

Truly Subquadratic Time Algorithms for Diameter and Related Problems in Graphs of Bounded VC-dimension

Timothy M. Chan* Hsien-Chih Chang† Jie Gao‡ Sándor Kisfaludi-Bak§
Hung Le¶ Da Wei Zheng||

January 4, 2026

Abstract

We give the first truly subquadratic time algorithm, with $\tilde{O}(n^{2-1/18})$ running time, for computing the diameter of an n -vertex unit-disk graph, resolving a central open problem in the literature. Our result is obtained as an instance of a general framework, applicable to different graph families and distance problems. Surprisingly, our framework completely bypasses sublinear separators (or r -divisions) which were used in all previous algorithms. Instead, we use *low-diameter decompositions* in their most elementary form. We also exploit *bounded VC-dimension* of set systems associated with the input graph, as well as new ideas on *geometric data structures*. Among the numerous applications of the general framework, we obtain:

1. An $\tilde{O}(mn^{1-1/(2d)})$ time algorithm for computing the diameter of m -edge sparse unweighted graphs with constant VC-dimension d . The previously known algorithms by Ducoffe, Habib, and Viennot [SODA 2019] and Duraj, Konieczny, and Potępa [ESA 2024] are truly subquadratic only when the diameter is a small polynomial. Our result thus generalizes truly subquadratic time algorithms known for planar and minor-free graphs (in fact, it slightly improves the previous time bound for minor-free graphs).
2. An $\tilde{O}(n^{2-1/12})$ time algorithm for computing the diameter of intersection graphs of axis-aligned squares with arbitrary size. The best-known algorithm by Duraj, Konieczny, and Potępa [ESA 2024] only works for unit squares and is only truly subquadratic in the low-diameter regime.
3. The first algorithms with truly subquadratic complexity for other distance-related problems, including all-vertex eccentricities, Wiener index, and exact distance oracles. In particular, we obtain the first exact distance oracle with truly subquadratic space and $\tilde{O}(1)$ query time for any sparse graph with bounded VC-dimension, again generalizing previous results for planar and minor-free graphs.

*Siebel School of Computing and Data Science, University of Illinois at Urbana-Champaign. Email: tmc@illinois.edu. Supported by NSF grant CCF-2224271.

†Department of Computer Science, Dartmouth College. Email: hsien-chih.chang@dartmouth.edu.

‡Department of Computer Science, Rutgers University. Email: jg1555@rutgers.edu. Gao would like to acknowledge NSF support through CNS-2515159, IIS-2229876, DMS-2220271, DMS-2311064, CCF-2208663, CCF-2118953.

§Department of Computer Science, Aalto University, Finland. Email: sandorkisfaludi-bak@aalto.fi. Supported by the Research Council of Finland, Grant 363444.

¶Manning CICS, UMass Amherst. Email: hungle@cs.umass.edu. Supported by NSF grants CCF-2517033 and CCF-2121952, NSF CAREER Award CCF-2237288, and a Google Faculty Research Award.

||Institute of Science and Technology Austria, Klosterneuburg, Austria. Work in this paper was done while at the Siebel School of Computing and Data Science, University of Illinois at Urbana-Champaign.

29 **Contents**

30	1 Introduction	1
31	1.1 Main Results on Diameter	2
32	1.2 Technical Overview	2
33	1.3 Other Distance-related Problems	7
34	2 Preliminaries	9
35	2.1 Graphs and Low-diameter Decomposition	9
36	2.2 VC-dimension	10
37	2.3 Stabbing Path and Interval Representation	12
38	2.4 Geometric Data Structures	12
39	2.5 Handling of Small Pieces	13
40	3 Framework for Diameter and Eccentricities	14
41	4 Diameter/Eccentricities in Sparse Graphs of Bounded VC-dimension	16
42	5 Diameter/Eccentricities in Square Graphs	17
43	6 Diameter/Eccentricities in Unit-square Graphs	20
44	7 Diameter/Eccentricities in Unit-disk Graphs	20
45	7.1 Restriction to Fixed Types	21
46	7.2 Implementation of the Neighborhood Growing Step	22
47	7.3 Analysis for Eccentricities	26
48	7.4 Analysis for Diameter	27
49	8 Framework for Distance Oracles (and Wiener Index)	28
50	9 Distance Oracles for Sparse Graphs of Bounded VC-dimension	30
51	10 Distance Oracles for Square Graphs	31
52	11 Distance Oracles for Unit-square Graphs	32
53	12 Distance Oracles for Unit-disk Graphs	33
54	13 Conclusion and Open Questions	34
55	A Low-diameter Decompositions	39
56	A.1 Sparse Graphs	40
57	A.2 Geometric Intersection Graphs	40
58	B VC-dimension Lemma	42
59	B.1 Type-1 M -Walks	43
60	B.2 Type-2 M -Walks	45

61	C Geometric Data Structures	45
62	C.1 Reductions Between Data Structure Problems	45
63	C.2 Data Structure for Square Graphs	48
64	C.3 Data Structure for Unit Disks	50
65	C.4 Data Structure for Unit Squares	52
66	D Switching Interval Representation between Different Stabbing Paths	56
67	E Handling Small Pieces	58
68	E.1 Patterns	58
69	E.2 Diameter and Eccentricities using Patterns	58
70	E.3 Boundary Weighted BFS in Geometric Intersection Graphs	59
71	E.4 Exact Distance Oracles	60

72 1 Introduction

73 A simple algorithm for computing the diameter of an unweighted n -vertex graph is to run a BFS from
74 every vertex of the graph. For sparse graphs or intersection graphs of various classes of geometric objects
75 (such as unit disks), BFS can be implemented in $\tilde{O}(n)$ time, leading to an algorithm to compute the
76 graph diameter in $\tilde{O}(n^2)$ time¹. Can we beat this naïve quadratic-time algorithm? More precisely, can we
77 compute the diameter in *truly subquadratic time* $O(n^{2-\varepsilon})$ for some fixed constant $\varepsilon > 0$ for these graphs?
78 This simple question has motivated the development of a broad range of techniques that have driven
79 algorithmic research for decades.

80 For general sparse graphs, even distinguishing the diameter between 2 and 3 in truly subquadratic
81 time is impossible, assuming the Strong Exponential Time Hypothesis (SETH) [RW13]. (For dense
82 undirected graphs, one can exploit the matrix multiplication subroutine to compute the diameter in
83 $O(n^\omega)$ time [Sei95], where $\omega < 2.371339$ is the matrix multiplication exponent [ADW⁺25].) Given the
84 negative result, it is natural to consider more structured classes of sparse graphs, such as planar and
85 minor-free graphs. For planar graphs, Cabello [Cab18] designed the first truly subquadratic algorithm
86 for the diameter problem by introducing a new technique based on abstract Voronoi diagrams. This
87 technique heavily exploits planarity and hence fails for minor-free graphs. Then Ducoffe, Habib, and
88 Viennot [DHV22] devised a new technique based on VC-dimension to compute the diameter of minor-
89 free graphs in truly subquadratic time. Both the Voronoi diagram and VC-dimension techniques are
90 major milestones in algorithm design for planar and minor-free graphs, opening the door for solving
91 other distance-related problems in truly subquadratic complexity (time or space), such as designing
92 compact (exact) distance oracles and computing eccentricities or Wiener index in planar and minor-free
93 graphs [Cab18, GKM⁺21, LP19, DHV22, LW24, KZ25].

94 For geometric intersection graphs of objects in the plane, designing a truly subquadratic time
95 algorithm for the diameter problem has been much more challenging. A *geometric intersection graph* is
96 a graph whose vertices are associated with objects in the plane, and edges correspond to object pairs
97 that intersect.² Unit-disk graphs—the intersection graphs of unit disks—are among the most basic
98 and well-studied graphs in the geometric setting; alternatively, this is equivalent to constructing an
99 unweighted graph based on a set of points in a metric space by connecting pairs of points whose distance
100 is below some fixed threshold.

101 While truly subquadratic algorithms have been ruled out for intersection graphs of unit segments, unit
102 equilateral triangles, or unit balls (in 3D) under standard fine-grained complexity assumptions [BKK⁺22],
103 the lower bound techniques for these objects fail for unit disks. Therefore, computing diameter for unit-
104 disk graphs in truly subquadratic time has become a central open problem raised by many authors [CS16,
105 BKK⁺22, DKP24, CGL24]. Such an algorithm points to a larger landscape where truly subquadratic
106 results for basic geometric intersection graphs are possible. We note that even distinguishing the diameter
107 between 2 and 3 in truly subquadratic time for unit-disk graphs remains open.

108 **Question 1.1.** *Can one compute the diameter of unit-disk graphs in truly-subquadratic time?*

109 Unlike planar graphs which are sparse, unit-disk graphs (and intersection graphs in general) can be
110 dense: they can contain cliques of arbitrary size. Even computing the BFS tree in $\tilde{O}(n)$ time becomes non-
111 trivial [CJ15]. Recently, Chang, Gao, and Le [CGL24] ported the VC-dimension technique for computing
112 diameter of minor-free graphs to unit-disk graphs; similar to planar graphs, on a unit-disk graph, the
113 radius- r balls for all integer values r also have bounded VC-dimension. A one-sentence summary of their
114 technique is that they treated a (possibly large) clique as a single vertex, and designed a clique-based

¹Throughout this paper, $\tilde{O}(\cdot)$ notation hides polylogarithmic factors, and $O^*(\cdot)$ hides $n^{o(1)}$ factors.

²We represent an intersection graph by the objects themselves, so the input size is $O(n)$ even if the graph could be dense.

separator hierarchy [dKMT23]. As a result, they obtained a subquadratic ($\tilde{O}(n^{2-1/18})$ -time) algorithm that could only compute an approximation of the diameter with an additive error at most 1 in unit-disk graphs. While the additive error is very small, their algorithm falls short of distinguishing between diameters 2 and 3. This suggests that computing the diameter *exactly* for unit-disk graphs requires a very different approach. (There are many examples in the general graph literature where allowing a small constant additive approximation can make the problem significantly easier to solve; for example, see [ACIM99].) For exact algorithms, Duraj, Konieczny, and Potępa [DKP24] adapted the technique by Ducoffe, Habib, and Viennot [DHV22], which is also based on VC-dimension and a *stabbing path data structure*, to the intersection graph of *unit squares*.³ However, their technique only works when the true diameter is small $D = O(n^{1/4-\epsilon})$ [DKP24] and more importantly, their stabbing path data structure does not work for unit disks, or even (non-unit) square graphs, as they heavily exploit the nice geometry of unit squares. (In fact, they explicitly asked, even when the diameter is a constant, if the diameter of a unit-disk graph can be computed in truly subquadratic time.)

1.1 Main Results on Diameter

In this paper, we give the first truly subquadratic algorithm for computing the diameter in unit-disk graphs, resolving Question 1.1 affirmatively. Moreover, our framework has many other applications and yields the first truly subquadratic algorithms for the intersection graph of axis-aligned (arbitrarily sized) squares, as well as arbitrary sparse graphs with bounded VC-dimension.

Theorem 1.2. *Let G be a graph on n vertices. We can compute the diameter of G by Las Vegas randomized algorithms in:*

- $O^*(n^{2-1/18})$ time if G is the intersection graph of unit disks, and
- $\tilde{O}(n^{2-1/12})$ time if G is the intersection graph of axis-aligned squares. For unit-square graphs, the running time is $O^*(n^{2-1/8})$.
- $\tilde{O}(mn^{1-1/(2d)})$ time if G has m edges and VC-dimension d . For the special case of K_h -minor-free graphs for a fixed h , the running time becomes $\tilde{O}(n^{2-1/(2h-2)})$.

See Table 1 for the summary of our results on the diameter problem in comparison with previous work. (Incidentally, our result even slightly improves previous time bounds in the special case of K_h -minor-free graphs. The fact that the exponent of our algorithm for unit disks is the same as in Chang, Gao, and Le's +1-approximation algorithm [CGL24] is a complete coincidence—the algorithms are very different.)

1.2 Technical Overview

All previous subquadratic diameter algorithms for planar and minor-free graphs for arbitrary diameters [Cab18, GKM⁺21, LW24, DHV22, CGL24] use sublinear separators (or r -divisions), which are not available for geometric intersection graphs that could be dense. A key highlight of our framework is that we completely bypass sublinear separators! Instead, we use *low-diameter decompositions (LDD)*. LDDs have been used in recent breakthrough results, such as negative-weight shortest paths [BNW22] and $(2 - \epsilon)$ -approximation for vertex cover on string graphs [LPS⁺24] (see the references in [BNW22] for more background). We stress that we only need the most elementary, non-probabilistic form of LDDs (dating back to [Awe85]), which are constructible simply by a number of “truncated” BFSes, and do not require expanders or flows. In some ways, they are even simpler than planar-graph separators or r -divisions.

³All squares are axis-aligned in this paper.

graph class	best previous	new
planar	$\tilde{O}(n^{5/3})$	[Cab18, GKM ⁺ 21]
K_h -minor-free	$\tilde{O}(n^{2-1/(3h-1)})$	[DHV22, LW24]
VC-dim.-bounded	$\tilde{O}(\min\{Dmn^{1-1/d}, mn\})$	[DHV22, DKP24]
unit-square	$\tilde{O}(\min\{Dn^{7/4}, n^2\})$	[DKP24]
square	$\tilde{O}(n^2)$	[CS19]
unit-disk	$O(n^2 \sqrt{\frac{\log \log n}{\log n}})$	[CS16]
		$O^*(n^{2-1/8})$
		$\tilde{O}(n^{2-1/12})$
		$O^*(n^{2-1/18})$

Table 1. Time bounds of exact diameter algorithms for different classes of unweighted graphs. Here, n is the number of vertices, m is the number of edges, D is the diameter, and d is the (generalized distance) VC-dimension. Squares are axis-aligned.

155 In addition to LDD, our framework incorporates many new ideas about the usage of *bounded VC-
156 dimension* as well as the design of *geometric data structures*. We will describe all three components of our
157 framework in a little more detail below.

158 **Component 1: Low-diameter decomposition.** For a given parameter $\Delta > 0$, a *low-diameter decomposition*
159 (*LDD*) decomposes the input graph into *pieces* of diameter at most Δ such that the total number of
160 boundary vertices of all the pieces is $\tilde{O}(n/\Delta)$. (It is helpful to imagine choosing $\Delta = n^\delta$ for some small
161 constant δ , and hence the number of boundary vertices is truly sublinear.) The ability to control the
162 total number of boundary vertices is reminiscent of *r-division* [Fre87] used for diameter computation in
163 planar [Cab18, GKM⁺21] and minor-free graphs [LW22], but an important difference is that a piece in
164 an LDD could have up to $\Omega(n)$ vertices, while in an *r-division*, every piece has truly sublinear size (for a
165 typical choice of r). LDDs can be computed in $\tilde{O}(m)$ time for general graphs and $\tilde{O}(n)$ time for many
166 classes of intersection graphs, as we will show (in Appendix A).

167 **Component 2: Bounded VC-dimension and stabbing paths.** Since any sparse graph has a good
168 low-diameter decomposition, an LDD itself is not sufficient for constructing truly subquadratic algorithms
169 due to the aforementioned conditional lower bound based on SETH [RW13]. A recent line of work on
170 the diameter problem has hinted at bounded VC-dimension as an overarching property: planar graphs
171 (more generally, minor-free graphs) [CEV07, BT15, LP19, DHV22, LW24] and intersection graphs of
172 pseudo-disks [ACM⁺21, DKP24, CGL24] (in particular, disks and squares) have bounded VC-dimension.
173 Thus, we also assume that the input graph has a bounded VC-dimension.

174 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
175 the cardinality of the largest $S \subseteq U$ such that S is *shattered* by \mathcal{F} —for every $S' \subseteq S$, there is some $X \in \mathcal{F}$
176 such that $X \cap S = S'$. Given a graph G , there are several different ways to form a set system of bounded
177 VC-dimension; see Section 2. The simplest one is the set system of neighborhood balls $(V_G, \{N^r[v]\}_{r \geq 0})$:
178 we say that a graph G has *VC-dimension*⁴ at most d if its system of neighborhood balls has VC-dimension
179 at most d . ($N^r[v]$ is the set of all vertices that are at a distance at most r from v , including v itself.) It
180 was known that planar graphs have VC-dimension at most 4; K_h -minor-free graphs have VC-dimension
181 at most $h-1$; and intersection graphs of pseudo-disks have VC-dimension at most 4 [CGL24].

182 There are two main ways that VC-dimension was used in the diameter computation: (1) *stabbing*
183 *path*: constructing a path that stabs each neighborhood ball $N^r[v]$ a sublinear number of times (in the
184 worst case or on average), and (2) *distance compression*: showing that there are few different distance

⁴A more precise terminology is *distance VC-dimension* at most d ; see Section 2 for clarification.

185 vectors to a fixed set of important vertices (i.e., the boundary of a piece in an r -division). The first
 186 approach has been very successful in the *low-diameter regime*: computing the diameter in time $\tilde{O}(Dn^{2-\varepsilon_d})$
 187 where ε_d is a constant depending on the VC-dimension d [DHV22, DKP24]. The second approach works
 188 for the arbitrary-diameter regime, but either requires sublinear separators [LW24] or allows distance
 189 approximation [CGL24]. We overcome the limitation and inherent obstacles from both approaches and
 190 devise a method in the presence of low-diameter decomposition to compute stabbing paths even when
 191 the graph diameter is large. (In certain applications, we also manage to perform distance compression
 192 exactly without the presence of separators.)
 193

194 The basic idea of the stabbing path approach is to order the vertices from 1 to n , in such a way that
 195 each neighborhood ball $N^r[v]$ of radius r can be represented as a union of $\tilde{O}(n^{1-1/d})$ many intervals on
 196 the stabbing path.⁵ The existence of a spanning path with $O(n^{1-1/d})$ stabbing (or “crossing”) number
 197 was first shown in a seminal paper by Chazelle and Welzl [CW89], and had found numerous applications
 198 in computational geometry, for example, in geometric range searching. Constructing a good stabbing
 199 path may seem to require knowledge of the entire set system of balls $N^r[v]$ in the first place (which we
 200 do not have, since our problem is to compute all $N^r[v]!$). Fortunately, it turns out that by known random
 201 sampling techniques⁶, we only need to evaluate a small subset of balls to compute a good stabbing path;
 202 for example, in the unit disk or square case, the construction time is $\tilde{O}(n^{1+1/d})$ (more generally, the
 construction time is $\tilde{O}(n\rho)$ for stabbing number $\tilde{O}(n/\rho + \rho^{d-1})$ for a trade-off parameter ρ).
 203

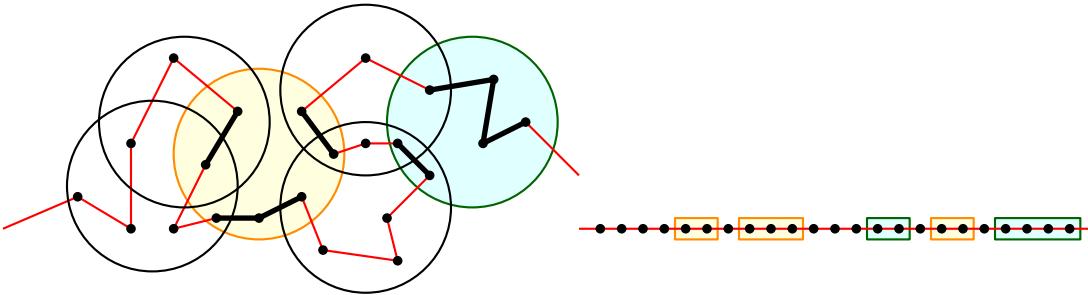


Figure 1. Stabbing path and interval representation of disks. The yellow disk is represented by three yellow intervals, and the green disk is represented by two intervals. The intervals representing different disks could overlap.

203 Let’s say v is one vertex in a diametral pair, whose shortest path distance realizes the diameter D . Given
 204 the interval representation, we can check in time linear to the number of intervals ($\tilde{O}(n^{1-1/d})$) whether
 205 the union of all the intervals (and hence $N^r[v]$) covers $[1 : n]$. If the answer is yes, then r is at least
 206 the diameter D . By iterating through every vertex v as a potential endpoint of a diametral path, we can
 207 check if r is greater than the true diameter D in $\tilde{O}(n^{2-1/d})$ time. To compute the interval representations
 208 of $N^r[v]$ for all vertices in V , the pioneering work of Ducoffe, Habib, and Viennot [DHV22] introduced a
 209 *ball growing* process: For each r , one computes the interval representations of $\{N^r[v] : v \in V\}$ from the
 210 interval representations of $\{N^{r-1}[v] : v \in V\}$ via the identity $N^r[v] = \bigcup_{u \in N^r[v]} \hat{N}^{r-1}[u]$. (For the base
 211 case $r = 0$, $N^0[v] = \{v\}$.) This approach leads to running time $\tilde{O}(Dmn^{1-1/d})$ for sparse graphs [DKP24].
 212 For geometric intersection graphs, we cannot afford to access the neighbors of every vertex as it would
 213 result in $\Omega(n^2)$ time, and hence, different ideas are needed to avoid explicitly accessing neighbors.
 214 For the intersection graphs of unit squares, Duraj, Konieczny, and Potępa [DKP24] devised a certain
 215 “neighbor-set data structure” to achieve total running time $\tilde{O}(Dn^{2-1/4})$. The factor of D in the running
 216 time seems inherent to this approach, and D could be as big as $\Omega(n)$. Furthermore, their data structure
 217 does not work for arbitrary squares or unit disks.

5For a simpler exposition, we assume the worst-case bound on $\tilde{O}(n^{1-1/d})$ on the number of intervals representing $N^r[v]$. In
 the detailed implementation, we work with an *amortized bound* which is faster to compute.

6This is where all our algorithms use Las Vegas randomization.

218 To handle possibly large D , our new approach is to *combine with low-diameter decomposition*. First,
 219 by computing BFS trees from the boundary vertices of the LDD with parameter Δ , we could compute an
 220 estimated value $\tau \in [D - \Delta : D]$. Note that there are only a truly sublinear number of boundary vertices,
 221 and hence, the total running time of this step remains truly subquadratic. (In the case of geometric
 222 intersection graphs, we can implement BFS in $\tilde{O}(n)$ time using known techniques such as bichromatic
 223 intersections [CS16, Klo23, GGL24]. See Appendix A.2 for details.) If we have $\{N^\tau[v] : v \in V\}$, then
 224 we only need to grow balls for another Δ iterations (here $\Delta \ll D$). However, we do not have access to
 225 $\{N^\tau[v] : v \in V\}$ explicitly. Our key idea here is to define a *modified* neighborhood ball $\hat{N}^\tau[v]$ in a way that
 226 we can initialize $\hat{N}^\tau[v] = \emptyset$ to kick-start the ball growing process, and at the same time the information
 227 computed is sufficient to answer the question about diameter D . Therefore, the precise definition of
 228 $\hat{N}^\tau[v]$ is somewhat tricky, tying directly to the pieces in the LDD; see Section 3 for the details. For sparse
 229 graphs (with bounded VC-dimension), we can afford to access the neighbors of every vertex explicitly.
 230 Hence, we could simply grow the modified balls $\{\hat{N}^\tau[v]\}$ in $O(\Delta)$ rounds in $\tilde{O}(\Delta \cdot mn^{1-1/d})$ time. By
 231 choosing $\Delta = n^{1/2d}$ (to balance with the $\tilde{O}(mn/\Delta)$ running time of BFS computation), this leads to a
 232 relatively simple diameter algorithm running in $\tilde{O}(mn^{1-1/(2d)})$ time, which is truly subquadratic for the
 233 arbitrary-diameter regime.

234 **Component 3: Geometric data structures.** For geometric intersection graphs that are not sparse, we
 235 cannot afford to access the neighbors directly even with the modified neighborhoods, so more ideas
 236 are needed. The data structure subproblem for the ball growing step we need to solve is the following:
 237 assume the interval representations of $\hat{N}^{r-1}[v]$ for every vertex v are precomputed and stored; given a
 238 query object s , compute the union of intervals in $\hat{N}^{r-1}[v]$ for all objects v that intersect s . We can reduce
 239 this problem to the following:

240 **Problem 1.3 (Interval Cover).** *Given a set of N objects \mathcal{O} and each object $o \in \mathcal{O}$ is associated with an
 241 interval $I_o \subseteq [1 : n]$. Design a data structure to answer the following query:*

- 242 • **COVERS?**(q, I): *Given a query object q and a query interval $I \subseteq [1 : n]$, decide whether the union
 243 of intervals associated with the objects intersecting⁷ q in \mathcal{O} covers the whole I .*

244 Problem 1.3 can be viewed as a generalization of *range searching* [AE99]: given a query object q , find
 245 the objects intersecting q . In the computational geometry literature, colored variants of range searching
 246 have been studied [GJS95, KRSV08, GJRS18, CHN20]. The above problem is an even more challenging
 247 variant, where each object is equipped with not a color but an interval. This interesting generalization
 248 has not been considered before, to the best of our knowledge. (There have been some prior works on
 249 *time-windowed* geometric data structures [BDG⁺14, BCE15, CHP19], but typically queries are associated
 250 with a time interval but not the input objects; even more crucially, the queries in those works are mainly
 251 about whether a property is true for *some* time value in a query interval I , rather than for *all* time
 252 values in I .) One reason the problem is more challenging than standard range searching is that it is not
 253 *decomposable* (if the input set is divided into two subsets, knowing the answers of a query for the subsets
 254 does not necessarily help with the overall answer).

255 We note that en route to their unit square result, Duraj, Konieczny, and Potępa [DKP24] also
 256 formulated a non-standard geometric “neighbor-set data structure” problem, but their formulation
 257 appears more complicated, as they (and Ducoffe, Habib, and Viennot [DHV22] earlier) worked with
 258 symmetric differences of neighborhood sets. Our approach using interval representations is in some
 259 sense “dual” to these previous approaches, and is more natural, leading to a geometric data structure
 260 problem that is simpler to state.

⁷Here we mean the objects intersect, not their associated intervals.

261 For unit squares, we give a solution to Problem 1.3 with $N^{1+o(1)}$ preprocessing time and $N^{o(1)}$ query
 262 time. Our data structure is deterministic, in contrast to Duraj et al.’s, which uses hashing techniques
 263 and inherently requires Monte Carlo randomization. For arbitrary squares, Duraj et al.’s data structure
 264 approach does not work at all. Although we are not able to obtain $N^{o(1)}$ query time for Problem 1.3
 265 either, we propose a simple method which divides the range $[1 : n]$ into blocks of size b , and builds
 266 a data structure for *rainbow colored intersection searching* (a version of colored range searching) for
 267 each block. (See Appendix C.1 for details.) This yields $\tilde{O}(N \cdot b)$ preprocessing and $\tilde{O}(L/b)$ query time,
 268 where L is the length of the query interval. This trade-off turns out to be sufficient to obtain a truly
 269 subquadratic algorithm in the end, for an appropriate choice of the parameter b .

270 For unit disks, Problem 1.3 is related to the well-known *Hopcroft problem*⁸, and hence a query time
 271 $o(N^{1/3})$ appears unlikely [Eri96]; however, to obtain a truly subquadratic time for diameter, we need
 272 $O(N^\delta)$ query time for tiny $\delta > 0$ (since the total number of input and query intervals is $\Omega(n^{2-1/4})$ or
 273 worse in our application). We circumvent this issue entirely by partitioning the given set of unit disks into
 274 constantly many *modulo classes* (i.e., we partition the plane into cells of constant side-length (say 1/2),
 275 and take modulo classes of the index pairs of the cells). This way, if we take one cell \square , the collection
 276 of disks from the same modulo class intersecting \square forms a *pseudoline arrangement*. When input disks
 277 are restricted to one modulo class, we are able to solve Problem 1.3 with $N^{1+o(1)}$ preprocessing time
 278 and $N^{o(1)}$ query time.⁹ These data structure results may be of independent interest to computational
 279 geometers. They do not follow directly from existing techniques. Instead, we propose a clever recursion,
 280 repeatedly and alternately taking *lower envelopes* and *upper envelopes* of pseudo-segments [AS00, Pet15].
 281 (Experts in geometric data structure may find this part interesting, and are encouraged to read Section C.3
 282 for the details.)

283 **Additional complications for unit disks.** The fact that we have efficient data structures for unit disks
 284 only when restricted to a fixed modulo class creates a number of extra technical challenges:

- 285 • Because we can only take union over balls from a fixed modulo class, the intermediate sets are no
 286 longer neighborhood balls, i.e., we need to work with a new set system. Fortunately, we can still
 287 prove that the (dual) VC-dimension of the new set system is at most 4, but *only when the balls
 288 have the same radius r* .
- 289 • This condition in turn forces us to change the stabbing path—and all of its associated interval
 290 representations—every time we increment r . Fortunately, we show that the interval representations
 291 can be updated efficiently using random sampling techniques (with slightly worse amortized
 292 stabbing number $\tilde{O}(n/\rho + \rho^d)$).
- 293 • At intermediate steps, we may now need to work with balls from two or three set systems across
 294 different types. Fortunately, the combined set systems still have bounded VC-dimension (at most 8).
- 295 • The extra overhead in switching stabbing paths is too costly since each stabbing path computation
 296 costs $O(n\rho)$ time, and we have to compute for $\tilde{O}(n/\Delta)$ pieces and $O(\Delta)$ rounds. To achieve overall
 297 subquadratic time, we only work with pieces larger than a certain threshold A ; for small pieces,
 298 we need to switch to a different method (based on distance compression), which achieves running
 299 time $\tilde{O}(n \cdot |\partial P| + |P| \cdot (|P| + (|\partial P| \Delta)^d))$ for each piece P of size at most A with boundary ∂P .

300 All these details are explained in Section 7, but to illustrate the intricacies of the overall algorithm to the
 301 curious readers, the time bound for diameter for unit-disk graphs has the following form, where the

⁸The Hopcroft problem tests, for a given system of points and lines in the Euclidean plane, whether any point lies on any line. The total number of points and lines is assumed to be n .

⁹We do not break the Hopcroft problem’s lower bound as we only solve the data structure problem for one modulo class.

graph class	best previous	new
planar	$O^*(n^{3/2})$, $O^*(n)$	[CGL ⁺ 23]
K_h -minor-free	$\tilde{O}(n^{2-1/(3h-1)})$	[LW24]
VC-dim.-bounded	$O(mn)$, $O(n^2)$	folklore
unit-square	$\tilde{O}(n^2)$	[CS19]
square	$\tilde{O}(n^2)$	[CS19]
unit-disk	$O(n^2 \sqrt{\frac{\log \log n}{\log n}})$	[CS16]
		$\tilde{O}(mn^{1-1/(4d+1)})$

Table 2. Construction time and space bounds of exact distance oracles for different classes of unweighted undirected graphs, with $\tilde{O}(1)$ query time. We write out both construction time and space bounds only when they are different.

302 sums are over all pieces P of the LDD (which satisfies $\sum_P |P| = O(n)$ and $\sum_P |\partial P| = \tilde{O}(n/\Delta)$):

$$303 O^* \left(\Delta \cdot n\rho + \sum_{P: |P| > A} \left(|\partial P| \cdot n + \Delta \cdot (n + |P| \cdot (n/\rho + \rho^8)) \right) + \sum_{P: |P| \leq A} \left(n \cdot |\partial P| + |P| \cdot (|P| + (|\partial P| \Delta)^4) \right) \right).$$

304 Balancing cost by setting parameters $\Delta = n^{1/18}$ and $\rho = A = \Delta^2$ then yields $O^*(n^{2-1/18})$. (Other variants
305 of the algorithm for different graph classes and other related problems will have different expressions
306 and different settings of parameters.)

307 1.3 Other Distance-related Problems

308 Our framework for computing diameter naturally opens up the possibility of solving other distance-related
309 problems. Here, we focus on three well-studied problems: all-vertex eccentricities, exact distance oracles,
310 and Wiener index.

311 **Eccentricities.** To highlight the new challenges beyond diameter computation, let us begin with
312 eccentricities. The *eccentricity* of a vertex v , denoted by $\text{ecc}(v)$, is the maximum distance from v to any
313 other vertex in G . Our goal is to compute $\text{ecc}(v)$ for every $v \in V_G$ in truly subquadratic time. Observe
314 that the diameter is the maximum eccentricity and hence, computing all eccentricities is often more
315 difficult.

316 For computing diameter, we kick-start the ball growing process with radius $\tau \in [D - \Delta : D]$ and
317 therefore we only need to grow in $O(\Delta)$ interactions. The key challenge in computing eccentricities
318 is that $\text{ecc}(v)$ of some vertex v could be as small as $D/2$, and hence any ball growing process has to
319 cover radii in the entire range $[D/2 : D]$, which can be as large as $\Omega(n)$. Interestingly, our framework for
320 the diameter problem also points us to a way to resolve this issue. Specifically, we grow the modified
321 neighborhood ball $\hat{N}^r[v]$ only for vertices in the same piece P of the low-diameter decomposition. The
322 observation is that for any two vertices u and v in P , $|\text{ecc}(u) - \text{ecc}(v)| = O(\Delta)$. Hence, to restrict to
323 computing eccentricities of vertices in P , it suffices to grow modified neighborhood balls in $O(\Delta)$ steps.
324 (For different pieces, the ranges of radii could be vastly different.) As the range of radii is piece-specific,
325 the stabbing path data structure also has to be piece-specific instead of being “global” as in the case of
326 computing graph diameter. Our results are summarized in the following theorem.

327 **Theorem 1.4.** *Let G be a graph on n vertices. We can compute all-vertex eccentricities of G by Las Vegas
328 randomized algorithms in:*

- 329 • $O^*(n^{2-1/20})$ time if G is the intersection graph of unit disks, and

- $\tilde{O}(n^{2-1/12})$ time if G is the intersection graph of axis-aligned squares. For unit-square graphs, the running time is $O^*(n^{2-1/8})$.
- $\tilde{O}(mn^{1-1/(2d)})$ time if G has m edges and (generalized distance) VC-dimension d .

