Generalized Delaunay triangulations: Graph-theoretic properties and Algorithms

by María del Pilar Cano Vila

A thesis to obtain the

Doctoral degree in Applied Mathematics

from Universitat Politècnica de Catalunya

Supervised by

Prosenjit Bose, and
Rodrigo Silveira

Departament de Matemàtiques
Barcelona 2020

Generalized Delaunay triangulations: Graph-theoretic properties and Algorithms

by María del Pilar Cano Vila

A thesis presented to obtain the

Doctoral degree in Computer Science

from Carleton University

Supervised by

Prosenjit Bose, and
Rodrigo Silveira

School of Computer Science
Ottawa 2020

To all women in science.

"...a scientist must also be absolutely like a child. If he sees a thing, he must say that he sees it, whether it was what he thought he was going to see or not. See first, think later, then test. But always see first. Otherwise you will see only what you were expecting."

Douglas Adams - The Hitchhicker's Guide to the Galaxy

Abstract

This thesis studies different generalizations of *Delaunay triangulations*, both from a combinatorial and algorithmic point of view. The Delaunay triangulation of a point set S, denoted $DT(S)$, has vertex set S. An edge uv is in $DT(S)$ if it satisfies the *empty circle property*: there exists a circle with u and v on its boundary that does not enclose points of S. Due to different optimization criteria, many generalizations of the $DT(S)$ have been proposed. Several properties are known for $DT(S)$, yet, few are known for its generalizations. The main question we explore is: to what extent can properties of $DT(S)$ be extended for generalized Delaunay graphs?

First, we explore the connectivity of the *flip graph* of *higher order Delaunay triangulations* of a point set S in the plane. The order-k flip graph might be disconnected for $k \geq 3$, yet, we give upper and lower bounds on the flip distance from one order-k triangulation to another in certain settings.

Later, we show that there exists a length-decreasing sequence of plane spanning trees of S that converges to the *minimum spanning tree* of S with respect to an arbitrary convex distance function. Each pair of consecutive trees in the sequence is contained in a *constrained convex shape Delaunay graph*. In addition, we give a linear upper bound and specific bounds when the convex shape is a square.

With focus still on convex distance functions, we study *Hamiltonicity* in *k-order convex shape Delaunay graphs*. Depending on the convex shape, we provide several upper bounds for the minimum k for which the k-order convex shape Delaunay graph is always Hamiltonian. In addition, we provide lower bounds when the convex shape is in a set of certain regular polygons.

Finally, we revisit an *affine invariant triangulation*, which is a special type of convex shape Delaunay triangulation. We show that many properties of the standard Delaunay triangulations carry over to these triangulations. Also, motivated by this affine invariant triangulation, we study different triangulation methods for producing other affine invariant geometric objects.

Resumen

Esta tesis estudia diferentes generalizaciones de la *triangulación de Delaunay*, tanto desde un punto de vista combinatorio como algorítmico. La triangulación de Delaunay de un conjunto de puntos S, denotada $DT(S)$, tiene como conjunto de vértices a S. Una arista uv está en $DT(S)$ si satisface la *propiedad del círculo vacío*: existe un círculo con u y v en su frontera que no contiene ningún punto de S en su interior. Debido a distintos criterios de optimización, se han propuesto varias generalizaciones de la $DT(S)$. Hoy en día, se conocen bastantes propiedades de la $DT(S)$, sin embargo, poco se sabe sobre sus generalizaciones. La pregunta principal que exploramos es: ¿Hasta qué punto las propiedades de la $DT(S)$ se pueden extender para generalizaciones de gráficas de Delaunay?

Primero, exploramos la conectividad de la *gráfica de flips* de las *triangulaciones de Delaunay de orden alto* de un conjunto de puntos S en el plano. La gráfica de flips de triangulaciones de orden $k \geq 3$ podría ser disconexa, sin embargo, nosotros damos una cota superior e inferior para la distancia en flips de una triangulación de orden k a alguna otra cuando S cumple con ciertas características.

Relacionado con transformaciones entre dos triangulaciones, está el problema de transformar un árbol generador de S a otro. Nosotros probamos que existe una secuencia de árboles generadores sin cruces tal que la suma total de la longitud de las aristas con respecto a una distancia convexa arbitraria es decreciente y converge al árbol generador mínimo con respecto a la distancia correspondiente. Cada par de árboles consecutivos en la secuencia se encuentran en una *triangulación de Delaunay con restricciones*. Adicionalmente, damos una cota superior lineal para la longitud de la secuencia y cotas específicas cuando el conjunto convexo es un cuadrado.

Aún concentrados en distancias convexas, estudiamos *hamiltonicidad* en las *gráficas de Delaunay de distancia convexa de k-orden*. Dependiendo en la distancia convexa, exhibimos diversas cotas superiores para el mínimo valor de k que satisface que la gráfica de Delaunay de distancia convexa de orden-k es hamiltoniana. También damos cotas inferiores para k cuando el conjunto convexo pertenece a un conjunto de ciertos polígonos regulares.

Finalmente, re-visitamos una *triangulación afín invariante*, la cual es un

caso especial de triangulación de Delaunay de distancia convexa. Probamos que muchas propiedades de la triangulación de Delaunay estándar se preservan en estas triangulaciones. Además, motivados por esta triangulación afín invariante, estudiamos diferentes algoritmos que producen otros objetos geométricos afín invariantes.

