

32 PROXIMITY ALGORITHMS

Joseph S. B. Mitchell and Wolfgang Mulzer

INTRODUCTION

The notion of distance is fundamental to many aspects of computational geometry. A classic approach to characterize the distance properties of planar (and high-dimensional) point sets that has been studied since the early 1980s uses *proximity graphs* (Section 32.1). Proximity graphs are geometric graphs in which two vertices p, q are connected by an edge (p, q) if and only if a certain *exclusion region* for p, q contains no points from the vertex set. Depending on the specific exclusion region, many variants of proximity graphs can be defined, such as relative neighborhood graphs, Delaunay triangulations, β -skeletons, empty-strip graphs, etc. Since proximity graphs encode interesting information on the intrinsic structure of the point set, they have found many applications. From an algorithmic point of view, it is extremely useful to have a compact representation of the distance structure of a point set. The *well-separated pair decomposition* (WSPD) offers one way to achieve this (Section 32.2). WSPDs have numerous algorithmic applications, and the notion generalizes to certain non-Euclidean metrics. Furthermore, several variants of the WSPD have been developed to address its shortcomings, e.g., *semi-separated pair decompositions* and (α, β) -*pair decompositions*. *Geometric spanners* provide another means to approximate the complete Euclidean metric (Section 32.3). Here, the distance function is approximated by the shortest path distance in a sparse geometric graph. There are four basic constructions for geometric spanners: the greedy spanner, the Yao graph, the Θ -graph and the WSPD-spanner. To optimize various parameters, many variants have been defined, and the notion can be generalized beyond the Euclidean setting. Finally, we discuss work on making proximity structures *dynamic*, allowing for insertions and deletions of points (Section 32.4). The fundamental problem here is the *dynamic nearest neighbor problem*, which serves as a starting point for other structures. Additionally, there are several results on making geometric spanners dynamic.

32.1 PROXIMITY GRAPHS

GLOSSARY

Geometric graph: A graph $G = (V, E)$ together with an embedding in \mathbb{R}^d that maps V to points and E to straight line segments that do not pass through nonincident vertices. (See Chapter 10.)

Planar straight-line graph (PSLG): A geometric graph $G = (V, E)$ embedded in \mathbb{R}^2 with noncrossing edges.

$\delta(p, q)$: The distance between two points p and q .

Diameter of a point set V : The maximum distance $\delta(p, q)$ between two points $p, q \in V$. A pair that achieves the diameter is called a *diametric pair*.

Closest pair of a point set V : A pair $\{p, q\}$ of two distinct points in V with minimum distance $\delta(p, q)$. The distance $\delta(p, q)$ is called the *closest pair distance*.

Spread $\Phi(V)$ of a point set $V \subset \mathbb{R}^d$: The ratio between the diameter and the closest pair distance of V : $\Phi(V) = \max_{s, t \in V} \delta(s, t) / \min_{s \neq t \in V} \delta(s, t)$.

L_p -metric: Let $x = (x_1, \dots, x_d)$ and $y = (y_1, \dots, y_d)$ be two d -dimensional points. For $p \in [1, \infty]$, we set $\delta_p(x, y) = (\sum_{i=1}^d |x_i - y_i|^p)^{1/p}$.

In particular, we have $\delta_1(x, y) = \sum_{i=1}^d |x_i - y_i|$ and $\delta_\infty(x, y) = \max_{i=1}^d |x_i - y_i|$.

Ball $B(x, r)$: Let x be a point and $r \geq 0$. Then, we define the open ball $B(x, r) = \{y \mid \delta(x, y) < r\}$.

Nearest-neighbor graph $\text{NNG}(V)$: The directed graph with vertex set V and an edge (p, q) if and only if $B(p, \delta(p, q)) \cap V = \emptyset$.

Lune $L(p, q)$: For two points p and q , we set

$$L(p, q) = B(p, \delta(p, q)) \cap B(q, \delta(p, q)).$$

Some authors prefer the term *lens* instead of lune.

Relative neighborhood graph $\text{RNG}(V)$: The graph with vertex set V and an edge (p, q) if and only if $L(p, q) \cap V = \emptyset$. Thus, the edge is present if and only if

$$\delta(p, q) = \min_{v \in V} \max\{\delta(p, v), \delta(q, v)\}.$$

Gabriel graph $\text{GG}(V)$: The graph with vertex set V and an edge (p, q) if and only if

$$B\left(\frac{p+q}{2}, \frac{\delta(p, q)}{2}\right) \cap V = \emptyset.$$

β -lune $L_\beta(p, q)$: Let p and q be two points. For $\beta = 0$, $L_\beta(p, q)$ is the open line segment pq .

For $\beta \in (0, 1)$, $L_\beta(p, q)$ is the intersection of the two open disks of radius $\delta(p, q)/(2\beta)$ having bounding circles passing through both p and q .

For $\beta \geq 1$, we set

$$L_\beta(p, q) = B\left(p\left(1 - \frac{\beta}{2}\right) + q\frac{\beta}{2}, \frac{\beta}{2}\delta(p, q)\right) \cap B\left(q\left(1 - \frac{\beta}{2}\right) + p\frac{\beta}{2}, \frac{\beta}{2}\delta(p, q)\right).$$

Lune-based β -skeleton $G_\beta^l(V)$: Let $\beta \geq 0$. We define $G_\beta^l(V)$ as the graph with vertex set V and an edge (p, q) if and only if $L_\beta(p, q) \cap V = \emptyset$.

Circle-based β -skeleton $G_\beta^c(V)$: For $\beta = 0$, we define $G_\beta^c(V)$ as the graph with vertex set V and an edge (p, q) if and only if the open line segment pq contains no other points from V .

For $\beta \in (0, 1)$, we define $G_\beta^c(V)$ as the graph with vertex set V and an edge (p, q) if and only if the intersection of the two open disks with radius $\delta(p, q)/(2\beta)$ passing through p and q does not contain any other points from V .

For $\beta \geq 1$, we define $G_\beta^c(V)$ as the graph with vertex set V and an edge (p, q) if and only if the union of the two open disks with radius $\beta\delta(p, q)/2$ passing through p and q does not contain any other points from V .

L_p -Delaunay triangulation $D_p(V)$: The graph with vertex set V that is the straight-line dual of the Voronoi diagram of V with respect to the L_p -norm.

Empty strip graph: The graph with vertex set V and an edge (p, q) if and only if the open infinite strip bounded by the two lines through p and through q orthogonal to the line segment pq contains no points from V .

Sphere of influence graph $SIG(V)$: Let C_p be a circle centered at p with radius equal to the distance to a nearest neighbor of p . Then, $SIG(V)$ is the graph with vertex set V and an edge (p, q) if and only if C_p and C_q intersect in at least two points.

Minimum-weight triangulation $MWT(V)$: A geometric triangulation (i.e., an edge maximal planar straight-line graph) with vertex set V and the minimum total edge length.

BASIC STRUCTURES

Let V be a finite set in the Euclidean plane. The *nearest neighbor graph* $NNG(V)$ connects each point in V to its nearest neighbor. It is usually defined as a directed graph, but some authors treat it as undirected. In general, $NNG(V)$ is not connected, but each point in V has at least one incident edge.

The *relative neighborhood graph* $RNG(V)$ connects two points p and q if and only if the *lune* $L(p, q)$ is empty of points from V . It was defined by Toussaint [Tou80]. The *Gabriel graph* $GG(V)$ was first introduced by Gabriel and Sokal [GS69]. It is defined similarly as $RNG(V)$, but two points p and q are connected by an edge if and only if their *diameter sphere* (i.e., the sphere with diameter pq) is empty. The RNG and the GG are always connected. (The RNG can be disconnected if one defines the exclusion region as a *closed* set.)

The β -skeletons are a continuous generalization of the Gabriel graph and the relative neighborhood graph [KR85]. They come in two variants, *circle-based* and *lune-based*, depending on the region that needs to be empty for an edge to be present. Both circle- and lune-based β -skeletons depend on a parameter $\beta \geq 0$. In circle-based β -skeletons, the union of two open generalized diameter circles needs to be empty of other points from V . In lune-based β -skeletons, an open β -lune needs to be empty; see Figure 32.1.1. The lune-based β -skeleton can be defined for any L_p -metric. Unless stated otherwise, we refer to the Euclidean case. For $\beta \in [0, 1]$, the circle-based and the lune-based β -skeleton coincide. For $\beta = 0$, the β -skeleton is the complete graph, provided that no three points of V lie on a line. For $\beta = 1$, we have $G_1^c(V) = G_1^l(V) = GG(V)$. For $\beta > 1$, the circle-based β -skeleton is a subgraph of the (Euclidean) lune-based β -skeleton. For $\beta = 2$, the lune-based β -skeleton coincides with the relative-neighborhood graph. For $\beta = \infty$, the circle-based β -skeleton becomes the empty graph and the lune-based β -skeleton becomes the *empty-strip graph*. For $0 \leq \beta_1 \leq \beta_2 \leq \infty$, we have $G_{\beta_2}^l(V) \subseteq G_{\beta_1}^l(V)$ and $G_{\beta_2}^c(V) \subseteq G_{\beta_1}^c(V)$.

These graph definitions capture the internal structure of a point set and are motivated by various applications, such as computer vision, texture discrimination, geographic analysis, pattern analysis, cluster analysis, and others. The following theorem states some relationships between proximity graphs. A version of this theorem was first established by Toussaint [Tou80]; see also [KR85, O'R82, MS80].




FIGURE 32.1.1

The exclusion regions for the circle-based and Euclidean lune-based β -skeleton for various values of β (cf. [Vel91]). For $\beta = 1$, the exclusion region coincides with the exclusion region for the Gabriel graph. For $\beta = 2$, the exclusion region for the lune-based β -skeleton is the exclusion region of the relative neighborhood graph. For $\beta = \infty$, the exclusion region of the lune-based β -skeleton is the exclusion region of the empty strip graph and the exclusion region of the circle-based β -skeleton is the whole plane.

THEOREM 32.1.1 Hierarchy Theorem

In any L_p metric, $p \in (1, \infty)$, for any finite point set V and for any $1 \leq \beta \leq 2$, we have

$$\text{RNG} \subseteq \text{MST}_p \subseteq \text{RNG} \subseteq G_\beta^l \subseteq \text{GG} \subseteq \text{DT}_p,$$

where MST_p is a minimum spanning tree of V in the L_p -norm and DT_p is the Delaunay triangulation of V .

O'Rourke showed that for $p = 1$ and $p = \infty$, the inclusion $\text{RNG} \subseteq \text{DT}_p$ does not necessarily hold; however, with a slightly different definition of Delaunay triangulation (in terms of empty open balls, instead of being the dual of the Voronoi diagram), the inclusion can be rescued [O'R82]. The MST is always connected, so by the Hierarchy Theorem, this also holds for the RNG, the GG and the lune-based β -skeleton with $\beta \in [0, 2]$. In general, the circle-based β -skeleton is not connected for $\beta > 1$.

