

Towards Sub-Quadratic Diameter Computation in Geometric Intersection Graphs

Karl Bringmann ✉

Universität des Saarlandes, Saarbrücken, Germany

Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany

Sándor Kisfaludi-Bak ✉

Aalto University, Espoo, Finland

Marvin Künnemann ✉

Institute for Theoretical Studies, ETH Zürich, Switzerland

André Nusser ✉ 

BARC, University of Copenhagen, Denmark

Zahra Parsaeian ✉

Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany

Abstract

We initiate the study of diameter computation in geometric intersection graphs from the fine-grained complexity perspective. A geometric intersection graph is a graph whose vertices correspond to some shapes in d -dimensional Euclidean space, such as balls, segments, or hypercubes, and whose edges correspond to pairs of intersecting shapes. The diameter of a graph is the largest distance realized by a pair of vertices in the graph.

Computing the diameter in near-quadratic time is possible in several classes of intersection graphs [Chan and Skrepetos 2019], but it is not at all clear if these algorithms are optimal, especially since in the related class of planar graphs the diameter can be computed in $\tilde{O}(n^{5/3})$ time [Cabello 2019, Gawrychowski et al. 2021].

In this work we (conditionally) rule out sub-quadratic algorithms in several classes of intersection graphs, i.e., algorithms of running time $\mathcal{O}(n^{2-\delta})$ for some $\delta > 0$. In particular, there are no sub-quadratic algorithms already for fat objects in small dimensions: unit balls in \mathbb{R}^3 or congruent equilateral triangles in \mathbb{R}^2 . For unit segments and congruent equilateral triangles, we can even rule out strong sub-quadratic approximations already in \mathbb{R}^2 . It seems that the hardness of approximation may also depend on dimensionality: for axis-parallel unit hypercubes in \mathbb{R}^{12} , distinguishing between diameter 2 and 3 needs quadratic time (ruling out $(3/2-\varepsilon)$ -approximations), whereas for axis-parallel unit squares, we give an algorithm that distinguishes between diameter 2 and 3 in near-linear time.

Note that many of our lower bounds match the best known algorithms up to sub-polynomial factors. Ultimately, this fine-grained perspective may enable us to determine for which shapes we can have efficient algorithms and approximation schemes for diameter computation.

2012 ACM Subject Classification Theory of computation \rightarrow Computational geometry

Keywords and phrases Hardness in P, Geometric Intersection Graph, Graph Diameter, Orthogonal Vectors, Hyperclique Detection

Digital Object Identifier 10.4230/LIPIcs.SoCG.2022.21

Related Version *Full Version:* <https://arxiv.org/abs/2203.03663>

Funding *Karl Bringmann:* This work is part of the project TIPEA that has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant agreement No. 850979).

Sándor Kisfaludi-Bak: Part of this research was conducted while the author was at the Max Planck Institute for Informatics, and part of it while he was at the Institute for Theoretical Studies, ETH Zürich.



© Karl Bringmann, Sándor Kisfaludi-Bak, Marvin Künnemann, André Nusser, and Zahra Parsaeian;

licensed under Creative Commons License CC-BY 4.0


38th International Symposium on Computational Geometry (SoCG 2022).

Editors: Xavier Goaoc and Michael Kerber; Article No. 21; pp. 21:1–21:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



Marvin Künnemann: Research supported by Dr. Max Rössler, by the Walter Haefner Foundation, and by the ETH Zürich Foundation. Part of this research was conducted while the author was at the Max Planck Institute for Informatics.

André Nusser: Part of this research was conducted while the author was at Saarbrücken Graduate School of Computer Science and Max Planck Institute for Informatics. The author is supported by the VILLUM Foundation grant 16582.

1 Introduction

The diameter of a simple graph $G = (V, E)$ is the largest distance realized by a pair of its vertices; formally, it is $\text{diam}(G) = \max_{u, v \in V} \text{dist}_G(u, v)$, where $\text{dist}_G(u, v)$ is the number of edges on a shortest path from u to v . It is one of the crucial parameters of a graph that can be computed in polynomial time. Geometric intersection graphs are the standard model for wireless communication networks [34], but more abstractly, they can be used to represent networks where the connection of nodes relies on proximity in some metric space. For a (slightly oversimplified) example, consider a set of devices in the plane capable of receiving and transmitting information in a range of radius 2. These devices form a communication network that is a unit disk graph. Indeed, two devices can communicate with each other if and only if their distance is at most 2, i.e., if the unit disks centered at the devices have a non-empty intersection. For our purposes, the underlying metric space will be d -dimensional Euclidean space (henceforth denoted by \mathbb{R}^d), and we will consider intersection graphs of common objects such as balls and segments. For a set F of objects in \mathbb{R}^d (that is, $F \subset 2^{\mathbb{R}^d}$), the corresponding intersection graph $G[F]$ has vertex set F and edge set $\{uv \mid u, v \in F, u \cap v \neq \emptyset\}$.

Computing the diameter in geometric intersection graphs is an important task: if the graph represents a communication network, then the diameter of the network can help estimate the time required to spread information in the network, as the information needs to go through up to $\text{diam}(G)$ links to reach its destination. In large networks, it is also indispensable to have near-linear time algorithms; it is therefore natural to study if a given class of geometric intersection graphs admits a near-linear time algorithm for exact or approximate diameter computation.

The extensive literature on diameter computation serves as a good starting point. The diameter of an n -vertex (unweighted) graph can be computed in $\mathcal{O}(n^\omega \log n)$ expected time, where $\omega < 2.37286$ is the exponent of matrix multiplication [41]. If the graph has m edges, then the diameter can also be computed in $\mathcal{O}(mn)$ time [42], which gives a near-quadratic running time of $\tilde{\mathcal{O}}(n^2)$ in case of sparse graphs, i.e., when $m = \tilde{\mathcal{O}}(n)$. In fact, these algorithms are capable of computing not only the diameter, but also all pairwise distances in a graph, known as the all pairs shortest paths problem.

On the negative side, we know that computing the diameter of a graph cannot be done in $\mathcal{O}(n^{2-\varepsilon})$ time under the Orthogonal Vectors Hypothesis¹ (OV); in fact, deciding if the diameter of a sparse graph is at most 2 or at least 3 requires $n^{2-o(1)}$ time under OV [40]², which rules out sub-quadratic $(3/2 - \varepsilon)$ -approximations for all $\varepsilon > 0$.

In special graph classes however it is possible to compute the diameter in sub-quadratic time. In planar graphs, an algorithm with running time $\mathcal{O}(n^2)$ is very easy: one can just run n breadth-first searches, each of which take linear time because the number of edges is $\mathcal{O}(n)$.

¹ See Section 2 for the definitions and some background on the hypotheses used in our lower bounds.

² More precisely, Roditty and Vassilevska-Williams [40] give a reduction from k -Dominating Set, which can be adapted to a reduction from OV as described in the beginning of Section 4.

It has been a long-standing open problem whether a truly sub-quadratic algorithm exists for diameter computation, until the breakthrough of Cabello [12], who used Voronoi diagrams in planar graphs. The technique was later improved by Gawrychowski *et al.* [29], who obtained a running time of $\tilde{O}(n^{5/3})$.

Certain geometric intersection graphs often behave similarly to planar graphs. The most widely studied classes, (unit) disk and ball graphs admit approximation schemes for maximum independent set, maximum dominating set, and several other problems [30, 31, 14], with techniques similar to planar graphs. Unlike planar graphs, geometric intersection graphs can have arbitrarily large cliques, but at least the maximum clique can be approximated efficiently [7]. In fact, planar graphs are special disk intersection graphs by the circle packing theorem [33]. When it comes to computing the diameter, the similarity with planar graphs is not so easy to see. Even getting near-quadratic diameter algorithms is non-trivial, as geometric intersection graphs can be arbitrarily dense.

Chan and Skrepetos [16] provide near-quadratic ($\tilde{O}(n^2)$) APSP algorithms for several graph classes, including disks, axis-parallel segments, and fat triangles in the plane, and cubes and boxes in constant-dimensional space. Unit disk graphs have a “weakly” sub-quadratic algorithm (that is poly-logarithmically faster than $\mathcal{O}(n^2)$) [15]. We are not aware of any $\mathcal{O}(n^{2-\epsilon})$ algorithms for computing the diameter in intersection graphs of any planar shape.