Acknowledgements

First of all, I would like to thank Jit and Rodrigo. Both of you have been excellent supervisors and more than I could have wished for. I am so grateful for all of your lessons, your patience, your jokes, and your support. For teaching me that science can be hard but also exciting and fun, for believing in me and for your work. For the Argentinian and the Quebecois accents, and the funny questions about Mexican slang. For all this journey, thank you! Hopefully we can have our solomillo con foie soon.

Thanks to all of my co-authors during this time. For taking your time to work with me, for having fun together while doing science, for all the new things I learned from you, and for correcting my writing. Special thanks to Maria, for receiving me in Prague, for all the discussions and for taking your time on reading every sentence I wrote in our work together. In addition, thanks to Elena, for your friendship, for being the best presentation's partner, and for all of our fruitful work together.

I want to thank the people at UPC that were with me in some or other way. To all the members of the DCCG group, specially to Vera, who believed in me from the beginning and introduced me to my supervisors. To my "academic family in combinatorics": Vasiliki, Max, Christoph and Clement. Thanks for the 13:00 lunch, the after coffee, and the late night olives and vermut.

Gràcies a la meva primera i millor companya de despatx, Anna. Gracias por escucharme cada vez que dudaba de mí, por hacerme ver que no soy la única y por ser una gran amiga. Merci! - Pronto tendremos que ir a México.

I would also like to thank the people at Carleton U. Thanks to all the members of the Computational Geometry lab, for making Carleton such a nice place to be in. To the 12:00 lunch time, the coffee, the chocolates and all the beer every Friday at Mike's place. I am especially grateful for all the former and current students and postdocs that made the CGlab so fun: Sander, Anthony, Darryl, Hugo, Chris, and Saeed, and sometimes Luis.

Thanks to all the wonderful friends I made in Barcelona along this journey. A Juan y Gaby, por siempre darme asilo, cuidar mis cosas y ser tan lindos conmigo. A mis chicas cangrejo – Ko y Lili – gracias por esos primeros meses de mi doctorado, por los videos de youtube, la cumbia y el vinito. A los amigos

del barrio – Alba, Agathe, Isa y Jaime – por hacer de ese piso un hogar. Y finalmente, a Kat y Lucho, por siempre invitarme a sus ricas comidas y buenas conversaciones.

Thanks to all the wonderful friends I made in Ottawa along this journey. To Andrea, thanks for all the dancing, the Canadian experiences and your friendship. To my very first friend in Ottawa, Tanvir, thanks for all the rides, the boardgames, the trivia and food. A mi "familia" mexicana – Gera – gracias por ser tan buen amigo, por los brunches, por hablar español, por los chistes, por los memes y tanto más. To my big family at 380 Lewis: Devon, Gita, Audrey and Priscilla.

A mi física favorita, Diani, gracias por escucharme siempre, en cualquier momento, por terapearme, por alojarme y por siempre, siempre estar ahí.

Y por último pero no menos importante, quiero agradecer a mi familia. A mis padres, Ricardo y Pilar, por apoyarme siempre. A mi hermano Javi, por alentarme a hacer este doctorado, y bien o mal, escucharme. A mis hermanos, Ricardín, Bernie y Pau. Y finalmente, a Stephen, por todo el apoyo, el amor, for correcting my English, los viajes, bailes y caminatas, por todo, ¡Gracias!.

Thanks to you, for taking your time to read this thesis.

This work has been made thanks to the support of CONACyT (Consejo Nacional de Ciencia y Tecnología, México).

Thank you! ¡Gracias! Gràcies! Merci!

Contents

Introduction 1

Computational geometry is an area of discrete mathematics and computer science dealing with the study of geometric problems. For instance, the real world problem of finding the closest hospital to home can be rephrased as: given a point p and a set of points S in the plane, find the closest point of S to p. The field of computational geometry is concerned with the systematic study of algorithms and data structures associated with geometric objects, with an emphasis on exact algorithms that are asymptotically fast. An important feature of this discipline is that often the classical mathematical characterization of geometric objects does not lead to the design of efficient algorithms. As such, a focus on computational aspects is a prerequisite for guaranteeing efficient algorithms.

Combinatorial geometry and *discrete geometry* are branches of geometry that deal with combinations and arrangements of geometric objects and discrete properties of these objects, respectively. Combinatorial geometry includes aspects of topology, graph theory, number theory, and other disciplines. In particular, *geometric graph theory* lies in the union of the fields of computational geometry, combinatorial geometry and graph theory. More precisely, geometric graph theory is concerned with graphs defined by geometric means. For instance, once we find which hospital is the closest to home, the next problem is to find the fastest route from home to the hospital. This problem can be translated in the following fashion: given two vertices in a connected graph G in which each edge has weight equal to the distance between its endpoints, find the path with smallest weight in G from p to q.

A *geometric graph* has a vertex set, a set of points in the plane and edges are pairs of points joined by straight line segments. This definition of edge is in contrast with the definition in the combinatorial setting where an edge is defined as a binary relation rather than as a geometric object. The study of

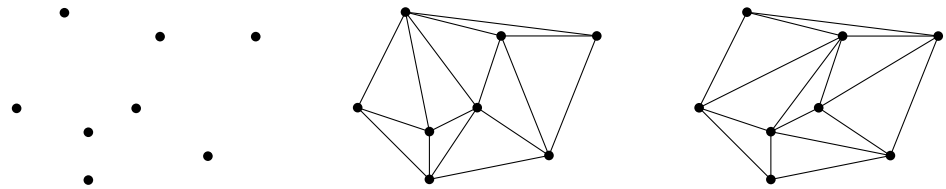

Figure 1.1: Left: a point set in the plane. Center and right: two different triangulations of the point set.

graph-theoretic properties and algorithms of geometric graphs is much closer in spirit to our research.