Exact distance oracle. An *exact distance oracle* is a data structure that, when given a pair of vertices, returns the shortest path distance of the vertices quickly. Our goal is to construct an oracle with a truly subquadratic space. For geometric intersection graphs, known oracles with a truly subquadratic space can answer a distance query approximately within an additive error of 1 [AdT24, CGL24]; for sparse graphs, the query time is close to linear ($\Omega(n^{1-\varepsilon_d})$) for some small constant $\varepsilon_d = 1/2^{O(d)}$ depending on the VC-dimension d [DHV22].

For square and unit-disk graphs and sparse graphs with bounded VC-dimension, we provide an exact oracle with a truly subquadratic space and polylogarithmic query time. Furthermore, our oracle can be constructed in truly subquadratic time; therefore, our result can be interpreted as solving the all-pairs shortest-path problem in truly subquadratic time. (Of course any such algorithm has to output an implicit representation of the shortest distances since the explicit output size is $\Omega(n^2)$.)

Constructing an exact distance oracle is more difficult than computing all-vertex eccentricities: the queried distance range is $[0 : n]$. In computing diameter and eccentricities, the modified ball $\hat{N}^r[s]$ is a subset of the true neighborhood ball $N^r[s]$ and we compute $\hat{N}^r[s]$ for all $r \in [\text{ecc}(s) - O(\Delta) : \text{ecc}(s) + \Delta]$. However, if we query distance between s and t where $d_G(s, t) \ll \text{ecc}(s) - O(\Delta)$, then knowing the true neighborhood ball $N^r[s]$ for $r \geq \text{ecc}(s) - O(\Delta)$ (let alone its subset) does not tell us anything about $d_G(s, t)$. Our idea is to assign weights to vertices of G appropriately and incorporate vertex weight in the definition of $\hat{N}^r[s]$, so that every vertex t belongs to $\hat{N}^r[s]$ for some value of $r \in [-\Delta : \Delta]$; the radius could be negative, which is somewhat counterintuitive. As the range of the (weighted) radii is now $\Theta(\Delta)$, the ideas we develop for computing the diameter and eccentricities now can be applied here. As distances with vertex weights are closely connected to the notion of *generalized* VC-dimension (formally defined in Section 2), we assume the input graph has a bounded generalized VC-dimension in the case of sparse graphs. All other graphs, such as geometric intersection graphs and minor-free graphs, have their generalized VC-dimension equal to the regular VC-dimension (of the neighborhood ball system).

All of these ideas lead to our exact distance oracles for various types of graphs. See Table 2 for a comparison of existing results and ours.

Theorem 1.5. *Let G be a graph on n vertices. We can compute an exact distance oracles for G (by randomized Las Vegas algorithms) with the following guarantees:*

- $O^*(n^{2-1/20})$ construction time and size and $\tilde{O}(1)$ query time if G is a unit-disk graph.
- $\tilde{O}(n^{2-1/20})$ construction time and size and $\tilde{O}(1)$ query time if G is a square graph. For unit-square graphs, the construction time and size can be improved to $\tilde{O}(n^{2-1/16})$.
- $\tilde{O}(mn^{1-1/(4d+1)})$ construction time, $\tilde{O}(n^{2-1/(4d+1)})$ size, and $\tilde{O}(1)$ query time if G has m edges and (generalized distance) VC-dimension d .

Interestingly, if we ignore construction time, the above theorem implies the existence of subquadratic-size, $\tilde{O}(1)$ -time distance oracles for all (not necessarily sparse) graphs with bounded (generalized distance) VC-dimension, in particular, all pseudo-disk graphs.

Wiener index. The Wiener index of a graph G is the sum of the distances between all pairs of vertices. Computing the Wiener index has been studied [CK97, CK09, WN09]; truly subquadratic algorithms are only known for planar and minor-free graphs [Cab18, GKM⁺21, LW24]. Here we provide the first such

372 algorithms for graphs with bounded generalized VC-dimension and several geometric intersection graphs.
 373 Indeed, the algorithms for Wiener index are simple corollaries of our algorithms for exact distance oracles
 374 in Theorem 1.5 and therefore have the same running time guarantees.

375 **Theorem 1.6.** *Let G be a graph on n vertices. We can compute the Wiener index of G (by randomized
 376 Las Vegas algorithms) in:*

- 377 • $O^*(n^{2-1/20})$ time if G is the intersection graph of unit disks.
- 378 • $\tilde{O}(n^{2-1/20})$ time if G is the intersection graph of axis-aligned squares. For unit-square graphs, the
 379 running time is $\tilde{O}(n^{2-1/16})$.
- 380 • $\tilde{O}(mn^{1-1/(4d+1)})$ time if G has m edges and (generalized distance) VC-dimension d .

381 2 Preliminaries

382 2.1 Graphs and Low-diameter Decomposition

383 **Graph notation.** Let $G = (V_G, E_G)$ be an unweighted undirected graph with n vertices and m edges.
 384 For two vertices $u, v \in V$, let $d_G(u, v)$ denote the distance between u and v in G . Often we will omit
 385 the subscript and simply write $d(u, v)$ when the graph G is clear. The **neighborhood** of a vertex $v \in V_G$
 386 is the set of vertices that are distance at most 1 to v , denoted by $N[v] := \{u \in V_G : d(u, v) \leq 1\}$. The
 387 **k -neighborhood ball** of a vertex $v \in V_G$ is the set of vertices with distance at most k from v , denoted by
 388 $N^k[v] := \{u \in V : d(u, v) \leq k\}$. (Notice that $N[v] = N^1[v]$ and $N^k[v] = N[N^{k-1}[v]]$.) Define the set of
 389 k -neighborhood balls as $\mathcal{N}_G^k := \{N^k[v] : v \in V\}$, and the set of all neighborhoods balls as $\mathcal{B}_G := \bigcup_k \mathcal{N}_G^k$.

390 **Geometric intersection graphs.** Consider a set S of n geometric objects in the plane. We define the
 391 **geometric intersection graph** G of S as the graph obtained by creating a vertex for every geometric object,
 392 and connecting two geometric objects if they intersect. When S consists of unit disks, i.e., disks of radius
 393 1, we refer to the geometric intersection graph G as a **unit-disk graph**. If S consists of axis-aligned unit
 394 squares, we refer to the geometric intersection graph G as a **unit-square graph**. We will also consider
 395 when S consists of axis-aligned squares (of arbitrary size). We refer to such graphs as **square graphs**. In
 396 Appendix A.2, we describe a near-linear time algorithm for computing a BFS tree for square graphs, as
 397 stated below; the algorithm for unit-disk graphs is known [Klo23].

398 **Lemma 2.1.** *Let G be the geometric intersection graph of squares or unit disks with n vertices. We can
 399 compute a BFS tree from any given vertex of G in $\tilde{O}(n)$ time.*

400 **Low-diameter decomposition.** Let G be a graph with n vertices and $\Delta > 0$ be a diameter parameter.
 401 A **low-diameter decomposition** (LDD) of G with parameter Δ is a decomposition of the vertex set V into
 402 disjoint sets $V = V_1 \cup \dots \cup V_k$ and corresponding induced subgraphs $P_i := G[V_i]$ called **pieces**, such that:

- 403 • **Low diameter:** Piece P_i is a single connected component of (strong) diameter¹⁰ at most Δ .
- 404 • **Small boundary:** Denote the boundary vertices of P_i as ∂P_i , that is, the subset of vertices of P_i
 405 that has an edge to a vertex in $V_G \setminus V_i$. The decomposition satisfies $\sum_{i=1}^k |\partial P_i| = \tilde{O}(n/\Delta)$.
- 406 • **No small pieces:** Each piece has size at least $\tilde{\Omega}(\Delta)$.

407 We show in Appendix A that such a decomposition always exists. Furthermore, in Appendix A.1, we
 408 show an efficient algorithm for computing this decomposition.

¹⁰By strong diameter we mean that the shortest path between any two vertices in P_i within the subgraph P_i is at most Δ .

409 **Theorem 2.2.** Let G be a graph with n vertices and m edges. For any parameter $24 \log n < \Delta \leq n$, we
410 we can compute a low-diameter decomposition for G in $O(m + n)$ time.

411 For unit-disk graphs and square graphs, we prove in Appendix A.2 that the low-diameter decomposi-
412 tion is efficiently computable in near-linear time.

413 **Theorem 2.3.** Let G be an intersection graph of n unit disks or an intersection graph of n axis-aligned
414 squares. For any parameter $24 \log n < \Delta \leq n$, we can compute a low-diameter decomposition for G in
415 $\tilde{O}(n)$ time.

416 2.2 VC-dimension

417 A **set system** is a pair (X, \mathcal{S}) , consisting of a ground set X and a collection of ranges that are subsets of X ;
418 in notation, $\mathcal{S} \subseteq 2^X$. A subset $Y \subseteq X$ is said to be **shattered** by \mathcal{S} if the collection $\{Y \cap S : S \in \mathcal{S}\} = 2^Y$,
419 that is, all possible subsets of Y can be obtained by \mathcal{S} . The **shatter function**, denoted by $\pi_{(X, \mathcal{S})}(k)$ is the
420 largest number of sets that is created by the set system when restricted to $Y \subseteq X$ of size k . Formally it is:

$$421 \pi_{(X, \mathcal{S})}(k) = \max_{\substack{Y \subseteq X \\ |Y|=k}} |\{Y \cap S : S \in \mathcal{S}\}|.$$

422 The **shatter dimension** of a set system is the smallest value d such that $\pi_{(X, \mathcal{S})}(k) = O(k^d)$ for all k . The
423 **VC-dimension** of a set system (X, \mathcal{S}) is the size of the largest subset of $Y \subseteq X$ that can be shattered
424 by \mathcal{S} . The **dual set system** of (X, \mathcal{S}) is the set system (\mathcal{S}^*, X^*) , where the ground set $\mathcal{S}^* = \{w_S : S \in \mathcal{S}\}$
425 consists of elements indexed by \mathcal{S} , and each $s \in S$ induces a range $s^* = \{w_S \in \mathcal{S}^* : S \ni s\}$ in \mathcal{S}^* . The **dual
426 VC-dimension** of a (X, \mathcal{S}) is the VC-dimension of the dual set system, and analogously the **dual shatter
427 dimension** is the shatter dimension of the dual set system. We state some well-known results [Har11].

428 **Lemma 2.4.** Let (X, \mathcal{S}) be a set system of VC-dimension d . The following is true:

- 429 1. The dual set system (\mathcal{S}^*, X^*) has VC-dimension at most $2^{d+1} - 1$.
- 430 2. For $Y \subseteq X$, the set system (Y, \mathcal{S}) has VC-dimension at most d .
- 431 3. (Sauer-Shelah Lemma [She72, Sau72].) If $|X| \leq n$ then $|\mathcal{S}| \leq O(n^d)$, so the shatter dimension of
432 (X, \mathcal{S}) is at most d .

433 **VC-dimension in graphs.** The **k -distance VC-dimension** of a graph $G = (V_G, E_G)$ is the VC-dimension
434 of the set system of k -neighborhood balls (V_G, \mathcal{N}_G^k) . (Sometimes in the literature, e.g., [DHV22], the
435 VC-dimension of G is defined to be the 1-distance VC-dimension.) The **distance VC-dimension** of G is the
436 VC-dimension of the set system of balls (V_G, \mathcal{B}_G) . Observe that the k -neighborhood set system (V_G, \mathcal{N}_G^k)
437 is equivalent to its dual, so the dual VC-dimension is the same as the primal. This is not the case for the
438 set system of arbitrary balls since the ground set and the set of ranges have different sizes.

439 Karczmarz and Zheng [KZ25] introduced¹¹ a natural generalization: a set system (U, \mathcal{GB}_G) whose
440 ground set is $U = V_G \times \mathbb{Z} = \{(u, r) : u \in V_G, r \in \mathbb{Z}\}$, and the ranges \mathcal{GB}_G consists of **generalized
441 neighborhood balls** for $v \in V_G$ and $k \in \mathbb{Z}$ of the form:

$$442 \tilde{N}^k[v] := \{(u, r) \in V_G \times \mathbb{Z} : d(u, v) \leq r + k\}.$$

443 Note that values of r and k are allowed to be negative. We call the VC-dimension of (U, \mathcal{GB}_G) the
444 **generalized distance VC-dimension** of a graph. It can be observed that this set system is equivalent to its
445 dual. Furthermore, we can observe the following relationship between these VC-dimensions.

¹¹[KZ25] consider what they call a *multiball* set system where the ground set is $V_G \times M$ for a set of real weights $M \subseteq \mathbb{R}$.

446 **Observation 2.5.** k -distance VC-dimension of $G \leq$ distance VC-dimension of $G \leq$ generalized distance
 447 VC-dimension of G .

448 Throughout this paper, when we refer to graphs of bounded VC-dimension, we will be referring to
 449 families of graphs whose generalized distance VC-dimension of the graph is bounded by an absolute
 450 constant. Many of our results can also be adapted with more work to graphs that have bounded k -distance
 451 VC-dimension for all k , or graphs with bounded distance VC-dimension. We will focus on generalized
 452 distance VC-dimension as it holds for minor-free graphs and the geometric intersection graphs we care
 453 about, and also leads to the simplest exposition of our ideas.

454 **Connection to distance encoding VC-dimension.** Distance encodings were used by Li and Parter
 455 [LP19] to compute the diameter in a planar graph. This was later modified to a more general setting by
 456 Le and Wulff-Nilsen [LW24], whose definition we present below (restricted to unweighted graphs).

457 **Definition 2.6.** Let $G = (V_G, E_G)$ be an undirected unweighted graph. Let $M \subseteq \mathbb{Z}$ be a set of integers.
 458 Let $S \subseteq V_G$ be an ordered set of ℓ vertices $S = \{s_0, s_1, \dots, s_{\ell-1}\}$. For every vertex $v \in V_G$ define the set:

$$459 X_{S,M}(v) := \{(s_i, \delta) : s_i \in S, \delta \in M, d(v, s_i) - d(v, s_0) \leq \delta\}.$$

460 Let $X_{S,M} := \{X_{S,M}(v) : v \in V\}$ be the set of subsets of the ground set $S \times M$. The *distance encoding*
 461 *VC-dimension* of G is the maximum VC-dimension of set systems of the form $(S \times M, X_{S,M})$ for all possible
 462 S and M .

463 Observe that the set $X_{S,M}(v)$ is isomorphic to $\tilde{N}^{d(v,s_0)}[v] \cap (S \times M)$. Restricting the ground set of the
 464 set system (U, \mathcal{GB}_G) to $(S \times M, \mathcal{GB}_G)$ does not increase the VC-dimension by Lemma 2.4, so we conclude
 465 the following observation.

466 **Observation 2.7.** Distance encoding VC-dimension of $G \leq$ generalized distance VC-dimension of G .

467 **Graphs of bounded generalized distance VC-dimension.** It was shown by Chepoi, Estellon, and
 468 Vaxes [CEV07] that planar graphs have distance VC-dimension at most 4 by explicitly constructing a K_5
 469 minor (by contradiction). This argument was extended by Bousquet and Thomassé [BT15] to show that
 470 K_h -minor-free graphs have distance VC-dimension at most $h - 1$. Le and Wulff-Nilsen [LW24] used a
 471 variation of this argument to show that K_h -minor-free graphs have distance encoding VC-dimension at
 472 most $h - 1$. This argument was adapted by Karczmarz and Zheng [KZ25] to show that K_h -minor-free
 473 graphs have generalized distance VC-dimension at most $h - 1$ as well.

474 **Theorem 2.8** ([KZ25]). Any K_h -minor-free graph has generalized distance VC-dimension at most $h - 1$.

475 For unit-disk graphs, it was shown by Abu-Affash *et al.* [ACM⁺21] that the distance VC-dimension
 476 is 4. Later, by Chang, Gao, and Le [CGL24], the intersection graph of pseudo-disks has distance VC-
 477 dimension and distance encoding VC-dimension of 4 as well. The bound on distance VC-dimension was
 478 also independently shown by Duraj, Konieczny, and Potępa [DKP24] for intersection graphs of fixed
 479 translates of geometric objects in the plane. The proof in [CGL24] can be easily adapted to also bound
 480 the generalized distance VC-dimension.

481 **Theorem 2.9.** Any geometric intersection graph of pseudo-disks in the plane has generalized distance
 482 VC-dimension at most 4.

483 2.3 Stabbing Path and Interval Representation

484 Let (X, \mathcal{S}) be a set system with $|X| \leq n$ and $|\mathcal{S}| \leq m$. Let λ be an ordering of the elements of X . Given a
 485 set $S \in \mathcal{S}$, define the **λ -interval representation** $\text{Rep}_\lambda(S)$ (λ -representation for short) as the collection of
 486 maximal contiguous subsequences of λ —called **intervals**—whose union is S . The size of the representation
 487 $|\text{Rep}_\lambda(S)|$ refers to the number of such intervals. For a parameter $1 \leq \rho \leq m$, a **ρ -stabbing path** λ of a set
 488 system (X, \mathcal{S}) of dual VC-dimension d is an ordering of X such that $\sum_{S \in \mathcal{S}} |\text{Rep}_\lambda(S)| = \tilde{O}(mn/\rho + m\rho^{d-1})$.
 489 Observe that if $n^{1/d} \leq m$, this quantity is minimized when $\rho = n^{1/d}$ so $\sum_{S \in \mathcal{S}} |\text{Rep}_\lambda(S)| = \tilde{O}(mn^{1-1/d})$.
 490 We will sometimes refer to an $n^{1/d}$ -stabbing path λ simply as a **stabbing path**. We assume the existence
 491 of an **element reporting oracle** that, given $S \in \mathcal{S}$, can enumerate all elements of S in $T_0(n)$ time, where
 492 $T_0(n) \geq n$.

493 In Appendix D, we show the following lemma to construct ρ -stabbing paths with high probability¹².

494 **Lemma 2.10.** *Let (X, \mathcal{S}) be a set system with $|X| \leq n$ and $|\mathcal{S}| \leq m$ with dual shatter dimension at most
 495 d . For any parameter $1 \leq \rho \leq m$, we can construct a stabbing path of (X, \mathcal{S}) , that is, an ordering λ of X
 496 such that $\sum_{S \in \mathcal{S}} |\text{Rep}_\lambda(S)| = \tilde{O}(mn/\rho + m\rho^{d-1})$ in $\tilde{O}(T_0(n) \cdot \rho)$ time with high probability.*

497 2.4 Geometric Data Structures

498 We consider three different geometric data structures in decreasing difficulty that we will use in our
 499 algorithms, and the reduction from difficult problems to easier problems. We provide the details of the
 500 reductions in Appendix C.1.

501 **Interval searching.** The interval searching problem directly captures the ball growing process for
 502 various objects in the frameworks we study in Section 3 and Section 8.

503 **Problem 2.11 (Interval Searching).** *Let \mathcal{O}_{IS} be a given set of objects, where each object $o \in \mathcal{O}_{IS}$ is
 504 associated with a set of intervals (of integer points) of $[1 : n]$, denoted by \mathcal{I}_o . Design a data structure
 505 that answers the following query:*

- 506 • INTERVALSEARCH(q): *Given an object q , return (the interval representation of) all the integer
 507 points in $[1 : n]$ associated with objects in \mathcal{O}_{IS} that intersect q .*

508 For each query object q , let $\mathcal{I}_{out}(q)$ be the set of intervals representing all the integer points of
 509 the **output**. Ideally, we want to construct a data structure for the interval searching problem that has
 510 near-linear preprocessing time and poly-logarithmic query time. However, this is difficult even when the
 511 objects are unit-disk graphs.

512 In our context, we will be querying interval searching for each object in \mathcal{O}_{IS} , and therefore, we
 513 will be solving the offline version of Problem 2.11. Let $N_{IS} := \sum_{o \in \mathcal{O}_{IS}} (|\mathcal{I}(o)| + |\mathcal{I}_{out}(o)|)$ be the total
 514 number of input and output intervals. Let $L_{IS} := \sum_{o \in \mathcal{O}_{IS}} \sum_{I \in \mathcal{I}(o) \cup \mathcal{I}_{out}(o)} |I|$ be the total length of the input
 515 and output intervals. We want to construct a data structure \mathcal{D}_{IS} for solving Problem 2.11 that has a
 516 small total run time as a function of N_{IS} and L_{IS} . Here, the **total run time** of the \mathcal{D}_{IS} includes: (1) the
 517 preprocessing time and (2) the total time to answer all the queries.¹³

¹²In this paper we say an event E happens **with high probability** if $\Pr[E] \geq 1 - n^{-c}$ for some big enough constant c .

¹³Alternatively, the offline version of Problem 2.11 is equivalent to computing a Boolean matrix product $C = AB$, where A is the adjacency matrix of an intersection graph, and B and C are Boolean matrices whose 1 entries can be covered by a small number of row intervals. We will not adopt this viewpoint here.

518 **Interval cover.** This is the data structure Problem 1.3. Recall that N is the number of input objects.
 519 Let N_{IC} be the total number of input objects and query objects, and L_{IC} be the total length of the input
 520 and query intervals. Similar to the interval searching problem, we want to construct a data structure \mathcal{D}_{IC}
 521 for solving Problem 1.3 with small total running time, as a function of N_{IC} and L_{IC} . We will show (in
 522 Appendix C.1) that if we can solve the interval cover problem efficiently, then we can solve the interval
 523 searching efficiently.

524 **Lemma 2.12.** *If one can construct a data structure \mathcal{D}_{IC} for solving Problem 1.3 with total run time
 525 $T(N_{IC}, n, L_{IC})$ (for some polynomial function T), then we can construct a data structure \mathcal{D}_{IS} for solving
 526 Problem 2.11 in total run time $\tilde{O}(T(N_{IS}, n, L_{IS}))$. Furthermore, if \mathcal{D}_{IC} has preprocessing time $P(N)$ and
 527 query time $Q(N)$, then \mathcal{D}_{IS} has preprocessing time $\tilde{O}(P(\tilde{N}_{IS}))$ and query time $\tilde{O}(Q(\tilde{N}_{IS}) \cdot |\mathcal{I}_{out}(q)|)$ where
 528 $\tilde{N}_{IS} := \sum_{o \in \mathcal{O}_{IS}} |\mathcal{I}(o)|$ is the total number of input intervals and $\mathcal{I}_{out}(q)$ is the set of output intervals from
 529 the interval search query of q to \mathcal{D}_{IS} .*

530 **Rainbow colored intersection search.** On the surface, the next problem we present seems to be a
 531 strict special case of Problem 1.3 by requiring the interval to be a singleton. However, we will show that
 532 a solution to this problem gives us solutions to the two other problems.

533 **Problem 2.13 (Rainbow Colored Intersection Searching).** *Given a set of objects \mathcal{O}_{RC} , each object
 534 $o \in \mathcal{O}_{RC}$ is associated with a color. Design a data structure to answer the following query:*

- 535 • RAINBOWCOVER?(q): *Given a query object q , decide whether all the colors appear in the set of
 536 objects intersecting q .*

537 In Appendix C.1, we show how to use a data structure \mathcal{D}_{RC} for solving Problem 2.13 to design a data
 538 structure \mathcal{D}_{IC} for solving Problem 1.3.

539 **Lemma 2.14.** *If we can construct in $\tilde{O}(|\mathcal{O}_{RC}|)$ time a data structure \mathcal{D}_{RC} with $\tilde{O}(1)$ query time for solving
 540 Problem 2.13, then for any parameter $b \in [1, n]$, we can construct a data structure \mathcal{D}_{IC} for solving
 541 Problem 1.3 that has total run time $\tilde{O}(N_{IC} \cdot b + L_{IC}/b)$.*

542 This together with Lemma 2.12 implies a solution to the interval searching problem, in particular, a
 543 data structure \mathcal{D}_{IS} for solving Problem 2.11 that has total run time $\tilde{O}(N_{IS} \cdot b + L_{IS}/b)$.

544 2.5 Handling of Small Pieces

545 While the low-diameter decomposition guarantees that all pieces have size at least $\tilde{\Omega}(\Delta)$, sometimes this
 546 guarantee is not enough, and we will switch to a different algorithm.

547 **Diameter and eccentricities.** For computing diameter and eccentricities, we use the notion of *pat-
 548 terns* [LP19], and present the following lemma implicit in the work of Le and Wulff-Nilsen [LW24].

549 **Lemma 2.15.** *Let G be a graph on n vertices with distance encoding VC-dimension d . Let P be a piece in
 550 G with boundary ∂P and diameter Δ . If distances from ∂P to all vertices of G are known, the eccentricity
 551 of all vertices in P can be computed in $O(n \cdot |\partial P| + (|P| + |\partial P|^d \Delta^d) \cdot T(P))$ where $T(P)$ is the time it
 552 takes to run boundary weighted BFS on P with weights at most Δ .*

553 We add one small optimization to the result of Le and Wulff-Nilsen [LW24] using the notion of
 554 boundary weighted BFS, a BFS where boundary vertex distances are initialized. This boundary weighted
 555 BFS can be performed in time linear in the number of edges of the piece for sparse graphs, and in time
 556 near-linear in the number of vertices of P for geometric intersection graphs. See Appendix E.2 for further
 557 details and a complete proof of Lemma 2.15.

558 **Distance oracles.** Similarly, for distance oracles, we will use the following lemma, also implicit in the
 559 work of Le and Wulff-Nilsen [LW24].

560 **Lemma 2.16 (Section 4.3.1 of [LW24]).** *Let $G = (V_G, E_G)$ be a graph with bounded distance VC-
 561 dimension d , and P be an induced subgraph of G with boundary ∂P and diameter Δ . There exists a
 562 distance oracle that answers distances from any vertex $s \in P$ and any vertex $t \in V_G$ with $O(n \cdot |\partial P| + |P|^d)$
 563 space and $O(\log |\partial P|)$ query time.*

564 *Furthermore, if G also has bounded generalized distance VC-dimension d and distances from ∂P
 565 to all vertices of G , the distance oracle can be computed in $O(n \cdot |\partial P| + (|\partial P|^d \Delta^d + |P|) \cdot T(P))$ time,
 566 where $T(P)$ is the time it takes to run vertex weighted BFS on P with weights at most Δ .*

567 For completeness, we provide the proof of Lemma 2.16 in Appendix E.4.

568 3 Framework for Diameter and Eccentricities

569 In this section, we outline the algorithmic framework for computing the all-vertex eccentricity of different
 570 graph families in truly subquadratic time. (Recall that the *eccentricity* of a vertex u is defined to be
 571 $\text{ecc}(u) := \max_{v \in V_G} d(u, v)$.) Note that as diameter is the maximum eccentricity of any vertex in the graph,
 572 we can also compute diameter in truly subquadratic time. Our framework can be tweaked for other
 573 problems, such as constructing distance oracles and Weiner index; we defer to Section 8.

574 Now we formally set up the framework, which consists of a few high-level instructions, with the goal
 575 to compute for every vertex u , the r -neighborhood balls $N^r[u]$ iteratively for growing values of r . This is
 576 enough to answer the diameter problem exactly because a graph G has diameter at most D if and only if
 577 every radius- D neighborhood ball contains all vertices in G .

578 Let G be the input graph, given either explicitly using adjacency lists or implicitly as the intersection
 579 graph of objects. Our framework has three steps.

580 **Step 1: Low-diameter decomposition (LDD).** Compute a low-diameter decomposition \mathcal{L} of G with
 581 a diameter parameter $\Delta > 0$. \mathcal{L} has $\tilde{O}(n/\Delta)$ pieces, each of strong diameter at most Δ . Furthermore,
 582 $\sum_{P \in \mathcal{L}} |\partial P| = \tilde{O}(n/\Delta)$. The vertices in $\bigcup_P \partial P$ are called the *boundary vertices*.

583 **Step 2: Shortest-path computations.** For each boundary vertex v in $\bigcup_P \partial P$, compute a breadth-first
 584 search tree in G rooted at v . We obtain $\text{ecc}(v)$ as a byproduct. Define

$$585 \partial\text{ecc} := \max_{\text{boundary vertex } v} \text{ecc}(v).$$

586 **Step 3: Growing neighborhood balls.** Consider one piece P in the low-diameter decomposition \mathcal{L} .
 587 Our goal is to compute *some modified version of $N^r[s]$* for every vertex s in P and *necessary values of r* .
 588 Fix an arbitrary vertex s_p in ∂P , and define ecc_P as the corresponding eccentricity $\text{ecc}(s_p)$. (Notice that
 589 $\text{ecc}(s_p)$ is known after the shortest-path computation in Step 2 because s_p is a boundary vertex of P .)
 590 We set the *modified neighborhood ball* for each vertex s in P to be

$$591 \hat{N}^r[s] := N^r[s] \cap R_p, \text{ with } R_p := \{t \in V_G : \text{dist}(s_p, t) \geq \text{ecc}_P - 2\Delta\},$$

592 where Δ was defined to be the strong diameter bound of the pieces in the LDD and R_p is called the
 593 *relevant region* for the eccentricity computation for P . We will compute $\hat{N}^r[s]$ for every s in $P \setminus \partial P$
 594 iteratively using the inductive formula

$$595 \hat{N}^r[s] = \bigcup_{v \in N[s]} \hat{N}^{r-1}[v]. \tag{1}$$

596 We emphasize that while the notation seems to suggest otherwise, the definition of modified neighbor-
 597 hood balls $\hat{N}^r[s]$ depends on the piece P . Define the set of *relevant balls* to be

$$598 \quad \mathcal{S}_P := \{\hat{N}^r[v] : v \in P, r \in [\text{ecc}_P - 3\Delta, \text{ecc}_P + \Delta]\}.$$

599 Let $\mathcal{S} := \bigcup_{P \in \mathcal{L}} \mathcal{S}_P$. The *ball growing process* consists the following substeps:

- 600 3.1. For every $s \in \partial P$ and every $r \in [\text{ecc}_P - 3\Delta, \text{ecc}_P + \Delta]$, compute modified balls $\hat{N}^r[s]$ using Step 2.
- 601 3.2. As a base case, we initialize $\hat{N}^r[s] = \emptyset$ for every $s \in P \setminus \partial P$ when $r = \text{ecc}_P - 3\Delta - 1$.
- 602 3.3. For other values of $r \in [\text{ecc}_P - 3\Delta, \text{ecc}_P + \Delta]$, compute $\hat{N}^r[s]$ using the inductive formula (1).

603 Then $\text{ecc}(s)$ is the smallest value r such that $\hat{N}^r[s]$ is the whole relevant region R_p . Therefore, we can
 604 compute $\text{ecc}(s)$ from $\{\hat{N}^r[s] : r \in [\text{ecc}_P - 3\Delta - 1, \text{ecc}_P + \Delta]\}$.

605 Note that we will assume that the diameter of the graph is at least 4Δ , otherwise the entire graph G
 606 is a low-diameter decomposition of parameter 4Δ , and we can simply apply Step 3 with the relevant
 607 neighborhood balls being all balls, the relevant region R_p being V_G , and the modified neighborhood balls
 608 being normal neighborhood balls.

609 **Correctness.** To show that our algorithmic framework is correct, we show that we correctly computed
 610 all (modified) neighborhood balls, assuming the ball growing process is correct. First note that for $s \in \partial P$,
 611 we have correctly computed the modified neighborhood balls in Step 3.1. If $s \in P \setminus \partial P$, given the pair
 612 (s, t) realizing $\text{ecc}(s)$, we can guarantee that the vertex t must lie in the relevant region R_p : Denote t_p
 613 to be the vertex that has distance ecc_P to s_p , then because $\text{dist}(s, s_p) \leq \Delta$, we have

$$614 \quad \text{dist}(s_p, t) \geq \text{dist}(s, t) - \Delta \geq \text{dist}(s, t_p) - \Delta \geq \text{dist}(s_p, t_p) - 2\Delta,$$

615 and thus t can be found in $\hat{N}^r[s]$ if $\hat{N}^r[s]$ is a relevant ball in \mathcal{S}_P . Furthermore, again by triangle inequality,
 616 $\text{ecc}(s)$ is at least $\text{ecc}_P - \Delta$ and at most $\text{ecc}_P + \Delta$. Thus it is sufficient to initialize r to be $\text{ecc}_P - 3\Delta - 1$ (in
 617 which case $N^r[s] \cap R_p = \emptyset$), so the initialization in Step 3.2 is correct. Thus, assuming the ball expansion
 618 step is correct, all modified neighborhood balls are computed correctly in Step 3.3.

619 **VC-dimension of neighborhood balls and stabbing paths.** We cannot afford to store the (modified)
 620 neighborhood balls $\hat{N}^r[v] \in \mathcal{S}_P$ explicitly. Instead, we will rely on a compact representation of a set system
 621 with bounded VC-dimension to store the neighborhood balls *implicitly* in a data structure. Given the
 622 (modified) neighborhood ball system (V_G, \mathcal{S}_P) , we are responsible for bounding the (dual) VC-dimension
 623 of (V_G, \mathcal{S}_P) to be a constant d . Then we compute stabbing path λ for (V_G, \mathcal{S}_P) , such that the interval
 624 representation $\text{Rep}_\lambda(\hat{N}^r[v])$ of set $\hat{N}^r[v]$ has sublinear size. (See Section 2.3 for definition.)

625 **Implementing the ball growing process.** To implement the ball growing process, we will use a
 626 stabbing path λ for the modified neighborhood balls to ensure we can compactly store all such balls. The
 627 exact details on how we implement the process will depend on the type of graph we are dealing with.

628 In a sparse graph G , we will show how to implement the ball expansion data structure in G directly
 629 by explicitly considering the neighbors $N[v]$ of each vertex v in the graph G .

