

Empty Squares in Arbitrary Orientation Among Points

Sang Won Bae 

Division of Computer Science and Engineering, Kyonggi University, Suwon, South Korea
swbae@kgu.ac.kr

Sang Duk Yoon 

Department of Service and Design Engineering, Sungshin Women's University, Seoul, South Korea
sangduk.yoon@sungshin.ac.kr

Abstract

This paper studies empty squares in arbitrary orientation among a set P of n points in the plane. We prove that the number of empty squares with four contact pairs is between $\Omega(n)$ and $O(n^2)$, and that these bounds are tight, provided P is in a certain general position. A contact pair of a square is a pair of a point $p \in P$ and a side ℓ of the square with $p \in \ell$. The upper bound $O(n^2)$ also applies to the number of empty squares with four contact points, while we construct a point set among which there is no square of four contact points. We then present an algorithm that maintains a combinatorial structure of the L_∞ Voronoi diagram of P , while the axes of the plane continuously rotate by 90 degrees, and simultaneously reports all empty squares with four contact pairs among P in an output-sensitive way within $O(s \log n)$ time and $O(n)$ space, where s denotes the number of reported squares. Several new algorithmic results are also obtained: a largest empty square among P and a square annulus of minimum width or minimum area that encloses P over all orientations can be computed in worst-case $O(n^2 \log n)$ time.

2012 ACM Subject Classification Theory of computation \rightarrow Computational geometry

Keywords and phrases empty square, arbitrary orientation, Erdős–Szekeres problem, L_∞ Voronoi diagram, largest empty square problem, square annulus

Digital Object Identifier 10.4230/LIPIcs.SoCG.2020.13

Related Version A full version [8] of the paper is available at <https://arxiv.org/abs/1911.12988>.

Funding Sang Won Bae: Supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (2018R1D1A1B07042755).

Sang Duk Yoon: Supported by “Cooperative Research Program for Agriculture Science & Technology Development (Project No. PJ01526903)” Rural Development Administration, Republic of Korea.

1 Introduction

We start by posing the following combinatorial question:

Given a set P of n points in a proper general position in \mathbb{R}^2 , how many empty squares in arbitrary orientation whose boundary contains four points in P can there be?

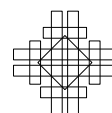
For a square, its contact pair is a pair of a point $p \in P$ and a side ℓ of it such that $p \in \ell$, regarding ℓ as a segment including its endpoints. An analogous question asks the number of empty squares with four contact pairs. These questions can be seen as a new variant of the Erdős–Szekeres problem [22, 23] for empty squares. Let $s = s(P)$ and $s^* = s^*(P)$ be the number of empty squares with four contact points and with four contact pairs, respectively. The difference between contact points and contact pairs is that the first counts the number of points, while the second counts the number of incidences; in particular contact points that are corners count as two contact pairs. In this paper, we prove that $0 \leq s < c_1 n^2$ and $c_2 n < s^* < c_3 n^2$ for some constants $c_1, c_2, c_3 > 0$. These lower and upper bounds are tight



© Sang Won Bae and Sang Duk Yoon;
licensed under Creative Commons License CC-BY
36th International Symposium on Computational Geometry (SoCG 2020).
Editors: Sergio Cabello and Danny Z. Chen; Article No. 13; pp. 13:1–13:17



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



by the existence of point sets with the asymptotically same number of such squares. These questions and results are shown to be intrinsic to several computational problems on empty squares, implying new algorithmic results.

For the purpose, we provide a solid understanding of empty squares in arbitrary orientation among P by establishing a geometric and topological relation among those squares. The family of axis-parallel empty squares is well understood by the L_∞ Voronoi diagram of P . We extend this knowledge to those in arbitrary orientation by investigating the L_∞ Voronoi diagrams of P with the axes rotated. Our proof for the above combinatorial results is based on new observations made upon the consideration of the Voronoi diagram in this way. This also motivates the problem of maintaining the L_∞ diagram of P while the axes continuously rotate. A noteworthy observation states that every combinatorial change of the diagram during the rotation of the axes corresponds to an empty square with four contact pairs. Hence, the total amount of changes of the diagram is bounded by $\Theta(s^*) = O(n^2)$.

Based on the observations, we then present an output-sensitive algorithm that finds all empty squares with four contact pairs in $O(s^* \log n)$ time using $O(n)$ space. Our algorithm indeed maintains a combinatorial description of the L_∞ Voronoi diagram as the axes continuously rotates by 90 degrees by capturing every occurrence of such special squares and handling them so that the diagram is correctly maintained as the invariants. To our best knowledge, there was no such algorithmic result in the literature. Our algorithm also applies to two other geometric problems, achieving new algorithmic results:

- (1) A *largest empty square* in arbitrary orientation can be found in $O(n^2 \log n)$ time. This improves the previous $O(n^3)$ -time algorithm by Bae [6]. We also solve some query versions of this problem.
- (2) A *square annulus* is the closed region between a pair of concentric and parallel squares. A square annulus in arbitrary orientation with minimum width or area can be computed in $O(n^2 \log n)$ time, improving the previous $O(n^3 \log n)$ and $O(n^3)$ -time algorithms [5, 6].

Our combinatorial problem on the number of empty squares in arbitrary orientation can be viewed as a new variant of the *Erdős–Szekeres problem*, which has a long history since 1935 [22] with consistent effort [23, 41] on its original version and many other variants and generalizations [9–11, 18, 21, 24, 26, 28, 36, 38, 42]. Among them the most relevant to us is the problem of bounding the number of empty convex k -gons whose corners are chosen from P , called *k-holes*. The maximum number of k -holes among n points is proven to be $\Theta(n^k)$ for $k \geq 3$ and sufficiently large n by Bárány and Valtr [10]. Its minimum is known to be $\Theta(n^2)$ for $3 \leq k \leq 4$; $\Omega(n)$ and $O(n^2)$ for $5 \leq k \leq 6$ [9]; and zero for $k \geq 7$ [28]. For more details on this subject, see survey papers by Morris and Soltan [33] and [34].

Since any four points in P lying on the boundary of an empty square form a 4-hole, s cannot exceed the number of 4-holes among P . This, however, gives only trivial upper and lower bounds on s and s^* . In this paper, we give asymptotically tight bounds on s and s^* since we rather focus on algorithmic applications of the bounds.

Dobkin, Edelsbrunner, and Overmars [16] presented an algorithm that enumerates all convex k -holes for $3 \leq k \leq 6$. Their algorithm in particular for $k = 3, 4$ is indeed output-sensitive in time proportional to the number of reported k -holes, which is comparable to our algorithm that enumerates all squares with four contact pairs. Rote et al. [39], Rote and Woeginger [40], and Mitchell et al. [32] considered the problem of exactly counting convex k -gons and convex k -holes faster than enumerating them. Some minimization and maximization problems have also been considered in the literature [2, 13, 17, 19, 20].

The problem of maintaining the L_∞ (or, equivalently, L_1) Voronoi diagram while the axes rotate cannot be found in the literature. A similar paradigm about rotating axes can be seen in Bae et al. [7] in which the authors show how to maintain the *orthogonal convex*

hull of P in $O(n^2)$ time. Alegría-Galicia et al. [4] recently improved it into $O(n \log n)$ time using $O(n)$ space. As will be seen in the following, the orthogonal convex hull of P indeed describes the unbounded edges of the L_∞ Voronoi diagram of P .

The problem of finding a largest empty square is a square variant of the well-known *largest empty rectangle problem*. The largest empty rectangle problem is one of the most intensively studied problems in early time of computational geometry. After early results [15, 31, 35, 37], the currently fastest algorithm that finds a largest empty axis-parallel rectangle among P runs in $O(n \log^2 n)$ time by Aggarwal and Suri [3]. Mckenna et al. [31] proved a lower bound of $\Omega(n \log n)$ for this problem. Chauduri et al. [14] showed that there are $O(n^3)$ combinatorially different classes of maximal empty rectangles over all orientations and presented an $O(n^3)$ -time algorithm that computes a largest empty rectangle in arbitrary orientation by enumerating all those classes.

The largest empty square problem, however, has attained relatively less interest. It is obvious that a largest empty axis-parallel square can be found in $O(n \log n)$ time by computing the L_∞ Voronoi diagram of P [29, 30], as also remarked in early papers, including Naamad et al. [35] and Chazelle et al. [15]. It is rather surprising, however, that no further result about the largest empty square problem in arbitrary orientation has been come up with for over three decades, to our best knowledge. From an easy observation that any maximal empty square in arbitrary orientation is contained in a maximal empty rectangle, Bae [6] recently showed how to compute a largest empty square in arbitrary orientation in $O(n^3)$ time. In this paper, we improve this to $O(n^2 \log n)$ time.

