Lines: 500

Charting the Diameter Computation Landscape of Intersection Graphs in 3D and Above

Anonymous author(s)
Anonymous affiliation(s)

Abstract

Recent research on computing the diameter of geometric intersection graphs has made significant strides, primarily focusing on the 2D case [24, 17, 15] where truly subquadratic-time algorithms were given for simple objects such as unit-disks and (axis-aligned) squares. However, in three or higher dimensions, there is no known truly subquadratic-time algorithm for any *intersection graph of* non-trivial objects, even basic ones such as unit balls or (axis-aligned) unit cubes. This was partially explained by the pioneering work of Bringmann et al. [10] which gave several truly subquadratic lower bounds, notably for unit balls or unit cubes in 3D when the graph diameter Δ is at least $\Omega(\log n)$, hinting at a pessimistic outlook for the complexity of the diameter problem in higher dimensions. In this paper, we substantially extend the landscape of diameter computation for objects in three and higher dimensions, giving a few positive results. Our highlighted findings include:

1. A truly subquadratic-time algorithm for deciding if the diameter of unit cubes in 3D is at most 3 (Diameter-3 hereafter), the first algorithm of its kind for objects in 3D or higher dimensions. Our algorithm is based on a novel connection to pseudolines, which is of independent interest.
2. A truly subquadratic time lower bound for Diameter-3 of unit balls in 3D under the Orthogonal Vector (OV) hypothesis, giving the first separation between unit balls and unit cubes in the small diameter regime. Previously, computing the diameter for both objects was known to be truly subquadratic hard when the true diameter is $\Omega(\log n)$ [10].
3. A near-linear-time algorithm for Diameter-2 of unit cubes in 3D, generalizing the previous result for unit squares in 2D [10].
4. A truly subquadratic-time algorithm and lower bound for Diameter-2 and Diameter-3 of rectangular boxes (of arbitrary dimension and sizes), respectively.

This paper is one of the two papers in which we undertake a comprehensive study of computing the diameter of geometric intersection graphs for various types of objects. The other paper focuses solely on the 2D case, while this paper is devoted to higher-dimensional cases.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Graph Diameter, Geometric Intersection Graphs, Unit Ball Graphs

1 Introduction

Computing the diameter of sparse graphs is *quadratic hard*: under the Strong Exponential Time Hypothesis (SETH), there is no $O(n^{2-\varepsilon})$ algorithm for distinguishing diameter 2 vs. 3 of graphs with n vertices and $\tilde{O}(n)$[1] edges [32]. That is, the trivial algorithm for computing the diameter by doing BFS from every vertex is essentially optimal. Since then, the major research focus has been on substantially beating the BFS-based algorithm for structural families of graphs, especially planar/minor-free graphs and geometric intersection graphs. For planar/minor-free graphs, truly subquadratic algorithms were known [11, 25, 22, 28, 15]. For geometric intersection graphs, the complexity of computing the diameter remains poorly understood, due to the sheer diversity of geometric objects underlying the graphs and the fact that geometric intersection graphs can have $\Omega(n^2)$ edges.

[1] $\tilde{O}()$ hides polylogarithmic factors

42nd Conference on Very Important Topics (CVIT 2016).
Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:32
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The pioneering work of Bringmann *et al.* [10] studied diameter of intersection graphs of several types of objects. In dimension 3 and above, they considered (axis-parallel) hypercubes[2] and balls. Let *Diameter-Δ* be the problem to decide if a given input graph has diameter at most Δ or at least $\Delta + 1$. They showed that it is quadratic hard to solve: (1) Diameter-2 for hypercubes in $\mathbb{R}^{12}$ under the (3-uniform 6-)Hyperclique Hypothesis; (2) Diameter-3 for (not necessarily axis-aligned) unit segments and equilateral triangles in $\mathbb{R}^2$ under Orthogonal Vector (OV) Hypothesis; and (3) Diameter-$\Omega(\log n)$ for unit balls and unit cubes in $\mathbb{R}^3$, and axis-parallel line segments in $\mathbb{R}^2$, all under the OV Hypothesis.

The hardness results can be interpreted as the radius-Δ neighborhood balls of each corresponding class of geometric shapes are sufficiently complex and expressive, capable of encoding the hard instances used in various fine-grained reductions. One possible way to quantify the complexity of the set system of neighborhood balls is the notion of *VC-dimension*. Given a set system $(U, \mathcal{F})$ with a ground set U and a family $\mathcal{F}$ of subsets of U, its *VC-dimension* is the cardinality of the largest $S \subseteq U$ such that S is *shattered* by $\mathcal{F}$—for every $S' \subseteq S$, there is some $X \in \mathcal{F}$ such that $X \cap S = S'$. We say that a graph G has *distance VC-dimension* at most d if the set system of neighborhood balls $(V_G, \{N^r[v]\}_{r \geq 0})$ has VC dimension at most d. Planar graphs (more generally, minor-free graphs) [19, 9, 29, 23, 28] and intersection graphs of pseudo-disks [2, 24, 17] both have bounded distance VC-dimension, whereas intersection graphs of unit segments and equilateral triangles do not. Recently Chan *et al.* [15] gave truly-subquadratic-time diameter algorithms for squares and unit-disks in the plane, both being instances of pseudo-disks, thus with finite distance VC-dimension. Together, these results hinted at *finite VC-dimension* as an overarching property.

We undertake a comprehensive study of computing diameter on geometric intersection graphs. Our results are presented in a pair of papers. While the companion paper [7] focuses solely on the 2D case, this paper is devoted to three or higher dimensions. In the higher-dimensional case, there is no known truly subquadratic time algorithm for any non-trivial type of objects, even for basic ones such as unit balls or unit cubes. Upper bound techniques for the 2D cases [10, 24, 17, 15] heavily rely on planarity in various places, most notably in bounding the VC dimension using the non-planarity of K_5 [24, 17, 15]. These planarity-specific techniques, in addition to the negative results by Bringmann *et al.* [10], hint at a pessimistic outlook for the complexity of diameter problem in higher dimensions.

▶ **Question 1.** *Does a truly subquadratic time algorithm for computing the diameter exist for any natural class of geometric intersection graphs in 3D or higher dimensions?*

As a starting point, the hardness landscape for unit cubes seems peculiar. Recall that Bringmann *et al.* [10] showed that in 12D, distinguishing diameter 2 versus 3 is already quadratic hard (under the Hyperclique Hypothesis). However for 3D, the quadratic hardness does not kick in until $\Delta \geq \Omega(\log n)$. Duraj, Konieczny, and Potępa [24] solves the 2D case in time $\tilde{O}(\Delta n^{7/4})$ for Diameter-Δ, which was later extended by [15] to solve Diameter in $\tilde{O}(n^{2-1/8})$ time for general Δ. What about unit cubes in 3D? The same question can be asked about unit balls, which typically behave similarly to unit cubes.

▶ **Question 2.** *What is the complexity landscape of Diameter-Δ for unit cubes and unit balls in constant dimensions? Do radius-Δ neighborhood balls of their intersection graphs have bounded VC-dimension, in particular in 3D?*

[2] In this paper, hypercubes and boxes are axis-parallel by default, unless noted otherwise.

Graph class		Lower bound	Upper bound
Unit balls	2D		$O^*(n^{2-1/18})$ for general Δ [15] $O^*(n^{2-1/9})$ for $\Delta = O(1)$ [15] $\tilde{O}(n^{4/3})$ for $\Delta = 2$ (2D Paper [7])
	3D	$\Omega^*(n^2)$ for general Δ (OV) [10] $\Omega^*(n^2)$ for $\Delta = 3$ (OV) (Thm. 8)	
	dD	$\Omega^*(n^2)$ for $\Delta = 2$ & $d = 7$ (combK4) (Thm. 41)	
Unit hypercubes	2D		$O(n \log n)$ for $\Delta = 2$ [10] $O^*(n^{7/4})$ for $\Delta = O^*(1)$ [24] $\tilde{O}(n^{2-1/8})$ for general Δ [15] $O^*(n)$ for any constant Δ (2D Paper [7])
	3D	$\Omega^*(n^2)$ for general Δ (OV) [10]	$\tilde{O}(n)$ for $\Delta = 2$ (Thm. 47) $\tilde{O}(n^{2-1/13})$ for $\Delta = 3$ (Thm. 57)
	4D	$\Omega^*(n^2)$ for $\Delta = 3$ (OV) (Thm. 26)	
	dD	$\Omega^*(n^2)$ for $\Delta = 2$ & $d = 12$ (3H6) [10] $\Omega^*(n^2)$ for $\Delta = 2$ & $d = 10$ (3H6) (Thm. 32) $\Omega^*(n^2)$ for $\Delta = 2$ & $d = 6$ (combK4) (Thm. 35)	
Cubes	2D		$O^*(n^{2-1/12})$ for general Δ [15]
	3D	$\Omega^*(n^2)$ for general Δ (OV) [10] $\Omega^*(n^2)$ for $\Delta = 3$ (OV) (Thm. 29)	$\tilde{O}(n^{2-1/5})$ for $\Delta = 2$ (Thm. 62)
Boxes	2D	$\Omega^*(n^2)$ for $\Delta = 3$ (OV) (Thm 30)	$\tilde{O}(n^{7/4})$ for $\Delta = 2$ (Thm 61)
	3D		$\tilde{O}(n^{2-1/6})$ for $\Delta = 2$ (Thm. 60)

Table 1. Previous and new time bounds for deciding whether an intersection graph of geometric objects has diameter at most Δ. All hypercubes/boxes are axis-aligned. New results are shown in bold. Conditional lower bounds marked "(OV)", "(3H6)", and "(combK4)" assume the Orthogonal Vectors hypothesis, the 3-uniform 6-hyperclique hypothesis, and the combinatorial 4-clique hypothesis, respectively, where the latter is for combinatorial algorithms (all upper bounds are obtained from combinatorial algorithms). See Appendix A for definitions. "2D Paper" refers to the companion paper in the pair [7].

1.1 Our Contributions

We provide a comprehensive study on basic objects in higher dimensions; our results are summarized in Table 1. As corollaries, we make significant progress towards Question 1 and Question 2. Below we discuss each of the contribution in more details.

Unit cubes versus unit balls. Our first result is a separation between unit (hyper)cubes and unit balls: Diameter-3 for unit balls is hard, while Diameter-3 for unit cubes admits a truly subquadratic time algorithm.

▶ **Theorem 3.** *There is no $O(n^{2-\varepsilon})$ algorithm for Diameter-3 for unit balls in $\mathbb{R}^3$ under the OV hypothesis for any $\varepsilon > 0$, while for unit cubes in $\mathbb{R}^d$ for $d \geq 3$, Diameter-3 and Diameter-2 can be solved in $\tilde{O}(n^{2-1/13})$ and $\tilde{O}(n)$ time, respectively.*

Our lower bound for Diameter-3 of unit balls is by reducing from Diameter-2 of sparse tripartite graphs $G = (A \cup B \cup C, E)$, which is quadratic hard under the OV hypothesis. Our first idea for the reduction would be creating a unit ball for each vertex in $A \cup B \cup C$ and encode each edge by intersection: a vertex $u \in A \cup B \cup C$ is incident to edge $uv \in E$ if and only if the unit balls corresponding to u and uv intersect. Then a length-2 path $(a \to b \to c)$, in G will correspond to a length-4 path in the intersection graph $((\!|a|\!) \to (\!|ab|\!) \to (\!|b|\!) \to (\!|bc|\!) \to (\!|c|\!))$, where $(\!|x|\!)$ is the unit ball corresponding to (a vertex or an edge) x. With this idea, we could obtain hardness of Diameter-4. To obtain a lower bound for Diameter-3, we do not create unit balls for the middle set of vertices B; the path would then be $(⦇a⦈ \to ⦇ab⦈ \to ⦇bc⦈ \to ⦇c⦈)$. The difficulty is to guarantee that $⦇ab⦈$ only intersects $⦇bc⦈$, not $⦇\tilde{b}c⦈$ for some other vertex $\tilde{b} \in B$; in the Diameter-4 case, this is enforced by the ball $⦇b⦈$. Here, we achieve the guarantee by introducing another angle parameter δ into the parameterization of the unit balls. This new parameter gives us finer control over the intersections of the balls. See §2.

The algorithms for Diameter-3 and Diameter-2 for unit cubes in Theorem 3 are perhaps most interesting. First, we show that *the VC-dimension of the 2- and 3-neighborhoods of 3D unit-cube graphs is bounded by a constant.* Here we depart from existing techniques for the 2D case [24, 17, 15], which rely on the non-planarity of K_5. In a small region, unit cubes behave like orthants, and so intersection between unit cubes corresponds to a *dominance* relation. We divide into a number of special cases, bound the VC-dimension in each case by 1 or 2 via direct arguments, and then bound the overall VC-dimension by carefully combining subsystems together. See §3.1 and §4.1 for details.

Our VC-dimension bound, in combination with existing techniques [15], is enough to imply truly subquadratic time algorithms for Diameter-2 and Diameter-3, but because of large constants, the running time would be awfully close to quadratic (about $O(n^{1.999998})$ for Diameter-3). We further build on the proof ideas to give more efficient algorithms: In case of Diameter-2, because the subsystems from various cases actually all have VC-dimension 1, we are able to use orthogonal range searching, in combination with divide-and-conquer, to achieve the surprising near-linear time. (Large constants appear in the logarithmic factors, instead of in the main polynomial factor.) For Diameter-3, because the subsystems from various cases actually turn out to be special kinds of 2-dimensional systems corresponding to *pseudoline arrangements*, we are able to use known range searching techniques, in combination with a grid approach, to achieve a more reasonable running time $\tilde{O}(n^{2-1/13})$. (We find it surprising that pseudolines turn out to be relevant for a problem about 3D unit cubes!) The technical details are intricate, as we need to work with abstract pseudolines that are only implicitly represented via certain oracles. See §3 and Appendix H for details.

Our constant bound on the VC-dimension of the 2- and 3-neighborhoods of 3D unit cube graphs suggests an intriguing open question of independent interest to combinatorial geometers: do r-neighborhoods similarly have bounded VC-dimension for all constant $r > 3$?

▶ **Conjecture 4** (VC-Dimension of Unit-Cube Graphs)**.** *There exist a function f such that the VC dimension of $(V, \{N^r[v]\}_{v\in V, r\geq 0})$ is at most $f(r)$, where $N^r[v]$ is the r-neighborhood of v in the unit cube graph $G = (V, E)$.*

Combining with [15], Conjecture 4 would imply that unit-cube graphs of constant diameter admit a truly subquadratic time algorithm for computing diameter, which will be the first natural class of intersection graphs asked in Question 1.

Unit-hypercubes in 4⁺D. Our next result shows that for unit hypercubes in 4D, Diameter-3 is quadratic hard. Therefore, there is no longer a meaningful separation between unit hypercubes and unit balls in dimension 4. When the dimension is higher, say at least 6, we can get hardness results for Diameter-2, under the *3-uniform 6-hyperclique hypothesis (3H6)* (Definition 24) or *combinatorial 4-clique (combK4)* hypothesis (Definition 25), by refining Bringmann *et al.*'s previous proof [10] which required 12 dimensions.

▶ **Theorem 5.** *Let $\varepsilon \in (0, 1)$ be any given parameter. There is no $O(n^{2-\varepsilon})$ algorithm for:*
1. *Diameter-3 for* unit hypercubes in $\mathbb{R}^4$ *under the OV hypothesis.*
2. *Diameter-2 of* unit balls in $\mathbb{R}^7$ *under the CombK4 hypothesis.*
3. *Diameter-2 for* unit hypercubes in $\mathbb{R}^6$ and $\mathbb{R}^{10}$ *under CombK4 / 3H6 hypotheses, resp.*

Cubes and Boxes. In 2D, the intersection graphs of (non-unit) squares admit a truly subquadratic-time algorithm for computing the diameter (of any value) [15]. In 3D, our next result shows that for (non-unit) cubes, Diameter-3 is quadratic hard; this, along with the algorithms in Theorem 3, gives a separation between cubes and their unit counterpart. Somewhat surprisingly, we show that Diameter-2 of (non-unit) cubes is solvable in truly subquadratic time. Even for the more general case of 3D boxes, we also obtain a truly subquadratic time algorithm. (Note that this result is new even for 2D rectangles!) Instead of relying on VC-dimension, this subquadratic algorithm uses a different, grid-based approach. Grid-based approaches have been used before for solving other problems about boxes [14, 16], but one challenge arises from the fact that 3D boxes (or cubes) have quadratic union complexity, unlike 3D unit cubes or orthants. We show that in the case of diameter 2, some of the sides are irrelevant and some boxes can actually be replaced by orthants. See Appendix I.

▶ **Theorem 6.** *There is no $O(n^{2-\varepsilon})$ algorithm for Diameter-3 for (non-unit) cubes under the OV hypothesis for any $\varepsilon > 0$. On the other hand, Diameter-2 can be solved in time $\tilde{O}(n^{2-1/5})$ for cubes and $\tilde{O}(n^{2-1/6})$ for boxes.*

Finite VC-dimension as an if-and-only-if condition? Our collection of results for many types of objects in 2D (in the companion paper [7]) and higher-dimensional cases reinforces the same pattern: truly subquadratic-time algorithms are mostly for graphs with bounded VC dimension. We conjecture that these are special cases of a broader phenomenon.

▶ **Conjecture 7** (VC-Dimension vs Diameter Conjecture)**.** *Diameter-Δ of the intersection graph of low-complexity geometric objects (regardless of the dimension) can be solved in truly subquadratic time if the VC dimension of Δ-neighborhoods is bounded by a constant.*

Here, by low-complexity, we mean that the geometric object has constant description complexity. A positive resolution of Conjecture 7 would provide a powerful algorithmic tool for solving the diameter problem on geometric intersection graphs. A non-trivial test case for Conjecture 7 is (non-unit) disks in 2D: the distance VC dimension is at most 4 [17] whereas currently there is no known truly subquadratic algorithm even for Diameter-2.