630 In a geometric intersection graph G , we instead implement the ball growing process for a piece $P \in \mathcal{L}$
 631 with a data structure for the interval searching problem defined in Problem 2.11. Each vertex v in P is
 632 associated with a geometric object o_v . Let \mathcal{O}_P denote the set of these geometric objects. Suppose we
 633 have computed compact interval representations $\text{Rep}_\lambda(\hat{N}^{r-1}[v])$ for every vertex v in $P \setminus \partial P$, so we can
 634 associate these intervals to o_v . Using a data structure \mathcal{D}_{IC} for Problem 2.11, the union of intervals of

635 objects in \mathcal{O}_p that intersect with o_v is exactly $\hat{N}^r[v]$ by Equation (1). Thus, we can implement the ball
 636 growing process in a geometric intersection graph if we have an offline data structure for Problem 2.11.
 637 The efficiency of the algorithm will depend on the number of intervals in the representation with respect
 638 to the stabbing path λ .

639 **Organization.** In the next four sections, we will apply our framework to devise algorithms for diameter
 640 and eccentricities for different graph classes: sparse graphs of bounded VC-dimension (Section 4),
 641 arbitrary-square graphs (Section 5), unit-square graphs (Section 6), and unit-disk graphs (Section 7).
 642 Sections 4, 5–6, and 7 can be read independently, depending on the interest of the reader. The sparse
 643 graph case is perhaps the simplest, not requiring geometric data structures. The unit-disk case is the
 644 most involved and requires overcoming a number of (interesting) technical challenges.

645 4 Diameter/Eccentricities in Sparse Graphs of Bounded VC-dimension

646 We begin by applying the framework in Section 3 to sparse graphs of bounded VC-dimension. In
 647 this setting, the low-diameter decomposition could be constructed in $O(m)$ time using Theorem 2.2.
 648 Computing the BFS tree for every boundary vertex takes $\tilde{O}(mn/\Delta)$ total time where Δ is the diameter
 649 parameter in the low-diameter decomposition. Thus we focus on the third step of performing ball
 650 expansion.

651 To begin, we construct a global ordering λ on all the vertices for our stabbing path data structure.
 652 The following is analogous to Corollary 15 in [DKP24] that we tailor to our setting.

653 **Lemma 4.1.** *We can compute in $\tilde{O}(mn^{1/d})$ time an ordering λ of the vertices in V such that for the
 654 system $\mathcal{S} = \bigcup_{P \in \mathcal{L}} \mathcal{S}_P$ such that $\sum_{P \in \mathcal{L}} \sum_{s \in P} \sum_{r=ecc_p-3\Delta}^{ecc_p+\Delta} \deg(s) \cdot |\text{Rep}_\lambda(N^r[s])| = \tilde{O}(\Delta mn^{1-1/d})$.*

655 **Proof:** Let $\check{\mathcal{S}}$ be the set obtained by taking each set $N^r[s]$ in \mathcal{S} and adding $\deg(s)$ copies of $N^r[s]$ to $\check{\mathcal{S}}$.
 656 Observe that:

$$657 |\check{\mathcal{S}}| = \sum_{P \in \mathcal{L}} \sum_{s \in P} \sum_{r=ecc_p-3\Delta}^{ecc_p+\Delta} \deg(s) = O(\Delta \cdot m)$$

658 We then apply Lemma 2.10 to $X = V(G)$ and $\check{\mathcal{S}}$ with $\rho = n^{1/d}$. Since we can implement the element
 659 reporting oracle in $O(m)$ time via BFS, the result follows. \square

660 Now for every relevant piece P in the low-diameter decomposition \mathcal{L} , we will restrict our attention
 661 to only the relevant region R_P for the eccentricity computation. To do so, we consider the ordering λ_P of
 662 R_P obtained from λ by restricting to the vertices of R_P . Observe that doing so does not increase the size
 663 of the interval representation of any sets.

664 **Observation 4.2.** *Let R be a subset of V . Let λ be an ordering of the vertices V , and λ' be an ordering
 665 of R obtained by restricting λ to the vertices in R . Then for any set $S \subseteq V$, $|\text{Rep}_{\lambda'}(S \cap R)| \leq |\text{Rep}_\lambda(S)|$.*

666 **Proof:** For any interval $I \in \text{Rep}_\lambda(S)$, $I \cap R$ is also an interval in λ' . \square

667 **Ball expansion data structure.** To implement the ball expansion data structure, we will store each
 668 neighborhood ball in interval form. For a boundary vertex $s \in \partial P$, we can compute $\hat{N}^r[s]$ for all
 669 $r \in [ecc_p - 3\Delta, ecc_p + \Delta]$ in $O(n)$ time using the BFS tree we have computed from step 2, and in addition
 670 represent these balls in interval form.

671 Next we describe how to perform the ball expansion operation. For vertices $s \in P \setminus \partial P$, each neighbor
 672 $v \in N[s]$ is also in P and we have a compact interval representation for $\hat{N}^{r-1}[s]$. We can take the union
 673 of the set of intervals by doing a line sweep in time $\tilde{O}\left(\sum_{v \in N[s]} |\text{Rep}_{\lambda_p}(\hat{N}^{r-1}[v])|\right)$.

674 Furthermore, it is easy to detect if $\hat{N}^r[s] = R_p$ if the interval representation is all of λ_p .

675 **Time analysis.** The amount of time taken for computing the global ordering in Lemma 4.1 is $\tilde{O}(mn^{1/d})$.
 676 The runtime for ball expansion of the boundary vertices is:

$$677 \sum_{P \in \mathcal{L}} \sum_{s \in \partial P} O(n) = \tilde{O}(n^2/\Delta) \quad (2)$$

678 By Observation 4.2, the ball expansion for a non-boundary vertex $s \in P$ and $s \notin \partial P$ from radius $r-1$ to
 679 r takes time

$$680 \tilde{O}\left(\sum_{v \in N[s]} |\text{Rep}_{\lambda_p}(\hat{N}^{r-1}[v])|\right) \leq \tilde{O}\left(\sum_{v \in N[s]} |\text{Rep}_{\lambda}(\hat{N}^{r-1}[v])|\right).$$

681 The total time taken for all ball expansion steps for non-boundary vertices across all the pieces is at most:

$$682 \sum_{P \in \mathcal{L}} \sum_{s \in P} \sum_{r=ecc_p-3\Delta}^{ecc_p+\Delta} \tilde{O}\left(\sum_{v \in N[s]} |\text{Rep}_{\lambda}(\hat{N}^{r-1}[v])|\right) = \tilde{O}\left(\sum_{P \in \mathcal{L}} \sum_{s \in P} \sum_{r=ecc_p-3\Delta}^{ecc_p+\Delta} \deg(s) \cdot |\text{Rep}_{\lambda}(\hat{N}^r[s])|\right) \\ 683 = \tilde{O}(\Delta mn^{1-1/d}).$$

684 The last equality follows from Lemma 4.1.

685 Recall that the first two steps of the framework can be implemented in $\tilde{O}(mn/\Delta)$ time. The total
 686 runtime for all three parts is:

$$687 \tilde{O}(mn/\Delta + mn^{1/d} + n^2/\Delta + \Delta mn^{1-1/d}) = \tilde{O}(mn/\Delta + \Delta mn^{1-1/d}).$$

688 Setting $\Delta = O(n^{1/(2d)})$ yields a that this algorithm runs in $\tilde{O}(mn^{1-1/(2d)})$ time.

689 **Theorem 4.3.** *The diameter problem in a sparse undirected graph G with n vertices and m edges and
 690 general distance VC-dimension at most d can be solved in $\tilde{O}(mn^{1-1/(2d)})$ time.*

691 **Remark 4.4.** We can also obtain similar results (albeit with possibly worse exponents) for other VC-
 692 dimension bounds. If the distance VC-dimension is bounded by d or even if the k -neighborhood
 693 VC-dimension is bounded by d for all k , we can follow an approach similar to what we have for unit-disk
 694 graphs (see Section 7). The main difference is an extra step to reorder the vertices when we transition
 695 from $k-1$ -neighborhoods to k -neighborhoods using Appendix D.

696 5 Diameter/Eccentricities in Square Graphs

697 Next, we apply the framework in Section 3 to intersection graphs of squares. In step 1, we apply
 698 Theorem 2.3 to obtain our low-diameter decomposition \mathcal{L} of G in $\tilde{O}(n)$ time into pieces of size $\log n \leq$
 699 $\Delta \leq n$, where Δ is a parameter we will choose later. In step 2, we compute BFS trees from each
 700 $v \in \bigcup_P \partial P$ using Lemma 2.1. The algorithm takes $\tilde{O}(n)$ time per vertex, so this step takes $\tilde{O}(n^2/\Delta)$ time.
 701 Note that we also explicitly store all distances from these vertices, which takes $\tilde{O}(n^2/\Delta)$ space.

702 Recall that when $s \in \partial P$ then we can explicitly compute $\hat{N}^r[s]$ for all values of r in $O(n)$ time using
 703 the distances computed in step 2 of our framework.

704 **Stabbing path.** We now compute a global stabbing path λ . We use the following lemma.

705 **Lemma 5.1.** *Let G be a graph with generalized distance VC-dimension d , and a single-source distance*
 706 *finding algorithm with running time $T(n)$. Then the modified neighborhood ball system has a path λ*
 707 *such that we have*

$$708 \sum_{P \in \mathcal{L}} \sum_{s \in P} \sum_{r=ecc_p-3\Delta}^{ecc_p+\Delta} |\text{Rep}_\lambda(\hat{N}^r[s])| = \tilde{O}(\Delta \cdot n^{2-1/d})$$

709 with high probability, i.e., λ is a stabbing path of the modified r -balls. Furthermore, λ can be computed
 710 with a randomized algorithm in $\tilde{O}(n^{1/d} T(n))$ time.

711 **Proof:** We apply Lemma 2.10 with $\rho = n^{1/d}$ for the set system

$$712 S := \{N^r[s] : P \in \mathcal{L}, s \in P \setminus \partial P, r \in [ecc_p - 3\Delta, ecc_p + \Delta]\}.$$

713 Notice that the system has size at most $|S| = 3\Delta n$, and we can use the BFS algorithm to report the
 714 squares in the modified ball. By Observation 4.2, as $\hat{N}^r[s] = R_p \cap N^r[s]$, we obtain the bound

$$715 \sum_{P \in \mathcal{L}} \sum_{s \in P} \sum_{r=ecc_p-3\Delta}^{ecc_p+\Delta} |\text{Rep}_\lambda(\hat{N}^r[s])| \leq \sum_{P \in \mathcal{L}} \sum_{s \in P} \sum_{r=ecc_p-3\Delta}^{ecc_p+\Delta} |\text{Rep}_\lambda(N^r[s])| = \tilde{O}(\Delta \cdot n^{2-1/d}) \quad \square$$

717 In all of the intersection graphs studied in this paper, we have $T(n) = \tilde{O}(n)$ and $d = 4$. This leads
 718 to a stabbing path λ that is computed in $\tilde{O}(n^{5/4})$ time and has the property that the total size of the
 719 representation is $\tilde{O}(\Delta \cdot n^{7/4})$. Given that there are $O(\Delta \cdot n)$ modified balls we consider, the *amortized*
 720 interval count to represent a single modified ball is $O(n^{3/4})$.

721 **Growing balls.** To grow the modified neighborhood balls, we will design a data structure for solving
 722 the interval searching problem (Problem 2.11) for squares, which we restate here: we are given a set
 723 of square S , where each square $s \in S$ is associated with a set of intervals (of integer points) of $[1 : n]$,
 724 denoted by \mathcal{I}_s . Design a data structure that answers the following query:

725 • **INTERVALSEARCH(q):** Given a square q , return (the interval representation of) all the integer points
 726 in $[1 : n]$ associated with squares in S that intersect q .

727 We will be querying the data structure once for each square $s \in S$. Therefore, we are interested in
 728 minimizing the total query time. Let $\mathcal{I}_{out}(s)$ be the set of output intervals for a query square s . Let $\mathbf{N} :=$
 729 $\sum_{s \in S} (|\mathcal{I}(s)| + |\mathcal{I}_{out}(s)|)$ be the total number of input and output intervals. Let $\mathbf{L} := \sum_{s \in S} \sum_{I \in \mathcal{I}(s) \cup \mathcal{I}_{out}(s)} |I|$
 730 be the total length of the input and output intervals.

731 **Lemma 5.2.** *For any parameter $b \in [1 : n]$, we can construct a data structure \mathcal{D} for solving the interval*
 732 *searching problem for squares such that the total time to (i) construct \mathcal{D} and (ii) answer $|S|$ queries, one*
 733 *for each square $s \in S$, is $\tilde{O}(N \cdot b + L/b)$.*

734 **Proof:** By Lemma 2.14, it suffices to construct a data structure \mathcal{D}_{RC} for the rainbow colored intersection
 735 searching for squares that has nearly linear preprocessing time and poly-logarithmic query time. We
 736 provide such a data structure in Appendix C.2. \square

737 Next, we present a simpler (but slower) algorithm for computing all eccentricities. Then we show
 738 how to improve the running time.

739 **First version.** To compute all eccentricities, for each piece P we restrict λ to R_P in $O(n)$ time, and
 740 denote the resulting ordering by λ_P . Next, we set $r = \text{ecc}_P - 3\Delta$ and compute the balls $\{\hat{N}^r[s]\}$ for each
 741 s . In general, once the representations of $\hat{N}^{r-1}[s]$ are known, the data structure to set up for computing
 742 modified balls of radius r will have $\sum_{s \in P} |\text{Rep}_{\lambda_P}(\hat{N}^{r-1}[s])|$ intervals in it.

743 To compute the representations of $\{\hat{N}^r[s]\}_{s \in P}$, we setup the data structure \mathcal{D} that takes as input: (i) a
 744 set of squares corresponding to vertices of P and (ii) the interval representation $\{\text{Rep}_{\lambda_P}(\hat{N}^{r-1}[s])\}_{s \in P}$ for
 745 radius $r-1$. Then we apply $|P|$ queries $\{\text{INTERVALSEARCH}(s) : s \in P\}$ to output the interval representations
 746 of $\{\hat{N}^r[s]\}_{s \in P}$. Observe that the total length of all the intervals is at most $2|P| \cdot |R_P| = O(|P| \cdot n)$. Thus,
 747 the total running time for each r is:

$$748 \tilde{O}\left(b \cdot \sum_{s \in P} (|\text{Rep}_{\lambda_P}(\hat{N}^{r-1}[s])| + |\text{Rep}_{\lambda_P}(\hat{N}^r[s])|) + |P|n/b\right)$$

749 Therefore, the total running time of computing all-vertex eccentricities, including the running time
 750 of the first two steps in the framework, is:

$$751 \begin{aligned} & \tilde{O}(n^2/\Delta + n^{5/4}) + \sum_{P \in \mathcal{L}} \sum_{r=\text{ecc}_P-3\Delta}^{\text{ecc}_P+\Delta} \tilde{O}\left(b \cdot \sum_{s \in P} (|\text{Rep}_{\lambda_P}(\hat{N}^{r-1}[s])| + |\text{Rep}_{\lambda_P}(\hat{N}^r[s])|) + |P|n/b\right) \\ &= \tilde{O}(n^2/\Delta + n^{5/4}) + \tilde{O}(b) \left(\sum_{P \in \mathcal{L}} \sum_{r=\text{ecc}_P-3\Delta}^{\text{ecc}_P+\Delta} |\text{Rep}_{\lambda_P}(\hat{N}^r[s])| \right) + \tilde{O}(n^2\Delta/b) \\ &= \tilde{O}(n^2/\Delta + n^{5/4}) + \tilde{O}(b\Delta \cdot n^{7/4}) + \tilde{O}(n^2\Delta/b) \quad (\text{by Lemma 5.1 and } d=4) \\ &= \tilde{O}(n^2/\Delta + b\Delta \cdot n^{7/4} + n^2\Delta/b) \\ &= \tilde{O}(n^{2-1/16}). \quad (\text{for optimal choices of } b = \Delta^2 \text{ and } \Delta = n^{1/16}) \end{aligned} \quad (3)$$

752 **Improved version.** We improve the running time by reducing the $\tilde{O}(n^2\Delta/b)$ in Equation (3), which
 753 is the *total length of the intervals*, to $\tilde{O}(n^2/b)$ by keeping track of the sets $\hat{N}^r[s] \setminus \hat{N}^{r-1}[s]$ instead of
 754 $\hat{N}^r[s]$. Notice that the eccentricity of s will be the largest r where $\hat{N}^r[s] \setminus \hat{N}^{r-1}[s]$ is non-empty. Let
 755 $\hat{N}^{=r}[s] := \hat{N}^r[s] \setminus \hat{N}^{r-1}[s]$. Then $|\text{Rep}_{\lambda_P}(\hat{N}^{=r}[s])| \leq |\text{Rep}_{\lambda_P}(\hat{N}^r[s])| + |\text{Rep}_{\lambda_P}(\hat{N}^{r-1}[s])|$ and therefore,
 756 $\{\hat{N}^{=r}[s]\}_{s \in P, P \in \mathcal{L}}$ has a compact representation:

$$757 \sum_{P \in \mathcal{L}} \sum_{r=\text{ecc}_P-3\Delta}^{\text{ecc}_P+\Delta} |\text{Rep}_{\lambda_P}(\hat{N}^{=r}[s])| \leq \sum_{P \in \mathcal{L}} \sum_{r=\text{ecc}_P-3\Delta}^{\text{ecc}_P+\Delta} (|\text{Rep}_{\lambda_P}(\hat{N}^{r-1}[s])| + |\text{Rep}_{\lambda_P}(\hat{N}^r[s])|) = \tilde{O}(\Delta \cdot n^{7/4}).$$

758 Observe that

$$759 \hat{N}^{=r}[s] = \left(\bigcup_{v \in N[s]} \hat{N}^{=r-1}[v] \right) \setminus \left(\hat{N}^{=r-1}[s] \cup \hat{N}^{=r-2}[s] \right).$$

760 Thus, we could apply the same growing ball process for $\hat{N}^r[s]$. More precisely, we compute the
 761 interval representation of $\bigcup_{v \in N[s]} \hat{N}^{=r-1}[v]$ by querying the interval searching data structure, and then
 762 remove elements from $\hat{N}^{=r-1}[s] \cup \hat{N}^{=r-2}[s]$ using the interval representations of $\hat{N}^{=r-1}[s]$ and $\hat{N}^{=r-2}[s]$
 763 computed from the previous iterations. On the other hand, the intervals in this representation are
 764 disjoint, so we can bound the total length L over all 3Δ iterations as $L \leq O(|P|n)$ instead of $O(\Delta|P|n)$.
 765 Therefore, by applying the same calculation in Equation (3), the final running time is:

$$766 \tilde{O}(n^2/\Delta + n^{5/4}) + \tilde{O}(b) \cdot N + \tilde{O}(n^2/b) = \tilde{O}(n^2/\Delta + b\Delta n^{7/4} + n^2/b) = \tilde{O}(n^{2-1/12})$$

767 for $b = \Delta = n^{1/12}$.

768 **Theorem 5.3.** *Computing the diameter and all-vertex eccentricities of square graphs with n vertices
 769 can be done in $\tilde{O}(n^{2-1/12})$ time.*

770 6 Diameter/Eccentricities in Unit-square Graphs

771 For unit squares, we can obtain a slightly faster algorithm. The first two steps and the stabbing path
 772 computation are the same as the algorithm for square graphs. The total running time of these steps is
 773 $\tilde{O}(n^2/\Delta + \Delta n^{5/4})$. The growing ball step is more efficient since we can develop a better geometric data
 774 structure for unit squares.

775 **Growing balls.** For unit squares, we design a more efficient data structure for the interval cover problem
 776 (Problem 1.3), and as a result, we obtain a more efficient data structure for the interval searching problem.
 777 Let S be a given set of unit squares where each square $q \in S$ is associated with a set of intervals of $[1 : n]$.
 778 Each query INTERVALSEARCH(q) returns the intervals of integer points in $[1 : n]$ associated with unit
 779 squares intersecting q . Let $\tilde{N} := \sum_{s \in S} |\mathcal{I}(s)|$ be the number of input intervals.

780 **Lemma 6.1.** *We can construct in $\tilde{N}^{1+o(1)}$ time a data structure $\tilde{\mathcal{D}}$ for solving the interval searching
 781 problem for unit squares that can answer each query INTERVALSEARCH(q) in $\tilde{N}^{o(1)} \cdot |\mathcal{I}_{out}(q)|$ time.*

782 **Proof:** In Appendix C.2 (and more specifically Theorem C.5), we construct a data structure for the
 783 interval cover problem (Problem 1.3) for unit square with $\tilde{N}^{1+o(1)}$ preprocessing time and $\tilde{O}(1)$ query
 784 time. Then by Lemma 2.12, we obtain the preprocessing time and query time as in the lemma. \square

785 **The algorithm.** The algorithm is essentially the same for the square graphs in the previous section:
 786 restricting the ordering λ to R_p to get λ_p , and growing balls $\{\hat{N}^r[s]\}_{s \in P}$ for $r = \text{ecc}_p - 3\Delta$ to $\text{ecc}_p + \Delta$ by
 787 applying a query INTERVALSEARCH(s) for each unit square $s \in P$ to the interval searching data structure
 788 in Lemma 6.1 built for $\hat{N}^{r-1}[s]$. Let $\tilde{N}_{r-1} := \sum_{s \in P} (|\text{Rep}_{\lambda_p}(\hat{N}^{r-1}[s])|)$ be the number of input intervals to
 789 the data structure; the total output size is \tilde{N}_n . Since $\tilde{N}_{r-1}^{o(1)} = n^{o(1)}$, the total running time to grow all the
 790 balls per piece is

$$791 n^{o(1)} \cdot \sum_{r=\text{ecc}_p-3\Delta}^{\text{ecc}_p+\Delta} |\text{Rep}_{\lambda}(\hat{N}^r[s])|.$$

792 The total running time to compute all eccentricities is:

$$793 \begin{aligned} & \tilde{O}(n^2/\Delta + \Delta n^{5/4}) + n^{o(1)} \sum_{P \in \mathcal{L}} \sum_{r=\text{ecc}_p-3\Delta}^{\text{ecc}_p+\Delta} |\text{Rep}_{\lambda}(\hat{N}^r[s])| \\ &= O^*(n^2/\Delta + \Delta n^{5/4} + \Delta n^{7/4}) \quad (\text{by Lemma 5.1 and } d = 4) \\ &= O^*(n^{2-1/8}) \quad (\text{for } \Delta = n^{1/8}). \end{aligned}$$

794 **Theorem 6.2.** *Computing the diameter and all-vertex eccentricities in unit square graphs with n vertices
 795 can be done in $O^*(n^{2-1/8})$ time.*

796 7 Diameter/Eccentricities in Unit-disk Graphs

797 We now describe how to adapt the framework in Section 3 to the more complicated setting of computing
 798 diameter and eccentricities for unit-disks. The computation of LDD and BFS remains unchanged because
 799 unit-disks are fat pseudo-disks; we follow Step 1 and Step 2 of the framework (Section 3). For LDD
 800 we use Theorem A.3; for BFS we use Lemma 2.1. However, Step 3 requires drastic changes in order to
 801 implement the ball expansion step.

802 In this section we assume the unit-disk graphs are in *center-disk intersection model*: Unlike a typical
 803 geometric intersection graph where we create an edge between two objects if they intersect, here we add
 804 an edge between two unit-disks if the center of one disk lies in the other disk. It is straightforward to see
 805 that the two models are equivalent by doubling the radii of all unit disks. For the sake of simplicity, we
 806 will scale the disks so that the radius is still one unit.

807 7.1 Restriction to Fixed Types

808 **Partition into modulo classes.** We partition the plane into square *cells*: every unit-square is divided
 809 into 2×2 many cells, each of side length $1/2$. Each cell is indexed by the coordinates of its bottom-left
 810 corner modulo 3; notice that the coordinates are multiples of $1/2$ and thus there are 6 modulo classes
 811 per coordinate. We collect all cells of the same index (i, j) into a set $Cell_{i,j}$; in other words,

$$812 Cell_{i,j} := \{ \text{cell } \square : \text{square } \square \text{ is located at } (x, y) \text{ where } x \equiv i \text{ and } y \equiv j \pmod{3} \}.$$

813 We then classify the set of unit-disks \mathcal{D} based on the cell classes where the center of the unit-disk lies:

$$814 \mathcal{D}_{i,j} := \{ D \in \mathcal{D} : \text{disk } D \text{ has its center located in some cell in } Cell_{i,j} \}.$$

815 Notice that $\{\mathcal{D}_{i,j}\}_{i,j}$ is a partition of \mathcal{D} . We say a disk in $\mathcal{D}_{i,j}$ has *type* (i, j) . Denote the number of types
 816 to be σ ; there are exactly $\sigma = 36$ types.

817 The disks intersecting a query disk D_q whose center point q lies inside a cell \square come in two flavors:
 818 those disks that completely contain the cell \square , and those that partially intersect the cell. We call the cells
 819 where the centers of these intersecting disks belong *relevant* to \square ; among them, we call those cells with
 820 disks partially intersecting \square *perimetric*. Observe that there are only constantly many cells relevant to
 821 any fixed cell \square , because we set the side length of each cell to be $1/2$.

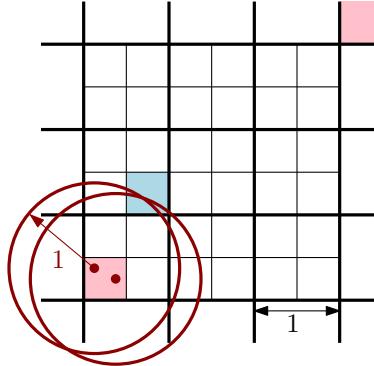


Figure 2. The 36 cells formed by partitioning a 3×3 square. Cells of the same color are of the same type. The disks in $\mathcal{D}_{0,0}$ (in pink) intersects the blue cell as a pseudoline arrangement.

822 Fixed an arbitrary cell \square . A *pseudoline arrangement* \mathcal{L} inside \square is a collection of boundary-to-boundary
 823 simple curves in \square , such that every pair of curves in \mathcal{L} intersect each other at most once. We now establish
 824 the main combinatorial property for disks of the same type: while two unit-disks intersect up to two
 825 times in the plane, if we focus on a single cell \square and two disks whose centers are in some other cells of
 826 the same type, at most one intersection will appear in \square . (Indeed, based on the way we partition cells
 827 into modulo classes, at most one cell of each type is relevant to \square .)

828 **Lemma 7.1.** *Let \mathcal{D} be any set of unit-disks, partitioned into types as described above. Given any cell \square
 829 and a fixed type (i, j) , the boundary of the disks in $\mathcal{D}_{i,j}$ intersects inside \square as a pseudoline arrangement.*

830 **Proof:** Assume for contradiction that there are two intersecting unit-disks D and D' with two intersection
 831 points inside \square simultaneously. As \square has side-length $1/2$ and diameter at most $\sqrt{2}/2 \leq 0.71$, the centers
 832 of D and D' must be at least $2 \cdot (1^2 - (\sqrt{2}/4)^2)^{1/2} = \sqrt{14}/2 \geq 1.87$ units away. Thus the two centers
 833 cannot be in the same cell (which has diameter at most 0.71). On the other hand, we reach a contradiction
 834 as any two distinct cells of the same type must be at least 2.5 units away (because we index the cells by
 835 modulo 3), while the centers of the intersecting disks D and D' can be at most 2 units away. \square

836 If we focus on one perimetric cell \square' of \square and rotate the plane so that \square' lie about vertically above \square
 837 (the cells might not be parallel to the axis anymore), we can safely assume each pseudoline formed by
 838 the partial intersection by a disk in \square' with \square to be ***x-monotone***, that is, any vertical line intersects the
 839 pseudoline at most once.

840 7.2 Implementation of the Neighborhood Growing Step

841 We first describe how to implement the ball growing process (Step 3.3) in the framework using the
 842 inductive formula (1), which we recall here:

$$843 \hat{N}^r[s] = \bigcup_{v \in N[s]} \hat{N}^{r-1}[v]. \quad (1)$$

844 Fix a piece P from the LDD, and some vertex s in $P \setminus \partial P$. Each modified neighborhood ball $\hat{N}^r[s]$ can
 845 be written as the union of a collection of r -balls with restrictions on the type of the second vertex in
 846 length- r paths. More precisely, recall that a disk D is of type (i, j) if the center of the disk D is located in
 847 some cell whose bottom-left corner has coordinates in the modulo class (i, j) . We arbitrarily order the
 848 constantly many types and label them from 1 to σ . We say that a path in the intersection graph G from
 849 a vertex s to a vertex v is a ***τ -path*** if the disk corresponding to the vertex following s (the second vertex)
 850 in the path is of type $\tau \in [1 : \sigma]$. For each vertex s and an arbitrary subset $M \subseteq [1 : \sigma]$, define

$$851 \hat{N}_M^r[s] := \{v \in V_P : \text{there is a } \tau\text{-path from } s \text{ to } v \text{ of length at most } r, \text{ for some } \tau \in M\}. \quad (4)$$

852 We often use the \leq_T subscript to represent the subset $[1 : T]$ below. Notice that by restricting to τ -paths
 853 from s to v in the definition, we can implement the inductive formula using

$$854 \hat{N}_T^r[s] = \bigcup_{v \in N[s] \cap \text{Cell}_T} \hat{N}_{\leq \sigma}^{r-1}[v].$$

855 We prove the following lemma in Appendix B (Lemma B.2).

856 **Lemma 7.2.** *For every r and every $T \in [1 : \sigma]$, both set systems*

$$857 (V_G, \{\hat{N}_T^r[s] : s \in V_P\}) \quad \text{and} \quad (V_G, \{\hat{N}_{\leq T}^r[s] : s \in V_P\})$$

858 *have dual VC-dimension at most 4.*

859 Notice that $\hat{N}^r[s] = \hat{N}_{\leq \sigma}^r[s]$ and we can compute $\hat{N}_{\leq T}^r[s]$ iteratively by

$$860 \hat{N}_{\leq T}^r[s] = \hat{N}_{\leq T-1}^r[s] \cup \hat{N}_T^r[s] = \hat{N}_{\leq T-1}^r[s] \cup \bigcup_{v \in N[s] \cap \text{Cell}_T} \hat{N}_{\leq \sigma}^{r-1}[v]. \quad (5)$$

861 Now the strategy should be clear: We will compute $\hat{N}_T^r[s]$ from the previously stored $\hat{N}_{\leq \sigma}^{r-1}[v]$ for every
 862 1-neighbor v of s , then take the union with $\hat{N}_{\leq T-1}^r[s]$. The first operation is done with the help of the
 863 geometric data structure; to do so, one has to first switch the interval representation for the relevant

neighborhood balls to be with respect to a unifying stabbing path over some *combined* set system that has a bounded VC-dimension. After $\hat{N}_T^r[s]$ is computed, we switch the interval representation back then proceed to compute $\hat{N}_T^r[s] = \hat{N}_{\leq T-1}^r[s] \cup \hat{N}_T^r[s]$, this time switching the interval representations to be with respect to another unifying stabbing path over some *auxiliary* set system. Both set systems must have $O(1)$ (dual) VC-dimension. This is the main technical hurdle which we will explain next.

Geometric data structure. Fix a cell \square containing s and a perimetric cell \square' of \square which lies vertically above \square (after a rotation). By Lemma 7.1, the collection of disks whose center lies in \square' intersects \square as a collection of N pseudolines \mathcal{L} . Assume each pseudoline ℓ_v in \mathcal{L} has an associated interval I_v in some given stabbing path λ . Our next goal is to describe how to build a stabbing path data structure \mathcal{D}_λ that answers the $\text{COVERS?}(s, I)$ query: whether the union of intervals I_v for every object v intersecting s covers the whole I . (This data structure would then be used to solve the interval searching problem, which gives us the interval representation of the modified neighborhood of s .) A disk D_v intersecting a query disk D_s with center s in \square corresponds to a pseudoline ℓ_v that lies below the center s of the query disk. Therefore it is equivalent if we can support the following query:

- $\text{COVERS?}(s, I)$: Given a query point s and an interval I , test whether $\bigcup_{\substack{\ell_v \in \mathcal{L} \\ \ell_v \text{ below } s}} I_v$ contains I .

Lemma 7.3. Fix a radius r . Let λ^b be a stabbing path defined for the union of set systems

$$\left\{ \hat{N}_{\leq \sigma}^{r-1}[v] : v \in V_P \right\} \cup \left\{ \hat{N}_T^r[s] : s \in V_P \right\}.$$

A stabbing path data structure \mathcal{D}_T^r (with respect to λ^b) can be constructed in $n^{1+o(1)}$ time, and support each $\text{COVERS?}(s, I)$ query in $n^{o(1)}$ time for any s .

A proof of Lemma 7.3 can be found in Appendix C.3. By the same argument in Lemma 6.1, we can augment \mathcal{D}_T^r to construct the interval representation of $\hat{N}_T^r[s]$ with respect to λ_T^r for every s by calling an interval search query $\text{INTERVALSEARCH}(s)$ to \mathcal{D}_T^r . The readers might have noticed that the stabbing order maintained by the data structure \mathcal{D}_T^r is not the same as the stabbing order we would like to represent the neighborhoods $\hat{N}_T^r[s]$ in. This discrepancy leads to the need for tool that allows the switch between different interval representations.

Switching between interval representations. We gain the ability to switch between different interval representations by computing a special kind of stabbing paths that “respect” some common ρ -sampling of the set system (X, \mathcal{S}) . This requires us to compute stabbing paths not using Lemma 2.10, but something more sophisticated. Ultimately we will be able to shrink from \mathcal{S} to a subcollection \mathcal{S}' of \mathcal{S} , and vice versa. First we set up the terminologies.