The square annulus problem asks to find a square annulus of minimum width or area that encloses a given set P of points. Its fixed-orientation version is solved in $O(n \log n)$ time [1, 25]. Bae [5] presented the first $O(n^3 \log n)$ -time algorithm for the problem in arbitrary orientation, and improved it to $O(n^3)$ time [6]. In this paper, we further improve it to $O(n^2 \log n)$ time.

2 Preliminaries

We consider the standard coordinate system with the (horizontal) x -axis and the (vertical) y -axis in the plane \mathbb{R}^2 . We mean by the *orientation* of any line, half-line, or line segment ℓ a unique real number $\theta \in [0, \pi)$ such that ℓ is parallel to a counter-clockwise rotated copy of the x -axis by θ .

For any square S in \mathbb{R}^2 , the *orientation* of S is a real number $\theta \in [0, \pi/2)$ such that the orientation of each side of S is either θ or $\theta + \pi/2$. We regard the set $\mathbb{O} := [0, \pi/2)$ of all orientations of squares as a topological space homeomorphic to a circle.

Each side of a square, as a subset of \mathbb{R}^2 , is assumed to include its incident corners. We identify the four sides of a square S by the *top*, *bottom*, *left*, and *right* sides, denoted by $\top(S)$, $\perp(S)$, $\leftarrow(S)$, and $\rightarrow(S)$, respectively. This identification is clear, regardless of the orientation of S since it is chosen from $\mathbb{O} = [0, \pi/2)$. The *center* of a square is the intersection point of its two diagonals, and its *radius* is half its side length.

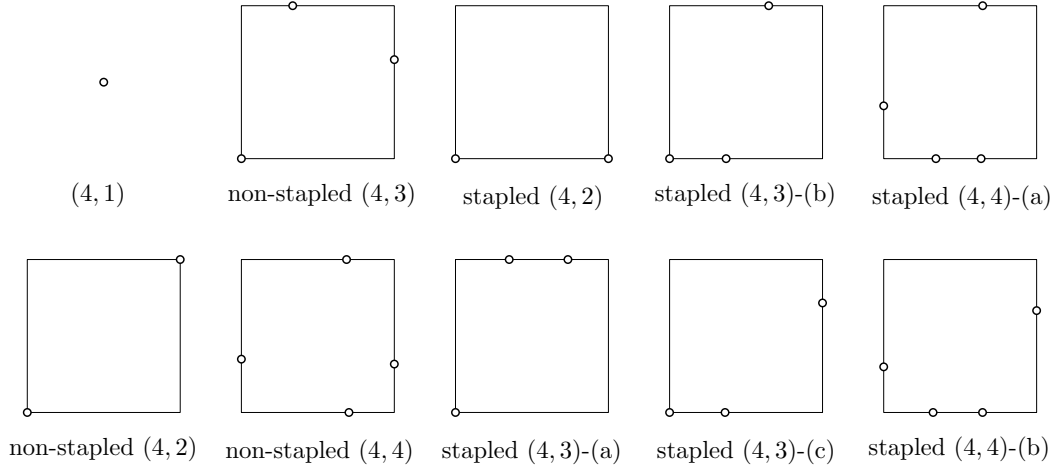
Let P be a set of n points in \mathbb{R}^2 . A square is called *empty* if no point in P lies in its interior. An empty square may contain some points in P on its boundary. A pair (p, \square) of a point $p \in P$ and a side identifier $\square \in \{\top, \perp, \leftarrow, \rightarrow\}$ is called a *contact pair* of S if $p \in \square(S)$. A set of contact pairs is called a *contact type*. If κ is the set of all contact pairs of S , then we say that κ is the *contact type* of S . A *contact point* $p \in P$ of S is a point on a side of S , that is, one belonging to a contact pair of S . Each contact point $p \in P$ of S may lie either on the relative interior of a side of S or at a corner of S . In the former case, the contact point p contributes to one contact pair of S , while in the latter case, it contributes to two contact pairs.

13:4 Empty Squares in Arbitrary Orientation Among Points

Throughout the paper, we assume that P is in *general position* in the following sense:

There is no square in arbitrary orientation with five or more contact pairs among P .

Note that the assumption implies the number of empty squares with four contact pairs is finite.



■ **Figure 1** Illustration of 10 types of 4-squares and the type names. Small circles on the boundary of each square depict its contact points in P . The four left ones are non-stapled types and the six right ones are stapled types.

Consider any empty square S of contact type κ . We call S an m -square or (m, k) -square if $m = |\kappa|$ and k is the number of contact points in κ . A side of S is called *pinned* if it contains a point in P , so it is involved in some contact pair in κ ; or *stapled* if it contains two distinct points in P . If S has a stapled side, then S is called *stapled*. From the general position assumption, there are no three or more contact pairs in κ involving a common side of S , and there is at most one stapled side of S .

We then classify the 4-squares into 10 types under the symmetry group of the square. See Figure 1 for an illustration to the 10 types of 4-squares with type names.

Throughout the paper, we are less interested in trivial $(4, 1)$ -squares in most cases. Hereafter, we thus mean by a 4-square a nontrivial 4-square, that is, a $(4, k)$ -square with $k \geq 2$, unless stated otherwise.

3 Empty Squares and the Voronoi Diagram

The empty squares among P are closely related to the Voronoi diagram of P under the L_∞ distance. In this section, we define the Voronoi diagram for every $\theta \in \mathbb{O}$ based on empty squares and collect several essential properties of empty squares in terms of the Voronoi diagram, based on which we will be able to bound the number of 4-squares and to present an efficient algorithm that computes all 4-squares.

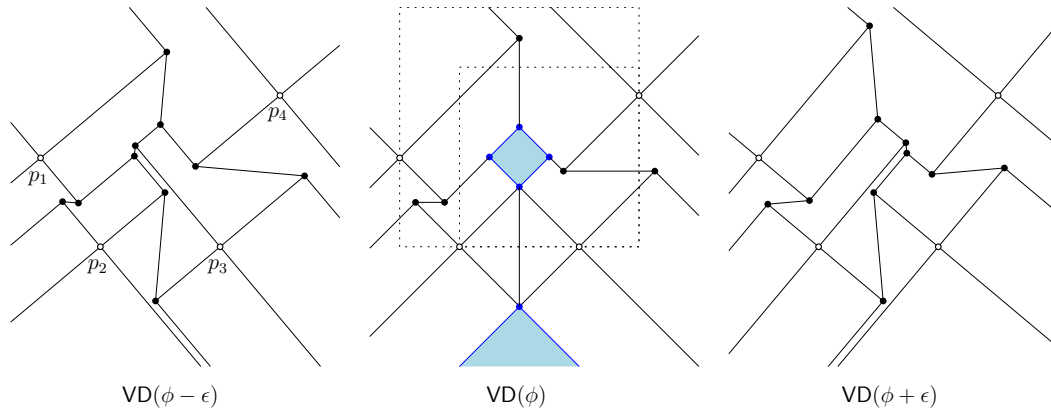
3.1 Definition of Voronoi diagrams

Let P be a given set of n points in general position as discussed in Section 2. For each $\theta \in \mathbb{O}$, we define $\text{VD}(\theta)$ to be the L_∞ Voronoi diagram of P with the axes rotated by θ , or equivalently, the Voronoi diagram of P under the symmetric convex distance function d_θ based on a unit square whose orientation is θ . The *Voronoi region* of $p \in P$ in θ is

$$VR_p(\theta) := \{x \in \mathbb{R}^2 \mid d_\theta(x, p) < d_\theta(x, q), q \in P \setminus \{p\}\}.$$

The diagram $\text{VD}(\theta)$ can also be defined in terms of empty squares. More precisely, we view $\text{VD}(\theta)$ as a plane graph whose vertices $\hat{V}(\theta)$ and edges $\hat{E}(\theta)$ are determined as follows:

- The vertex set $\hat{V}(\theta)$ consists of the centers of all empty squares in orientation θ with three or four pinned sides, and a point *at infinity*, denoted by $\hat{\infty}$.
- An edge is contained in $\hat{E}(\theta)$ if and only if it is a maximal set of centers of all empty squares in orientation θ having a common contact type with two pinned sides. Each edge in $\hat{E}(\theta)$ is either a half-line or a line segment, called *unbounded* or *bounded*, respectively.



