) such that P is covered by $r_1 \cup r_2 \cup r_3$, in $\widetilde{O}(1)$ time.*

Proof. By orthogonal range searching [4, 26] on P, we can find the minimum/maximum x- and y-values among the points of P in the complement of $r_1 \cup r_2$ in $\widetilde{O}(1)$ time (since the complement can be expressed as a union of $O(1)$ orthogonal ranges). As a result, we obtain the minimum bounding box b enclosing $P \setminus (r_1 \cup r_2)$. To finish, we find a minimum-weight rectangle in R enclosing b; this is a "rectangle enclosure" query on R and can be solved in $\widetilde{O}(1)$ time, since it also reduces to orthogonal range searching (the rectangle $[x^-, x^+] \times [y^-, y^+]$ encloses the rectangle $[\xi^-, \xi^+] \times [\eta^-, \eta^+]$ in $\mathbb{R}^2$ iff the point (x^-, x^+, y^-, y^+) lies in the box $(-\infty, \xi^-] \times [\xi^+, \infty) \times (-\infty, \eta^-] \times [\eta^+, \infty)$ in $\mathbb{R}^4$). ◀

▶ **Theorem 2.** *Given a set P of n points and a set R of n weighted rectangles in $\mathbb{R}^2$, we can find 3 rectangles $r_1^*, r_2^*, r_3^* \in R$ of minimum total weight (if they exist), such that P is covered by $r_1^* \cup r_2^* \cup r_3^*$, in $\widetilde{O}(n^{16/9})$ time.*

Proof. Let B_0 be the minimum bounding box enclosing P (which touches 4 points). If a rectangle of R has an edge outside of B_0, we can eliminate that edge by extending the rectangle, making it unbounded.

Let g be a parameter to be determined later. Form a $g \times g$ (non-uniform) grid, where each column/row contains $O(n/g)$ rectangle vertices.

Step 1. For each pair of rectangles $r_1, r_2 \in R$ that have vertical edges in a common column or horizontal edges in a common row, we query the data structure in Lemma 1 to find a minimum-weight rectangle $r_3 \in R$ (if exists) such that $P \subset r_1 \cup r_2 \cup r_3$, and add the triple $r_1r_2r_3$ to a list L. The number of queried pairs r_1r_2 is $O(g \cdot (n/g)^2) = O(n^2/g)$, and so this step takes $\widetilde{O}(n^2/g)$ total time.

Step 2. For each rectangle $r_1 \in R$ and each of its horizontal (resp. vertical) edges e_1, define $\gamma^-(e_1)$ and $\gamma^+(e_1)$ to be the leftmost and rightmost (resp. bottommost and topmost) grid cell that intersects e_1 and contains a point of P not covered by r_1. We can naively find $\gamma^-(e_1)$ and $\gamma^+(e_1)$ by enumerating the $O(g)$ grid cells intersecting e_1 and performing $O(g)$ orthogonal range queries; this takes $\widetilde{O}(gn)$ total time. For each rectangle $r_2 \in R$ that has an edge intersecting $\gamma^-(e_1)$ or $\gamma^+(e_1)$, we query the data structure in Lemma 1 to find a minimum-weight rectangle $r_3 \in R$ (if exists) such that $P \subset r_1 \cup r_2 \cup r_3$, and add the triple $r_1r_2r_3$ to the list L. The total number of queried pairs r_1r_2 is $O(n \cdot n/g) = O(n^2/g)$, and so this step again takes $\widetilde{O}(n^2/g)$ total time. (This entire Step 2, and the definition of $\gamma^-(\cdot)$ and $\gamma^+(\cdot)$, might appear mysterious at first, but their significance will be revealed later in Step 3.)

Step 3. We guess the column containing each of the vertical edges of r_1^*, r_2^*, r_3^* and the row containing each of the horizontal edges of r_1^*, r_2^*, r_3^*; there are at most 12 edges and so $O(g^{12})$ choices. Actually, 4 of the 12 edges are eliminated after extension, and so the number of choices can be lowered to $O(g^8)$.

After guessing, we know which grid cells are completely inside r_1^*, r_2^*, r_3^* and which grid cells intersect which edges of r_1^*, r_2^*, r_3^*. We may assume that the vertical edges from different rectangles in $\{r_1^*, r_2^*, r_3^*\}$ are in different columns, and the horizontal edges from different rectangles in $\{r_1^*, r_2^*, r_3^*\}$ are in different rows: if not, $r_1^*r_2^*r_3^*$ would have already been found in Step 1. In particular, we know combinatorially what the arrangement of r_1^*, r_2^*, r_3^* looks like, even though we do not know the precise coordinates and identities of r_1^*, r_2^*, r_3^*.

We classify each grid cell γ into the following types (see Figure 1):

- Type A: γ is completely contained in some r_j^* ($j \in \{1, 2, 3\}$). Here, we *assign* γ to each such r_j^*.
- Type B: γ is not of type A, and intersects an edge of exactly one rectangle r_j^*. We *assign* γ to this r_j^*. Observe that points in $P \cap \gamma$ can only be covered by r_j^*.

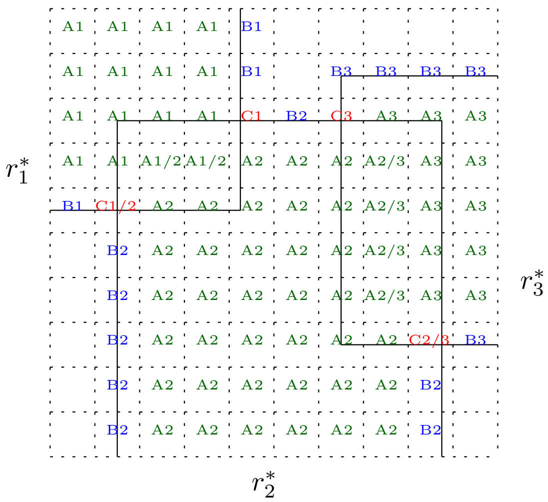

Figure 1 Proof of Theorem 2: grid cells in Step 3. The letter in a cell indicates its type (A, B, or C), and the number (or numbers) in a cell indicates the index (or indices) $j \in \{1, 2, 3\}$ of the rectangle r_j^* that the cell is assigned to.