We are given a set system (X, \mathcal{S}) with at most $n = |X|$ elements and $m = |\mathcal{S}|$ sets with dual shatter dimension of (X, \mathcal{S}) is d . We fix a *unique ρ -sampling* \mathcal{R} of \mathcal{S} , where each set in \mathcal{S} chosen with probability ρ/m . (Later on we will restrict \mathcal{R} to subcollection \mathcal{S}' of \mathcal{S} and obtain \mathcal{R}' ; we can still think of \mathcal{R}' as obtained from \mathcal{S}' by sampling each element with probability ρ/m , even though we do not explicitly sample from \mathcal{S}' . Notice that the parameter m does not change even if \mathcal{S}' gets smaller.)

Let λ be an ordering of X . We say that a set S *crosses* a pair (x, y) if $x \in S$ and $y \notin S$, or vice versa. The number of consecutive pairs in λ crossed by S is at most twice the size $|\text{Rep}_\lambda(S)|$. For any collection \mathcal{R} , define the *equivalence relation* $\equiv_{\mathcal{R}}$ over X , where $x \equiv_{\mathcal{R}} y$ if and only if no set in \mathcal{R} crosses (x, y) . (In other words, $\{S \in \mathcal{R} : x \in S\} = \{S \in \mathcal{R} : y \in S\}$.) Then $\equiv_{\mathcal{R}}$ has $O(|\mathcal{R}|^d)$ equivalence classes since the dual shatter dimension is at most d .

904 Let \mathcal{S}' be an arbitrary subcollection of \mathcal{S} . Denote the restriction of the unique ρ -sampling \mathcal{R} of \mathcal{S} in
905 \mathcal{S}' as \mathcal{R}' ; in notation, $\mathcal{R}' := \mathcal{S}' \cap \mathcal{R}$. Notice that \mathcal{R}' is also a ρ -sampling. Given any set system (X, \mathcal{S}) , a
906 stabbing path λ of (X, \mathcal{S}) is **\mathcal{R}' -respecting** if each equivalence class of $\equiv_{\mathcal{R}'}$ appears contiguously in λ for
907 the restriction \mathcal{R}' . (The equivalence classes of $\equiv_{\mathcal{R}'}$ are defined with respect to the restriction \mathcal{R}' , not \mathcal{R} .)
908

We compute specialized stabbing paths using the following lemma.

909 **Lemma 7.4.** *Assume the existence of an element reporting oracle that, given $S \in \mathcal{S}$, can enumerate
910 all elements of S in $T_0(n)$ time. Consider a fixed ρ -sampling \mathcal{R} of \mathcal{S} . We can compute the equivalence
911 classes of $\equiv_{\mathcal{R}}$ and construct an \mathcal{R} -respecting stabbing path λ of (X, \mathcal{S}) such that $\sum_{S \in \mathcal{S}} |\text{Rep}_{\lambda}(S)| =$
912 $\tilde{O}(mn/\rho + m\rho^{d-1})$ in $\tilde{O}(T_0(n) \cdot \rho)$ time with high probability. In other words, one can compute a sampled
913 ρ -stabbing path λ of (X, \mathcal{S}) and the equivalence classes of $\equiv_{\mathcal{R}}$ as byproducts.*

914 For the case of unit-disks, we have the element reporting oracle with query time $T_0(n) = \tilde{O}(n)$ by
915 computing a BFS tree using Lemma 2.1. Once both stabbing paths (and their corresponding equivalence
916 classes) were computed, respecting a common ρ -sampling (and its restriction), we can convert one
917 interval representation to the other efficiently.

918 **Lemma 7.5.** *[Conversion of interval representations.] Let (X, \mathcal{S}) be a set system with $|X| \leq n$ and $|\mathcal{S}| \leq m$.
919 Let \mathcal{S}' be a subcollection of \mathcal{S} and \mathcal{T} be a subcollection of \mathcal{S}' . Let \mathcal{R} be the unique ρ -sampling of \mathcal{S} , and
920 \mathcal{R}' be its restriction in \mathcal{S}' . We are given an \mathcal{R} -respecting stabbing path λ of (X, \mathcal{S}) , and an \mathcal{R}' -respecting
921 ordering λ' of (X, \mathcal{S}') (along with the equivalence classes of $\equiv_{\mathcal{R}}$ and $\equiv_{\mathcal{R}'}$).*

922 (1) *[Shrinking from \mathcal{S} to \mathcal{S}' .] Given $\text{Rep}_{\lambda}(S)$ for all $S \in \mathcal{T}$, we can compute $\text{Rep}_{\lambda'}(S)$ for all $S \in \mathcal{T}$ in
923 $\tilde{O}(mn/\rho + m\rho^d)$ total time with high probability.*
924 (2) *[Expanding from \mathcal{S}' to \mathcal{S} .] Given $\text{Rep}_{\lambda'}(S)$ for all $S \in \mathcal{T}$, we can compute $\text{Rep}_{\lambda}(S)$ for all $S \in \mathcal{T}$ in
925 $\tilde{O}(mn/\rho + m\rho^d)$ total time with high probability.*

926 The proof of the two lemmas can be found in Appendix D.

927 **Neighborhood growing algorithm.** Assuming we are equipped with the geometric data structure
928 (Lemma 7.3) and the ability to switch between interval representations with respect to different stabbing
929 paths (Lemma 7.5), we can now formally describe the algorithm.

930 Fix a piece P . For simplicity of the proof, we use \mathcal{S}^r to denote the collection $\{\hat{N}^r[v] : v \in V_P\}$ and
931 \mathcal{S}_M^r to denote $\{\hat{N}_M^r[v] : v \in V_P\}$ for any subset $M \subseteq [1 : \sigma]$. (Recall that the modified balls are defined
932 by intersecting with the relevant region R_P , and thus are dependent on P .) Similarly, we define λ^r be a
933 stabbing path for the (V_G, \mathcal{S}^r) , and λ_M^r be a stabbing path for (V_G, \mathcal{S}_M^r) for any subset $M \subseteq [1 : \sigma]$. Since
934 all the set systems we need here are with respect to the same ground set V_G , we will slightly abuse the
935 notation and use the shorthand \mathcal{S} to denote the set system (V_G, \mathcal{S}) , and use $\mathcal{S}_1 \cup \mathcal{S}_2$ to denote the union
936 of the two set systems $(V_G, \mathcal{S}_1 \cup \mathcal{S}_2)$.

937 The algorithm has an *outer-loop* and an *inner-loop*. The *outer-loop* has $4\Delta + 1$ rounds, iterating
938 over every relevant radii $r \in [\text{ecc}_P - 3\Delta : \text{ecc}_P + \Delta]$; at the start of round r , we maintain the following
939 *invariants* that we have computed,

940 (1) an $\mathcal{R}_{\leq \sigma}^{r-1}$ -respecting ρ -stabbing path $\lambda_{\leq \sigma}^{r-1}$ for the set system $\mathcal{S}_{\leq \sigma}^{r-1}$ and its ρ -sampling $\mathcal{R}_{\leq \sigma}^{r-1}$; and
941 (2) $\lambda_{\leq \sigma}^{r-1}$ -representation of the modified neighborhood balls $\hat{N}_{\leq \sigma}^{r-1}[v]$ for every vertex v in P .

942 For each round with radius r , our algorithm now performs an *inner-loop* by repeating the following
943 steps for σ *iterations*, where T ranges from 1 to σ ; in iteration T we take into account type- T shortest
944 paths using Eq. (5), until we include all σ types and thus finish computing the λ^r -representation of \mathcal{S}^r .
945 At the start of iteration T , we maintain the following *invariants* that we have computed, for each type T ,

946 (i) an $\mathcal{R}_{\leq T-1}^r$ -respecting ρ -stabbing path $\lambda_{\leq T-1}^r$ for the set system $\mathcal{S}_{\leq T-1}^r$ and its ρ -sampling $\mathcal{R}_{\leq T-1}^r$;
 947 (ii) $\lambda_{\leq T-1}^r$ -representation of the modified neighborhood balls $\hat{N}_{\leq T-1}^r[s]$ for every vertex s in $P \setminus \partial P$.

948 At every iteration T , we perform the following steps in order. For the base case when $T = 1$, objects with
 949 the $\leq T-1$ subscript are considered to be null, which we omit from the algorithm.

950 1. Consider the *combined set system* $\mathcal{S}_{\leq \sigma}^{r-1} \cup \mathcal{S}_T^r$. Compute a ρ -sampling \mathcal{R}_T^r of \mathcal{S}_T^r , and take union with
 951 the ρ -sampling $\mathcal{R}_{\leq \sigma}^{r-1}$ of $\mathcal{S}_{\leq \sigma}^{r-1}$ from invariant (1) to form a 2ρ -sampling \mathcal{R}^b of $\mathcal{S}_{\leq \sigma}^{r-1} \cup \mathcal{S}_T^r$.
 952 2. Compute an \mathcal{R}^b -respecting 2ρ -stabbing path λ^b along with the equivalence classes of \equiv_{λ^b} for the
 953 combined set system using Lemma 7.4.
 954 3. Convert the $\lambda_{\leq \sigma}^{r-1}$ -representation of $\hat{N}_{\leq \sigma}^{r-1}[v]$ for every v in P given by invariant (2) into λ^b -
 955 representation, using Lemma 7.5(2) and the fact that $\mathcal{S}_{\leq \sigma}^{r-1}$ is a subcollection of the combined
 956 set system.
 957 4. Compute the geometric data structure \mathcal{D}_T^r with respect to the ρ -stabbing path λ^b using Lemma 7.3.
 958 5. Compute the λ^b -representation of

$$\hat{N}_T^r[s] = \bigcup_{v \in N[s] \cap \text{Cell}_T} \hat{N}_{\leq \sigma}^{r-1}[v]$$

959 for every vertex s in $P \setminus \partial P$ with the help of geometric data structure \mathcal{D}_T^r .

960 6. Convert the λ^b -representation of $\hat{N}_T^r[s]$ for every s in $P \setminus \partial P$ back into λ_T^r -representation, using
 961 Lemma 7.5(1).

962 (At this point, we have successfully computed the spanning path λ_T^r for \mathcal{S}_T^r and its interval representation
 963 with respect to λ_T^r . We now proceed to take union with $\mathcal{S}_{\leq T-1}^r$, currently in $\lambda_{\leq T-1}^r$ -representation.)

964 7. Define the *auxiliary set system*:

$$965 \mathcal{S}_{\leq T-1}^r \cup \mathcal{S}_T^r \cup \mathcal{S}_{\leq T}^r = (V_G, \{\hat{N}_{\leq T-1}^r[v] : v \in V_P\} \cup \{\hat{N}_T^r[v] : v \in V_P\} \cup \{\hat{N}_{\leq T}^r[v] : v \in V_P\}). \quad (6)$$

966 Compute a ρ -sampling $\mathcal{R}_{\leq T}^r$ of $\mathcal{S}_{\leq T}^r$, and take union with the ρ -sampling $\mathcal{R}_{\leq T-1}^r$ of $\mathcal{S}_{\leq T-1}^r$ from
 967 invariant (i) and ρ -sampling \mathcal{R}_T^r of \mathcal{S}_T^r computed in Step 1 to form a 3ρ -sampling \mathcal{R}^\sharp of the auxiliary
 968 set system $\mathcal{S}_{\leq T-1}^r \cup \mathcal{S}_T^r \cup \mathcal{S}_{\leq T}^r$.

969 8. Compute a \mathcal{R}^\sharp -respecting 3ρ -stabbing path λ^\sharp along with the equivalence classes of \equiv_{λ^\sharp} for the
 970 auxiliary set system, using Lemma 7.4.
 971 9. Convert $\lambda_{\leq T-1}^r$ -representation of $\hat{N}_{\leq T-1}^r[s]$ for every s in $P \setminus \partial P$ given by invariant (ii) to λ^\sharp -
 972 representations, using Lemma 7.5(2).
 973 10. Convert λ_T^r -representation of $\hat{N}_T^r[s]$ for every s in $P \setminus \partial P$ given by Step 6 to λ^\sharp -representations,
 974 using Lemma 7.5(2).
 975 11. Compute $\hat{N}_{\leq T}^r[s]$ by taking the union of $\hat{N}_{\leq T-1}^r[s]$ and $\hat{N}_T^r[s]$ as λ^\sharp -representations for every s in
 976 $P \setminus \partial P$. The output $\hat{N}_{\leq T}^r[s]$ is again in the auxiliary set system and thus have λ^\sharp -representation.
 977 12. Compute an $\mathcal{R}_{\leq T}^r$ -respecting ρ -stabbing path $\lambda_{\leq T}^r$ for set system $\mathcal{S}_{\leq T}^r$, by restricting λ^b to $\mathcal{S}_{\leq T}^r$.
 978 Convert the λ^\sharp -representation of $\hat{N}_{\leq T}^r[s]$ for every s in $P \setminus \partial P$ into $\lambda_{\leq T}^r$ -representation, using
 979 Lemma 7.5(1) and the fact that $\mathcal{S}_{\leq T}^r$ is a subcollection of the auxiliary set system.

980 Notice that Step 12 of the algorithm maintains invariants (i) and (ii). After σ iterations, the inner-loop
 981 ends. We perform one extra step:

983 13. Convert the λ^\flat -representation of $\hat{N}_{\leq\sigma}^{r-1}[s]$ for every s in ∂P computed in Step 3 into $\lambda_{\leq\sigma}^r$ -representation,
 984 using Lemma 7.5(1). Insert elements in the difference $\hat{N}_{\leq\sigma}^r[s] \setminus \hat{N}_{\leq\sigma}^{r-1}[s]$ to create $\lambda_{\leq\sigma}^r$ -representation
 985 of $\hat{N}_{\leq\sigma}^r[s]$ for every s in ∂P .

986 We then proceed to the next round of the outer-loop. Notice that invariant (1) follows directly from
 987 invariant (i), and invariant (2) follows from invariants (ii) together with Step 13 from the previous
 988 round.

989 **Handling small pieces.** The most time-consuming part of our algorithm is to compute the ρ -stabbing
 990 paths; each computation takes $\tilde{O}(n\rho)$ time. But we have to compute $O(1)$ many stabbing paths *for*
 991 *each piece and each round of the outer-loop*; there are $\tilde{O}(n/\Delta)$ many pieces (remember the modified
 992 neighborhood balls were defined differently for each piece P), but also for $O(\Delta)$ many rounds. Therefore
 993 the computation of stabbing paths alone already takes $\tilde{O}(n\rho \cdot (n/\Delta) \cdot \Delta) = O(n^2\rho)$ time.

994 To handle the issue, we only apply the above algorithm to *large* pieces whose size is above certain
 995 threshold $A \geq \Delta$. This way, the number of such pieces is at most $\tilde{O}(n/A)$ instead of $\tilde{O}(n/\Delta)$. (We
 996 eventually set $A = \Delta^{O(1)}$.) To compute diameter or eccentricities for small pieces, we use Lemma 2.15.

997 7.3 Analysis for Eccentricities

998 We make the following observations about the shatter dimension of unions of set systems.

999 **Observation 7.6.** Let X be a ground set and let \mathcal{S}_1 and \mathcal{S}_2 be two set systems on X . Let us denote the
 1000 set of ranges obtained by taking unions of ranges from \mathcal{S}_1 and \mathcal{S}_2 by $\hat{\mathcal{S}} := \{S_1 \cup S_2 : S_1 \in \mathcal{S}_1, S_2 \in \mathcal{S}_2\}$.
 1001 Suppose that the shatter dimension of \mathcal{S}_1 is d_1 and the shatter dimension of \mathcal{S}_2 is d_2 . Then the shatter
 1002 dimension of $\mathcal{S}_1 \cup \mathcal{S}_2$ is at most $d_1 + d_2$, and the shatter dimension of $\mathcal{S}_1 \cup \mathcal{S}_2 \cup \hat{\mathcal{S}}$ is also at most $d_1 + d_2$.

1003 The observation shows that the combined set system $\mathcal{S}_T^{r-1} \cup \mathcal{S}_T^r$ has dual shatter dimension 8, because
 1004 both \mathcal{S}_T^{r-1} and \mathcal{S}_T^r individually has dual VC-dimension (and thus dual shatter dimension) at most 4 by
 1005 Lemma 7.2. It also shows that the dual shatter dimension of the auxiliary set system $\mathcal{S}_{\leq T-1}^r \cup \mathcal{S}_T^r \cup \mathcal{S}_{\leq T}^r$ is
 1006 also 8 again as the individual dual VC-dimensions are at most 4 Lemma 7.2. Thus we set $d = 8$ for the
 1007 time analysis that follows.

1008 **Lemma 7.7.** Fix a piece P , a radius r , and some arbitrary parameter ρ . We can maintain invariants (1)
 1009 and (2) between iterations $T-1$ and T , in time $O^*(n \cdot \rho + |P| \cdot (n/\rho + \rho^8))$, by computing (1) an \mathcal{R}^r -
 1010 respecting ρ -stabbing path λ_T^r for the set system \mathcal{S}_T^r for some ρ -sampling \mathcal{R}^r ; and (2) λ_T^r -representation
 1011 of the modified neighborhood balls $\hat{N}_T^r[v]$ for every vertex v .

1012 **Proof:** Steps 1 and 7 take $O(n)$ time to compute ρ -samplings. Steps 2 and 8 take $\tilde{O}(n \cdot \rho)$ time to
 1013 compute $O(\rho)$ -stabbing paths. Steps 3, 6 and 13 take $\tilde{O}(|P| \cdot (n/\rho + \rho^8))$ time to convert interval
 1014 representations, because the combined set system $\mathcal{S}_{\leq\sigma}^{r-1} \cup \mathcal{S}_T^r$ has size $2|P|$. Step 4 takes $O(n^{1+o(1)})$ time
 1015 to compute the geometric data structure \mathcal{D}_T^r . Step 5 takes $O(n \cdot n^{o(1)})$ time to compute the union using
 1016 \mathcal{D}_T^r . Steps 9, 10, 12 take $\tilde{O}(|P| \cdot (n/\rho + \rho^8))$ time to convert interval representations, because the
 1017 auxiliary set system $\mathcal{S}_{\leq T-1}^r \cup \mathcal{S}_T^r \cup \mathcal{S}_{\leq T}^r$ has size $3|P|$. Step 11 takes $O(n)$ time to compute the union of
 1018 two sets with the same λ^\sharp -representation. Overall the neighborhood growing step can be implemented
 1019 in $O^*(n \cdot \rho + |P| \cdot (n/\rho + \rho^8))$ time per piece per radius. \square

1020 To analyze the total running time, we separate the pieces of LDD into large and small, based on
 1021 whether the size of the piece is at least A or not for some parameter A . For large piece P of size
 1022 at least A , we run the ball growing algorithm in Section 7.2. The outer-loop is executed for $O(\Delta)$

1023 rounds, each taking $O^*(n \cdot \rho + |P| \cdot (n/\rho + \rho^8))$ time by Lemma 7.7, followed by Step 13, which takes
 1024 $\tilde{O}(\sum_{s \in \partial P} |\hat{N}_{\leq \sigma}^r[s] \setminus \hat{N}_{\leq \sigma}^{r-1}[s]|)$ time to compute $\hat{N}_{\leq \sigma}^r[s]$ for every s in ∂P using binary search in the
 1025 $\lambda_{\leq \sigma}^r$ -representation. There are $O(n/A)$ large pieces. Overall, for each large piece, this takes time

$$\begin{aligned} 1026 \quad & \tilde{O}\left(\sum_r \sum_{s \in \partial P} |\hat{N}_{\leq \sigma}^r[s] \setminus \hat{N}_{\leq \sigma}^{r-1}[s]|\right) + O(\Delta) \cdot O^*(n \cdot \rho + |P| \cdot (n/\rho + \rho^8)) \\ 1027 \quad & = O^*(|\partial P| \cdot n + \Delta(n \cdot \rho + |P| \cdot (n/\rho + \rho^8))). \end{aligned}$$

1028 By Lemma 2.15, for each small piece P of size less than A , it takes $\tilde{O}(n \cdot |\partial P| + |P| \cdot (|P| + (|\partial P| \Delta)^4))$ time
 1029 to compute eccentricity of all vertices in P . There are at most $\tilde{O}(n/\Delta)$ small pieces.

1030 The final running time of the all-vertex eccentricities algorithm for unit-disk graphs is:

$$\begin{aligned} 1031 \quad & O^*\left(\sum_{P: |P| > A} (|\partial P| \cdot n + \Delta \cdot (n \cdot \rho + |P| \cdot (n/\rho + \rho^8))) + \sum_{P: |P| \leq A} (n \cdot |\partial P| + |P| \cdot (|P| + (|\partial P| \Delta)^4))\right) \\ 1032 \quad & \leq O^*\left(n^2/\Delta + \left(\frac{n}{A} \cdot \Delta \cdot n\rho\right) + \sum_{P: |P| > A} (|P| \cdot (n/\rho + \rho^8)) + \sum_{P: |P| \leq A} (n \cdot |\partial P| + |P|^2) + \sum_{P: |P| \leq A} (|\partial P| \cdot A^4 \Delta^4)\right) \\ 1033 \quad & = O^*(n^2/\Delta + \Delta n^2 \rho/A + \Delta n^2/\rho + \Delta n \rho^8 + n^2/\Delta + nA^2/\Delta + \Delta^3 A^4 n). \end{aligned}$$

1034 Balancing cost by setting parameters $\Delta = n^{1/20}$, $\rho = \Delta^2$ and $A = \Delta^4$ then yields $O^*(n^{2-1/20})$.

1035 **Theorem 7.8.** *Computing eccentricities of an n -node unit-disk graph can be done in $O^*(n^{2-1/20})$ time.*

1036 7.4 Analysis for Diameter

1037 For the special case of computing the diameter of unit-disk graphs, we can get a slight improvement in
 1038 running time by making the following observation:

1039 In the analysis for the all-vertex eccentricities algorithm, computing ρ -stabbing paths in Step
 1040 2 and Step 8 using Lemma 7.4 takes $\tilde{O}(n \cdot \rho)$ time per piece per r , which is the bottleneck.
 1041 We can instead compute a *global* ρ -stabbing path per type for both the combined set system
 1042 and the auxiliary set system at the start of each iteration of the inner-loop, then restrict these
 1043 stabbing paths to each piece P .

1044 More specifically, at the start of iteration T :

1045 0.1. Consider the *global combined set system* $(V_G, \{N_{\leq \sigma}^{r-1}[v] : v \in V_G\} \cup \{N_T^r[v] : v \in V_G\})$. This set
 1046 system differs from $\mathcal{S}_{\leq \sigma}^{r-1} \cup \mathcal{S}_T^r$ in two places: the neighborhood balls are *not* modified, and the ball
 1047 centers range over all vertices in G , not just in P .

1048 Compute a ρ -sampling $\check{\mathcal{R}}_T^r$ of $(V_G, \{N_T^r[v] : v \in V_G\})$, then take union with the ρ -sampling of
 1049 $(V_G, \{N_{\leq \sigma}^{r-1}[v] : v \in V_G\})$ computed from the previous round $r-1$ to form a ρ -sampling $\check{\mathcal{R}}^r$ of the
 1050 global combined set system.

1051 0.2. Compute an $\check{\mathcal{R}}^r$ -respecting ρ -stabbing path $\check{\lambda}^r$ along with the equivalence classes of $\equiv_{\check{\lambda}^r}$ for the
 1052 global combined set system using Lemma 7.4.

1053 0.3. Consider the *global auxiliary set system*

$$1054 \quad (V_G, \{N_{\leq T-1}^r[v] : v \in V_G\} \cup \{N_T^r[v] : v \in V_G\} \cup \{N_{\leq T}^r[v] : v \in V_G\}).$$

1055 Compute a ρ -sampling $\check{\mathcal{R}}_{\leq T}^r$ of $(V_G, \{N_{\leq T}^r[v] : v \in V_G\})$, then take union with the ρ -sampling
 1056 of $(V_G, \{N_{\leq T-1}^r[v] : v \in V_G\})$ computed from the previous iteration $T-1$ and ρ -sampling of
 1057 $(V_G, \{N_T^r[v] : v \in V_G\})$ from Step 0.1 to form a ρ -sampling $\check{\mathcal{R}}^\sharp$ of the global auxiliary set system.

1058 0.4. Compute an $\check{\mathcal{R}}^\sharp$ -respecting ρ -stabbing path $\check{\lambda}^\sharp$ along with the equivalence classes of $\equiv_{\check{\lambda}^\sharp}$ for the
 1059 global auxiliary set system using Lemma 7.4.

1060 Then, for each piece P , we modify the following steps in iteration T :

1061 2. Restrict the $\check{\mathcal{R}}^\flat$ -respecting ρ -stabbing path $\check{\lambda}^\flat$ to another stabbing path λ^\flat for the combined set
 1062 system $\mathcal{S}_{\leq \sigma}^{r-1} \cup \mathcal{S}_T^r$ for piece P . This is done by first removing every neighborhood ball $N_T^r[v]$ not
 1063 centered in P , then taking intersection between $N_T^r[v]$ and the relevant region R_P to form $\hat{N}_T^r[v]$.
 1064 8. Restrict the $\check{\mathcal{R}}^\sharp$ -respecting ρ -stabbing path $\check{\lambda}^\sharp$ to another stabbing path λ^\sharp for the auxiliary set
 1065 system $\mathcal{S}_{\leq \sigma}^{r-1} \cup \mathcal{S}_T^r$ for piece P .

1066 **Analysis.** We only count for the new changes in the diameter case; for the remaining steps, see the
 1067 time analysis for computing eccentricities.

1068 In the new Step 2, the removal of balls not centered in P does not increase the stabbing number
 1069 of λ^\flat . Taking intersection with R_P does not change the stabbing number, because this is equivalent to
 1070 restricting the stabbing path range from $[1 : n]$ to R_P (in the same order), and what was one interval in
 1071 $[1 : n]$ remains one interval in R_P . So λ^\flat is still a ρ -stabbing path. The ρ -sampling \mathcal{R}^\flat can be obtained
 1072 by restricting $\check{\mathcal{R}}^\flat$ to $\mathcal{S}_{\leq \sigma}^{r-1} \cup \mathcal{S}_T^r$. The removal of balls not centered in P only decreases the number of sets
 1073 in consideration and thus makes \equiv_{λ^\flat} coarser than $\equiv_{\check{\lambda}^\flat}$. Taking intersection with R_P does not change the
 1074 status of $\hat{N}_T^r[v]$ being chosen in the sample or not. Thus λ^\flat remains \mathcal{R}^\flat -respecting. As a result, λ^\flat is an
 1075 \mathcal{R}^\flat -respecting ρ -stabbing path.

1076 For the new Step 8, using similar reasoning, λ^\sharp is an \mathcal{R}^\sharp -respecting ρ -stabbing path.

1077 Steps 0.2 and 0.4 take $\tilde{O}(n \cdot \rho)$ time. Steps 2 and 8 now take $O(n)$ time to carry out the restriction.
 1078 (Only the part about restricting $[1 : n]$ to R_P needs to be implemented, not the removal of balls centered
 1079 outside P .) Overall the neighborhood growing step can be implemented in $O^*(n + |P| \cdot (n/\rho + \rho^8))$ time
 1080 per piece, plus another $\tilde{O}(n\rho)$ time across all pieces.

1081 The final running time of the diameter algorithm for unit-disk graphs is:

$$1082 O^* \left(\Delta \cdot n\rho + \sum_{P: |P| > A} \left(|\partial P| \cdot n + \Delta \cdot (n + |P| \cdot (n/\rho + \rho^8)) \right) + \sum_{P: |P| \leq A} \left(n \cdot |\partial P| + |P| \cdot (|P| + (|\partial P| \Delta)^4) \right) \right) \\ 1083 = O^* \left(n^2/\Delta + \Delta n\rho + \Delta n^2/A + \Delta n^2/\rho + \Delta n\rho^8 + n^2/\Delta + nA^2/\Delta + \Delta^3 A^4 n \right).$$

1084 Balancing cost by setting parameters $\Delta = n^{1/18}$ and $\rho = A = \Delta^2$ then yields $O^*(n^{2-1/18})$.

1085 **Theorem 7.9.** *Computing diameter of an n -node unit-disk graph can be done in $O^*(n^{2-1/18})$ time.*

1086 8 Framework for Distance Oracles (and Wiener Index)

1087 Our algorithm for computing the Wiener index is a simple extension of our algorithm for computing the
 1088 exact distance oracle. Therefore, in the following, we focus exclusively on describing the framework for
 1089 computing an exact distance oracle. Then we give more details on how to compute the Wiener index
 1090 with the same running time.

For distance oracles, the first two steps are the same as in the framework for eccentricities described in Section 3. Nonetheless, we present a full description of the framework since there are significant differences in step 3. The key difference is that instead of a specific set of relevant vertices R_P for piece P , we need to consider all vertices V_G . The distances we need to consider will vary depending on the vertex $t \in V_G$, and therefore, we could not use the same definition of $\hat{N}^r[s]$ in the diameter computation for distance oracles. However, we observe that since we have a good additive estimate \hat{d} that is within $\pm \Delta$ of the true distance between $t \in V_G$ and a vertex $s \in P$, we only need to consider distances in an $O(\Delta)$ range around \hat{d} . Our idea is to add a weight to each vertex t and use vertex weights to define $\hat{N}^r[s]$ (Equation (7)).

In our oracle construction, it is important to distinguish between large and small pieces (determined by some size threshold) in the LDD \mathcal{L} . For large pieces, we will use the interval representation. For small pieces, we use the oracle construction of Lemma 2.16.

Oracle construction. Let A be the parameter chosen later. For each piece P in an LDD \mathcal{L} .

1. Compute a low-diameter decomposition \mathcal{L} of G with a diameter parameter $\Delta > 0$.
2. For each vertex $v \in \bigcup_{P \in \mathcal{L}} \partial P$ compute a breath-first search tree in G rooted at v .
3. For each piece $P \in \mathcal{L}$ where $|P| > A$, let s_P be an arbitrary vertex of ∂P . For every vertex $v \in V_G$, we compute and store a weight $w_P(v) = d_G(v, s_P)$. Observe by the triangle inequality that for any vertex $s \in P$:

$$w_P(v) - \Delta \leq d_G(s, v) \leq w_P(v) + \Delta$$

Now we define an adjusted neighborhood ball as follows:

$$\hat{N}^r[s] := \{v \in V : d_G(v, s) \leq r + w_P(v)\} \quad \forall r \in [-\Delta, \Delta] \quad (7)$$

Then we compute $\hat{N}^r[s]$ with the ball expansion data structure \mathcal{D} and store all intermediate balls in the following procedure:

- 3.1 For every $s \in \partial P$, we can explicitly compute the modified balls $\hat{N}^r[s]$ for all $r \in [-\Delta, \Delta]$ as well as compute and store a compact interval representation with respect to an ordering λ .
- 3.2 As a base case, we initialize $\hat{N}^r[s] = \emptyset$ for every $s \in P$ when $r = -\Delta - 1$.
- 3.3 For other values of $r \in [-\Delta, \Delta]$, compute $\text{Rep}_\lambda(\hat{N}^r[s])$ using the inductive formula

$$\hat{N}^r[s] = \bigcup_{v \in N[s]} \hat{N}^{r-1}[v] \quad (8)$$

by taking the union of the intervals.

4. For each piece $P \in \mathcal{L}$ with $|P| < A$, construct the distance oracle of Lemma 2.16.

Correctness of ball expansion initialization. Since $d_G(s_P, t) - \Delta \leq d_G(s, t) \leq d_G(s_P, t) + \Delta$ for every $t \in V$, $t \notin \hat{N}^{-\Delta-1}[s]$ and $t \in \hat{N}^\Delta[s]$. Thus, the initialization is correct, and we have correctly computed the desired modified neighborhood balls.

1124 **Answering queries.** Suppose we get a distance query between a vertex s that is in a piece $P \in \mathcal{L}$, and
 1125 any other vertex $t \in G$. If $|P| < A$, we query the distance oracle for small pieces, and by Lemma 2.16 the
 1126 query time is $O(\log n)$. Otherwise, for any $r \in [-\Delta, \Delta]$, we can detect if $t \in \hat{N}^r[s]$ by checking if t lies in
 1127 an interval of $\text{Rep}_\lambda(\hat{N}^r[s])$ by binary search in $O(\log n)$ time¹⁴. Thus, we can binary search for the first
 1128 radius r_t such that $t \in \hat{N}^{r_t}[s]$ and $t \notin \hat{N}^{r_t-1}[s]$. By the definition of \hat{N} of Equation (7), we can conclude:

$$1129 \quad d_G(s, t) = d_G(t, s_p) + r_t.$$

1130 In either case, we spend $\tilde{O}(1)$ time.

1131 **Computing the Wiener index.** In the oracle construction, we compute and store $\hat{N}^r[s]$ for every s
 1132 in a large piece P . For every vertex $t \in \hat{N}^r[s] \setminus \hat{N}^{r-1}[s]$, the exact distance from s to t is $d_G(t, s_p) + r$,
 1133 and hence $\sum_{t \in \hat{N}^r[s] \setminus \hat{N}^{r-1}[s]} d_G(s, t) = \sum_{t \in \hat{N}^r[s] \setminus \hat{N}^{r-1}[s]} d_G(t, s_p) + |\hat{N}^r[s] \setminus \hat{N}^{r-1}[s]| \cdot r$. This allows us to
 1134 compute $\sum_{v \in V} d_G(s, v)$ in the same running time as it takes to construct the interval representation
 1135 of $\{\hat{N}^r[s]\}_{r=-\Delta}^{\Delta}$. For small pieces, Le and Wulff-Nilsen [LW24] provided an algorithm for computing
 1136 $\sum_{v \in V} d_G(s, v)$ that has the same running time as the construction time for exact oracles of small pieces.
 1137 Therefore, the time to compute the Wiener index is the same as the time to construct an exact distance
 1138 oracle.