We define a *planar graph* as a geometric graph where no pair of edges properly intersect, i.e., their interiors are disjoint. A *triangulation* of a point set S in the plane is a planar graph with vertex set S where each of its faces define a triangle except maybe its outer face. Triangulations are one of the most studied objects in discrete and computational geometry due to their many applications in different fields like mesh generation, computer aided geometric design, and geographic information science [4, 18, 44, 95], among others.

A point set in the plane can have many different triangulations, as shown in Figure 1.1. More precisely, there are point sets with n points that have $\Omega(8.65^n)$ distinct triangulations [52], while the best upper bound is currently 30^n [94]. In the literature, much effort has been dedicated to study the efficient computation of triangulations that are optimal with respect to a desired property or criteria. Examples of properties studied include: having the minimum sum of edge-length among all possible triangulations [53, 77, 83], or maximizing the minimum angle among all possible triangulations [73, 96], among others. The *Delaunay triangulation* is the most widely used and studied type of triangulation.

The study of the Delaunay triangulation comes from two different paths that are joined. The oldest path comes from the study of *Voronoi diagrams*, also known as *Dirichlet tessellations*, whose origins date back to the 17th century by René Descartes [47]. Roughly speaking, a Voronoi diagram of a set of sites (points) in a certain space, is defined as a partition of the space into regions, where each region "belongs" to a site, and for each site p, its region consists of elements that are "influenced" by p in some way. See Figure 1.2. For instance, the sites can represent the hospitals in Barcelona and the influence is the walking distance to each hospital. Voronoi diagrams have not only been important in mathematics but in other areas, such as physics, mineralogy, geography [107, 101, 106], among a host of others; for this reason we can find them in the literature with different names. When the space is the plane and the influence is measured using the Euclidean distance, we obtain the name *Voronoi diagrams*

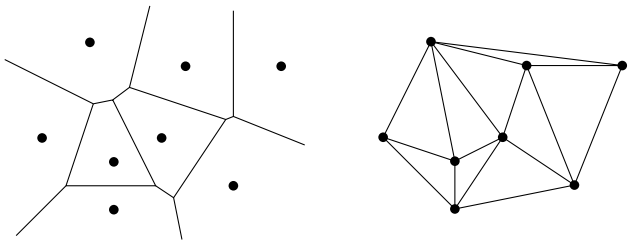

Figure 1.2: Left: Voronoi diagram of a point set. Right: Dual of the Voronoi diagram.

after Georgi Feodosjewitsch Voronoi [105] who formally introduced this concept.

Voronoi [105] also defines the geometric dual of the Voronoi diagram, where two points will be connected if and only if their Voronoi regions have a boundary in common. The second path comes from Boris Delaunay [46] who defines a geometric graph of a point set S in the plane by joining two points p and q by an edge if they satisfy the *empty circle property*, that is there exists a circle that contains p and q on its boundary and no other point of S lies in its interior. This structure is called *Delaunay graph* or *Delaunay tessellation* after Boris Delaunay. This structure is a triangulation and unique, called Delaunay triangulation and denoted $DT(S)$, when the point set S is in *general position*, that is, no three points are collinear and no four points are cocircular. The Delaunay triangulation turns out to be the dual of the Voronoi diagram.

As mentioned before, the Delaunay triangulation is one of the most studied triangulations due to its fascinating properties. One such property follows directly from one of its definitions, that it is the dual of the Voronoi diagram, which by itself is already interesting and complex. A second property, is that the Delaunay triangulation maximizes the minimum angle of all angles in the triangles among all triangulations of a given point set. This property describes "well shaped" triangles, which is one of the main reasons for its popularity. In particular, in mesh generation this well shaped property is important for modeling objects in 3D because it allows some kind of smoothness. For a survey, we refer to [25, 95].

Another interesting property of Delaunay triangulations is that it contains different important *proximity graphs*. A proximity graph is simply a graph where a pair of vertices are joined by an edge if they satisfy a certain geometric property. In particular, for the Delaunay triangulation, the context of proximity is associated with a distance function. It is also a 1.99-spanner [108] (i.e., for any pair of vertices x, y, the shortest path between x and y in the Delaunay triangulation has length that is at most 1.99 times the Euclidean distance from x to y). See [14, 87] for an encyclopedic treatment of this structure and its

many properties.

For these reasons and more, the Delaunay triangulation is one of the most studied structures in computational geometry. However, for different applications there are certain desired properties that the Delaunay triangulation almost but does not quite possess. Thus, several generalizations of the Delaunay graph have been introduced.

In this thesis we explore different combinatorial properties of some generalizations of the Delaunay graph, together with some of their different proximity subgraphs. In the remainder of this chapter, we first give some basic notions about graphs that will be fundamental throughout this thesis, followed by a brief introduction to the hierarchy of subgraphs of the Delaunay triangulation. Then, we introduce different generalizations of the Delaunay graph. Later, we give an introduction to convex distances together with a discussion about their relation with convex Gabriel graphs, and finally, we give the outline of this thesis.