■ **Figure 2** Illustration of $\text{VD}(\phi - \epsilon)$, $\text{VD}(\phi)$, and $\text{VD}(\phi + \epsilon)$ for $P = \{p_1, \dots, p_4\}$ and some $\phi \in \mathbb{O}$.

See Figure 2. From our definition, there are four more edges incident to each $p \in P$ such that each of them corresponds to 2-squares with one corner anchored at p . In this way, each point $p \in P$ is also a vertex in $\hat{V}(\theta)$ since p is the center of a trivial $(4, 1)$ -square with its four sides pinned. The diagram $\text{VD}(\theta)$ as a plane graph divides the plane into its *faces*. Each face of $\text{VD}(\theta)$ is the locus of centers of empty squares having a common contact type with *one pinned side*; hence, for every contact pair (p, \square) , there exists a unique face of $\text{VD}(\theta)$ consisting of the centers of 1-squares with contact type $\{(p, \square)\}$. Therefore, the Voronoi region $\text{VR}_p(\theta)$ of each $p \in P$ includes exactly four faces of $\text{VD}(\theta)$. On the other hand, there may exist some *neutral faces* of $\text{VD}(\theta)$ that do not belong to any Voronoi region $\text{VR}_p(\theta)$, if it corresponds to 2-squares with a stapled side (see the two shaded faces of $\text{VD}(\phi)$ in color lightblue in Figure 2). This is obviously a degenerate case which has been avoided from most discussions about Voronoi diagrams in the literature. In this way, our definition of $\text{VD}(\theta)$ completely represents all cases of point set P , even if there are four equidistant points in P under d_θ or there are two points in P such that pq is in orientation θ or $\theta + \pi/2$.

The *combinatorial structure* of $\text{VD}(\theta)$ is represented by its underlying graph $\text{VG}(\theta) = (V(\theta), E(\theta))$, called the *Voronoi graph*; conversely, $\text{VD}(\theta)$ is a plane embedding of $\text{VG}(\theta)$. More precisely, the vertices and edges of $\text{VG}(\theta)$ are described and identified as follows:

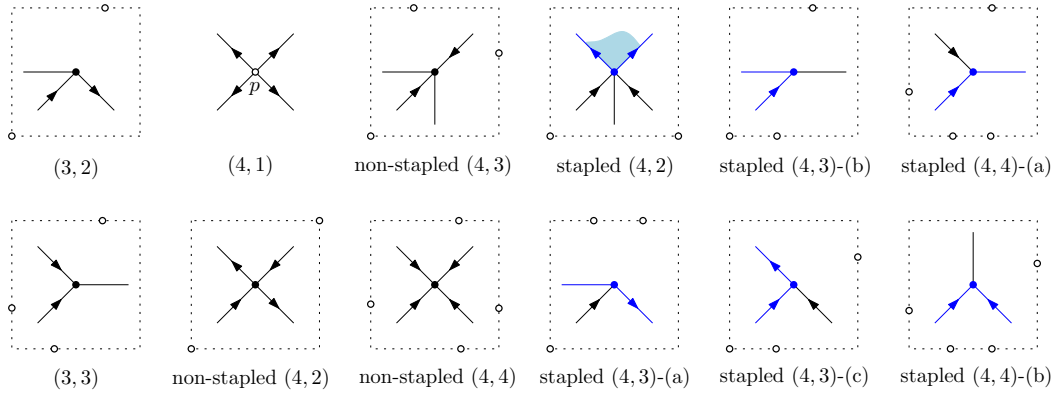
- Each vertex $v \in V(\theta)$ corresponds to $\hat{v} \in \hat{V}(\theta)$. In particular, the vertex at infinity, denoted by $\infty \in V(\theta)$, corresponds to $\hat{\infty} \in \hat{V}(\theta)$. Each $v \in V(\theta) \setminus \{\infty\}$ is identified by the contact type κ_v of the square defining $\hat{v} \in \hat{V}(\theta)$. For completeness, we define $\kappa_\infty := \emptyset$.
- Each edge $e \in E(\theta)$ corresponds to $\hat{e} \in \hat{E}(\theta)$. Each edge $e = uv \in E(\theta)$ for $u, v \in V(\theta)$ is identified by a triple $(\kappa_u, \kappa_v; \kappa_e)$, where κ_e is the contact type of the squares defining $\hat{e} \in \hat{E}(\theta)$. If e is bounded, then we have $\kappa_e = \kappa_u \cap \kappa_v$.

Hence, for $\theta, \theta' \in \mathbb{O}$, two vertices $v \in V(\theta)$ and $v' \in V(\theta')$ are the same if $\kappa_v = \kappa_{v'}$; two edges $uv \in E(\theta)$ and $u'v' \in E(\theta')$ are the same if $(\kappa_u, \kappa_v; \kappa_{uv}) = (\kappa_{u'}, \kappa_{v'}; \kappa_{u'v'})$. We say that $\text{VD}(\theta)$ and $\text{VD}(\theta')$ are *combinatorially equivalent* if $\text{VG}(\theta) = \text{VG}(\theta')$.

3.2 Basic properties

For each vertex $v \in V(\theta)$, we call v and its embedding $\hat{v} \in \hat{V}(\theta)$ *regular* if $|\kappa_v| = 3$, that is, its corresponding empty square is a 3-square. For each edge $e \in E(\theta)$, we call e and its embedding $\hat{e} \in \hat{E}(\theta)$ *regular* if $|\kappa_e| = 2$.

Each edge $e \in E(\theta)$ is called *sliding* if the two pinned sides in κ_e are parallel, or *growing*, otherwise. From the properties of the L_∞ Voronoi diagram, we observe that any sliding edge is in orientation θ or $\theta + \pi/2$, while any growing edge is in orientation $\theta + \pi/4$ or $\theta + 3\pi/4$ (modulo π). For $e \in E(\theta)$, we regard each growing edge to be directed in which its corresponding square is growing.



■ **Figure 3** Illustration of the 12 vertex types of $VD(\theta)$ labeled with their type names. Dotted squares and small circles on them depict an empty square corresponding to each vertex type and its contact points. Dots and line segments depict vertices and all incident edges to each vertex; in black if regular, or in blue, otherwise. The arrows on edges depict the direction in which the corresponding square grows. The shaded area in the stapled (4, 2)-type depicts a neutral face which does not belong to any Voronoi region $VR_p(\theta)$ for $p \in P$. (We keep the above convention on every figure in this paper.) Note that the right six types are stapled, while the left six are not.

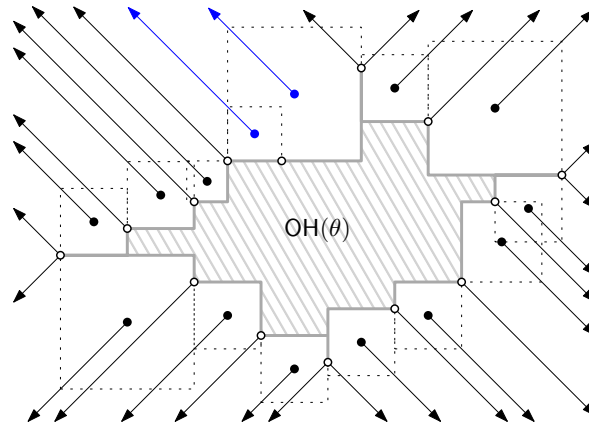
For each vertex $v \in V(\theta)$ with $v \neq \infty$, the local structure of the diagram $VD(\theta)$ around \hat{v} is completely determined by its contact type κ_v . From the possible contact types of 3- and 4-squares, we classify all vertices in $V(\theta) \setminus \{\infty\}$ into 12 vertex types.

► **Lemma 3.1.** *There are 12 types for the vertices of $VD(\theta)$ as illustrated in Figure 3. For any vertex $v \in V(\theta) \setminus \{\infty\}$, the contact types κ_e of all edges $e \in E(\theta)$ incident to v can be obtained just from the contact type κ_v without knowing the other incident vertex of e .*

Consider all squares with contact type κ_e defining $\hat{e} \in \hat{E}(\theta)$. If e is sliding, then all these squares have the same radius; if e is growing, then their radius grows along \hat{e} . The following lemma is an immediate observation.