1139 **Organization.** In the next four sections, we will apply our framework to devise algorithms for exact
 1140 distance oracles (and thus Wiener index) for different graph classes: sparse graphs of bounded VC-
 1141 dimension (Section 9), arbitrary-square graphs (Section 10), unit-square graphs (Section 11), and
 1142 unit-disk graphs (Section 12). There will be similarities with earlier sections on diameter (Sections 4–7).

1143 9 Distance Oracles for Sparse Graphs of Bounded VC-dimension

1144 We begin by considering sparse graphs of bounded VC-dimension.

1145 **Stabbing path construction.** For a piece $P \in \mathcal{L}$, let $\text{vol}(P) = \sum_{s \in P} \deg(s)$ be the total degree of
 1146 vertices in P , i.e., the *volume* of P . We will construct a stabbing path λ_p for each piece $P \in \mathcal{L}$ satisfying
 1147 $\tilde{O}(1) \cdot \sum_{r=-\Delta}^{\Delta} \sum_{s \in P} \deg(s) \cdot |\text{Rep}_{\lambda_p}(\hat{N}^r[s])| = \tilde{O}(\Delta \text{vol}(P)(n/\rho + \rho^{d-1}))$ for a parameter ρ to be specified
 1148 later using Lemma 7.4.

1149 **Construction time.** Computing the low diameter decomposition and the boundary distances stored in
 1150 step 2 takes $\tilde{O}(mn/\Delta)$ time. For large pieces, the total construction time involves computing the ordering
 1151 λ in $\tilde{O}(m\rho)$ time (by Lemma 7.4) and the ball expansion procedure which takes $O(\Delta \text{vol}(P) \cdot (n/\rho + \rho^{d-1}))$
 1152 time per piece. Thus, the total running time is:

$$1153 \quad \sum_{\substack{P \in \mathcal{L} \\ |V_p| \geq A}} \tilde{O}(m\rho + \Delta \text{vol}(P) \cdot (n/\rho + \rho^{d-1})) = \tilde{O}(nm\rho/A + \Delta mn/\rho + \Delta m\rho^{d-1}).$$

1154 For small pieces, we observe that we can compute a vertex weighted BFS on P with weights at most
 1155 Δ in time $\tilde{O}(\text{vol}(P))$. Therefore, in Lemma 2.16, $T(P) = \tilde{O}(\text{vol}(P))$, giving the total running time for all
 1156 the small pieces:

$$1157 \quad \begin{aligned} \sum_{\substack{P \in \mathcal{L} \\ |V_p| \leq A}} O(n|\partial P| + (|\partial P|^d \Delta^d + |P|) \cdot T(P)) &= \tilde{O}(n^2/\Delta) + \tilde{O}(A^d \Delta^d + A) \cdot \sum_{P \in \mathcal{L}} \text{vol}(P) \\ &= \tilde{O}(n^2/\Delta + mA^d \Delta^d). \end{aligned}$$

¹⁴We can reduce this running time to $O(1)$ by using the fractional cascading technique; this would complicate the details.

1158 The total running time for the algorithm is:

$$1159 \quad \tilde{O}(mn/\Delta + nmp/A + \Delta mn/\rho + \Delta m\rho^{d-1} + mA^d\Delta^d) = O(mn^{1-1/(4d+1)})$$

1160 for $\Delta = n^{1/(4d+1)}$, $\rho = \Delta^2$, and $A = \Delta^3$.

1161 **Space usage.** The boundary distances that we store in step 2 take $\tilde{O}(n^2/\Delta)$ space. For large pieces, in
1162 step 3, we use $\tilde{O}(\Delta \text{vol}(P)(n/\rho + \rho^{d-1}))$ space to store compact representations of the neighborhood
1163 balls and $O(n)$ space to store distances from each vertex to s_p . We also use $O(n)$ space per boundary
1164 vertex to store $\hat{N}^r[s]$ for all $r \in [-\Delta, \Delta]$ in step 3(0) by storing $\{\hat{N}^r[s] \setminus \hat{N}^{r-1}[s]\}$ for every r . Thus, the
1165 total space is:

$$1166 \quad \sum_{\substack{P \in \mathcal{L} \\ |V_P| \geq |A|}} \tilde{O}(n \cdot |\partial P| + \Delta \text{vol}(P)(n/\rho + \rho^{d-1})) = \tilde{O}(n^2/\Delta + \Delta m(n/\rho + \rho^{d-1})) \\ 1167 \quad = \tilde{O}(n^2/\Delta + mn/\Delta + m\Delta^{2d-1}) \quad (\text{since } \rho = \Delta^2).$$

1168 For each small piece, step 4 requires $O(n|\partial P| + |V_P|^d)$ space by Lemma 2.16. The total space required for
1169 all small pieces is

$$1170 \quad \sum_{\substack{P \in \mathcal{L} \\ |V_P| \leq |A|}} O(n|\partial P| + |V_P|^d) = \tilde{O}(n^2/\Delta + nA^{d-1}) = \tilde{O}(n^2/\Delta + n\Delta^{3d-3}) \quad \text{space as } A = \Delta^3.$$

1171 Therefore, the total space of our oracle is:

$$1172 \quad \tilde{O}(n^2/\Delta + mn/\Delta + m\Delta^{2d-1} + n\Delta^{3d-3}) = \tilde{O}(mn^{1-1/(4d+1)})$$

1173 for $\Delta = n^{1/(4d+1)}$.

1174 **Theorem 9.1.** *Given undirected graph G with n vertices and m edges that has generalized distance VC-
1175 dimension at most d , we can construct in $\tilde{O}(mn^{1-1/(4d+1)})$ time an exact distance oracle of $\tilde{O}(mn^{1-1/(4d+1)})$
1176 space and $\tilde{O}(1)$ query time.*

1177 **Remark 9.2.** We chose our parameters to minimize the construction time. We can trade off between
1178 space and query time. In the extreme, if construction time does not matter, we can apply the large piece
1179 solution to all pieces to obtain a distance oracle using $\tilde{O}(mn^{1-1/(2d)})$ space.

1180 **Remark 9.3.** In this section, we assumed the graph has a bounded distance VC-dimension. The exponent
1181 can be slightly optimized when the time it takes to perform BFS in a piece P is $O(|P|)$ instead of $O(\text{vol}(P))$.
1182 This is the case for minor-free graphs, where the space can be improved to $\tilde{O}(mn^{1-1/(4d)})$. We can also
1183 obtain similar results (albeit with worse exponents) if we make other bounds on VC-dimension, such
1184 as the distance VC-dimension, and even if we only assume that the k -neighborhood VC-dimension is
1185 bounded by d for all k .

1186 10 Distance Oracles for Square Graphs

1187 For square graphs, we follow the construction of the oracle in Section 8. We note that the VC-dimension
1188 $d = 4$ in this case. We only analyze the construction time since space is bounded by it.

1189 **Stabbing path construction.** For a piece $P \in \mathcal{L}$. We will construct a stabbing path λ_P for each piece
 1190 $P \in \mathcal{L}$ satisfying:

$$1191 \quad \tilde{O}(1) \cdot \sum_{r=-\Delta}^{\Delta} \sum_{s \in P} |\text{Rep}_{\lambda_P}(\hat{N}^r[s])| = \tilde{O}(\Delta|P|(n/\rho + \rho^3)) \quad (9)$$

1192 for a parameter ρ to be specified later using Lemma 7.4. The running time is $\tilde{O}(\rho n)$ as we show in
 1193 Appendix A that we can find a BFS tree in square graphs in $\tilde{O}(n)$ time.

1194 Given the interval representation $\{\hat{N}^{r-1}[s]\}_{s \in P}$ for radius $r-1$, we compute the interval representation
 1195 of $\{\hat{N}^r[s]\}_{s \in P}$ using the data structure \mathcal{D} for the interval search problem for squares (Lemma 5.2) in
 1196 the eccentricities computation with the same setup: the input contains a set of squares corresponding
 1197 to vertices of P and the interval representation $\{\text{Rep}_{\lambda_P}(\hat{N}^{r-1}[s])\}_{s \in P}$ for radius $r-1$. The queries are
 1198 $\{\text{INTERVALSEARCH}(s) : s \in P\}$ whose outputs are the interval representations of $\{\hat{N}^r[s]\}_{s \in P}$. The total
 1199 time to grow balls for all radii, using the same efficient encoding as in the improved algorithm for
 1200 computing eccentricities in Section 5, is:

$$1201 \quad \begin{aligned} & \sum_{r=-\Delta}^{\Delta} \tilde{O}(b \cdot \sum_{s \in P} (|\text{Rep}_{\lambda_P}(\hat{N}^{r-1}[s])| + |\text{Rep}_{\lambda_P}(\hat{N}^r[s])|)) + \tilde{O}(|P|n/b) \\ &= \tilde{O}(b\Delta|P|(n/\rho + \rho^3) + |P|n/b) \quad (\text{by Equation (9)}). \end{aligned} \quad (10)$$

1202 **Construction time.** For small pieces, we show (in Lemma E.3 in the appendix) that we can compute a
 1203 vertex weighted BFS on P with weights at most Δ in time $\tilde{O}(|P|)$. Therefore, in Lemma 2.16, $T(P) =$
 1204 $\tilde{O}(|P|)$, giving the total running time for all the small pieces:

$$1205 \quad \begin{aligned} \sum_{\substack{P \in \mathcal{L} \\ |V_P| \leq |A|}} \tilde{O}(n|\partial P| + (|\partial P|^4 \Delta^4 + |P|) \cdot |P|) &= \sum_{P \in \mathcal{L}} \tilde{O}(n|\partial P| + A^4 \Delta^4 \cdot |\partial P|) \\ &= \tilde{O}(n^2/\Delta + nA^4 \Delta^3). \end{aligned}$$

1206 For large pieces, the running time to grow balls (Equation (10)) plus the running time of $\tilde{O}(n\rho)$ to
 1207 compute λ_P for each piece P is:

$$1208 \quad \sum_{\substack{P \in \mathcal{L} \\ |V_P| \geq |A|}} \tilde{O}(n\rho + b\Delta|P|(n/\rho + \rho^3) + |P|n/b) = \tilde{O}(n^2\rho/A + b\Delta n(n/\rho + \rho^3) + n^2/\rho).$$

1209 Therefore, the total running time to construct the oracle is:

$$1210 \quad \tilde{O}(n^2/\Delta + nA^4 \Delta^3 + n^2\rho/A + b\Delta n(n/\rho + \rho^3) + n^2/\rho) = \tilde{O}(n^{2-1/20})$$

1211 by setting $b = \Delta = n^{1/20}$, $\rho = \Delta^3$, $A = \Delta^4$.

1212 **Theorem 10.1.** Given a square graph with n vertices, we can construct in $\tilde{O}(n^{2-1/20})$ time an exact
 1213 distance oracle of $\tilde{O}(n^{2-1/20})$ space and $\tilde{O}(1)$ query time.

11 Distance Oracles for Unit-square Graphs

1215 For unit square graphs, we follow the oracle construction for square graphs above. The only difference
 1216 is that we use Lemma 6.1 for solving the interval searching problem for unit squares. Therefore, the
 1217 running time to construct all the intervals is within an $n^{o(1)}$ factor of the total number of intervals. Since
 1218 the stabbing path λ_P for each piece P still satisfies Equation (9), the total running time to grow all

1219 the balls for each large piece is $O^*(\Delta|P|(n/\rho + \rho^3))$. Therefore, the construction time for large pieces
 1220 becomes:

$$1221 \sum_{\substack{P \in \mathcal{L} \\ |V_p| \geq A}} O^*(n\rho + \Delta|P|(n/\rho + \rho^3)) = O^*(n^2\rho/A + \Delta n^2/\rho + \Delta n\rho^3).$$

1222 The construction time for small pieces is the same: $\tilde{O}(n^2/\Delta + nA^4\Delta^3)$. Thus, the total construction time
 1223 of the oracle is:

$$1224 O^*(n^2/\Delta + nA^4\Delta^3 + n^2\rho/A + \Delta n^2/\rho + \Delta n\rho^3) = O^*(n^{2-1/16})$$

1225 for $\Delta = n^{1/16}, A = \Delta^3, \rho = \Delta^2$.

1226 **Theorem 11.1.** *Given a unit square with n vertices, we can construct in $\tilde{O}(n^{2-1/16})$ time an exact
 1227 distance oracle of $\tilde{O}(n^{2-1/16})$ space and $\tilde{O}(1)$ query time.*

12 Distance Oracles for Unit-disk Graphs

1229 For unit-disk graphs, we follow the same strategy in Section 7, adapted to the distance oracle framework
 1230 Section 8.

- 1231 • We partition the neighborhood balls into types, so that within any cell, balls of a fixed type intersect
 1232 the cell as a pseudoline arrangement.
- 1233 • We use the same geometric data structure (Appendix C.3) and interval representation switching
 1234 technique (Appendix D), to implement the ball growing step (Step 3.3) in the framework using
 1235 the inductive formula (1).
- 1236 • We switch to a different distance oracle construction using Lemma 2.16 when the piece has size at
 1237 most A .

1238 We only analyze the construction time since space is bounded by it.

1239 **Ball expansion step.** We now need to deal with modified balls of a fixed radius r for different types
 1240 with vertex weights on their endpoints. To be precise:

$$1241 \hat{N}_M^r[s] := \{(v, w(v)) : v \in V \text{ where the } \tau\text{-walk from } s \text{ to } v \text{ is at most } r + w(v) \text{ for } \tau \in M\}$$

1242 We bound the dual VC-dimension of the set system $((v, w(v))_{v \in V}, \{N_M^r[s]\}_{s \in P})$ in Lemma B.3.

1243 For each type T , given the λ_T^{r-1} -representation for every modified balls in the set system \mathcal{S}_T^{r-1} , we
 1244 compute the λ_T^r -representation for every modified balls in the set system \mathcal{S}_T^r , using the same interval
 1245 representation switching strategy and the data structure \mathcal{D}_T^r for the interval cover problem for unit-disks,
 1246 similar to Section 7.2. The total time to grow balls for all radii is $n \cdot \rho + |P| \cdot (n/\rho + \rho^8)$.

1247 **Construction time.** For small pieces, we show (in Observation E.2 in the appendix) that we can
 1248 compute a vertex weighted BFS on P with weights at most Δ in time $\tilde{O}(|P|)$. Therefore, in Lemma 2.16,
 1249 $T(P) = \tilde{O}(|P|)$, giving the total running time for each small piece to be $n \cdot |\partial P| + |P| \cdot (|P| + (|\partial P| \Delta)^4)$.

1250 For large pieces, we grow the balls for $O(\Delta)$ rounds, each taking time $n \cdot \rho + |P| \cdot (n/\rho + \rho^8)$. Therefore,
 1251 the total running time to construct the oracle is:

$$1252 O^* \left(n^2/\Delta + \sum_{P: |P| > A} \Delta(n \cdot \rho + |P| \cdot (n/\rho + \rho^8)) + \sum_{P: |P| \leq A} (n \cdot |\partial P| + |P| \cdot (|P| + (|\partial P| \Delta)^4)) \right)$$

$$= O^*(n^2/\Delta + \Delta n^2 \rho/A + \Delta n^2/\rho + \Delta n \rho^8 + n^2/\Delta + nA^2/\Delta + \Delta^3 A^4 n).$$

Balancing cost by setting parameters $\Delta = n^{1/20}$, $\rho = \Delta^2$ and $A = \Delta^4$ then yields $\tilde{O}(n^{2-1/20})$.

Theorem 12.1. *Given a unit-disk graph with n vertices, we can construct in $\tilde{O}(n^{2-1/20})$ time an exact distance oracle of $\tilde{O}(n^{2-1/20})$ space and $\tilde{O}(1)$ query time.*

13 Conclusion and Open Questions

In this paper, we have presented the first truly subquadratic algorithms for diameter and related problems for many classes of geometric intersection graphs. Naturally, many open questions follow, for example, improving the exponents of the time bounds of any of our algorithms. More intriguingly:

- Is there a truly subquadratic algorithm for computing the diameter of arbitrary disk graphs? Our algorithm can be extended to the case when the number of different radii is $n^{o(1)}$, but the general case appears more difficult.
- Could we prove any conditional lower bound on the running time of the form $\Omega(n^{1+\delta})$ for computing the diameter of unit-disk graphs? Bringmann *et al.* [BKK⁺22] proved a near-quadratic conditional lower bound for 3D unit-ball graphs under the orthogonal vector (OV) hypothesis.

If one considers more difficult problems than diameter, e.g., counting the number of pairs with shortest-pair distance at most r (which can be solved by our algorithms in subquadratic time), an $\Omega(n^{4/3})$ conditional lower bound follows for unit-disk graphs if one believes certain offline range searching problems similar to Hopcroft’s problem require $\Omega(n^{4/3})$ time (namely, counting the number of pairs of points with Euclidean distance at most 1 in \mathbb{R}^2).

- Is there a near-linear-time algorithm for distinguishing between diameter 2 vs. 3 for unit-disk graphs? Bringmann *et al.* [BKK⁺22] proved a near-quadratic conditional lower bound for 12D unit-hypercube graphs under the hyperclique hypothesis, and obtained an $O(n \log n)$ -time algorithm for unit-square graphs.

There are a few specific open questions related to our algorithms. For example:

- Is the VC-dimension of the set system in Lemma 7.2 bounded when we do not restrict to a fixed r and T ? If so, this might simplify our algorithms for unit disks.
- Could we solve the interval cover data structure problem (Problem 1.3) for arbitrary squares with $N^{1+o(1)}$ preprocessing time and $N^{o(1)}$ query time? If so, this would improve the exponent for our algorithms for arbitrary squares. This appears difficult.
- Less importantly, on the interval cover problem data structure problem for unit disks from a fixed modulo class, could the extra $2^{O(\sqrt{\log N \log \alpha(N)})} \leq N^{o(1)}$ factors be reduced to polylogarithmic? A related question is to determine tight bounds on the combinatorial complexity of the “generalized envelopes” from Appendix C.3.

Besides unit squares, Duraj, Konieczny, and Potępa [DKP24] also considered translates of a convex polygon with constant complexity. It is not difficult to similarly extend our algorithms for unit/arbitrary squares to translates/homothets of other convex polygonal shapes with constant complexity (and our algorithms for unit disks to translations of fat convex non-polygonal shapes with constant complexity).

1290 References

[ACIM99] Donald Aingworth, Chandra Chekuri, Piotr Indyk, and Rajeev Motwani. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM Journal On Computing*, 28(4):1167–1181, 1999. [doi:10.1137/S0097539796303421](https://doi.org/10.1137/S0097539796303421).

[ACM⁺21] 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](https://doi.org/10.1007/s00454-021-00327-y).

[AdT24] Boris Aronov, Mark de Berg, and Leonidas Theobarakos. A clique-based separator for intersection graphs of geodesic disks in R^2 . In *40th International Symposium on Computational Geometry (SoCG)*, volume 293, pages 9:1–9:15, 2024. [doi:10.4230/LIPIcs.SoCG.2024.9](https://doi.org/10.4230/LIPIcs.SoCG.2024.9).

[ADW⁺25] Josh Alman, Ran Duan, Virginia Vassilevska Williams, Yinzhan Xu, Zixuan Xu, and Renfei Zhou. More asymmetry yields faster matrix multiplication. In *36th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2005–2039, 2025. [doi:10.1137/1.9781611978322.63](https://doi.org/10.1137/1.9781611978322.63).

[AE99] 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.

[AKPW95] Noga Alon, Richard M. Karp, David Peleg, and Douglas West. A graph-theoretic game and its application to the k -server problem. *SIAM Journal on Computing*, 24(1):78–100, 1995. [doi:10.1137/s0097539792224474](https://doi.org/10.1137/s0097539792224474).

[AS00] Pankaj K. Agarwal and Micha Sharir. Davenport-Schinzel sequences and their geometric applications. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 1–47. North-Holland, Amsterdam, 2000. [doi:10.1016/B978-044482537-7/50002-4](https://doi.org/10.1016/B978-044482537-7/50002-4).

[Awe85] Baruch Awerbuch. Complexity of network synchronization. *Journal of The ACM*, 32(4):804–823, 1985. [doi:10.1145/4221.4227](https://doi.org/10.1145/4221.4227).

[BCE15] Drago Bokal, Sergio Cabello, and David Eppstein. Finding all maximal subsequences with hereditary properties. In *31st International Symposium on Computational Geometry (SoCG)*, volume 34 of *LIPIcs*, pages 240–254, 2015. [doi:10.4230/LIPIcs.SOCG.2015.240](https://doi.org/10.4230/LIPIcs.SOCG.2015.240).

[BDG⁺14] Michael J. Bannister, William E. Devanny, Michael T. Goodrich, Joseph A. Simons, and Lowell Trott. Windows into geometric events: Data structures for time-windowed querying of temporal point sets. In *26th Canadian Conference on Computational Geometry (CCCG)*, 2014. URL: <http://www.cccg.ca/proceedings/2014/papers/paper02.pdf>.

[BKK⁺22] Karl Bringmann, Sándor Kisfaludi-Bak, Marvin Künemann, 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.

[BKK⁺24] Alexander Baumann, Haim Kaplan, Katharina Kloß, Kristin Knorr, Wolfgang Mulzer, Liam Roditty, and Paul Seiferth. Dynamic connectivity in disk graphs. *Discrete & Computational Geometry*, 71(1):214–277, 2024. [doi:10.1007/s00454-023-00621-x](https://doi.org/10.1007/s00454-023-00621-x).

[BNW22] Aaron Bernstein, Danupon Nanongkai, and Christian Wulff-Nilsen. Negative-weight single-source shortest paths in near-linear time. In *63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 600–611, 2022. [doi:10.1109/focs54457.2022.00063](https://doi.org/10.1109/focs54457.2022.00063).

[BT15] 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](https://doi.org/10.1016/j.disc.2015.05.026).

[Cab18] 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.

[CEV07] 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](https://doi.org/10.1007/s00454-006-1260-0).

1336 [CG86a] Bernard Chazelle and Leonidas J Guibas. Fractional cascading: I. A data structuring technique.
 1337 *Algorithmica*, 1(1-4):133–162, 1986.

1338 [CG86b] Bernard Chazelle and Leonidas J Guibas. Fractional cascading: II. Applications. *Algorithmica*,
 1339 1(1-4):163–191, November 1986.

1340 [CGL⁺23] Panagiotis Charalampopoulos, Paweł Gawrychowski, Yaowei Long, Shay Mozes, Seth Pettie, Oren
 1341 Weimann, and Christian Wulff-Nilsen. Almost optimal exact distance oracles for planar graphs.
 1342 *Journal of the ACM*, 70(2):1–50, 2023. [doi:10.1145/3580474](https://doi.org/10.1145/3580474).

1343 [CGL24] Hsien-Chih Chang, Jie Gao, and Hung Le. Computing diameter+2 in truly-subquadratic time for unit-
 1344 disk graphs. In *40th International Symposium on Computational Geometry (SoCG)*, pages 38:1–38:14,
 1345 2024. [doi:10.4230/LIPIcs.SoCG.2024.38](https://doi.org/10.4230/LIPIcs.SoCG.2024.38).

1346 [Cha86] Bernard Chazelle. Filtering search: A new approach to query-answering. *SIAM J. Comput.*, 15(3):703–
 1347 724, 1986. [doi:10.1137/0215051](https://doi.org/10.1137/0215051).

1348 [CHN20] Timothy M. Chan, Qizheng He, and Yakov Nekrich. Further results on colored range searching.
 1349 In *36th International Symposium on Computational Geometry (SoCG)*, volume 164 of *LIPIcs*, pages
 1350 28:1–28:15, 2020. [doi:10.4230/LIPIcs.SOCG.2020.28](https://doi.org/10.4230/LIPIcs.SOCG.2020.28).

1351 [CHP19] Timothy M. Chan, John Hershberger, and Simon Pratt. Two approaches to building time-
 1352 windowed geometric data structures. *Algorithmica*, 81(9):3519–3533, 2019. [doi:10.1007/S00453-019-00588-3](https://doi.org/10.1007/S00453-019-00588-3).

1354 [CJ15] Sergio Cabello and Miha Jejčič. Shortest paths in intersection graphs of unit disks. *Computational
 1355 Geometry*, 48(4):360–367, 2015. [doi:10.1016/j.comgeo.2014.12.003](https://doi.org/10.1016/j.comgeo.2014.12.003).

1356 [CK97] V. Chepoi and S. Klavžar. The Wiener index and the Szeged index of benzenoid systems in linear
 1357 time. *Journal of Chemical Information and Computer Sciences*, 37(4):752–755, 1997. [doi:10.1021/ci9700079](https://doi.org/10.1021/ci9700079).

1359 [CK09] Sergio Cabello and Christian Knauer. Algorithms for graphs of bounded treewidth via orthogonal
 1360 range searching. *Computational Geometry*, 42(9):815–824, 2009. [doi:10.1016/j.comgeo.2009.02.001](https://doi.org/10.1016/j.comgeo.2009.02.001).

1362 [CS16] Timothy M. Chan and Dimitrios Skrepetos. All-pairs shortest paths in unit-disk graphs in slightly
 1363 subquadratic time. In *27th International Symposium on Algorithms and Computation (ISAAC)*,
 1364 volume 64, pages 24:1–24:13, 2016.

1365 [CS19] Timothy M. Chan and Dimitrios Skrepetos. All-pairs shortest paths in geometric intersection graphs.
 1366 *J. Comput. Geom.*, 10(1):27–41, 2019. [doi:10.20382/JOCG.V10I1A2](https://doi.org/10.20382/JOCG.V10I1A2).

1367 [CW89] Bernard Chazelle and Emo Welzl. Quasi-optimal range searching in spaces of finite VC-dimension.
 1368 *Discrete & Computational Geometry*, 4(5):467–489, 1989. [doi:10.1007/BF02187743](https://doi.org/10.1007/BF02187743).

1369 [dBCvKO08] Mark de Berg, Otfried Cheong, Marc J. van Kreveld, and Mark H. Overmars. *Computational Geometry:
 1370 Algorithms and Applications*. Springer, 3rd edition, 2008. [doi:10.1007/978-3-540-77974-2](https://doi.org/10.1007/978-3-540-77974-2).

1371 [DHV22] Guillaume Ducoffe, Michel Habib, and Laurent Viennot. Diameter, eccentricities and distance oracle
 1372 computations on h -minor free graphs and graphs of bounded (distance) Vapnik–Chervonenkis
 1373 dimension. *SIAM Journal on Computing*, 51(5):1506–1534, 2022. [doi:10.1137/20M136551X](https://doi.org/10.1137/20M136551X).

1374 [dKMT23] Mark de Berg, Sándor Kisfaludi-Bak, Morteza Monemizadeh, and Leonidas Theobarakos.
 1375 Clique-based separators for geometric intersection graphs. *Algorithmica. An International Journal in
 1376 Computer Science*, 85(6):1652–1678, 2023.

1377 [DKP24] Lech Duraj, Filip Konieczny, and Krzysztof Potępa. Better diameter algorithms for bounded VC-
 1378 dimension graphs and geometric intersection graphs. In *32nd Annual European Symposium on
 1379 Algorithms (ESA)*, volume 308, pages 51:1–51:18, 2024.

1380 [EIK01] Alon Efrat, Alon Itai, and Matthew J Katz. Geometry helps in bottleneck matching and related
 1381 problems. *Algorithmica*, 31(1):1–28, 2001.

1382 [Eri96] Jeff Erickson. New lower bounds for Hopcroft’s problem. *Discrete & Computational Geometry*,
 1383 16(4):389–418, 1996. [doi:10.1007/bf02712875](https://doi.org/10.1007/bf02712875).

1384 [For87] Steven Fortune. A sweepline algorithm for Voronoi diagrams. *Algorithmica*, 2:153–174, 1987.
 1385 [doi:10.1007/BF01840357](https://doi.org/10.1007/BF01840357).

1386 [Fre87] Greg N. Frederickson. Fast algorithms for shortest paths in planar graphs, with applications. *SIAM
 1387 Journal on Computing*, 16(6):1004–1022, 1987. [doi:10.1137/0216064](https://doi.org/10.1137/0216064).

1388 [GJRS18] Prosenjit Gupta, Ravi Janardan, Saladi Rahul, and Michiel H. M. Smid. Computational geometry:
 1389 Generalized (or colored) intersection searching. In *Handbook of Data Structures and Applications*,
 1390 chapter 67, pages 1042–1057. CRC Press, 2nd edition, 2018. URL: <https://www-users.cs.umn.edu/~sala0198/Papers/ds2-handbook.pdf>.

1391 [GJS95] Prosenjit Gupta, Ravi Janardan, and Michiel H. M. Smid. Further results on generalized intersection
 1392 searching problems: Counting, reporting, and dynamization. *Journal of Algorithms*, 19(2):282–317,
 1393 1995. [doi:10.1006/jagm.1995.1038](https://doi.org/10.1006/jagm.1995.1038).

1394 [GKM⁺21] Paweł Gawrychowski, Haim Kaplan, Shay Mozes, Micha Sharir, and Oren Weimann. Voronoi diagrams
 1395 on planar graphs, and computing the diameter in deterministic $\tilde{O}(n^{5/3})$ time. *SIAM Journal on
 1396 Computing*, 50(2):509–554, 2021. [doi:10.1137/18M1193402](https://doi.org/10.1137/18M1193402).

1397 [Har11] Sariel Har-Peled. *Geometric Approximation Algorithms*. Number 173. American Mathematical Soc.,
 1398 2011.

1400 [Kle89] Rolf Klein. *Concrete and Abstract Voronoi Diagrams*, volume 400 of *Lecture Notes in Computer Science*.
 1401 Springer, 1989. [doi:10.1007/3-540-52055-4](https://doi.org/10.1007/3-540-52055-4).

1402 [Klo23] Katharina Kloost. An algorithmic framework for the single source shortest path problem with
 1403 applications to disk graphs. *Computational Geometry*, 111:101979, 2023.

1404 [KRSV08] Haim Kaplan, Natan Rubin, Micha Sharir, and Elad Verbin. Efficient colored orthogonal range
 1405 counting. *SIAM Journal on Computing*, 38(3):982–1011, 2008. [doi:10.1137/070684483](https://doi.org/10.1137/070684483).

1406 [KZ25] Adam Karczmarz and Da Wei Zheng. Subquadratic algorithms in minor-free digraphs: (Weighted)
 1407 distance oracles, decremental reachability, and more. In *36th Annual ACM-SIAM Symposium on
 1408 Discrete Algorithms (SODA)*, pages 4338–4351, 2025. [doi:10.1137/1.9781611978322.147](https://doi.org/10.1137/1.9781611978322.147).

1409 [LP19] Jason Li and Merav Parter. Planar diameter via metric compression. In *51st Annual ACM Symposium
 1410 on Theory of Computing (STOC)*, pages 152–163, Phoenix, AZ, USA and New York, NY, USA, 2019.

1411 [LPS⁺24] Daniel Lokshtanov, Fahad Panolan, Saket Saurabh, Jie Xue, and Meirav Zehavi. A 1.9999-
 1412 approximation algorithm for vertex cover on string graphs. In *40th International Symposium
 1413 on Computational Geometry (SoCG)*, volume 293, pages 72:1–72:11, 2024. [doi:10.4230/LIPIcs.SoCG.2024.72](https://doi.org/10.4230/LIPIcs.SoCG.2024.72).

1414 [LW22] Hung Le and Christian Wulff-Nilsen. Optimal approximate distance oracle for planar graphs. In
 1415 *62nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 363–374, 2022.

1416 [LW24] Hung Le and Christian Wulff-Nilsen. VC set systems in minor-free (di)graphs and applications.
 1417 In *35th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 5332–5360, 2024.
 1418 [doi:10.1137/1.9781611977912.192](https://doi.org/10.1137/1.9781611977912.192).

1419 [Pet15] Seth Pettie. Sharp bounds on Davenport-Schinzel sequences of every order. *Journal of The ACM*,
 1420 62(5), 2015. [doi:10.1145/2794075](https://doi.org/10.1145/2794075).

1421 [RW13] Liam Roditty and Virginia Vassilevska Williams. Fast approximation algorithms for the diameter and
 1422 radius of sparse graphs. In *45th Annual ACM Symposium on Theory of Computing (STOC)*, pages
 1423 515–524, 2013.

1424 [Sau72] N. Sauer. On the density of families of sets. *J. Comb. Theory Ser. A*, 13(1):145–147, 1972. [doi:10.1016/0097-3165\(72\)90019-2](https://doi.org/10.1016/0097-3165(72)90019-2).

1427 [Sei95] Raimund Seidel. On the all-pairs-shortest-path problem in unweighted undirected graphs. *Journal*
1428 *of Computer and System Sciences*, 51(3):400–403, 1995. [doi:10.1006/jcss.1995.1078](https://doi.org/10.1006/jcss.1995.1078).

1429 [She72] Saharon Shelah. A combinatorial problem; stability and order for models and theories in infinitary
1430 languages. *Pacific Journal of Mathematics*, 41(1):247–261, 1972. [doi:10.2140/pjm.1972.41.247](https://doi.org/10.2140/pjm.1972.41.247).

1431 [SJ05] Qingmin Shi and Joseph JaJa. Novel transformation techniques using q-heaps with applications
1432 to computational geometry. *SIAM Journal On Computing*, 34(6):1474–1492, 2005. [doi:10.1137/S0097539703435728](https://doi.org/10.1137/S0097539703435728).

1433 [ST86] Neil Sarnak and Robert E. Tarjan. Planar point location using persistent search trees. *Communications*
1434 *of The ACM*, 29(7):669–679, July 1986. [doi:10.1145/6138.6151](https://doi.org/10.1145/6138.6151).

1435 [WN09] Christian Wulff-Nilsen. Wiener index and diameter of a planar graph in subquadratic time. In
1436 *Proceedings of the 25th European Workshop on Computational Geometry*, pages 25–28, 2009.