1.1 Basic notions

We assume that the reader has some knowledge on the basic notions of computational geometry. If this is not the case, we refer to the books by Shamos and Preparata [89], and by de Berg et al. [43]. In this section we give some basic definitions that will be used in this thesis. Most of them are standard in graph theory but we include their definitions for the sake of completeness. See Bondy and Murty [23] for a comprehensive overview of the terminology used on graph theory.

A *graph* $G = (V, E)$ is a set V of *vertices* and a set E defined by a subset of elements in $V \times V$ called *edges*. If it is not clear from the context, the vertex set of a graph G is denoted by $V(G)$ and edge set by $E(G)$. All graphs considered are undirected, finite and simple, unless stated otherwise.

A graph H is a *subgraph* of G, denoted as $H \subseteq G$, if $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$. If $V(H) = V(G)$ we say that H is a *spanning subgraph* of G. Let $A \subseteq V$, the subgraph of G *induced* by A, denoted $G[A]$, is the graph with vertex set A where two vertices are joined by an edge if and only if they are joined by an edge in G. A graph is *complete* if for each pair of vertices u and v there exists the edge uv in $E(G)$.

A *path* in G is a finite non-empty sequence of pairwise distinct vertices $P = v_0 v_1 \dots v_{n-1} v_n$, such that $v_i v_{i+1}$ is an edge in E for all $i \in \{0, \dots, n-1\}$. A *cycle* is the union of a path and the edge joining the first and last vertex of the path. An *acyclic* graph is one that contains no cycles.

Two vertices u and v are *connected* in G if there exists a path between them

in G. A graph G is *connected* if every pair of vertices in G is connected. A *tree* is a connected acyclic graph. A *spanning tree* of G is a spanning subgraph of G that is a tree. Consider a subset A of $V(G)$, we denote by $G \setminus A$ the induced subgraph $G[V(G) \setminus A]$. The graph G is *k-connected* if there exists a subset C of $V(G)$ of size k such that $G \setminus C$ is disconnected. For instance, a tree is 1-connected, a cycle is 2-connected and a triangulation can be 2-, 3- and 4-connected.

A path that contains every vertex of G is called a *Hamiltonian path*; similarly, a *Hamiltonian cycle* of G is a cycle that contains every vertex of G. If a graph G contains a Hamiltonian cycle then G is called *Hamiltonian*. A graph is 1-*tough* if, for any k, removing k vertices from the graph splits the graph into at most k components. An important relation between the concept of 1-toughness and Hamiltonian graphs is that every Hamiltonian graph is 1-tough. However, not every 1-tough graph is Hamiltonian.

A graph G is *weighted* if each edge e in $E(G)$ is associated to a real number $w(e)$, called its weight. The *weight* of G is the total sum of the weight of all its edges. Given a weighted graph $G = (V, E)$ and a real number $t \geq 1$, a *t-spanner* of G is a spanning subgraph G' such that for every edge uv in G, there exists a path from u to v in G' whose weight is no more than t times the weight of the edge uv in G. When G' is a t-spanner of the complete graph and each edge is weighted with its Euclidean length, we denote $t = sr(G')$ and call $sr(G')$ the *spanning ratio* of G'. We say that G' is a t-spanner —or simply *spanner*— if $sr(G') = t$ is a constant.

Throughout this thesis we will assume that all graphs are geometric graphs, i.e., the edges are line segments.

Let S be a point set in the plane. Let $G = (S, E)$ be a plane graph with vertex set $V(G) = S$ and edges $E(G)$. We say that the edge pq is *visible* in G if and only if one of the following holds: (i) $pq \in E(G)$, or (ii) pq does not cross edges in $E(G)$. When pq is visible in G, we say that p and q *see each other* in G or that p is *visible to* q in G and vice versa. Let e be an edge in $E(G)$, we say that e is *blocking* a point $p \in S$ from a point $q \in S$, if the line segment pq crosses e in its interior. We define the *visibility graph* of S in G as the graph with vertices S and the set of all visible edges in G as the edge set. Notice that the visibility graph is always a connected graph.

1.2 A hierarchy of subgraphs

As mentioned before, one of the nice properties of the Delaunay triangulation is the containment of some important proximity graphs. In this section we introduce such graphs.

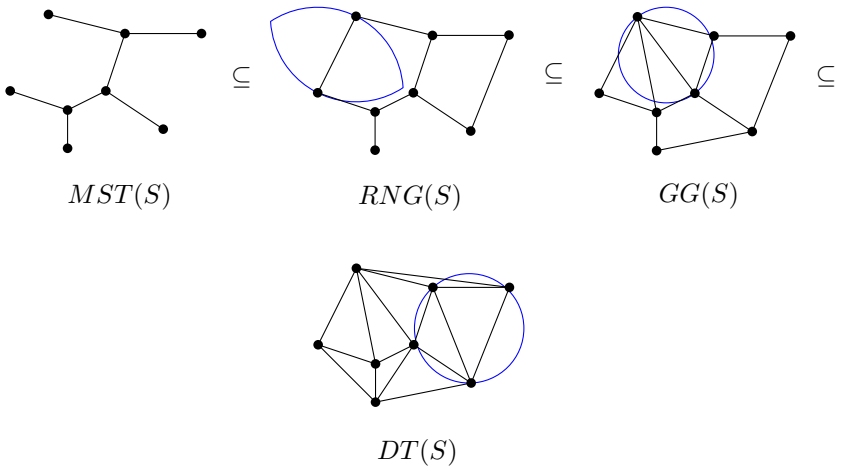

Figure 1.3: The Euclidean minimum spanning tree, the relative neighborhood graph, the Gabriel graph and the Delaunay triangulation of a point set.