- Type C: γ is not of type A, and intersects edges from two different rectangles in $\{r_1^*, r_2^*, r_3^*\}$. W.l.o.g., suppose that γ intersects a horizontal edge e_1^* of r_1^* and a vertical edge e_2^* of r_2^*; note that the intersection point $v^* = e_1^* \cap e_2^*$ lies on the boundary of the union $r_1^* \cup r_2^* \cup r_3^*$. By examining the arrangement of $\{r_1^*, r_2^*, r_3^*\}$, we know that at least one of the following is true: (i) we can walk horizontally from v^* to an endpoint of e_1^* (or a point at infinity) while staying on the boundary of $r_1^* \cup r_2^* \cup r_3^*$, or (ii) we can walk vertically from v^* to an endpoint of e_2^* (or a point at infinity) while staying on the boundary of $r_1^* \cup r_2^* \cup r_3^*$.

  If (i) is true, we *assign* γ to r_1^*. Observe that if there is a point in $P \cap \gamma$ not covered by r_1^* (and if the guesses are correct), then γ must be equal to $\gamma^-(e_1^*)$ or $\gamma^+(e_1^*)$ (as defined in Step 2), and so $r_1^* r_2^* r_3^*$ would have already been found in Step 2. This is because except for γ, the grid cells encountered while walking from v^* to that endpoint of e_1^* can intersect only r_1^* and so points in those cells can only be covered by r_1^*.

  If (ii) is true, we *assign* γ to r_2^* for a similar reason.

Note that there are at most $O(1)$ grid cells γ of type C; and the grid cells γ of type B form $O(1)$ contiguous blocks. Let ρ_j be the union of all grid cells assigned to r_j^*. Then ρ_j is a rectilinear polygon of $O(1)$ complexity. We compute the minimum/maximum x- and y-values of the points in $P \cap \rho_j$, by orthogonal range searching in $\widetilde{O}(1)$ time. As a result, we obtain the minimum bounding box b_j enclosing $P \cap \rho_j$. We find a minimum-weight rectangle $r_j \in R$ enclosing b_j, by a rectangle enclosure query (reducible to orthogonal range searching, as before). If $P \setminus (r_1 \cup r_2 \cup r_3) = \emptyset$ (testable by orthogonal range searching), we add the triple $r_1 r_2 r_3$ (which should coincide with $r_1^* r_2^* r_3^*$, if it has not been found earlier and if the guesses are correct) to L. The total time over all guesses is $\widetilde{O}(g^8)$.

At the end, we return a minimum-weight triple in L. The overall running time is $\widetilde{O}(g^8 + n^2/g + gn)$. Setting $g = n^{2/9}$ yields the theorem. ◀

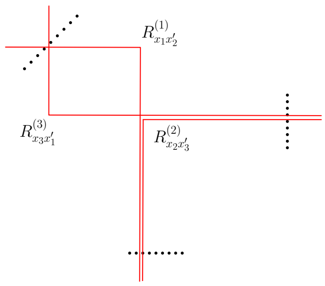

Figure 2 Reduction from the minimum-weight triangle problem to weighted size-3 set cover for orthants in $\mathbb{R}^2$.

3 Conditional Lower Bounds for Size-3 Set Cover for Boxes

In this section, we prove conditional lower bounds for size-3 set cover for boxes in certain dimensions (rectilinear discrete 3-center is related to size-3 set cover for unit hypercubes). We begin with the weighted version, which is more straightforward and has a simple proof, and serves as a good warm-up to the more challenging, unweighted version later.

3.1 Weighted size-3 set cover for unit squares in $\mathbb{R}^2$

An *orthant* (also called a dominance range) refers to a d-sided box in $\mathbb{R}^d$ which is unbounded along each of the d dimensions. (Note that orthants may be oriented in 2^d ways.) To obtain a lower bound for the unit square or unit hypercube case, it suffices to obtain a lower bound for the orthant case, since we can just replace each orthant with a hypercube with a sufficiently large side length M, and then rescale by a $1/M$ factor.

▶ **Theorem 3.** *Given a set P of n points and a set R of n weighted orthants in $\mathbb{R}^2$, finding 3 orthants in R of minimum total weight that cover P requires $\Omega(n^{3/2-\delta})$ time for any constant $\delta > 0$, assuming the APSP Hypothesis.*

Proof. The APSP Hypothesis is known to be equivalent [47] to the hypothesis that finding a minimum-weight triangle in a weighted graph with n vertices requires $\Omega(n^{3-\delta})$ time for any constant $\delta > 0$. We will reduce the minimum-weight triangle problem on a graph with n vertices and m edges ($m \in [n, n^2]$) to the weighted size-3 set cover problem for $O(m)$ points and orthants in $\mathbb{R}^2$. Thus, if there is an $O(m^{3/2-\delta})$-time algorithm for the latter problem, there would be an algorithm for the former problem with running time $O(m^{3/2-\delta}) \le O(n^{3-2\delta})$, refuting the hypothesis.

Let $G = (V, E)$ be the given weighted graph with n vertices and m edges. Without loss of generality, assume that all edge weights are in $[0, 0.1]$, and that $V \subset [0, 0.1]$, i.e., vertices are labelled by numbers that are rescaled to lie in $[0, 0.1]$. Assume that $0 \in V$ and $0.1 \in V$.

The reduction. For each vertex $t \in V$, we create three points $(t, 1+t)$, $(2, t)$, and $(1+t, -1)$ (call them of type 1, 2, and 3, respectively).

Create the following orthants in $\mathbb{R}^2$:

$$\begin{aligned}
\forall x_1x_2' \in E:\ R^{(1)}_{x_1x_2'} &= (-\infty, 1+x_2') \times (-\infty, 1+x_1] && \text{(type 1)}\\
\forall x_2x_3' \in E:\ R^{(2)}_{x_2x_3'} &= [1+x_2, \infty) \times (-\infty, x_3') && \text{(type 2)}\\
\forall x_3x_1' \in E:\ R^{(3)}_{x_3x_1'} &= (x_1', \infty) \times [x_3, \infty) && \text{(type 3)}
\end{aligned}$$

The weight of each orthant is set to be the number of points it covers plus the weight of the edge it represents. The total number of points and orthants is $O(n)$ and $O(m)$ respectively. The reduction is illustrated in Figure 2.