1437 [WN09]

1438 A Low-diameter Decompositions

1439 In this section, we construct the low-diameter decomposition of sparse graphs and geometric intersection
 1440 graphs. Recall that we use $N^r[v] = \{u : d_G(u, v) \leq r\}$ to denote the set of vertices in the neighborhood
 1441 ball of radius r centered at v .

1442 We give an algorithm for computing a low-diameter decomposition as claimed in Section 2.1. Our
 1443 low-diameter decomposition for graphs is perhaps most similar to the low-diameter decomposition
 1444 in [AKPW95]; we are not aware of any work stating the exact guarantees with our definition of LDD.
 1445 We first present the general algorithm and the properties of the LDD. We will discuss the detailed
 1446 implementation and running time for sparse graphs and geometric intersection graphs separately.

1447 **Basic algorithm.** Let ϕ be a parameter with $0 < \phi < 1$. We will choose $\phi := 24 \log n / \Delta$. Start with
 1448 the entire graph $G_1 = G$. Pick an arbitrary vertex $u \in V$, perform a BFS to compute neighborhood balls
 1449 centered at v : $N^1[v], N^2[v], \dots, N^\ell[v]$. Stop when $|N^\ell[v]| / |N^{\ell-2}[v]| \leq 1 + \phi$, and set $V_1 = N^{\ell-1}[v]$
 1450 as one piece in the decomposition. Note that this is guaranteed to eventually happen because when
 1451 $N^{\ell-2}[v]$ is the entire connected component of v , then $N^\ell[v] = N^{\ell-1}[v] = N^{\ell-2}[v]$. We mark the vertices
 1452 in $N^\ell[v] \setminus N^{\ell-2}[v]$ as boundary vertices¹⁵. Repeat this procedure on $G_2 = G_1 \setminus V_1$ to find V_2 , then on
 1453 $G_3 = G_2 \setminus V_2$, and so on.

1454 **Low diameter property.** First we bound the strong diameter of one such ball $N^{\ell-1}[v]$ that we included
 1455 in our low diameter decomposition. Observe that for all $2 \leq i < \ell$, the ball $N^i[v]$ of radius i satisfies
 1456 $|N^i[v]| > (1 + \phi)|N^{i-2}[v]|$. Since the size of the largest ball is at most n , if ℓ is odd, we have that:

$$1457 n \geq |N^\ell(v)| \geq |N^{\ell-1}(v)| > (1 + \phi) \cdot |N^{\ell-3}(v)| > (1 + \phi)^{(\ell-1)/2} \cdot |N^0(v)| = (1 + \phi)^{(\ell-1)/2}$$

1458 Taking logarithms on both sides, and using the fact that $\phi < 1$, we obtain:

$$1459 \log n > \frac{\ell - 1}{2} \cdot \log(1 + \phi) \geq \frac{\ell - 1}{2} \cdot (\phi - \phi^2/2) > \frac{(\ell - 1) \cdot \phi}{4} = \frac{6(\ell - 1) \cdot \log n}{\Delta}$$

1460 Rearranging the inequality yields $\ell \leq \Delta/6$. The diameter is at most $2\ell \leq \Delta/3$.

1461 **Small boundary property.** Let $N_{G_1}^{r_1}[v_1]$ be the ball of largest radii we compute in G_1 , $N_{G_2}^{r_2}[v_2]$ the
 1462 ball in G_2 , $\dots, N_{G_k}^{r_k}[v_k]$ the ball in G_k . Observe that when we choose $P_i = G[N_{G_i}^{r_i-1}[v_i]]$, the vertices
 1463 of $N_{G_i}^{r_i-1}[v_i] \setminus N_{G_i}^{r_i-2}[v_i]$ are potentially boundary vertices ∂P_i , and a vertex in $N_{G_i}^{r_i}[v_i] \setminus N_{G_i}^{r_i-1}[v_i]$ is a
 1464 boundary vertex in ∂P_j for some piece P_j with $j > i$. Assume we have a total of k pieces in the LDD.
 1465 Thus it can be seen that:

$$\begin{aligned} 1466 \sum_{i=1}^k |\partial P_i| &\leq \sum_{i=1}^k |N_{G_i}^{r_i}[v_i] \setminus N_{G_i}^{r_i-2}[v_i]| \\ 1467 &\leq \sum_{i=1}^k \phi \cdot |N_{G_i}^{r_i-2}[v_i]| && (\text{since } |N_{G_i}^{r_i}[v_i]| \leq (1 + \phi)|N_{G_i}^{r_i-2}[v_i]|) \\ 1468 &\leq \sum_{i=1}^k \phi \cdot |N_{G_i}^{r_i-1}[v_i]| && (\text{since } N_{G_i}^{r_i-2}[v_i] \subseteq N_{G_i}^{r_i-1}[v_i]) \\ 1469 &\leq \phi \cdot n = 24n \log n / \Delta. \end{aligned}$$

¹⁵The vertices in $N^\ell[v] \setminus N^{\ell-1}[v]$ are boundary vertices of later pieces constructed in the process. The vertices in $N^{\ell-1}[v] \setminus N^{\ell-2}[v]$ are boundary vertices of V_1 , although there may be other boundary vertices in $N^{\ell-2}(v)$ that we accounted for earlier in the process.

1470 **No small pieces.** To ensure that no piece is small, we will do some post-processing of the pieces
 1471 obtained from the basic algorithm. We use the following claim.

1472 **Lemma A.1.** *Let P_i be a piece found in the basic LDD algorithm found by taking the vertices $N^{r-1}[v]$ in
 1473 G_i . Either P_i has at least $\Omega(\Delta/\log n)$ vertices, or P_i is an entire connected component of G_i .*

1474 **Proof:** Suppose that $|N^r(v)| > |N^{r-2}(v)|$. Then as $(1 + \phi) \cdot |N^{r-2}(v)| \geq |N^r(v)| \geq |N^{r-2}(v)| + 1$, we
 1475 conclude that $|N^{r-2}(v)| \geq 1/\phi = O(\Delta/\log n)$, so P has size at least $O(\Delta/\log n)$. Otherwise $|N^r(v)| =$
 1476 $|N^{r-1}(v)| = |N^{r-2}(v)|$ and thus P_i is an entire connected component of G_i . \square

1477 In our post-processing, we will merge P_i with an arbitrary neighboring component. Observe that since P_i
 1478 is an entire connected component of G_i , no later piece P_j with $j > i$ will merge into P_i . Now consider a
 1479 piece P_j with multiple pieces $P_{i_1}, P_{i_2}, \dots, P_{i_t}$ merging into it in the post-processing step, $j < i_1, \dots, j < i_t$.
 1480 Since all pieces have diameter at most $\Delta/3$, the resulting merged P_j has diameter at most Δ .

1481 A.1 Sparse Graphs

1482 Consider the standard BFS algorithm that computes $N^r[v]$ by adding all neighbors incident to $N^{r-1}[v]$
 1483 into a queue. For every vertex v , the basic algorithm will add all its neighbors into a queue at most
 1484 once, so the basic algorithm can be implemented in $O(m + n)$ time. The post-processing step involving
 1485 merging components can also be done in $O(m + n)$ time.

1486 **Theorem 2.2.** *Let G be a graph with n vertices and m edges. For any parameter $24 \log n < \Delta \leq n$, we
 1487 can compute a low-diameter decomposition for G in $O(m + n)$ time.*

1488 A.2 Geometric Intersection Graphs

1489 Here we consider geometric intersection graphs of fat pseudo-disks of similar size and squares of varying
 1490 sizes. A family of objects are called *pseudo-disks* if each one is the interior of a simple closed Jordan
 1491 curve and two objects are either disjoint, have one object fully inside the other, or properly intersect
 1492 each other at two boundary points. Disks are by definition pseudo-disks. The geometric intersection
 1493 graph of a family of pseudo-disks can be considered, combinatorially, as a set of vertices V representing
 1494 the pseudo disks and two vertices are connected if their corresponding pseudo-disks have non-empty
 1495 intersection. For the algorithm below, we consider fat pseudo-disks that are of roughly the same size and
 1496 have constant complexity. Specifically, a fat pseudo-disk is sandwiched between two disks of the same
 1497 center p of radius r and R with two fixed constants r, R and $r \leq R$ and the boundary can be described by
 1498 a constant number of algebraic curves. We call this pseudo-disk centered at p as C_p . The input to our
 1499 algorithm consists of the description of a family of n fat pseudo-disks with input size $O(n)$. We assume
 1500 that one can compute in $O(1)$ time whether two pseudo-disks have an edge or not. The geometric
 1501 intersection graph of such pseudo-disks can be dense (i.e., having edges of size $\Theta(n^2)$). We show that
 1502 the low diameter decomposition can still be computed in near linear time, similar to the running time
 1503 for sparse graphs (Appendix A.1).

1504 Recall that the basic idea is to perform BFS from a vertex v to compute balls centered at v : $N^0[v] = \{v\}$,
 1505 $N^1[v], N^2[v], \dots, N^\ell[v]$, and stop when $|N^\ell[v]|/|N^{\ell-2}[v]| \leq 1 + \phi$. Let $V_1 = N^{\ell-1}[v]$. Then repeat this
 1506 procedure on $G_2 = G_1 \setminus V_1$ to find V_2 , then on $G_3 = G_2 \setminus V_2$, and so on.

1507 We have to be careful in implementing the basic idea: we do not want to spend $\tilde{O}(n)$ time per
 1508 iteration as the number of iterations could be $\Omega(n)$. This is achievable by not explicitly constructing
 1509 all the edges, an idea that is generally adopted for computing a breadth-first search tree for geometric
 1510 intersection graphs (for fat objects of similar sizes) [EIK01, CJ15, CS16]. We use the same algorithm

as in [CGL24] for pseudo-disks of similar sizes. The core step in the BFS is to find the vertices that are exactly j -hops away from v , denoted by Y_j – from the vertices that are exactly $j - 1$ hops away from v , Y_{j-1} . Put a grid of size $r\sqrt{2}$. Two pseudo-disks with centers in the same grid cell are connected by an edge for sure. Thus, if a pseudo disk centered at p in one cell appears in Y_{j-1} , all pseudo-disks centered in the same cell will be included in Y_j if they are not yet covered in $B^{j-1}(v)$. In addition, the other vertices to be included in Y_j will come from cells that have distance at most $2R$ away from cells touched by Y_{j-1} . Since R^2/r^2 is a constant, we only need to check for each cell touched by Y_{j-1} , at most a constant number of nearby cells. This step can be implemented by using an operation called the red-blue intersection problem, which finds all the blue pseudo-disks that intersect at least one red pseudo-disks, where all red pseudo-disks and blue pseudo-disks are separated by a horizontal (or a vertical) line. We use the following lemma from [CGL24].

Lemma A.2 ([CGL24]). *In time $O(n_b \log n_b + n_r \alpha(n_r) \log n_r + n_r 2^{\alpha(n_r)})$, we can solve the red-blue intersection problem of n_r pseudo-disks and n_b blue pseudo-disks. Here $\alpha(n)$ is the inverse Ackermann function.*

With this Lemma we can conclude the following theorem.

Theorem A.3. *Let G be the intersection graphs of n fat pseudo-disks of similar size. For any parameter $24 \log n < \Delta \leq n$, we can compute a low-diameter decomposition for G in $\tilde{O}(n)$ time.*

Proof: We argue that for the entire algorithm, a non-empty cell in the grid of size $r\sqrt{2}$ is only visited a constant number of times. First, if cell c has a vertex $p \in V_i$ and p is not on the boundary of this piece V_i , then all pseudo-disks centered in the cell will be included in V_i . After V_i is removed, cell c becomes empty and will not be visited again in later iterations. Therefore, a cell c visited by V_i is only visited again by pieces V_j with $j > i$ if c has only vertices that are at the boundary of V_i . That says, the cell c has a neighboring cell c' (within distance $2R$ from c), such that c' contains a vertex p of V_i and p is not on the boundary of V_i . In this case all vertices in c' are entirely in V_i (or earlier pieces). Thus, after the i -th iteration, at least one of the neighboring cells of c is wiped out. Since c has only a constant of such neighboring cells, c is only visited a constant number of times. This finishes the argument. \square

As a corollary, since unit squares and unit disks are fat pseudo-disks of similar size, we conclude that a low diameter decomposition can be computed for these classes of intersection graphs in $\tilde{O}(n)$ time.

Axis-aligned squares. We will need a similar theorem for axis-aligned squares (which might not be of similar size.) A BFS on the intersection graph of axis-parallel squares can be done in time $O(n \log n)$ [Klo23], by using data structures developed in [BKK⁺24]. Again we focus on how to find the objects of j -hops away from a starting vertex v from the objects of $j - 1$ hops away. When the squares have different sizes, instead of a grid of a single size, one can use a hierarchical structure such as the (compressed) quadtree. Each square is associated with a quad whose size is comparable with its size. Further the compressed quadtree can be decomposed into $O(n)$ canonical paths such that each root to leaf path can be represented by $O(\log n)$ disjoint canonical paths. A canonical path has a smallest cell σ and largest cell τ , and is associated with a constant number of regions, classified as inner, middle and outer regions. The inner region is a disk centered at the smallest cell σ of the canonical path. Further, each region A is associated with two sets, the first type $S_1(A)$ contains a collection of objects centered inside A that form a clique, and the second type $S_2(A)$ contains objects that intersect with at least one site in $S_1(A)$. A similar red-blue intersection problem can be solved in linear time for axis-parallel squares, assuming sorting along x and Y coordinates is performed already, as shown in [Klo23]. In summary, to implement a BFS step, for each region A touched by the vertices in Y_{i-1} , include all objects that are

1554 in $S_1(A)$ and then perform red-blue intersection modules with A against a constant number of other
 1555 regions. Since each object stays in at most $O(\log n)$ sets of the first type and at most $O(\log n)$ sets of the
 1556 second type, the total running time carries an extra $O(\log n)$ factor. We can use this algorithm for the
 1557 low-diameter decomposition and obtain the following.

1558 **Theorem A.4.** *Let G be the intersection graphs of n axis-aligned squares. For any parameter $24\log n <$
 1559 $\Delta \leq n$, we can compute a low-diameter decomposition for G in $\tilde{O}(n)$ time.*

1560 **Proof:** The same argument as in Theorem A.3 applies here: for each region A , either the type one objects
 1561 $S_1(A)$ are completely included in a piece V_i and this region disappears; or, one of the (constantly many)
 1562 nearby regions are completely included in V_i and disappears. By a charging argument, each region is
 1563 only touched a constant number of times. Thus the total running time is in the order of $\tilde{O}(n)$. \square

1564 B VC-dimension Lemma

1565 In this section, we prove a lemma bounding the VC-dimension of certain set systems (Lemma 7.2) from
 1566 Section 7, which is needed in our algorithms for unit-disk graphs.

1567 Let G be the geometric intersection graphs of unit-disk graphs. Let M be a subset of vertices, called a
 1568 *type*. We say that a walk W from a vertex v to a vertex u is a *Type-1 M -walk* if the vertex preceding u
 1569 (the second to last vertex) in the walk is in M . We say that the walk is a *Type-2 M -walk* if the vertex
 1570 following v (the second to first vertex) in the walk is in M .

1571 For a technical reason explained later, we assume that every vertex in G has a self-loop attached to it.
 1572 For every vertex v , define:

$$1573 \begin{aligned} B_M^{(1)}(v, r) &= \{u \in V \mid \text{there is a Type-1 } M\text{-walk from } v \text{ to } u \text{ of length exactly } r\} \\ B_M^{(2)}(v, r) &= \{u \in V \mid \text{there is a Type-2 } M\text{-walk from } v \text{ to } u \text{ of length exactly } r\} \end{aligned} \quad (11)$$

1574 The reason for attaching a self-loop to every vertex is that if $d_G(v, x) \leq r - 1$ for some vertex x in M ,
 1575 then $x \in B_M^{(1)}(v, r)$ since we can make a Type-1 M -walk of length r by traversing from v to x along the
 1576 shortest path (of length at most $r - 1$) and then along the self-loop to get a walk of length at most r . The
 1577 second to last vertex of the walk is x itself, which is in M . Furthermore, if there is a Type-1 M -walk from
 1578 v to u of length less than r , then there is a Type-1 M -walk from v to u of length exactly r by traveling
 1579 the self-loop attached to the vertex in M preceding u . The same holds for Type-2 M -walk. The main
 1580 result of this section is to show that the system of balls deriving from Type-1 M -walk has a bounded
 1581 VC-dimension.

1582 **Lemma B.1.** *$(V, \{B_M^{(1)}(v, r)\}_{r \in \mathbb{R}, v \in V})$ has VC-dimension at most 4.*

1583 Observe that Type-1 and Type-2 M -walks are dual to each other: a Type-1 M -walk from v to u of
 1584 length r is a Type-2 M -walk from u to v of length r . Therefore, for a given r , $(V, \{B_M^{(2)}(v, r)\}_{v \in V})$ is the
 1585 dual set system of $(V, \{B_M^{(1)}(v, r)\}_{v \in V})$, and therefore, has VC-dimension at most $2^4 = 16$. By modifying
 1586 the proof of Lemma B.1, get an improved bound for balls from Type-2 M -walk:

1587 **Lemma B.2.** *For any $r \in \mathbb{N}$, $(V, \{B_M^{(2)}(v, r)\}_{v \in V})$ has VC-dimension at most 4.*

1588 The set system in Lemma B.2 only includes balls of fixed radius. It is possible that the more general set
 1589 system $(V, \{B_M^{(2)}(v, r)\}_{r \in \mathbb{R}, v \in V})$, which includes all balls of all radii, has VC-dimension at most 4. However,
 1590 for a technical reason, our proof of Lemma B.1 does not extend to this general case. See Remark B.5 for
 1591 more details.

1592 For the distance oracle application, we will need to handle vertices with weights. So we define the
 1593 following set system for weighted vertices. Suppose each vertex u has a weight $w(u)$, and the ground set
 1594 is $\{(u, w(u))\}_{u \in V}$. Recall for the distance oracle, we maintain the adjusted neighborhood ball as follows
 1595 (see Section 8 Equation (7), copied below):

$$1596 \hat{N}^r[s] := \{v \in V : d_G(v, s) \leq r + w_P(v)\} \quad \forall r \in [-\Delta, \Delta]$$

1597 Further, the path connecting s (the center of the neighborhood ball) to v has the second vertex (adjacent
 1598 to s) of a special type. Thus, we consider a Type-2 walk from s to v . Therefore the set system we work
 1599 with will be $(\{(u, w(u))\}_{u \in V}, \{B_{M,w}^{(2)}(v, r)\}_{v \in V})$ where

$$1600 B_{M,w}^{(2)}(s, r) = \{v \in V \mid \text{there is a Type-2 } M\text{-walk from } s \text{ to } v \text{ of length exactly } r + w(v)\} \quad (12)$$

1601 Take this set system as the primal system, we can define the dual system as follows. Specifically,
 1602 $v \in B_{M,w}^{(2)}(s, r)$ if and only if $s \in B_{M,w}^{(1)}(v, r)$ where

$$1603 B_{M,w}^{(1)}(v, r) = \{s \in V \mid \text{there is a Type-1 } M\text{-walk from } v \text{ to } s \text{ of length exactly } r + w(v)\} \quad (13)$$

1604 Notice that $B_{M,w}^{(1)}(v, r) = B_M^{(1)}(v, r + w(v))$. Therefore, the VC-dimension bound we need is provided
 1605 precisely by Lemma B.1 for Type-1 walks, which fortunately works for neighborhood balls of varying
 1606 radii. With this we immediately have the following.

1607 **Lemma B.3.** *For any $r \in \mathbb{N}$, $(\{(u, w(u))\}_{u \in V}, \{B_{M,w}^{(2)}(v, r)\}_{v \in V})$ has dual VC-dimension at most 4.*

1608 B.1 Type-1 M -Walks

1609 In this section, we prove Lemma B.1. As all M -walks in this section are of Type-1, we will drop the prefix
 1610 Type-1, and only refer to Type-1 M -walks as M -walk. We also call the last edge of an M -walk to u as an
 1611 ***M-edge***.

1612 **Proof (Sketch Proof of Lemma B.1):** The strategy is basically the same as [CGL24]¹⁶. We only show
 1613 the steps needed for adapting the proof here. Consider four vertices a, b, c, d representing four disks
 1614 D_a, D_b, D_c, D_d and assume that there are two (Type 1) M -walks $P(b, a)$ from b to a and $P(c, d)$ from c
 1615 to d . (The vertices preceding a and d in the walks are in M .) We define a ***local crossing pattern*** to be
 1616 four distinct vertices a', b', c', d' with a', b' on $P(a, b)$ (with a' closer to a than b') and c', d' on $P(c, d)$
 1617 (with c' closer to c than d') such that one of the four vertices a', b', c', d' has edges to all the other three
 1618 vertices; see Figure 3. The central claim is the following; if the claim holds, then the rest of the argument
 1619 is standard.

1620 **Claim B.4.** *Either there is an M -walk $P'(c, a)$ whose hop length is at most $|P(c, d)|$ or there is an M -walk
 1621 $P'(b, d)$ whose hop length is at most $|P(b, a)|$.*

1622 We consider a case study depending on whether the local crossing pattern involves an M -edge. In
 1623 the first case when the local crossing pattern does not involve the last edge (from a vertex in M to the
 1624 endpoint of the walk) of the two M -walks (see Figure 3 (a) for an example), the proof follows exactly the
 1625 same as that of [CGL24]. The second case, which is also easy, is when the local crossing pattern involves
 1626 two M -edges. In this case, both $c', b' \in M$. Either we have the edge $c'a'$ or the edge $b'd'$. In both cases
 1627 the claim is true. Figure 3 (b) shows the case with edge $c'a'$ present. In this case, we can find an M -walk

¹⁶We refer to <https://arxiv.org/pdf/2401.12881.pdf>.

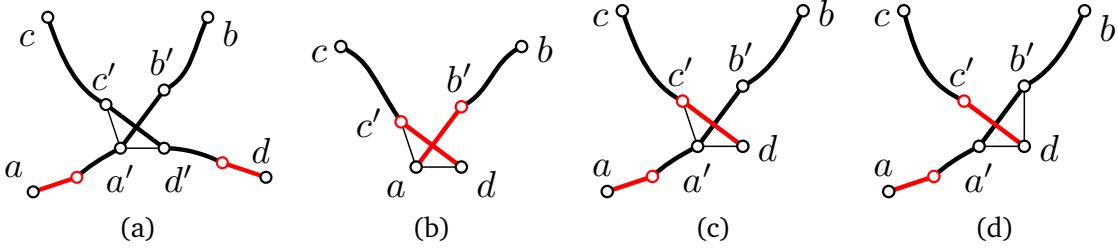


Figure 3. If two M -walk $P(b, a)$ and $P(c, d)$ intersect with a local crossing pattern a', b', c', d' , then there is an M -walk from c to a that are no longer than $|P(c, d)|$ or there is an M -walk from b to d that is no longer than $|P(b, a)|$. The vertices in M are highlighted red.

1628 $P'(c, a)$ through $P(c, c')$ and then take edge $c'a$, which is not longer than path $P(c, d)$. If $b'd$ is present,
1629 then the path that follows $P(b, b')$ and then edge $b'd$ is an M -walk and not longer than $P(b, a)$.

1630 The difficult case for the proof of the claim is when the local crossing pattern happens at an M -edge of
1631 one path with the non M -edge part of the other path. Without loss of generality, assume that the M -edge
1632 involved in a local crossing pattern is the edge $c'd$. See case (c) and (d) in Figure 3 for an example.

1633 We first consider the case when $c'a'$ and $a'd$ are present. We prove by contradiction. Consider an
1634 M -path from b to d :

$$P'(b, d) = P(b, b') \circ (b'a') \circ (a'c') \circ (c'd').$$

1635 Since the claim does not hold, the following holds:

$$|P'(b, d)| > |P(b, a)| \Leftrightarrow |P(b, b')| + 3 > |P(b, b')| + 1 + |P(a, a')| \Leftrightarrow |P(a', a)| < 2.$$

1638 If $|P(a', a)| = 0$, then $a = a'$ and b' must be a vertex in M . This is a contradiction, as the crossing
1639 occurs between two M -edges.

1640 If $|P(a, a')| = 1$, then $a' \in M$. This means $a'a$ is an M -edge. Now we define another walk $Q(b, d) =$
1641 $P(b, b') \circ (b'a') \circ (a'd)$. Since $a' \in M$, $Q(b, d)$ is an M -walk and, furthermore, $|Q(b, d)| = |P(b, a)|$. Thus,
1642 $Q(b, d)$ is the M -walk that satisfies the claim. Now, if $Q(b, d)$ is still longer than $P(b, a)$, by the same
1643 analysis we have $|P(a', a)| < 1$. This leads to a contradiction.

1644 The next case we consider is where the edges $a'd, b'd$ are present. See Figure 3 (d). Consider an
1645 M -path from b to d :

$$P'(b, d) = P(b, b') \circ (b'd) \circ (dc') \circ (c'd).$$

1647 Notice that this is an M -walk. If it is longer than $|P(b, a)|$, we have

$$|P'(b, d)| > |P(b, a)| \Leftrightarrow |P(b, b')| + 3 > |P(b, b')| + 1 + |P(a, a')| \Leftrightarrow |P(a', a)| < 2.$$

1649 For the same reason as explained earlier, $|P(a', a)| = 0$ is not possible and $|P(a', a)| = 1$ means that
1650 $a' \in M$ and we now find an M -walk $Q(b, d)$ by following $P(b, b')$ and then edges $b'a'$ and $a'd$. This
1651 path is one shorter than $P'(b, d)$ and this again gives a contradiction.

1652 The other two cases, when either edges $c'b', c'a'$ or edges $c'b', b'd$ are present, are easy. Basically
1653 the edge $b'c'$ provides an M -walk from b to d which is not longer than $P(b, a)$. \square

1654 **Remark B.5.** If we apply the same proof to Type-2 M -walks, Claim B.4 remains true. However, what we
1655 need is a slightly different version: Either there is a Type-2 M -walk $P'(a, c)$ whose hop length is at most
1656 $|P(a, b)|$ or there is a Type-2 M -walk $P'(d, b)$ whose hop length is at most $|P(d, c)|$. The proof does not
1657 extend to show this version. On the other hand, if we fix a radius r , then everything goes through; see
1658 the next section.

1659 **B.2 Type-2 M -Walks**

1660 M -walks in this section are referred to Type-2 M -walks.

1661 **Proof (Proof of Lemma B.2):** We follow the same setup in the proof of Lemma B.1. Assume that there
1662 are two (Type 2) M -walks $P(a, b)$ from a to b and $P(d, c)$ from d to c . (We switch the roles of a and b ,
1663 and of c and d , so that we can reuse Figure 3.) The following claim implies the lemma:

1664 **Claim B.6.** *Either there is an M -walk $P'(a, c)$ whose hop length is at most $|P(a, b)|$ or there is an M -walk
1665 $P'(d, b)$ whose hop length is at most $|P(d, c)|$.*

1666 Observe that $|P(a, b)|$ and $|P(d, c)|$ have length exactly r each since they are from balls of radius
1667 exactly r . Therefore, $|P(a, b)| = |P(d, c)|$, and hence Claim B.6 follows directly from Claim B.4. \square

1668 **C Geometric Data Structures**

1669 In this section, we describe how to solve the interval searching problem (Problem 2.11), the main
1670 geometric data structure problem used by our diameter algorithms and distance oracles, for different
1671 types of geometric objects. In Appendix C.1, we first describe how to reduce the interval cover to the
1672 rainbow colored intersection searching (Problem 2.13) and then describe how to reduce the interval
1673 searching problem to the interval cover problem (Problem 1.3), though with some loss of efficiency.
1674 For squares, we solve the rainbow colored intersection searching problem in Appendix C.2. For unit
1675 disks of a fixed modulo class and for unit disks, we solve the interval cover problem directly (without
1676 going through rainbow colored searching), and thus more efficiently, in Appendix C.3 and Appendix C.4
1677 respectively, using an interesting recursive approach.

1678 **C.1 Reductions Between Data Structure Problems**

1679 In this subsection, we provide the reductions between the data structure problems in Section 2.4, and in
1680 particular, proving Lemma 2.12 and Lemma 2.14.

1681 **Interval cover to rainbow colored intersection searching.** We reduce the interval cover problem
1682 (Problem 1.3) to the rainbow colored intersection searching problem (Problem 2.13).

1683 **Lemma 2.14.** *If we can construct in $\tilde{O}(|\mathcal{O}_{RC}|)$ time a data structure \mathcal{D}_{RC} with $\tilde{O}(1)$ query time for solving
1684 Problem 2.13, then for any parameter $b \in [1, n]$, we can construct a data structure \mathcal{D}_{IC} for solving
1685 Problem 1.3 that has total run time $\tilde{O}(N_{IC} \cdot b + L_{IC}/b)$.*

1686 **Proof:** Consider an instance of Problem 1.3. Divide the range $[1, n]$ into n/b blocks of length b , denoted
1687 by intervals $B_1, \dots, B_{n/b}$, with $[1, n] = \bigcup_{k=1}^{n/b} B_k$. Denote by \mathcal{S} the set of all intervals associated with
1688 objects in \mathcal{O} and $\mathcal{S}(q)$ the set of intervals associated with objects intersecting q . Consider the query
1689 interval I . It intersects with a set of blocks B_i, \dots, B_j such that I overlaps with at most two of the
1690 blocks partially, namely, the two blocks at the end (B_i or B_j), and fully contains all the middle chunks
1691 B_{i+1}, \dots, B_{j-1} . To verify if the union of the intervals of $\mathcal{S}(q)$ covers the query interval I , we need to check
1692 for each of the blocks, B_k , $i \leq k \leq j$, if $B_k \cap I$ is covered by the union of the intervals of $\mathcal{S}(q)$, limited
1693 within block B_k . If for each B_k the answer is true, we answer Yes. Otherwise, we answer No. In the
1694 following we focus on answering the coverage query for a fixed block B and check if the union of the
1695 intervals $\{I_s \cap B \mid I_s \in \mathcal{S}(q)\}$ covers $I \cap B$.

1696 Now fix a block B . Take $I' = I \cap B$. Similarly, for each object s we restrict the interval I_s within B
 1697 and take $I'_s = I_s \cap B$ and take $\mathcal{S}' = \{I'_s \mid I_s \in \mathcal{S}\}$. Each interval I_s fully covers a set of middle chunks and
 1698 only partially covers at most two extreme blocks at the end of I_s . Thus we can write $\mathcal{S}' = \mathcal{S}'_1 \cup \mathcal{S}'_2$, with
 1699 the first category \mathcal{S}'_1 containing the intervals $I'_s = B$ (i.e., I_s fully covers B) and the second category \mathcal{S}'_2
 1700 containing the intervals $I'_s \neq B$ (i.e., I_s partially covers B). We perform two queries for I' against \mathcal{S}'_1 and
 1701 \mathcal{S}'_2 respectively.

1702 For \mathcal{S}'_1 , all the intervals are given the same color and we just check if at least one of them is associated
 1703 with an object intersecting q . We solve this problem by issuing RAINBOWCOVER?(q) against the objects
 1704 whose intervals appear in \mathcal{S}'_1 . If this rainbow query returns a positive answer, I' is covered and we are
 1705 done. Otherwise, we check for I' against \mathcal{S}'_2 . This query is more complicated since the intervals $I'_s \in \mathcal{S}'_2$
 1706 only partially cover B . We give each of the elements in B a unique color. There are at most b colors. Also,
 1707 for each object s with $I'_s \in \mathcal{S}'_2$, we make a colored copy of the object s for each element in I'_s ; the color
 1708 of the copy is equal to the color of the corresponding element. Now we discuss the case when $I' = B$
 1709 and when $I' \subset B$ separately. When $I' = B$, i.e., B is an ‘internal’ block, we build a rainbow colored query
 1710 structure for all color/elements in B and issue a query RAINBOWCOVER?(q) to see if all colors show up.
 1711 If the query returns no, we return negative to the interval cover query. In the case when I' is a boundary
 1712 block ($I' \subset B$), we build the rainbow colored query structure for each color/element in B . To answer the
 1713 query for I' , we issue RAINBOWCOVER?(q) for each color in I' against the corresponding data structure to
 1714 see if this color appears among objects that intersect q . The total number of such queries is the number
 1715 of elements in I' and is at most b . If all the rainbow queries return True – that all colors in I' appear –
 1716 then all elements in I' are covered by the union of intervals in \mathcal{S}'_2 for those objects intersecting q . If any
 1717 rainbow query returns a no, we return a negative answer for the interval cover query.

1718 To analyze the total running time, we need to account for the preprocessing time and the total query
 1719 time for all the n/b blocks. Recall that we solve the interval cover problem in an ‘off-line’ version and
 1720 assume all input intervals and query intervals are given. N_{IC} is the total number of input objects and
 1721 query objects, and L_{IC} is the total length of the input and query intervals. We issue a total of $O(L_{IC}/b)$
 1722 rainbow queries in the first category since for each query (q, I) we only consider the blocks that overlap
 1723 with I . For the second category we issue a total of $O(L_{IC}/b)$ rainbow queries for the blocks that are
 1724 internal to the interval queries and $O(N_{IC}b)$ rainbow queries for the boundary blocks. Thus the total
 1725 query time is $\tilde{O}(N_{IC}b + L_{IC}/b)$.