It is tempting to strengthen Conjecture 7 to "if and only if"; however, our truly subquadratic algorithms for Diameter-2 of boxes in Theorem 6 provides a counter-example: the intersection graphs of diameter 2 for boxes and rectangles have unbounded VC-dimension. (For proof of unbounded VC-dimension, see Remark 31 in Appendix D.)

Graph preliminaries. Throughout this paper, $G = (V, E)$ be a geometric intersection graph on n geometric objects $\mathcal{O}$. The input will be $\mathcal{O}$ and the graph will be implicit. We use $d(u, v)$ to denote the distance between two vertices $u, v \in V$ in G. We denote the 1-hop *neighborhood* of a vertex $v \in V$ by $N[v]$, and the *r-neighborhood ball* of v by $N^r[v] := \{u \in V : d(u, v) \leq r\}$.

Shatter Dimension. For a set system $(U, \mathcal{F})$, define $Sh(n)$:= $\max_{|S|=n} |\{S \cap F : F \in \mathcal{F}\}|$. The *shatter dimension* of the set system is the minimum d such that $Sh(n) = O(n^d)$. By Sauer-Shelah lemma, If the VC-dimension is d, then shatter dimension is d. The converse partially holds: if the shatter dimension is d, then the VC-dimension is $O(d \log d)$.

2 Diameter-3 for Unit Balls in 3D

▶ **Theorem 8.** *Assuming the OV hypothesis, there is no $O(n^{2-\varepsilon})$ time algorithm for deciding if the intersection graph of a given set of n unit balls in $\mathbb{R}^3$ has diameter at most 3.*

The reduction is from diameter-2 for n-vertex *sparse tripartite graphs* $G = (A \cup B \cup C, E)$, where the number of edges m is $\tilde{O}(n)$. It is well-known that solving Diameter-2 for sparse tripartite graph in truly subquadratic time would break OV hypothesis [32].

▶ **Lemma 9** (Diameter-2 sparse tripartite graphs)**.** *There is no $O(n^{2-\varepsilon})$ time algorithm for deciding if a given sparse tripartite graph has diameter at most 2 assuming OV.*

Let $G = (A \cup B \cup C, E)$ be a tripartite graph with n vertices and $m = \tilde{O}(n)$ edges. It will be more convenient for our construction to work with balls of radius $1/2$ instead of 1, and therefore, all balls in this section have radius $1/2$. Let $\varepsilon = \Theta(1/n^4)$ and $\delta = \Theta(1/n^3)$.

First, we map every vertex v in $A \cup C$ to a distinct number, also denoted by v, in $[\varepsilon, 2\varepsilon]$, such that $\min_{a_1,a_2\in A:\, a_1\neq a_2} |a_1 - a_2| \geq \Omega(\varepsilon/n)$ and $\min_{c_1,c_2\in C:\, c_1\neq c_2} |c_1 - c_2| \geq \Omega(\varepsilon/n)$. Vertices of B are mapped to distinct numbers in $[0.9, 1]$ such that $\min_{b_1,b_2\in B:\, b_1\neq b_2} |b_1 - b_2| \geq \Omega(1/n)$. Then we create four sets of balls as follows.

1. For every vertex $a \in A$, we add a ball centered at $s_a = (1 + a\cos\delta, a\sin\delta, 0)$.
2. For every edge $(a, b) \in (A \times B) \cap E$, we add a ball centered at $p_{ab} = (1 + a\cos\delta - b\sin\delta, a\sin\delta + b\cos\delta, \sqrt{1 - b^2})$.
3. For every edge $(b, c) \in (B \times C) \cap E$, we add a ball centered at $q_{bc} = (-c, b, \sqrt{1 - b^2})$.
4. For every vertex $c \in C$, we add a ball centered at $t_c = (-c, 0, 0)$.

Let $\mathcal{B}_i$ be the set of balls created in step i for $i \in [4]$ and $\mathcal{B} := \cup_i \mathcal{B}_i$. Observe that:

▶ **Observation 10.** *Let $(\!|p|\!)$ be the ball of radius $1/2$ centered at a point p. We have:*
1. *$(\!|s_a|\!) \cap (\!|p_{\tilde{a}b}|\!) \neq \emptyset$ if and only if $a = \tilde{a}$.*
2. *$(\!|p_{ab}|\!) \cap (\!|q_{\tilde{b}c}|\!) \neq \emptyset$ if and only if $b = \tilde{b}$.*
3. *$(\!|p_{bc}|\!) \cap (\!|t_{\tilde{c}}|\!) \neq \emptyset$ if and only if $c = \tilde{c}$.*
4. *Every two balls in $\mathcal{B}_i$ intersect for every $i \in [4]$.*
5. *For any two balls $b_1 \in \mathcal{B}_i$ and $b_2 \in \mathcal{B}_j$ such that $|i - j| \geq 2$, $b_1 \cap b_2 = \emptyset$.*

Proof. For item 1, observe that

$$\begin{aligned}\|s_a - p_{\tilde{a}b}\|_2^2 &= ((a-\tilde{a})\cos\delta + b\sin\delta)^2 + ((a-\tilde{a})\sin\delta - b\cos\delta)^2 + 1 - b^2 \\ &= (a-\tilde{a})^2(\cos^2\delta + \sin^2\delta) + b^2(\sin^2\delta + \cos^2\delta) + 1 - b^2 \ = \ (a-\tilde{a})^2 + 1,\end{aligned}$$

which is exactly 1 if $a = \tilde{a}$, and at least $1 + \Omega(1/n^{10})$ if $|a - \tilde{a}| \geq \Omega(\varepsilon/n)$.

For item 2, observe that

$$\begin{aligned}\|p_{ab} - q_{\tilde{b}c}\|_2^2 &= (1 + a\cos\delta - b\sin\delta + c)^2 + (a\sin\delta + b\cos\delta - \tilde{b})^2 + (\sqrt{1-b^2} - \sqrt{1-\tilde{b}^2})^2 \\ &= (1 + \Theta(\varepsilon) - \Theta(\delta) + \Theta(\varepsilon))^2 + (\Theta(\varepsilon\delta) + b - \Theta(\delta^2) - \tilde{b})^2 + (\sqrt{1-b^2} - \sqrt{1-\tilde{b}^2})^2 \\ &= (1 - \Theta(1/n^3))^2 + (|b - \tilde{b}| \pm \Theta(1/n^6))^2 + (\sqrt{1-b^2} - \sqrt{1-\tilde{b}^2})^2,\end{aligned}$$

which is $1 - \Theta(1/n^3)$ if $b = \tilde{b}$, and at least $1 + \Omega(1/n^2)$ if $|b - \tilde{b}| \geq \Omega(1/n)$.

For item 3, observe that $\|q_{bc} - t_{\tilde{c}}\|_2^2 = (\tilde{c} - c)^2 + 1$, which is exactly 1 if $c = \tilde{c}$, and at least $1 + \Omega(1/n^{10})$ if $|c - \tilde{c}| \geq \Omega(\varepsilon/n)$. Item 4 is easy to check. For item 5, observe that $\|s_a - t_c\|_2 \geq 1 + a\cos\delta + c \geq 1 + \Omega(1/n^4)$, whereas $\|s_a - q_{bc}\|_2$ and $\|p_{ab} - t_c\|_2$ are both $\sqrt{2} \pm o(1)$. ◀

In the real RAM model, balls $\mathcal{B}$ can clearly be constructed in $O(n+m) = \tilde{O}(n)$ time. In the RAM model, it suffices to approximate all coordinates up to $O(1/n^{10})$ additive error, which can be done in $\tilde{O}(1)$ time (we could actually pick "nice" choices of $b \in B$ so that b and $\sqrt{1-b^2}$ are both rational, and also a nice δ so that $\cos\delta$ and $\sin\delta$ are both rational). Thus, the reduction can be done in $\tilde{O}(n)$ time. Theorem 8 follows from the following lemma.

▶ **Lemma 11.** *Let K be the intersection graph of the set of balls $\mathcal{B}$. Then G has diameter at most 2 if and only if K has diameter at most 3.*

Proof. Observe by item 4 of Observation 10 that the balls in the same set $\mathcal{B}_i$ induce a clique in K. By item 5 of Observation 10, there is no edge between two balls in $\mathcal{B}_i$ and $\mathcal{B}_j$ when $|j-i| \geq 2$. Thus, *K has diameter at most 3 if and only if for every ball $[\![s_a]\!] \in \mathcal{B}_1$ and every ball $[\![t_c]\!] \in \mathcal{B}_4$, $d_K(s_a, t_c) \leq 3$.* Here we slightly abuse the notation by using the center of the ball to denote the corresponding vertex in K.

Suppose that G has diameter at most 2. Let $P = (a, b, c)$ be any path of length 2 in G. Then by Observation 10, $(s_a, p_{ab}, q_{bc}, t_c)$ is a path of length 3 in K. Thus, K has diameter at most 3 by the above observation. For the other direction, if K has diameter at most 3, then for any two vertices s_a and t_c (such that $[\![s_a]\!] \in \mathcal{B}_1$ and $[\![t_c]\!] \in \mathcal{B}_4$), a path of length 3 between s_a and t_c in K must be of the form $(s_a, p_{\tilde{a}\tilde{b}}, q_{\tilde{b}c}, t_c)$. Then again by Observation 10, $a = \tilde{a}$ and $b = \tilde{b}$, and therefore, (a, b, c) is a path of length 2 in G. ◀

3 Near-Linear Diameter-2 Algorithm for 3D Unit Cubes

In this section, we present an $\tilde{O}(n)$-time algorithm for testing whether the diameter of a 3D unit cube graph is at most 2. This shows that the near-linear diameter-2 result of Bringmann *et al.* [10] for 2D unit squares surprisingly extends to 3D (ignoring logarithmic factors). It also complements our lower bound result from Theorem 39, which shows conditionally that there are no similar, near-linear combinatorial diameter-2 algorithms for 4D unit hypercubes.

In Section 3.1, we study the combinatorial problem of bounding the VC-dimension of the distance-2 neighborhoods of 3D unit cube graphs as a warm-up. In Appendix G, we use ideas from the combinatorial proof to design a near-linear-time divide-and-conquer algorithm.

For a point $p \in \mathbb{R}^3$, denote $[\![p]\!]$ the unit cube centered at p. For simplicity, we assume that the input is in general position, e.g., all coordinate values are distinct. We will solve the problem in a slightly more general setting for 3 point sets P, Q, R, testing whether all distances between $[\![p]\!]$ and $[\![r]\!]$ for $(p, r) \in P \times R$ are exactly 2 in the tripartite intersection graph of the unit cubes centered at P, Q, R. The original problem reduces to the case when P, Q, R are identical sets (or almost identical, if we want to ensure general position).

3.1 VC-dimension bound

For two points $p, q \in \mathbb{R}^3$, write $p \prec_x q$ if p has smaller x-coordinate than q, and $p \succ_x q$ if p has larger x-coordinate than q. Define $\prec_y$, $\succ_y$, $\prec_z$, $\succ_z$ similarly. Define the trivial relation $\mathbf{T}$ which is always true. A *generalized dominance relation* is a relation $\lhd$ where $p \lhd q$ iff $p \lhd_x q$ and $p \lhd_y q$ and $p \lhd_z q$, for some choices of $\lhd_x \in \{\prec_x, \succ_x, \mathbf{T}\}$, $\lhd_y \in \{\prec_y, \succ_y, \mathbf{T}\}$, and $\lhd_z \in \{\prec_z, \succ_z, \mathbf{T}\}$.

The motivation for considering generalized dominance relations is this: Consider a uniform grid with unit side length. Inside a grid cell, any unit cube is equivalent to an orthant. If p lies in a grid cell and q lies in a neighboring grid cell, then the condition that $[\![p]\!]$ and $[\![q]\!]$ intersect corresponds precisely to $p \lhd q$ for some generalized dominance relation.

Let P, Q, R be 3 point sets in $\mathbb{R}^3$. Let $\lhd_1, \lhd_2$ be 2 generalized dominance relations in $\mathbb{R}^3$. For each $r \in R$, we can write $N^2[r]$ as $\{p \in P : \exists q \in Q \text{ with } p \lhd_1 q \text{ and } q \lhd_2 r\}$. Define the set system $\mathcal{S}_{\lhd_1,\lhd_2}(P, Q, R) \coloneqq (P, \{N^2[r] : r \in R\})$. We first prove that this set system has bounded VC-dimension by bounding its shatter dimension.

Previous proofs of bounded distance VC-dimension usually involve planarity arguments (avoidance of K_5), but for the 3D problem here, we use a different strategy. We divide into easier cases based on the separability of the given point sets—in the following, we say that two sets are *x-separated* if they are separated by a plane orthogonal to the x-axis; we define *y-* and *z-separation* similarly. We first show how to handle the case when P and Q are both x- and y-separated by a simple direct argument:

▶ **Lemma 12.** *If P and Q are both x- and y-separated, then $\mathcal{S}_{\lhd_1,\lhd_2}(P, Q, R)$ has VC-dimension* 1*.*

Proof. First, for any $\lhd_i$, we define $\lhd_{ix}$, $\lhd_{iy}$, and $\lhd_{iz}$ to be the x-, y-, and z-coordinate dominance relation, respectively.

Consider 2 points $p, p' \in P$. Suppose $\{p\}$ and $\{p'\}$ are both shattered. Then there exist $r, r' \in R$ with $p \in N^2[r]$, $p' \notin N^2[r]$, $p' \in N^2[r']$, $p \notin N^2[r']$. Let $q, q' \in Q$ with $p \lhd_1 q \lhd_2 r$ and $p' \lhd_1 q' \lhd_2 r'$.

Because P and Q are x- and y-separated, we already know that $p \lhd_{1x} q'$, $p \lhd_{1y} q'$, and $p' \lhd_{1x} q$, $p' \lhd_{1y} q'$. Note that $p \lhd_{1z} q'$ or $p' \lhd_{1z} q$, because otherwise, $q' \lhd_{1z} p \lhd_{1z} q \lhd_{1z} p' \lhd_{1z} q'$: a contradiction. Thus, $p \lhd_1 q'$ or $p' \lhd_1 q$, implying $p \in N^2[r']$ or $p' \in N^2[r]$. ◀

We then handle the case when P and R are separated along all 3 axes.

▶ **Lemma 13.** *If P and R are x-, y-, and z-separated, then $\mathcal{S}_{\lhd_1,\lhd_2}(P, Q, R)$ has shatter dimension at most* 6*.*

Proof. W.l.o.g., say $P \subset (-\infty, 0)^3$ and $R \subset (0, \infty)^3$. Consider the following 6 set systems of the form $\mathcal{S}_{\lhd_1,\lhd_2}(P, Q \cap \Lambda, R)$, where Λ can be one of

$$((0,\infty) \times (0,\infty) \times \mathbb{R}), \quad (\mathbb{R} \times (0,\infty) \times (0,\infty)), \quad ((0,\infty) \times \mathbb{R} \times (0,\infty)),$$

$$((-\infty,0) \times (-\infty,0) \times \mathbb{R}), \quad (\mathbb{R} \times (-\infty,0)) \times (-\infty,0)), \quad ((-\infty,0) \times \mathbb{R} \times (-\infty,0)).$$

Each set system has VC-dimension at most 1 by Lemma 12 (or symmetric variants).

Each set in $\mathcal{S}_{\lhd_1,\lhd_2}(P, Q, R)$ can be expressed as the union of 6 sets (one of each) from these 6 set systems (since a point $q \in Q$ must has two positive coordinates or two negative coordinates). So, the number of distinct sets in $\mathcal{S}_{\lhd_1,\lhd_2}(P, Q, R)$ is at most $O(|P|^6)$. ◀

It remains to reduce the general case to the case when P and R are separated along all 3 axes. We accomplish this by a simple grid idea:

▶ **Lemma 14.** *For any $P, Q, R \subset \mathbb{R}^3$, $\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P, Q, R)$ has shatter dimension at most* 51*.*

Proof. Build a (non-uniform) grid formed by the x-, y-, and z-coordinates of all the points of P. The grid has $O(|P|^3)$ cells. For each grid cell γ, consider the 8 octants at an arbitrary point inside γ; for each such octant τ, the set system $\mathcal{S}_{\lhd_1,\lhd_2}(P \cap \tau, Q, R \cap \gamma)$ has shatter dimension at most 6 by Lemma 13. For each $r \in R \cap \gamma$, the set $N^3[r]$ is the union of 8 sets (one of each) from the set systems for the 8 octants. So for each γ, the number of distinct sets in $\{N^3[r] : r \in R \cap \gamma\}$ is $O((|P|^6)^8)$, giving $|\{N^3[r] : r \in R\}| = O(|P|^3 \cdot (|P|^6)^8)$. ◀

We can now bound the shatter dimension for unit cubes.

▶ **Lemma 15.** *For any $P, Q, R \subset \mathbb{R}^3$ with $R \subset (0,1)^3$, rewrite $N^2[r]$ as*

$$\Big\{p \in P : \exists q \in Q \text{ with } [\![q]\!] \text{ intersecting both } [\![p]\!] \text{ and } [\![r]\!]\Big\}.$$

Then the set system $(P, \{N^2[r] : r \in R\})$ has shatter dimension at most 4131.