Correctness. We prove that the minimum-weight triangle in G has weight w (where $w \in [0, 0.3]$) iff the optimal weighted size-3 set cover has weight $3n + w$.

Any feasible solution (if exists) must use an orthant of each type, since the point $(0, 1)$ of type 1 (resp. the point $(2, 0.1)$ of type 2, and the point $(1.1, -1)$ of type 3) can only be covered by an orthant of type 1 (resp. 3 and 2). So, the three orthants in the optimal solution must be of the form $R^{(1)}_{x_1x'_2}$, $R^{(2)}_{x_2x'_3}$ and $R^{(3)}_{x_3x'_1}$ for some $x_1x'_2, x_2x'_3, x_3x'_1 \in E$.

If $x_1 < x'_1$, some point (of type 1) would be uncovered; on the other hand, if $x_1 > x'_1$, some point (of type 1) would be covered twice, and the total weight would then be at least $3n + 1$. Thus, $x_1 = x'_1$. Similarly, $x_2 = x'_2$ and $x_3 = x'_3$. So, $x_1x_2x_3$ forms a triangle in G. We conclude that the minimum-weight solution $R^{(1)}_{x_1x_2}$, $R^{(2)}_{x_2x_3}$ and $R^{(3)}_{x_3x_1}$ correspond to the minimum-weight triangle $x_1x_2x_3$ in G. ◀

3.2 Unweighted size-3 set cover for boxes in $\mathbb{R}^3$

Our preceding reduction uses weights to ensure equalities of two variables representing vertices. For the unweighted case, this does not work. We propose a different way to force equalities, by using an extra dimension and extra sides (i.e., using boxes instead of orthants), with some carefully chosen coordinate values.

▶ **Theorem 4.** *Given a set P of n points and a set R of n unweighted axis-aligned boxes in $\mathbb{R}^3$, deciding whether there exist 3 boxes in R that cover P requires $\Omega(n^{4/3-\delta})$ time for any constant $\delta > 0$, assuming the Sparse Triangle Hypothesis.*

Proof. We will reduce the triangle detection problem on a graph with m edges to the unweighted size-3 set cover problem for $O(m)$ points and boxes in $\mathbb{R}^3$. Thus, if there is an $O(m^{4/3-\delta})$-time algorithm for the latter problem, there would be an algorithm for the former problem with running time $O(m^{4/3-\delta})$, refuting the hypothesis.

Let $G = (V, E)$ be the given unweighted sparse graph with n vertices and m edges ($n \le m$). Without loss of generality, assume that $V \subset [0, 0.1]$, and $0 \in V$ and $0.1 \in V$.

The reduction. For each vertex $t \in V$, create six points

$$\begin{array}{rrrl}
(-1+t, & 0, & 2+t) & \text{(type 1)}\\
(1+t, & 0, & -2+t) & \text{(type 2)}\\
(2+t, & -1+t, & 0) & \text{(type 3)}\\
(-2+t, & 1+t, & 0) & \text{(type 4)}\\
(0, & 2+t, & -1+t) & \text{(type 5)}\\
(0, & -2+t, & 1+t) & \text{(type 6)}
\end{array}$$

Create the following boxes in $\mathbb{R}^3$:

$$\begin{array}{lcccccc}
\forall x_1x'_2 \in E: R^{(1)}_{x_1x'_2} & = & (-1+x_1, 1+x_1) & \times & [-2+x'_2, 2+x'_2] & \times & \mathbb{R}\\
\forall x_2x'_3 \in E: R^{(2)}_{x_2x'_3} & = & [-2+x'_3, 2+x'_3] & \times & \mathbb{R} & \times & (-1+x_2, 1+x_2)\\
\forall x_3x'_1 \in E: R^{(3)}_{x_3x'_1} & = & \mathbb{R} & \times & (-1+x_3, 1+x_3) & \times & [-2+x'_1, 2+x'_1]
\end{array}$$

(call them of type 1, 2, and 3, respectively).

Correctness. We prove that a size-3 set cover exists iff a triangle exists in G.

Any feasible solution (if exists) must use a box of each type, since the point $(-1,0,2)$ of type 1 (resp. the point $(2,-1,0)$ of type 3, and the point $(0,2,-1)$ of type 5) can only be covered by a box of type 3 (resp. 2 and 1). So, the three boxes in a feasible solution must be of the form $R^{(1)}_{x_1x'_2}$, $R^{(2)}_{x_2x'_3}$ and $R^{(3)}_{x_3x'_1}$ for some $x_1x'_2, x_2x'_3, x_3x'_1 \in E$.

Consider points of type 1 with the form $(-1+t, 0, 2+t)$. The box $R^{(2)}_{x_2x'_3}$ cannot cover any of them due to the third dimension. The box $R^{(1)}_{x_1x'_2}$ covers all such points corresponding to $t > x_1$, and the box $R^{(3)}_{x_3x'_1}$ covers all such points corresponding to $t \le x'_1$. So, all points of type 1 are covered iff $x_1 \le x'_1$. Similarly, all points of type 2 are covered iff $x'_1 \le x_1$. Thus, all points of type 1–2 are covered iff $x_1 = x'_1$. By a symmetric argument, all points of type 3–4 are covered iff $x_3 = x'_3$; and all points of type 5–6 are covered iff $x_2 = x'_2$. We conclude that a feasible solution exists iff a triangle $x_1x_2x_3$ exists in G. ◀

We remark that the boxes above can be made fat, with side lengths between 1 and a constant (by replacing $\mathbb{R}$ with an interval of a sufficiently large constant length).