1726 For the preprocessing time, we consider the time spent to prepare for the rainbow query in the first
 1727 and second category separately. For the second category, we have a total of $2N_{IC}$ intervals since each
 1728 interval I_s of an input object s only contributes at most two boundary intervals. Each interval generates
 1729 at most b colored objects so we have a total of $O(N_{IC}b)$ objects, over all the n/b blocks. We build the
 1730 rainbow colored query data structure for each block separately. The total preprocessing time for rainbow
 1731 query in the second category is thus $\tilde{O}(N_{IC}b)$. For the rainbow query in the first category, we perform a
 1732 linear scan of the blocks and only update the rainbow query data structure \mathcal{D}_{RC} when needed – an input
 1733 object appears (starts to fully cover a new block) or disappears (stops covering the current block). Each
 1734 input object only triggers two updates. In fact, for each update, we simply rebuild the rainbow query
 1735 data structure from scratch. For each input interval I_s , the amortized run time attributed to I_s in these
 1736 preprocessing and rebuilding efforts is $\tilde{O}(|I_s|/b)$ and therefore the total running time remains $\tilde{O}(L_{IC}/b)$,
 1737 where L_{IC} is the total length of the input and query intervals. Therefore, the total run time is bounded
 1738 by $\tilde{O}(N_{IC}b + L_{IC}/b)$. This finishes the proof. \square

1739 We also need a data structure that can answer interval avoidance queries. Specifically,

1740 **Problem C.1 (Interval Avoidance Problem).** *Given a set of N objects \mathcal{O} and each object $o \in \mathcal{O}$ is
 1741 associated with an interval $I_o \subseteq [1 : n]$. Design a data structure to answer the following query:*

- **AVOIDS?** (q, I) : Given a query object q and a query interval $I \subseteq [1 : n]$, decide whether the union of intervals associated with the objects intersecting¹⁷ q in \mathcal{O} is disjoint from the interval I .

The interval avoidance problem is easier than the interval cover problem, as it is decomposable – we can partition the input objects into two sets and check the query (q, I) against each set for avoidance separately.

Lemma C.2. *If we can construct in $\tilde{O}(|\mathcal{O}_{RC}|)$ time a data structure \mathcal{D}_{RC} with $\tilde{O}(1)$ query time for solving Problem 2.13, then we can construct a data structure \mathcal{D}_{IA} for solving Problem C.1 that has a preprocessing time of $\tilde{O}(N_{IC})$ such that each interval avoidance query takes time $\tilde{O}(1)$.*

Proof: An interval I_o intersects I if either at least one endpoint of I_o is inside I or one endpoint of I is inside I_o . Therefore, to answer the interval avoidance query, we run two types of queries. In the first type we verify if I includes any endpoints of intervals whose associated objects intersect q . If yes, we immediately return no to the interval avoidance query. If not, we proceed to the second type of queries where we check if an endpoint of I stabs any intervals whose associated objects intersect q . The first type is a range query, and the second type is an interval stabbing query. We explain the two operations separately.

For the range query, we take the set \mathcal{S} of all intervals associated with objects in \mathcal{O} and build a binary tree \mathcal{T} on all the $2|\mathcal{S}|$ endpoints of the intervals. Further, for each node v on the tree \mathcal{T} we build a rainbow colored query structure on the objects in \mathcal{O} whose associated intervals have at least one endpoint in the subtree of v . In particular, the data structure at the root of \mathcal{T} includes all objects in \mathcal{O} . The total preprocessing time for these query data structures is $\tilde{O}(N_{IC})$, since each object in \mathcal{O} only appears in $O(\log N_{IC})$ of the rainbow colored query structures. Next we run a standard range query with I on tree \mathcal{T} to find a set $Q(I)$ of $O(\log |\mathcal{S}|)$ vertices of \mathcal{T} such that each vertex $v \in Q(I)$ has the entire subtree fully inside I , but its parent does not meet this condition. We issue a query of q on the rainbow colored structure at each vertex in $Q(I)$. If any query returns a positive answer (indicating intersection), then I does not avoid the objects intersecting q . We issue at most $O(\log N_{IC})$ rainbow colored range queries with a total cost of $\tilde{O}(1)$.

For the interval stabbing query, we build an interval tree on the intervals \mathcal{S} . Specifically, we have a binary tree \mathcal{Y} where the root r is associated with value $\ell(r) = \lfloor n/2 \rfloor$ (the median of $[1, n]$) as well as a subset of intervals $S(r)$ – all the intervals in \mathcal{S} that are stabbed by $\ell(r)$. Recursively, we build the left (right) subtree by using all the intervals to the left (right) of $\ell(r)$ respectively. Further, for each node v in the interval tree, we build two binary trees, $\mathcal{Z}_1(v)$ on the left endpoints of the intervals in $S(v)$ (that are all smaller than or equal to $\ell(v)$), and $\mathcal{Z}_2(v)$ on the right endpoints of the intervals in $S(v)$ (that are all greater than or equal to $\ell(v)$). For each node u on a tree $\mathcal{Z}_i(v)$, $i = 1, 2$, we build a rainbow colored query structure for all the intervals in the subtree of u . Again, these objects are given the same color.

The total preprocessing time for these query data structures is $\tilde{O}(N_{IC})$, since each interval in \mathcal{S} only appears in the set $S(v)$ of one vertex v on tree \mathcal{Y} and then at most $O(\log N_{IC})$ vertices in the secondary level trees $\mathcal{Z}_i(v)$.

Next we take one endpoint p of I and issue a stabbing query on \mathcal{Y} . We first issue stabbing query against the root vertex r of \mathcal{Y} and depending on whether p is less than or greater than $\ell(r)$, recursively query either the left subtree or the right subtree of \mathcal{Y} . We just explain how to query p against a node v of \mathcal{Y} . The total query cost is just an extra log factor more. Specifically, if $p \leq \ell(v)$, we issue a query to $\mathcal{Z}_1(v)$; if $p \geq \ell(v)$, we issue a query to $\mathcal{Z}_2(v)$. Suppose we query p on $\mathcal{Z}_1(v)$. The other case is symmetric. We take all the vertices $Z_1(p)$ of $\mathcal{Z}_1(v)$: $u \in Z_1(p)$ if all vertices in the subtree of u are completely to the left of p but u 's parent fails to meet this condition. $|Z_1(p)| = O(\log N_{IC})$. Now we query q again the

¹⁷Here we mean the objects intersect, not their associated intervals.

rainbow colored range query structure for all vertices in $Z_1(p)$. If any of these queries return a Yes, the interval avoidance query is negative. In total the query cost adds a total factor of $O(\log^2 N_{IC})$ on top of the cost of a single rainbow colored range query.

In summary, we only add extra poly-logarithmic factors on top of the rainbow colored range query structure. Thus, we can implement the interval avoidance query with a preprocesing time of $\tilde{O}(N_{IC})$ such that each interval avoidance query takes time $\tilde{O}(1)$. \square

Interval searching to interval cover. We reduce the interval searching problem (Problem 2.11) to the interval cover problem (Problem 1.3) with polylogarithmic loss.

Lemma 2.12. *If one can construct a data structure \mathcal{D}_{IC} for solving Problem 1.3 with total run time $T(N_{IC}, n, L_{IC})$ (for some polynomial function T), then we can construct a data structure \mathcal{D}_{IS} for solving Problem 2.11 in total run time $\tilde{O}(T(N_{IS}, n, L_{IS}))$. Furthermore, if \mathcal{D}_{IC} has preprocessing time $P(N)$ and query time $Q(N)$, then \mathcal{D}_{IS} has preprocessing time $\tilde{O}(P(\tilde{N}_{IS}))$ and query time $\tilde{O}(Q(\tilde{N}_{IS}) \cdot |\mathcal{I}_{out}(q)|)$ where $\tilde{N}_{IS} := \sum_{o \in \mathcal{O}_{IS}} |\mathcal{I}(o)|$ is the total number of input intervals and $\mathcal{I}_{out}(q)$ is the set of output intervals from the interval search query of q to \mathcal{D}_{IS} .*

Proof: We take an instance of Problem 2.11. For each object $s \in \mathcal{O}_{IS}$ we duplicate it to k copies if s is associated with k intervals. Each copy is now associated with a single interval of \mathcal{I}_o . This creates a total of $\tilde{N}_{IS} = \sum_{o \in \mathcal{O}_{IS}} |\mathcal{I}(o)|$ objects. Now we build a data structure \mathcal{D}_{IC} to solve Problem 1.3 on this set of objects with preprocessing time $\tilde{O}(P(\tilde{N}_{IS}))$. For each query INTERVALSEARCH(q), we recursively issue queries to \mathcal{D}_{IC} . Specifically, we start with $I = [1, n]$. If I is completely covered by the union of the intervals associated with objects in \mathcal{O}_{IS} that intersect q (which is checked by a query to \mathcal{D}_{IC} with q and I), we output I and we are done. Otherwise, if I is completely avoided, we output \emptyset and we are also done. For the other case, we will recurse. We divide I into two intervals of equal length, I_1 and I_2 , and issue queries (q, I_1) and (q, I_2) with \mathcal{D}_{IC} . In the end, we will output the union of all the intervals that are fully covered by the intervals associated with objects in \mathcal{O}_{IS} that intersect q .

The running time for a query q is dependent on the total number of queries issued to \mathcal{D}_{IC} recursively. Notice that all query intervals are dyadic intervals. In addition, recursion stops when an interval I is completely covered by the union of intervals $\mathcal{S}(q)$ or completely avoided. Thus only the dyadic intervals whose parent partially overlaps with a query output interval will ever trigger a query. The total number of such intervals is in the order of $O(|\mathcal{I}_{out}(q)| \cdot \log n)$. Recall that each query to \mathcal{D}_{IC} takes time $Q(\tilde{N}_{IS})$. Summing up everything, we have the claim in the Lemma. \square

C.2 Data Structure for Square Graphs

We now solve the rainbow colored intersection searching problem (Problem 2.13) for a set of axis-parallel squares of possibly different size. We use an approach that can be commonly found in previous work on colored range searching [GJS95, GJRS18]: for each color class, we build a set of new objects, so that colored range searching reduces to standard range searching on all the new objects.

Consider the input squares as being in the plane $z = 0$ in 3D. For each square s of center $(x, y, 0)$ and side length $2r$ (or ℓ_∞ -radius r) consider the point $a_s = (x, y, -r) \in \mathbb{R}^3$ and the cone C_s with apex a_s whose intersection with the plane $z = 0$ is the square s . (If we imagine the z axis pointing vertically up, then this cone opens upward.) See Figure 4 for an example. For a collection S of squares and the corresponding cones, consider a new square q with center $(x_q, y_q, 0)$ and side length $2r_q$. Notice that the normals of the planes bounding any cone C_s is the intersection of four upper half-spaces, and the normals of these upper-half-spaces are $(1, 0, 1), (0, 1, 1), (0, -1, 1)$ and $(-1, 0, 1)$.

Observation C.3. *The square q intersects s if and only if $q := (x_q, y_q, r_q) \in C_s$.*

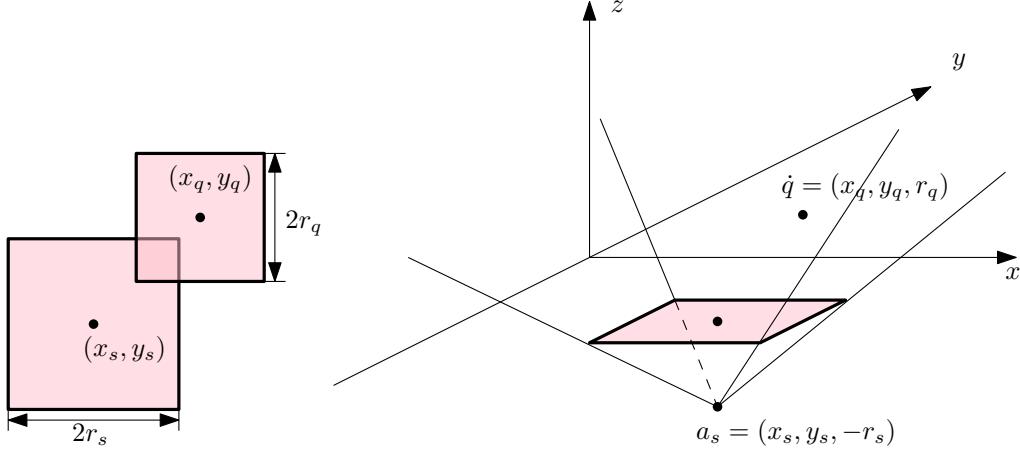


Figure 4. Left: a square centered at (x_s, y_s) with side length $2r_s$ and a query square centered at (x_q, y_q) with side length $2r_q$. Right: the cone C_s and point $\dot{q} = (x_q, y_q, r_q)$.

Proof: For a point $(x_q, y_q, 0)$ outside s we have that its ℓ_∞ distance to the square ∂s is $\min(|x_s - x_q|, |y_s - y_q|) - r_s$. On the other hand, the vertical line through $(x_q, y_q, 0)$ intersects the cone at exactly

$$(x_q, y_q, \max(|x_s - x_q|, |y_s - y_q|) - r_q),$$

that is, the signed vertical distance from $(x_q, y_q, 0)$ to ∂C_s is equal to the ℓ_∞ distance from $(x_q, y_q, 0)$ to s . In particular, the square of ℓ_∞ radius r_q centered at $(x_q, y_q, 0)$ intersects s if and only if (x_q, y_q, r_q) is above ∂C_s , i.e., if and only if $(x_q, y_q, r_q) \in C_s$. \square

As a consequence of the above observation, the square q intersects some square among some set S of squares if and only if $\dot{q} \in \bigcup_{s \in S} C_s$. In particular, if we have a convenient data structure to represent $\bigcup_{s \in S} C_s$, then we can quickly answer the query: given an axis-aligned square q , does it intersect at least one square from S ?

Detecting intersection with some square from S . We will now work on representing $U_S := \bigcup_{s \in S} C_s$. Observe that U_S is the union of translates of a fixed convex cone of constant complexity, thus it has linear union complexity. Indeed, each face f of ∂U_S is bounded from below, and the bottommost vertex (i.e., the vertex of minimum z -coordinate) on f cannot be the intersection of a cone edge and a cone face nor the intersection of three cone faces, as a simple case distinction shows that all such vertices have an incident edge in f where this vertex is strictly above the other endpoint. Thus the bottommost vertex of f is the apex of some cone. On the other hand, each cone apex can be assigned to at most 4 faces (as there cannot be two faces f, f' of ∂U_S within the same cone face). We conclude that there are at most $4|S|$ faces in ∂U_S . By Euler's formula we have that ∂U_S has complexity $O(|S|)$.

Consider the vertical projection U_S^0 of ∂U_S into the plane $z = 0$. Notice that this is exactly an additively weighted ℓ_∞ Voronoi diagram (where weights are the radii of the squares). Using standard techniques [For87, Kle89] this diagram and ∂U_S itself can be computed in $\tilde{O}(|S|)$ time.

We obtain a planar subdivision where edges are either axis-aligned or they are aligned with a 45 degree rotation of the axes. We decompose this subdivision into $O(|S|)$ trapezoids with two vertical sides (or right-angle isosceles triangles with axis-aligned legs, as well as some unbounded polygons with at most two non-vertical sides) using the standard trapezoidation used for point location data structures [ST86] in $O(|S| \log |S|)$ time; let T_S denote the resulting subdivision of size $O(|S|)$. More precisely, in order to get a *partition* of the plane into faces, on boundary edges with normals $(0, 1), (1, 0), (1, 1), (-1, 1)$

1857 we require weak inequalities, while we require strong inequalities for boundary edges with normals
 1858 $(0, -1), (-1, 0), (-1, -1), (1, -1)$.

1859 We project T_S vertically to get a 3-dimensional subdivision \mathcal{T} of U_S into convex vertical slabs: here
 1860 each region is a vertical slab bounded by one face of U_S from below and $\partial f \times \mathbb{R}$ on the sides, where f is
 1861 a face of T_S .

1862 Note that the complexity of \mathcal{T} is $O(|S|)$ and it was computed in $O(|S| \log |S|)$ time. Moreover, each
 1863 slab $T \in \mathcal{T}$ is bounded by faces whose normals can have 12 possible directions: there are 4 possible
 1864 normals for faces coming from ∂U_S , and $4 \cdot 2$ for the vertical faces, as each of these are parallel to one of
 1865 four directions in the plane $z = 0$.

1866 To check whether a point \dot{q} lies in some region $T \in \mathcal{T}$, we need to verify if it is contained in each
 1867 half-space given by ∂T . Each such condition is of the form $\langle \dot{q}, \nu_j \rangle \leq c_T^j$ (or $< c_T^j$) where ν_j is one of 12
 1868 possible normals and c_T^j is a constant that depends only on T . (We define $c_T^j = \infty$ if T does not have a
 1869 face with normal direction ν_j .) For a fixed region T all of these linear conditions can be written as

$$1870 \dot{T} := (c_T^1, c_T^2, \dots, c_T^{12}) \in \text{ort}_q := ((-\infty, \langle \dot{q}, \nu_1 \rangle] \times (-\infty, \langle \dot{q}, \nu_2 \rangle] \times (-\infty, \langle \dot{q}, \nu_3 \rangle) \cdots \times (-\infty, \langle \dot{q}, \nu_{12} \rangle)).$$

1871 Thus, our problem of deciding if q intersects at least one square from S is reduced to the following:
 1872 given a query square q , we compute a 12-dimensional orthogonal range that contains *exactly one* point
 1873 among $\{\dot{T} \mid T \in \mathcal{T}\}$ if and only if q intersects at least one square from S . This problem can be solved with
 1874 12-dimensional orthogonal range searching [dBCvKO08], which requires $\tilde{O}(|\mathcal{T}|) = \tilde{O}(|S|)$ pre-processing
 1875 time and space and $\tilde{O}(1)$ query time, to decide if the query range ort_q contains some point \dot{T} .

1876 **Solving rainbow colored intersection searching.** Suppose now that we are given a set of objects \mathcal{O} ,
 1877 each associated with some color; let \mathcal{S} be the partition of \mathcal{O} into its color classes. For each color class
 1878 $S \in \mathcal{S}$ we set up the subdivision \mathcal{T}_S and compute the corresponding points $\{\dot{T} \mid T \in \mathcal{T}_S\}$. Then we set up
 1879 a standard orthogonal range counting data structure on the 12-dimensional point set $T_S := \bigcup_{T \in \mathcal{T}_S} \dot{T}$ |
 1880 $T \in \mathcal{T}_S$. This takes $\sum_{S \in \mathcal{S}} \tilde{O}(|S|) = \tilde{O}(|\mathcal{O}|)$ preprocessing time and space, and for any orthogonal query
 1881 we can return the number of points in the range in $\tilde{O}(1)$ time.

1882 Given a query square q we can compute the orthant query ort_q and observe that the number of points
 1883 in ort_q is equal to the number of classes $S \in \mathcal{S}$ such that q intersects at least one square from S . Thus, q
 1884 intersects all color classes if and only if ort_q contains exactly $|\mathcal{S}|$ points from T_S .

1885 C.3 Data Structure for Unit Disks

1886 In this subsection, we directly solve the interval cover problem for unit disks restricted to a fixed modulo
 1887 class, as needed in our diameter algorithm and distance oracle for unit-disk graphs. (We do so without
 1888 going through rainbow colored intersection searching, to get better time bounds.) As noted in Section 7,
 1889 this problem reduces to a corresponding interval cover problem about pseudolines:

1890 **Problem C.4.** We are given an input set S of N pseudolines¹⁸ in the plane, where each pseudoline $s \in S$
 1891 has an associated interval I_s . We want to build a data structure to answer the following type of queries:
 1892 given a query point q and interval I , test whether $\bigcup_{\substack{s \in S \\ s \text{ below } q}} I_s$ contains I .¹⁹

¹⁸We assume $O(1)$ time oracle access to deciding if a point is above/on/below a pseudoline, as well as to find the intersection of a pair of pseudolines (or determine that no intersection exists).

¹⁹One way to interpret the problem is to think of each pseudoline s as being “active” for a time window I_s ; a query is to determine whether a given point q stays above the upper envelope of the active pseudolines for the entire duration of the time window I . We will not need this viewpoint for our algorithm.

1893 The rest of this subsection is dedicated to showing that Problem C.4 can be constructed in $N \cdot$
 1894 $2^{O(\sqrt{\log N \log \alpha(N)})}$ preprocessing time and answering a query takes $2^{O(\sqrt{\log N \log \alpha(N)})}$ time, where $\alpha(\cdot)$ is
 1895 the slow-growing inverse Ackermann function. This is sufficient to prove Lemma 7.3.

1896 To appreciate the difficulty of the problem, the reader may first consider the special case when $I = \mathbb{R}$,
 1897 which is already nontrivial. Our idea is to explicitly construct the region of all query points q for which
 1898 the answer is yes. Interestingly, we are able to prove that this region has near-linear combinatorial
 1899 complexity. After constructing the region, answering queries in the case when $I = \mathbb{R}$ would become easy.

1900 To prove this combinatorial fact and at the same time design a data structure for general I , we will
 1901 use a divide-and-conquer strategy.

1902 **Decomposing intervals into canonical intervals.** Assume that the endpoints of all intervals I_s , as
 1903 well as I , are integers bounded by $O(N)$ (by replacing numbers by their ranks). Fix a parameter b . A
 1904 *canonical interval* refers to an interval of the form $[j \cdot b^i, (j+1) \cdot b^i]$ for some i and j . Any interval can
 1905 be expressed as a union of $O(b \log_b N)$ canonical intervals. This is a well-known fact (e.g., in analyzing
 1906 a b -ary range tree [AE99, dBcKO08]). For completeness, we include a quick proof in the following
 1907 paragraph:

1908 Let $J = [x, y]$ be an original interval, and suppose that the largest canonical interval covered by J
 1909 has size $b^k \leq N$. Remove the maximum number of such intervals. Notice that this operation removes
 1910 some middle part M of J consisting of at most b intervals of size b^{k_1} , and leaves an interval J_1 on the left
 1911 of M and J_2 to the right of M , both having size less than b^k and one endpoint that is an integer power
 1912 of b^k . Now if $J_1 = [x, \ell_x \cdot b^k]$, then we can shift it to the interval $J_1 = [x - \ell_x \cdot b^k, 0]$. If $(-z_k \dots z_1 z_0)_b$ is
 1913 the base- b representation of the left endpoint of this interval, then it naturally decomposes this interval
 1914 into $\sum_i z_i \leq 1 + b \log_b N$ intervals. All of these intervals can be shifted by $\ell_y \cdot b^k$ to get a decomposition
 1915 of J_1 . Similarly, we can decompose $J_2 = [\ell_y \cdot b^k, y]$ by considering the base- b representation of
 1916 $J'_2 = [0, y - \ell_y \cdot b^k]$. The resulting representation has at most $b + 2 + 2b \log_b N = O(b \log_b N)$ intervals,
 1917 and it can be found in $O(b \log_b N)$ time.

1918 We replace each interval in the input and queries by canonical intervals. If we do this procedure for
 1919 all of our N intervals then we end up with $O(N \cdot b \log_b N)$ canonical input intervals. For each pseudoline
 1920 s whose original interval J_s has been decomposed into k_s canonical intervals, we will have k_s copies of s
 1921 instead, each associated with one such canonical interval. Thus, we have $N' = O(N b \log_b N)$ pseudolines,
 1922 each associated with a single canonical interval. With slight abuse of notation, we will keep using S for
 1923 this set of pseudolines (where a single pseudoline may appear several times as long as their associated
 1924 canonical intervals are different). Similarly, when a query interval is decomposed into canonical intervals,
 1925 the query cost goes up by at most an $O(b \log_b N)$ factor.

1926 **Preprocessing.** Let $\text{LE}(X)$ and $\text{UE}(X)$ denote the *lower and upper envelope* of a set X of x -monotone
 1927 pseudolines, respectively.

1928 For each canonical interval I , let $S_{\leq I} := \{s \in S : I_s \subseteq I\}$ and $S_I := \{s \in S : I_s = I\}$. Let $\mathcal{E}_{\leq I}$ be the
 1929 boundary of the region of all points $q \in \mathbb{R}^2$ such that

$$\bigcup_{\substack{s \in S_{\leq I} \\ s \text{ below } q}} I_s = I.$$

1931 Then $\mathcal{E}_{\leq I}$ is an x -monotone chain in the arrangement of S —we can view this as a kind of “generalized
 1932 envelope”. We will show that this generalized envelope has near-linear combinatorial complexity and
 1933 can be computed in near-linear time for a sufficiently large choice of b .

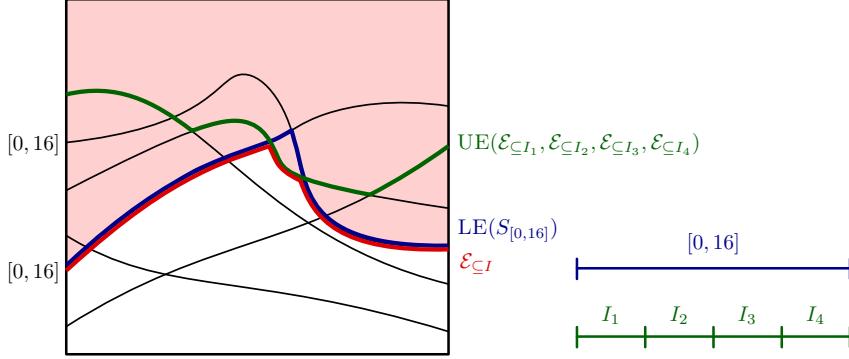


Figure 5. The region $\mathcal{E}_{\subseteq I}$ of points q such that the intervals associated with the pseudolines under them covers $I = [0, 16]$ (red shaded region). The boundary of this region is below all pseudolines associated with $[0, 16]$ (blue envelope) and below the upper envelope of the regions associated with the canonical child intervals (green envelope).

Decompose I into b “child” canonical intervals I_1, \dots, I_b . Note that $S_{\subseteq I} = S_I \cup S_{\subseteq I_1} \cup \dots \cup S_{\subseteq I_b}$. Then we have the following recursive formula for the generalized envelope $\mathcal{E}_{\subseteq I}$ (see Figure 5):

$$\mathcal{E}_{\subseteq I} = \text{LE}(\{\text{LE}(S_I), \text{UE}(\{\mathcal{E}_{\subseteq I_1}, \dots, \mathcal{E}_{\subseteq I_b}\})\}).$$

Let $|\mathcal{E}_{\subseteq I}|$ denote the combinatorial complexity (number of arcs) of $\mathcal{E}_{\subseteq I}$. The upper envelope of f pseudo-segments is known to have combinatorial complexity $O(f \cdot \alpha(f))$, through Davenport-Schinzel sequences [AS00, Pet15]. Thus, $\text{UE}(\{\mathcal{E}_{\subseteq I_1}, \dots, \mathcal{E}_{\subseteq I_b}\})$ has combinatorial complexity $O(|\mathcal{E}_{\subseteq I_1}| + \dots + |\mathcal{E}_{\subseteq I_b}|) \cdot \alpha(N)$. Now, $\text{LE}(S_I)$ has combinatorial complexity $O(|S_I|)$. Thus, $|\mathcal{E}_{\subseteq I}| \leq O(|S_I| + (|\mathcal{E}_{\subseteq I_1}| + \dots + |\mathcal{E}_{\subseteq I_b}|) \cdot \alpha(N))$. The maximum combinatorial complexity, $E(n)$, of $\mathcal{E}_{\subseteq I}$ among those with $|S_{\subseteq I}| = n$, satisfies the recurrence

$$E(n) \leq \max_{n_0, \dots, n_b: n_0 + \dots + n_b = n} (O(\alpha(N)) \cdot (E(n_1) + \dots + E(n_b)) + O(n_0)),$$

which solves to $E(n) = n \cdot \alpha(N)^{O(\log_b N)}$.

For the data structure, we store $\mathcal{E}_{\subseteq I}$ as well as $\text{LE}(S_I)$ for each canonical interval I . The preprocessing time satisfies the recurrence

$$T(n) \leq \max_{n_0, \dots, n_b: n_0 + \dots + n_b = n} (T(n_1) + \dots + T(n_b) + \tilde{O}(E(n_1) + \dots + E(n_b) + n_0)),$$

which solves to $T(n) = \tilde{O}(n \cdot \alpha(N)^{O(\log_b N)})$.

Querying. Given a query point q and a canonical interval I , we check that q is above $\mathcal{E}_{\subseteq I}$ by binary search in the generalized envelope, or that q is above $\text{LE}(S_{I'})$ for some “ancestor” canonical interval $I' \supset I$ (there are $O(\log_b N)$ such intervals I'). The query time is $\tilde{O}(1)$.

Conclusion. After including the $O(b \log_b N)$ factor, the overall preprocessing time is $\tilde{O}(bN \cdot \alpha(N)^{O(\log_b N)})$ and query time is $\tilde{O}(b)$. Setting $b := 2\sqrt{\log N \log \alpha(N)}$, we get $N 2^{O(\sqrt{\log N \log \alpha(N)})} \leq N^{1+o(1)}$ preprocessing time and $2^{O(\sqrt{\log N \log \alpha(N)})} \leq N^{o(1)}$ query time. This concludes the proof.

C.4 Data Structure for Unit Squares

In this subsection, we directly solve the interval cover problem for unit squares. (Again, we do so without going through rainbow colored intersection searching, to get better time bounds.)

1958 **Theorem C.5.** *There is a data structure $\mathcal{D}_{\text{square}}$ that solves the interval cover problem (Problem 1.3) for axis-aligned unit square objects, each associated with a single interval with $N \cdot 2^{O(\sqrt{\log N})} = N^{1+o(1)}$ preprocessing time and $2^{O(\sqrt{\log N})} = N^{o(1)}$ query time.*

1961 Observe that for the geometric intersection graph of unit side-length squares, we can replace each of
 1962 the squares with squares of side-length 2 with the same center such that a pair s, t of original squares
 1963 intersect if and only if the center of s is contained in the scaled square t' . As a result, the data structure
 1964 problem is modified as follows: given a set S of squares, where each $s \in S$ is associated with an interval
 1965 I_s , we need a data structure to decide if the intervals of the squares containing the query point q will
 1966 cover the query interval I .

1967 Instead of the above variant, we overlay a grid of side length 2 (such that no grid line is collinear
 1968 with any square of S); let Γ denote the set of grid cells. Notice that if q is in a given grid cell $\square \in \Gamma$, then
 1969 for each square $s \in S$ we have that $s \cap \square$ appears as an *orthant*, i.e., a rectangle²⁰ containing exactly one
 1970 vertex of \square . Thus, in each cell \square we have the following data structure problem. See Figure 6 for an
 1971 example.

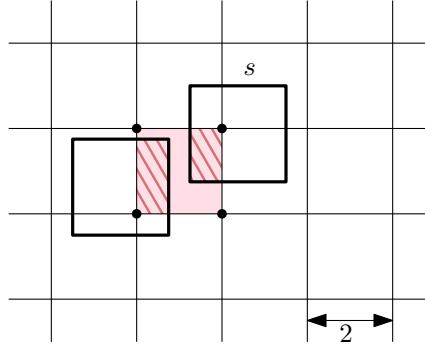


Figure 6. A square s of side length of 2 intersecting cell \square and the intersection (shaded) is an orthant containing exactly one vertex of \square .

1972 **Problem C.6.** *We are given an input set S of N orthants in a cell \square where each orthant covers exactly
 1973 one vertex of \square in \mathbb{R}^2 , where each orthant $s \in S$ has an associated interval I_s . We want to build a data
 1974 structure to answer the following type of queries: given a query point q and interval I , test whether
 1975 $\bigcup_{s \in S: q \in s} I_s$ contains I .*²¹

1976 The rest of this subsection is dedicated to showing that Problem C.6 can be constructed in $N \cdot 2^{O(\sqrt{\log N})}$
 1977 preprocessing time and answering a query takes $2^{O(\sqrt{\log N})}$ time. This is sufficient to prove Theorem C.5,
 1978 as we can use this data structure in each cell: the preprocessing time is $\sum_{\square \in \Gamma} N_{\square} \cdot 2^{O(\sqrt{\log N_{\square}})} = N 2^{O(\sqrt{\log N})}$
 1979 where N_{\square} is the number of orthants in cell \square and $\sum_{\square} N_{\square} = 4N$. To answer queries, we switch to the cell \square_q
 1980 containing q in $O(1)$ time and answer the query using the data structure of \square_q in $2^{O(\sqrt{\log N_{\square}})} \leq 2^{O(\sqrt{\log N})}$
 1981 time.

²⁰We use the term *orthant* to distinguish these rectangles from other rectangles in the proof.

²¹Alternatively, if we think of intervals as living in a third dimension, the problem is equivalent to the following: given a set of axis-aligned boxes in \mathbb{R}^3 where the xy -projection of each box is an orthant, determine whether a query line segment parallel to the z -axis is completely contained in the union of the boxes. We will not need this viewpoint for our algorithm (though this type of 3-dimensional data structure problem seems interesting in its own right). As mentioned, unlike traditional range searching, this problem is not decomposable.

1982 We will use a divide-and-conquer strategy like in Appendix C.3, but the combinatorial complexity of
1983 the regions we want may no longer have near linear complexity (because we do not restrict orthants to a
1984 fixed type), so extra ideas are needed.

1985 Let $\text{UNION}(X)$ denote the union of a set X of rectangles. As seen in Appendix C.3, we will later set
1986 some number b and use canonical intervals of the form $[j \cdot b^i, (j+1) \cdot b^i)$ for some i and j . As seen
1987 in Appendix C.3, for each orthant s whose interval I_s has been decomposed to k_s canonical intervals,
1988 we will have k_s copies of s instead, each associated with one such canonical interval. Thus, we have
1989 $N' = O(N b \log_b N)$ objects, each associated with a single canonical interval. With slight abuse of notation,
1990 we will keep using S for this set of objects. For an interval I we again denote by $S_{\subseteq I}$ and S_I the set of
1991 orthants whose intervals are subsets of I or equal to I , respectively.

1992 Preprocessing.

1993 Let $\mathcal{Z}_{\subseteq I}$ be the region of all points $q \in \mathbb{R}^2$ such that $\bigcup_{s \in S_{\subseteq I}: q \in s} I_s \neq I$.