Clearly, neighborhood graphs on n vertices can have at most $\binom{n}{2}$ edges. In many cases, this is also attained, for example, for the L_1 and L_∞ metric [Kat88], for Gabriel graphs in three and more dimensions [CEG⁺94], and for RNGs in four dimensions and higher. In the plane, the Euclidean RNG has at most $3n - 8$ edges [BDH⁺12]; an earlier upper bound of $3n - 10$ [Urq83] turned out to be incorrect [BDH⁺12]. The planar Gabriel graph has at most $3n - 8$ edges [MS80]. For $1 < p < \infty$, the fact that the RNG is contained in the Delaunay triangulation yields an upper bound of $3n - 6$ edges [JT92]. In three dimensions, Euclidean RNGs have size at most $O(n^{4/3})$ [AM92]. No matching lower bound is known. See Table 32.1.1. Bose et al. [BDH⁺12] give bounds on a large number of parameters in various proximity graphs.

TABLE 32.1.1 Size of RNGs and Gabriel graphs.

DIM	METRIC	SIZE	REFERENCE	COMMENT
2	L_2	$\leq 3n - 8$	[BDH ⁺ 12]	Gabriel Graphs
2	L_2	$\leq 3n - 8$	[MS80]	
2	$L_p, 1 < p < \infty$	$\in [n - 1, 3n - 6]$	[JT92]	
≥ 2	L_1, L_∞	$\Theta(n^2)$	[Kat88]	
3	L_2	$O(n^{4/3})$	[AM92]	Gabriel Graphs
≥ 3	L_2	$\Theta(n^2)$	[CEG ⁺ 94]	
≥ 4	L_2	$\Theta(n^2)$	[JT92]	

TABLE 32.1.2 RNG algorithms.

DIM	METRIC	COMPLEXITY	REFERENCE	COMMENT
2	L_2	$O(n \log n)$	[Sup83]	arbitrary points
2	$L_p, 1 < p < \infty$	$O(n^2)$	[Tou80, Kat88]	arbitrary points
2	L_1, L_∞	$O(n \log n)$	[Lee85]	general position
2	L_1, L_∞	$O(n \log n + m)$	[Kat88]	m output size arbitrary points
3	L_2	$O(n^{3/2+\epsilon})$	[AM92]	general position
3	L_2	$O(n^{7/4+\epsilon})$	[AM92]	arbitrary points
3	$L_p, 1 < p < \infty$	$O(n^2)$	[Sup83, JK87, KN87]	general position
3	L_1, L_∞	$O(n \log^2 n)$	[Smi89]	general position
d	L_2	$O(n^{2(1-1/(d+1))+\epsilon})$	[AM92]	general position
d	L_1, L_∞	$O(n \log^{d-1} n)$	[Smi89]	general position

ALGORITHMS

It is an interesting algorithmic problem to construct proximity graphs efficiently. Using the definition, $O(n^3)$ time complexity is trivial. In the case of general L_p -metrics, for $1 < p < \infty$, the fact that the Delaunay triangulation is a superset of the RNG leads to an $O(n^2)$ time algorithm in the plane. A faster algorithm for the Euclidean case was given by Supowit, who showed that in this case the RNG can be computed in $O(n \log n)$ time [Sup83]. For the L_1 and L_∞ metric, Lee described an $O(n \log n)$ time algorithm for planar point sets in general position, improving a previous $O(n^2 \log n)$ algorithm by O'Rourke [O'R82]. Katajainen [Kat88] gives an output-sensitive algorithm for the L_1 and L_∞ that achieves $O(n \log n + m)$ time, where m is the size of the resulting RNG. In three dimensions, Agarwal and Matoušek obtain $O(n^{3/2+\epsilon})$ time for computing the Euclidean RNG of points in general position, and $O(n^{7/4+\epsilon})$ for arbitrary points [AM92]. Their approach generalizes to higher dimensions, yielding time $O(n^{2(1-1/(d+1))+\epsilon})$ for d -dimensional point sets in general position. For the L_p -norm, $1 < p < \infty$, several algorithms with running time $O(n^2)$ are available for three-dimensional point sets in general position [Sup83, JK87, KN87]. Finally, in the L_1 and L_∞ -norm, the d -dimensional RNG for points in general position can be found in time $O(n \log^{d-1} n)$ [Smi89]. See Table 32.1.2.

There are also many results that describe algorithmic relationships between different proximity structures. Since planar graphs are closed under the minors relation, one can use the Borůvka-Sollin algorithm to find an MST of a planar point set V in $O(n)$ time, once $DT(V)$ is available [CT76]. This works in any L_p -norm, $1 < p < \infty$. Similarly, given $DT(V)$, we can compute the lune-based β -skeleton for V with $\beta \in [1, 2]$ in $O(n)$ additional time [Lin94]. This result holds in any L_p -metric, for $1 < p < \infty$ [Lin94]. By setting $\beta = 1$ and $\beta = 2$, this result applies in particular to $GG(V)$ and $RNG(V)$. A linear time algorithm to construct $GG(V)$ from $DT(V)$ was also described by Matula and Sokal [MS80]. The circle-based β -skeleton for a planar point set V with $\beta \geq 1$ can be found in $O(n)$ additional time given $DT(V)$ [KR85, Vel91]. The fastest algorithm for computing a β -skeleton for a planar point set with $\beta \in [0, 1]$ takes $O(n^2)$ time [HLM03]. In some cases, the β -skeleton for $\beta \in [0, 1]$ can have $\Theta(n^2)$ edges. Given the Euclidean MST (or any connected subgraph of $DT(V)$), we can compute $DT(V)$ in $O(n)$ additional time [CW98, KL96]. There exists a general reduction from computing Delaunay triangulations to computing NNGs: Suppose we are given an algorithm that computes the NNG of any planar m -point set in $T(m)$ time, where $m \mapsto T(m)/m$ is monotonically increasing. Then, we can compute the Delaunay triangulation of a planar n -point set in $O(T(n))$ expected time [BM11]. This reduction is useful in settings where faster NNG algorithms are available, e.g., in transdichotomous models that allow manipulations at the bit-level [BM11].

APPLICATIONS AND VARIANTS

An important connection between the circle-based β -skeleton and the minimum-weight triangulation (MWT) was discovered by Keil: for $\beta = \sqrt{2} \approx 1.41421$, we have $G_\beta^c(V) \subseteq \text{MWT}(V)$ [Kei94]. Cheng and Xu later improved this to $\beta = \sqrt{1 + \sqrt{4/27}} \approx 1.17682$ [CX01]. For $\beta = \sqrt{5/4 + \sqrt{1/108}} \approx 1.16027$, the circle-based β -skeleton need not be a subgraph of the MWT [WY01]. Even though it is NP-hard to compute the minimum weight triangulation [MR08], the β -skeleton provides a good heuristic for well-behaved point sets.

Just as Delaunay triangulations/Voronoi diagrams have been generalized to k th-order diagrams (see Chapter 27), the relative neighborhood graph and the Gabriel graph have k th-order generalizations, k -RNG and k -GG, in which the exclusion region may contain up to k points of V . The k -GG has $O(k(n-k))$ edges and can be constructed in time $O(k^2 n \log n)$ [SC90]; the k -GG is $(k+1)$ -connected [BCH⁺13], and the 10-GG is Hamiltonian (while the 1-GG is not necessarily Hamiltonian) [KSVC15]. The 17-RNG contains the Euclidean bottleneck matching, which leads to an efficient (roughly $O(n^{1.5})$) algorithm for computing a bottleneck matching [CTL92].

There are many ways to generalize the proximity graphs described in this section. The *sphere-of-influence* (SIG) graph was defined by Toussaint as a graph-theoretical “primal sketch” [Tou88]. In the SIG, two points are connected by an edge if and only if their nearest-neighbor circles intersect. In the plane, the SIG has at most $15n = O(n)$ edges [Sos99], and it can be computed in $\Theta(n \log n)$ time [AH85]. Veltkamp defines a family of γ -neighborhood graphs [Vel91]. Veltkamp’s graphs are parameterized by two parameters γ_0 and γ_1 , and they provide a common generalization for the Delaunay triangulation, the convex hull, the Gabriel graph, and

the circle-based β -skeleton [Vel91]. *Empty-ellipse graphs* [DEG08] are a more recent variant of proximity graphs. They were defined by Devillers, Erickson and Goaoc to study the local behavior of Delaunay triangulations on surfaces in three-dimensional space. Here, two points are connected if and only if they lie on an axis-aligned ellipse with no other points from V in its interior. Devillers et al. show that the empty ellipse graph for a point set with stretch Φ has $O(n\Phi)$ edges [DEG08]. Cardinal, Collette and Langerman reverse the viewpoint of previous work on proximity graphs [CCL09]: in an *empty region graph*, two points form an edge if and only if some neighborhood around them—derived from a *template region*—is empty. Instead of analyzing the properties of certain given proximity graphs, Cardinal et al. start with certain desirable graph properties, such as connectivity, planarity, bipartiteness, or cycle-freeness, and they characterize maximal and minimal template regions that ensure these properties.

The field of *proximity drawings* studies which graphs can be represented as proximity graphs. For example, a tree can be represented as an RNG if and only if it has maximum vertex degree at most 5 [BLL96]. There is also a characterization of the trees representable as Gabriel graphs [BLL96]. Refer to [DBLL94, Lio13] for many more results. Chapter 54 discusses applications of proximity graphs in pattern recognition.

OPEN PROBLEMS

1. The complexity of Euclidean RNGs in \mathbb{R}^3 has still not been settled. Agarwal and Matoušek showed an upper bound of $O(n^{4/3})$, where n is the number of points [AM92]. No super-linear lower bound is known.
2. What is the complexity of the SIG? The best upper bound for n vertices is $15n$, but no lower bound exceeding $9n$ is known [Sos99]. The SIG has a linear number of edges in any fixed dimension [GPS94], and bounds on the expected number of edges are known [Dwy95]. However, there are no tight results.

32.2 QUADTREES AND WSPDS

GLOSSARY

Quadtree T associated with a set $S \subset \mathbb{R}^d$: A tree in which each inner node has exactly 2^d children, with each node ν having an associated subset $S(\nu) \subseteq S$ and an axis-parallel bounding hypercube $R(\nu)$ for $S(\nu)$, such that (i) $|S(\nu)| = 1$ if ν is a leaf; and (ii) for each internal node ν , the hypercubes for the 2^d children of ν constitute a partition of $R(\nu)$ into 2^d congruent hypercubes.

Compressed quadtree T associated with a set $S \subset \mathbb{R}^d$: A tree in which each inner node has exactly one or 2^d children, with each node ν having an associated subset $S(\nu) \subseteq S$ and an axis-parallel bounding hypercube $R(\nu)$ for $S(\nu)$, such that (i) $|S(\nu)| = 1$ if ν is a leaf; (ii) if ν has 2^d children, the hypercubes for the 2^d children partition $R(\nu)$ into 2^d congruent hypercubes; and (iii) if ν has

one child ν' , then $R(\nu') \subset R(\nu)$ and $R(\nu')$ is smaller than $R(\nu)$ by at least a constant factor. Usually, a compressed quadtree is obtained from a quadtree by contracting long paths in which each node has only one non-empty child square.