The *minimum spanning tree* was introduced by Borůvka [26, 27, 84] in order to design an efficient electric distribution network of Moravia. A minimum spanning tree of a weighted graph G is a spanning tree of G with minimum weight. The minimum spanning tree has been widely studied due to its several applications in the design of networks, in the approximation of the traveling salesman problem, cluster analysis [13, 42, 61, 70, 90], and more. Consider a point set S in the plane and a graph G with vertex set S. The Euclidean minimum spanning tree of G, denoted $MST(G)$, is the minimum spanning tree of G where the weight of each edge e in G is the Euclidean distance between the endpoints of e. When G is the complete graph, it is called the *Euclidean minimum spanning tree of* S and it is denoted by $MST(S)$.

In order to describe geographical variation data and develop statistical methods for categorizing sets of populations sampled from different localities, Gabriel with Sokal [106] introduced the following geometric graph. Given a point set S in the plane, two points p and q in S are joined by an edge if the smallest circle that contains p and q on its boundary does not contain other points of S in its interior. This graph is called the *Gabriel graph* after K. Ruben Gabriel and it is denoted $GG(S)$.

Motivated by pattern recognition, Toussaint introduced the *relative neighborhood graph* [102], slightly changing the notion of "relatively close" neighbours defined by Lankford [71]. A *relative neighborhood graph* of a point set S in the plane, denoted $RNG(S)$, is a geometric graph with vertex set S, where a pair of points p and q in S are connected by an edge if and only if no other point of S is at lower Euclidean distance to both of them. Geometrically speaking, two

points p and q in S are connected by an edge in $RNG(S)$ if the lune defined by the intersection of the disks centered at p and q with radius the Euclidean distance $d(p, q)$ does not contain points of S in its interior. Toussaint showed that for any point set S, the relative neighborhood graph of S is a subgraph of the Delaunay triangulation of S [102]. In fact, one of the nicest properties about all of these graphs is the following hierarchical order in containment:

$$MST(S) \subseteq RNG(S) \subseteq GG(S) \subseteq DT(S) \tag{1.1}$$

Figure 1.3 depicts an example of this hierarchy of subgraphs.

Finally, note that the definitions of the Relative neighborhood graph, Gabriel graph and Delaunay triangulation satisfy an *empty region property*. That is, a pair of points of S are joined by an edge if there exists certain closed region that contains the two points but no other point of S in its interior.

1.3 Generalized Delaunay graphs

The Delaunay triangulation, as mentioned before, is an important triangulation due to its several applications in mesh generation, finite element methods, GIS, graphics, and more. For a survey we refer to [14, 37, 25, 87, 95]. However, the fact that the Delaunay triangulation of a point set S in general position is unique can become an issue in certain application domains where extra flexibility is needed.

For instance, triangulations are often used to model terrains. In this case, the points in S are samples of a 3D surface, thus they also have an elevation. Yet, the Delaunay triangulation ignores the elevation information, potentially resulting in poor terrain models where important terrain features, such as valleys or ridge lines, are ignored [44, 63]. Then it would be nice if the triangulation could preserve the additional information as well. This leads us to consider the construction of a "nice" triangulation for a set of points and line segments. This motivated Lee [74, 75] to define the *"Generalized Delaunay Triangulation"*, such graph was also called *"obstacle triangulation"* by Chew [41] due to its applications in motion planning with obstacles. Finally, Chew [40] introduced such graph as the *constrained Delaunay triangulation*.

Let S be a point set in the plane and G be a geometric graph with vertex set S and edges called *constraints* The constrained Delaunay triangulation of G, denoted $DG(G)$, is a geometric graph with vertex set S where each pair of points p and q in S are joined by an edge if either pq is a constraint or there exists a circle that contains p and q that does not enclose points of S that are visible to both p and q. Note that when G has no edges, then we obtain the Delaunay triangulation. Moreover, since this definition depends on the empty

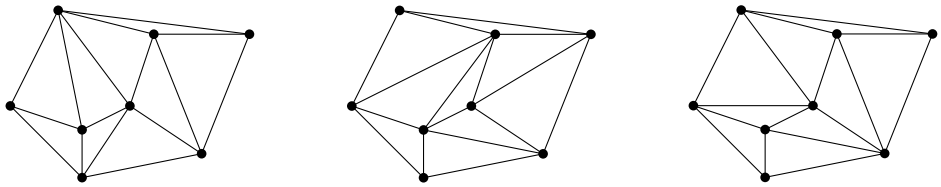

Figure 1.4: Three different order-2 triangulations of a point set.

region property, then one can similarly define the *constrained Gabriel Graph* and *constrained Relative neighborhood graph.*

Gudmundsson et al. [62] were also motivated by the terrain interpolation problem. They were trying to describe a nice triangulation that contains certain constraints, in which case the constrained Delaunay triangulation is not always the optimal. Thus, they proposed the *higher order Delaunay triangulations*, a generalization of the Delaunay triangulation that intends to provide well shaped triangles, while giving flexibility to choose from a larger set of triangulations. A triangulation T of S is an *order-k Delaunay triangulation* —or, simply, *order-k* triangulation— if the circumcircle of each triangle of T encloses at most k points of S in its interior. Note that an order-0 triangulation is a standard Delaunay triangulation. Order-k triangulations have been used for terrain modeling, minimum interference networks and triangulation of polygons [17, 97, 91]. Note that for $k \geq 1$, there might be more than one order-k triangulation. See Figure 1.4.