► **Lemma 3.2** (Boissonnat et al. [12]). *Let $e = uv \in E(\theta)$ be any bounded edge for any $\theta \in \mathbb{O}$, and S_u and S_v be the empty squares in θ with contact types κ_u and κ_v , respectively. Then, the union R of all squares in θ with contact type κ_e is equal to $S_u \cup S_v$. More specifically, if e is sliding, then $R = S_u \cup S_v$ forms a rectangle; if e is growing, then one of S_u and S_v completely contains the other, so R is a square.*

The above lemma indeed extends to the case of unbounded edges. Consider any unbounded edge $e \in E(\theta)$ and the union R of all squares as declared in Lemma 3.2. Observe that R forms an empty unbounded quadrant rotated by θ and the empty quadrant R has two or three



■ **Figure 4** Illustration to Lemma 3.3. The orthogonal convex hull $\text{OH}(\theta)$ of P (shaded region) and the four staircases (gray thick lines) describe all unbounded edges (arrows) and their incident vertices of $\text{VD}(\theta)$ (small circles and dots).

contact points in P . This tells us a relation between unbounded edges and the orthogonal convex hull of P . The *orthogonal convex hull* of point set P is defined to be the minimal subset of \mathbb{R}^2 such that any vertical or horizontal line intersects it in at most one connected component. It is known that the orthogonal convex hull is obtained by subtracting all empty quadrants (of four directions) from the whole plane \mathbb{R}^2 and its boundary is represented by four monotone chains, called the *staircases*. For $\theta \in \mathbb{O}$, let $\text{OH}(\theta)$ denote the orthogonal convex hull of P with the axes rotated by θ . See Figure 4 and Bae et al. [7] for more details on $\text{OH}(\theta)$, including the precise definition of $\text{OH}(\theta)$ and the staircases.

► **Lemma 3.3.** *For any $\theta \in \mathbb{O}$, all the unbounded edges and their incident vertices of $\text{VD}(\theta)$ are explicitly described by $\text{OH}(\theta)$, in the sense that if $v \in V(\theta) \setminus \{\infty\}$ is a vertex incident to an unbounded edge, then either*

- (i) *we have $\hat{v} = p \in P$ and p coincides with a vertex of $\text{OH}(\theta)$, or*
- (ii) *its contact points in κ_v appear consecutively in a staircase of $\text{OH}(\theta)$.*

Another implication of Lemma 3.1 is that the degree of every vertex, except $\infty \in V(\theta)$, is at least three and at most five. This implies the linear complexity of $\text{VD}(\theta)$ for any $\theta \in \mathbb{O}$.


► **Lemma 3.4.** *For any $\theta \in \mathbb{O}$, the number of vertices, edges, and faces of $\text{VD}(\theta)$ is $\Theta(n)$.*

3.3 Combinatorial changes of $\text{VD}(\theta)$ and 4-squares

Let \mathfrak{S}_4 be the set of all nontrivial 4-squares among P . By our general position assumption on P , we know that \mathfrak{S}_4 is finite. Let $s_4 := |\mathfrak{S}_4|$ be the number of 4-squares among P . For any orientation $\theta \in \mathbb{O}$, we call θ *regular* if there is no nontrivial 4-square in orientation θ , or *degenerate*, otherwise. Since there are only finitely many (exactly s_4) 4-squares, all orientations $\theta \in \mathbb{O}$ but at most s_4 of them are regular.

In a regular orientation θ , the diagram $\text{VD}(\theta)$ has the following properties.

► **Lemma 3.5.** *For any regular orientation $\theta \in \mathbb{O}$, every vertex in $V(\theta)$, except those in P and ∞ , is regular, and every edge in $E(\theta)$ is regular. There are five types of bounded edges, as shown in Figure 5, and two types of unbounded edges.*



■ **Figure 5** Illustration of five types of bounded edges of $VD(\theta)$ for a regular orientation $\theta \in \mathbb{O}$, depending on the types of incident vertices. The first three are sliding and the last two are growing.

Consider any contact type κ with three contact pairs and three pinned sides. For any $\theta \in \mathbb{O}$, let $S_\kappa(\theta)$ be the square in orientation θ whose contact type includes the three contact pairs in κ , regardless of its emptiness. Note that $S_\kappa(\theta)$ is well defined only in a closed interval of \mathbb{O} , or is never defined for any $\theta \in \mathbb{O}$. For each θ for which $S_\kappa(\theta)$ is well defined, we call θ *valid for κ* if $S_\kappa(\theta)$ is empty and its contact type is exactly κ . Note that if θ is valid for κ , then there exists a vertex $v \in V(\theta)$ with $\kappa_v = \kappa$.

► **Lemma 3.6.** *For any contact type κ with three contact pairs and three pinned sides, the set of all valid orientations for κ forms zero or more open intervals $I \subset \mathbb{O}$ such that $S_\kappa(\phi)$ is a 4-square for any endpoint ϕ of I .*

We call each of these open intervals described in Lemma 3.6 a *valid interval* for κ . Lemma 3.6 also implies that any endpoint of a valid interval for κ is a degenerate orientation.

For any degenerate orientation $\phi \in \mathbb{O}$, let $\mathfrak{S}_4(\phi) \subseteq \mathfrak{S}_4$ be the set of 4-squares whose orientation is ϕ . Note that $\mathfrak{S}_4(\phi)$ is nonempty and consists of at most $O(n)$ squares by Lemma 3.4. We then observe the following.

► **Lemma 3.7.** *Let $S \in \mathfrak{S}_4(\phi)$ be any 4-square in orientation ϕ and κ be its contact type. Then, there are exactly two or four contact types κ' with three contact pairs and three pinned sides such that $S_{\kappa'}(\phi) = S$ and either $\phi - \epsilon$ or $\phi + \epsilon$ is valid for κ' for sufficiently small $\epsilon > 0$. Specifically, the number of such κ' is four if S is non-stapled, or two if S is stapled.*

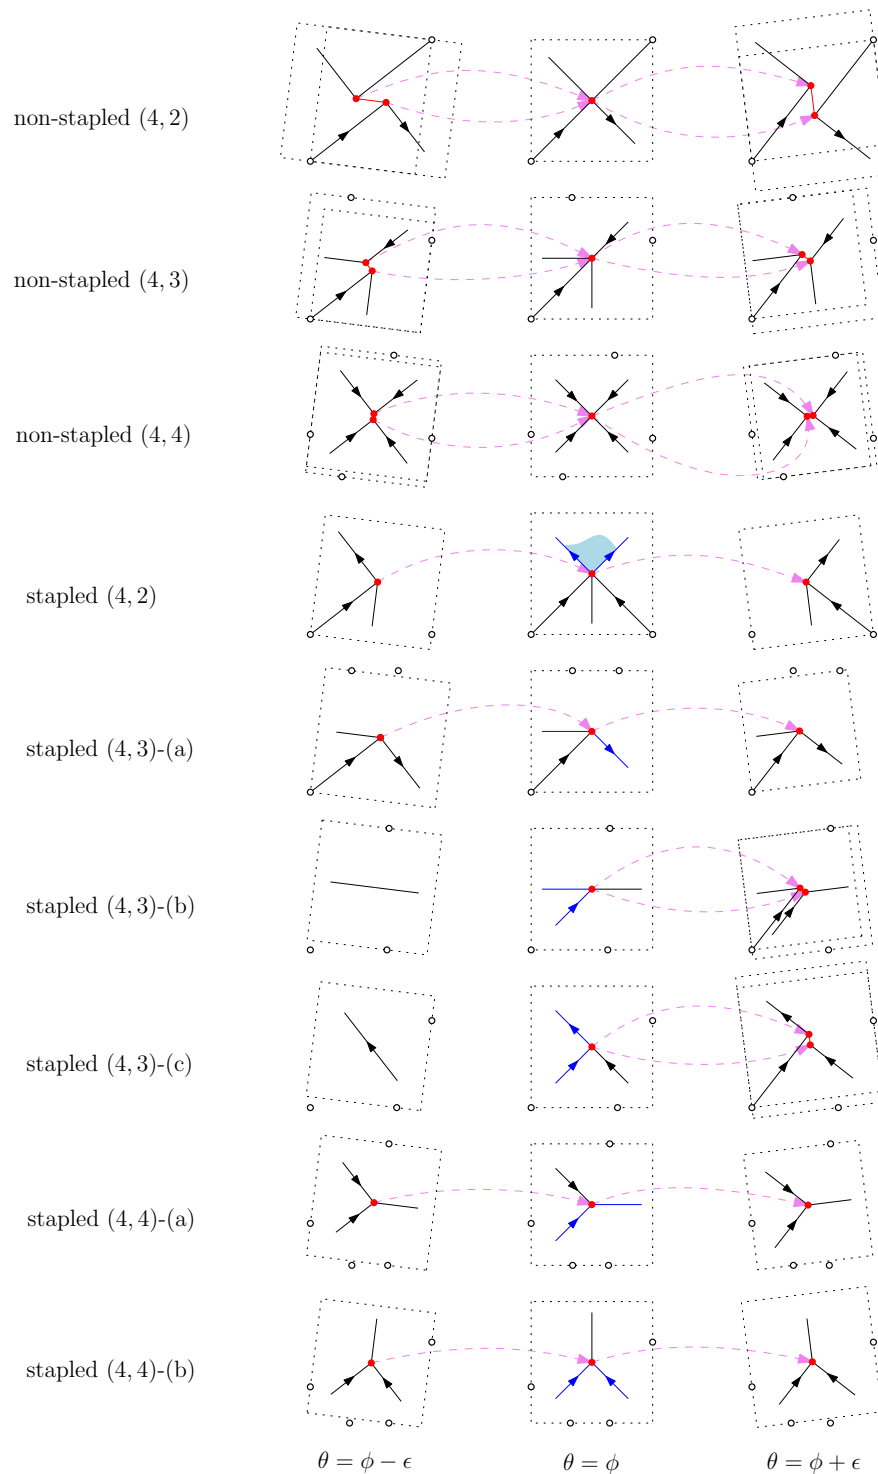
Figure 6 illustrates transitions of a non-regular vertex, which corresponds to a 4-square in $\mathfrak{S}_4(\phi)$, from and to regular vertices, locally at $\theta = \phi$, so almost describes our proof for Lemma 3.7. Observe that any non-stapled 4-square is relevant to an *edge flip* of $VD(\theta)$, and that only stapled (4,3)-(b)- or (c)-type squares show a bit different behavior; such a non-regular vertex suddenly appears on a regular edge and splits to two regular vertices, or, reversely, two regular vertices are merged into such a non-regular one and soon disappear.