Proof. Fix $\alpha = (\alpha_P, \alpha_Q) \in (\mathbb{Z}^3)^2$ with $\|\alpha_R\|_\infty, \|\alpha_P - \alpha_Q\|_\infty \le 1$. For $p \in \alpha_P + (0,1)^3$ and $q \in \alpha_Q + (0,1)^3$, $[\![p]\!]$ intersects $[\![q]\!]$ iff $p - \alpha_P \lhd_1 q - \alpha_Q$, where we define $\lhd_{1x}$ to be $\succ_x$ if $x(\alpha_Q) = x(\alpha_P) + 1$, and $\prec_x$ if $x(\alpha_Q) = x(\alpha_P) - 1$, and $\mathbf{T}$ if $x(\alpha_Q) = x(\alpha_P)$, and we define $\lhd_{1y}$ and $\lhd_{1z}$ in the same way. Similarly, for $q \in \alpha_Q + (0,1)^3$ and $r \in (0,1)^3$, $[\![q]\!]$ intersects $[\![r]\!]$ iff $q - \alpha_Q \lhd_2 r$ for some generalized dominance relation $\lhd_2$. Define the set system $\mathcal{S}_\alpha = \mathcal{S}_{\lhd_1,\lhd_2}(P \cap (\alpha_P + (0,1)^3), (Q \cap (\alpha_Q \cap (0,1)^3)) + \alpha_P - \alpha_Q, R + \alpha_P)$.

Each set $N^2[r]$ is the union of at most 9^2 sets (one of each) from the set systems $\mathcal{S}_\alpha$ for the at most 9^2 choices of α. By Lemma 14, we have $|\{N^2[r] : r \in R\}| = O((|P|^{51})^{9^2})$. ◀

▶ **Theorem 16.** *For any $P, Q, R \subset \mathbb{R}^3$, the set system $(P, \{N^2[r] : r \in R\})$ has shatter dimension at most 4131.*

Proof. Build a uniform grid of side length 1. For each grid cell $\alpha_R + (0,1)^3$ (with $\alpha_R \in \mathbb{Z}^3$), the number of different $N^2[r]$ sets over all $r \in R \cap (\alpha_R + (0,1)^3)$ is $O(|P \cap (\alpha_R + (-2,3)^3)|^{4131})$ by Lemma 15. Since each point $p \in P$ belongs to $O(1)$ number of expanded cells $\alpha_R + (-2,3)^3$, the sum over all α_R is $O(|P|^{4131})$. ◀

In the interest of simplicity, we have not optimized the above constant (which admittedly is quite large, but fortunately will not matter in our final algorithm).

3.2 Algorithm

Knowing that the VC-dimension is bounded, we could at this point apply the framework of Chan et al. [15] to obtain a subquadratic algorithm for Diameter-2 of 3D unit cubes. However, the exponent would be very close to 2 (much worse than our more general result for arbitrary 3D boxes in Theorem 60). We present a faster, direct algorithm that runs in *near-linear* time. The key observation is that although the final shatter dimension bound is large, the proof in Section 3.1 tells us that the set system is in some sense "made up of" a constant number of simpler subsystems with much smaller VC-dimension, namely, VC-dimension 1 (Lemma 12)! We can't just solve the problem for each subsystem separately—the diameter problem isn't decomposable that way. Instead, we "encode" each subsystem by adding a single coordinate value to each element, and in the end reduce the whole problem to an orthogonal range searching problem for vectors in a sufficiently large constant dimension. Due to lack of space, the details of the algorithm will be given in Appendix G.

4 Subquadratic Diameter-3 Algorithm for 3D Unit Cubes

In this section, we present a subquadratic algorithm for testing whether the diameter of a 3D unit cube graph is at most 3. This complements our lower bound results, which show conditionally that there are no similar diameter-3 algorithms for 3D unit ball graphs (Theorem 8) nor for 4D unit hypercube graphs (Theorem 26). This result is more challenging than our result for diameter 2 in Section 3. As before, we warm up by studying the corresponding combinatorial problem of bounding the VC-dimension of the 3-neighborhoods of 3D unit cube graphs (Section 4.1). The combinatorial proof will reveal a surprising connection to 2D pseudolines. We review known techniques about pseudolines in Appendix H.1, as preparation towards our final subquadratic algorithm, presented in Appendix H.2.

For a point $p \in \mathbb{R}^3$, recall that $[\![p]\!]$ denotes the unit cube centered at p. For simplicity, we assume that all coordinate values are distinct. Like before, we will solve the problem in a slightly more general setting for 4 point sets P, Q, R, S, testing whether all distances between $[\![p]\!]$ and $[\![s]\!]$ for $(p, s) \in P \times S$ are exactly 3 in the 4-partite intersection graphs of the unit cubes centered at P, Q, R, S.

4.1 VC-dimension bound

Following Section 3.1, we begin with the corresponding problem for generalized dominance relations. Let P, Q, R, S be 4 point sets in $\mathbb{R}^3$. Let $\lhd_1, \lhd_2, \lhd_3$ be 3 generalized dominance relations in $\mathbb{R}^3$. For each $s \in S$, we can write $N^3[s]$ as $\{p \in P : \exists (q, r) \in Q \times R$ with $p \lhd_1 q$ and $q \lhd_2 r$ and $r \lhd_3 s\}$. Define the set system $\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P, Q, R, S) := (P, \{N^3[s] : s \in S\})$. We first prove that this set system has bounded shatter dimension.

As in Section 3.1, we do not use planarity arguments like previous proofs. Instead, we divide into cases based on separability of the given point sets. As we go from 2- to 3-neighborhoods, we face more challenges. The first case turns out to be the most crucial (and interesting), where as the one that follows it is similar to Lemma 13:

▶ **Lemma 17.** *If P and Q are x-separated, Q and R are y-separated, and R and S are z-separated, then the set system $\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P, Q, R, S)$ has VC-dimension at most 2.*

Proof. Suppose $p \lhd_j q$ iff $p \lhd_{jx} q$ and $p \lhd_{jy} q$ and $p \lhd_{jz} q$ with $\lhd_{jx} \in \{\prec_x, \succ_x, \mathbf{T}\}$, $\lhd_{jy} \in \{\prec_y, \succ_y, \mathbf{T}\}$, and $\lhd_{jz} \in \{\prec_z, \succ_z, \mathbf{T}\}$. We may assume that $\lhd_{1x} \neq \mathbf{T}$ (because if not, we can shift all the x-coordinates of P sufficiently far downward and replace $\lhd_{1x}$ with $\prec_x$), and similarly, all other $\lhd_{jx}, \lhd_{jy}, \lhd_{jz}$ are not $\mathbf{T}$. We first prove the following property:

Property: Let $p, p' \in P$ and $s, s' \in S$. If $p \in N^3[s]$, $p' \in N^3[s']$, $p \notin N^3[s']$, $p' \notin N^3[s]$, then $p \prec_y p'$ implies $s' \blacktriangleleft_y s$, where $\blacktriangleleft_y \in \{\prec_y, \succ_y\}$ is a relation determined solely from the choices of $\lhd_1, \lhd_2, \lhd_3$.

Proof: Let $q, q' \in Q$ and $r, r' \in R$ with $p \lhd_1 q \lhd_2 r \lhd_3 s$ and $p' \lhd_1 q' \lhd_2 r' \lhd_3 s'$. We may assume neither $p \lhd_1 q'$ nor $p' \lhd_1 q$, because otherwise, $p \in N^3[s']$ or $p' \in N^3[s]$. Because P and Q are x-separated, we already know $p \lhd_{1x} q'$ and $p' \lhd_{1x} q$. So, there are two remaining possibilities: (i) $p \lhd_{1y} q \lhd_{1y} p' \lhd_{1y} q'$ and $p' \lhd_{1z} q' \lhd_{1z} p \lhd_{1z} q$, or (ii) $p' \lhd_{1y} q' \lhd_{1y} p \lhd_{1y} q$ and $p \lhd_{1z} q \lhd_{1z} p' \lhd_{1z} q'$. (See Figure 1.)

Define $\blacktriangleleft_{1z} \in \{\prec_z, \succ_z\}$ to be $\lhd_{1z}$ iff $\lhd_{1y}$ is $\prec_y$. Then $p \prec_y p'$ implies $q' \blacktriangleleft_{1z} q$.

Similarly, we may assume neither $q \lhd_2 r'$ nor $q' \lhd_2 r$, because otherwise, $p \in N^3[s']$ or $p' \in N^3[s]$. Because Q and R are y-separated, $q \lhd_{2y} r'$ and $q' \lhd_{2y} r$. There are two remaining possibilities: (i) $q \lhd_{2z} r \lhd_{2z} q' \lhd_{2z} r'$ and $q' \lhd_{2x} r' \lhd_{2x} q \lhd_{2x} r$, or (ii) $q' \lhd_{2z} r' \lhd_{2z} q \lhd_{2z} r$ and $q \lhd_{2x} r \lhd_{2x} q' \lhd_{2x} r'$.

Define $\blacktriangleleft_{2x} \in \{\prec_x, \succ_x\}$ to be $\lhd_{2x}$ iff $\lhd_{2z}$ is $\blacktriangleleft_{1z}$. Then $q' \blacktriangleleft_{1z} q$ implies $r \blacktriangleleft_{2x} r'$.

Similarly (and lastly), we may assume neither $r \lhd_3 s'$ nor $r' \lhd_3 s$, because otherwise, $p \in N^3[s']$ or $p' \in N^3[s]$. Because R and S are z-separated, $r \lhd_{3z} s'$ and $r' \lhd_{3z} s$. There are two remaining possibilities: (i) $r \lhd_{3x} s \lhd_{3x} r' \lhd_{3x} s'$ and $r' \lhd_{3y} s' \lhd_{3y} r \lhd_{3y} s$, or (ii) $r' \lhd_{3x} s' \lhd_{3x} r \lhd_{3x} s$ and $r \lhd_{3y} s \lhd_{3y} r' \lhd_{3y} s'$.

Define $\blacktriangleleft_y \in \{\prec_y, \succ_y\}$ to be $\lhd_{3y}$ iff $\lhd_{3x}$ is $\blacktriangleleft_{2x}$. Then $r \blacktriangleleft_{2x} r'$ implies $s' \blacktriangleleft_y s$. ☐

(i) (ii)

Figure 1. In the yz plane, if p is dominated by q and p' is dominated by q', but p is not dominated by q' and p' is not dominated by q, then (i) and (ii) depict the only two possibilities.

Now, consider 3 points $p, p', p'' \in P$, with $p \prec_y p' \prec_y p''$. Suppose $\{p, p''\}$ and $\{p'\}$ are both shattered. Then there exist $s, s' \in S$ with $p, p'' \in N^3[s]$, $p' \notin N^3[s]$, $p' \in N^3[s']$, $p, p'' \notin N^3[s']$. Applying the property twice yields $s' \blacktriangleleft_y s$ and $s \blacktriangleleft_y s'$: a contradiction. ◀

▶ **Lemma 18.** *If (a) P and Q are both x- and y-separated, or (b) Q and R are both x- and y-separated, or (c) R and S are both x- and y-separated, then $\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P,Q,R,S)$ has VC-dimension 1.*

Proof. Consider 2 points $p, p' \in P$. Suppose $\{p\}$ and $\{p'\}$ are both shattered. Then there exist $s, s' \in S$ with $p \in N^3[s]$, $p' \notin N^3[s]$, $p' \in N^3[s']$, $p \notin N^3[s']$. Let $q, q' \in Q$ and $r, r' \in R$ with $p \lhd_1 q \lhd_2 r \lhd_3 s$ and $p' \lhd_1 q' \lhd_2 r' \lhd_3 s'$.

For (a), $p \lhd_{1x} q'$, $p \lhd_{1y} q'$, and $p' \lhd_{1x} q$, $p' \lhd_{1y} q'$ by the separation assumptions. Note that $p \lhd_{1z} q'$ or $p' \lhd_{1z} q$, because otherwise, $q' \lhd_{1z} p \lhd_{1z} q \lhd_{1z} p' \lhd_{1z} q'$: a contradiction. Thus, $p \lhd_1 q'$ or $p' \lhd_1 q$, implying $p \in N^3[s']$ or $p' \in N^3[s]$.

Claims (b) and (c) follow from a similar argument. ◀

We now handle the case when P and S are separated along all 3 axes, by invoking the previous cases some constant number of times:

▶ **Lemma 19.** *If P and S are x-, y-, and z-separated, then $\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P,Q,R,S)$ has shatter dimension at most 21.*

Proof. W.l.o.g., say $P \subset (-\infty, 0)^3$ and $S \subset (0, \infty)^3$. Consider the following 15 set systems.

$$\begin{array}{llll}
\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P, & Q \cap ((0,\infty) \times (-\infty,0) \times \mathbb{R}), & R \cap (\mathbb{R} \times (0,\infty) \times (-\infty,0)), & S) \\
\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P, & Q \cap (\mathbb{R} \times (0,\infty) \times (-\infty,0)), & R \cap ((-\infty,0) \times \mathbb{R} \times (0,\infty)), & S) \\
\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P, & Q \cap ((-\infty,0) \times \mathbb{R} \times (0,\infty)), & R \cap ((0,\infty) \times (-\infty,0) \times \mathbb{R}), & S) \\
\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P, & Q \cap ((-\infty,0) \times (0,\infty) \times \mathbb{R}), & R \cap ((0,\infty) \times \mathbb{R} \times (-\infty,0)), & S) \\
\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P, & Q \cap (\mathbb{R} \times (-\infty,0) \times (0,\infty)), & R \cap ((-\infty,0) \times (0,\infty) \times \mathbb{R}), & S) \\
\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P, & Q \cap ((0,\infty) \times \mathbb{R} \times (-\infty,0)), & R \cap (\mathbb{R} \times (-\infty,0) \times (0,\infty)), & S) \\
\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P, & Q \cap ((0,\infty) \times (0,\infty) \times \mathbb{R}), & R, & S) \\
\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P, & Q \cap ((0,\infty) \times \mathbb{R} \times (0,\infty)), & R, & S) \\
\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P, & Q \cap (\mathbb{R} \times (0,\infty) \times (0,\infty)), & R, & S) \\
\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P, & Q \cap ((-\infty,0) \times (-\infty,0) \times \mathbb{R}), & R \cap ((0,\infty) \times (0,\infty) \times \mathbb{R}), & S) \\
\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P, & Q \cap ((-\infty,0) \times \mathbb{R} \times (-\infty,0)), & R \cap ((0,\infty) \times \mathbb{R} \times (0,\infty)), & S) \\
\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P, & Q \cap (\mathbb{R} \times (-\infty,0) \times (-\infty,0)), & R \cap (\mathbb{R} \times (0,\infty) \times (0,\infty)), & S)
\end{array}$$

$$\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P,\quad Q,\quad R\cap((-\infty,0)\times(-\infty,0)\times\mathbb{R}),\quad S)$$
$$\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P,\quad Q,\quad R\cap((-\infty,0)\times\mathbb{R}\times(-\infty,0)),\quad S)$$
$$\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P,\quad Q,\quad R\cap(\mathbb{R}\times(-\infty,0]\times(-\infty,0)),\quad S).$$

The first 6 set systems have VC-dimension at most 2 by Lemma 17 (or symmetric variants); the next 9 have VC-dimension 1 by Lemma 18(a,b,c) (or symmetric variants).

Each set in $\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P,Q,R,S)$ can be expressed as the union of 15 sets one from each of these 15 set systems.[3] So, $|\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P,Q,R,S)| = O((|P|^2)^6 \cdot |P|^9)$. ◀

We can now reduce the general case to the case when P and S are separated along all 3 axes.

▶ **Lemma 20.** *For any $P,Q,R,S \subset \mathbb{R}^3$, $\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P,Q,R,S)$ has shatter dimension ≤ 171.*

Proof. Build a (non-uniform) grid formed by the x-, y-, and z-coordinates of all the points of P. The grid has $O(|P|^3)$ cells. For each grid cell γ, consider the 8 octants at an arbitrary point inside γ; for each such octant τ, the set system $\mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P\cap\tau,Q,R,S\cap\gamma)$ has shatter dimension at most 18 by Lemma 19. For each $s \in S\cap\gamma$, the set $N^3[s]$ is the union of 8 sets (one of each) from the set systems for the 8 octants. So for each γ, the number of distinct sets in $\{N^3[s] : s \in S\cap\gamma\}$ is $O((|P|^{21})^8)$. Thus, we have $|\{N^3[s] : s \in S\}| = O(|P|^3 \cdot (|P|^{21})^8)$. ◀

As before, unit cubes reduce to generalized dominance relations:

▶ **Lemma 21.** *For any $P,Q,R,S \subset \mathbb{R}^3$ with $S \subset (0,1)^3$, rewrite $N^3[s]$ (again) as*

$$\Big\{p \in P : \exists (q,r) \in Q\times R \text{ s.t. } [\![p]\!] \text{ intersects } [\![q]\!],\ [\![q]\!] \text{ intersects } [\![r]\!], \text{ and } [\![r]\!] \text{ intersects } [\![s]\!]\Big\}.$$

Then the set system $(P, \{N^3[s] : s \in S\})$ has shatter dimension at most 124659.