3.3 Unweighted size-3 set cover for unit hypercubes in $\mathbb{R}^4$

Our preceding lower bound for unweighted size-3 set cover for boxes in $\mathbb{R}^3$ immediately implies a lower bound for orthants (and thus unit hypercubes) in $\mathbb{R}^6$, since the point (x, y, z) is covered by the box $[a^-, a^+] \times [b^-, b^+] \times [c^-, c^+]$ in $\mathbb{R}^3$ iff the point (x, x, y, y, z, z) is covered by the orthant $[a^-, \infty) \times (-\infty, a^+] \times [b^-, \infty) \times (-\infty, b^+] \times [c^-, \infty) \times (-\infty, c^+]$ in $\mathbb{R}^6$.

A question remains: can the dimension 6 be lowered? Intuitively, there seems to be some wastage in the above construction: there are several 0's in the coordinates of the points, and several $\mathbb{R}$'s in the definition of the boxes, and these get doubled after the transformation to 6 dimensions. However, it isn't clear how to rearrange coordinates to eliminate this wastage: we would have to give up this nice symmetry of our construction, and there are too many combinations to try. We ended up writing a computer program to exhaustively try all these different combinations, and eventually find a construction that lowers the dimension to 4! Once it is found, correctness is straightforward to check, as one can see in the proof below.

▶ **Theorem 5.** *Given a set P of n points and a set R of n unweighted orthants in $\mathbb{R}^4$, deciding whether there exists a size-3 set cover requires $\Omega(n^{4/3-\delta})$ time for any constant $\delta > 0$, assuming the Sparse Triangle Hypothesis.*

Proof. We will reduce the triangle detection problem on a graph with m edges to the unweighted size-3 set cover problem for $O(m)$ points and orthants in $\mathbb{R}^4$.

Let $G = (V, E)$ be the given unweighted graph with n vertices and m edges ($n \le m$). Without loss of generality, assume that $V \subset [0, 0.1]$, and $0 \in V$ and $0.1 \in V$.

The reduction. For each vertex $t \in V$, create six points

$$
\begin{array}{rrrrl}
(0+t, & 2+t, & -0.5, & -0.5) & \quad\text{(type 1)}\\
(2-t, & 0-t, & -0.5, & -0.5) & \quad\text{(type 2)}\\
(1-t, & 0.5, & 1+t, & 1.5) & \quad\text{(type 3)}\\
(0.5, & 1+t, & 0.5, & 2-t) & \quad\text{(type 4)}\\
(-0.5, & -0.5, & 2-t, & 0-t) & \quad\text{(type 5)}\\
(-0.5, & -0.5, & 0+t, & 1+t) & \quad\text{(type 6)}
\end{array}
$$

Create the following orthants in $\mathbb{R}^4$:

$$\begin{array}{llcl}
\forall x_1x_2' \in E: & R^{(1)}_{x_1x_2'} & = & [0+x_1,+\infty) \times [0-x_1,+\infty) \times (-\infty,1+x_2') \times (-\infty,2-x_2') \\
\forall x_2x_3' \in E: & R^{(2)}_{x_2x_3'} & = & (-\infty,1-x_2] \times (-\infty,1+x_2] \times (0+x_3',+\infty) \times (0-x_3',+\infty) \\
\forall x_3x_1' \in E: & R^{(3)}_{x_3x_1'} & = & (-\infty,2-x_1') \times (-\infty,2+x_1') \times (-\infty,2-x_3] \times (-\infty,1+x_3]
\end{array}$$

(call them of type 1, 2, and 3, respectively).

Correctness. We prove that a size-3 set cover exists iff a triangle exists in G.

Any feasible solution (if exists) must use an orthant of each type, as one can easily check (like before). So, the three orthants in a feasible solution must be of the form $R^{(1)}_{x_1x_2'}$, $R^{(2)}_{x_2x_3'}$ and $R^{(3)}_{x_3x_1'}$ for some $x_1x_2', x_2x_3', x_3x_1' \in E$.

Consider points of type 1 with the form $(0+t, 2+t, -0.5, -0.5)$. The orthant $R^{(2)}_{x_2x_3'}$ cannot cover any of them due to the third dimension. The orthant $R^{(1)}_{x_1x_2'}$ covers all such points corresponding to $t \geq x_1$, and the orthant $R^{(3)}_{x_3x_1'}$ covers all such points corresponding to $t < x_1'$. So, all points of type 1 are covered iff $x_1 \leq x_1'$. By similar arguments, it can be checked that all points of type 2 are covered iff $x_1 \geq x_1'$; all points of type 3 are covered iff $x_2 \leq x_2'$; all points of type 4 are covered iff $x_2 \geq x_2'$; all points of type 5 are covered iff $x_3 \leq x_3'$; all points of type 6 are covered iff $x_3 \geq x_3'$. We conclude that a feasible solution exists iff a triangle $x_1x_2x_3$ exists in G. ◀

In the computer search, we basically tried different choices of points with coordinate values of the form $c \pm t$ or c for some constant c, and orthants defined by intervals of the form $[c \pm x_j, +\infty)$ or $(-\infty, c \pm x_j]$ (closed or open) for some variable x_j (or x_j'). The constraints are not exactly easy to write down, but are self-evident as we simulate the correctness proof above. Naively, the number of cases is in the order of 10^{14}, but can be drastically reduced to about 10^7 with some optimization and careful pruning of the search space. The C++ code is not long (under 150 lines) and, after incorporating pruning, runs in under a second.

It is now straightforward to modify the above lower bound proof for unweighted orthants (or unit hypercubes) in $\mathbb{R}^4$ to the rectilinear discrete 3-center problem in $\mathbb{R}^4$. In the full version of the paper, we also prove a higher conditional lower bound for weighted size-6 set cover for rectangles in $\mathbb{R}^2$.

4 Conditional Lower Bound for Euclidean Discrete 2-Center

In this section, we prove our conditional lower bound for the Euclidean discrete 2-center problem in a sufficiently large constant dimension. The general structure of our proof is inspired by Bringmann et al.'s recent conditional hardness proof [12] for the problem of computing diameter of box intersection graphs in $\mathbb{R}^{12}$, specifically, testing whether the diameter is more than 2. (Despite the apparent dissimilarities of the two problems, what led us to initially suspect that the ideas there might be useful is that both problems are concerned with paths of length 2 in certain geometrically defined graphs, and both problems have a similar "quantifier structure", after unpacking the problem definitions.) Extra ideas are needed, as we are dealing with the Euclidean metric instead of boxes; we end up needing an extra dimension, with carefully tuned coordinate values, to make the proof work.