The above discussions provide us a thorough view on the 3-squares and the 4-squares. Consider the graph \mathfrak{G} whose vertex set is the set \mathfrak{S}_4 of all 4-squares and whose edge set consists of edges (S, S') for $S, S' \in \mathfrak{S}_4$ such that there is a valid interval $I = (\phi, \phi')$ for some contact type κ such that $S_\kappa(\phi) = S$ and $S_\kappa(\phi') = S'$. The graph \mathfrak{G} is well defined by Lemma 3.6 and the degree of every vertex in \mathfrak{G} is exactly two or four by Lemma 3.7.

4 Number of 4-Squares

In this section, we prove asymptotically tight upper and lower bounds on the number of 4-squares and $(4, k)$ -squares for each $2 \leq k \leq 4$. Following are two main theorems.

► **Theorem 4.1.** *The number of 4-squares among n points in general position is always between $\Omega(n)$ and $O(n^2)$. These lower and upper bounds are asymptotically tight.*



■ **Figure 6** Transitions between regular (3-squares) and non-regular vertices (4-squares).

13:10 Empty Squares in Arbitrary Orientation Among Points

► **Theorem 4.2.** *Among n points in general position, the number of empty squares whose boundary contains four points of P is between 0 and $O(n^2)$. These lower and upper bounds are asymptotically tight.*


For any positive integers m and $k \leq m$, let $s_m(P)$ and $s_{m,k}(P)$ be the number of m -squares and (m, k) -squares, respectively, among P . We then define

$$\begin{aligned} \sigma_m(n) &:= \min_{|P|=n} s_m(P), & \sigma_{m,k}(n) &:= \min_{|P|=n} s_{m,k}(P), \\ \Sigma_m(n) &:= \max_{|P|=n} s_m(P), & \text{and} & \\ \Sigma_{m,k}(n) &:= \max_{|P|=n} s_{m,k}(P). \end{aligned}$$

In the following, we show the asymptotically tight bounds on these quantities for $m = 4$. It is obvious that $\sigma_{4,1}(n) = \Sigma_{4,1}(n) = n$ and $\Sigma_{4,k}(n) \leq \binom{n}{k}$ for $2 \leq k \leq 4$. So, we have $\sigma_4(n) = \Omega(n)$ and $\Sigma_4(n) = O(n^4)$.

4.1 Upper bounds

Here, we prove the upper bounds $\Sigma_{4,k}(n)$ for $2 \leq k \leq 4$ and $\Sigma_4(n)$ are quadratic in n . We first show that there exists a point set P with $n = |P|$ having $\Omega(n^2)$ many 4-squares.



■ **Figure 7** Illustration of a point set P_n with $\Omega(n^2)$ many 4-squares.

► **Lemma 4.3.** *For any $n \geq 4$, there exists a set P_n of n points such that $s_{4,k}(P_n) = \Omega(n^2)$ for each $2 \leq k \leq 4$ and thus $s_4(P_n) = \Omega(n^2)$.*

This already proves that $\Sigma_{4,2}(n) = \Theta(n^2)$. In the following, we prove the matching upper bounds of $\Sigma_{4,3}(n)$ and $\Sigma_{4,4}(n)$.

Upper bound of $\Sigma_{4,3}(n)$

Any $(4, 3)$ -square is of one of the four types: non-stapled $(4, 3)$ and stapled $(4, 3)$ -(a–c) types, see Figure 1. We bound the number of $(4, 3)$ -squares of each type, separately.

First, consider the stapled $(4, 3)$ -(b–c) types. Note that if S is in this case with two contact points $p, q \in P$ on its stapled side, then one of p and q also lie on a corner of S .

► **Lemma 4.4.** *Let $S \in \mathfrak{S}_4(\phi)$ be a square of stapled $(4, 3)$ -(b)- or (c)-type with contact type κ , and $v \in V(\phi)$ be the vertex with $\kappa_v = \kappa$. Then, there exists a stapled $(4, 2)$ -square $S' \in \mathfrak{S}_4(\phi)$ with contact type κ' such that v is adjacent to $v' \in V(\phi)$ with $\kappa_{v'} = \kappa'$ in $VG(\phi)$.*

From Lemma 3.1, we know that each vertex $v' \in V(\phi)$ of stapled (4, 2)-type has two growing edges whose growing direction is outwards from v' . Hence, for each vertex of stapled (4, 2)-type, there can be at most two adjacent vertices whose corresponding square is larger and has three contact points. This implies that the number of stapled (4, 3)-(b)- and (c)-type squares is at most twice the number of stapled (4, 2)-squares, which is $O(n^2)$, by Lemma 4.4.

Next, we consider the other two types of (4, 3)-squares: non-stapled (4, 3)-type and stapled (4, 3)-(a)-type. Consider any (4, 3)-square $S \in \mathfrak{S}_4(\phi)$ whose type is one of the above two. Without loss of generality, assume that a contact point $p \in P$ lies on the bottom-left corner of S . Then, regardless of its specific type, the other two contact points $q_1, q_2 \in P$ lie on either the top or the right side of S . So, $\kappa = \{(p, \ominus), (p, \omin�), (q_1, \boxplus_1), (q_2, \boxplus_2)\}$ for $\boxplus_1, \boxplus_2 \in \{\boxplus, \boxminus\}$. See Figure 1. Notice that q_1 and q_2 are two equidistantly closest points from p under the distance function d_ϕ among those points in the quadrant with apex p in orientation ϕ .

In the following, we count the number of those pairs (q_1, q_2) with the discussed property for each fixed $p \in P$. This bounds the number of those (4, 3)-squares with p on the bottom-left corner. More precisely, let ℓ_θ be the upward half-line from p in orientation $\theta \in [0, \pi)$, and \angle_θ be the cone with apex p and angle span $\pi/4$ defined by two half-lines ℓ_θ and $\ell_{\theta+\pi/4}$. Then, for any $\theta \in \mathbb{O}$, consider the two subsets

$$Q_{\boxplus}(\theta) := P \setminus \{p\} \cap \angle_\theta \quad \text{and} \quad Q_{\boxminus}(\theta) := P \setminus \{p\} \cap \angle_{\theta+\pi/4}.$$

For any $q \in P \setminus \{p\}$, define a function $f_q: \mathbb{O} \rightarrow \mathbb{R}$ to be

$$f_q(\theta) := \begin{cases} d_\theta(p, q) & \text{if } q \in Q_{\boxplus}(\theta) \cup Q_{\boxminus}(\theta) \\ \infty & \text{otherwise} \end{cases}.$$

Then, our task is reduced to find the complexity of the lower envelope F of the functions f_q for all $q \in P \setminus \{p\}$. By a trick similar to that used in Hershberger [27], we show the following.