Further related work. While computing the diameter is known to require time $n^{2\pm o(1)}$ already on sparse graphs (assuming the OV Hypothesis), an extensive line of research including [3, 40, 19, 13, 5, 25, 9, 24, 37, 8, 23] studies the (non-)existence of faster approximation algorithms. On the positive side, this includes in particular a folklore 2-approximation in time $\tilde{O}(m)$ and a 3/2-approximation in time $\tilde{O}(m^{3/2})$ [3, 40, 19], both already for weighted digraphs. Remarkably, these algorithms can be shown to be tight: [40, 5] establish that the 3/2-approximation in time $\tilde{O}(m^{3/2})$ cannot be improved in either approximation guarantee or running time (assuming the k -OV Hypothesis), already for unweighted undirected graphs. The near-linear time 2-approximation is conditionally optimal as well: For unweighted directed graphs, this has been proven independently in [24, 37]. For unweighted undirected graphs, following further work [8], a resolution has been announced only very recently [23]. Thus, approximating the diameter in sparse graphs is quite well understood, including detailed insights into the full accuracy-time trade-off. In the context of our work, the challenge is to obtain a similar understanding for our setting of unweighted, undirected *geometric* graphs, which are non-sparse in general.

Note that for graph classes that are non-sparse, a natural question is whether diameter can be computed in $\mathcal{O}(m+n)$ time, i.e., linear time in the number of edges plus vertices. The question has been studied by several authors: using a variant of breadth-first search called *lexicographic breadth-first search*, one can find a vertex of very large eccentricity. In some classes, we now know that there is an $\mathcal{O}(m+n)$ algorithm for diameter: notably, this holds in interval graphs as well as {claw,asteroidal triple}-free graphs [27, 10]. In many other graph classes (such as chordal graphs and asteroidal-triple-free graphs) we can get approximations for the diameter that differ only by a small additive constant from the optimum [27, 26, 21]. See [22] for an overview on the connection of lexicographic BFS and diameter, and see [20] for a survey on lexicographic BFS.

Another related direction is to consider edge weighted graph classes. In some classes of geometric intersection graphs there is a natural weighting to consider: for example in ball graphs, it is customary to draw the graph edges with straight segments that connect the centers of the two adjacent disks. The edges then have a natural weighting by their Euclidean

length. This was considered for unit disk graphs in the plane by Gao and Zhang [28], who obtained a $(1 + \varepsilon)$ -approximation for DIAMETER in $\mathcal{O}(n^{3/2})$ time for any fixed $\varepsilon > 0$. A faster $(1 + \varepsilon)$ -approximation with running time $\mathcal{O}(n \log^2 n)$ for any fixed $\varepsilon > 0$ was given by Chan and Skrepetos [17]. Since the underlying graph is not changed by this weighting, it is natural to think that similar results should be possible also for unweighted unit disk graphs. It remains an open question whether the complexity of diameter computation is influenced by the presence of these Euclidean weights.

Our results. In this article, we show that most of the results of Chan and Skrepetos [16] cannot be significantly improved under standard complexity-theoretic assumptions, even if we are only interested in the diameter instead of all pairs shortest paths. In particular, we rule out sub-quadratic diameter algorithms for fat triangles and axis-aligned segments in the plane, as well as for unit cubes in \mathbb{R}^3 , leaving only their $\tilde{\mathcal{O}}(n^{7/3})$ algorithm for arbitrary segments in \mathbb{R}^2 as well as their $\tilde{\mathcal{O}}(n^2)$ algorithm for disks without a matching lower bound.

The DIAMETER problem has as input a set of geometric objects in \mathbb{R}^d and a number k ; the goal is to decide whether the diameter of the intersection graph of the objects is at most k . The DIAMETER- t problem is the same problem, but with k set to the constant number t . We show the following lower bounds.

- **Theorem 1.** *For all $\delta > 0$ there is no $\mathcal{O}(n^{2-\delta})$ time algorithm for*
 - DIAMETER-3 in intersection graphs of unit segments in \mathbb{R}^2 under the OV Hypothesis.
 - DIAMETER-3 in intersection graphs of congruent equilateral triangles in \mathbb{R}^2 under the OV Hypothesis.
 - DIAMETER in intersection graphs of unit balls in \mathbb{R}^3 under the OV Hypothesis.
 - DIAMETER in intersection graphs of axis-parallel unit cubes in \mathbb{R}^3 under the OV Hypothesis.
 - DIAMETER in intersection graphs of axis-parallel line segments in \mathbb{R}^2 under the OV Hypothesis.
 - DIAMETER-2 in intersection graphs of axis-parallel hypercubes in \mathbb{R}^{12} under the Hyperclique Hypothesis.

Our results imply lower bounds for approximations. (See Section 4.2 for a short proof.)

- **Corollary 2.** *Under the Orthogonal Vectors and Hyperclique Hypotheses, for all $\delta, \varepsilon > 0$ there is no $\mathcal{O}(n^{2-\delta})$ time $(4/3 - \varepsilon)$ -approximation for DIAMETER in intersection graphs of unit segments or congruent equilateral triangles in \mathbb{R}^2 , and no $(3/2 - \varepsilon)$ -approximation in intersection graphs of axis-parallel hypercubes in $\mathbb{R}^{\geq 12}$. Furthermore, for all $\delta > 0$ there is no $\mathcal{O}(n^{2-\delta} \text{poly}(1/\varepsilon))$ time approximation scheme that provides a $(1 + \varepsilon)$ -approximation for DIAMETER for any $\varepsilon > 0$ in intersection graphs of axis-parallel unit segments in \mathbb{R}^2 , or unit balls or axis-parallel unit cubes in \mathbb{R}^3 .*

Theorem 1 shows that sub-quadratic algorithms in many intersection graphs classes are unlikely to exist; one must wonder if such algorithms are possible at all? A notable case missing from our lower bounds are the case of unit disks; indeed, it is possible that unit disk graphs enjoy sub-quadratic diameter computation. More generally, it is an interesting open question whether intersection graphs of so-called pseudodisks admit sub-quadratic diameter algorithms. (Pseudodisks are objects bounded by Jordan curves such that the boundaries of any pair of objects have at most two intersection points.) We make a step towards resolving this problem with the following theorem for intersection graphs of axis-parallel unit squares – since axis-parallel unit squares are pseudodisks.

► **Theorem 3.** *There is an $\mathcal{O}(n \log n)$ algorithm for DIAMETER-2 in unit square graphs.*

The algorithm is based on the insight that the problem can be simplified to the following: given *skylines* A, B and a list of axis-parallel squares S , check whether each pair $(a, b) \in A \times B$ is covered by some square $s \in S$. Since any axis-parallel square $s \in S$ covers *intervals* in A and B , this problem in turn reduces to checking whether the union of $|S|$ rectangles covers the $A \times B$ grid. Using near-linear skyline computation [35], and a line sweep for the grid covering problem, we obtain a surprisingly simple $\mathcal{O}(n \log n)$ time algorithm (in contrast to the quadratic-time hardness in higher dimensions).

Organization. After some preliminaries and the introduction of the complexity-theoretic hypotheses used in the paper, we present our algorithm for unit squares in Section 3. Section 4 showcases our lower bound techniques. The lower bounds for unit segments, congruent equilateral triangles as well as for axis-parallel unit segments have a structure similar to two other lower bounds in Section 4, and they can be found in the full version.

2 Preliminaries

Let $G = (V, E)$ be a graph, and u and v be vertices in G . The distance from u to v is denoted by $\text{dist}_G(u, v)$, and equals the number of edges on the shortest path from u to v in G . The diameter of G is denoted by $\text{diam}(G)$ and equals to $\max_{u, v \in V} \text{dist}_G(u, v)$. The open and closed neighborhood of a vertex v are $N(v) = \{u \in V \mid uv \in E\}$ and $N[v] = \{v\} \cup N(v)$, respectively. Let $A, B \subseteq V$ be sets of vertices. The diameter of A and B is denoted by $\text{diam}_G(A, B) = \max_{(a, b) \in A \times B} \text{dist}_G(a, b)$. Finally, let $[n]$ denote the set $\{1, \dots, n\}$.

2.1 Hardness assumptions


We use two hypotheses from fine-grained complexity theory for our lower bounds. For an overview of this field, we refer to the survey [44].

Orthogonal Vectors Hypothesis. Let OV denote the following problem: Given sets A, B of n vectors in $\{0, 1\}^d$, determine whether there exists an *orthogonal pair* $a \in A, b \in B$, i.e., for all $i \in [d]$ we have $(a)_i = 0$ or $(b)_i = 0$. Exhaustive search yields an $\mathcal{O}(n^2 d)$ algorithm, which can be improved for small dimension $d = c \log n$ to $\mathcal{O}(n^{2-1/O(\log(c))})$ [2, 18]. For larger dimensions $d = \omega(\log n)$, it is known [45] that no $\mathcal{O}(n^{2-\epsilon})$ -time algorithm can exist unless the Strong Exponential Time Hypothesis [32] fails. Thus, the Strong Exponential Time Hypothesis implies the following (so-called “moderate-dimensional”) OV Hypothesis.