▶ **Theorem 6.** *For any constant $\delta > 0$, there is no $O(n^{2-\delta})$-time algorithm for Euclidean discrete 2-center in $\mathbb{R}^{13}$, assuming the Hyperclique Hypothesis.*

Proof. We will reduce the problem of detecting a 6-hyperclique in a 3-uniform hypergraph with n vertices, to the Euclidean discrete 2-center problem on $N = O(n^3)$ points in $\mathbb{R}^{13}$. Thus, if there is an $O(N^{2-\delta})$-time algorithm for the latter problem, there would be $O(n^{6-3\delta})$-time algorithm for the former problem, refuting the Hyperclique Hypothesis.

Let $G = (V, E)$ be the given 3-uniform hypergraph. By a standard color-coding technique [7], we may assume that G is 6-partite, i.e., V is partitioned into 6 parts $V_1, \ldots, V_6$, and each edge in E consists of 3 vertices from 3 different parts. The goal is to decide the existence of a 6-hyperclique, i.e., $(x_1, \ldots, x_6) \in V_1 \times \cdots \times V_6$ such that $\{x_i, x_j, x_k\} \in E$ for all distinct $i, j, k \in \{1, \ldots, 6\}$.

Without loss of generality, assume that $V \subset [0, 1]$, i.e., vertices are labelled by numbers that are rescaled to lie in $[0, 1]$. Let $f, g : [0, 1] \to [0, 1]$ be some injective functions satisfying $f(x)^2 + g(x)^2 = 1$. For example, we can take $f(x) = \cos x$ and $g(x) = \sin x$; or alternatively, avoiding trigonometric functions, $f(x) = x$ and $g(x) = \sqrt{1 - x^2}$; or avoiding irrational functions altogether, $f(x) = 2x/(x^2 + 1)$ and $g(x) = (x^2 - 1)/(x^2 + 1)$. (With the last two options, by rounding to $O(\log n)$ bits of precision, it is straightforward to make our reduction work in the standard integer word RAM model.)

The reduction. We construct the following set S of $O(n^3)$ points in $\mathbb{R}^{13}$:

1. For each $(x_1, x_2, x_3) \in V_1 \times V_2 \times V_3$ with $\{x_1, x_2, x_3\} \in E$, create the point
$$p_{x_1x_2x_3} \;=\; (f(x_1), g(x_1), f(x_2), g(x_2), f(x_3), g(x_3), 0, 0, 0, 0, 0, 0, 1)\,.$$
2. Similarly, for each $(x_4, x_5, x_6) \in V_4 \times V_5 \times V_6$ such that $\{x_4, x_5, x_6\} \in E$, create the point
$$q_{x_4x_5x_6} \;=\; (0, 0, 0, 0, 0, 0, f(x_4), g(x_4), f(x_5), g(x_5), f(x_6), g(x_6), -1)\,.$$
3. For each $(v_i, v_j, v_k) \in V_i \times V_j \times V_k$ with distinct i, j, k such that $\{v_i, v_j, v_k\} \notin E$, $\{i, j, k\} \neq \{1, 2, 3\}$, and $\{i, j, k\} \neq \{4, 5, 6\}$, create a point $z_{v_iv_jv_k}$: the coordinates in dimensions $2i-1, 2i$ are $-f(v_i), -g(v_i)$, and similarly the coordinates in dimensions $2j-1, 2j, 2k-1, 2k$ are $-f(v_j), -g(v_j), -f(v_k), -g(v_k)$, respectively; the 13-th coordinate is
$$\phi_{ijk} \;=\; |\{1, 2, 3\} \cap \{i, j, k\}| - 1.5 \;\in\; \{-0.5, 0.5\}\,;$$
and all other coordinates are 0. For example, if $i = 1, j = 2, k = 4$,
$$z_{v_1v_2v_4} \;=\; (-f(v_1), -g(v_1), -f(v_2), -g(v_2), 0, 0, -f(v_4), -g(v_4), 0, 0, 0, 0, 0.5)\,.$$
4. Finally, add two auxiliary points $s_\pm = (0, \ldots, 0, \pm 3.5)$.

We solve the discrete 2-center problem on the above point set S, and return true iff the minimum radius is strictly less than $\sqrt{10.25}$.

Correctness. Suppose there exists a 6-hyperclique $(x_1, \ldots, x_6) \in V_1 \times \cdots \times V_6$ in G. We claim that every point of S has distance strictly less than $\sqrt{10.25}$ from $p_{x_1x_2x_3}$ or $q_{x_4x_5x_6}$. Thus, S can be covered by 2 balls centered at $p_{x_1x_2x_3}$ and $q_{x_4x_5x_6}$ with radius less than $\sqrt{10.25}$. To verify the claim, consider a point $z_{v_1v_2v_4} \in S$ for a triple $(v_1, v_2, v_4) \in V_1 \times V_2 \times V_4$ with $\{v_1, v_2, v_4\} \notin E$. Observe that the distance between the points $(f(v_\ell), g(v_\ell))$ and $(-f(x_\ell), -g(x_\ell))$ in $\mathbb{R}^2$ is at most 2, with equality iff $v_\ell = x_\ell$. On the other hand, the distance between $(f(v_\ell), g(v_\ell))$ and $(0, 0)$ is 1, and the distance between $(0, 0)$ and $(-f(x_\ell), -g(x_\ell))$ is 1. Thus,

$$\|z_{v_1v_2v_4} - p_{x_1x_2x_3}\|^2 \;\leq\; 2^2 + 2^2 + 1 + 1 + 0 + 0 + (0.5 - 1)^2 \;\leq\; 10.25,$$