Proof. Fix $\alpha = (\alpha_P,\alpha_Q,\alpha_R) \in (\mathbb{Z}^3)^3$ with $\|\alpha_R\|_\infty, \|\alpha_Q-\alpha_R\|_\infty, \|\alpha_P-\alpha_Q\|_\infty \leq 1$. For $p \in \alpha_P + (0,1)^3$ and $q \in \alpha_Q + (0,1)^3$, $[\![p]\!]$ intersects $[\![q]\!]$ iff $p-\alpha_P \lhd_1 q-\alpha_Q$, for some generalized dominance relation $\lhd_1$ as in the proof of Lemma 15. Similarly, for $q \in \alpha_Q+(0,1)^3$ and $r \in \alpha_R + (0,1)^3$, $[\![q]\!]$ intersects $[\![r]\!]$ iff $q-\alpha_Q \lhd_2 r-\alpha_R$ for some generalized dominance relation $\lhd_2$. Similarly, for $r \in \alpha_R + (0,1)^3$ and $s \in (0,1)^3$, $[\![r]\!]$ intersects $[\![s]\!]$ iff $r-\alpha_R \lhd_3 s$ for some generalized dominance relation $\lhd_3$. Define the set system $\mathcal{S}_\alpha = \mathcal{S}_{\lhd_1,\lhd_2,\lhd_3}(P\cap(\alpha_P+(0,1)^3), (Q\cap(\alpha_Q\cap(0,1)^3))+\alpha_P-\alpha_Q, (R\cap(\alpha_R\cap(0,1)^3))+\alpha_P-\alpha_R, S+\alpha_P)$.

Each set $N^3[s]$ is the union of at most 9^3 sets (one of each) from the set systems $\mathcal{S}_\alpha$ for the at most 9^3 choices of α. By Lemma 20, we have $|\{N^3[s] : s \in S\}| = O((|P|^{171})^{9^3})$. ◀

▶ **Theorem 22.** *For any $P,Q,R,S \subset \mathbb{R}^3$, the set system $(P,\{N^3[s] : s \in S\})$ has shatter dimension at most 124659.*

Proof. Build a uniform grid of side length 1. For each grid cell $\alpha_S + (0,1)^3$ (with $\alpha_S \in \mathbb{Z}^3$), the number of different $N^3[s]$ sets over all $s \in S\cap(\alpha_S+(0,1)^3)$ is $O(|P\cap(\alpha_S+(-3,4)^3)|^{124659})$ by Lemma 21. Since each point $p \in P$ belongs to $O(1)$ number of expanded cells $\alpha_S+(-3,4)^3$, the sum over all α_S is $O(|P|^{124659})$. ◀

---

[3] In any point sequence $\langle p,q,r,s\rangle$ with $p \in (-\infty,0)^3$ and $s \in (0,\infty)^3$, there must be a consecutive pair going from negative to positive x-coordinate, and a consecutive pair going from negative to positive y-coordinate, and a consecutive pair going from negative to positive z-coordinate; all 3 such pairs could be distinct, or 2 of them could be the same.

4.2 Algorithm

Knowing that the VC-dimension is bounded, we could at this point apply the framework of [15] to obtain a subquadratic algorithm for diameter-3 in 3D unit cubes. However, the exponent would be *extremely* close to 2 (around $2 - 1/(4 \cdot 124660) \approx 1.999998$). We present a faster, direct algorithm with a more reasonable exponent smaller than 2. The key observation is that although the final shatter dimension bound is large, the proof in Section 4.1 tells us that the set system is in some sense "made up of" a (very large but) constant number of simpler subsystems with much smaller VC-dimension, of 1 and 2 (Lemma 17 and Lemma 18). As in our diameter-2 algorithm, we can't just solve the problem for each subsystem separately, because the diameter problem isn't decomposable. Instead, we "encode" each subsystem as an extra "2-dimensional constraint", and in the end reduce the whole problem to a "multi-level" range searching problem in a sufficiently large constant dimension.

Multi-level range searching [3, 31] is usually solved by sampling-based geometric divide-and-conquer techniques, e.g., via so-called "cuttings", but unfortunately analogs of cuttings provably do not exist for abstract set systems, even with VC-dimension 2 (e.g., see [8, Remark 4.7]). Fortunately, our proof of VC-dimension 2 in Lemma 17 reveals a stronger property, namely, that the 3-neighborhoods there actually form a *point-pseudoline system*—see Appendix H.1. (It is quite unexpected that 3D unit cubes could naturally produce pseudoline arrangements.) As noted in Lemma 50, cuttings and multi-level range searching exists for pseudolines. Due to lack of space, details of the algorithm are given in Appendix H.

References

1 Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. If the current clique algorithms are optimal, so is Valiant's parser. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 98–117. IEEE, October 2015.

2 A. Karim Abu-Affash, Paz Carmi, Anil Maheshwari, Pat Morin, Michiel Smid, and Shakhar Smorodinsky. Approximating maximum diameter-bounded subgraph in unit disk graphs. *Discrete & Computational Geometry*, 66(4):1401–1414, 2021. doi:10.1007/s00454-021-00327-y.

3 Pankaj K Agarwal. Range searching. In *Handbook of discrete and computational geometry*, pages 1057–1092. Chapman and Hall/CRC, 2017.

4 Pankaj K. Agarwal and Jeff Erickson. Geometric range searching and its relatives. In B. Chazelle, J. E. Goodman, and R. Pollack, editors, *Advances in Discrete and Computational Geometry*, pages 1–56. AMS Press, 1999.

5 Pankaj K. Agarwal and Micha Sharir. Pseudo-line arrangements: Duality, algorithms, and applications. *SIAM J. Comput.*, 34(3):526–552, 2005. doi:10.1137/S0097539703433900.

6 Pankaj K. Agarwal, Micha Sharir, and Alex Steiger. Decomposing the complement of the union of cubes and boxes in three dimensions. *Discret. Comput. Geom.*, 72(2):407–450, 2024. URL: https://doi.org/10.1007/s00454-024-00632-2, doi:10.1007/S00454-024-00632-2.

7 Anonymous Author(s). Charting the diameter computation landscape on intersection graphs in 2D, 2025. In submission to SoCG 2026.

8 Sujoy Bhore, Timothy M. Chan, Zhengcheng Huang, Shakhar Smorodinsky, and Csaba D. Tóth. Sparse bounded hop-spanners for geometric intersection graphs. In *Proc. 41st International Symposium on Computational Geometry (SoCG)*, volume 332 of *LIPIcs*, pages 17:1–17:15, 2025. doi:10.4230/LIPICS.SOCG.2025.17.

9 Nicolas Bousquet and Stéphan Thomassé. VC-dimension and Erdős–Pósa property. *Discrete Mathematics*, 338(12):2302–2317, 2015. doi:10.1016/j.disc.2015.05.026.

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

11 Sergio Cabello. Subquadratic algorithms for the diameter and the sum of pairwise distances in planar graphs. *ACM Transactions on Algorithms*, 15(2):1–38, December 2018.

12 Timothy M Chan. A (slightly) faster algorithm for Klee's measure problem. In *Proceedings of the twenty-fourth annual symposium on Computational geometry*, pages 94–100, June 2008.

13 Timothy M. Chan. Near-optimal randomized algorithms for selection in totally monotone matrices. In Dániel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 1483–1495. SIAM, 2021. doi:10.1137/1.9781611976465.89.

14 Timothy M. Chan. Finding triangles and other small subgraphs in geometric intersection graphs. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1777–1805, 2023. doi:10.1137/1.9781611977554.CH68.

15 Timothy M. Chan, Hsien-Chih Chang, Jie Gao, Sándor Kisfaludi-Bak, Hung Le, and Da Wei Zheng. Truly subquadratic time algorithms for diameter and related problems in graphs of bounded VC-dimension. In *66th Annual Symposium on Foundations of Computer Science*, FOCS '25, 2025. URL: https://arxiv.org/abs/2510.16346.

16 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. In *Proc. 50th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 34:1–34:19, 2023. doi:10.4230/LIPICS.ICALP.2023.34.

17 Hsien-Chih Chang, Jie Gao, and Hung Le. Computing diameter+2 in truly-subquadratic time for unit-disk graphs. In *40th International Symposium on Computational Geometry (SoCG)*, pages 38:1–38:14, 2024. Arxiv version https://arxiv.org/abs/2401.12881. doi:10.4230/LIPIcs.SoCG.2024.38.

18 Bernard Chazelle. Cutting hyperplanes for divide-and-conquer. *Discret. Comput. Geom.*, 9:145–158, 1993. doi:10.1007/BF02189314.

19 Victor Chepoi, Bertrand Estellon, and Yann Vaxes. Covering planar graphs with a fixed number of balls. *Discrete & Computational Geometry*, 37(2):237–244, 2007. doi:10.1007/s00454-006-1260-0.

20 Kenneth L. Clarkson and Peter W. Shor. Application of random sampling in computational geometry, II. *Discret. Comput. Geom.*, 4:387–421, 1989. doi:10.1007/BF02187740.

21 Mark de Berg, Otfried Cheong, Marc J. van Kreveld, and Mark H. Overmars. *Computational Geometry: Algorithms and Applications.* Springer, 3rd edition, 2008. doi:10.1007/978-3-540-77974-2.

22 Guillaume Ducoffe, Michel Habib, and Laurent Viennot. Diameter computation on h-minor free graphs and graphs of bounded (distance) VC-dimension. In *30th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Proceedings, pages 1905–1922. 2019.

23 Guillaume Ducoffe, Michel Habib, and Laurent Viennot. Diameter, eccentricities and distance oracle computations on h-minor free graphs and graphs of bounded (distance) Vapnik–Chervonenkis dimension. *SIAM Journal on Computing*, 51(5):1506–1534, 2022. doi:10.1137/20M136551X.

24 Lech Duraj, Filip Konieczny, and Krzysztof Potępa. Better diameter algorithms for bounded VC-dimension graphs and geometric intersection graphs. In *32nd Annual European Symposium on Algorithms (ESA)*, volume 308, pages 51:1–51:18, 2024.

25 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 Journal on Computing*, 50(2):509–554, 2021. doi:10.1137/18M1193402.

26 Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, March 2001.

27 Balázs Keszegh and Dömötör Pálvölgyi. An abstract approach to polychromatic coloring: shallow hitting sets in aba-free hypergraphs and pseudohalfplanes. *J. Comput. Geom.*, 10(1):1–26, 2019. URL: https://doi.org/10.20382/jocg.v10i1a1, doi:10.20382/JOCG.V10I1A1.
28 Hung Le and Christian Wulff-Nilsen. VC set systems in minor-free (di)graphs and applications. In *35th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 5332–5360, 2024. doi:10.1137/1.9781611977912.192.
29 Jason Li and Merav Parter. Planar diameter via metric compression. In *51st Annual ACM Symposium on Theory of Computing (STOC)*, pages 152–163, Phoenix, AZ, USA and New York, NY, USA, 2019.
30 Andrea Lincoln, Virginia Vassilevska Williams, and Ryan Williams. Tight hardness for shortest cycles and paths in sparse graphs. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1236–1252. Society for Industrial and Applied Mathematics, Philadelphia, PA, January 2018.
31 Jiří Matoušek. Range searching with efficient hierarchical cuttings. *Discret. Comput. Geom.*, 10:157–182, 1993. doi:10.1007/BF02573972.
32 Liam Roditty and Virginia Vassilevska Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In *45th Annual ACM Symposium on Theory of Computing (STOC)*, pages 515–524, 2013.
33 Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theor. Comput. Sci.*, 348(2-3):357–365, December 2005.

A Fine Grained Complexity Hypothesis

Recent development of fine-grained complexity has identified a few hypothesis that we include here for completeness. We use them to prove our lower bounds.

▶ **Definition 23** (Orthogonal Vectors (OV) hypothesis)**.** *Given sets A, B of n vectors in $\{0,1\}^d$, $d = \omega(\log n)$, deciding whether there exists an orthogonal pair $a \in A$, $b \in B$ requires $n^{2-o(1)}$ time.*

The OV-hypothesis is implied [33] by the Strong Exponential Time Hypothesis [26].

▶ **Definition 24** (3-uniform 6-hyperclique (3H6) hypothesis)**.** *Given a 6-partite 3-uniform hypergraph $G = (V, E)$ where V is the disjoint union of vertex set $V^{(1)}, \ldots, V^{(6)}$, each containing n vertices, and $E \subseteq \binom{V}{3}$ such that each edge connects three vertices from different vertex sets. The problem is to decide whether there are 6 vertices $S = \{v_1, v_2, \ldots, v_6\}$ with $v_i \in V^{(i)}$, $i = 1, \ldots, 6$, forming a 6-clique, i.e., $\{v_i, v_j, v_k\} \in E$ for all $\{i, j, k\} \in \binom{S}{3}$. The 3H6 hypothesis says that the problem requires $n^{6-o(1)}$ time.*

More information about the hyperclique hypothesis can be found in [30].

▶ **Definition 25** (Combinatorial 4-clique (combK4) hypothesis)**.** *Any combinatorial algorithm detecting whether a graph of n vertices contains a 4-clique requires $n^{4-o(1)}$ time.*

More information about the combinatorial 4-clique problem and connections to other hypothesis can be found in [12, 1].

B Diameter-3 Lower Bounds for Unit Hypercubes in 4D

In this and the next two sections, we follow the same proof approach as in Section 2 to establish conditional lower bounds for Diameter-3 for other types of objects: 4D unit hypercubes, 3D cubes, and 2D rectangles. The details are a little simpler than in the proof for 3D unit balls.

▶ **Theorem 26.** *Assuming the OV hypothesis, there is no $O(n^{2-\varepsilon})$ time algorithm for deciding if the intersection graph of a given set of n unit hypercubes in $\mathbb{R}^4$ has diameter at most 3.*

We provide a $\tilde{O}(n)$-time reduction from Diameter-2 for sparse tripartite graphs. Let $G = (A \cup B \cup C, E)$ be a sparse tripartite graph. For every point c, we denote by $[\![c]\!]$ the unit hypercube centered at c. First, we map vertices of $A \cup B \cup C$ to (arbitrary) distinct numbers in $[0, 0.1]$. We abuse the notation by using v to denote the value that a vertex v is mapped to. Then we create hypercubes in the following steps:

1. For every vertex $a \in A$, we add a hypercube centered at $s_a = \{a, -a, 0, 0\}$.
2. For every edge $(a, b) \in (A \times B) \cap E$, we add a hypercube centered at $p_{ab} = \{1 + a, 1 - a, 0.5 + b, 0.5 - b\}$.
3. For every edge $(b, c) \in (B \times C) \cap E$, we add a hypercube centered at $q_{bc} = \{1 + c, 1 - c, 1.5 + b, 1.5 - b\}$.
4. For every $c \in C$, we add a hypercube $t_c = \{c, -c, 2, 2\}$.

Let $\mathcal{Q}_i$ be the set of hypercubes created in step i for $i \in [4]$. Let $\mathcal{Q} = \mathcal{Q}_1 \cup \mathcal{Q}_2 \cup \mathcal{Q}_3 \cup \mathcal{Q}_4$ be the set of resulting hypercubes. The following observation explains our design of $\mathcal{Q}$.

▶ **Observation 27.** *We have:*

1. $[\![s_{\tilde{a}}]\!] \cap [\![p_{ab}]\!] \neq \emptyset$ if and only if $\tilde{a} = a$.
2. $[\![p_{a\tilde{b}}]\!] \cap [\![q_{bc}]\!] \neq \emptyset$ *if and only if* $b = \tilde{b}$.
3. $[\![t_{\tilde{c}}]\!] \cap [\![q_{bc}]\!] \neq \emptyset$ *if and only if* $\tilde{c} = c$.
4. *Any two cubes in $\mathcal{Q}_i$ intersects for every $i \in [4]$.*
5. *Any two cubes $(x, y) \in \mathcal{Q}_i \times \mathcal{Q}_j$ for $i, j \in [4]$ such that $|i - j| \geq 2$ are disjoint: $x \cap y = \emptyset$.*

Clearly $\mathcal{Q}$ can be constructed in $O(n + m) = \tilde{O}(n)$ time. Thus, Theorem 26 follows from the following lemma.

▶ **Lemma 28.** *Let K be the intersection graph of $\mathcal{Q}$. Then G has diameter at most 2 if and only if K has diameter at most 3.*

Proof. By the same argument in Lemma 11 using Observation 27. ◀

C Diameter-3 Lower Bound for 3D Cubes

▶ **Theorem 29.** *Assuming the OV hypothesis, there is no $O(n^{2-\varepsilon})$ time algorithm for deciding if the intersection graph of a given set of n cubes (all "close" to unit) in $\mathbb{R}^3$ has diameter at most 3.*

The reduction is again from Diameter-2 for sparse tripartite graphs. Let $G = (A \cup B \cup C, E)$ be a tripartite graph with n vertices and $m = \tilde{O}(n)$ edges.

First, we map every vertex v in $A \cup C$ to a distinct number, also denoted by v, in $(0, \varepsilon]$, and every vertex v in B to a distinct number in $[1 - \varepsilon, 1)$. Then we create sets of cubes as follows.

1. For every vertex $a \in A$, we add a cube $s_a = [a, a + 1] \times [a - 1, a] \times [1, 2]$.
2. For every edge $(a, b) \in (A \times B) \cap E$, we add a cube $p_{ab} = [2a - b, a] \times [a, b] \times [b, 2b - a]$ (with side length $b - a$).
3. For every edge $(b, c) \in (B \times C) \cap E$, we add a cube $q_{bc} = [2c - b, c] \times [b, 2b - c] \times [c, b]$ (with side length $b - c$).
4. For every vertex $c \in C$, we add a cube $t_c = [c, c + 1] \times [1, 2] \times [c - 1, c]$.

The result follows as before.

D Diameter-3 Lower Bound for 2D Rectangles

▶ **Theorem 30.** *Assuming the OV hypothesis, there is no $O(n^{2-\varepsilon})$ time algorithm for deciding if the intersection graph of a given set of n rectangles (all "close" to unit squares) in $\mathbb{R}^2$ has diameter at most 3.*

The reduction is again from Diameter-2 for sparse tripartite graphs. Let $G = (A \cup B \cup C, E)$ be a tripartite graph with n vertices and $m = \tilde{O}(n)$ edges.