► **Hypothesis 4 (Orthogonal Vectors Hypothesis).** *For no $\epsilon > 0$, there is an algorithm that solves OV in time $\mathcal{O}(\text{poly}(d)n^{2-\epsilon})$.*

By now, there is an extensive list of problems with tight lower bounds (including sub-quadratic equivalences) based on this assumption, see [44].

Hyperclique Hypothesis. For $k \geq 4$, let 3-UNIFORM k -HYPERCLIQUE denote the following problem: Given a 3-uniform hypergraph $G = (V, E)$, determine whether there exists a *hyperclique* of size k , i.e., a set $S \subseteq V$ such that for all $e \in \binom{S}{3}$, we have $e \in E$. By exhaustive search, we can solve this problem in time $\mathcal{O}(n^k)$ where $n = |V|$. Unlike the usual k -CLIQUE problem in graphs, for which a $\mathcal{O}(n^{\omega k/3 + \mathcal{O}(1)})$ algorithm exists [39], no techniques are known that would beat exhaustive search by a polynomial factor for the problem in hypergraphs. This has led to the hypothesis that exhaustive search is essentially best possible.



■ **Figure 1** The skylines (or fronts) of a point set P . In figure (a), the points on the blue curve are $\text{BRF}(P)$ and on the red curve are $\text{TLF}(P)$. In figure (b), the points on the green curve are $\text{TRF}(P)$ and on the orange curve are $\text{BLF}(P)$.

▶ **Hypothesis 5 (Hyperclique Hypothesis).** *For no $\epsilon > 0$ and $k \geq 4$, there is an algorithm that would solve 3-UNIFORM k -HYPERCLIQUE in time $\mathcal{O}(n^{k-\epsilon})$.*

See [38] for a detailed description of the plausibility of this hypothesis. Tight conditional lower bounds (including fine-grained equivalences) have been obtained, e.g., in [1, 11, 36, 4].

3 Solving the Diameter-2 problem on unit square graphs

In this section, we are going to present an algorithm with running time $\mathcal{O}(n \log n)$ for the DIAMETER-2 problem for unit square graphs. For each unit square $v \in V$, we consider the center of v , denoted \dot{v} , as the point representing v in the plane; for a square set $X \subset V$, we use \dot{X} to denote the set of corresponding centers. Let $\dot{G} = (\dot{V}, E)$ denote the graph on centers of squares in G . Hence, for all $\{u, v\} \in E(G)$, there is an edge between \dot{u} and \dot{v} . Note that we will often use \dot{G} and G interchangeably.

Notice that a graph has diameter at most two if and only if for every pair of vertices $u, v \in V : N[u] \cap N[v] \neq \emptyset$, i.e., there is a square w that both u and v have an intersection with or they intersect each other. Equivalently, the square of side length 2 centered at \dot{w} must cover both \dot{u} and \dot{v} . For a square w , let w^2 denote the side-length-2 square of center \dot{w} . Thus, in order to decide whether $\text{diam}(G) \leq 2$, it is sufficient to check whether for every $u, v \in V$ there exists $w \in V$ such that $\dot{u}, \dot{v} \in w^2$.

For a set of points P we define the *top-left front*, $\text{TLF}(P)$, and *bottom-right front*, $\text{BRF}(P)$ as follows (see Figure 1a).


$$\begin{aligned} \text{TLF}(P) &= \{p \in P \mid \forall q \in P: p_x \leq q_x \text{ or } p_y \geq q_y\} \\ \text{BRF}(P) &= \{p \in P \mid \forall q \in P: p_x \geq q_x \text{ or } p_y \leq q_y\} \end{aligned}$$

Similarly, we define the *top-right front*, $\text{TRF}(P)$, and *bottom-left front*, $\text{BLF}(P)$ as follows (see Figure 1b).

$$\begin{aligned} \text{TRF}(P) &= \{p \in P \mid \forall q \in P: p_x \geq q_x \text{ or } p_y \geq q_y\} \\ \text{BLF}(P) &= \{p \in P \mid \forall q \in P: p_x \leq q_x \text{ or } p_y \leq q_y\} \end{aligned}$$

▶ **Lemma 6.** *The graph G has diameter at most 2 if and only if*

$$\max \left(\text{diam}_{\dot{G}}(\text{BLF}(\dot{V}), \text{TRF}(\dot{V})), \text{diam}_{\dot{G}}(\text{TLF}(\dot{V}), \text{BRF}(\dot{V})) \right) \leq 2.$$



■ **Figure 2** Any square covering a' and b' also covers a and b .

Proof. If G has diameter at most two, then clearly any pair of subsets of \dot{V} have diameter at most 2 in \dot{G} . For the other direction, consider any pair $a, b \in \dot{V}$, and assume that $a_x \leq b_x$ and $a_y \leq b_y$. We prove that $\text{dist}_{\dot{G}}(a, b) \leq 2$.

Select $a' \in \text{BLF}(\dot{V})$ such that $a'_x \leq a_x$ and $a'_y \leq a_y$, see Figure 2. Similarly, select $b' \in \text{TRF}(\dot{V})$ such that $b_x \leq b'_x$ and $b_y \leq b'_y$. Then we can observe that the minimum bounding box of $\{a', b'\}$ covers the minimum bounding box of $\{a, b\}$. Since $\text{dist}_{\dot{G}}(a', b') \leq \text{diam}_{\dot{G}}(\text{BLF}(\dot{V}), \text{TRF}(\dot{V})) \leq 2$, there exists a square $w \in V$ such that w^2 covers $\{a', b'\}$. Consequently, w^2 also covers $\{a, b\}$, and thus $\text{dist}_{\dot{G}}(a, b) \leq 2$.

Finally, the case $a_x > b_x$ and $a_y > b_y$ is symmetric, and the cases $a_x > b_x, a_y \leq b_y$ and $a_x \leq b_x, a_y > b_y$ are analogous with TLF and BRF instead of TRF and BLF. ◀

Using Lemma 6, we are able to prove Theorem 3.

Proof of Theorem 3. We start our algorithm by computing $\text{TLF}(\dot{V}), \text{TRF}(\dot{V}), \text{BLF}(\dot{V})$, and $\text{BRF}(\dot{V})$ in $\mathcal{O}(n \log n)$ time [35]. Let $\dot{P} = \text{BLF}(\dot{V})$ and $\dot{Q} = \text{TRF}(\dot{V})$. By Lemma 6, it is sufficient to show that in $\mathcal{O}(n \log n)$ time we can decide whether $\text{diam}_{\dot{G}}(\dot{P}, \dot{Q}) \leq 2$; using the same algorithm for $\text{BRF}(\dot{V})$ and $\text{TLF}(\dot{V})$ will then get the desired running time.

In order to check whether $N[\dot{p}] \cap N[\dot{q}] \neq \emptyset$ for all $(\dot{p}, \dot{q}) \in \dot{P} \times \dot{Q}$, we do the following: Consider $\dot{P} = \{\dot{p}_1, \dots, \dot{p}_{|\dot{P}|}\}$ and $\dot{Q} = \{\dot{q}_1, \dots, \dot{q}_{|\dot{Q}|}\}$ in x -order. Also, let $\text{GRID} = [|\dot{P}|] \times [|\dot{Q}|]$ be a grid where \dot{p}_i corresponds to the i -th row and \dot{q}_j corresponds to the j -th column.

For each square $v \in V$, recall that v^2 denotes the square with the same center but twice the side length. For each $v \in V$, define $I_v \subseteq \{1, 2, \dots, |\dot{P}|\}$ such that $i \in I_v$ iff v^2 contains \dot{p}_i . Similarly, $J_v \subseteq \{1, 2, \dots, |\dot{Q}|\}$ such that $j \in J_v$ iff v^2 contains \dot{q}_j . Since v^2 is an axis-parallel square, it covers intervals from both \dot{P} and \dot{Q} , thus I_v and J_v consist of consecutive integers. Therefore, we can think of the sets $I_v \times J_v$ as rectangles in GRID .

▷ **Claim 7.** We have $N[\dot{p}] \cap N[\dot{q}] \neq \emptyset$ for all $(\dot{p}, \dot{q}) \in \dot{P} \times \dot{Q}$ if and only if the union of $I_v \times J_v$ over all squares $v \in V$ covers GRID .