Fair-split tree T associated with a set $S \subset \mathbb{R}^d$: A binary tree, where each node ν has an associated subset $S(\nu) \subseteq S$ and the axis-parallel bounding box $R(\nu)$ of $S(\nu)$, such that (i) $|S(\nu)| = 1$ if ν is a leaf; and (ii) for each internal node ν , let ν_1 and ν_2 be the two children of ν . Then, there exists a hyperplane h_ν orthogonal to the longest edge, ξ , of $R(\nu)$ separating $S(\nu_1)$ and $S(\nu_2)$ such that h_ν is at distance at least $|\xi|/3$ from each of the sides of $R(\nu)$ parallel to it.

s -well-separated pair: Let $s \geq 1$ be a fixed *separation constant*. Two non-empty point sets X and Y constitute an s -well-separated pair if and only if there are two radius- r enclosing balls, $B_X \supset X$ and $B_Y \supset Y$, such that the distance between B_X and B_Y is at least sr .

Well-separated pair decomposition (WSPD) of a set $S \subset \mathbb{R}^d$ of points for a fixed separation constant $s \geq 1$: A set, $\{\{A_1, B_1\}, \{A_2, B_2\}, \dots, \{A_m, B_m\}\}$, of pairs of nonempty subsets of S such that (i) $A_i \cap B_i = \emptyset$, for $i = 1, 2, \dots, m$; (ii) each pair of distinct elements $\{a, b\} \subseteq S$ has a unique pair $\{A_i, B_i\}$ with $a \in A_i$, $b \in B_i$; and (iii) A_i and B_i are s -well-separated. The **size** of the WSPD is m .

Doubling dimension of a metric space (S, δ) : The doubling parameter $\lambda \in \mathbb{N}$ is the smallest integer such that for every $r \geq 0$ and every $p \in S$, the ball $B(p, r)$ can be covered by at most λ balls of radius $r/2$. The *doubling dimension* of (S, δ) is $\log \lambda$. A family of metric spaces has *bounded doubling dimension* c if the doubling dimension of all spaces in the family is at most c .

(Unit) disk graph: The graph with vertex set V where each $p \in V$ has an *associated radius* $r_p > 0$. There is an edge (p, q) if and only if $\delta(p, q) \leq r_p + r_q$, i.e., if the closed balls with radius r_p around p and with radius r_q around q intersect. The graph is called a *unit disk graph* if $r_p = 1/2$, for all $p \in V$.

s -semi-separated pair: Let $s > 1$ be a fixed *separation constant*. Two non-empty point sets X and Y constitute an s -semi-separated pair if and only if there are two enclosing balls, $B_X \supset X$ and $B_Y \supset Y$ with radius r_A and r_B , respectively such that the distance between B_X and B_Y is at least $s \min\{r_A, r_B\}$.

Semi-separated pair decomposition (SSPD) of a set $S \subset \mathbb{R}^d$ of points for a fixed separation constant $s > 1$: A set, $\{\{A_1, B_1\}, \{A_2, B_2\}, \dots, \{A_m, B_m\}\}$, of pairs of nonempty subsets of S such that (i) $A_i \cap B_i = \emptyset$, for $i = 1, 2, \dots, m$; (ii) each pair of distinct elements $\{a, b\} \subseteq S$ has a unique pair $\{A_i, B_i\}$ with $a \in A_i$, $b \in B_i$; and (iii) A_i and B_i are s -semi-separated. The **size** of the SSPD is m .

QUADTREES

For every set S of n points, there is a quadtree with $O(n \log \Phi(S))$ nodes and depth $O(\log \Phi(S))$, and it can be computed in the same time [BCKO08, Sam90]. In general, the size and depth of a (regular) quadtree can be unbounded in n . To address this issue, one can define *compressed quadtrees*. The precise definition of a compressed quadtree varies in the literature [HP11, LM12], but the essential idea is to take a (regular) quadtree and to contract long paths in which each node has




FIGURE 32.2.1

A regular and a compressed quadtree for a planar point set.

only a single non-empty child into single edges; see Figure 32.2.1. For every set S of n points, there is a compressed quadtree with $O(n)$ nodes. The depth may also be $\Theta(n)$. A compressed quadtree can be computed in $O(n \log n)$ time [HP11, LM12]. In fact, given the Delaunay triangulation $DT(S)$ of a planar point set S , we can find a compressed quadtree for S in $O(n)$ additional time [KL98]. Conversely, given a suitable compressed quadtree for a planar point set S , we can find $DT(S)$ in $O(n)$ additional time [BM11, LM12].

COMPUTATIONAL MODELS

When dealing with compressed quadtrees, it is important to keep the computational model in mind. Algorithms that use compressed quadtrees often rely on the real RAM model of computation and require the floor function $x \mapsto \lfloor x \rfloor$, $x \in \mathbb{R}$ [PS85]. In fact, Har-Peled pointed out that if we want the squares of a compressed quadtree to be aligned to a grid, some kind of non-standard operation is inevitable [HP11]. Nonetheless, the floor function provides unexpected computational power. It can be used to circumvent established lower bounds in the algebraic decision tree model [Ben83]. For example, using the floor function, we can use Rabin's algorithm to find the closest pair of a set of n points in $O(n)$ expected time, despite an $\Omega(n \log n)$ lower bound for algebraic decision trees [Ben83]. Not only that, the floor function lets us solve PSPACE-complete problems in polynomial time [Sch79]. Despite these issues, algorithms that use the floor function are often simple, efficient and practical. There is also a way to define compressed quadtrees in a way that is compatible with the algebraic decision model, but this usually comes at the cost of increased algorithmic complexity [BLMM11, LM12]. When comparing results that involve quadtrees, we should be aware of the details of the underlying computational model.

WELL-SEPARATED PAIR DECOMPOSITION

Callahan and Kosaraju [CK95] defined the notion of a well-separated pair decomposition (WSPD) for a point set S . They also showed the remarkable theorem that a WSPD of size $O(n)$ can be constructed in time $O(n)$, given a fair split tree of an




FIGURE 32.2.2

Let $s > 1$. Any WSPD for the shortest path metric of the star graph with separation parameter s has size $\Omega(n^2)$. For every pair of distinct vertices v_i, v_j , there must be a distinct s -well-separated pair $\{A, B\}$ with $v_i \in A, v_j \in B$: if there were another v_k with, say, $v_k \in A$, then A would have diameter 2 and distance at most 2 from B , making $\{A, B\}$ not s -well-separated.

input set S of n points in \mathbb{R}^d , for any fixed dimension d and separation constant $s \geq 1$. (More precisely, the size of the WSPD is $O(s^d n)$.) A fair split tree can be constructed using quadtree methods in time $O(n \log n)$ for any fixed dimension. Alternatively, the WSPD can be computed from a (compressed) quadtree in $O(n)$ additional time [BM11, Cha08, HP11, LM12]. In the algebraic decision tree model, it takes $\Omega(n \log n)$ steps to compute the WSPD.

Well-separated pair decompositions have countless applications in proximity problems [Smi07]. For example, let S be a set of n points. Given a WSPD for S of size m with separation parameter $s > 2$, we can find a closest pair in S in $O(m)$ additional time, since the closest pair occurs as a well-separated pair in the decomposition [CK95]. In fact, given a fair-split tree or a compressed quadtree that represents a WSPD for S with separation parameter $s > 2$, we can compute $\text{NNG}(S)$ in $O(n)$ additional time [CK95]. This fact, together with the connection between fast algorithms for $\text{NNG}(S)$ and fast algorithms for $\text{DT}(S)$, can be used to obtain improved running times for computing Delaunay triangulations in various models of computation, such as the word RAM [BM11]. WSPDs are also extremely useful in the context of approximation algorithms. For example, they can be used to approximate the diameter and the minimum-spanning tree of high-dimensional point sets. As we will see in the next section, they also play an important role in spanner construction. The survey of Smid contains further applications [Smi07].

Not every metric space admits a WSPD of subquadratic size. For example, in the shortest path metric of the unweighted star graph, every WSPD with separation parameter $s > 1$ must have $\Omega(n^2)$ pairs; see Figure 32.2.2. Notwithstanding, there are large families of finite metrics that have WSPDs with linear or near-linear size.

A family of metric spaces has *bounded doubling dimension* c if the doubling dimension of all spaces in the family is at most c . This notion was defined by Gupta, Krauthgamer and Lee [GKL03], following earlier work by Assouad [Ass83]. The family of finite subsets of \mathbb{R}^d has bounded doubling dimension $\Theta(d)$. Talwar gave an algorithm to compute a WSPD with separation parameter $s \geq 1$ of size $O(s^{\log \lambda} n \log \Phi(S))$, where $\Phi(S)$ is the spread of S [Tal04]. This was improved by Har-Peled and Mendel, who showed how to compute a WSPD of size $O(s^{\log \lambda} n)$ in time $O(\lambda n \log n + s^{\log \lambda} n)$, asymptotically matching the bounds for the Euclidean case [HP11, HPM06].

Another interesting metric is given by *unit disk graphs*. In a unit disk graph,

there is an edge (p, q) between two vertices if and only if $\delta(p, q) \leq 1$. Even though the shortest path metric in unit disk graphs (with Euclidean edge lengths) does not have bounded doubling dimension, Gao and Zhang showed that it admits WSPDs of near-linear size [GZ05]. More specifically, they showed that the planar unit disk graphs with n vertices have WSPDs of size $O(s^4 n \log n)$ that can be computed in the same time. In dimension $d \geq 3$, there is a WSPD of size $O(n^{2-2/d})$ that can be found in time $O(n^{4/3} \log^{O(1)} n)$ for $d = 3$ and in time $O(n^{2-2/d})$ for $d \geq 4$. The bound on the size is tight for $d \geq 4$. The result of Gao and Zhang can be extended to general disk graphs where the ratio between the largest and the smallest radius is bounded [GZ05, Wil16]. For an unbounded radius ratio, however, no WSPD of subquadratic size is possible in general. Generalizations to the shortest path metric in *unweighted* disk graphs (also known as *hop distance*) are also possible, albeit with somewhat weaker results [GZ05].

For Euclidean point sets, there is always a WSPD with a linear number of pairs, but the total number of points in the sets of the pairs may be quadratic. In many applications, this is not an issue, because the sets can be represented implicitly and the algorithms do not need to inspect all sets in the WSPD. However, if this becomes necessary, a quadratic running time becomes hard to avoid. To address this, Varadarajan [Var98] introduced the notion of *s-semi-separated pair decomposition* (*s-SSPD*). In d dimensions, an *s-SSPD* with $O(s^d n)$ pairs whose sets contain $O(s^d n \log n)$ points in total can be computed in $O(s^d n \log n)$ time [ABFG09, AH12]. SSPDs have numerous applications, e.g., in computing the min-cost perfect matching of a planar point set [Var98] or in constructing spanners with certain properties [ABFG09, AH12, ACFS13]. Abu-Affash et al. [AACKS14] introduce the (α, β) -*pair decomposition*, another variant of the WSPD. They provide several applications, including its application to the Euclidean bottleneck Steiner path problem.