First, we map every vertex v in $A \cup B \cup C$ to a distinct number, also denoted by v, in $(0, \varepsilon]$. Then we create sets of rectangles as follows.

1. For every vertex $a \in A$, we add a rectangle $s_a = [-1 + a, a] \times [2 + a, 3 + a]$.
2. For every edge $(a, b) \in (A \times B) \cap E$, we add a rectangle $p_{ab} = [a, 1 + b] \times [1 + b, 2 + a]$.
3. For every edge $(b, c) \in (B \times C) \cap E$, we add a rectangle $q_{bc} = [1 + b, 2 + c] \times [c, 1 + b]$.
4. For every vertex $c \in C$, we add a rectangle $t_c = [2 + c, 3 + c] \times [-1 + c, c]$.

▶ Remark 31. A slight modification of this construction shows that diameter-2 rectangles have unbounded VC-dimension. Consider any set system $(\mathcal{S}, X)$ and define a bipartite graph $G = (A \cup B, E)$ where $A = \mathcal{S}$, $B = X$, and there is an edge from the vertex of A corresponding to every set $S \in \mathcal{S}$ to the vertex of B corresponding to elements $x \in S$. Consider the rectangles s_a and p_{ab} from the first two steps of the above construction, and additionally rectangles $q_b = [1 + b, 2 + b] \times [b, 1 + b]$ for every $b \in B$. The diameter-2 balls centered at s_a restricted to the squares q_b corresponds exactly to $(\mathcal{S}, X)$.

E Diameter-2 Lower Bound for Hypercubes

We next turn to conditional lower bounds for Diameter-2, using the hyperclique hypothesis instead of OV. We revisit Bringmann *et al.*'s result for 12D hypercubes [10], and show how to slightly lower the number of dimensions to 10, by a more careful construction that uses some of the dimensions in a more economical way.

▶ **Theorem 32.** *Assuming the 3-uniform 6-hyperclique hypothesis, there is no $O(n^{2-\varepsilon})$ time algorithm for deciding if the intersection graph of a given set of n unit hypercubes in $\mathbb{R}^{10}$ has diameter at most 2.*

We provide a reduction from 3-uniform 6-hyperclique to Diameter-2 for 10D unit hypercubes. Let $G = (V_A \cup V_B \cup V_C \cup V_D \cup V_E \cup V_F, E)$ be a 6-partite 3-uniform hypergraph with $n^{1/3}$ vertices.

First, we map every vertex v in $V_A \cup V_B \cup V_C$ to a distinct number, also denoted by v, in $[0.3, 0.4]$, and every vertex v in $V_D \cup V_E \cup V_F$ to a distinct number in $[0, 0.1]$. Then we create the following sets of $O(n)$ unit hypercubes in $\mathbb{R}^{10}$ of side length 1 as follows.

1. For every hyperedge $(a, b, c) \in (V_A \times V_B \times V_C) \cap E$, we add a hypercube centered at

$$s_{abc} = (a, 1 + a,\ b, 1 + b,\ c, 1 + c,\ 0.5, 0.5,\ 0.5, 0.5).$$

2. For every hyperedge $(d, e, f) \in (V_D \times V_E \times V_F) \cap E$, we add a hypercube centered at

$$t_{def} = (1 + d, d,\ 1 + d, d,\ 1 + d, d,\ 1 + e, e,\ 1 + f, f).$$

3. For every non-hyperedge $(a, b, d) \in (V_A \times V_B \times V_D) \setminus E$, we add a hypercube centered at

$$p_{abd} = (1 + a, a,\ 1 + b, b,\ d, 1 + d,\ 0.5, 0.5,\ 0.5, 0.5).$$

4. For every non-hyperedge $(a, b, e) \in (V_A \times V_B \times V_E) \setminus E$, we add a hypercube centered at

$$p_{abe} = (1 + a, a,\ 1 + b, b,\ 0.5, 0.5,\ e, 1 + e,\ 0.5, 0.5).$$

5. For every non-hyperedge $(a, b, f) \in (V_A \times V_B \times V_F) \setminus E$, we add a hypercube centered at

$$p_{abf} = (1 + a, a,\ 1 + b, b,\ 0.5, 0.5,\ 0.5, 0.5,\ f, 1 + f).$$

6. Similarly do the last 3 steps with a, b replaced by b, c or a, c.
7. For every non-hyperedge $(a, d, e) \in (V_A \times V_D \times V_E) \setminus E$, we add a hypercube centered at

$$p_{ade} = (1 + a, a,\ 0.5, 0.5,\ d, 1 + d,\ e, 1 + e,\ 0.5, 0.5).$$

8. For every non-hyperedge $(a, d, f) \in (V_A \times V_D \times V_F) \setminus E$, we add a hypercube centered at

$$p_{adf} = (1 + a, a,\ 0.5, 0.5,\ d, 1 + d,\ 0.5, 0.5,\ f, 1 + f).$$

9. For every non-hyperedge $(a, e, f) \in (V_A \times V_E \times V_F) \setminus E$, we add a hypercube centered at

$$p_{aef} = (1 + a, a,\ 0.5, 0.5,\ 0.5, 0.5,\ e, 1 + e,\ f, 1 + f).$$

10. Similarly do the last 3 steps with a replaced by b or c.
11. We add an extra hypercube centered at $z_0 = (0.8, 0.2,\ 0.8, 0.2,\ 0.8, 0.2,\ 0.8, 0.2,\ 0.8, 0.2)$.

It is not difficult to verify the following:

▶ **Observation 33.** *We have:*
1. $[\![s_{abc}]\!] \cap [\![p_{\tilde{a}\tilde{b}d}]\!] \neq \emptyset$ if and only if $\tilde{a} = a$ *and* $\tilde{b} = b$.
2. $[\![t_{def}]\!] \cap [\![p_{ab\tilde{d}}]\!] \neq \emptyset$ *if and only if* $\tilde{d} = d$.
3. $[\![p_{abd}]\!] \cap [\![z_0]\!] \neq \emptyset$.
4. $[\![s_{abc}]\!] \cap [\![t_{def}]\!] = \emptyset$ *and* $[\![s_{abc}]\!] \cap [\![z_0]\!] = \emptyset$.

Similar statements hold for p_{abe}*, etc.*

▶ **Lemma 34.** *G contains a 6-hyperclique if and only if the intersection graph of the above boxes has diameter greater than* 2.

Proof. If G contains a 6-hyperclique $(a, b, c, d, e, f) \in V_A \times V_B \times V_C \times V_D \times V_E \times V_F$, then $[\![s_{abc}]\!]$ and $[\![t_{def}]\!]$ have distance more than 2 in the intersection graph, since according to Observation 33 the only possible common neighbors of $[\![s_{abc}]\!]$ and $[\![t_{def}]\!]$ are $[\![p_{abd}]\!]$, $[\![p_{abe}]\!]$, $[\![p_{abf}]\!]$, etc., but all these cubes are not present because $(a, b, d), (a, b, e), (a, b, f), \ldots \in E$.

Conversely, if two vertices have distance more than 2, then according to Observation 33 they must be of the form $[\![s_{abc}]\!]$ and $[\![t_{def}]\!]$ with $(a, b, c), (d, e, f) \in E$. Since $[\![p_{abd}]\!]$, $[\![p_{abe}]\!]$, $[\![p_{abf}]\!]$, etc. are not common neighbors of $[\![s_{abc}]\!]$ and $[\![t_{def}]\!]$, we have $(a, b, d), (a, b, e), (a, b, f), \ldots \in E$, and so (a, b, c, d, e, f) is a 6-hyperclique in G. ◀

Thus, if Diameter-2 for 10D unit hypercubes could be solved in truly subquadratic time, then we would have an algorithm for 6-hyperclique in a 3-uniform graph with $n^{1/3}$ vertices running in $O((n^{1/3})^{6-\varepsilon})$ time, violating the 3H6 hypothesis.

E.1 Variant 1

We can modify the above proof to lower the dimension further to 6 for combinatorial algorithms, if we work under the combinatorial 4-clique hypothesis instead:

▶ **Theorem 35.** *Assuming the combinatorial 4-clique hypothesis, there is no $O(n^{2-\varepsilon})$ time combinatorial algorithm for deciding if the intersection graph of a given set of n unit hypercubes in $\mathbb{R}^6$ has diameter at most 2.*

We reduce from the 4-clique problem instead. THe reduction can be viewed as a simplification of our previous reduction from 3-uniform 6-hyperclique. Let $G = (A \cup B \cup C \cup D, E)$ be a 4-partite graph with $\sqrt{n}$ vertices.

First, we map every vertex v in $A \cup B$ to a distinct number, also denoted by v, in $[0.3, 0.4]$, and every vertex v in $C \cup D$ to a distinct number in $[0, 0.1]$. Then we create the following sets of $O(n)$ unit hypercubes in $\mathbb{R}^6$ of side length 1 as follows.

1. For every edge $(a, b) \in (A \times B) \cap E$, we add a hypercube centered at
$$s_{ab} = (a, 1 + a,\ b, 1 + b,\ 0.5, 0.5).$$
2. For every edge $(c, d) \in (C \times D) \cap E$, we add a hypercube centered at
$$t_{cd} = (1 + c, c,\ 1 + c, c,\ 1 + d, d).$$
3. For every non-edge $(a, c) \in (A \times C) \setminus E$, we add a hypercube centered at
$$p_{ac} = (1 + a, a,\ c, 1 + c,\ 0.5, 0.5).$$
4. For every non-edge $(a, d) \in (A \times D) \setminus E$, we add a hypercube centered at
$$p_{ad} = (1 + a, a,\ 0.5, 0.5,\ d, 1 + d).$$
5. For every non-edge $(b, c) \in (B \times C) \setminus E$, we add a hypercube centered at
$$p_{bc} = (c, 1 + c,\ 1 + b, b,\ 0.5, 0.5).$$
6. For every non-edge $(b, d) \in (B \times D) \setminus E$, we add a hypercube centered at
$$p_{bd} = (0.5, 0.5,\ 1 + b, b,\ d, 1 + d).$$
7. We add an extra hypercube centered at $z_0 = (0.8, 0.2,\ 0.8, 0.2,\ 0.8, 0.2)$.

It is not difficult to verify the following:

▶ **Observation 36.** *We have:*
1. $[\![s_{ab}]\!] \cap [\![p_{\tilde{a}c}]\!] \neq \emptyset$ if and only if $\tilde{a} = a$.
2. $[\![t_{cd}]\!] \cap [\![p_{c\tilde{d}}]\!] \neq \emptyset$ *if and only if* $\tilde{d} = d$.
3. $[\![p_{ac}]\!] \cap [\![z_0]\!] \neq \emptyset$.
4. $[\![s_{ab}]\!] \cap [\![t_{cd}]\!] = \emptyset$ *and* $[\![s_{ab}]\!] \cap [\![z_0]\!] = \emptyset$.

Similar statements hold for p_{ad}, p_{bc}, and p_{bd}.

▶ **Lemma 37.** *G contains a 4-clique if and only if the intersection graph of the above boxes has diameter greater than* 2.

Proof. Similar to the proof of Lemma 34, using Observation 36. ◀

Thus, if Diameter-2 for 6D unit hypercubes has a truly subquadratic time, combinatorial algorithm, then we would have a combinatorial algorithm for 6-hyperclique in a 3-uniform graph with $\sqrt{n}$ vertices running in $O((\sqrt{n})^{4-\varepsilon})$ time, violating the combinatorial 4-clique hypothesis.

▶ Remark 38. For non-combinatorial algorithms, the above reduction still implies a non-trivial $\Omega(n^{3/2-\varepsilon})$ lower bound, under the hypothesis that 4-clique for graphs with n vertices requires near-cubic time. (The fastest 4-clique algorithm known needs at least cubic time, even if the matrix multiplication exponent were to be 2.)

E.2 Variant 2

Even more substantially, we can lower the dimension down to 4, but with a catch: the lower bound is weaker (though still superlinear for combinatorial algorithms).

▶ **Theorem 39.** *Assuming the combinatorial 3-clique (or equivalently combinatorial BMM) hypothesis, there is no $O(n^{3/2-\varepsilon})$ time combinatorial algorithm for deciding if the intersection graph of a given set of n unit hypercubes in $\mathbb{R}^4$ has diameter at most 2. The same holds for deciding if the maximum eccentricity of a given subset of size $\sqrt{n}$ in the intersection graph is at most 2.*

We reduce from the 3-clique problem, i.e., triangle detection. The reduction can be viewed as a further simplification of the reduction from the previous subsection. Let $G = (A \cup B \cup C, E)$ be a tripartite graph with $\sqrt{n}$ vertices.

First, we map every vertex v in $A \cup B$ to a distinct number, also denoted by v, in $[0.3, 0.4]$, and every vertex v in C to a distinct number in $[0, 0.1]$. Then we create the following sets of $O(n)$ unit hypercubes in $\mathbb{R}^4$ of side length 1 as follows.

1. For every edge $(a, b) \in (A \times B) \cap E$, we add a hypercube centered at $s_{ab} = (a, 1+a,\ b, 1+b)$.
2. For every vertex $c \in C$, we add a hypercube centered at $t_c = (1 + c, c,\ 1 + c, c)$.
3. For every non-edge $(a, c) \in (A \times C) \setminus E$, we add a hypercube centered at $p_{ac} = (1 + a, a,\ c, 1 + c)$.
4. For every non-edge $(b, c) \in (B \times C) \setminus E$, we add a hypercube centered at $p_{bc} = (c, 1 + c,\ 1 + b, b)$.
5. We add an extra hypercube centered at $z_0 = (0.8, 0.2,\ 0.8, 0.2)$.

As before, we can then show that G contains a triangle if and only if the intersection graph of the above boxes has diameter greater than 2. The second statement in the theorem also follows because the number of the t_c's is only $\sqrt{n}$.

▶ Remark 40. Interestingly, although the first statement of Theorem 39 may not be tight, the second is: we can compute the eccentricity of $\sqrt{n}$ vertices by running BFS $\sqrt{n}$ times, in total $\tilde{O}(n^{3/2})$ time (since one can implement BFS in the intersection graph of unit hypercubes in near linear time via orthogonal range searching data structures).

F Diameter-2 Lower Bound for Unit Balls

We can adapt our conditional lower bound proof for Diameter-2 for 6D unit hypercubes in Appendix E.1 under the combinatorial 4-clique hypothesis, to the case of 7D unit balls:

▶ **Theorem 41.** *Assuming the combinatorial 4-clique hypothesis, there is no $O(n^{2-\varepsilon})$ time combinatorial algorithm for deciding if the intersection graph of a given set of n unit balls in $\mathbb{R}^7$ has diameter at most 2.*

Let $G = (A \cup B \cup C \cup D, E)$ be a 4-partite graph with $\sqrt{n}$ vertices.

Let $\varepsilon = \Theta(1/n^3)$. First, we map every vertex v in $A \cup B \cup C$ to a distinct number, also denoted by v, in $[0, 0.1]$, such that $\min_{a_1,a_2 \in A:\, a_1 \neq a_2} |a_1 - a_2| \geq \Omega(1/n)$, $\min_{b_1,b_2 \in B:\, b_1 \neq b_2} |b_1 - b_2| \geq \Omega(1/n)$, and $\min_{c_1,c_2 \in C:\, c_1 \neq c_2} |c_1 - c_2| \geq \Omega(1/n)$. Vertices of D are mapped to distinct numbers in $[1-\varepsilon, 1]$ such that $\min_{d_1,d_2 \in D:\, d_1 \neq d_2} |d_1 - d_2| \geq \Omega(\varepsilon/n)$. Then we create the following sets of $O(n)$ balls in $\mathbb{R}^7$ of radius $\sqrt{2}$ as follows.

1. For every edge $(a, b) \in (A \times B) \cap E$, we add a ball centered at
$$s_{ab} = (a, \sqrt{1-a^2},\ b, \sqrt{1-b^2},\ 0, 0,\ 0).$$
2. For every edge $(c, d) \in (C \times D) \cap E$, we add a ball centered at
$$t_{cd} = (0, 0,\ 0, 0,\ c, \sqrt{1-c^2},\ d).$$
3. For every non-edge $(a, c) \in (A \times C) \setminus E$, we add a ball centered at
$$p_{ac} = (a, \sqrt{1-a^2},\ 0, 0,\ c, \sqrt{1-c^2},\ 0).$$
4. For every non-edge $(a, d) \in (A \times D) \setminus E$, we add a ball centered at
$$p_{ad} = (a, \sqrt{1-a^2},\ 0, 0,\ 0, 0,\ d).$$
5. For every non-edge $(b, c) \in (B \times C) \setminus E$, we add a ball centered at
$$p_{bc} = (0, 0,\ b, \sqrt{1-b^2},\ c, \sqrt{1-c^2},\ 0).$$
6. For every non-edge $(b, d) \in (B \times D) \setminus E$, we add a ball centered at
$$p_{bd} = (0, 0,\ b, \sqrt{1-b^2},\ 0, 0,\ d).$$
7. We add an extra ball centered at $z_1 = (0, 0,\ 0, 0,\ 0.5, 0.5,\ 0.5)$.

It is not difficult to verify an analog of Observation 36, and so the rest of the proof proceeds as before.