Proof. If the union of all rectangles covers the whole grid, then for any pair $(\dot{p}_i, \dot{q}_j) \in \dot{P} \times \dot{Q}$ of centers, there is a rectangle $I_v \times J_v$ that covers (i, j) . Therefore, v^2 covers both \dot{p}_i and \dot{q}_j . Thus, v is a shared neighbor of \dot{p} and \dot{q} .

If $N[\dot{p}] \cap N[\dot{q}] \neq \emptyset$ for all $(\dot{p}, \dot{q}) \in \dot{P} \times \dot{Q}$, then for each pair (\dot{p}_i, \dot{q}_j) there is at least one square v_{ij} such that v_{ij}^2 contains both \dot{p}_i and \dot{q}_j . Hence, $(i, j) \in I_{v_{ij}} \times J_{v_{ij}}$ for each $(i, j) \in \text{GRID}$. As a result, the union of $I_v \times J_v$ over all squares v^2 covers GRID . ◀

Note that the problem in Claim 7 corresponds to determining whether a union of rectangles covers the full grid. This problem can be solved in $\mathcal{O}(n \log n)$ time with a plane sweep [6, 43]. The time needed to construct the rectangles in GRID is $\mathcal{O}(n \log n)$ as there are $\mathcal{O}(n)$ rectangles. This concludes the proof of Theorem 3. ◀

4 Lower bounds based on the Orthogonal Vectors Hypothesis

In this section, we prove lower bounds for finding the diameter in various intersection graphs.


For a comparison to similar results on sparse graphs, let us briefly describe the result ruling out a $(3/2 - \epsilon)$ -approximation in time $\mathcal{O}(n^{2-\delta})$, for any $\epsilon, \delta > 0$, due to Roditty and Vassilevska-Williams [40]. While it is originally stated as a reduction from k -Dominating Set, we adapt it to give a reduction from OV: Given sets $A, B \subseteq \{0, 1\}^d$, introduce *vector nodes* for each $a \in A$ and $b \in B$ as well as *coordinate nodes* for $k \in [d]$. Without loss of generality (see Section 4.1), one may assume that all vectors $a \in A$ have $(a)_{d-1} = 1$ and all vectors $b \in B$ have $(b)_d = 1$. We connect each vector node $v \in A \cup B$ to the coordinate node $k \in [d]$ iff $(v)_k = 1$, and make all coordinate nodes a clique by adding all possible edges between coordinate nodes. The important observation is that (1) a pair $a \in A, b \in B$ has distance at most 2 iff there is a $k \in [d]$ such that $(a)_k = (b)_k = 1$, i.e., a, b do *not* form an orthogonal pair, and (2) all other types of node pairs have distance at most 2. Thus, A, B contains an orthogonal pair iff the diameter of the constructed graph is at least 3. Since the reduction produces a sparse graph with $\mathcal{O}(n + d)$ nodes and $\mathcal{O}(nd)$ edges in time $\mathcal{O}(nd)$, any $\mathcal{O}(m^{2-\delta})$ -time algorithm distinguishing between diameter 2 and 3 would give a $\mathcal{O}(n^{2-\delta} \text{poly}(d))$ -time OV algorithm, refuting the OV Hypothesis.

Generally speaking, implementing this reduction using low-dimensional geometric graphs is problematic: we must be able to implement an arbitrary bipartite graph on a vertex set $L \times R$ where $|L| = n$ and $|R| = d$. Instead, in this section we implement two different types of reductions via geometric graphs; the main ideas are as follows:

Diameter-3 graphs (Section 4.1 and full version). Instead of *coordinate nodes*, we introduce *1-entry nodes* $(v)_k$ for all $v \in A \cup B, k \in [d]$ with $(v)_k = 1$. This increases the number of nodes only to $\mathcal{O}(nd)$, while allowing us to geometrically implement edges of the form $\{v, (v)_k\}$ for all $v \in A \cup B, k \in [d]$ with $(v)_k = 1$ and $\{(v)_k, (v')_k\}$ for all $v, v' \in A \cup B, k \in [d]$ with $(v)_k = (v')_k = 1$. Now, a witness of non-orthogonality of a, b is a 3-path $a - (a)_k - (b)_k - b$. By showing that all other distances are bounded by 3, we obtain hardness for the DIAMETER-3 problem. See Section 4.1 and the full version for details, including the use of an additional node to make all 1-entry nodes sufficiently close in distance.

(Non-sparse) Diameter- $\Theta(d)$ graphs (Section 4.2 and full version). Instead of *coordinate nodes* or *1-entry nodes*, we introduce *vector-coordinate nodes* $(v)_k$ for all $v \in A \cup B, k \in [d]$, irrespective of whether $(v)_k = 1$. As opposed to previously, we do not create a constant diameter instance: The idea is to create an instance where the most distant pairs are of the form $(a)_1, (b)_d$ for $a \in A, b \in B$, and a non-orthogonality witness is a path of the form $(a)_1 \rightsquigarrow \dots \rightsquigarrow (a)_k \rightsquigarrow (b)_k \rightsquigarrow \dots \rightsquigarrow (b)_d$ with $(a)_k = (b)_k = 1$. This construction requires us to implement perfect matchings between vector-coordinate gadgets $(a)_k$ for $a \in A$ and $(a')_{k+1}$ for $a' \in A$ if $a = a'$, as well as a gadget for implementing short connections for $(a)_k \rightsquigarrow (b)_k$ that check whether $(a)_k = (b)_k = 1$. Interestingly, this type of reduction generally produces dense graphs with $\Omega(n^2)$ edges, so this approach crucially exploits the expressive power of geometric graphs to give a subquadratic reduction. See Section 4.2 and the full version for details, including a description of auxiliary nodes not mentioned here.

Finally, we remark that the reduction for unit hypercubes given in Section 5 has the most similar structure to the reduction by Roditty and Vassilevska-Williams [40], despite starting from a different hypothesis, and has similarities to [4, Theorem 14]. We crucially exploit properties of the hyperclique problem to implement it using hypercube graphs.



■ **Figure 3** Reducing orthogonal vectors to DIAMETER-3 in intersection graphs of line segments.

4.1 The Diameter-3 problem for line segment intersection graphs

In this section, we are going to present a lower bound on the running time of the algorithm for the DIAMETER-3 problem for line segment intersection graphs, such that vertices are line segments with any length, and there is an edge between a pair of line segments if they intersect. This serves as a warm-up for the slightly more complicated reductions below.

► **Theorem 8.** *For all $\epsilon > 0$, there is no $\mathcal{O}(n^{2-\epsilon})$ time algorithm for the DIAMETER-3 problem for line segment intersection graphs, unless the OV Hypothesis fails.*

Let $A = \{a_1, a_2, \dots, a_n\}$ and $B = \{b_1, b_2, \dots, b_n\}$ be two sets of n vectors in $\{0, 1\}^d$. We construct a set of segments such that the diameter of the corresponding intersection graph is at most 3 if and only if there is no orthogonal pair $(a, b) \in A \times B$.

Without loss of generality, we assume that for each $a_i \in A$ and $b_j \in B$, $((a_i)_{d-1}, (a_i)_d) = (1, 0)$ and $((b_j)_{d-1}, (b_j)_d) = (0, 1)$, by adding two coordinates to the ends of the vectors. Note that adding these coordinates does not change whether vectors a, b are orthogonal or not.


For each vector $a_i \in A$, let \bar{a}_i denote a zero-length line segment from $(i, 1)$ to $(i, 1)$. Analogously, for each vector $b_j \in B$, let \bar{b}_j denote a line segment from $(j, -1)$ to $(j, -1)$. Furthermore, let ℓ be a line segment from $(1, 0)$ to $(d, 0)$, and let $\{w_1, w_2, \dots, w_d\}$ be d different points on ℓ such that for all $k \in [d]$, w_k is located at $(k, 0)$. Moreover, for each $a_i \in A$, if $(a_i)_k = 1$, we define a line segment $e_{i,k}$ from \bar{a}_i to w_k (i.e., from $(i, 1)$ to $(k, 0)$). Analogously, for each $b_j \in B$, if $(b_j)_{k'} = 1$, we define a line segment $e'_{j,k'}$ from \bar{b}_j to $w_{k'}$ (i.e., from $(j, -1)$ to $(k', 0)$). Let \bar{V} be the set of constructed line segments, and let G be their intersection graph (see Figure 3).