In general, the WSPD is the method of choice when we need to represent approximately the pairwise distances in a point set. The SSPD offers weaker guarantees, but it can be useful if we need to inspect all sets of the decomposition explicitly. The (α, β) -pair decomposition does not give a general approximation of distances, but it works only at a fixed scale. It is simpler to work with than the other two decompositions and it can provide stronger guarantees for certain problems.

OPEN PROBLEMS

1. What is the right bound of the size for WSPDs for unit disk graphs in the plane? Gao and Zhang's result gives an upper bound of $O(n \log n)$ [GZ05], but no super-linear lower bound is known.
2. Disk graphs with an unbounded radius ratio generally do not admit a WSPD of subquadratic size. Is there another way to represent the pairwise distances in these graphs compactly?

32.3 GEOMETRIC SPANNERS

GLOSSARY

Euclidean graph: A geometric graph with Euclidean lengths associated with the edges.

Complete Euclidean graph \mathcal{E}_d : A d -dimensional Euclidean graph (V, E) whose edge set E joins each pair of points of $V \subset \mathbb{R}^d$.

Yao-graph Y_k : For integer $k \geq 2$, a geometric graph in which each $v \in V$ is joined by an edge to the closest point $u \in V \cap C_i$, where, in dimension $d = 2$, each C_i is one of $k = 2\pi/\theta$ equal-sized **sectors** (cones) with apex v and angle θ .

Theta-graph Θ_k : For integer $k \geq 2$, a geometric graph, similar to the Yao-graph, in which each $v \in V$ is joined by an edge to the “closest” point $u \in V \cap C_i$, where “closest” is based on the projections of the points $V \cap C_i$ onto a ray with apex v within C_i , a sector of angle $2\pi/k$; typically, the ray is the bisector of C_i .

t -Spanner: A subgraph $G' = (V, E')$ of a graph $G = (V, E)$ such that for any $u, v \in V$ the distance $\delta_{G'}(u, v)$ within G' is at most t times the distance $\delta_G(u, v)$ within G . We focus on **Euclidean t -spanners** for which the underlying graph G is the complete Euclidean graph in \mathbb{R}^d .

Plane t -spanner: A Euclidean t -spanner that is a PSLG in \mathbb{R}^2 . (Also known as a *planar t -spanner*.)

Stretch factor, t^* , of a Euclidean graph $G = (V, E)$:

$$t^* = \max_{u, v \in V, u \neq v} \left\{ \frac{\delta_G(u, v)}{\delta_2(u, v)} \right\}$$

where $\delta_2(u, v)$ is the Euclidean distance between u and v . Thus, t^* is the smallest value of t for which G is a Euclidean t -spanner. The stretch factor is also known as the *spanning ratio* or the *dilation* of G .

Size of a Euclidean graph $G = (V, E)$: The number of edges, $|E|$.

Weight of a Euclidean graph $G = (V, E)$: The sum of the Euclidean lengths of all edges $e \in E$.

Degree of a graph $G = (V, E)$: The maximum number of edges incident on a common vertex $v \in V$.

k -vertex fault-tolerant t -spanner: A t -spanner with the property that the removal of any subset of at most k nodes, along with the incident edges, results in a subgraph that remains a t -spanner on the remaining set of points.

t -SPANNERS

A natural greedy algorithm, similar to Kruskal’s minimum spanning tree algorithm, can be used to construct t -spanners:

Given an input geometric graph $G = (V, E)$ and a real number $t > 1$. Initialize edge set $E' \leftarrow \emptyset$. For each edge $(u, v) \in E$, considered in

nondecreasing order of length $\delta_2(u, v)$, if $\delta_{G'}(u, v) > t \cdot \delta_2(u, v)$, then $E' \leftarrow E' \cup \{(u, v)\}$. Output the graph $G' = (V, E')$.

The greedy algorithm results in a t -spanner of size $O(n)$, weight $O(\log n) \cdot |\text{MST}|$, and degree $O(1)$, for any fixed dimension d and spanning ratio $t > 1$ [ADD⁺93, CDNS95]. It can be applied also to general (nongeometric) graphs with weighted edges. The greedy spanner can be computed in $O(n^2 \log n)$ time with $O(n^2)$ space [BCF⁺10] and in $O(n^2 \log^2 n)$ time with $O(n)$ space [ABBB15].

The Yao-graphs Y_k and theta-graphs Θ_k explicitly take advantage of geometry, and each yields a t -spanner with spanning ratio $t = 1 + O(1/k)$ arbitrarily close to 1, for sufficiently large k [ADD⁺93, Cla87, Kei92, RS91, Yao82]. The Yao-graphs Y_k and theta-graphs Θ_k are connected for $k \geq 2$; see [ABB⁺14].

TABLE 32.3.1 Spanning ratios of Yao-graphs and Θ -graphs in \mathbb{R}^2 .

GRAPH	SPANNING RATIO	REFERENCE
Y_2, Y_3	unbounded	[EM09]
Y_4	$\leq 8(29 + 23\sqrt{2}) \approx 663$	[BDD ⁺ 12]
Y_5	$\in [2.87, 2 + \sqrt{3} \approx 3.74]$	[BBD ⁺ 15]
Y_6	$\in [2, 5.8]$	[BBD ⁺ 15]
$Y_k, k \geq 7$	lower/upper bounds depending on $k \pmod{4}$	[BBD ⁺ 15]
Θ_2, Θ_3	unbounded	[EM09]
Θ_4	[7, 237]	[BBC ⁺ 13]
Θ_5	[3.79, 9.96]	[BMRV15]
$\Theta_k, k \geq 6, k \equiv 2 \pmod{4}$	$1 + 2 \sin(\pi/k)$ (tight)	[BDCM ⁺ 16, BGHI10]
$\Theta_k, k \geq 7$	lower/upper bounds depending on $k \pmod{4}$	[BDCM ⁺ 16]

The known bounds on spanning ratios of Yao-graphs and Θ -graphs are shown in Table 32.3.1; see also the detailed tables in [BDCM⁺16] for Θ -graphs and in [BBD⁺15] for Yao-graphs. Note that Y_5 has spanning ratio $3.74 < 3.79$, making it the only known case ($k = 5$) in which there is a strict separation between the spanning ratio of a theta-graph Θ_k and that of a Yao-graph Y_k ; for other values of $k \geq 4$, it is not known which graph (Y_k or Θ_k) has a smaller spanning ratio. The spanning ratio for Θ -graphs does not necessarily decrease with an increase in k ; for $k = 6$, the spanning ratio is 2 (and this is tight), while for $k = 8$, it is known that the spanning ratio is at least 2.17 [BDCM⁺16].

By selecting a representative edge from each pair in a WSPD, one obtains a t -spanner of size $O(n)$ with a spanning ratio that can be made arbitrarily close to 1, depending on the separation constant s .

One can in fact obtain t -spanners for n points in \mathbb{R}^d that are simultaneously good with respect to size, weight, and degree—size $O(n)$, weight $O(|\text{MST}|)$, and bounded degree (independent of the dimension d). Gudmundsson et al. [GLN02] show that such spanners can be computed in time $O(n \log n)$, improving the previous bound of $O(n \log^2 n)$ [DN97] and re-establishing the time bound claimed in Arya et al. [ADM⁺95] (which was found to be flawed). $\Omega(n \log n)$ time is required for constructing *any* t -spanner for n points in \mathbb{R}^d in the algebraic decision tree

model [CDS01].

It was shown by Levcopoulos, Narasimhan, and Smid [LNS02] that k -vertex fault-tolerant spanners of size $O(k^2n)$ can be constructed in time $O(n \log n + k^2n)$; alternatively, spanners of size $O(kn \log n)$ can be constructed in time $O(kn \log n)$. Lukovszki [Luk99] and Czumaj and Zhao [CZ04] showed how to obtain even smaller, degree-bounded low-weight k -vertex fault-tolerant spanners; degree $O(k)$ and weight $O(k^2|MST|)$ can be obtained, and these bounds are asymptotically optimal.

The spanning ratio of a given graph $G = (V, E)$ can be computed exactly in worst-case time $O(n^2 \log n + n|E|)$ using an all-pairs shortest path computation. Given a Euclidean graph with n vertices and m edges, its spanning ratio (stretch factor) can be $(1 + \varepsilon)$ -approximated in time $O(m + n \log n)$ [GLNS08]. Narasimhan and Smid [NS02] have studied the *bottleneck stretch factor problem*, in which the goal is to be able to compute quickly, for any given $b > 0$, an approximate stretch factor of the *bottleneck graph* $G_b = (V, E_b)$ whose edge set E_b consists of those edges of the complete graph whose length is at most b . We say that t is a (c_1, c_2) -approximate stretch factor of a graph if the true stretch factor, t^* , satisfies $t/c_1 \leq t^* \leq c_2t$. A data structure of size $O(\log n)$ can be constructed that supports $O(\log \log n)$ -time queries, for any $b > 0$, yielding a (c_1, c_2) -approximate stretch factor of G_b . The construction of the data structure, which is based on a WSPD, is done using a randomized algorithm with expected running time that is slightly subquadratic.

Spanners can be computed for geodesic distances in a polygonal domain P : a $(1 + \varepsilon)$ -spanner of the visibility graph $VG(P)$ can be computed in time $O(n \log n)$, for any $\varepsilon > 0$ [ACC⁺96]. Geometric spanners can be used to obtain very efficient approximate two-point shortest path distance queries: for any constant $t > 1$, a t -spanner G for n points in \mathbb{R}^d with m edges can be processed in time $O(m \log n)$, building a structure of size $O(n \log n)$, to support $(1 + \varepsilon)$ -approximate shortest path (in G) distance queries in $O(1)$ time between any two vertices of G . (A path can be reported in additional time proportional to the number of its edges.) Then, if the visibility graph $VG(P)$ is a t -spanner of the vertices of P , for some constant t , one obtains $O(1)$ -time (resp., $O(\log n)$ -time) $(1 + \varepsilon)$ -approximate shortest path distance queries between any two vertices (resp., points) of P . The assumption on $VG(P)$ holds if P has the “ t -rounded” property for some t : the shortest path distance between any pair of vertices is at most t times the Euclidean distance between them; such is the case if the obstacles are *fat*, as shown by Chew et al. [CDKK02].

PLANE t -SPANNERS

For finite point sets in the plane it is natural to consider constructing *plane* t -spanner networks, whose edges do not cross. One cannot hope, in general, to obtain plane t -spanners with t arbitrarily close to 1: four points at the corners of a square have no plane t -spanner with $t < \sqrt{2}$.