1994 Unfortunately, the combinatorial complexity of $\mathcal{Z}_{\subseteq I}$ may be quadratic. Instead, we will maintain a set of
1995 rectangles $Z_{\subseteq I}$ with $\text{UNION}(Z_{\subseteq I}) = \mathcal{Z}_{\subseteq I}$. In other words, instead of maintaining the region $\mathcal{Z}_{\subseteq I}$ explicitly,
1996 we implicitly represent $\mathcal{Z}_{\subseteq I}$ as a union of (possibly overlapping) rectangles. We will show that a near
1997 linear number of rectangles suffices (for a sufficiently large b).

1998 With this representation scheme, we can union two regions trivially. However, intersection is a
1999 trickier operation. In the lemma below, we show how to perform intersection with $\text{UNION}(S)^c$ for a set
2000 S of orthants, which is sufficient for our purposes. Here, for a region U , we let $U^c := \square \setminus U$ denote the
2001 complement of U in \square .

2002 **Lemma C.7.** *Given a set S of orthants and a set Z of rectangles in \square we can construct a set Z' of
2003 $O(|S| + |Z|)$ rectangles in $\tilde{O}(|S| + |Z|)$ time, such that $\text{UNION}(S)^c \cap \text{UNION}(Z) = \text{UNION}(Z')$.*

2004 **Proof:** A *tallest-edge data structure* solves the following problem. We are given a set Z of axis-aligned
2005 rectangles in the plane. Then, given a query segment e , we want to find the rectangle $z^* \in Z$ where
2006 the top side of z^* has the maximal y coordinate (i.e., z^* is the tallest) among the rectangles $z \in Z$
2007 covering e . We also allow queries in the other three axis directions, i.e., instead of the tallest reaching
2008 rectangle covering e , we also want to be able to find the leftmost, rightmost, or bottommost reaching
2009 rectangle covering e . Such queries can be answered using range trees in $\text{poly}(\log |Z|)$ query time and
2010 $\tilde{O}(|Z|)$ preprocessing [dBcvKO08]. We start our construction by making a tallest-edge data structure
2011 $\mathcal{D}_{\text{tall}}$ for Z .

2012 Let $S_1 \cup S_2 \cup S_3 \cup S_4$ be the partition of S according to the vertex of \square covered by the orthants. The
2013 *staircase* i for $i = 1, 2, 3, 4$ is the polygonal path $(\partial \bigcup_{s \in S_i} s) \cap \square$.

2014 Set $Z_0 := Z$ and $\mathcal{D}_{\text{tall}}^0 := \mathcal{D}_{\text{tall}}$. For each set S_i we will do the following computation in the order of
2015 their indices ($i = 1, \dots, 4$). Suppose without loss of generality that S_i covers the bottom left corner of \square ;
2016 the other cases will be obtained from this via rotation. Observe that $\text{UNION}(S_i)^c$ is the region above a
2017 staircase. The staircase has $O(|S_i|)$ edges. Intersect the staircase with the boundaries of the rectangles of
2018 Z . Subdivide the edges of the staircase at those intersection points. Note that the edge of the staircase
2019 intersected by a given edge of a rectangle z can be found with a simple binary search. The staircase now
2020 has $O(|S_i| + |Z_{i-1}|)$ edges, and it has been constructed in $\tilde{O}(|S_i| + |Z_{i-1}|)$ time.

2021 For each edge e of the staircase, we query $\mathcal{D}_{\text{tall}}^i$ to find the rectangle $z_e \in Z_{i-1}$ containing e with the
2022 highest top side in $\text{poly}(\log |Z_{i-1}|)$ time. Define z'_e to be the rectangle with bottom side e and top side
2023 touching the top side of z_e . Add z'_e to Z_i .

2024 For each rectangle $z \in Z_{i-1}$, if the bottom side of z intersects the staircase at a point p_z , define z' to
2025 be the part of z to the right of p_z . Add this rectangle z' to Z_i . If z is completely above the staircase, add
2026 z to Z_i .

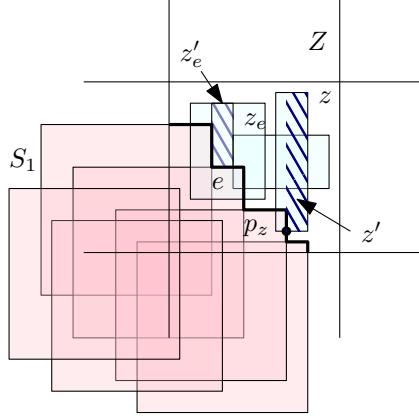


Figure 7. S_1 is the set of squares that cover the bottom left vertex of \square . The staircase of S_1 is shown in a solid polygonal path. The figure shows the rectangles z'_e and z' added to Z_i .

2027 Then Z_i has $O(|S_i| + |Z_{i-1}|)$ rectangles and satisfies the stated property. Finally, we set up a tallest-cover
2028 data structure for Z_i in $\tilde{O}(|Z_i|)$ time. The total time for step i is therefore $\tilde{O}(|S_i| + |Z_{i-1}|)$.

2029 We can handle each of the sets S_i one after another, and we set $Z' := Z_4$. The resulting number of
2030 rectangles is

$$2031 |Z_4| = O(|S_4| + |Z_3|) = O(|S_4| + O(|S_3| + |Z_2|)) = \dots = O\left(\sum_i |S_i| + |Z_0|\right) = O(|S| + |Z|).$$

2032 The total running time is $\sum_i \tilde{O}(|S_i| + |Z_{i-1}|) = \tilde{O}(|S| + |Z|)$. \square

2033 Recall that the canonical interval I can be decomposed into b “child” canonical intervals I_1, \dots, I_b .
2034 Suppose that there are n_j orthants in $S_{\subseteq I_j}$ and n_0 orthants in S_I . We can compute $Z_{\subseteq I}$ using the following
2035 recursive formula:

$$2036 Z_{\subseteq I} = \text{UNION}(S_I)^c \cap (Z_{\subseteq I_1} \cup \dots \cup Z_{\subseteq I_b}).$$

2037 We can apply the lemma to find a set $Z_{\subseteq I}$ of $O(|S_I| + |Z_{\subseteq I_1} \cup \dots \cup Z_{\subseteq I_b}|)$ rectangles with $\text{UNION}(Z_{\subseteq I}) = Z_{\subseteq I}$.

2038 The number of rectangles in $Z_{\subseteq I}$, assuming $|S_{\subseteq I}| = n$, satisfies the recurrence

$$2039 E(n) \leq \max_{n_0, \dots, n_b: n_0 + \dots + n_b = n} (O(1) \cdot (E(n_1) + \dots + E(n_b)) + O(n_0)),$$

2040 which solves to $E(n) = O(n \cdot 2^{O(\log_b N)})$.

2041 To construct the data structure $\mathcal{D}_{\text{square}}$, we store $Z_{\subseteq I}$ and S_I in individual rectangle stabbing data
2042 structures [Cha86, SJ05], for each canonical interval I . The data structure for a given canonical interval
2043 I can therefore be made in $\tilde{O}(E(n))$ time.

2044 Consequently, the preprocessing time satisfies the recurrence

$$2045 T(n) \leq \max_{n_0, \dots, n_b: n_0 + \dots + n_b = n} (T(n_1) + \dots + T(n_b) + \tilde{O}(E(n_1) + \dots + E(n_b) + n_0)),$$

2046 which solves to $T(n) = \tilde{O}(n \cdot 2^{O(\log_b N)})$.

2047 **Querying.** Given a query point q and a canonical interval I , we check that q is not stabbing any
2048 rectangle in $Z_{\subseteq I}$, or that q stabs some orthant in $S_{I'}$ for some “ancestor” canonical interval $I' \supset I$. Since
2049 there are $O(\log_b N) = O(\log N)$ ancestor canonical intervals, the query time is $\tilde{O}(1)$.

2050 To answer the query about the original interval J , we make individual queries on each of the
2051 $O(b \log_b N)$ canonical intervals in its decomposition, and answer “yes” if and only if each canonical
2052 interval was covered.

2053 **Conclusion.** After including the $O(b \log_b N)$ factor, the overall preprocessing time is $\tilde{O}(bN \cdot 2^{O(\log_b N)})$
 2054 and the query time is $\tilde{O}(b)$. Setting $b = 2^{\sqrt{\log N}}$, we get $N2^{O(\sqrt{\log N})} \leq N^{1+o(1)}$ preprocessing time and
 2055 $2^{O(\sqrt{\log N})} \leq N^{o(1)}$ query time, and conclude the proof of Theorem C.5.

2056 D Switching Interval Representation between Different Stabbing Paths

2057 We are given a set system (X, \mathcal{S}) with at most $n = |X|$ elements and $m = |\mathcal{S}|$ sets with dual shatter
 2058 dimension of (X, \mathcal{S}) is d . Throughout the rest of the section we assume the existence of an *element*
 2059 *reporting oracle* that, given $S \in \mathcal{S}$, can enumerate all elements of S in $T_0(n)$ time, where $T_0(n) \geq n$.

2060 Let λ be an ordering of X . We say that a set S *crosses* a pair (x, y) if $x \in S$ and $y \notin S$, or vice versa.
 2061 The number of consecutive pairs in λ crossed by S is at most twice the size $|\text{Rep}_\lambda(S)|$. For any collection
 2062 \mathcal{R} , define the *equivalence relation* $\equiv_{\mathcal{R}}$ over X , where $x \equiv_{\mathcal{R}} y$ if and only if no set in \mathcal{R} crosses (x, y) . (In
 2063 other words, $\{S \in \mathcal{R} : x \in S\} = \{S \in \mathcal{R} : y \in S\}$.) Then $\equiv_{\mathcal{R}}$ has $O(|\mathcal{R}|^d)$ equivalence classes since the
 2064 dual shatter dimension is at most d . For every x and y in X , the *crossing number* $c_{\mathcal{S}}(x, y)$ is the number
 2065 of sets in \mathcal{S} crossing (x, y) . (Notice that $c_{\mathcal{S}}(\cdot, \cdot)$ forms a pseudometric.)

2066 For the purpose of the remaining section, we will fix a *ρ -sampling* \mathcal{R} of \mathcal{S} , where each set in \mathcal{S} chosen
 2067 with probability ρ/m . (Later on we will restrict \mathcal{R} to subcollection \mathcal{S}' of \mathcal{S} and obtain \mathcal{R}' ; we can still
 2068 think of \mathcal{R}' as obtained from \mathcal{S}' by sampling each element with probability ρ/m , even though we do not
 2069 explicit sample from \mathcal{S}' . Notice that the parameter m does *not* change even if \mathcal{S}' gets smaller.)

2070 Our first goal is to prove that any ρ -sampling of \mathcal{S} has low crossing number and thus can be used to
 2071 construct a stabbing path for (X, \mathcal{S}) .

2072 **Lemma D.1.** *Let \mathcal{R} be ρ -sampling of \mathcal{S} . Then for every $x, y \in X$ with $x \equiv_{\mathcal{R}} y$, crossing number $c_{\mathcal{S}}(x, y)$
 2073 is at most $O((m/\rho) \log n)$ with high probability.*

2074 **Proof:** This follows by a standard hitting set argument. Consider any two elements $x, y \in X$ with
 2075 crossing number $c_{\mathcal{S}}(x, y)$. By standard Chernoff bounds, if $c_{\mathcal{S}}(x, y) = \Omega((m/\rho) \log n)$, we would have
 2076 sampled one of the sets in \mathcal{R} that cross (x, y) with high probability, i.e., probability $1 - 1/n^c$ for a large
 2077 constant c , but $x \equiv_{\mathcal{R}} y$ which is a contradiction. The conclusion follows after taking a union bound over
 2078 the n^2 pairs of elements. \square

2079 Let \mathcal{S}' be an arbitrary subcollection of \mathcal{S} . Denote the restriction of the fixed ρ -sampling \mathcal{R} of \mathcal{S} in
 2080 \mathcal{S}' as \mathcal{R}' ; in notation, $\mathcal{R}' := \mathcal{S}' \cap \mathcal{R}$. Notice that \mathcal{R}' is also a ρ -sampling. Given any set system (X, \mathcal{S}) , a
 2081 stabbing path λ of (X, \mathcal{S}) is *\mathcal{R}' -respecting* if each equivalence class of $\equiv_{\mathcal{R}'}$ appears contiguously in λ for
 2082 the restriction \mathcal{R}' . (The equivalent classes of $\equiv_{\mathcal{R}'}$ is with respect to the restriction \mathcal{R}' , not \mathcal{R} .) The above
 2083 proof can be adapted so that the resulting stabbing path is \mathcal{R} -respecting (by choosing $\mathcal{S}' = \mathcal{S}$):

2084 **Lemma 7.4.** *Assume the existence of an element reporting oracle that, given $S \in \mathcal{S}$, can enumerate
 2085 all elements of S in $T_0(n)$ time. Consider a fixed ρ -sampling \mathcal{R} of \mathcal{S} . We can compute the equivalence
 2086 classes of $\equiv_{\mathcal{R}}$ and construct an \mathcal{R} -respecting stabbing path λ of (X, \mathcal{S}) such that $\sum_{S \in \mathcal{S}} |\text{Rep}_\lambda(S)| =$
 2087 $\tilde{O}(mn/\rho + m\rho^{d-1})$ in $\tilde{O}(T_0(n) \cdot \rho)$ time with high probability. In other words, one can compute a sampled
 2088 ρ -stabbing path λ of (X, \mathcal{S}) and the equivalence classes of $\equiv_{\mathcal{R}}$ as byproducts.*

2089 **Proof:** Let \mathcal{R} be a ρ -sampling of \mathcal{S} ; then $|\mathcal{R}| = \tilde{O}(\rho)$ with high probability. We first enumerate the
 2090 elements in all $R \in \mathcal{R}$ in $\tilde{O}(T_0(n) \cdot \rho)$ time, and compute the $O(\rho^d)$ equivalence classes of $\equiv_{\mathcal{R}}$ in $\tilde{O}(n\rho)$
 2091 time. (Each class will appear contiguously in the stabbing path λ to be constructed.) Within each
 2092 equivalence class C_i , we order its elements $x_1^{(C_i)}, \dots, x_{|C_i|}^{(C_i)}$ arbitrarily. We recursively compute an ordering

2093 of $\{x_1^{(C_i)} : i \in [1 : \rho^d]\}$ by invoking the main statement of the lemma itself (which is inductively \mathcal{R} -
 2094 respecting), with run time $\tilde{O}(T_0(n) \cdot (\rho^d)^{1/d}) = \tilde{O}(T_0(n) \cdot \rho)$. We then order the classes C_i (as intervals
 2095 of elements) according to the order of $\{x_1^{(C_i)} : i \in [1 : \rho^d]\}$.

2096 Since $2 \cdot |\text{Rep}_\lambda(S)|$ is equal to the number of consecutive pairs in λ crossed by S , with high probability

$$\begin{aligned} 2097 \quad & \sum_{S \in \mathcal{S}} 2 \cdot |\text{Rep}_\lambda(S)| \\ 2098 \quad &= \sum_{C_i} (c_S(x_1^{(C_i)}, x_2^{(C_i)}) + \dots + c_S(x_{|C_i|}^{(C_i)}, x_1^{(C_{i+1})})) \\ 2099 \quad &\leq \sum_{C_i} (2c_S(x_1^{(C_i)}, x_2^{(C_i)}) + \dots + c_S(x_1^{(C_i)}, x_1^{(C_{i+1})})) \\ 2100 \quad &\leq \tilde{O}(mn/\rho) + \sum_{C_i} c_S(x_1^{(C_i)}, x_1^{(C_{i+1})}), \end{aligned}$$

2101 where the first inequality follows from applying the triangle inequality of $c_S(\cdot, \cdot)$ (because $c_S(\cdot, \cdot)$ forms a
 2102 pseudometric) on the last term $c_S(x_{|C_i|}^{(C_i)}, x_1^{(C_{i+1})})$, and the second inequality is from Lemma D.1. Since we
 2103 recurse on the first element of every equivalence class, by recursion we have

$$2104 \quad \sum_{S \in \mathcal{S}} |\text{Rep}_\lambda(S)| \leq \tilde{O}(mn/\rho + m(\rho^d)^{1-1/d}) = \tilde{O}(mn/\rho + m\rho^{d-1}).$$

2105 The ordering is clearly \mathcal{R} -respecting. The total running time is $\tilde{O}(T_0(n) \cdot \rho)$. \square

2106 **Lemma D.2.** Consider a fixed ρ -sampling \mathcal{R} of \mathcal{S} . We are given two \mathcal{R} -respecting stabbing paths λ and
 2107 λ' of (X, \mathcal{S}) (along with the equivalence classes of $\equiv_{\mathcal{R}}$). Let \mathcal{T} be an arbitrary subcollection of \mathcal{S} . Given
 2108 $\text{Rep}_\lambda(S)$ for all $S \in \mathcal{T}$, we can compute $\text{Rep}_{\lambda'}(S)$ for all $S \in \mathcal{T}$ in $\tilde{O}(mn/\rho + m\rho^d)$ total time with high
 2109 probability.

2110 **Proof:** Given $S \in \mathcal{T}$ and an equivalent class C of $\equiv_{\mathcal{R}}$, we compute the part of $\text{Rep}_{\lambda'}(S)$ within C as
 2111 follows. Fix one representative element $x_C \in C$.

- 2112 • Case 1: $x_C \notin S$. We enumerate all $x \in C$ in S , by examining the union of intervals of $\text{Rep}_\lambda(S)$. We
 2113 then concatenate $\langle x \rangle$ (singletons) over all such x in the order determined by λ' .
- 2114 • Case 2: $x_C \in S$. We enumerate all $x \in C$ not in S , by examining the complement of the union of
 2115 intervals of $\text{Rep}_\lambda(S)$. We then concatenate $\langle x \rangle$ (singletons) over all such x in the order determined
 2116 by λ' , and take the complement of the resulting union of intervals.

2117 In both cases, the run time is linear in the number of $x \in C$ such that S crosses (x_C, x) . So, the total
 2118 run time over all $S \in \mathcal{T}$ is upper-bounded by $\sum_C \sum_{x \in C} c_S(x_C, x) = \tilde{O}(mn/\rho)$ with high probability by
 2119 Lemma D.1.

2120 Finally, we concatenate the different parts of $\text{Rep}_{\lambda'}(S)$ over all the classes, in the order determined
 2121 by λ' . This takes additional $O(m\rho^d)$ total time. \square

2122 **Lemma 7.5.** [Conversion of interval representations.] Let (X, \mathcal{S}) be a set system with $|X| \leq n$ and $|\mathcal{S}| \leq m$.
 2123 Let \mathcal{S}' be a subcollection of \mathcal{S} and \mathcal{T} be a subcollection of \mathcal{S}' . Let \mathcal{R} be the unique ρ -sampling of \mathcal{S} , and
 2124 \mathcal{R}' be its restriction in \mathcal{S}' . We are given an \mathcal{R} -respecting stabbing path λ of (X, \mathcal{S}) , and an \mathcal{R}' -respecting
 2125 ordering λ' of (X, \mathcal{S}') (along with the equivalence classes of $\equiv_{\mathcal{R}}$ and $\equiv_{\mathcal{R}'}$).

2126 (1) [Shrinking from \mathcal{S} to \mathcal{S}' .] Given $\text{Rep}_\lambda(S)$ for all $S \in \mathcal{T}$, we can compute $\text{Rep}_{\lambda'}(S)$ for all $S \in \mathcal{T}$ in
 2127 $\tilde{O}(mn/\rho + m\rho^d)$ total time with high probability.

2128 (2) [Expanding from S' to S .] Given $\text{Rep}_{\lambda'}(S)$ for all $S \in \mathcal{T}$, we can compute $\text{Rep}_{\lambda}(S)$ for all $S \in \mathcal{T}$ in
 2129 $\tilde{O}(mn/\rho + m\rho^d)$ total time with high probability.

2130 **Proof:** For (1), notice that the equivalence classes for $\equiv_{\mathcal{R}}$ are refinements of the equivalence classes for
 2131 $\equiv_{\mathcal{R}'}$. Let λ'' be an ordering obtained by taking λ , and re-ordering the classes for $\equiv_{\mathcal{R}}$ so that classes inside
 2132 a common class of $\equiv_{\mathcal{R}'}$ appear contiguously, which takes $O(m\rho^d)$ time. This way, λ'' is both \mathcal{R} -respecting
 2133 and \mathcal{R}' -respecting. Now, we can apply Lemma D.2 twice, to convert from $\text{Rep}_{\lambda}(S)$ to $\text{Rep}_{\lambda''}(S)$ and from
 2134 $\text{Rep}_{\lambda''}(S)$ to $\text{Rep}_{\lambda}(S)$ for all $S \in \mathcal{T}$. This takes $\tilde{O}(mn/\rho + m\rho^d)$ time.

2135 Similarly for (2), Let λ'' be an ordering obtained by taking λ , and re-ordering the classes for $\equiv_{\mathcal{R}}$ so
 2136 that classes inside a common class of $\equiv_{\mathcal{R}'}$ appear contiguously. This way, λ'' is both \mathcal{R} -respecting and
 2137 \mathcal{R}' -respecting. Now, we can apply Lemma D.2 twice, to convert from $\text{Rep}_{\lambda'}(S)$ to $\text{Rep}_{\lambda''}(S)$ and from
 2138 $\text{Rep}_{\lambda''}(S)$ to $\text{Rep}_{\lambda}(S)$. \square

2139 E Handling Small Pieces

2140 E.1 Patterns

2141 Let P be a piece in some LDD of G with diameter Δ . Recall that the set of boundary vertices of P is
 2142 denoted by ∂P . Fix an arbitrary sequence of vertices $\sigma_P = \langle s_1, s_2, \dots, s_{|\partial P|} \rangle$. For each vertex $v \in V(G)$,
 2143 let $d(v, P)$ denote the distance between v and any vertex of P . We denote a *pattern* of v with respect to
 2144 the ordering σ_P , denoted by \mathbf{p}_v to be the following $|\partial P|$ dimensional vector:

$$2145 \mathbf{p}_v[i] = d(v, s_i) - d(v, P) \quad \text{for every } 1 \leq i \leq |\partial P|.$$

2146 We remark that instead of subtracting by an offset of $d(v, P)$, we could have subtracted by any other
 2147 offset. For example, [LW24] instead use the offset of $d(v, s_1)$.

2148 Le and Wulff-Nilsen [LW24] showed a bound on the total number of patterns with respect to σ_P
 2149 if the distance encoding VC-dimension is bounded. The proof also works for generalized distance
 2150 VC-dimension.

2151 **Lemma E.1.** *Let P be a piece in a graph G with general distance VC-dimension d and σ_P an arbitrary
 2152 ordering on ∂P . Let $\mathbf{P} = \{\mathbf{p}_v \mid v \in V(G)\}$ be the set of patterns with respect to σ_P . Then $|\mathbf{P}| = O(|\partial P|^d \Delta^d)$.*

2153 **Proof:** Consider the set system $(V_G \times \mathbb{Z}, \mathcal{GB})$ of generalized neighborhood balls, and the set system
 2154 where we restrict the ground set $(\partial P \times [\Delta], \mathcal{GB})$. This restriction of the ground set does not increase
 2155 the VC-dimension of the set system. There is a clear bijection between $\mathbf{p}_v \in \mathbf{P}$ and the generalized
 2156 neighborhood ball: $\tilde{N}^{d(v, P)}[v] \cap (\partial P \times [\Delta]) = \{(u, r) : u \in \partial P, r \in [\Delta], d(u, v) \leq d(u, P) + r\}$. So the
 2157 number of patterns is bounded by the number of unique sets of $(\partial P \times [\Delta], \mathcal{GB})$. By the Sauer-Shelah
 2158 Lemma (see Lemma 2.4), this is at most $O(|\partial P|^d |\Delta|^d)$. \square

2159 E.2 Diameter and Eccentricities using Patterns

2160 The following algorithm computes the eccentricities of all vertices in a piece P of the graph G .

- 2161 1. Compute the pairwise distance between pairs of vertices in P . Let d'_v denote the distance to the
 2162 farthest vertex from v that is within P .
- 2163 2. Compute all patterns \mathbf{P} for P , and for each pattern $\mathbf{p} \in \mathbf{P}$ find the farthest vertex $u \in V(G)$ that
 2164 attains that pattern. Let $d_{\mathbf{p}}$ be $d(u, s_0)$, the base distance for the pattern.

2165 3. For each pattern $\mathbf{p} \in \mathbf{P}$, compute the distance $d(\mathbf{p}, v)$ from the pattern to each $v \in P$ by doing
 2166 a **boundary weighted BFS**, i.e., a BFS where the boundary vertex distances are initialized to the
 2167 values of the pattern \mathbf{p} , and for each vertex $v \in P$ compute $d_v = \max_{\mathbf{p} \in \mathbf{P}} d(\mathbf{p}, v) + d_p$.
 2168 4. Return $\max\{d_v, d'_v\}$.

2169 Step 1 can be implemented by running a BFS within P from each vertex. By Lemma E.1 the number
 2170 of patterns computed in step 2 is at most $O(|\partial P|^d \Delta^d)$, and it takes $O(n|\partial P|)$ time to consider all distances
 2171 to compute the pattern. Running a BFS for each pattern in step 3 takes time $T(P)$ per pattern where
 2172 $T(P)$ is the time it takes to run a boundary weighted BFS in P .

2173 **Lemma 2.15.** *Let G be a graph on n vertices with distance encoding VC-dimension d . Let P be a piece in
 2174 G with boundary ∂P and diameter Δ . If distances from ∂P to all vertices of G are known, the eccentricity
 2175 of all vertices in P can be computed in $O(n \cdot |\partial P| + (|P| + |\partial P|^d \Delta^d) \cdot T(P))$ where $T(P)$ is the time it
 2176 takes to run boundary weighted BFS on P with weights at most Δ .*

2177 E.3 Boundary Weighted BFS in Geometric Intersection Graphs

2178 One approach to compute a shortest path tree in a unit disk graph of n disks uses a semi-dynamic data
 2179 structure, developed in [EIK01], that in $O(\log n)$ amortized time finds a disk containing a query point
 2180 and deletes it from the set. Thus one can repeatedly apply the data structure to find the disks at the
 2181 $(i + 1)$ -hop frontier in a BFS tree from the i th hop frontier— for each disk at i -hop away from the root,
 2182 repeated query the center of the disk to look for disks that intersect with it until such disks are exhausted.
 2183 This gives a running time of $O(n \log n)$ to compute a BFS tree, since each disk is only deleted once. The
 2184 semi-dynamic data structure uses a grid of side length $1/2$. For each cell Q of the grid, maintain the set
 2185 of disks whose center lies in Q . Furthermore, maintain the upper envelope S_1 of the disks that intersect
 2186 Q with centers below the line through the lower boundary of Q , and similarly maintain the envelopes
 2187 S_2, S_3, S_4 for the other three boundaries. Therefore, if a query point q lies in a cell Q , all the disks that
 2188 are centered inside Q would contain q and can be returned. Further, query q against the upper envelope
 2189 S_1 (check if q is below S_1) to look for additional candidates. And repeat the same procedure for the
 2190 other three envelopes. The upper envelope is maintained by a binary tree similar to a segment tree.

2191 The boundary weighted BFS problem in a unit disk graph can be solved by a slight modification of
 2192 this procedure: vertices on the boundary appear as query points when the shortest path tree has reached
 2193 a sufficient depth. Therefore we have the following observation.

2194 **Observation E.2.** *The boundary weighted BFS problem in a unit disk graph can be implemented in
 2195 $O(|P| \log |P|)$ time.*

2196 For boundary weighted BFS in the intersection graph of axis aligned squares (of varying sizes), we
 2197 can use the same idea above. We need the following semi-dynamic data structure for a set of axis-parallel
 2198 squares: given a query square q return a square r that intersects q , and then delete r . If q and r intersect,
 2199 either some corners of q is inside r or some corners of r are inside q . Thus, the above query can be
 2200 implemented by running an orthogonal range query of q on the set of corner points of current set of
 2201 squares, as well as a point enclosure query [Cha86] (also called a rectangle stabbing query [SJ05])
 2202 of each of the corners of q against the set of current squares. These queries can be answered by 2D
 2203 orthogonal range trees or 2D segment trees. By using dynamic fractional cascading with deletion only,
 2204 both query and deletion can be handled in $O(\log s)$ amortized time if we have s squares [CG86b, CG86a].
 2205 Therefore, we have the following lemma.

2206 **Lemma E.3.** *The boundary weighted BFS problem in the intersection graph of axis-aligned squares can
 2207 be implemented in $O(|P| \log |P|)$ time.*

2208 E.4 Exact Distance Oracles

2209 The lemmas in this section are implicit in the distance oracles of [LW24], but we present their proofs in
 2210 full to keep our exposition self-contained.

2211 **Lemma E.4 (Section 3.2.3 of [LW24]).** *Let G be a graph on n vertices with bounded generalized
 2212 distance VC-dimension d and P be a piece in G with boundary ∂P and diameter Δ . There exists an exact
 2213 distance oracle for queries in which at least one end point lies within P with $O(n + |\partial P|^d \Delta^d |P| + |P|^2)$
 2214 space and $O(1)$ query time.*

2215 *Furthermore if distances from ∂P to all vertices of G are known, the distance oracle can be computed
 2216 in $O(n|\partial P| + (|\partial P|^d \Delta^d + |P|) \cdot T(P))$ precomputation time, where $T(P)$ is the time it takes to run vertex
 2217 weighted BFS on P with weights at most Δ .*

2218 **Proof:** For each vertex $v \in P$, store the distances to all other vertices in P . Every other vertex of the
 2219 graph $u \in G \setminus P$ stores a pointer to their respective pattern \mathbf{p}_u , and the distance $d(u, P)$. Also store the
 2220 distance $d(\mathbf{p}, v)$ for each pattern $\mathbf{p} \in \mathbf{P}$ to each $v \in P$.

2221 To handle a query between two vertices of P , we can look up the distance between the vertices in
 2222 constant time. For one vertex $v \in P$, and another vertex $u \in G \setminus P$, we know that:

$$2223 \quad d(u, v) = d(u, P) + d(\mathbf{p}_u, v).$$

2224 and we can look up $d(u, P)$, \mathbf{p}_u , and $d(\mathbf{p}_u, v)$ in constant time.

2225 The total space needed for the oracle is $O(|P|^2)$ for the distances between pairs of vertices in P , $O(n)$
 2226 for the pointers from vertices $u \in G \setminus P$ to their respective patterns, and $O(|\partial P|^d \Delta^d |P|)$ to store the
 2227 pattern to P distances.

2228 The precomputation time is the same as in Lemma 2.15 for eccentricities. \square

2229 **Lemma 2.16 (Section 4.3.1 of [LW24]).** *Let $G = (V_G, E_G)$ be a graph with bounded distance VC-
 2230 dimension d , and P be an induced subgraph of G with boundary ∂P and diameter Δ . There exists a
 2231 distance oracle that answers distances from any vertex $s \in P$ and any vertex $t \in V_G$ with $O(n \cdot |\partial P| + |P|^d)$
 2232 space and $O(\log |\partial P|)$ query time.*

2233 *Furthermore, if G also has bounded generalized distance VC-dimension d and distances from ∂P
 2234 to all vertices of G , the distance oracle can be computed in $O(n \cdot |\partial P| + (|\partial P|^d \Delta^d + |P|) \cdot T(P))$ time,
 2235 where $T(P)$ is the time it takes to run vertex weighted BFS on P with weights at most Δ .*

2236 **Proof:** Store the distances between pairs of vertices in P . For every other vertex $u \in G \setminus P$, consider the
 2237 sequence of balls $B(u, r_1), \dots, B(u, r_k)$ such that $B(u, r_1)$ is the smallest ball that contains at least one
 2238 vertex of ∂P , and $B(u, r_i)$ is the smallest ball containing at least one vertex of $\partial P \setminus B(u, r_{i-1})$ (note that
 2239 $k \leq |\partial P|$). Store a pointer to each of these balls, and the set of vertices $Y_i = B(u, r_i) \cap P$ (and $Y_0 = \emptyset$) in
 2240 a data structure that allows for $O(1)$ time membership lookup. For each relevant Y_i , store the distance
 2241 $d(Y_i, v) := \min_{s \in Y_i} d(s, v)$.

2242 If two vertices u and v are within P , we can look up their distance in $O(1)$ time. Otherwise, if $v \in P$
 2243 and $u \in G \setminus P$, then we can binary search over Y_0, Y_1, \dots, Y_k to find the first Y_i where $v \notin Y_i$ and $v \in Y_{i+1}$ in
 2244 $O(\log k) = O(\log |\partial P|)$ time. Then, we can look up the distance $d(Y_i, v)$ in constant time and return the
 2245 distance:

$$2246 \quad d(u, v) = d(u, r_i) + d(Y_i, v).$$

2247 The space required to store distances between pairs of vertices in P is at most $O(|P|^2)$. The space
 2248 required is $O(n|\partial P|)$ to store the pointers between $u \in G \setminus P$ and their respective Y_0, Y_1, \dots, Y_k since
 2249 $k \leq |\partial P|$. The total number of balls is at most $O(|P|^d)$ by Lemma 2.4 (Sauer's lemma).

2250 To compute this distance oracle, we need to compute Y_1, \dots, Y_k for each vertex $u \in G \setminus P$. To do so,
2251 we can cluster these vertices into vertices with the same pattern \mathbf{p}_u , and consider Y_1, \dots, Y_k with respect
2252 to each pattern. This can be done as the BFS to compute $d(\mathbf{p}_u, v)$ for every vertex $v \in P$ also implicitly
2253 computes the balls Y_1, \dots, Y_k , as well as the distances $d(Y_i, v)$. To compute a pointer from u to Y_i , we can
2254 look up the balls we computed by storing all Y_i s in a data structure that supports $O(1)$ lookup for sets
2255 (e.g. a hashing based data structure). The precomputation time analysis is the same as in Lemma 2.15
2256 for eccentricities. \square