► **Lemma 4.5.** *The complexity of the lower envelope F of f_q for all $q \in P \setminus \{p\}$ is $O(n)$.*

This proves that the number of non-stapled (4, 3)-squares and stapled (4, 3)-(a)-squares among P is at most $O(n^2)$. Combining the above discussions about stapled (4, 3)-(b–c) squares and Lemma 4.3, we conclude that $\Sigma_{4,3}(n) = \Theta(n^2)$.

Upper bound of $\Sigma_{4,4}(n)$

Now, we prove the matching upper bound on the number of (4, 4)-squares. Any (4, 4)-square is one of the three types: non-stapled (4, 4) and stapled (4, 4)-(a–b) types. See Figure 1.

► **Lemma 4.6.** *Let $S \in \mathfrak{S}_4(\phi)$ be a stapled (4, 4)-square with contact type κ , and $v \in V(\phi)$ be the vertex with $\kappa_v = \kappa$. Then, there exists a stapled (4, 3)-square $S' \in \mathfrak{S}_4(\phi)$ with contact type κ' such that v is adjacent to $v' \in V(\phi)$ in $VG(\phi)$ with $\kappa_{v'} = \kappa'$.*

From Lemma 3.1, we know that each vertex $v' \in V(\phi)$ of stapled (4, 3)-type has at most one growing edge whose growing direction is outwards from v' . Hence, for each vertex of stapled (4, 3)-type, there can be at most one adjacent vertex whose corresponding square is larger and has four contact points. This implies that the number of stapled (4, 4)-squares is at most the number of stapled (4, 3)-squares, which is $O(n^2)$ by Lemmas 4.4 and 4.6.

Lastly, we bound the number of non-stapled (4, 4)-squares. Recall that we have so far proved that the number of 4-squares whose type is not non-stapled (4, 4)-type is $O(n^2)$.

► **Lemma 4.7.** *In the graph \mathfrak{G} defined in Section 3, any non-stapled (4, 4)-square is adjacent to at least one 4-square that is not of non-stapled (4, 4)-type.*

13:12 Empty Squares in Arbitrary Orientation Among Points


Lemma 3.7 states that the degree of each 4-square in \mathfrak{S}_4 is of degree at most four, and hence is adjacent to at most four non-stapled $(4, 4)$ -squares. Since the number of 4-squares that is not of non-stapled $(4, 4)$ -type is $O(n^2)$ and each non-stapled $(4, 4)$ -square is adjacent to one of them by Lemma 4.7, we conclude that the number of non-stapled $(4, 4)$ -squares is also $O(n^2)$. This proves the upper bound of Theorem 4.2.

Since $\Sigma_4(n) \leq \sum_{1 \leq k \leq 4} \Sigma_{4,k}(n)$, we have that $\Sigma_4(n) = O(n^2)$. By Lemma 4.3, we have $\Sigma_4(n) = \Omega(n^2)$, completing our proof for the claimed upper bound of Theorem 4.1.

4.2 Lower bounds

We then turn into proving the lower bounds. As discussed above, we already have $\sigma_4(n) = \Omega(n)$, which matches the claimed lower bound in Theorem 4.1. Here, we show $\Omega(n)$ lower bounds for $\sigma_{4,2}(n)$ and $\sigma_{4,3}(n)$, and then construct a point set with $O(n)$ 4-squares.

► **Lemma 4.8.** *For any integer $n \geq 3$, $\sigma_{4,2}(n) = \Omega(n)$ and $\sigma_{4,3}(n) = \Omega(n)$.*



■ **Figure 8** Illustration of a point set P'_n with $O(n)$ many 4-squares.

We finally construct a point set having a small number of 4-squares.

► **Lemma 4.9.** *For any integer $n \geq 1$, there exists a set P'_n of n points such that $s_{4,2}(P'_n) = O(n)$, $s_{4,3}(P'_n) = O(n)$, $s_{4,4}(P'_n) = 0$, and thus $s_4(P'_n) = O(n)$.*

Consequently, by Lemma 4.9, we have $\sigma_{4,2}(n) = \Theta(n)$, $\sigma_{4,3}(n) = \Theta(n)$, and thus $\sigma_4(n) = \Theta(n)$, while $\sigma_{4,4}(n) = 0$. This proves the lower bounds in Theorems 4.1 and 4.2.

5 Maintaining the L_∞ Voronoi Diagram under Rotation

Lemma 3.6 implies that for any two consecutive degenerate orientations $\phi_1, \phi_2 \in \mathbb{O}$ the vertex set $V(\theta)$ stays the same for all $\phi_1 < \theta < \phi_2$. By Lemma 3.2, a change in the edge set $E(\theta)$ happens when and only when its incident vertices change, so $\text{VD}(\theta)$ for all $\phi_1 < \theta < \phi_2$ are combinatorially equivalent. So, the combinatorial change of $\text{VD}(\theta)$ happens at every degenerate orientation $\theta = \phi$ only.

In the following, $\epsilon \in \mathbb{R}$ denotes arbitrarily small positive real. For any degenerate orientation $\phi \in \mathbb{O}$ and a 4-square $S \in \mathfrak{S}_4(\phi)$ with contact type κ , let $V_S^- \subseteq V(\phi - \epsilon)$ and $V_S^+ \subseteq V(\phi + \epsilon)$ be the sets of regular vertices v in $V(\phi - \epsilon)$ and $V(\phi + \epsilon)$, respectively, such that $\kappa_v \subset \kappa$. For each degenerate orientation $\phi \in \mathbb{O}$, define

$$V^-(\phi) := V(\phi - \epsilon) \setminus V(\phi + \epsilon) \quad \text{and} \quad V^+(\phi) := V(\phi + \epsilon) \setminus V(\phi - \epsilon)$$

to be the sets of vertices to be deleted and inserted, respectively, as θ goes through ϕ . Similarly, define

$$E^-(\phi) := E(\phi - \epsilon) \setminus E(\phi + \epsilon) \quad \text{and} \quad E^+(\phi) := E(\phi + \epsilon) \setminus E(\phi - \epsilon).$$

Lemmas 3.6 and 3.7 imply the following.

► **Lemma 5.1.** *For any degenerate orientation $\phi \in \mathbb{O}$, the following hold:*

- (i) $V^-(\phi) = \bigcup_{S \in \mathfrak{S}_4(\phi)} V_S^-$ and $V^+(\phi) = \bigcup_{S \in \mathfrak{S}_4(\phi)} V_S^+$.
- (ii) $E^-(\phi)$ and $E^+(\phi)$ consist of edges incident to a vertex in $V^-(\phi)$ and $V^+(\phi)$, respectively.
- (iii) $|V^-(\phi)| + |V^+(\phi)| + |E^-(\phi)| + |E^+(\phi)| = \Theta(|\mathfrak{S}_4(\phi)|)$.

5.1 Events

Thus, every combinatorial change of $\text{VD}(\theta)$ can be specified by finding all degenerate orientations and all 4-squares. For the purpose, our algorithm handles *events*.

- An *edge event* is a pair (e, ϕ) for a bounded edge $e \in E(\phi - \epsilon)$ and an orientation $\phi \in \mathbb{O}$ such that the embedding $\hat{e} \in \hat{E}(\phi - \epsilon)$ of e is about to collapse into a point in orientation ϕ . As a result, the two vertices u, v incident to e are merged into one with contact type $\kappa_u \cup \kappa_v$, and there is a unique 4-square $S \in \mathfrak{S}_4(\phi)$ such that $S = S_{\kappa_u}(\phi) = S_{\kappa_v}(\phi)$ and its contact type is $\kappa_u \cup \kappa_v$. We call S the *relevant square* to this edge event.
- An *align event* is a triple (p, q, ϕ) for distinct $p, q \in P$ and $\phi \in \mathbb{O}$ such that the orientation of segment pq is either ϕ or $\phi + \pi/2$ and there is a 4-square in $\mathfrak{S}_4(\phi)$ one of whose sides contains both p and q . Each such 4-square is called *relevant* to this align event.

We then observe the following lemmas.

► **Lemma 5.2.** *For any degenerate orientation $\phi \in \mathbb{O}$, an event occurs at ϕ . More precisely, every stapled 4-square in $\mathfrak{S}_4(\phi)$ is relevant to an align event at ϕ and every non-stapled 4-square in $\mathfrak{S}_4(\phi)$ is relevant to an edge event at ϕ .*

We call an align event (p, q, ϕ) an *outer align event* if both p and q appear consecutively on the boundary of $\text{OH}(\phi)$ or, otherwise, an *inner align event*. By Lemma 3.3, we can precompute all outer align events using the algorithm by Alegría-Galicia et al. [4].