► **Lemma 9.** *The sets A and B contain an orthogonal pair if and only if $\text{diam}(G) \geq 4$.*

Let \bar{A} be the set of line segments corresponding to vectors in A . Analogously, let \bar{B} be the set of line segments corresponding to vectors in B . To prove the lemma, we show that each pair of vertices is within distance at most 3, unless it is in $\bar{A} \times \bar{B}$ (see Claim 10 below and see the full version for its proof). The pairs in $\bar{A} \times \bar{B}$ have distance 4 or 3 depending on whether their corresponding vectors in $A \times B$ are orthogonal or not.

▷ **Claim 10.** $\text{dist}(\bar{u}, \bar{v}) \leq 3$ for all $(\bar{u}, \bar{v}) \in (\bar{V} \times \bar{V}) \setminus (\bar{A} \times \bar{B} \cup \bar{B} \times \bar{A})$.

Proof of Lemma 9. If a_i and b_j are not orthogonal, then there is at least one $k \in [d]$ such that $(a_i)_k = (b_j)_k = 1$. Hence, the path $\bar{a}_i - e_{i,k} - e'_{j,k} - \bar{b}_j$ exists, and it has length 3. If a_i and b_j are orthogonal, then there is no index k such that $(a_i)_k = (b_j)_k = 1$. Consequently, there is no path of length 3 from \bar{a}_i to \bar{b}_j , and $\text{dist}(\bar{a}_i, \bar{b}_j) \geq 4$. Together with Claim 10 this proves the lemma. ◀



■ **Figure 4** Schematic picture of the graph $G(A)$.

Proof of Theorem 8. The above reduction creates a set of $N = \mathcal{O}(nd)$ segments in $\mathcal{O}(nd)$ time. If there is an algorithm solving DIAMETER-3 in $\mathcal{O}(N^{2-\delta})$ time in segment intersection graphs, then combining this algorithm with the reduction would solve OV in time $\mathcal{O}(nd) + \mathcal{O}((nd)^{2-\delta}) = \mathcal{O}(n^{2-\delta}\text{poly}(d))$, refuting the OV Hypothesis. ◀

4.2 The diameter problem for unit ball graphs

► **Theorem 11.** *For all $\epsilon > 0$, there is no $\mathcal{O}(n^{2-\epsilon})$ time algorithm for solving DIAMETER in unit ball graphs in \mathbb{R}^3 under the Orthogonal Vectors Hypothesis.*

Let $A = \{a_1, a_2, \dots, a_n\}$ be a given set of vectors from $\{0, 1\}^d$. First, we construct graph $G(A)$ and show that $G(A)$ has diameter $\geq 2d + 5$ if and only if there is an orthogonal pair of vectors in A . Next, we show how $G(A)$ can be realized as an intersection graph of unit balls in \mathbb{R}^3 . Without loss of generality, assume that the all-one vector is an element of A (if it is not in A , then adding the all-one vector does not change whether there is an orthogonal pair.)

We construct a graph $G(A)$ as follows. Let C_1^T, \dots, C_{2d}^T and C_1^B, \dots, C_{2d}^B be cliques, such that for all $k \in [2d]$, $C_k^T = \{v_{k,1}^T, \dots, v_{k,n}^T\}$, $C_k^B = \{v_{k,1}^B, \dots, v_{k,n}^B\}$, and $v_{k,i}^T$ and $v_{k,i}^B$ correspond to a_i for all $i \in [n]$, see Figure 4. We add a perfect matching between each pair C_k^T and C_{k+1}^T for all $k \in [2d-1]$ such that there is an edge incident to $v_{k,i}^T$ and $v_{k+1,i}^T$ for all $i \in [n]$. Analogously, there is a perfect matching between each pair C_k^B and C_{k+1}^B .

Let M_1^T, \dots, M_d^T be cliques such that if $(a_i)_k = 1$, then there is a vertex $m_{k,i}^T$ in M_k^T that is adjacent to $v_{k,i}^T$. Similarly, let M_1^B, \dots, M_d^B be cliques such that if $(a_i)_k = 1$, then there is a vertex $m_{k,i}^B$ in M_k^B that is adjacent to $v_{k,i}^B$. Notice that because of the addition of the all ones vector, the cliques M_k^T and M_k^B are all non-empty.

Finally, let $Q = \{q_1, q_2, \dots, q_d\}$ be a set of vertices such that q_k has edges to all vertices in M_k^T and M_k^B for all $k \in [d]$.

► **Lemma 12.** *The graph $G(A)$ has diameter at most $2d + 4$ iff A has no orthogonal pair.*

Proof. Assume that there is an orthogonal pair $(a_i, a_j) \in A$ such that $i \neq j$. Hence, $\sum_{k=1}^n (a_i)_k (a_j)_k = 0$, which means that there is no $k \in [n]$ such that $(a_i)_k = (a_j)_k = 1$. Consequently, for all $k \in [d]$, the distance from $v_{2k,i}^T \in C_{2k}^T$ to $v_{2k-1,j}^B \in C_{2k-1}^B$ is at least 5. Therefore, $2d + 4 < \text{dist}(v_{1,i}^T, v_{2d,j}^B) \leq \text{diam}(G(A))$.

Now suppose that A has no orthogonal pair. We want to prove that $\text{diam}(G(A)) \leq 2d + 4$. Since A has no orthogonal pair, for each pair (a_i, a_j) there is at least one $k \in [n]$ such that $(a_i)_k = (a_j)_k = 1$. Therefore, there are cliques M_k^T and M_k^B that have the vertices $m_{k,i}^T$ and $m_{k,j}^B$, respectively. Since all vertices in M_k^T and M_k^B have an edge to q_k , we can reach q_k from $v_{1,i}^T$ by a path of length $2k - 1 + 2$. Simultaneously, we can reach q_k from $v_{2d,j}^B$ by a path of length $2d - 2k + 1 + 2$. In total, this gives a path of length $2d + 4$ between $v_{1,i}^T$ and $v_{2d,j}^B$. Furthermore, it is easy check that the distance of any pair of vertices where at least one vertex is outside $C_1^T \cup C_{2d}^B$ is at most $2d + 4$. As a result, $\text{diam}(G(A)) \leq 2d + 4$. ◀

► **Lemma 13.** $G(A)$ can be realized as an intersection graph of unit balls in \mathbb{R}^3 .

Proof. For converting $G(A)$ into an intersection graph of unit balls, we should consider each vertex in $G(A)$ as the center of a unit diameter ball, and for those vertices that are adjacent, their corresponding unit balls should intersect. To this end, we choose the following coordinates for the centers of the unit balls in \mathbb{R}^3 :

- For all $k \in [2d]$ and $i \in [n]$, the center point of $v_{k,i}^T \in C_k^T$ is $(k, \frac{i}{n}, 0)$.
- For all $k \in [d]$ and $i \in [n]$, if $m_{k,i}^T \in M_k^T$ exists, then its center point is $(2k, \frac{i}{n}, -1)$.
- For all $k \in [d]$, the center point of $q_k \in Q$ is $(2k, \frac{1}{2}, -1.6)$.
- For all $k \in [d]$ and $i \in [n]$, if $m_{i,j}^B \in M_i^B$ exists, its center point is $(2k, \frac{i}{n}, -2.2)$.
- For all $k \in [2d]$ and $i \in [n]$, the center point of $v_{k,i}^B \in C_k^B$ is $(k, \frac{i}{n}, -3.2)$.

The distance between center points that correspond to adjacent vertices should be at most 1. For each two vertices in the same clique in C_k^T , C_k^B , M_k^T , and M_k^B their center points differ only in the y -coordinate. Since this difference is at most $1 - 1/n < 1$, they form a clique. For each two adjacent vertices in two different cliques, their center points differ either only in the x -, or only in the z -coordinate, by exactly 1, hence, they intersect. For a vertex in Q and M_k^T , if $m_{k,i}^T$ exists, the distance between $m_{k,i}^T$ and q_k is

$$\sqrt{(2k - 2k)^2 + (\frac{i}{n} - \frac{1}{2})^2 + (-1 - (-1.6))^2} = \sqrt{(\frac{i}{n} - \frac{1}{2})^2 + (0.6)^2} \leq \sqrt{(\frac{1}{2})^2 + (0.6)^2} < 1$$

The same argument holds for adjacent vertices in Q and M^B . One can easily check that the non-adjacent vertices have distance strictly greater than 1. ◀

Proof of Theorem 11. The construction creates a set of $N = \mathcal{O}(nd)$ balls in $\mathcal{O}(nd)$ time. If there is an algorithm to solve DIAMETER in $\mathcal{O}(N^{2-\delta})$ time in ball graphs, then we could combine this construction with the algorithm, and solve the ORTHOGONAL VECTORS problem in $\mathcal{O}(nd) + \mathcal{O}((nd)^{2-\delta}) = \mathcal{O}(n^{2-\delta} \text{poly}(d))$ time. This contradicts the Orthogonal Vectors Hypothesis, and concludes the theorem. ◀

A simple transformation of this construction shows that we can realize $G(A)$ also as an intersection graph of axis-parallel unit cubes.