G Near-Linear Diameter-2 Algorithm for 3D Unit Cubes (Continued)

As in Section 3.1, we begin with the corresponding problem for generalized dominance relations. As in Lemma 12, we first consider the special case when P and Q are both x- and y-separated:

▶ **Lemma 42.** *We can preprocess a point set $Q \subset \mathbb{R}^3$ in $\tilde{O}(|Q|)$ time so that the following holds. Given point sets $P \subset (-\infty, \mu_x) \times (-\infty, \mu_y) \times \mathbb{R}$ and $R \subset \mathbb{R}^3$ for some μ_x, μ_y, and given generalized dominance relations $\lhd_1$ and $\lhd_2$, we can compute mappings $\phi : P \to \mathbb{R}$ and $\psi : R \to \mathbb{R}$ in $\tilde{O}(|P| + |R|)$ time, satisfying the following property for every $(p, r) \in P \times R$:*
$$(\exists q \in Q \cap ((\mu_x, \infty) \times (\mu_y, \infty) \times \mathbb{R}) \text{ with } p \lhd_1 q \text{ and } q \lhd_2 r) \iff \phi(p) < \psi(r).$$

Proof. We may assume that $\lhd_{1x} \in \{\prec_x, \mathbf{T}\}$ and $\lhd_{1y} \in \{\prec_y, \mathbf{T}\}$ (because if not, we can trivially set $\phi = 1$ and $\psi = 0$). We may assume that $\lhd_{1z} \neq \mathbf{T}$ (because if not, we can replace all z-coordinates of P with a sufficiently small negative number and replace $\lhd_{2z}$ with $\prec_z$). If $\lhd_{1z} = \prec_z$, we define $\phi(p)$ to be the z-coordinate of p, and define $\psi(r)$ to be the largest z-coordinate among all points $q \in Q \cap ((\mu_x, \infty) \times (\mu_y, \infty) \times \mathbb{R})$ with $q \lhd_2 r$. The property is obviously satisfied. Furthermore, ψ can be evaluated in $\tilde{O}(1)$ time each by an orthogonal range max query [3], assuming that Q has been preprocessed in $\tilde{O}(|Q|)$ time (and ϕ is trivial to evaluate). The case when $\lhd_{1z} = \succ_z$ is similar (by negating all z-coordinates). ◀

Next, as in Lemma 13, we consider the case when P and R are separated along all 3 axes:

▶ **Lemma 43.** *We can preprocess a point set $Q \subset \mathbb{R}^3$ in $\tilde{O}(|Q|)$ time so that the following holds. Given point sets $P \subset (-\infty, \mu_x) \times (-\infty, \mu_y) \times (-\infty, \mu_z)$ and $R \subset (\mu_x, \infty) \times (\mu_y, \infty) \times (\mu_z, \infty)$ for some μ_x, μ_y, μ_z, and given generalized dominance relations $\lhd_1$ and $\lhd_2$, we can compute mappings $\phi : P \to \mathbb{R}^6$ and $\psi : R \to \mathbb{R}^6$ in $\tilde{O}(|P| + |R|)$ time, satisfying the following property for every $(p, r) \in P \times R$:*

$$(\exists q \in Q \text{ with } p \lhd_1 q \text{ and } q \lhd_2 r) \iff \phi(p) \text{ does not dominate } \psi(r).$$

Proof. We compute mappings $\phi_1, \ldots, \phi_6 : P \to \mathbb{R}$ and $\psi_1, \ldots, \psi_6 : R \to \mathbb{R}$ satisfying the following properties:

1. $(\exists q \in Q \cap ((\mu_x, \infty) \times (\mu_y, \infty) \times \mathbb{R})$ with $p \lhd_1 q$ and $q \lhd_2 r) \iff \phi_1(p) < \psi_1(r)$;
2. $(\exists q \in Q \cap (\mathbb{R} \times (\mu_y, \infty) \times (\mu_z, \infty))$ with $p \lhd_1 q$ and $q \lhd_2 r) \iff \phi_2(p) < \psi_2(r)$;
3. $(\exists q \in Q \cap ((\mu_x, \infty) \times \mathbb{R} \times (\mu_z, \infty))$ with $p \lhd_1 q$ and $q \lhd_2 r) \iff \phi_3(p) < \psi_3(r)$;
4. $(\exists q \in Q \cap ((-\infty, \mu_x) \times (-\infty, \mu_y) \times \mathbb{R})$ with $p \lhd_1 q$ and $q \lhd_2 r)) \iff \phi_4(p) < \psi_4(r)$.
5. $(\exists q \in Q \cap (\mathbb{R} \times (-\infty, \mu_y) \times (-\infty, \mu_z))$ with $p \lhd_1 q$ and $q \lhd_2 r) \iff \phi_5(p) < \psi_5(r)$;
6. $(\exists q \in Q \cap ((-\infty, \mu_x) \times \mathbb{R} \times (-\infty, \mu_z))$ with $p \lhd_1 q$ and $q \lhd_2 r) \iff \phi_6(p) < \psi_6(r)$.

Each such mapping can be computed by Lemma 42 (possibly with x-, y-, z-coordinates permuted and/or negated, and/or P and R swapped). Finally, we define $\phi(p) = (\phi_1(p), \ldots, \phi_6(p))$ and $\psi(r) = (\psi_1(r), \ldots, \psi_6(r))$. ◀

We now transform the result from dominance to unit cubes:

▶ **Lemma 44.** *We can preprocess a point set $Q \subset \mathbb{R}^3$ in $\tilde{O}(|Q|)$ time so that the following holds. Given point sets $P \subset \alpha_P + ((0, \mu_x) \times (0, \mu_y) \times (0, \mu_z))$ and $R \subset \alpha_R + ((\mu_x, 1) \times (\mu_y, 1) \times (\mu_z, 1))$ for some $\mu_x, \mu_y, \mu_z \in (0, 1)$ and $\alpha_P, \alpha_R \in \mathbb{Z}^3$, we can compute mappings $\phi : P \to \mathbb{R}^{54}$ and $\psi : R \to \mathbb{R}^{54}$ in $\tilde{O}(|P| + |R|)$ time, satisfying the following property for every $(p, r) \in P \times R$:*

$$(\exists q \in Q \text{ with } [\![q]\!] \text{ intersecting both } [\![p]\!] \text{ and } [\![r]\!]) \iff \phi(p) \text{ does not dominate } \psi(r).$$

Proof. For each $\alpha_Q \in \mathbb{Z}^3$ with L_∞-distance at most 1 from both α_P and α_R, we compute a mapping $\phi^{(\alpha_Q)}$ satisfying the following property:

$$(\exists q \in Q \cap (\alpha_Q + (0, 1)^3) \text{ with } [\![q]\!] \text{ intersecting both } [\![p]\!] \text{ and } [\![r]\!]) \iff \phi^{(\alpha_Q)}(p) < \psi^{(\alpha_Q)}(r).$$

Each such mapping can be computed by Lemma 43. This is because for $p \in \alpha_P + (0, 1)^3$ and $q \in \alpha_Q + (0, 1)^3$, $[\![p]\!]$ intersects $[\![q]\!]$ iff $p - \alpha_P \lhd_1 q - \alpha_Q$, for some generalized dominance relation $\lhd_1$ as in the proof of Lemma 15. Similarly, for $q \in \alpha_Q + (0, 1)^3$ and $r \in \alpha_R + (0, 1)^3$, $[\![q]\!]$ intersects $[\![r]\!]$ iff $q - \alpha_Q \lhd_2 r - \alpha_R$ for some generalized dominance relation $\lhd_2$. Finally, we define ϕ and ψ as the Cartesian products of $\phi^{(\alpha_Q)}$ and $\psi^{(\alpha_Q)}$ respectively over all at most 9 choices of α_Q. ◀

We then solve the problem by orthogonal range searching:

▶ **Lemma 45.** *We can preprocess a point set $Q \subset \mathbb{R}^3$ in $\tilde{O}(|Q|)$ time so that the following holds. Given point sets $P \subset \alpha_P + ((0, \mu_x) \times (0, \mu_y) \times (0, \mu_z))$ and $R \subset \alpha_R + ((\mu_x, 1) \times (\mu_y, 1) \times (\mu_z, 1))$ for some $\mu_x, \mu_y, \mu_z \in (0, 1)$ and $\alpha_P, \alpha_R \in \mathbb{Z}^3$, we can decide whether for all $(p, r) \in P \times R$, there exists $q \in Q$ with $[\![q]\!]$ intersecting both $[\![p]\!]$ and $[\![r]\!]$, in $\tilde{O}(|P| + |R|)$ time.*

Proof. After constructing the mappings from Lemma 44, we can test for the non-existence of a pair $(p, r) \in P \times R$ with $\phi(p)$ dominating $\phi(r)$ and p not intersecting r, by 54-dimensional orthogonal range searching in $\tilde{O}(|P| + |R|)$ time. ◀

Having solved the problem for the case when P and R (modulo 1) are separated along all 3 axes, we still need to reduce the general problem to this case. To this end, we could mimic the grid approach in the proof of Lemma 14, but this would increase the running time (although this type of grid approach will still be useful in later sections, we more ambitiously aim for near-linear running time in this section). Instead, we adopt a *divide-and-conquer* algorithm to reduce to the separable case, which increases the running time only by a polylogarithmic factor. The divide-and-conquer is similar in style to that of range trees [21, 4], and may appear standard, but one unusual and interesting feature is that we are reducing only the sizes of P and R during recursion, while Q stays the same. This might seem problematic, but luckily in our setting, the cost at every recursive step is dependent only on $|P|$ and $|R|$ and not on $|Q|$, after an initial global preprocessing of Q (we are not so lucky in our later algorithms).

▶ **Lemma 46.** *We can preprocess a point set $Q \subset \mathbb{R}^3$ in $\tilde{O}(|Q|)$ time so that the following holds. Given point sets $P \subset \alpha_P + (0,1)^3$ and $R \subset \alpha_R + (0,1)^3$ for some $\alpha_P, \alpha_R \in \mathbb{Z}^3$, we can decide whether for all $(p,r) \in P \times R$, there exists $q \in Q$ with $[\![q]\!]$ intersecting both $[\![p]\!]$ and $[\![r]\!]$, in $\tilde{O}(|P| + |R|)$ time.*

Proof. We use "range-tree-style" divide-and-conquer. First consider the special case when $P \subset \alpha_P + ((0,\mu_x) \times (0,\mu_y) \times (0,1))$ and $R \subset \alpha_R + ((\mu_x,1) \times (\mu_y,1) \times (0,1))$, for a given μ_x, μ_y. To solve the problem in this special case:

1. Let μ_z be the median z-coordinate in $(P-\alpha_P)\cup(R-\alpha_R)$. Let $P^- = P \cap (\alpha_P + ((0,\mu_x) \times (0,\mu_y) \times (0,\mu_z)))$, $P^+ = P \cap (\alpha_P + (0,\mu_x) \times (0,\mu_y) \times (\mu_z,1)))$, $R^- = R \cap (\alpha_R + ((0,\mu_x) \times (0,\mu_y) \times (0,\mu_z)))$, and $R^+ = R \cap (\alpha_R + (0,\mu_x) \times (0,\mu_y) \times (\mu_z,1)))$.
2. Solve the problem for P^- and R^+ by Lemma 45.
3. Solve the problem for P^+ and R^- by Lemma 45 (after negating all z-coordinates).
4. Recursively solve the problem for P^- and R^-.
5. Recursively solve the problem for P^+ and R^+.

The running time for $m = |P| + |R|$ satisfies the recurrence $T_1(m) = 2\,T_1(m/2) + \tilde{O}(m)$, which solves to $T_1(m) = \tilde{O}(m)$. (Note that $|Q|$ is not reduced during recursion, but luckily, the complexity does not depend on $|Q|$, excluding the initial preprocessing.)

Next consider the special case when $P \subset \alpha_P + ((0,\mu_x) \times (0,1)^2)$ and $R \subset \alpha_R + ((\mu_x,1) \times (0,1)^2)$, for a given μ_x. By a similar recursive algorithm via the median y-coordinate, we obtain running time $T_2(m) = 2\,T_2(m/2) + \tilde{O}(T_1(m))$, which solves to $T_2(m) = \tilde{O}(m)$.

Finally, the general case can be solved by another similar recursive algorithm via the median x-coordinate, with running time $T(m) = 2\,T(m/2) + \tilde{O}(T_2(m))$, which solves to $T(m) = \tilde{O}(m)$. ◀

▶ **Theorem 47.** *Given 3 sets P, Q, R of $O(n)$ points in $\mathbb{R}^3$, we can decide whether for all $(p,r) \in P \times R$, there exists $q \in Q$ with $[\![q]\!]$ intersecting both $[\![p]\!]$ and $[\![r]\!]$, in $\tilde{O}(n)$ time.*

Proof. Build a uniform grid of side length 1. For each nonempty grid cell $\alpha_P + (0,1)^3$ (with $\alpha_P \in \mathbb{Z}^3$) and for each $\alpha_R \in \mathbb{Z}^3$ with L_∞-distance at most 2 from α_P, we solve the problem for $P \cap (\alpha_P + (0,1)^3)$, $Q \cap (\alpha_P + (-1,2)^3)$, and $R \cap (\alpha_R + (0,1)^3)$ by Lemma 46 in time near-linear in the sizes of these three subsets. Each point participates in only a constant number of subproblems. ◀

The number of logarithmic factors is admittedly huge (in the 50s), though we have not attempted to optimize it.

H Subquadratic Diameter-3 Algorithm for 3D Unit Cubes (Continued)

H.1 Tools about pseudolines

Before describing our subquadratic algorithm, we digress with some known combinatorial and computational facts about pseudoline arrangements.

▶ **Definition 48.** An *abstract point-pseudoline system* consists of a pair of sets $(\mathcal{P}, \mathcal{L})$. Each element $p \in \mathcal{P}$ is a point in $\mathbb{R}^2$, and each element $\ell \in \mathcal{L}$ is a curve, where the curves form a pseudoline family, i.e., each vertical line intersects each curve once, and each pair of curves intersect at most once. The curves are not explicitly given, but we assume that there are oracles to perform the following operations:

(O1) Given $p \in \mathcal{P}$ and $\ell \in \mathcal{L}$, decide whether p is below ℓ.

(O2) Given $p \in \mathcal{P}$ and $\ell_1, \ell_2 \in \mathcal{L}$, decide whether ℓ_1 is below ℓ_2 at the x-coordinate of p.

▶ **Lemma 49.** *Consider a set system $(X, \mathcal{S})$, where $X = \{p_1, \ldots, p_n\}$ and $\mathcal{S} = \{S_1, \ldots, S_m\}$, satisfying the following condition:*

($*$) *if $p_k \in S_j \setminus S_i$ and $p_h \in S_i \setminus S_j$ and $k < h$, then $i < j$.*

Then we can form an abstract point-pseudoline system $(\mathcal{P}, \mathcal{L})$ and mappings $\phi : X \to \mathcal{P}$ and $\psi : \mathcal{S} \to \mathcal{L}$, such that $p_k \in S_i$ iff the point $\phi(p_k)$ is below the pseudoline $\psi(S_i)$.

Furthermore, operations (O1) and (O2) reduce to operations (O1′) and (O2′) respectively:

(O1′) *Given i and k, decide whether $p_k \in S_i$.*

(O2′) *Given $i < j$, find the smallest index k with $p_k \in S_i \setminus S_j$.*

Proof. The connection between set systems satisfying ($*$) and pseudolines is known before; for example, see a work by Keszegh and Pálvölgyi [27], who called such systems *ABA-free hypergraphs*. For completeness, we sketch a quick proof, so that one can see how operation (O2′) relates to (O2). (A similar construction also appeared in a paper by Agarwal and Sharir [?], in a different context about dualization of pseudoline arrangements.)

We define the points simply by $\phi(p_k) = (k, 0)$. We construct the curves $\psi(S_1), \ldots, \psi(S_m)$ from left to right, while ensuring that $\psi(S_i)$ has positive y-coordinate at $x = k$ iff $p_k \in S_i$. In $(-\infty, 0] \times \mathbb{R}$, the curves $\psi(S_i)$ are non-intersecting and have negative y-coordinates in increasing order of i. For $k = 1, \ldots, n$, we do the following. Consider four groups of curves: $A_k = \{\psi(S_i) : p_{k-1}, p_k \in S_i\}$, $B_k = \{\psi(S_i) : p_{k-1} \in S_i,\ p_k \notin S_i\}$, $C_k = \{\psi(S_i) : p_{k-1} \notin S_i,\ p_k \in S_i\}$, and $D_k = \{\psi(S_i) : p_{k-1}, p_k \notin S_i\}$. Each group is non-intersecting in $[k-1, k] \times \mathbb{R}$. At $x = k$, we make D_k below B_k, below 0, below C_k, below A_k. In $[k-1, k] \times \mathbb{R}$, a curve $\beta \in B_k$ intersects only the curves in A_k that are below β at $x = k-1$, and intersects all the curves in C_k, but intersects none of the curves in D_k; similarly, a curve $\gamma \in C_k$ intersects only the curves in D_k that are above γ at $x = k-1$, and intersects all the curves in B_k, but intersects none of the curves in A_k.