The first result on plane t -spanners is due to Chew [Che86], who showed that the Delaunay triangulation in the L_1 metric is a $\sqrt{10}$ -spanner for the complete Euclidean graph. (It is a $\sqrt{5}$ -spanner for the complete graph whose edge lengths are measured in the L_1 metric.) Chew [Che89] improved this result, showing that the Delaunay triangulation in the convex distance function based on an equilateral triangle (also known as the triangular-distance Delaunay or TD-Delaunay graph) is a plane graph with spanning ratio at most 2; this bound is now known to be

tight [BGHI10]. This had been the best known spanning ratio for a plane t -spanner until the work of Xia [Xia13], who showed that the Euclidean Delaunay triangulation has a spanning ratio less than 1.998. The lower bound of $\sqrt{2}$ on the spanning ratio of a plane t -spanner, given by the four corners of a square, has been improved by Mulzer [Mul04] to 1.41611 (by considering vertices of a regular 21-gon), and then by Dumitrescu and Ghosh [DG16] to 1.4308 (by considering vertices of a regular 23-gon). For points in convex position a spanning ratio of 1.88 can always be achieved [ABB⁺16].

The spanning ratio, τ_{Del} , of the Euclidean (L_2) Delaunay triangulation is not less than $\pi/2$, as shown by the example of placing points around a circle. This lower bound has been improved recently to $1.5932 > \pi/2$ [XZ11]. Dobkin, Friedman, and Supowit [DFS90] were able to show that $\tau_{\text{Del}} \leq \phi\pi$, where $\phi = (1 + \sqrt{5})/2$ is the golden ratio. This upper bound was improved by Keil and Gutwin [KG92] to $\frac{4\pi}{3\sqrt{3}} \approx 2.42$. The current best known upper bound on τ_{Del} is 1.998, as shown by Xia [Xia13]. For the Delaunay triangulation in the L_1 metric, the original bound of $\sqrt{10}$ [Che86] on the spanning ratio has been improved to $\sqrt{4 + 2\sqrt{2}}$, and this is tight [BGHP15]. More generally, the Delaunay triangulation defined with respect to any convex distance function is a t -spanner [BCCS10].

Minimum spanning trees do not have a bounded spanning ratio. It is also known that lune-based β -skeletons, for any $\beta > 0$, can have an unbounded spanning ratio; see [Epp00]. Since for $\beta \geq 1$, the lune-based β -skeleton is a subgraph of the Gabriel graph ($\beta = 1$), it is a plane graph. In particular, the Gabriel graph ($\beta = 1$) and the relative neighborhood graph ($\beta = 2$) are not t -spanners for any constant t . Growth rates, as a function of n , for the spanning ratios of Gabriel graphs and β -skeletons for other values of β are given by [BDEK06].

The minimum weight triangulation and the greedy triangulation (see Chapter 31) are t -spanners for constant t . This follows from a more general result of Das and Joseph [DJ89], who show that a PSLG is a t -spanner if it has the “**diamond property**” and the “**good polygon property**.” This result is similar to the empty region graphs of Cardinal et al. discussed in Section 32.1, where certain graph properties were obtained by requiring that an edge is present if and only if certain template regions were empty [CCL09]. The difference is that the diamond property by Das and Joseph is only a necessary condition. A **fat** triangulation of S , for which the aspect ratio (ratio of the length of the longest side to the corresponding height) of every triangle is at most α , is known to be a 2α -spanner [KG01].

All of the plane spanners mentioned above have potentially unbounded degree. Bounded degree plane spanners are important in wireless network applications, especially in routing. One needs degree at least 3 to achieve a bounded spanning ratio (a Hamiltonian path has an unbounded spanning ratio). While Das and Heffernan [DH96] showed that t -spanners exist of maximum degree 3, their spanner is not necessarily plane. The best upper bounds on spanning ratio bounds currently known for plane spanners of degree at most δ are 20 for degree $\delta = 4$ [KPT16], 6 for $\delta = 6$ [BGHP10], 2.91 for $\delta = 14$ [KP08], and ≈ 4.414 for $\delta = 8$ [BHS16]; see [KPT16] for more details and a comprehensive table. While for points in convex position, degree-3 plane spanners are known [BBC⁺17, KPT16], it is not clear if there exist plane t -spanners of degree 3 for general point sets in the plane.

One can compute plane t -spanners of low weight. In linear time, for any $r > 0$, a plane t -spanner, with $t = (1+1/r)\tau_{\text{Del}}$, of weight at most $(2r+1)|\text{MST}|$ can be computed from a Delaunay triangulation, where τ_{Del} is the spanning ratio of the

Delaunay triangulation [LL92]. From any t -spanner, [GLN02] show that one can compute a subgraph of it that is low weight ($O(|\text{MST}|)$) and a t' -spanner (for a larger constant factor t'); thus, in order to find spanners that are of both bounded degree and of low weight, it suffices to focus on bounding the degree. One can compute in time $O(n \log n)$ a plane t -spanner that is simultaneously low weight ($O(|\text{MST}|)$) and low degree (degree at most k), with $t = (1 + 2\pi(k \cos(\pi/k))^{-1}) \cdot \tau_{\text{Del}}$ for any integer $k \geq 14$ [KPX10]; in particular, for degree 14, the stretch factor is at most ≈ 2.918 .

Planar t -spanners are also known for geodesic distances. A *conforming triangulation* for a polygonal domain P having triangles of aspect ratio at most α is a 2α -spanner for geodesic distances between vertices of P [KG01]. (A triangulation is *conforming* for P if each vertex of P is a vertex of the triangulation and each edge of P is the union of some edges of the triangulation.) The *constrained Delaunay triangulation* of P is a $\phi\pi$ -spanner [KG01].

NON-EUCLIDEAN METRICS

The WSPD-construction of Har-Peled and Mendel implies that for every $\varepsilon > 0$, a space of bounded doubling dimension d has a spanner with $n\varepsilon^{-O(d)}$ edges and stretch factor $1 + \varepsilon$ that can be found in $2^{O(d)}n \log n + n\varepsilon^{-O(d)}$ time [HPM06]. Independently, Chan et al. [CGMZ16] obtained a similar result. They showed the existence of a spanner with stretch factor $1 + \varepsilon$ in which every vertex has degree at most $\varepsilon^{-O(d)}$. Subsequently, several improved constructions of spanners for bounded doubling metrics were described [CLNS15]. The weight was also considered in this context. Smid showed that the greedy $(1 + \varepsilon)$ -spanner in spaces of bounded doubling dimension has $O(n)$ edges and weight $O(\log n |\text{MST}|)$ [Smi09]. Gottlieb provided an intricate construction of $(1 + \varepsilon)$ -spanners with weight $O(|\text{MST}|)$ and $O(n)$ edges [Got15]. Filtser and Solomon proved that this is also achieved by the greedy spanner: for any $\varepsilon > 0$, the greedy $(1 + \varepsilon)$ -spanner in a space of doubling dimension d has weight $(d/\varepsilon)^{O(d)}|\text{MST}|$ and $n(1/\varepsilon)^{O(d)}$ edges [FS16]. Moreover, an approximate version of the greedy spanner shows that for any $\varepsilon > 0$, one can construct in time $\varepsilon^{-O(d)}n \log n$ a $(1 + \varepsilon)$ -spanner with weight $(d/\varepsilon)^{O(d)}|\text{MST}|$ and degree $\varepsilon^{-O(d)}$ [FS16]. This matches the best result for the Euclidean case [GLN02].

Sparse spanners also exist for disk graphs. Fürer and Kasiviswanathan [FK12] used a modification of the Yao graph to show that for fixed $\varepsilon > 0$, every disk graph has a spanner with $O(n)$ edges and stretch factor $1 + \varepsilon$. They also described an algorithm to find such spanners in time $O(n^{4/3+\tau} \log^{2/3} \Psi)$, where $\tau > 0$ can be made arbitrarily small and Ψ is the *radius ratio* between the largest and the smallest radius of a vertex in V . Kaplan et al. extended this result to *transmission graphs*, a directed version of disk graphs in which each vertex $p \in V$ has an associated radius $r_p > 0$ and we have a directed edge (p, q) if and only if $\delta(p, q) \leq r_p$, i.e., if q lies in the closed radius- r_p disk around p [KMRS15]. For any fixed $\varepsilon > 0$, every transmission graph has a spanner with stretch factor $1 + \varepsilon$ and $O(n)$ edges. This spanner can be found in $O(n(\log n + \log \Psi))$ time, where Ψ is again the radius ratio [KMRS15, Sei16]. Alternatively, the spanner can be found in $O(n \log^5 n)$ time, independent of Ψ . The results of Kaplan et al. can also be applied to general disk graphs. Here, the time to construct the spanner described by Fürer and Kasiviswanathan [FK12] can be improved to $O(n2^{\alpha(n)} \log^{10} n)$ expected time, where $\alpha(n)$ is the inverse Ackermann function [KMR⁺17, Sei16].

OPEN PROBLEMS

1. What is the best possible spanning ratio for a plane t -spanner? It is known to be between 1.4308 and 1.998. For points in convex position an upper bound of 1.88 is known. The upper bound of 1.998 comes from the Euclidean Delaunay graph; is there a plane t -spanner with a spanning ratio better than the Delaunay?
2. Determine tight bounds for the spanning ratio of theta-graphs Θ_k for $k \geq 7$.
3. What exactly is the spanning ratio of the Euclidean Delaunay triangulation? It is known to be between 1.5932 and 1.998.
4. What are the best possible spanning ratios for bounded degree plane spanners, for various degree bounds? Are there plane spanners of bounded spanning ratio having degree at most 3?
5. For a given set of points in the plane, can one compute in polynomial time a plane graph having the minimum possible spanning ratio?
6. How efficiently can Yao and theta-graphs be constructed in higher dimensions?
7. Is it possible to compute the greedy spanner in subquadratic time?

32.4 DYNAMIC PROXIMITY ALGORITHMS

GLOSSARY

Additively weighted Euclidean distance: Let $S \subset \mathbb{R}^2$ be a set of sites, such that each $s \in S$ has an associated weight $w_s \in \mathbb{R}$. The additively weighted Euclidean distance $\delta : \mathbb{R}^2 \times S \rightarrow \mathbb{R}$ is defined as $\delta(p, s) = w_s + \delta_2(p, s)$.

Dynamic nearest neighbor: The problem of maintaining a set S of sites under the following operations: (i) insert a new site into S ; (ii) delete a site from S ; and (iii) given a query point q , find the site in S that minimizes the distance to q .

Dynamic bichromatic closest pair: The problem of maintaining two sets R and B of red and blue points such that points can be inserted into and deleted from R and B and such that we always have available a pair $(r, b) \in R \times B$ that minimizes the distance $\delta(r, b)$ among all pairs in $R \times B$.