with equality iff $v_1 = x_1$ and $v_2 = x_2$. Furthermore,

$$\|z_{v_1v_2v_4} - q_{x_4x_5x_6}\|^2 \ \le\ 1+1+0+2^2+1+1+(0.5+1)^2 \ \le\ 10.25,$$

with equality iff $v_4 = x_4$. Since $\{x_1, x_2, x_4\} \in E$, we cannot have simultaneously $v_1 = x_1$, $v_2 = x_2$, and $v_4 = x_4$. So, $z_{v_1v_2v_4}$ has distance strictly less than $\sqrt{10.25}$ from $p_{x_1x_2x_3}$ or $q_{x_4x_5x_6}$. Similarly, the same holds for $z_{v_iv_jv_k} \in S$ for all other choices of i, j, k. Points $p_{x'_1x'_2x'_3} \in S$ have distance at most $\sqrt{2+2+2+0+0+0+0} < \sqrt{10.25}$ from $p_{x_1x_2x_3}$, and similarly, points $q_{x'_4x'_5x'_6} \in S$ have distance less than $\sqrt{10.25}$ from $q_{x_4x_5x_6}$. Finally, the auxiliary point s_+ has distance at most $\sqrt{1+1+1+0+0+0+2.5^2} < \sqrt{10.25}$ from $p_{x_1x_2x_3}$, and similarly the point s_- has distance less than $\sqrt{10.25}$ from $q_{x_4x_5x_6}$.

On the reverse direction, suppose that the minimum radius for the discrete 2-center problem on S is strictly less than $\sqrt{10.25}$. Note that the distance between s_+ and $z_{v_iv_jv_k}$ is at least $\sqrt{1+1+1+0+0+0+3^2} > \sqrt{10.25}$, and the distance between s_+ and $q_{x_4x_5x_6}$ is at least $\sqrt{0+0+0+1+1+1+4.5^2} > \sqrt{10.25}$. Thus, in order to cover s_+, one of the two centers must be equal to $p_{x_1x_2x_3}$ for some $\{x_1, x_2, x_3\} \in E$. Similarly, in order to cover s_-, the other center must be equal to $q_{x_4x_5x_6}$ for some $\{x_4, x_5, x_6\} \in E$. Then for every $(v_1, v_2, v_4) \in V_1 \times V_2 \times V_4$ with $\{v_1, v_2, v_4\} \notin E$, the point $z_{v_1v_2v_4}$ has distance strictly less than $\sqrt{10.25}$ from $p_{x_1x_2x_3}$ or $q_{x_4x_5x_6}$. By the above argument, we cannot have $v_1 = x_1$ and $v_2 = x_2$ and $v_4 = x_4$. It follows that $\{x_1, x_2, x_4\} \in E$. Similarly, $\{x_i, x_j, x_k\} \in E$ for all other choices of i, j, k. We conclude that $\{x_1, \ldots, x_6\}$ is a 6-hyperclique. ◀

From the same proof (after rescaling), we immediately get a near-quadratic conditional lower bound for unweighted size-2 geometric set cover for unit balls in $\mathbb{R}^{13}$. In the full version of the paper, we extend the proof to Euclidean discrete k-center for larger constant k, with more technical effort and more delicate handling of the extra dimensions. This is interesting: discrete k-center seems even farther away from graph diameter, but in a way, our proof shows that discrete k-center is a better problem to illustrate the full power of Bringmann et al.'s technique [12].

In the full version of the paper, we also adapt the approach to prove a conditional lower bound for size-2 maximum coverage for boxes. The proof uses a different way to enforce conditions like $\{x_1, x_2, x_4\} \in E$, via an interesting counting argument – we encourage the readers to take a look at the full version.

5 Conclusions

In this paper, we have obtained a plethora of nontrivial new results on a fundamental class of problems in computational geometry related to discrete k-center and size-k geometric set cover for small values of k. (See Tables 1–3.) In particular, we have a few results where the upper bounds and conditional lower bounds are close:

- For weighted size-3 set cover for rectangles in $\mathbb{R}^2$, we have given the first subquadratic $\widetilde{O}(n^{7/4})$-time algorithm, and an $\Omega(n^{3/2-\delta})$ lower bound under the APSP Hypothesis.
- For Euclidean discrete k-center (or unweighted size-k set cover for unit balls) in $\mathbb{R}^{O(k)}$, we have proved an $\Omega(n^{k-\delta})$ lower bound under the Hyperclique Hypothesis, which is near optimal if $\omega = 2$.
- For size-2 maximum coverage for boxes in a sufficiently large constant dimension, we have proved an $\Omega(n^{2-\delta})$ lower bound under the Hyperclique Hypothesis, which is near optimal.

For all of our results, we have managed to find simple proofs (each with 1–3 pages). We view the simplicity and accessibility of our proofs as an asset – they would make good examples illustrating fine-grained complexity techniques in computational geometry. Generally speaking, there has been considerable development on fine-grained complexity in the broader algorithms community over the last decade [46], but to a lesser extent in computational geometry. A broader goal of this paper is to encourage more work at the intersection of these two areas. We should emphasize that while our conditional lower bound proofs may appear simple in hindsight, they are not necessarily easy to come up with; for example, see one of our proofs that require computer-assisted search (Theorem 5).

As many versions of the problems studied here still do not have matching upper and lower bounds, our work raises many interesting open questions. For example:

- Is it possible to make our subquadratic algorithm for rectilinear discrete 3-center in $\mathbb{R}^2$ work in dimension 3 or higher?
- Is it possible to make our conditional lower bound proof for rectilinear discrete 3-center in $\mathbb{R}^4$ work in dimension 2 or 3?
- Is it possible to make our conditional lower bound for Euclidean discrete 2-center in $\mathbb{R}^{13}$ work in dimension 3?
- Is it possible to make our conditional lower bound for size-2 maximum coverage for boxes in $\mathbb{R}^{12}$ work in dimension 2 or 3?
- Although we have ruled out subquadratic algorithms for Euclidean discrete 2-center in $\mathbb{R}^{13}$, could geometry still help in beating n^ω time if $\omega > 2$?