The construction satisfies the following property: if $\psi(S_i)$ is below $\psi(S_j)$ at $x = k-1$, then $\psi(S_i)$ intersects $\psi(S_j)$ in $[k-1, k] \times \mathbb{R}$ iff $p_k \in S_i \setminus S_j$. It follows that if $i < j$, then $\psi(S_i)$ stays below $\psi(S_j)$ at all integer x-coordinates up to the smallest k with $p_k \in S_i \setminus S_j$, then stays above at all larger x-coordinates, because of ($*$). Hence, the constructed curves $\psi(S_i)$ form a pseudoline family. Furthermore, for $i < j$, $\psi(S_i)$ is below $\psi(S_j)$ at x (operation (O2)) iff x is less than the smallest k with $p_k \in S_i \setminus S_j$ (the index found by (O2′)). ◀

▶ **Lemma 50.** *Let $(\mathcal{P}_1, \mathcal{L}_1), \dots, (\mathcal{P}_D, \mathcal{L}_D)$ be abstract point-pseudoline systems for a constant D. Given a set P of size m, a set S of size n, and mappings $\phi_j : P \to \mathcal{P}_j$ and $\psi_j : S \to \mathcal{L}_j$, we can decide whether there exists $(p, s) \in P \times S$ such that $\phi_j(p)$ is above $\psi_j(s)$ for all $j \in \{1, \dots, D\}$, in $O^*(n\sqrt{C_1 C_2 m} + C_2 m)$ expected time, where C_1 and C_2 are the costs of operation (O1) and (O2) respectively, with $C_1 \le C_2$. Furthermore, we can report all such (p, s) pairs in $O(K)$ additional time where K is the output size.*

Proof. This follows from standard "multi-level" range searching techniques via "cuttings" [18, 5] for pseudolines. We re-sketch the approach below, and in particular, work out the dependence on C_1 and C_2 (which will be important in our application):

Take a random sample $R \subset S$ of size $c_0 \rho \log \rho$, and consider the vertical decomposition of the arrangement of the curves in $\psi_D(R)$, which has $O(\rho^2 \log^2 \rho)$ size. Each cell in the decomposition is a pseudo-trapezoid whose left and right sides are vertical and the top and bottom sides are parts of the curves. We shrink the pseudo-trapezoid inward so that the left and right sides pass through points in $\phi_D(P)$. This simplifies primitive operations. (For example, two points in $\phi_D(P)$ are in the same pseudo-trapezoid iff they have the same set of curves below it, as well as the same curve immediately below it and the same curve immediately above it. Thus, we can assign points in $\phi_D(P)$ to pseudo-trapezoids using $\tilde{O}(\rho m)$ number of operations (O1) and (O2). Furthermore, we can test whether a curve in $\psi_D(R)$ intersects a pseudo-trapezoid by making $O(1)$ number of operations (O1) and (O2).)

Standard Clarkson–Shor analysis [20] tells us that with good probability, every pseudo-trapezoid intersects at most $|S|/\rho$ curves, assuming that c_0 is a sufficiently large constant (we can re-sample when the condition is not met, for an expected $O(1)$ number of trials). By additional vertical cuts, we can ensure that every pseudo-trapezoid contains at most $|P|/\rho^2$ points, while increasing the number of pseudo-trapezoids by $O(\rho^2)$. For each pseudo-trapezoid τ, we recurse for the subset of all points inside τ and the subset of all curves intersecting τ; we also recurse for the subset of all points inside τ and the subset of all curves completely below τ, but with D decremented. This yields the following recurrence for the expected running time:

$$T_D(m, n) \le O(\rho^2 \log^2 \rho) T_D(m/\rho^2, n/\rho) + \tilde{O}(\rho^2 T_{D-1}(m, n) + \rho^2 C_2(m + n)).$$

For the base case, when $m \le O(C_2/C_1)$, we switch to the naive bound $T_D(m, n) \le O(C_1 mn) \le O(C_2 n)$. Setting ρ to be an arbitrarily large constant, the recurrence solves to $T_D(n, n) = O^*(n\sqrt{C_1 C_2 m} + C_2 m)$ by induction on D.

(Derandomization is possible via known techniques. One could also get a better bound of the form $m^{2/3} n^{2/3} + m + n$, ignoring dependence on C_1 and C_2, if a dual pseudoline arrangement is available, but we will not need such an improvement.) ◀

H.2 Algorithm

As in Section 4.1, we begin with the corresponding problem for generalized dominance relations. As in Lemma 17, we first consider the most crucial special case when P and Q are x-separated, Q and R are y-separated, and R and S are z-separated:

▶ **Lemma 51.** *Given point sets P, Q, R, S in $\mathbb{R}^3$ of total size n, where P and Q are x-separated, Q and R are y-separated, and R and S are z-separated, and given generalized dominance relations $\lhd_1, \lhd_2, \lhd_3$, we can form an abstract point-pseudoline system $(\mathcal{P}, \mathcal{L})$ mappings $\phi : P \to \mathcal{P}$ and $\psi : S \to \mathcal{L}$, satisfying the following property for every $(p, s) \in P \times S$:*

$$(\exists (q, r) \in Q \times R : p \lhd_1 q \text{ and } q \lhd_2 r \text{ and } r \lhd_3 s) \iff \phi(p) \text{ is below } \psi(s).$$

After preprocessing P, Q, R, S in $\tilde{O}(n|P|^{9/10})$ expected time, operation (O1) takes $\tilde{O}(1)$ time and operation (O2) takes $\tilde{O}(|P|^{4/5})$ time.

Proof. The proof of Lemma 17 actually shows the property (∗) from Lemma 49, if we order P by increasing y-coordinate and order S by increasing/decreasing y-coordinate. So, by Lemma 49, we obtain the point-pseudoline system $(\mathcal{P}, \mathcal{L})$ and mappings ϕ, ψ with the desired property.

To implement the oracles, we apply the diameter computation technique by Chan et al. [15], knowing that the VC-dimension of the neighborhood sets is 2. For each $q \in Q$, let $N^1[q] = \{p \in P : p \lhd_1 q\}$. For each $r \in R$, let $N^2[r] = \{p \in P : p \lhd_1 q,\ q \lhd_2 r\}$. For each $s \in S$, let $N^3[s] = \{p \in P : p \lhd_1 q,\ q \lhd_2 r,\ r \lhd_3 s\}$. The algorithm computes the interval representations of the $N_1[q]$ sets, and then the interval representations of the $N^2[r]$ sets from all the $N^1[q]$ sets, and finally the interval representations of all the $N^3[s]$ sets from the all $N^2[r]$ sets. The adaptation of Chan et al.'s technique here is straightforward:

- The algorithm needs to work with intermediate set systems that have the VC-dimension doubled to 4, because our VC-dimension bound is for the neighborhood sets of a fixed radius, not all radii simultaneously.
- We use a version of the algorithm via rainbow colored intersection searching: for 3D dominance, it is straightforward to obtain such a data structure with near-linear preprocessing time and polylogarithmic query time, since union of orthants has linear complexity (this is similar to the rainbow intersection searching data structure for 2D squares from [15, Appendix C.2]).
- Plugging into the analysis from [15], we get the following expected time bound when the diameter is constant:

  $$\tilde{O}(n\rho + n(|P|/\rho + \rho^4)b + n|P|/b),$$

  for parameters b and ρ. Setting $\rho = |P|^{1/5}$ and $b = |P|^{1/10}$ yields $\tilde{O}(n|P|^{9/10})$.

At the end, we have obtained interval representations of all the $N^3[s]$ sets. The total size of the interval representations is $\tilde{O}(|S| \cdot (|P|/\rho + \rho^4)) = \tilde{O}(|S| \cdot |P|^{4/5})$. (With more steps, we could lower it to $\tilde{O}(|S| \cdot \sqrt{|P|})$ since the actual VC-dimension of the 3-neighborhoods is 2, but this will not be important.) Operation (O1′) is easy to support in $\tilde{O}(1)$ time: we just store the intervals of each $N^3[s]$ set in a binary search tree during preprocessing.

Operation (O2′) requires more effort. First, we store all the y-values of the elements along the stabbing path in a binary search tree for 1D range min/max queries. For each $N^3[s]$ set, we precompute the min/max y-value of the points in each interval of $N^3[s]$, and stores these values in another binary search tree; we do the same of the complement of $N^3[s]$. Note that the interval representation of $N^3[s]$ has size $\tilde{O}(|P|^{4/5})$ for at least half of the elements $s \in S$—call these elements *good*. We recursively solve the problem for the bad (i.e., non-good) elements of S (keeping P, Q, R the same). This increases running time by a logarithmic factor. (O2′) asks for the min/max y-value of elements in $N^3[s] \setminus N^3[s']$ for two given $s, s' \in S$. If both s and s' are bad, we recurse. Suppose one of s and s' is good—say it is s. For each interval I in $N^3[s]$, we find the min/max y-values for the intervals in the complement of $N^3[s]$ that are completely inside I; this takes $\tilde{O}(1)$ time. There may still be two remaining subintervals in $I \setminus N^3[s]$, and we can find the min/max y-values there in $\tilde{O}(1)$ time. The overall query time is $\tilde{O}(1)$ times the number of intervals, which is $\tilde{O}(|P|^{4/5})$ since s is good. The case when s' is good is similar. ◀

Next, we consider the (easier) cases from Lemma 18:

▶ **Lemma 52.** *Given point sets P, Q, R, S in $\mathbb{R}^3$ of total size n, where P and Q are both x- and y-separated, or Q and R are both x- and y-separated, or R and S are both x- and y-separated, and dominance relations $\lhd_1, \lhd_2, \lhd_3$, we can compute mappings $\phi : P \to \mathbb{R}$ and $\psi : S \to \mathbb{R}$ in $\tilde{O}(n)$ time, satisfying the following property for every $(p, s) \in P \times S$:*

$$(\exists (q, r) \in Q \times R : p \lhd_1 q \text{ and } q \lhd_2 r \text{ and } r \lhd_3 s) \iff \phi(p) < \psi(s).$$

Proof. This is similar to the proof of Lemma 42.

Consider the case when P and Q are both x- and y-separated. W.l.o.g., say $P \in (-\infty, 0)^2 \times \mathbb{R}$ and $Q \in (0, \infty)^2 \times \mathbb{R}$. We may assume that $\lhd_{1x} \in \{\prec_x, \mathbf{T}\}$ and $\lhd_{1y} \in \{\prec_y, \mathbf{T}\}$ (because if not, we can trivially set $\phi = 1$ and $\psi = 0$). We may assume that $\lhd_{1z} \neq \mathbf{T}$ (because if not, we can replace all z-coordinates of P with a sufficiently small negative number and replace $\lhd_{1z}$ with $\prec_z$). Suppose $\lhd_{1z} = \prec_z$. For each $p \in P$, we define $\phi(p)$ to be the z-coordinate of s. Next, we define the weight of each point $r \in R$ to be the largest z-coordinate among all points $q \in Q$ with $q \lhd_1 r$; these weights can be computed by $O(n)$ orthogonal range max queries; finally, for each $s \in S$, we define $\psi(s)$ to be the largest weight among all points $r \in R$ with $r \lhd_1 s$; again these values can be computed by orthogonal range max queries. The property is then satisfied. The case when $\lhd_{1z} = \succ_z$ is similar (by negating all z-coordinates).

The case when R and S are both x- and y-separated is similar.

Finally consider the case when Q and R are both x- and y-separated. W.l.o.g., say $Q \in (-\infty, 0)^2 \times \mathbb{R}$ and $R \in (0, \infty)^2 \times \mathbb{R}$. As before, we may assume that $\lhd_{2x} \in \{\prec_x, \mathbf{T}\}$ and $\lhd_{2y} \in \{\prec_y, \mathbf{T}\}$. We may assume that $\lhd_{2z} \neq \mathbf{T}$ (because if not, we can shift all the z-coordinates of R and S upward by a large number and replace $\lhd_{2z}$ with $\prec_z$). Suppose $\lhd_{2z} = \prec_z$. For each $p \in P$, we define $\phi(p)$ to be the smallest z-coordinates among all points $q \in Q$ with $p \lhd_1 q$; these values can be computed by orthogonal range min queries. For each $s \in S$, we define $\psi(s)$ to be the largest z-coordinates among all points $r \in R$ with $r \lhd_2 s$; again, these values can be computed by orthogonal range max queries. The property is then satisfied. The case when $\lhd_{2z} = \succ_z$ is similar (by negating all z-coordinates). ◀

Following Lemma 19, the natural next step would be to consider the case when P and S are separated along all 3 axes. However, we will actually need to handle a slightly more general case. For a box $\gamma = (\mu_x^-, \mu_x^+) \times (\mu_y^-, \mu_y^+) \times (\mu_z^-, \mu_z^+)$, define

$$\text{shadow}(\gamma) := ((\mu_x^-, \mu_x^+) \times \mathbb{R} \times \mathbb{R}) \cup (\mathbb{R} \times (\mu_y^-, \mu_y^+) \times \mathbb{R}) \cup (\mathbb{R} \times \mathbb{R} \times (\mu_z^-, \mu_z^+)).$$

The case we will consider is when $P \subset \gamma$ and not all of q, r, s are in shadow(γ):

▶ **Lemma 53.** *Given point sets P, Q, R, S in $\mathbb{R}^3$ of total size n, where $P \subset \gamma$ for some box γ, we can form $O(1)$ abstract point-pseudoline systems $(\mathcal{P}_j, \mathcal{L}_j)$ and mappings $\phi_j : P \to \mathcal{P}_j$ and $\psi_j : S \to \mathcal{L}_j$, satisfying the following property for every $(p, s) \in P \times S$:*

$$(\exists (q, r) \in Q \times R : p \lhd_1 q \text{ and } q \lhd_2 r \text{ and } r \lhd_3 s, \text{ and not all of } q, r, s \text{ are in shadow}(\gamma)$$
$$\iff \phi_j(p) \text{ is below } \psi_j(s) \text{ for some } j.$$

After preprocessing P, Q, R, S in $\tilde{O}(n|P|^{9/10})$ expected time, operation (O1) takes $\tilde{O}(1)$ time and operation (O2) takes $\tilde{O}(|P|^{4/5})$ time for each system.

Proof. The planes through the 6 faces of γ divide $\mathbb{R}^3$ into 9 box cells. Consider a triple $\tau = (\tau_Q, \tau_R, \tau_S)$ of cells such that not all of τ_Q, τ_R, τ_S are in shadow(γ); there are a (large) constant number of such triples. Then one of the pairs (τ_P, τ_Q), (τ_Q, τ_R), or (τ_Q, τ_S) must be x-separated, and one of them must be y-separated, and one of them must be z-separated. We form a point-pseudoline system $(\mathcal{P}_\tau, \mathcal{L}_\tau)$ and mappings $\phi_\tau : P \to \mathcal{P}_\tau$ and $\psi_\tau : (S \cap \tau_S) \to \mathcal{L}_\tau$, satisfying the following property for every $(p, s) \in P \times (S \cap \tau_S)$:

$(\exists (q, r) \in (Q \cap \tau_Q) \times (R \cap \tau_R) : p \lhd_1 q$ and $q \lhd_2 r$ and $r \lhd_3 s) \iff \phi_\tau(p)$ is below $\psi_\tau(s)$.

Such a system and mappings for each τ can be obtained from either Lemma 51 or Lemma 52 (in the latter case, numbers in $\mathbb{R}$ can be viewed as pseudolines in $\mathbb{R}^2$ trivially). ◀

We now transform the result from dominance to unit cubes:

▶ **Lemma 54.** *Given point sets P, Q, R, S in $\mathbb{R}^3$ of total size n, where $P \subset \gamma$ for some box $\gamma \subset (0, 1)^3$, we can form $O(1)$ abstract point-pseudoline systems $(\mathcal{P}_j, \mathcal{L}_j)$ and mappings $\phi_j : P \to \mathcal{P}_j$ and $\psi_j : S \to \mathcal{L}_j$, satisfying the following property for every $(p, s) \in P \times S$:*

$(\exists (q, r) \in Q \times R : [\![p]\!]$ intersects $[\![q]\!]$, $[\![q]\!]$ intersects $[\![r]\!]$, $[\![r]\!]$ intersects $[\![s]\!]$, and not all of q, r, s are in shadow(γ) modulo 1$)$
$\iff \phi_j(p)$ is below $\psi_j(s)$ for some j.

After preprocessing P, Q, R, S in $\tilde{O}(n|P|^{9/10})$ expected time, operation (O1) takes $\tilde{O}(1)$ time and operation (O2) takes $\tilde{O}(|P|^{4/5})$ time for each system.

Proof. Consider a triple $\alpha = (\alpha_Q, \alpha_R, \alpha_S) \in (\mathbb{Z}^3)^3$ with $\|\alpha_Q\|_\infty, \|\alpha_P - \alpha_Q\|_\infty, \|\alpha_R - \alpha_S\|_\infty \le 1$; there are a (large) constant number of such triples. We form point-pseudoline systems $(\mathcal{P}_j^{(\alpha)}, \mathcal{L}_j^{(\alpha)})$ and mappings $\phi_j^{(\alpha)} : P \to \mathcal{P}_j^{(\alpha)}$ and $\psi_j^{(\alpha)} : (S \cap (\alpha_S + (0, 1)^3)) \to \mathcal{L}_j^{(\alpha)}$, satisfying the following property for every $(p, s) \in P \times (S \cap (\alpha_S + (0, 1)^3))$:

$(\exists (q, r) \in (Q \cap (\alpha_Q + (0, 1)^3)) \times (R \cap (\alpha_R + (0, 1)^3)) : [\![p]\!]$ intersects $[\![q]\!]$, $[\![q]\!]$ intersects $[\![r]\!]$, $[\![r]\!]$ intersects $[\![s]\!]$, and not all q, r, s are in shadow(γ) modulo 1$)$
$\iff \phi_j^{(\alpha)}(p)$ is below $\psi_j^{(\alpha)}(s)$ for some j.