► **Lemma 5.3.** *There are $O(n)$ outer align events and we can compute them in $O(n \log n)$ time.*


We then observe the following for inner align events. See also Figure 9.

► **Lemma 5.4.** *If an inner align event occurs at $\phi \in \mathbb{O}$, then there is a stapled (4, 3)-square in $\mathfrak{S}_4(\phi)$ that is relevant to an edge event that occurs at ϕ .*

Lemma 5.4 implies that every inner align event can be noticed by handling an edge event whose relevant 4-square is of stapled (4, 3)-type. This, together with Lemma 5.3, allows us to maintain the diagram $\text{VD}(\theta)$ in an efficient and output-sensitive way, as we do not need to test all pairs of points $p, q \in P$ for potential align events.

In order to catch every edge event, we define the potential edge event as follows: for any regular orientation $\theta \in \mathbb{O}$ and any bounded edge $e = uv \in E(\theta)$, the *potential edge event* $w(e, \theta)$ for e and θ is a pair (e, ϕ) such that $\theta < \phi < \pi/2$ and $S_{\kappa_u}(\phi) = S_{\kappa_v}(\phi)$, regardless of its emptiness. If such ϕ does not exist, then $w(e, \theta)$ is undefined.

► **Lemma 5.5.** *The potential edge event $w(e, \theta)$ is uniquely defined, unless undefined. Given e and θ , one can decide if $w(e, \theta)$ is defined and compute it, if defined, in $O(1)$ time.*



■ **Figure 9** Illustration to the proof of Lemma 5.4. This figure shows the combinatorial changes around the vertex corresponding to a stapled $(4, 2)$ -square and incident edges (a) when p' lies on the left side of S' and (b) when p' lies on the top side of S' , at $\theta = \phi - \epsilon$ and $\theta = \phi$. The red edge \hat{e} in $\phi - \epsilon$ is collapsed into the red vertex in ϕ , and the blue edges are those that are non-regular.

5.2 Algorithm

Our algorithm maintains the combinatorial structure $VG(\theta)$ of the Voronoi diagrams $VD(\theta)$ as $\theta \in \mathbb{O}$ continuously increases from 0 to $\pi/2$. For the purpose, we increase θ and stop at every degenerate orientation ϕ to find all 4-squares in $\mathfrak{S}_4(\phi)$ and update $VG(\theta)$ according to the corresponding changes. For the purpose, we maintain data structures, keeping the invariants at the current orientation $\theta \in \mathbb{O}$ as follows.

- The *graph* $G = (V, E)$ stores the current Voronoi graph $VG(\theta) = (V(\theta), E(\theta))$.
- The *event queue* \mathcal{Q} is a priority queue that stores potential edge events $w(e, \theta)$ for all $e \in E$ and all outer align events that occur after θ , ordered by their associated orientations.
- The *search tree* \mathcal{T} is a balanced binary search tree on the set $K := P \times \{\bullet, \circ, \square, \diamond\}$ of all contact pairs indexed by any total order on K . Each node labeled by $(p, \square) \in K$ stores the set of all regular vertices $v \in V$ such that $(p, \square) \in \kappa_v$, denoted by $\mathcal{T}(p, \square)$, into a sorted list $L(p, \square)$ by the order along the boundary of the face of $VD(\theta)$ for (p, \square) .

Note that the structures we maintain stay the same between any two consecutive degenerate orientations. Also, by Lemmas 3.4 and 5.3, the space used by the data structures is bounded by $O(n)$ at any time by the invariants.

Our algorithm runs in two phases: the initialization and the main loop. Without loss of generality, we assume that $0 \in \mathbb{O}$ is a regular orientation. In the initialization phase, we initialize the data structures for $\theta = 0$. We first compute $VG(0)$ by any optimal algorithm computing the L_∞ Voronoi diagram [29, 30]. Then, for any regular vertex $v \in V(0)$, we insert v into V and \mathcal{T} ; for any bounded edge $e \in E(0)$, we insert e into E , we compute the potential edge event $w(e, 0)$, if defined, and insert it into \mathcal{Q} . Compute all outer align events by Lemma 5.3 and insert them into \mathcal{Q} .

We are then ready to run the main loop of our algorithm from the current orientation $\theta = 0$. In the main loop, we repeatedly recognize the next degenerate orientation $\phi > \theta$ by finding an event with a smallest orientation from \mathcal{Q} , collect all events that occur at ϕ by extracting them from \mathcal{Q} , and handle them by performing the following two steps: (1) computing all 4-squares in $\mathfrak{S}_4(\phi)$ and (2) updating our structures as θ proceeds over ϕ .

For the first step, let W be the set of all events in \mathcal{Q} whose associated orientation is ϕ . The set W can be obtained by repeatedly performing operations on the event queue \mathcal{Q} ; check if the associated orientation of the minimum element in \mathcal{Q} is exactly ϕ and extract it, if so.

For each event $w \in W$, we find all squares relevant to w , according to the type of w . We initialize $\mathfrak{S}_4(w)$ to be an empty set as a variable, and will finally consist of all 4-squares relevant to w . Consequently, we have the following.

► **Lemma 5.6.** *The first step of the main loop can be done in $O(|\mathfrak{S}_4(\phi)| \log n)$ time and we have $\mathfrak{S}_4(\phi) = \bigcup_{w \in W} \mathfrak{S}_4(w)$ in the end.*

In the second step, we first specify the sets $V^-(\phi)$, $V^+(\phi)$, $E^-(\phi)$, and $E^+(\phi)$, and then update our data structures to keep the invariants accordingly. For each $S \in \mathfrak{S}_4(\phi)$, we compute V_S^- and V_S^+ by Lemma 3.7 and its proof as illustrated in Figure 6. By Lemma 5.1(i), we obtain $V^-(\phi)$ and $V^+(\phi)$. By Lemma 5.1(ii), we can compute the edge sets $E^-(\phi)$ and $E^+(\phi)$ by searching the neighbors of $V^-(\phi)$ in $VG = VG(\theta)$.

► **Lemma 5.7.** *The sets $E^-(\phi)$ and $E^+(\phi)$ can be found in $O(|\mathfrak{S}_4(\phi)| \log |\mathfrak{S}_4(\phi)|)$ time.*

We are ready to update our structures for $\phi + \epsilon$ for arbitrarily small $\epsilon > 0$. Note that we currently have $V = V(\theta) = V(\phi - \epsilon)$ and $E = E(\theta) = E(\phi - \epsilon)$. We update V and E as follows: delete all vertices in $V^-(\phi)$ from V and all edges in $E^-(\phi)$ from E , and then insert all vertices in $V^+(\phi)$ into V and all edges in $E^+(\phi)$ into E . Then, update \mathcal{T} and \mathcal{Q} as follows: We delete each $v \in V^-(\phi)$ from \mathcal{T} and insert each $v \in V^+(\phi)$ into \mathcal{T} . For each $e \in E^-(\phi)$, we delete the potential edge event for e from \mathcal{Q} . For each $e \in E^+(\phi)$, we compute the potential edge event $w(e, \phi + \epsilon)$ by Lemma 5.5 and insert it into \mathcal{Q} , if defined. Lastly, set θ to be $\phi + \epsilon$.

We finally conclude the following.

► **Theorem 5.8.** *Given a set P of n points in general position, the total amount of combinatorial changes of the L_∞ Voronoi diagram of P while the axes rotate by $\pi/2$ is bounded by $\Theta(s_4)$, where s_4 denotes the number of 4-squares among P . The combinatorial structure of the Voronoi diagram can be maintained explicitly in total $O(s_4 \log n)$ time using $O(n)$ space.*

► **Corollary 5.9.** *Given a set P of n points in general position, we can compute all 4-squares among points in P in $O(s_4 \log n)$ time and $O(n)$ space.*

6 Maximal Empty Squares

It is now clear that our algorithm in the previous section collects a full description of all maximal empty squares in $O(n^2 \log n)$ time and its complexity is $O(n^2)$. Hence, it is not difficult to derive an algorithm that finds a largest empty square over all orientations.

► **Theorem 6.1.** *Given a set P of n points in the plane, a largest empty square among P in arbitrary orientation can be computed in worst-case $O(n^2 \log n)$ time.*

Some query versions of the problem can also be considered.