In dynamic proximity algorithms, we would like to maintain some proximity structure for a point set that changes through insertions and deletions. The quintessential problem is the *dynamic nearest neighbor* problem: maintain a point set S under the following operations: (i) insert a new point into S ; (ii) delete a point from S ; (iii) given a query point q , determine a point $p \in S$ with $\delta(p, q) = \min_{r \in S} \delta(r, q)$.

For the Euclidean case, the first solution to this problem is due to Agarwal and Matoušek [AM95], who obtained amortized update time $O(n^\varepsilon)$, for every $\varepsilon > 0$, and worst-case query time $O(\log n)$. This was dramatically improved more than ten years later by Chan [Cha10]. His data structure achieves worst-case query time $O(\log^2 n)$ with amortized insertion time $O(\log^3 n)$ and amortized deletion time $O(\log^6 n)$ (Chan's original data structure was randomized, but a recent result of Chan and Tsakalidis yields a deterministic structure [CT16]). Kaplan et al. provide a variant of Chan's data structure that improves the amortized deletion time to $O(\log^5 n)$ [KMR⁺17]. Similar results hold for more general metrics. Extending the results by Agarwal and Matoušek [AM95] for the Euclidean case, Agarwal, Efrat, and Sharir describe a dynamic nearest neighbor structure with amortized update time $O(n^\varepsilon)$, for any fixed $\varepsilon > 0$, and worst-case query time $O(\log n)$ [AES99]. The result by Agarwal, Efrat, and Sharir holds for a wide range of distance functions, including L_p -metrics and the additively weighted Euclidean distance. This result was improved by Kaplan et al. [KMR⁺17]. Kaplan et al. built on Chan's data structure [Cha10] to construct a dynamic nearest neighbor structure for general metrics with worst-case query time $O(\log^2 n)$, amortized expected insertion time $O(\log^5 n \lambda_{s+2}(\log n))$ and amortized expected deletion time $O(\log^9 n \lambda_{s+2}(\log n))$. Here, s is a constant that depends on the metric under consideration, and $\lambda_t(\cdot)$ is the function that bounds the maximum length of a Davenport-Schinzel sequence of order t [SA95]. See Section 28.10 for more details on Davenport-Schinzel sequences.

Eppstein describes several reductions that provide fast dynamic algorithms for proximity problems once an efficient dynamic nearest neighbor structure is at hand [Epp95]. In particular, he showed that if there is a dynamic nearest neighbor structure whose queries all run in time $T(n)$, where $T(n)$ is monotonically increasing and $T(3n) = O(T(n))$, then the *dynamic bichromatic nearest neighbor problem* can be solved with amortized insertion time $O(T(n) \log n)$ and amortized deletion time $O(T(n) \log^2 n)$ [Epp95]. This implies that the MST of a planar point set can be maintained in $O(T(n) \log^4 n)$ amortized time per update (this result is not stated in the original paper, since it also needs a fast data structure for maintaining an MST in a general dynamic graph [HdLT01] that was not available when Eppstein wrote his paper).

Dynamic algorithms for geometric spanners have also been considered. The first result in this direction is due to Arya et al. [AMS99] who show how to construct a data structure of size $O(n \log^d n)$ that maintains a d -dimensional Euclidean spanner in $O(\log^d n \log \log n)$ expected amortized time per insertion and deletion in a model of random updates. A dynamic spanner by Gao et al. [GGN06] can handle arbitrary update sequences, with the performance bounds depending on the spread of the point set. The first dynamic spanner whose performance depends only on the number of points is due to Roditty [Rod12]. This result was improved several times [GR08a, GR08b]. An optimal construction was eventually obtained by Gottlieb and Roditty [GR08b]. Their spanner has stretch factor $1 + \varepsilon$, for any $\varepsilon > 0$, constant degree $\varepsilon^{-O(d)}$, and update time $\varepsilon^{-O(d)} \log n$. It also works in general metric spaces of constant doubling dimension d .

OPEN PROBLEMS

1. Can the Euclidean dynamic nearest neighbor problem be solved with amortized update time $O(\log n)$?

2. Can Eppstein's reduction from the bichromatic nearest neighbor problem to the dynamic nearest neighbor problem be improved?

32.5 SOURCES AND RELATED MATERIAL

SURVEYS

Several other surveys offer a wealth of additional material and references:

- [BE97]: A survey of approximation algorithms for geometric optimization problems.
- [BS13]: A survey on plane geometric spanners, with 22 highlighted open problems.
- [Epp00]: A survey of results on spanning trees and t -spanners.
- [GK07]: A survey on geometric spanners and related problems.
- [HP11]: A book on geometric approximation algorithms.
- [JT92]: A survey on relative neighborhood graphs.
- [Lio13]: A survey on proximity drawings.
- [NS07]: A book on geometric spanners.
- [Sam90]: A book on quadtrees and related structures.
- [Smi07]: A survey on well-separated pair decompositions.
- [Tou14]: A recent survey on sphere-of-influence graphs.

RELATED CHAPTERS

- Chapter 27: Voronoi diagrams and Delaunay triangulations
- Chapter 31: Shortest paths and networks
- Chapter 38: Point location
- Chapter 39: Collision and proximity queries
- Chapter 40: Range searching
- Chapter 54: Pattern recognition

REFERENCES

- [AACKS14] A.K. Abu-Affash, P. Carmi, M.J. Katz, and M. Segal. The Euclidean bottleneck Steiner path problem and other applications of (α, β) -pair decomposition. *Discrete Comput. Geom.*, 51:1–23, 2014.
- [ABB⁺14] O. Aichholzer, S.W. Bae, L. Barba, P. Bose, M. Korman, A. van Renssen, P. Taslakian, and S. Verdonschot. Theta-3 is connected. *Comput. Geom.*, 47:910–917, 2014.
- [ABB⁺16] M. Amani, A. Biniaz, P. Bose, J.-D. De Carufel, A. Maheshwari, and M. Smid. A plane 1.88-spanner for points in convex position. *J. Comput. Geom.*, 7:520–539, 2016.

- [ABBB15] S.P.A. Alewijnse, Q.W. Bouts, A.P. ten Brink, and K. Buchin. Computing the greedy spanner in linear space. *Algorithmica*, 73:589–606, 2015.
- [ABFG09] M.A. Abam, M. de Berg, M. Farshi, and J. Gudmundsson. Region-fault tolerant geometric spanners. *Discrete Comput. Geom.*, 41:556–582, 2009.
- [ACC⁺96] S.R. Arikati, D.Z. Chen, L.P. Chew, G. Das, M. Smid, and C.D. Zaroliagis. Planar spanners and approximate shortest path queries among obstacles in the plane. In *Proc. 4th Eur. Sympos. Algorithms*, vol. 1136 of *LNCS*, pages 514–528, Springer, Berlin, 1996.
- [ACFS13] M.A. Abam, P. Carmi, M. Farshi, and M. Smid. On the power of the semi-separated pair decomposition. *Comput. Geom.*, 46:631–639, 2013.
- [ADD⁺93] I. Althöfer, G. Das, D.P. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete Comput. Geom.*, 9:81–100, 1993.
- [ADM⁺95] S. Arya, G. Das, D.M. Mount, J.S. Salowe, and M. Smid. Euclidean spanners: short, thin, and lanky. In *Proc. 27th ACM Sympos. Theory Comput.*, pages 489–498, 1995.
- [AES99] P.K. Agarwal, A. Efrat, and M. Sharir. Vertical decomposition of shallow levels in 3-dimensional arrangements and its applications. *SIAM J. Comput.*, 29:912–953, 1999.
- [AH85] D. Avis and J. Horton. Remarks on the sphere of influence graph. *Ann. New York Acad. Sci.*, 440:323–327, 1985.
- [AH12] M.A. Abam and S. Har-Peled. New constructions of SSPDs and their applications. *Comput. Geom.*, 45:200–214, 2012.
- [AM92] P.K. Agarwal and J. Matoušek. Relative neighborhood graphs in three dimensions. *Comput. Geom.*, 2:1–14, 1992.
- [AM95] P.K. Agarwal and J. Matoušek. Dynamic half-space range reporting and its applications. *Algorithmica*, 13:325–345, 1995.
- [AMS99] S. Arya, D.M. Mount, and M. Smid. Dynamic algorithms for geometric spanners of small diameter: Randomized solutions. *Comput. Geom.*, 13:91–107, 1999.
- [Ass83] P. Assouad. Plongements lipschitziens dans \mathbb{R}^n . *Bull. Soc. Math. France*, 111:429–448, 1983.
- [BBC⁺13] L. Barba, P. Bose, J.-D. de Carufel, A. van Renssen, and S. Verdonechot. On the stretch factor of the Theta-4 graph. In *Proc. 13th Int. Sympos. Algorithms Data Structures*, vol. 8037 of *LNCS*, pages 109–120, Springer, Berlin, 2013.
- [BBC⁺17] A. Biniaz, P. Bose, J.-D. de Carufel, C. Gavoille, A. Maheshwari, and M. Smid. Towards plane spanners of degree 3. *J. Comput. Geom.*, 8:11–31, 2017.
- [BBD⁺15] L. Barba, P. Bose, M. Damian, R. Fagerberg, W.L. Keng, J. O’Rourke, A. van Renssen, P. Taslakian, S. Verdonechot, and G. Xia. New and improved spanning ratios for Yao graphs. *J. Comput. Geom.*, 6:19–53, 2015.
- [BCCS10] P. Bose, P. Carmi, S. Collette, and M. Smid. On the stretch factor of convex Delaunay graphs. *J. Comput. Geom.*, 1:41–56, 2010.
- [BCF⁺10] P. Bose, P. Carmi, M. Farshi, A. Maheshwari, and M. Smid. Computing the greedy spanner in near-quadratic time. *Algorithmica*, 58:711–729, 2010.
- [BCH⁺13] P. Bose, S. Collette, F. Hurtado, M. Korman, S. Langerman, V. Sacristán, and M. Saumell. Some properties of k -Delaunay and k -Gabriel graphs. *Comput. Geom.*, 46:131–139, 2013.
- [BCKO08] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars. *Computational Geometry. Algorithms and Applications*, third edition. Springer-Verlag, Berlin, 2008.