► **Corollary 14.** For all $\epsilon > 0$, there is no $\mathcal{O}(n^{2-\epsilon})$ time algorithm for solving DIAMETER in intersection graphs of axis-parallel unit cubes in \mathbb{R}^3 under the Orthogonal Vectors Hypothesis.

Proof. Let P denote the set of centers constructed for unit balls. We rotate P by $\pi/4$ around the y axis, and scale P by a factor of $\sqrt{2}$. Let P' be the resulting set of points. Note that in P , all inter-clique edges were realized by a horizontal or vertical point pair of distance

exactly 1. In P' , the corresponding pairs are diagonal segments in some plane perpendicular to the y -axis, therefore the unit side-length cubes centered at the corresponding pair of points will have a touching edge. It is routine to check that the unit side-length cubes centered at P' realize the intersection graph $G(A)$. ◀

Proof of Corollary 2. The lower bounds regarding constant-approximations in sub-quadratic time are immediate consequences of our lower bounds for DIAMETER-2 and DIAMETER-3. Notice that our proofs for unit balls and axis-parallel unit cubes in \mathbb{R}^3 , as well as axis-parallel unit segments in \mathbb{R}^2 use a construction where the resulting intersection graph has diameter $d^* = \Theta(d)$. Under OV, there exists no $(1 + \varepsilon)$ -approximation for these problems that would run in $n^{2-\delta} \text{poly}(1/\varepsilon)$ time, as setting $\varepsilon = 1/d^* = \Theta(1/d)$ would enable us to decide OV in $n^{2-\delta} \text{poly}(d)$ time. ◀

5 The Diameter-2 problem for hypercube graphs: a hyperclique lower bound

► **Theorem 15.** *For all $\epsilon > 0$ there is no $O(n^{2-\epsilon})$ algorithm for DIAMETER-2 in unit hypercube graphs in \mathbb{R}^{12} , unless the Hyperclique Hypothesis fails.*

Proof. Observe that under the Hyperclique Hypothesis, it requires time $n^{6-o(1)}$ to find a hyperclique of size 6 in a given 3-uniform hypergraph $G = (V, E)$. In fact, using a standard color-coding argument, we can assume without loss of generality that G is 6-partite: We have $V = V_1 \cup \dots \cup V_6$ for disjoint sets V_i of size n each, and any 6-hyperclique must choose exactly one vertex from each V_i . By slight abuse of notation, we view each V_i as a disjoint copy of $[n]$, i.e., node $j \in [n]$ in V_i is different from node j in $V_{i'}$ with $i' \neq i$. Furthermore, by complementing the edge set, we arrive at the equivalent task of determining whether G has an *independent set* of size 6, i.e., whether there are $(v_1, \dots, v_6) \in V_1 \times \dots \times V_6$ such that $\{v_i, v_j, v_k\} \notin E$ for all distinct $i, j, k \in [6]$. Finally, for technical reasons, we assume without loss of generality that for each $v_i \in V_i$ and distinct $j, k \in [6] \setminus \{i\}$, there are $v_j \in V_j, v_k \in V_k$ with $\{v_i, v_j, v_k\} \in E$: To this end, simply add, for every $\ell \in [6]$, a dummy vertex v'_ℓ to V_ℓ , and add, for every i, j, k and $v_j \in V_j, v_k \in V_k$, the edge $\{v'_i, v_j, v_k\}$ to E , i.e., each dummy vertex is connected to all other pairs of vertices (including other dummy vertices). Observe that this yields an equivalent instance, since no dummy vertex can be contained in an independent set.

The reduction is given by constructing a set of $O(n^3)$ unit hypercubes in \mathbb{R}^{12} , which we specify by their centers. These (hyper)cubes are of three types: *left-half cubes* representing a choice of the vertices $(x_1, x_2, x_3) \in V_1 \times V_2 \times V_3$, *right-half cubes* representing a choice of the vertices $(y_1, y_2, y_3) \in V_4 \times V_5 \times V_6$ and *edge cubes* representing an edge $\{v_i, v_j, v_k\} \in E$. In particular, the choice of a vertex in V_i will be encoded in the dimensions $2i - 1$ and $2i$.

Specifically, for each $(x_1, x_2, x_3) \in V_1 \times V_2 \times V_3$ such that $\{x_1, x_2, x_3\} \notin E$, we define the center of the left-half cube X_{x_1, x_2, x_3} as

$$\left(\frac{x_1}{n+1}, 1 - \frac{x_1}{n+1}, \frac{x_2}{n+1}, 1 - \frac{x_2}{n+1}, \frac{x_3}{n+1}, 1 - \frac{x_3}{n+1}, 2, \dots, 2 \right).$$

Similarly, for each $(y_1, y_2, y_3) \in V_4 \times V_5 \times V_6$ such that $\{y_1, y_2, y_3\} \notin E$, we define the center of the right-half cube Y_{y_1, y_2, y_3} as

$$\left(2, \dots, 2, \frac{y_1}{n+1}, 1 - \frac{y_1}{n+1}, \frac{y_2}{n+1}, 1 - \frac{y_2}{n+1}, \frac{y_3}{n+1}, 1 - \frac{y_3}{n+1} \right).$$

Finally, for each edge $e = \{v_i, v_j, v_k\} \in E$ not already in $V_1 \times V_2 \times V_3 \cup V_4 \times V_5 \times V_6$, we define a corresponding edge cube E_{v_i, v_j, v_k} with the following center point: We set the $2i - 1$ -th coordinate to $1 + \frac{v_i}{n+1}$, the $2i$ -th coordinate to $2 - \frac{v_i}{n+1}$, and similarly we set the coordinates $2j - 1, 2j, 2k - 1, 2k$ to $1 + \frac{v_j}{n+1}, 2 - \frac{v_j}{n+1}, 1 + \frac{v_k}{n+1}, 2 - \frac{v_k}{n+1}$, respectively, and we set all remaining coordinates to 1. For example, if $i = 1, j = 2, k = 4$, the center point of E_{v_1, v_2, v_4} is

$$\left(1 + \frac{v_1}{n+1}, 2 - \frac{v_1}{n+1}, 1 + \frac{v_2}{n+1}, 2 - \frac{v_2}{n+1}, 1, 1, 1 + \frac{v_4}{n+1}, 2 - \frac{v_4}{n+1}, 1, 1, 1, 1\right).$$

Let S denote the set of all unit cubes $X_{x_1, x_2, x_3}, Y_{y_1, y_2, y_3}, E_{v_i, v_j, v_k}$ constructed above and let G_S denote the geometric intersection graph of the unit cubes. We prove that $\text{diam}(G_S) \leq 2$ if and only if there is no independent set $(v_1, \dots, v_6) \in V_1 \times \dots \times V_6$ in the 3-uniform hypergraph $G = (V_1 \cup \dots \cup V_6, E)$:

1. **Intra-set distances:** We have that the left- and right-half cubes as well as the edge cubes form cliques, i.e., $\text{dist}_{G_S}(X_{x_1, x_2, x_3}, X_{x'_1, x'_2, x'_3}) \leq 1$, $\text{dist}_{G_S}(Y_{y_1, y_2, y_3}, Y_{y'_1, y'_2, y'_3}) \leq 1$ and $\text{dist}_{G_S}(E_{v_1, v_2, v_3}, E_{v'_1, v'_2, v'_3}) \leq 1$: Observe that the center of each X_{x_1, x_2, x_3} is contained in $[0, 1]^6 \times \{2\}^6$ and thus in a hypercube of side length at most 1. Thus, all cubes X_{x_1, x_2, x_3} intersect each other, proving $\text{dist}_{G_S}(X_{x_1, x_2, x_3}, X_{x'_1, x'_2, x'_3}) \leq 1$. The remaining claims follow analogously by observing that the centers of Y_{y_1, y_2, y_3} and E_{v_1, v_2, v_3} are contained in $\{2\}^6 \times [0, 1]^6$ and $[1, 2]^{12}$, respectively, and thus also in hypercubes of side length at most 1.
2. **Equality checks:** Let $x_1 \in V_1, x_2 \in V_2, x_3 \in V_3$ and $v_i \in V_i, v_j \in V_j, v_k \in V_k$. Then $\text{dist}_{G_S}(X_{x_1, x_2, x_3}, E_{v_i, v_j, v_k}) = 1$ iff $v_\ell = x_\ell$ whenever $\ell \in \{1, 2, 3\} \cap \{i, j, k\}$: Consider $\ell \in \{1, 2, 3\} \cap \{i, j, k\}$. Then the dimensions $(2\ell - 1, 2\ell)$ of X_{x_1, x_2, x_3} and E_{v_i, v_j, v_k} are equal to $(\frac{x_\ell}{n+1}, 1 - \frac{x_\ell}{n+1})$ and $(1 + \frac{v_\ell}{n+1}, 2 - \frac{v_\ell}{n+1})$, respectively. Note that $(1 + \frac{v_\ell}{n+1}) - \frac{x_\ell}{n+1} \leq 1$ and $(2 - \frac{v_\ell}{n+1}) - (1 - \frac{x_\ell}{n+1}) \leq 1$ hold simultaneously iff $x_\ell = v_\ell$. All other dimensions $\ell' \notin \{1, 2, 3\} \cap \{i, j, k\}$ are trivially within distance 1, since dimensions $(2\ell' - 1, 2\ell')$ of X_{x_1, x_2, x_3} and E_{v_i, v_j, v_k} are $(2, 2)$ and in $[1, 2]^2$, respectively (if $\ell' \notin \{1, 2, 3\}$), or in $[0, 2]^2$ and $(1, 1)$, respectively (if $\ell' \notin \{i, j, k\}$). The analogous claim holds for distances between Y_{y_1, y_2, y_3} and E_{v_i, v_j, v_k} .
3. **Edge distances:** We have that $\text{dist}_{G_S}(X_{x_1, x_2, x_3}, E_{v_i, v_j, v_k}) \leq 2$: By our technical assumption, we have that there is an edge $\{x_1, v'_4, v'_5\} \in E$ for some vertices $v'_4 \in V_4$ and $v'_5 \in V_5$. Thus, by the previous properties, we obtain that

$$\text{dist}_{G_S}(X_{x_1, x_2, x_3}, E_{v_i, v_j, v_k}) \leq \text{dist}_{G_S}(X_{x_1, x_2, x_3}, E_{x_1, v'_4, v'_5}) + \text{dist}_{G_S}(E_{x_1, v'_4, v'_5}, E_{v_i, v_j, v_k}) \leq 2.$$

4. **Distances of left- and right-half cubes:** Let $x_1 \in V_1, x_2 \in V_2, x_3 \in V_3$ and $y_1 \in V_4, y_2 \in V_5, y_3 \in V_6$ such that $\{x_1, x_2, x_3\}, \{y_1, y_2, y_3\} \notin E$ (thus, the left-half/right-half cubes for $\{x_1, x_2, x_3\}, \{y_1, y_2, y_3\}$ exist). Then we have that $\text{dist}_{G_S}(X_{x_1, x_2, x_3}, Y_{y_1, y_2, y_3}) > 2$ iff $(x_1, x_2, x_3, y_1, y_2, y_3)$ is an independent set in G : If the tuple $(x_1, x_2, x_3, y_1, y_2, y_3)$ is not an independent set, then there must be an edge $\{x_i, y_j, y_k\}$ or $\{x_i, x_j, y_k\}$ with $i, j, k \in [3]$, since $\{x_1, x_2, x_3\}$ and $\{y_1, y_2, y_3\}$ are non-edges. Consider the first case, the other is symmetric. Then by the equality-check property, that $\text{dist}_{G_S}(X_{x_1, x_2, x_3}, E_{x_i, x_j, y_k}) = 1$ and $\text{dist}_{G_S}(E_{x_i, x_j, y_k}, Y_{y_1, y_2, y_3}) = 1$, which yields $\text{dist}_{G_S}(X_{x_1, x_2, x_3}, Y_{y_1, y_2, y_3}) \leq 2$. It remains to consider the case that the tuple $(x_1, x_2, x_3, y_1, y_2, y_3)$ is an independent set. Since there cannot be any edge between a left-half cube $X_{x'_1, x'_2, x'_3}$ – which is contained in $(0, 1)^6 \times \{2\}^6$ – and a right-half cube $Y_{y'_1, y'_2, y'_3}$ – which is contained in $\{2\}^6 \times (0, 1)^6$ –, the only way to reach Y_{y_1, y_2, y_3} from X_{x_1, x_2, x_3} via a path of length 2 would have to use some edge cube E_{v_i, v_j, v_k} . However, by the equality-check property, a path $X_{x_1, x_2, x_3} - E_{v_i, v_j, v_k} - Y_{y_1, y_2, y_3}$ would

imply that the vertices chosen by $(x_1, x_2, x_3, y_1, y_2, y_3)$ would agree with v_i, v_j, v_k in the sets V_i, V_j, V_k . Thus, we would have found an edge $\{v_i, v_j, v_k\}$ among $(x_1, x_2, x_3, y_1, y_2, y_3)$, contradicting the assumption that it is an independent set.

Finally, observe that given a 3-uniform hypergraph G , we can construct the corresponding cube set S , containing $O(n^3)$ nodes, in time $O(n^3)$. Thus, if we had an $O(N^{2-\epsilon})$ -time algorithm for determining whether an N -vertex unit cube graph G_S has a diameter of at most 2, we could detect existence of an independent set (or equivalently, hyperclique) of size 6 in G in time $O(n^{6-3\epsilon})$, which would refute the Hyperclique Hypothesis. ◀

References

- 1 Amir Abboud, Karl Bringmann, Holger Dell, and Jesper Nederlof. More consequences of falsifying SETH and the orthogonal vectors conjecture. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *Proc. 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2018)*, pages 253–266. ACM, 2018. doi:10.1145/3188745.3188938.
- 2 Amir Abboud, Ryan Williams, and Huacheng Yu. More applications of the polynomial method to algorithm design. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '15*, pages 218–230. SIAM, 2015.
- 3 Donald Aingworth, Chandra Chekuri, Piotr Indyk, and Rajeev Motwani. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM J. Comput.*, 28(4):1167–1181, 1999. doi:10.1137/S0097539796303421.
- 4 Haozhe An, Mohit Jayanti Gurumukhani, Russell Impagliazzo, Michael Jaber, Marvin Künnemann, and Maria Paula Parga Nina. The fine-grained complexity of multi-dimensional ordering properties. In Petr A. Golovach and Meirav Zehavi, editors, *Proc. 16th International Symposium on Parameterized and Exact Computation (IPEC 2021)*, volume 214 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 3:1–3:15, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.IPEC.2021.3.
- 5 Arturs Backurs, Liam Roditty, Gilad Segal, Virginia Vassilevska Williams, and Nicole Wein. Toward tight approximation bounds for graph diameter and eccentricities. *SIAM J. Comput.*, 50(4):1155–1199, 2021. doi:10.1137/18M1226737.
- 6 J. L. Bentley. Solutions to Klee’s rectangle problems. Unpublished manuscript, 1977.
- 7 Marthe Bonamy, Édouard Bonnet, Nicolas Bousquet, Pierre Charbit, Panos Giannopoulos, Eun Jung Kim, Pawel Rzazewski, Florian Sikora, and Stéphan Thomassé. EPTAS and subexponential algorithm for maximum clique on disk and unit ball graphs. *J. ACM*, 68(2):9:1–9:38, 2021. doi:10.1145/3433160.
- 8 Édouard Bonnet. 4 vs 7 sparse undirected unweighted diameter is SETH-hard at time $n^{4/3}$. In Nikhil Bansal, Emanuela Merelli, and James Worrell, editors, *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12–16, 2021, Glasgow, Scotland (Virtual Conference)*, volume 198 of *LIPIcs*, pages 34:1–34:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.ICALP.2021.34.
- 9 Édouard Bonnet. Inapproximability of diameter in super-linear time: Beyond the 5/3 ratio. In Markus Bläser and Benjamin Monmege, editors, *38th International Symposium on Theoretical Aspects of Computer Science, STACS 2021, March 16–19, 2021, Saarbrücken, Germany (Virtual Conference)*, volume 187 of *LIPIcs*, pages 17:1–17:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.STACS.2021.17.
- 10 Andreas Brandstädt and Feodor F. Dragan. On linear and circular structure of (claw, net)-free graphs. *Discret. Appl. Math.*, 129(2-3):285–303, 2003. doi:10.1016/S0166-218X(02)00571-1.
- 11 Karl Bringmann, Nick Fischer, and Marvin Künnemann. A fine-grained analogue of Schaefer’s theorem in P: dichotomy of $\exists^k\forall$ -quantified first-order graph properties. In Amir Shpilka, editor, *Proc. 34th Computational Complexity Conference (CCC 2019)*, volume 137 of *LIPIcs*, pages 31:1–31:27. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.CCC.2019.31.