► **Theorem 6.2.** *Given a set P of n points, in $O(n^2 \log n)$ time, we can preprocess P into a data structure of size $O(n^2 \alpha(n))$ that answers the following query in $O(\log n)$ time: given an orientation $\beta \in \mathbb{O}$, find a largest empty square in orientation β .*

► **Theorem 6.3.** *Given a set P of n points, in $O(s_4 \log n)$ time, we can preprocess P into a data structure of size $O(s_4)$ that answers the following query in $O(\log n)$ time: given a point $c \in \mathbb{R}^2$ and $\beta \in \mathbb{O}$, find a largest empty square centered at c in orientation β .*

By a similar approach to that used in Bae [6], we reduce the problem of computing a square annulus of minimum width or area to that of finding all maximal empty squares.

► **Theorem 6.4.** *Given a set P of n points, a square annulus of minimum width or minimum area in arbitrary orientation that encloses P can be computed in $O(n^2 \log n)$ time.*

References

- 1 M. Abellanas, Ferran Hurtado, C. Icking, L. Ma, B. Palop, and P.A. Ramos. Best fitting rectangles. In *Proc. Euro. Workshop Comput. Geom. (EuroCG 2003)*, 2003.
- 2 A. Aggarwal, M.M. Klawe, S. Moran, P. Shor, and R. Wilber. Geometric applications of a matrix-searching algorithm. *Algorithmica*, 2:195–208, 1987.
- 3 A. Aggarwal and S. Suri. Fast algorithm for computing the largest empty rectangle. In *Proc. 3rd ACM Sympos. Comput. Geom. (SoCG 1987)*, pages 278–290, 1987.
- 4 Carlos Alegría-Galicia, David Orden, Carlos Seara, and Jorge Urrutia. On the $O\beta$ -hull of a planar point set. *Comput. Geom.: Theory Appl.*, 68:277–291, 2018.
- 5 Sang Won Bae. Computing a minimum-width square annulus in arbitrary orientation. *Theoret. Comput. Sci.*, 718:2–13, 2018.
- 6 Sang Won Bae. On the minimum-area rectangular and square annulus problem, 2019. submitted to CGTA. [arXiv:1904.06832](https://arxiv.org/abs/1904.06832).
- 7 Sang Won Bae, Chunseok Lee, Hee-Kap Ahn, Sunghee Choi, and Kyung-Yong Chwa. Computing minimum-area rectilinear convex hull and L-shape. *Comput. Geom.: Theory Appl.*, 42(9):903–912, 2009.
- 8 Sang Won Bae and Sang Duk Yoon. Empty squares in arbitrary orientation among points, 2019. [arXiv:1911.12988](https://arxiv.org/abs/1911.12988).
- 9 I. Bárány and Z. Füredi. Empty simplices in Euclidean space. *Canad. Math. Bull.*, 30:436–445, 1987.
- 10 I. Bárány and P. Valtr. A positive fraction Erdős–Szekeres theorem. *Discrete Comput. Geom.*, 19(3):335–342, 1998.
- 11 I. Bárány and P. Valtr. Planar point sets with a small number of empty convex polygons. *Studia Sci. Math. Hungar.*, 41:243–269, 2005.
- 12 J.-D. Boissonnat, M. Sharir, B. Tagansky, and M. Yvinec. Voronoi diagrams in higher dimensions under certain polyhedral distance functions. *Discrete Comput. Geom.*, 19(4):485–519, 1998.
- 13 J.E. Boyce, D.P. Dobkin, R.L. Drysdale, and L.J. Guibas. Finding extreme polygons. *SIAM J. Comput.*, 14:134–147, 1985.
- 14 Jeet Chaudhuri, Subhas C. Nandy, and Sandip Das. Largest empty rectangle among a point set. *J. Algo.*, 46:54–78, 2003.
- 15 B. Chazelle, R.L. Drysdale, and D.T. Lee. Computing the largest empty rectangle. *SIAM J. Comput.*, 15:300–315, 1986.
- 16 David P. Dobkin, Herbert Edelsbrunner, and Mark H. Overmars. Searching for empty convex polygons. *Algorithmica*, 5(1):561–571, 1990.
- 17 R.L. Drysdale and J.W. Jaromczyk. A note on lower bounds for the maximum area and maximum perimeter k -gon problems. *Inform. Proc. Lett.*, 32(6):301–303, 1989.
- 18 A. Dumitrescu. Planar sets with few empty convex polygons. *Studia Sci. Math. Hungar.*, 36:93–109, 2000.
- 19 David Eppstein. New algorithms for minimum area k -gons. In *Proc. 3rd Annu. ACM-SIAM Sympos. Discrete Algo. (SODA '92)*, pages 83–88, 1992.
- 20 David Eppstein, Mark Overmars, Günter Rote, and Gerhard Woeginger. Finding minimum area k -gons. *Discrete Comput. Geom.*, 7(1):45–58, 1992.
- 21 Pual Erdős. Some more problems on elementary geometry. *Austral. Math. Soc. Gaz.*, 5:52–54, 1978.
- 22 Pual Erdős and George Szekeres. A combinatorial problem in geometry. *Compositio Math.*, 2:463–470, 1935.
- 23 Pual Erdős and George Szekeres. On some extremum problems in elementary geometry. *Ann. Univ. Sci. Budapest Eötvös Sect. Math.*, 3–4:53–62, 1961.
- 24 Tobias Gerken. Empty convex hexagons in planar point sets. *Discrete Comput. Geom.*, 39(1):239–272, 2008.

- 25 Olga N. Gluchshenko, Horst W. Hamacher, and Arie Tamir. An optimal $O(n \log n)$ algorithm for finding an enclosing planar rectilinear annulus of minimum width. *Operations Research Lett.*, 37(3):168–170, 2009.
- 26 H. Harborth. Kovexe Fünfecke in ebenen Punktmengen. *Elem. Math.*, 33:116–118, 1978.
- 27 J. Hershberger. Finding the upper envelope of n line segments in $O(n \log n)$ time. *Inform. Proc. Lett.*, 33:169–174, 1989.
- 28 J.D. Horton. Sets with no empty convex 7-gons. *Canad. Math. Bull.*, 26:482–484, 1983.
- 29 D. T. Lee. Two-dimensional Voronoi diagrams in the L_p -metric. *J. ACM*, 27:604–618, 1980.
- 30 D.T. Lee and C.K. Wong. Voronoi diagrams in $L_1(L_\infty)$ metrics with 2-dimensional storage applications. *SIAM J. Comput.*, 9:200–211, 1980.
- 31 M. McKenna, J. O'Rourke, and S. Suri. Finding the largest rectangle in an orthogonal polygon. In *Proc. 23rd Annual Allerton Conf. Comm. Control Comput.*, 1985.
- 32 J.S.B. Mitchell, G. Rote, G. Sundaram, and G. Woeginger. Counting convex polygons in planar point sets. *Inform. Proc. Lett.*, 56(1):45–49, 1995.
- 33 W. Morris and V. Soltan. The Erdős–Szekeres problem on points in convex position—a survey. *Bull. Amer. Math. Soc.*, 33:437–458, 2000.
- 34 Walter Morris and Valeriu Soltan. The Erdős–Szekeres problem. In John Forbes Nash, Jr. and Michael Th. Rassias, editors, *Open Problems in Mathematics*, pages 351–375. Springer International Publishing, Cham, 2016.
- 35 A. Naamad, D.T. Lee, and W.L. Hsu. On the maximum empty rectangle problem. *Discrete Appl. Math.*, 8:267–277, 1984.
- 36 C. M. Nicolás. The empty hexagon theorem. *Discrete Comput. Geom.*, 38(2):389–397, 2007.
- 37 M. Orlowski. A new algorithm for largest empty rectangle problem. *Algorithmica*, 5:65–73, 1990.
- 38 Rom Pinchasi, Radoš Radoičić, and Micha Sharir. On empty convex polygons in a planar point set. *J. Combinat. Theory, Series A*, 113(3):385–419, 2006.
- 39 G. Rote, Z. Wang, G. Woeginger, and B. Zhi. Counting k -subsets and convex k -gons in the plane. *Inform. Proc. Lett.*, 38(4):149–151, 1991.
- 40 G. Rote and G. Woeginger. Counting convex k -gons in planar point sets. *Inform. Proc. Lett.*, 41(4):191–194, 1992.
- 41 George Szekeres and Lindsay Peters. Computer solution to the 17-point Erdős–Szekeres problem. *ANZIAM J.*, 48(2):151–164, 2006.
- 42 P. Valtr. On the minimum number of empty polygons in planar point sets. *Studia Sci. Math. Hungar.*, 30:155–163, 1995.