A very similar generalization of the Delaunay triangulation is the *k-order Delaunay Graph.* A k-order Delaunay graph –or, simply, *k-Delaunay graph*– of a point set S in the plane, denoted k-$DG(S)$, is a geometric graph with vertex set S where each pair of points p and q of S are connected by an edge if and only if there exists a circle that contains both p and q such that the circle encloses at most k points of S. This graph was introduced by Tung-Hsin Su and Ruei-Chuan Chang [100], which they used to improve the running time of solving the Euclidean bottleneck biconnected edge subgraph problem and Euclidean bottleneck matching problem. Note that the k-$DG(S)$ is not necessarily a planar graph. Also, from the definition, one can easily show that for any $k \geq 0$, k-$DG(S) \subset (k+1)$-$DG(S)$, see Figure 1.5. In particular, note that the 0-Delaunay graph is the standard Delaunay triangulation. In addition, this generalization is again defined under the empty region property, then a similar definition exists for the *k-Gabriel Graph*, denoted k-$GG(S)$ and the *k-Relative neighborhood graph*, denoted k-$RNG(S)$. In fact, the hierarchy order k-$RNG(S) \subseteq k$-$GG(S) \subseteq k$-$DG(S)$ is still preserved.

For some applications, the Euclidean distance does not provide the desired

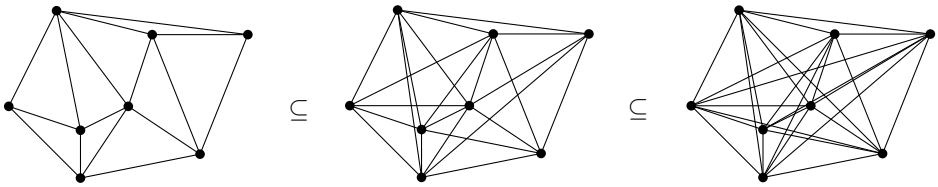

$DT(S)$ 1-$DG(S)$ 2-$DG(S)$

Figure 1.5: The 0-, 1- and 2-Delaunay graphs of a point set.

measure. In this thesis we are particularly interested in *convex distance* functions, which we will formally define in Section 1.4. Consider a compact convex set $\mathcal{C}$ that contains the origin as its center. Roughly speaking, the $\mathcal{C}$-distance from p to q, denoted $d_{\mathcal{C}}(p, q)$, is calculated in the following way. Center $\mathcal{C}$ at p and scale (expand or contract) $\mathcal{C}$ until its boundary touches q, then $d_{\mathcal{C}}(p, q)$ is the scaling factor of the resulted homothet [1] of $\mathcal{C}$.

There are two ways of defining the Delaunay triangulation — the geometric dual of the Voronoi Diagram, and with the empty circle property — these definitions are equivalent for the Euclidean distance. However, this might not be true for convex distance functions, since the symmetry property of a metric might not hold, i.e., $d_{\mathcal{C}}(p, q)$ is not necessarily equal to $d_{\mathcal{C}}(q, p)$. The definition of an edge in the empty circle property looks for the distance of the center towards the endpoints of the edge. This is in contrast to the definition of a bisector in the Voronoi diagram, which looks for the distance to each of the two sites towards points at same convex distance, i.e., for each $p, q \in S$, the bisector of p and q is defined by all points x such $d_{\mathcal{C}}(p, x) = d_{\mathcal{C}}(q, x)$. Let $\mathcal{C}'$ be the reflected shaped of $\mathcal{C}$ about the origin. Then, since $d_{\mathcal{C}}(p, q) = d_{\mathcal{C}'}(q, p)$, the dual of the Voronoi diagram in the $\mathcal{C}$-distance is equivalent to the structure with edges that satisfy the empty $\mathcal{C}'$ shape property. This leads to a surprising property, that the dual of a $\mathcal{C}$-distance Voronoi diagram is invariant under movements of the center of $\mathcal{C}$. See [14, 33, 79].

In this thesis we will study properties of the *Convex shape Delaunay graph* by the empty disk property. Thus, we define the Convex shape Delaunay graph –or, simply, *$\mathcal{C}$-Delaunay graph*– of a point set S in the plane, denoted $DG_{\mathcal{C}}(S)$, as the graph with vertex set S and for each pair p and q in S, the edge pq is in $DG_{\mathcal{C}}(S)$ provided that there exists a homothet of $\mathcal{C}$ that contains p and q on its boundary and no other point of S in its interior. Note that when $\mathcal{C}$ is a circle, then $DG_{\mathcal{C}}(S)$ is the standard $DT(S)$. It is known that the $\mathcal{C}$-Delaunay graph

[1] A *homothet* of $\mathcal{C}$ is obtained by scaling $\mathcal{C}$ with respect to the origin, followed by a translation.

is a geometric planar graph [33]. Several classes of Convex Delaunay graphs have been studied in the literature. For instance, Chew [41] showed that the $\triangle$-Delaunay graph (i.e., where the shape is an equilateral triangle instead of a disk), denoted $DG_{\triangle}(S)$, is a 2-spanner and that the $\square$-Delaunay graph (i.e., where the disk is replaced by a square), denoted $DG_{\square}(S)$, is a $\sqrt{10}$-spanner. Bose et al. [33] proved that the $\mathcal{C}$-Delaunay graph is a c-spanner where the constant c depends only on the perimeter and width of the convex shape $\mathcal{C}$.