- [BDCM⁺16] P. Bose, J.L. de Carufel, P. Morin, A. van Renssen, and S. Verdonschot. Towards tight bounds on Theta-graphs: More is not always better. *Theoret. Comput. Sci.*, 616:70–93, 2016.
- [BDD⁺12] P. Bose, M. Damian, K. Douïeb, J. O’Rourke, B. Seamone, M. Smid, and S. Wuhrer. $\pi/2$ -angle Yao graphs are spanners. *Internat. J. Comput. Geom. Appl.*, 22:61–82, 2012.
- [BDEK06] P. Bose, L. Devroye, W. Evans, and D.G. Kirkpatrick. On the spanning ratio of Gabriel graphs and beta-skeletons. *SIAM J. Discrete Math.*, 20:412–427, 2006.
- [BDH⁺12] P. Bose, V. Dujmović, F. Hurtado, J. Iacono, S. Langerman, H. Meijer, V. Sacristán, M. Saumell, and D.R. Wood. Proximity graphs: E , δ , Δ , χ and ω . *Internat. J. Comput. Geom. Appl.*, 22:439–470, 2012.
- [BE97] M. Bern and D. Eppstein. Approximation algorithms for geometric problems. In D. S. Hochbaum, editor, *Approximation Algorithms for NP-hard problems*, pages 296–345, PWS Publishing Company, 1997.
- [Ben83] M. Ben-Or. Lower bounds for algebraic computation trees. In *Proc. 15th ACM Sympos. Theory Comput.*, pages 80–86, 1983.
- [BGHI10] N. Bonichon, C. Gavoille, N. Hanusse, and D. Ilcinkas. Connections between Theta-graphs, Delaunay triangulations, and orthogonal surfaces. In *Proc. 36th Workshop Graph Theoretic Concepts in Comp. Sci.*, vol. 6410 of *LNCS*, pages 266–278, Springer, Berlin, 2010.
- [BGHP10] N. Bonichon, C. Gavoille, N. Hanusse, and L. Perkovic. Plane spanners of maximum degree six. In *Proc. 37th Internat. Colloq. Automata Lang. Program.*, vol. 6198 of *LNCS*, pages 19–30, Springer, Berlin, 2010.
- [BGHP15] N. Bonichon, C. Gavoille, N. Hanusse, and L. Perkovic. Tight stretch factors for L_1 - and L_∞ -Delaunay triangulations. *Comput. Geom.*, 48:237–250, 2015.
- [BHS16] P. Bose, D. Hill, and M. Smid. Improved spanning ratio for low degree plane spanners. In *Proc. 12th Latin American Symp. Theoret. Informatics*, vol. 9644 of *LNCS*, pages 249–262, Springer, Berlin, 2016.
- [BLL96] P. Bose, W. Lenhart, and G. Liotta. Characterizing proximity trees. *Algorithmica*, 16:83–110, 1996.
- [BLMM11] K. Buchin, M. Löffler, P. Morin, and W. Mulzer. Preprocessing imprecise points for Delaunay triangulation: Simplified and extended. *Algorithmica*, 61:674–693, 2011.
- [BM11] K. Buchin and W. Mulzer. Delaunay triangulations in $O(\text{sort}(n))$ time and more. *J. ACM*, 58:6, 2011.
- [BMRV15] P. Bose, P. Morin, A. van Renssen, and S. Verdonschot. The θ_5 -graph is a spanner. *Comput. Geom.*, 48:108–119, 2015.
- [BS13] P. Bose and M. Smid. On plane geometric spanners: A survey and open problems. *Comput. Geom.*, 46:818–830, 2013.
- [CCL09] J. Cardinal, S. Collette, and S. Langerman. Empty region graphs. *Comput. Geom.*, 42:183–195, 2009.
- [CDKK02] L.P. Chew, H. David, M.J. Katz, and K. Kedem. Walking around fat obstacles. *Inform. Process. Lett.*, 83:135–140, 2002.
- [CDNS95] B. Chandra, G. Das, G. Narasimhan, and J. Soares. New sparseness results on graph spanners. *Internat. J. Comput. Geom. Appl.*, 5:125–144, 1995.
- [CDS01] D.Z. Chen, G. Das, and M. Smid. Lower bounds for computing geometric spanners and approximate shortest paths. *Discrete Appl. Math.*, 110:151–167, 2001.

- [CEG⁺94] B. Chazelle, H. Edelsbrunner, L.J. Guibas, J. Hershberger, R. Seidel, and M. Sharir. Selecting heavily covered points. *SIAM J. Comput.*, 23:1138–1151, 1994.
- [CGMZ16] T.-H.H. Chan, A. Gupta, B. M. Maggs, and S. Zhou. On hierarchical routing in doubling metrics. *ACM Trans Algorithms*, 12:55:1–55:22, 2016.
- [Cha08] T.M. Chan. Well-separated pair decomposition in linear time? *Inform. Process. Lett.*, 107:138–141, 2008.
- [Cha10] T.M. Chan. A dynamic data structure for 3-D convex hulls and 2-D nearest neighbor queries. *J. ACM*, 57:16, 2010.
- [Che86] L.P. Chew. There is a planar graph almost as good as the complete graph. In *Proc. 2nd Sympos. Comput. Geom.*, pages 169–177, ACM Press, 1986.
- [Che89] L.P. Chew. There are planar graphs almost as good as the complete graph. *J. Comput. System Sci.*, 39:205–219, 1989.
- [CK95] P.B. Callahan and S.R. Kosaraju. A decomposition of multidimensional point sets with applications to k -nearest-neighbors and n -body potential fields. *J. ACM*, 42:67–90, 1995.
- [Cla87] K.L. Clarkson. Approximation algorithms for shortest path motion planning (extended abstract). In *Proc. 19th ACM Sympos. Theory Comput.*, pages 56–65, 1987.
- [CLNS15] T.H. Chan, M. Li, L. Ning, and S. Solomon. New doubling spanners: Better and simpler. *SIAM J. Comput.*, 44:37–53, 2015.
- [CT76] D. Cheriton and R.E. Tarjan. Finding minimum spanning trees. *SIAM J. Comput.*, 5:724–742, 1976.
- [CT16] T.M. Chan and K. Tsakalidis. Optimal deterministic algorithms for 2-d and 3-d shallow cuttings. *Discrete Comput. Geom.*, 56:866–881, 2016.
- [CTL92] M.-S. Chang, C.Y. Tang, and R.C.T. Lee. Solving the Euclidean bottleneck matching problem by k -relative neighborhood graphs. *Algorithmica*, 8:177–194, 1992.
- [CW98] F.Y.L. Chin and C.A. Wang. Finding the constrained Delaunay triangulation and constrained Voronoi diagram of a simple polygon in linear time. *SIAM J. Comput.*, 28:471–486, 1998.
- [CX01] S.-W Cheng and Y.-F. Xu. On β -skeleton as a subgraph of the minimum weight triangulation. *Theoret. Comput. Sci.*, 262:459–471, 2001.
- [CZ04] A. Czumaj and H. Zhao. Fault-tolerant geometric spanners. *Discrete Comput. Geom.*, 32:207–230, 2004.
- [DBLL94] G. Di Battista, W. Lenhart, and G. Liotta. Proximity drawability: a survey. In *Proc. DIMACS Int. Workshop on Graph Drawing*, vol. 894 of *LNCS*, pages 328–339, Springer, Berlin, 1994.
- [DEG08] O. Devillers, J. Erickson, and X. Goaoc. Empty-ellipse graphs. In *Proc. 19th ACM-SIAM Sympos. Discrete Algorithms*, pages 1249–1257, 2008.
- [DFS90] D.P. Dobkin, S.J. Friedman, and K.J. Supowit. Delaunay graphs are almost as good as complete graphs. *Discrete Comput. Geom.*, 5:399–407, 1990.
- [DG16] A. Dumitrescu and A. Ghosh. Lower bounds on the dilation of plane spanners. *Internat. J. Comput. Geom. Appl.*, 26:89–110, 2016.
- [DH96] G. Das and P.J. Heffernan. Constructing degree-3 spanners with other sparseness properties. *Int. J. Found. Comput. Sci.*, 7:121–136, 1996.
- [DJ89] G. Das and D. Joseph. Which triangulations approximate the complete graph? In *Proc. Int. Sympos. Optimal Algorithms*, pages 168–192, 1989.

- [DN97] G. Das and G. Narasimhan. A fast algorithm for constructing sparse Euclidean spanners. *Internat. J. Comput. Geom. Appl.*, 7:297–315, 1997.
- [Dwy95] R.A. Dwyer. The expected size of the sphere-of-influence graph. *Comput. Geom.*, 5:155–164, 1995.
- [EM09] N.M. El Molla. Yao spanners for wireless ad hoc networks. Master’s thesis, Villanova University, 2009.
- [Epp95] D. Eppstein. Dynamic Euclidean minimum spanning trees and extrema of binary functions. *Discrete Comput. Geom.*, 13:111–122, 1995.
- [Epp00] D. Eppstein. Spanning trees and spanners. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 425–461, Elsevier, Amsterdam, 2000.
- [FK12] M. Fürer and S.P. Kasiviswanathan. Spanners for geometric intersection graphs with applications. *J. Comput. Geom.*, 3:31–64, 2012.
- [FS16] A. Filtser and S. Solomon. The greedy spanner is existentially optimal. In *Proc. ACM Sympos. Principles Dist. Comput.*, pages 9–17, 2016.
- [GGN06] J. Gao, L.J. Guibas, and A.T. Nguyen. Deformable spanners and applications. *Comput. Geom.*, 35:2–19, 2006.
- [GK07] J. Gudmundsson and C. Knauer. Dilation and detours in geometric networks. In T. Gonzalez, editor, *Handbook of Approximation Algorithms and Metaheuristics*, chap. 52, CRC Press, Boca Raton, 2007.
- [GKL03] A. Gupta, R. Krauthgamer, and J.R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *Proc. 44th IEEE Sympos. Found. Comput. Sci.*, pages 534–543, 2003.
- [GLN02] J. Gudmundsson, C. Levcopoulos, and G. Narasimhan. Fast greedy algorithms for constructing sparse geometric spanners. *SIAM J. Comput.*, 31:1479–1500, 2002.
- [GLNS08] J. Gudmundsson, C. Levcopoulos, G. Narasimhan, and M. Smid. Approximate distance oracles for geometric graphs. *ACM Trans. Alg.*, 4:10, 2008.
- [Got15] L.-A. Gottlieb. A light metric spanner. In *Proc. 56th IEEE Sympos. Found. Comput. Sci.*, pages 759–772, 2015.
- [GPS94] L.J. Guibas, J. Pach, and M. Sharir. Sphere-of-influence graphs in higher dimensions. In K. Böröcký and G. Fejes Tóth, editors, *Intuitive Geometry*, vol. 63 of *Coll. Math. Soc. J. Bolyai*, pages 131–137, Budapest, 1994.
- [GR08a] L.-A. Gottlieb and L. Roditty. Improved algorithms for fully dynamic geometric spanners and geometric routing. In *Proc. 19th ACM-SIAM Sympos. Discrete Algorithms*, pages 591–600, 2008.
- [GR08b] L.-A. Gottlieb and L. Roditty. An optimal dynamic spanner for doubling metric spaces. In *Proc. 16th Eur. Sympos. Algorithms*, vol. 5193 of *LNCS*, pages 478–489, Springer, Berlin, 2008.
- [GS69] K.R. Gabriel and R.R. Sokal. A new statistical approach to geographic variation analysis. *Systematic Biology*, 18:259–278, 1969.
- [GZ05] J. Gao and L. Zhang. Well-separated pair decomposition for the unit-disk graph metric and its applications. *SIAM J. Comput.*, 35:151–169, 2005.
- [HdLT01] J. Holm, K. de Lichtenberg, and M. Thorup. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. *J. ACM*, 48:723–760, 2001.
- [HLM03] F. Hurtado, G. Liotta, and H. Meijer. Optimal and suboptimal robust algorithms for proximity graphs. *Comput. Geom.*, 25:35–49, 2003.