We should remark that some of these questions could be quite difficult. In fine-grained complexity, there are many examples of basic problems that still do not have tight conditional lower bounds (to mention one well-known geometric example, Künnemann's recent FOCS'22 paper [36] has finally obtained a near-optimal conditional lower bound for Klee's measure problem in $\mathbb{R}^3$, but tight lower bounds in dimension 4 and higher are still not known for non-combinatorial algorithms). Still, we hope that our work would inspire more progress in both upper and lower bounds for this rich class of problems.

References

1 Amir Abboud, Karl Bringmann, Seri Khoury, and Or Zamir. Hardness of approximation in P via short cycle removal: Cycle detection, distance oracles, and beyond. In *Proc. 54th Annual ACM Symposium on Theory of Computing (STOC)*, pages 1487–1500, 2022. doi:10.1145/3519935.3520066.

2 Pankaj K. Agarwal, Rinat Ben Avraham, and Micha Sharir. The 2-center problem in three dimensions. *Comput. Geom.*, 46(6):734–746, 2013. Preliminary version in SoCG'10. doi:10.1016/j.comgeo.2012.11.005.

3 Pankaj K. Agarwal, Alon Efrat, and Micha Sharir. Vertical decomposition of shallow levels in 3-dimensional arrangements and its applications. *SIAM J. Comput.*, 29(3):912–953, 1999. doi:10.1137/S0097539795295936.

4 Pankaj K. Agarwal and Jeff Erickson. Geometric range searching and its relatives. In *Advances in Discrete and Computational Geometry*, volume 223 of *Contemporary Mathematics*, pages 1–56. AMS Press, 1999. URL: http://jeffe.cs.illinois.edu/pubs/survey.html.

5 Pankaj K. Agarwal and Cecilia Magdalena Procopiuc. Exact and approximation algorithms for clustering. *Algorithmica*, 33(2):201–226, 2002. doi:10.1007/s00453-001-0110-y.

6 Pankaj K Agarwal, Micha Sharir, and Emo Welzl. The discrete 2-center problem. *Discrete & Computational Geometry*, 20(3):287–305, 1998. Preliminary version in SoCG'97.

7 Noga Alon, Raphael Yuster, and Uri Zwick. Color-coding. *J. ACM*, 42(4):844–856, 1995. doi:10.1145/210332.210337.

8 Noga Alon, Raphael Yuster, and Uri Zwick. Finding and counting given length cycles. *Algorithmica*, 17(3):209–223, 1997. doi:10.1007/BF02523189.

9 Ashwinkumar Badanidiyuru, Robert Kleinberg, and Hooyeon Lee. Approximating low-dimensional coverage problems. In *Proc. 28th ACM Symposium on Computational Geometry (SoCG)*, pages 161–170, 2012. doi:10.1145/2261250.2261274.

10 Sergei Bespamyatnikh and David G. Kirkpatrick. Rectilinear 2-center problems. In *Proc. 11th Canadian Conference on Computational Geometry (CCCG)*, 1999. URL: http://www.cccg.ca/proceedings/1999/fp55.pdf.

11 Sergei Bespamyatnikh and Michael Segal. Rectilinear static and dynamic discrete 2-center problems. In *Proc. 6th Workshop on Algorithms and Data Structures (WADS)*, pages 276–287, 1999. doi:10.1007/3-540-48447-7_28.

12 Karl Bringmann, Sándor Kisfaludi-Bak, Marvin Künnemann, André Nusser, and Zahra Parsaeian. Towards sub-quadratic diameter computation in geometric intersection graphs. In *Proc. 38th International Symposium on Computational Geometry (SoCG)*, pages 21:1–21:16, 2022. doi:10.4230/LIPIcs.SoCG.2022.21.

13 Karl Bringmann, Sándor Kisfaludi-Bak, Michal Pilipczuk, and Erik Jan van Leeuwen. On geometric set cover for orthants. In *Proc. 27th Annual European Symposium on Algorithms (ESA)*, pages 26:1–26:18, 2019. doi:10.4230/LIPIcs.ESA.2019.26.

14 Sergio Cabello, Panos Giannopoulos, Christian Knauer, Dániel Marx, and Günter Rote. Geometric clustering: Fixed-parameter tractability and lower bounds with respect to the dimension. *ACM Trans. Algorithms*, 7(4):43:1–43:27, 2011. Preliminary version in SODA'08. doi:10.1145/2000807.2000811.

15 Timothy M. Chan. Geometric applications of a randomized optimization technique. *Discret. Comput. Geom.*, 22(4):547–567, 1999. doi:10.1007/PL00009478.

16 Timothy M. Chan. More planar two-center algorithms. *Computational Geometry*, 13(3):189–198, 1999.

17 Timothy M. Chan. Finding triangles and other small subgraphs in geometric intersection graphs. In *Proc. 34th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2023. To appear. URL: https://arxiv.org/abs/2211.05345.

18 Timothy M. Chan and Qizheng He. Faster approximation algorithms for geometric set cover. In *Proc. 36th International Symposium on Computational Geometry (SoCG)*, pages 27:1–27:14, 2020. doi:10.4230/LIPIcs.SoCG.2020.27.

19 Timothy M. Chan, Qizheng He, and Yuancheng Yu. On the fine-grained complexity of small-size geometric set cover and discrete k-center for small k. arXiv preprint, 2023. arXiv:2305.01892.

20 Timothy M. Chan, Virginia Vassilevska Williams, and Yinzhan Xu. Hardness for triangle problems under even more believable hypotheses: Reductions from Real APSP, Real 3SUM, and OV. In *Proc. 54th ACM Symposium on Theory of Computing (STOC)*, pages 1501–1514, 2022. doi:10.1145/3519935.3520032.

21 Bernard Chazelle and Jirı Matoušek. On linear-time deterministic algorithms for optimization problems in fixed dimension. *Journal of Algorithms*, 21(3):579–597, 1996.

22 Rajesh Chitnis and Nitin Saurabh. Tight lower bounds for approximate & exact k-center in $\mathbb{R}^d$. In *Proc. 38th International Symposium on Computational Geometry (SoCG)*, pages 28:1–28:15, 2022. doi:10.4230/LIPIcs.SoCG.2022.28.