Similarly, the *$\mathcal{C}$-Gabriel graph* of a point set S, denoted $GG_{\mathcal{C}}(S)$, is defined as the graph with vertex set S where two points p and q of S are joined by an edge if and only if there exists a smallest homothet of $\mathcal{C}$ with p and q on its boundary such that no other point of S is in its interior. The *$\mathcal{C}$-Relative neighborhood graph* of a point set S, denoted $RNG_{\mathcal{C}}(S)$, is a graph with vertex set S and for each pair p and q in S, the edge pq is in $RNG_{\mathcal{C}}(S)$ if and only if the intersection of the homothets of $\mathcal{C}$ centered at p and q with scaling factor $d_{\mathcal{C}}(p,q)$ and $d_{\mathcal{C}}(q,p)$, respectively, contains no points of S in its interior. Finally, the *$d_{\mathcal{C}}$-minimum spanning tree* of a point set S, denoted $MST_{\mathcal{C}}(S)$ in the plane, is the minimum spanning tree of the complete subgraph of S with weight for each edge e equal to the $\mathcal{C}$-distance between its endpoints. Aurenhammer and Paulini [15] recently showed that the subgraph hierarchy also holds for any compact convex set $\mathcal{C}$, i.e., $MST_{\mathcal{C}}(S) \subseteq RNG_{\mathcal{C}}(S) \subseteq GG_{\mathcal{C}}(S) \subseteq DT_{\mathcal{C}}(S)$.

Note that the definition of the constrained Delaunay graph can be generalized to *constrained $\mathcal{C}$-Delaunay graph* by replacing the circle with a homothet of $\mathcal{C}$. Consider a planar graph G with vertex set S. The constrained $\mathcal{C}$-Delaunay graph of S, denoted $DT_{\mathcal{C}}(G)$, is formally defined as follows. For each pair u and v in S, an edge $uv \in DT_{\mathcal{C}}(G)$ if either $uv \in E(G)$ or there exists a homothet of $\mathcal{C}$ that contains u and v on its boundary and no point of S in its interior is visible to both u and v. Bose et al. [35] proved that these graphs are planar. Moreover, they proved that regardless of the shape there exists a constant t such that the constrained $\mathcal{C}$-Delaunay graph is a t-spanner.

Finally, we can define the most general of the generalized Delaunay graphs used in this thesis: *the k-order $\mathcal{C}$-Delaunay graph*. Similarly to the Euclidean distance (where the empty shape is a circle) a k-order $\mathcal{C}$-Delaunay graph of a point set S in the plane, denoted k-$DG_{\mathcal{C}}(S)$, is the graph with vertex set S such that, for each pair of points $p, q \in S$, the edge pq is in k-$DG_{\mathcal{C}}(S)$ if and only if there exists a $\mathcal{C}$-disk that has p and q on its boundary and contains at most k points of S different from p and q. When $k = 0$ and $\mathcal{C}$ is a circle, k-$DG_{\mathcal{C}}(S)$ is the standard *Delaunay triangulation*. In addition, we analogously define the *k-order $\mathcal{C}$-Gabriel graph* and the *k-order Relative neighborhood* of a point set S, denoted k-$GG_{\mathcal{C}}(S)$ and k-$RNG_{\mathcal{C}}(S)$, respectively. Except that in k-$GG_{\mathcal{C}}(S)$, the homothets considered are restricted to be smallest homothets of $\mathcal{C}$ with p and q on the boundary. In fact, $MST_{\mathcal{C}}(S) \subseteq k\text{-}RNG_{\mathcal{C}}(S) \subseteq k\text{-}GG_{\mathcal{C}}(S) \subseteq k\text{-}DG_{\mathcal{C}}(S)$.

The main topic in this thesis is the study of different combinatorial and algorithmic properties of these generalized Delaunay triangulations.

1.4 Convex distances and the $\mathcal{C}$-Gabriel graph

In this section we define the convex distance functions that will be used throughout this thesis.

A *norm* of $\mathbb{R}^2$ is a nonnegative function $\rho : \mathbb{R}^2 \to \mathbb{R}^+$ with the following properties, for all $\lambda \in \mathbb{R}^+$ and $u, v \in \mathbb{R}^2$: (a) $\rho(u+v) \leq \rho(u) + \rho(v)$, (b) $\rho(\lambda v) = \lambda\rho(v)$ and, (c) if $\rho(v) = 0$ then v is the zero vector. A *metric* is a distance function $d : \mathbb{R}^2 \to \mathbb{R}^+$ such that for all $u, v, w \in \mathbb{R}^2$ the following properties hold: (a) $d(u,v) = 0 \iff u = v$, (b) $d(u,v) = d(v,u)$ and, (c) $d(u,w) \leq d(u,v) + d(v,w)$. When the function $d(u,v) = \rho(u-v)$ is a norm, then it is called a normed metric.