Such systems and mappings for each α can be obtained from Lemma 53 by defining appropriate generalized dominance relations as in the proof of Lemma 15. ◀

We then solve the problem by multi-level range searching:

▶ **Lemma 55.** *Given point sets P, Q, R, S in $\mathbb{R}^3$ of total size n, where $P \subset \gamma$ for some box $\gamma \subset (0, 1)^3$, we can decide whether for all $(p, s) \in P \times S$, there exists $(q, r) \in Q \times R$ such that $[\![p]\!]$ intersects $[\![q]\!]$, $[\![q]\!]$ intersects $[\![r]\!]$, $[\![r]\!]$ intersects $[\![s]\!]$, and not all of q, r, s are in shadow(γ) modulo 1, in $O^*(n|P|^{9/10})$ expected time. Furthermore, we can report all pairs (p, s) not satisfying the property in $O(K)$ additional time, where K is the output size.*

Proof. The problem can be solved by combining Lemma 54 with Lemma 50, where $C_1 = \tilde{O}(1)$ and $C_2 = \tilde{O}(|P|^{4/5})$, in $O^*(n|P|^{9/10} + n\sqrt{|P|^{4/5}|P|}) = O^*(n|P|^{9/10})$ expected time. ◀

We remark that the way we use range searching above is rather unusual (and interesting), as we are dealing with ranges/pseudolines that are not explicitly generated but implicitly represented via oracles to (O1) and (O2). (For one prior example, an algorithm by Chan on selection in totally monotone matrices [13] similarly dealt with "abstract" pseudolines, but there oracle costs are $\tilde{O}(1)$. The closest example is perhaps Agarwal and Sharir's usage of abstract *pseudo-disks* to solve the 2D discrete 2-center problem.) In the recursive algorithm in the proof of Lemma 50, even though the number of objects corresponding to P and S may decrease, the oracle costs C_1 and C_2 stay fixed, as we cannot afford to re-preprocess (besides, Q and R never change). Even if there might be room for improvement in the range searching part, the bottleneck for the 9/10 exponent above lies in the simulation of the diameter algorithm of Chan *et al.* [15] in Lemma 51.

Finally, we reduce the general case to the case in Lemma 55. It is tempting to try a range-tree-style divide-and-conquer approach similar to the proof of Lemma 46, but it does not seem to work, primarily because the size of Q and R do not necessarily decrease during recursion. Instead, we switch to a grid approach (not exactly the same as that in Lemma 20); this worsens the exponent, but not by much (from $2 - 1/10$ to $2 - 1/13$).

▶ **Lemma 56.** *Given point sets P, Q, R, S in $\mathbb{R}^3$ of total size n, where $P \subset (0,1)^3$, we can decide whether for all $(p,s) \in P \times S$, there exists $(q,r) \in Q \times R$ such that $[\![p]\!]$ intersects $[\![q]\!]$, $[\![q]\!]$ intersects $[\![r]\!]$, and $[\![r]\!]$ intersects $[\![s]\!]$, in $O^*(n^{2-1/13})$ expected time.*

Proof. Build a $g \times g \times g$ (nonuniform) grid over $(0,1)^3$, where every two consecutive parallel grid planes contain $O(n/g)$ points in $P \cup Q \cup R \cup S$ modulo 1. For each grid cell γ, let $P_\gamma = P \cap \gamma$ and do the following:

1. Run the algorithm in Lemma 55 to report the list L_γ of all pairs $(p,s) \in P_\gamma \times S$ violating the following property: there exists $(q,r) \in Q \times R$ such that $[\![p]\!]$ intersects $[\![q]\!]$, $[\![q]\!]$ intersects $[\![r]\!]$, $[\![r]\!]$ intersects $[\![s]\!]$, and not all of q, r, s are in shadow(γ) modulo 1. This takes $O^*(n|P_\gamma|^{9/10} + |L_\gamma|)$ time.
2. Next, find the list L'_γ of all pairs $(p,s) \in P_\gamma \times S$ satisfying the following property: there exists $(q,r) \in Q \times R$ such that $[\![p]\!]$ intersects $[\![q]\!]$, $[\![q]\!]$ intersects $[\![r]\!]$, $[\![r]\!]$ intersects $[\![s]\!]$, and q, r, s are all in shadow(γ) modulo 1. Since shadow(γ) contains only $O(n/g)$ points, we can solve the problem naively, by computing the 1-neighborhoods of each point $p \in P$, then the 2-neighborhoods, and finally the 3-neighborhoods, via orthogonal range searching, in $\tilde{O}(|P_\gamma| \cdot n/g)$ total time.
3. Verify that $L_\gamma \subseteq L'_\gamma$, in $\tilde{O}(|L'_\gamma|) \leq \tilde{O}(|P_\gamma| \cdot n/g)$ time.

Note that since $|L'_\gamma| \leq O(|P_\gamma| \cdot n/g)$, as soon as the number of elements in L_γ we have found during step 1 exceeds a constant times $|P_\gamma| \cdot n/g$, we can stop and return false.

The total expected running time over all $O(g^3)$ grid cells γ is $O^*(\sum_\gamma (n|P_\gamma|^{9/10} + |P_\gamma| \cdot n/g)) = O^*(n^{19/10}(g^3)^{1/10} + n^2/g)$, which is $O^*(n^{2-1/13})$ by setting $g = n^{1/13}$. ◀

▶ **Theorem 57.** *Given point sets P, Q, R, S in $\mathbb{R}^3$ of total size n, we can decide whether for all $(p,s) \in P \times S$, there exists $(q,r) \in Q \times R$ such that $[\![p]\!]$ intersects $[\![q]\!]$, $[\![q]\!]$ intersects $[\![r]\!]$, and $[\![r]\!]$ intersects $[\![s]\!]$, in $O^*(n^{2-1/13})$ expected time.*

Proof. Build a uniform grid of side length 1. For each nonempty grid cell $\alpha_P + (0,1)^3$ (with $\alpha_P \in \mathbb{Z}^3$), we solve the problem for $P \cap (\alpha_P + (0,1)^3)$, $Q \cap (\alpha_P + (-1,2)^3)$, and $R \cap (\alpha_P + (-2,3)^3)$, and $S \cap (\alpha_P + (-3,4)^3)$, by Lemma 56. Each point participates in only a constant number of subproblems. ◀

▶ Remark 58. We do not know how to generalize our subquadratic algorithm for 3D unit cubes to diameter 4 and larger constants—this is perhaps one of the most intriguing questions we leave open. However, for the related problem of designing efficient distance oracles, one could obtain nontrivial results for distances up to 6: namely, there is a data structure with subquadratic preprocessing time and space, which can answer distance queries in sublinear (albeit, around $O(n^{0.999996})$) time. This follows from our VC-dimension bound for 3-neighborhoods and the interval representation techniques from [15], since two vertices have distance at most 6 iff their 3-neighborhoods intersect.

I Subquadratic Diameter-2 Algorithm for 3D Boxes

In this section, we present a subquadratic algorithm for testing whether the diameter of an intersection graph for 3D boxes is at most 2. This complements our lower bound result (Theorem 29), showing conditionally that there are no similar diameter-3 algorithms (this lower bound holds even in the special case of 3D cubes). The running time of our algorithm is $\tilde{O}(n^{11/6})$. Afterwards, we mention improvements in the special cases for 2D rectangles and 3D cubes (even in these special cases, subquadratic algorithms were not known before).

As before, we will solve the problem in a slightly more general setting for 3 sets P, Q, R of objects, testing whether all distances between P and R are exactly 2 in the tripartite intersection graph.

In past sections, we started by bounding the VC-dimension of the corresponding set system, as warm-up to the design of a subquadratic algorithm. However, for 3D boxes (or even 2D rectangles or 3D cubes), the VC-dimension is unbounded for distance-2 neighborhoods—thus, it is somewhat surprising that subquadratic algorithms are possible for objects as general as 3D boxes.

I.1 Algorithm

We use an approach based on a $g \times g \times g$ nonuniform grid for some choice of parameter g; this type of approach has been used before to obtain subquadratic algorithms for other problems (e.g., L_∞ discrete 3-center problem [16], or finding small-size independent set among rectangles or boxes [14]). For our problem, we face technical challenges due to the fact that 3D boxes may have *quadratic union complexity*. To resolve this issue, we show how to divide into a constant number of cases, so that in some cases, the boxes in the set R can be replaced by orthants, and in the remaining cases, the boxes in the set P can be replaced by orthants (exploiting the symmetry of the diameter problem). Orthants in 3D are known to have linear union complexity.

For a box $q \subset \mathbb{R}^3$, let $x^-(q)$, $y^-(q)$, and $z^-(q)$ denote its min x-, y-, and z-coordinate respectively, and $x^+(q)$, $y^+(q)$, and $z^+(q)$ denote its max x-, y-, and z-coordinate respectively. For simplicity, we assume that all coordinate values are distinct.

Map each box $q \subset \mathbb{R}^3$ to a point $\phi(q) = (x^-(q), -x^+(q), y^-(q), -y^+(q), z^-(q), -z^+(q)) \in \mathbb{R}^6$. Map each box $p \subset \mathbb{R}^3$ to another point $\psi(p) = (x^+(p), -x^-(p), y^+(p), -y^-(p), z^+(p), -z^-(p)) \in \mathbb{R}^6$. Observe that p and q intersect iff $\phi(q) \prec \psi(p)$, where $\prec$ denotes dominance.

For each $I \subset \{1, \ldots, 6\}$, let $\pi_I : \mathbb{R}^6 \to \mathbb{R}^{|I|}$ denote the projection map where we keep only the coordinate positions in I.

▶ **Lemma 59.** *Let $I \subset \{1, \ldots, 6\}$ be of size 3. Given 3 sets P, Q, R of $O(n)$ boxes in $\mathbb{R}^3$, we can decide whether for all $(p, r) \in P \times R$ with $\pi_I(\psi(p)) \prec \pi_I(\psi(r))$, there exists $q \in Q$ intersecting both p and r, in $\tilde{O}(n^{11/6})$ time.*

Proof. Form an $g \times g \times g$ (nonuniform) grid over $\mathbb{R}^3$, where there are $O(n/g)$ box vertices between any two consecutive parallel grid planes, for a parameter g to be set later. A *grid box* refers to a box whose sides lie on grid planes; there are $O(g^6)$ possible grid boxes. For each box p, let $\hat{p}$ be the largest grid box contained in p.

For each grid box $\hat{p}$, precompute a set

$$R(\hat{p}) = \{r \in R : \exists q \in Q \text{ intersecting both } \hat{p} \text{ and } r\}.$$

We can do so by first computing $Q(\hat{p}) = \{q \in Q : q \text{ intersects } \hat{p}\}$ naively in $O(n)$ time per $\hat{p}$, and then computing $R(\hat{p}) = \{r \in R : r \text{ intersects some } q \in Q(\hat{p})\}$ by performing $O(n)$

orthogonal range intersection queries on $Q(\hat{p})$, in $\tilde{O}(n)$ total time per $\hat{p}$. The total time so far is $\tilde{O}(g^6 n)$.

Fix $p \in P$. We want to check the following condition:

$$\forall r \in R \setminus R(\hat{p}) \text{ with } \pi_I(\psi(p)) \prec \pi_I(\psi(r)):\ \exists q \in Q \text{ intersecting both } p \text{ and } r.$$

Let $L(p) := \{q \in Q : q \text{ intersects } p \text{ but does not intersect } \hat{p}\} = \{q \in Q : \partial q \text{ intersects } p \setminus \hat{p}\}$, which has cardinality at most $O(n/g)$. The condition is then equivalent to:

$$\forall r \in R \setminus R(\hat{p}) \text{ with } \pi_I(\psi(p)) \prec \pi_I(\psi(r)):\ \exists q \in L(p) \text{ intersecting } r.$$

If we already know q intersects p, then q intersects $r \iff \phi(q) \prec \psi(r) \iff \pi_{I^c}(\phi(q)) \prec \pi_{I^c}(\psi(r))$, where $I^c = \{1,\ldots,6\} \setminus I$ (which has size 3). This is because $\pi_I(\phi(q)) \prec \pi_I(\psi(p)) \prec \pi_I(\psi(r))$. Thus, the above condition can be rewritten as:

$$\forall r \in R \setminus R(\hat{p}) \text{ with } \pi_I(\psi(p)) \prec \pi_I(\psi(r)):\ \pi_{I^c}(\psi(r)) \in \mathcal{U}(p), \text{ where}$$
$$\mathcal{U}(p) := \bigcup_{q \in L(p)} \{\xi \in \mathbb{R}^3 : \pi_{I^c}(\psi(q)) \prec \xi\}.$$

Now, $\mathcal{U}(p)$ is a union of $O(n/g)$ orthants in $\mathbb{R}^3$, forming a "staircase" polyhedron of $O(n/g)$ complexity. We can compute $\mathcal{U}(p)$ and decompose the complement of $\mathcal{U}(p)$ into $O(n/g)$ box cells, in $\tilde{O}(n/g)$ time [6]. To check the above condition, it suffices to examine each such cell γ and test whether there exists $r \in R \setminus R(\hat{p})$ with $\pi_I(\psi(p)) \prec \pi_I(\psi(r))$ and $\pi_{I^c}(\psi(r)) \in \gamma$. This can be done by performing $O(n/g)$ constant-dimensional orthogonal range queries, in $\tilde{O}(n/g)$ time, after preprocessing $R - R(\hat{p})$ in $\tilde{O}(n)$ time. The total preprocessing time for the range queries is $\tilde{O}(g^6 n)$, and the total query time over all $s \in S$ is $\tilde{O}(n \cdot n/g)$. Setting $g = n^{1/7}$ gives an $\tilde{O}(n^{13/7})$ time bound.

To improve the bound, we use *bottomless grid boxes*, i.e., grid boxes that are unbounded from below in the z-direction. There are only $O(g^5)$ bottomless grid boxes. For each box p, define $\hat{p}_\downarrow$ to be the bottomless grid box formed by extending $\hat{p}$ downward. For each bottomless grid box $\hat{p}_\downarrow$, compute a weighted point set $Q(\hat{p}_\downarrow) = \{q \in Q : q \text{ intersects } \hat{p}_\downarrow\}$ naively in $O(n)$ time, where the weight of q is the largest z such that q intersects $\hat{p}_\downarrow \cap (z, \infty)$. Compute a weighted point set $R(\hat{p}_\downarrow) = \{r \in R : r \text{ intersects some } q \in Q(\hat{p}_\downarrow)\}$, where the weight of r is the largest weight of all $q \in Q(\hat{p}_\downarrow)$ intersecting r; this can be done by performing orthogonal range max queries, in $\tilde{O}(n)$ time per $\hat{p}_\downarrow$. The time for this step is $\tilde{O}(g^5 n)$. These $O(g^5)$ weighted point sets provide an implicit representation of $R(\hat{p})$ for all $O(g^6)$ grid boxes $\hat{p}$, since $R(\hat{p})$ is just the subset of all points of $R(\hat{p}_\downarrow)$ with weight greater than $z^-(\hat{p})$. After preprocessing $R \setminus R(\hat{p}_\downarrow)$, we can proceed as before, testing each $p \in P$ using $\tilde{O}(n/g)$ range queries—these are now range min/max queries. The overall time bound is now $\tilde{O}(g^5 n + n \cdot n/g)$, which is $\tilde{O}(n^{11/6})$ by setting $g = n^{1/6}$. ◀

▶ **Theorem 60.** *Given 3 sets P, Q, R of $O(n)$ boxes in $\mathbb{R}^3$, we can decide whether for all $(p, r) \in P \times R$, there exists $q \in Q$ intersecting both p and r, in $\tilde{O}(n^{11/6})$ time.*

Proof. For any two boxes p and r, we must have (i) $\pi_I(\psi(p)) \prec \pi_I(\psi(r))$ for some $I \subset \{1,\ldots,6\}$ of size 3, or (ii) $\pi_I(\psi(r)) \prec \pi_I(\psi(p))$ for some $I \subset \{1,\ldots,6\}$ of size 3. (This is because if $\psi(p)$ is less than $\psi(r)$ in fewer than 3 coordinate positions, then $\psi(r)$ is less than $\psi(p)$ in more than 3 coordinate positions.) Thus, it suffices to run the algorithm in Lemma 59 for each I of size 3, for P and R, and also for P and R swapped. ◀

I.2 Special cases

▶ **Theorem 61.** *Given 3 sets P, Q, R of $O(n)$ rectangles in $\mathbb{R}^2$, we can decide whether for all $(p, r) \in P \times R$, there exists $q \in Q$ intersecting both p and r, in $\tilde{O}(n^{7/4})$ time.*

Proof. The algorithm is similar (and simpler), except that with a 2-dimensional grid, the number of bottomless grid rectangles is $O(g^3)$. The overall time bound is now $\tilde{O}(g^3 n + n \cdot n/g)$, which is $\tilde{O}(n^{7/4})$ by setting $g = n^{1/4}$. ◀

▶ **Theorem 62.** *Given 3 sets P, Q, R of $O(n)$ cubes in $\mathbb{R}^3$, we can decide whether for all $(p, r) \in P \times R$, there exists $q \in Q$ intersecting both p and r, in $\tilde{O}(n^{9/5})$ time.*

Proof. The algorithm is the same (but without the improvement via bottomless grid boxes $\hat{p}_\downarrow$), We observe that for boxes p that are cubes, the number of possible choices of grid boxes $\hat{p}$ is actually $O(g^4)$ instead of $O(g^6)$. To see this, map the boxes p to points $\phi(p)$ in $\mathbb{R}^6$. Boxes p with a common $\hat{p}$ map to points in a common grid cell in $\mathbb{R}^6$, for a (non-uniform) grid defined by $O(g)$ (axis-parallel) grid hyperplanes. Boxes p that are cubes map to points lying on a 4-dimensional flat $h = \{(x, -(x + w), y, -(y + w), z, -(z + w)) \in \mathbb{R}^6 : x, y, z, w \in \mathbb{R}\}$. The $O(g)$ grid hyperplanes form a 4-dimensional arrangement of complexity $O(g^4)$ inside h. Thus, the number of grid cells intersecting h is only $O(g^4)$. The overall time bound is now $\tilde{O}(g^4 n + n \cdot n/g)$, which is $\tilde{O}(n^{9/5})$ by setting $g = n^{1/5}$. ◀