23 Jongmin Choi and Hee-Kap Ahn. Efficient planar two-center algorithms. *Comput. Geom.*, 97:101768, 2021. doi:10.1016/j.comgeo.2021.101768.

24 Kenneth L. Clarkson. Las Vegas algorithms for linear and integer programming when the dimension is small. *Journal of the ACM*, 42(2):488–499, 1995.

25 Don Coppersmith. Rapid multiplication of rectangular matrices. *SIAM J. Comput.*, 11(3):467–471, 1982. doi:10.1137/0211037.

26 Mark de Berg, Otfried Cheong, Marc J. van Kreveld, and Mark H. Overmars. *Computational Geometry: Algorithms and Applications*. Springer, 3rd edition, 2008. URL: https://www.worldcat.org/oclc/227584184.

27 Martin E. Dyer. On a multidimensional search technique and its application to the Euclidean one-centre problem. *SIAM Journal on Computing*, 15(3):725–738, 1986.

28 Friedrich Eisenbrand and Fabrizio Grandoni. On the complexity of fixed parameter clique and dominating set. *Theor. Comput. Sci.*, 326(1-3):57–67, 2004. doi:10.1016/j.tcs.2004.05.009.

29 David Eppstein. Faster construction of planar two-centers. In *Proc. 8th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 131–138, 1997. URL: http://dl.acm.org/citation.cfm?id=314161.314198.

30 Greg N. Frederickson and Donald B. Johnson. The complexity of selection and ranking in $X + Y$ and matrices with sorted columns. *J. Comput. Syst. Sci.*, 24(2):197–208, 1982. doi:10.1016/0022-0000(82)90048-4.

31 R. Z. Hwang, R. C. Chang, and Richard C. T. Lee. The searching over separators strategy to solve some NP-hard problems in subexponential time. *Algorithmica*, 9(4):398–423, 1993. doi:10.1007/BF01228511.

32 R. Z. Hwang, Richard C. T. Lee, and R. C. Chang. The slab dividing approach to solve the Euclidean p-center problem. *Algorithmica*, 9(1):1–22, 1993. doi:10.1007/BF01185335.

33 Ce Jin and Yinzhan Xu. Removing additive structure in 3SUM-based reductions. *CoRR*, abs/2211.07048, 2022. doi:10.48550/arXiv.2211.07048.

34 Matthew J. Katz, Klara Kedem, and Michael Segal. Discrete rectilinear 2-center problems. *Comput. Geom.*, 15(4):203–214, 2000. doi:10.1016/S0925-7721(99)00052-8.

35 Matthew J. Katz and Frank Nielsen. On piercing sets of objects. In *Proc. 12th Annual Symposium on Computational Geometry (SoCG)*, pages 113–121, 1996. doi:10.1145/237218.237253.

36 Marvin Künnemann. A tight (non-combinatorial) conditional lower bound for Klee's measure problem in 3D. In *Proc. 63rd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 555–566, 2022. doi:10.1109/FOCS54457.2022.00059.

37 Andrea Lincoln, Virginia Vassilevska Williams, and R. Ryan Williams. Tight hardness for shortest cycles and paths in sparse graphs. In *Proc. 29th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1236–1252, 2018. doi:10.1137/1.9781611975031.80.

38 Dániel Marx. Efficient approximation schemes for geometric problems? In *Proc. 13th Annual European Symposium of Algorithms (ESA)*, pages 448–459, 2005. doi:10.1007/11561071_41.

39 Dániel Marx and Michal Pilipczuk. Optimal parameterized algorithms for planar facility location problems using voronoi diagrams. In *Proc. 23rd Annual European Symposium on Algorithms (ESA)*, pages 865–877, 2015. doi:10.1007/978-3-662-48350-3_72.

40 Jirí Matoušek. Reporting points in halfspaces. *Comput. Geom.*, 2:169–186, 1992. doi:10.1016/0925-7721(92)90006-E.

41 Nimrod Megiddo. Linear-time algorithms for linear programming in R^3 and related problems. *SIAM Journal on Computing*, 12(4):759–776, 1983.

42 Doron Nussbaum. Rectilinear p-piercing problems. In *Proc. International Symposium on Symbolic and Algebraic Computation (ISSAC)*, pages 316–323, 1997. doi:10.1145/258726.258828.

43 Mihai Pătraşcu. Towards polynomial lower bounds for dynamic problems. In *Proc. 42nd ACM Symposium on Theory of Computing (STOC)*, pages 603–610, 2010. doi:10.1145/1806689.1806772.

44 Micha Sharir. A near-linear algorithm for the planar 2-center problem. *Discrete & Computational Geometry*, 18(2):125–134, 1997. Preliminary version in SoCG'96.

45 Micha Sharir and Emo Welzl. Rectilinear and polygonal p-piercing and p-center problems. In *Proc. 12th Annual Symposium on Computational Geometry (SoCG)*, pages 122–132, 1996. doi:10.1145/237218.237255.

46 Virginia Vassilevska Williams. On some fine-grained questions in algorithms and complexity. In *Proceedings of the ICM*, volume 3, pages 3431–3472. World Scientific, 2018. URL: https://people.csail.mit.edu/virgi/eccentri.pdf.

47 Virginia Vassilevska Williams and R. Ryan Williams. Subcubic equivalences between path, matrix, and triangle problems. *J. ACM*, 65(5):27:1–27:38, 2018. Preliminary version in FOCS'10. `doi:10.1145/3186893`.

48 Virginia Vassilevska Williams and Yinzhan Xu. Monochromatic triangles, triangle listing and APSP. In *Proc. 61st IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 786–797, 2020. `doi:10.1109/FOCS46700.2020.00078`.

49 Haitao Wang. On the planar two-center problem and circular hulls. In *Proc. 36th International Symposium on Computational Geometry (SoCG)*, pages 68:1–68:14, 2020. `doi:10.4230/LIPIcs.SoCG.2020.68`.

50 Emo Welzl. Smallest enclosing disks (balls and ellipsoids). In *New Results and New Trends in Computer Science*, pages 359–370. Springer, 1991.