- [HP11] S. Har-Peled. *Geometric Approximation Algorithms*, vol. 173 of *Math. Surveys Monogr.*, AMS, Providence, 2011.
- [HPM06] S. Har-Peled and M. Mendel. Fast construction of nets in low-dimensional metrics and their applications. *SIAM J. Comput.*, 35:1148–1184, 2006.
- [JK87] J.W. Jaromczyk and M. Kowaluk. A note on relative neighborhood graphs. In *Proc. 3rd Sympos. Comput. Geom.*, pages 233–241, ACM Press, 1987.
- [JT92] J.W. Jaromczyk and G.T. Toussaint. Relative neighborhood graphs and their relatives. *Proc of IEEE*, 80:1502–1517, 1992.
- [Kat88] J. Katajainen. The region approach for computing relative neighbourhood graphs in the L_p metric. *Computing*, 40:147–161, 1988.
- [Kei92] J.M. Keil. Approximating the complete Euclidean graph. *Discrete Comput. Geom.*, 7:13–28, 1992.
- [Kei94] J.M. Keil. Computing a subgraph of the minimum weight triangulation. *Comput. Geom.*, 4:18–26, 1994.
- [KG92] J.M. Keil and C.A. Gutwin. Classes of graphs which approximate the complete Euclidean graph. *Discrete Comput. Geom.*, 7:13–28, 1992.
- [KG01] M.I. Karavelas and L.J. Guibas. Static and kinetic geometric spanners with applications. In *Proc. 12th ACM-SIAM Sympos. Discrete Algorithms*, pages 168–176, 2001.
- [KL96] R. Klein and A. Lingas. A linear-time randomized algorithm for the bounded Voronoi diagram of a simple polygon. *Internat. J. Comput. Geom. Appl.*, 6:263–278, 1996.
- [KL98] D. Krznanic and C. Levcopoulos. Computing a threaded quadtree from the Delaunay triangulation in linear time. *Nord. J. Comput.*, 5:1–18, 1998.
- [KMR⁺17] H. Kaplan, W. Mulzer, L. Roditty, P. Seiferth, and M. Sharir. Dynamic planar Voronoi diagrams for general distance functions and their algorithmic applications. In *Proc. 28th ACM-SIAM Sympos. Discrete Algorithms*, pages 2495–2504, 2017.
- [KMRS15] H. Kaplan, W. Mulzer, L. Roditty, and P. Seiferth. Spanners and reachability oracles for directed transmission graphs. In *Proc. 31st Sympos. Comput. Geom.*, vol. 34 of *LIPICs*, pages 156–170, Schloss Dagstuhl, 2015.
- [KN87] J. Katajainen and O. Nevalainen. An almost naive algorithm for finding relative neighbourhood graphs in L_p metrics. *ITA*, 21:199–215, 1987.
- [KP08] I.A. Kanj and L. Perkovic. On geometric spanners of Euclidean and unit disk graphs. In *Proc. 25th Sympos. Theoret. Aspects Comput. Sci.*, vol. 1 of *LIPICs* pages 409–420, Schloss Dagstuhl, 2008.
- [KPT16] I.A. Kanj, L. Perkovic, and D. Türkoglu. Degree four plane spanners: Simpler and better. In *Proc. 32nd Int. Sympos. Comput. Geom.*, vol. 51 of *LIPICs*, article 45, Schloss Dagstuhl, 2016.
- [KPX10] I.A. Kanj, L. Perkovic, and G. Xia. On spanners and lightweight spanners of geometric graphs. *SIAM J. Comput.*, 39:2132–2161, 2010.
- [KR85] D.G. Kirkpatrick and J.D. Radke. A framework for computational morphology. In G.T. Toussaint, editor, *Computational Geometry*, vol. 2 of *Machine Intelligence and Pattern Recognition*, pages 217–248. North-Holland, Amsterdam, 1985.
- [KSVC15] T. Kaiser, M. Saumell, and N. van Cleemput. 10-Gabriel graphs are Hamiltonian. *Inform. Process. Lett.*, 115:877–881, 2015.

- [Lee85] D.T. Lee. Relative neighborhood graphs in the L_1 -metric. *Pattern Recogn.*, 18:327–332, 1985.
- [Lin94] A. Lingas. A linear-time construction of the relative neighborhood graph from the Delaunay triangulation. *Comput. Geom.*, 4:199–208, 1994.
- [Lio13] G. Liotta. Proximity drawings. In R. Tamassia, editor, *Handbook on Graph Drawing and Visualization*, pages 115–154, CRC Press, Boca Raton, 2013.
- [LL92] C. Levkopoulos and A. Lingas. There are planar graphs almost as good as the complete graphs and almost as cheap as minimum spanning trees. *Algorithmica*, 8:251–256, 1992.
- [LM12] M. Löffler and W. Mulzer. Triangulating the square and squaring the triangle: Quadrees and Delaunay triangulations are equivalent. *SIAM J. Comput.*, 41:941–974, 2012.
- [LNS02] C. Levkopoulos, G. Narasimhan, and M. Smid. Improved algorithms for constructing fault-tolerant spanners. *Algorithmica*, 32:144–156, 2002.
- [Luk99] T. Lukovszki. New results of fault tolerant geometric spanners. In *Proc. 6th Int. Workshop Algorithms Data Structures*, vol. 1663 of *LNCS*, pages 193–204, Springer, Berlin, 1999.
- [MR08] W. Mulzer and G. Rote. Minimum-weight triangulation is NP-hard. *J. ACM*, 55:11, 2008.
- [MS80] D.W. Matula and R.R. Sokal. Properties of Gabriel graphs relevant to geographic variation research and the clustering of points in the plane. *Geogr. Anal.*, 12:205–222, 1980.
- [Mul04] W. Mulzer. Minimum dilation triangulations for the regular n -gon. Master’s thesis, Freie Universität Berlin, 2004.
- [NS02] G. Narasimhan and M. Smid. Approximation algorithms for the bottleneck stretch factor problem. *Nord. J. Comput.*, 9:13–31, 2002.
- [NS07] G. Narasimhan and M. Smid. *Geometric Spanner Networks*. Cambridge University Press, 2007.
- [O’R82] J. O’Rourke. Computing the relative neighborhood graph in the L_1 and L_∞ metrics. *Pattern Recogn.*, 15:189–192, 1982.
- [PS85] F.P. Preparata and M.I. Shamos. *Computational Geometry. An Introduction*. Springer-Verlag, New York, 1985.
- [Rod12] L. Roditty. Fully dynamic geometric spanners. *Algorithmica*, 62:1073–1087, 2012.
- [RS91] J. Ruppert and R. Seidel. Approximating the d -dimensional Euclidean graph. In *Proc. 3rd Canad. Conf. Comput. Geom.*, pages 207–210, 1991.
- [SA95] M. Sharir and P.K. Agarwal. *Davenport-Schönzel Sequences and Their Geometric Applications*. Cambridge University Press, 1995.
- [Sam90] H. Samet. *The Design and Analysis of Spatial Data Structures*. Addison-Wesley, Reading, 1990.
- [SC90] T.-H. Su and R.-C. Chang. The k -Gabriel graphs and their applications. In *SIGAL Int. Sympos. Algorithms*, pages 66–75. Springer-Verlag, Berlin, 1990.
- [Sch79] A. Schönhage. On the power of random access machines. In *Proc. 6th Int. Colloq. Automata Lang. Program.*, vol. 71 of *LNCS*, pages 520–529, Springer, Berlin, 1979.
- [Sei16] P. Seiferth. *Disk Intersection Graphs: Models, Data Structures, and Algorithms*. PhD thesis, Freie Universität Berlin, 2016.

- [Smi89] W.D. Smith. *Studies in Computational Geometry Motivated by Mesh Generation*. PhD thesis, Princeton University, 1989.
- [Smi07] M. Smid. The well-separated pair decomposition and its applications. In T. Gonzalez, editor, *Handbook of Approximation Algorithms and Metaheuristics*, chap. 53, CRC Press, Boca Raton, 2007.
- [Smi09] M. Smid. The weak gap property in metric spaces of bounded doubling dimension. In *Efficient Algorithms, Essays Dedicated to Kurt Mehlhorn on the Occasion of His 60th Birthday*, pages 275–289, Springer, Berlin, 2009.
- [Sos99] M.A. Soss. On the size of the Euclidean sphere of influence graph. In *Proc. 11th Canad. Conf. Comput. Geom.*, pages 43–46, 1999.
- [Sup83] K.J. Supowit. The relative neighborhood graph, with an application to minimum spanning trees. *J. ACM*, 30:428–448, 1983.
- [Tal04] K. Talwar. Bypassing the embedding: algorithms for low dimensional metrics. In *Proc. 36th ACM Sympos. Theory Comput.*, pages 281–290, 2004.
- [Tou80] G.T. Toussaint. The relative neighbourhood graph of a finite planar set. *Pattern Recogn.*, 12:261–268, 1980.
- [Tou88] G.T. Toussaint. A graph-theoretical primal sketch. In G.T. Toussaint, editor, *Computational Morphology*, pages 229–260, North-Holland, Amsterdam, 1988.
- [Tou14] G.T. Toussaint. The sphere of influence graph: Theory and applications. *Information Technology & Computer Science*, 14:2091–1610, 2014.
- [Urq83] R.B. Urquhart. Some properties of the planar Euclidean relative neighbourhood graph. *Pattern Recogn. Lett.*, 1:317–322, 1983.
- [Var98] K.R. Varadarajan. A divide-and-conquer algorithm for min-cost perfect matching in the plane. In *Proc. 39th IEEE Sympos. Found. Comput. Sci.*, pages 320–331, 1998.
- [Vel91] R.C. Veltkamp. The γ -neighborhood graph. *Comput. Geom.*, 1:227–246, 1991.
- [Wil16] M. Willert. Routing schemes for disk graphs and polygons. Master’s thesis, Freie Universität Berlin, 2016.
- [WY01] C.A. Wang and B. Yang. A lower bound for β -skeleton belonging to minimum weight triangulations. *Comput. Geom.*, 19:35–46, 2001.
- [Xia13] G. Xia. The stretch factor of the Delaunay triangulation is less than 1.998. *SIAM J. Comput.*, 42:1620–1659, 2013.
- [XZ11] G. Xia and L. Zhang. Toward the tight bound of the stretch factor of Delaunay triangulations. In *Proc. 23rd Canad. Conf. Comput. Geom.*, 2011.
- [Yao82] A.C.-C. Yao. On constructing minimum spanning trees in k -dimensional spaces and related problems. *SIAM J. Comput.*, 11:721–736, 1982.