- 12 Sergio Cabello. Subquadratic algorithms for the diameter and the sum of pairwise distances in planar graphs. *ACM Trans. Algorithms*, 15(2):21:1–21:38, 2019. doi:10.1145/3218821.
- 13 Massimo Cairo, Roberto Grossi, and Romeo Rizzi. New bounds for approximating extremal distances in undirected graphs. In Robert Krauthgamer, editor, *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 363–376. SIAM, 2016. doi:10.1137/1.9781611974331.ch27.
- 14 Timothy M. Chan. Polynomial-time approximation schemes for packing and piercing fat objects. *J. Algorithms*, 46(2):178–189, 2003. doi:10.1016/S0196-6774(02)00294-8.
- 15 Timothy M. Chan and Dimitrios Skrepetos. All-pairs shortest paths in unit-disk graphs in slightly subquadratic time. In Seok-Hee Hong, editor, *27th International Symposium on Algorithms and Computation, ISAAC 2016, December 12-14, 2016, Sydney, Australia*, volume 64 of *LIPICs*, pages 24:1–24:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPICs.ISAAC.2016.24.
- 16 Timothy M. Chan and Dimitrios Skrepetos. All-pairs shortest paths in geometric intersection graphs. *J. Comput. Geom.*, 10:27–41, 2019.
- 17 Timothy M. Chan and Dimitrios Skrepetos. Approximate shortest paths and distance oracles in weighted unit-disk graphs. *J. Comput. Geom.*, 10(2):3–20, 2019. doi:10.20382/jocg.v10i2a2.
- 18 Timothy M. Chan and Ryan Williams. Deterministic APSP, orthogonal vectors, and more: Quickly derandomizing Razborov-Smolensky. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '16*, pages 1246–1255. SIAM, 2016.
- 19 Shiri Chechik, Daniel H. Larkin, Liam Roditty, Grant Schoenebeck, Robert Endre Tarjan, and Virginia Vassilevska Williams. Better approximation algorithms for the graph diameter. In Chandra Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 1041–1052. SIAM, 2014. doi:10.1137/1.9781611973402.78.
- 20 Derek G. Corneil. Lexicographic breadth first search - A survey. In Juraj Hromkovic, Manfred Nagl, and Bernhard Westfechtel, editors, *Graph-Theoretic Concepts in Computer Science, 30th International Workshop, WG 2004, Bad Honnef, Germany, June 21-23, 2004, Revised Papers*, volume 3353 of *Lecture Notes in Computer Science*, pages 1–19. Springer, 2004. doi:10.1007/978-3-540-30559-0_1.
- 21 Derek G. Corneil, Feodor F. Dragan, Michel Habib, and Christophe Paul. Diameter determination on restricted graph families. *Discret. Appl. Math.*, 113(2-3):143–166, 2001. doi:10.1016/S0166-218X(00)00281-X.
- 22 Derek G. Corneil, Feodor F. Dragan, and Ekkehard Köhler. On the power of BFS to determine a graph's diameter. *Networks*, 42(4):209–222, 2003. doi:10.1002/net.10098.
- 23 Mina Dalirrooyfard, Ray Li, and Virginia Vassilevska Williams. Hardness of approximate diameter: Now for undirected graphs. *CoRR*, abs/2106.06026, 2021. arXiv:2106.06026.
- 24 Mina Dalirrooyfard and Nicole Wein. Tight conditional lower bounds for approximating diameter in directed graphs. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 1697–1710. ACM, 2021. doi:10.1145/3406325.3451130.
- 25 Mina Dalirrooyfard, Virginia Vassilevska Williams, Nikhil Vyas, and Nicole Wein. Tight approximation algorithms for bichromatic graph diameter and related problems. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece*, volume 132 of *LIPICs*, pages 47:1–47:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.ICALP.2019.47.
- 26 Feodor F. Dragan. Almost diameter of a house-hole-free graph in linear time via LexBFS. *Discret. Appl. Math.*, 95(1-3):223–239, 1999. doi:10.1016/S0166-218X(99)00077-3.
- 27 Feodor F. Dragan, Falk Nicolai, and Andreas Brandstädt. LexBFS-orderings and powers of graphs. In Fabrizio d'Amore, Paolo Giulio Franciosa, and Alberto Marchetti-Spaccamela, editors, *Graph-Theoretic Concepts in Computer Science*, pages 166–180. Berlin, Heidelberg, 1997. Springer Berlin Heidelberg.

- 28 Jie Gao and Li Zhang. Well-separated pair decomposition for the unit-disk graph metric and its applications. *SIAM J. Comput.*, 35(1):151–169, 2005. doi:10.1137/S0097539703436357.
- 29 Paweł Gawrychowski, Haim Kaplan, Shay Mozes, Micha Sharir, and Oren Weimann. Voronoi diagrams on planar graphs, and computing the diameter in deterministic $\tilde{O}(n^{5/3})$ time. *SIAM J. Comput.*, 50(2):509–554, 2021. doi:10.1137/18M1193402.
- 30 Dorit S. Hochbaum and Wolfgang Maass. Approximation schemes for covering and packing problems in image processing and VLSI. *J. ACM*, 32(1):130–136, 1985. doi:10.1145/2455.214106.
- 31 Harry B. Hunt III, Madhav V. Marathe, Venkatesh Radhakrishnan, S. S. Ravi, Daniel J. Rosenkrantz, and Richard Edwin Stearns. NC-approximation schemes for NP- and PSPACE-hard problems for geometric graphs. *J. Algorithms*, 26(2):238–274, 1998. doi:10.1006/jagm.1997.0903.
- 32 Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001.
- 33 Paul Koebe. *Kontaktprobleme der konformen Abbildung*. Hirzel, 1936.
- 34 Fabian Kuhn, Roger Wattenhofer, and Aaron Zollinger. Ad hoc networks beyond unit disk graphs. *Wirel. Networks*, 14(5):715–729, 2008. doi:10.1007/s11276-007-0045-6.
- 35 H. T. Kung, Fabrizio Luccio, and Franco P. Preparata. On finding the maxima of a set of vectors. *J. ACM*, 22(4):469–476, 1975. doi:10.1145/321906.321910.
- 36 Marvin Künnemann and Dániel Marx. Finding small satisfying assignments faster than brute force: A fine-grained perspective into boolean constraint satisfaction. In Shubhangi Saraf, editor, *Proc. 35th Computational Complexity Conference (CCC 2020)*, volume 169 of *LIPICs*, pages 27:1–27:28. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICs.CCC.2020.27.
- 37 Ray Li. Settling SETH vs. approximate sparse directed unweighted diameter (up to (NU)NSETH). In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21–25, 2021*, pages 1684–1696. ACM, 2021. doi:10.1145/3406325.3451045.
- 38 Andrea Lincoln, Virginia Vassilevska Williams, and R. Ryan Williams. Tight hardness for shortest cycles and paths in sparse graphs. In Artur Czumaj, editor, *Proc. 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2018)*, 2018. doi:10.1137/1.9781611975031.80.
- 39 Jaroslav Nešetřil and Svatopluk Poljak. On the complexity of the subgraph problem. *Commentationes Mathematicae Universitatis Carolinae*, 026(2):415–419, 1985.
- 40 Liam Roditty and Virginia Vassilevska Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1–4, 2013*, pages 515–524. ACM, 2013. doi:10.1145/2488608.2488673.
- 41 Raimund Seidel. On the all-pairs-shortest-path problem in unweighted undirected graphs. *J. Comput. Syst. Sci.*, 51(3):400–403, 1995. doi:10.1006/jcss.1995.1078.
- 42 Mikkel Thorup. Undirected single-source shortest paths with positive integer weights in linear time. *J. ACM*, 46(3):362–394, 1999. doi:10.1145/316542.316548.
- 43 Jan van Leeuwen and Derick Wood. The measure problem for rectangular ranges in d-space. *J. Algorithms*, 2(3):282–300, 1981. doi:10.1016/0196-6774(81)90027-4.
- 44 Virginia Vassilevska Williams. On some fine-grained questions in algorithms and complexity. In *Proceedings of the International Congress of Mathematicians, ICM '18*, pages 3447–3487, 2018.
- 45 Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theor. Comput. Sci.*, 348(2-3):357–365, 2005.