Let p and q be two points in the plane. Let $\mathcal{C}$ be a compact convex set that contains the origin, denoted $\bar{o}$, in its interior. We denote *the boundary of* $\mathcal{C}$ by $\partial\mathcal{C}$. The *convex distance* $d_{\mathcal{C}}(p,q)$ is defined as follows: If $p = q$, then $d_{\mathcal{C}}(p,q) = 0$. Otherwise, let $\mathcal{C}_p$ be the convex set $\mathcal{C}$ translated by the vector $\overrightarrow{p}$ and let q' be the intersection of the ray from p through q and $\partial\mathcal{C}_p$. Then, $d_{\mathcal{C}}(p,q) = \frac{d(p,q)}{d(p,q')}$ (see Figure 1.6) where d denotes the Euclidean distance.

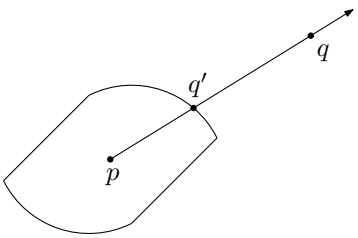

Figure 1.6: Convex distance from p to q.

The convex set $\mathcal{C}$ is the *unit $\mathcal{C}$-disk* of $d_{\mathcal{C}}$ with center $\bar{o}$, i.e., every point p in $\mathcal{C}$ satisfies that $d_{\mathcal{C}}(\bar{o}, p) \leq 1$. The *$\mathcal{C}$-disk with center c and radius r* is defined as the homothet of $\mathcal{C}$ centered at c and with scaling factor r. The *triangle inequality* holds: $d_{\mathcal{C}}(p,q) \leq d_{\mathcal{C}}(p,z) + d_{\mathcal{C}}(z,q), \forall p,q,z \in \mathbb{R}^2$. However, this distance may not define a metric when $\mathcal{C}$ is not *point-symmetric* about the origin,[2] since there may be points p, q for which $d_{\mathcal{C}}(p,q) \neq d_{\mathcal{C}}(q,p)$. When $\mathcal{C}$ is point-symmetric with respect to the origin, $d_{\mathcal{C}}$ is called a *symmetric convex distance function* and it is a metric. We will refer to such distance functions as *symmetric convex*.

[2] A shape $\mathcal{C}$ is point-symmetric with respect to a point $x \in \mathcal{C}$ provided that for every point $p \in \mathcal{C}$ there is a corresponding point $q \in \mathcal{C}$ such that $pq \in \mathcal{C}$ and x is the midpoint of pq.

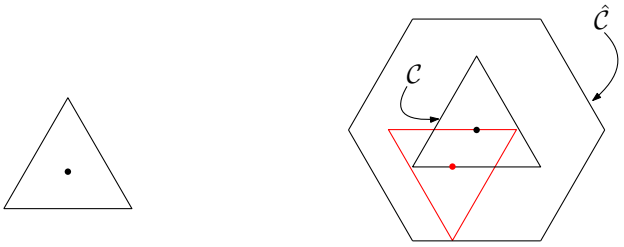

(a) (b)

Figure 1.7: (a) A triangle is a non-symmetric shape $\mathcal{C}$. (b) $\hat{\mathcal{C}}$ for this triangle is a hexagon.

Moreover, $d_{\mathcal{C}}(\bar{o}, p)$ defines a *norm* of a *metric space*. In addition, if a point p is on the line segment ab, then $d_{\mathcal{C}}(a, b) = d_{\mathcal{C}}(a, p) + d_{\mathcal{C}}(p, b)$ (see [14, Chapter 7]).

Recall that an edge pq is in the *k-order $\mathcal{C}$-Delaunay graph* of S, provided that there exists a $\mathcal{C}$-disk that has p and q on its boundary and contains at most k points of S different from p and q.

As mentioned earlier, the definition of Gabriel graphs requires the notion of a smallest homothet containing two points on its boundary. To be able to use our techniques, it is convenient to be able to associate a distance to the size of such smallest homothets, but $d_{\mathcal{C}}$ fails on defining such distance because $d_{\mathcal{C}}$ might not be symmetric when the shape is not point-symmetric. To circumvent this isssue, Aurenhammer and Paulini [15] showed how to define, from any convex shape $\mathcal{C}$, another shape that results in a distance function that is always symmetric: The set $\hat{\mathcal{C}}$ is defined as the Minkowski sum[3] of $\mathcal{C}$ and its shape reflected about its center. For an example, see Figure 1.7. The shape $\hat{\mathcal{C}}$ is point-symmetric and the $d_{\hat{\mathcal{C}}}$-distance from p to q is given by the scaling factor of a smallest homothet of $\mathcal{C}$ containing p and q on its boundary. The diameter and width of $\hat{\mathcal{C}}$ is twice the diameter and width of $\mathcal{C}$, respectively. Moreover, if $\mathcal{C}$ is point-symmetric, $d_{\hat{\mathcal{C}}}(p, q) = \frac{d_{\mathcal{C}}(p,q)}{2}$.

Thus, for any pair of points $p, q \in S$, edge pq is in the *k-order $\mathcal{C}$-Gabriel graph* of S provided that there exists a $\mathcal{C}$-disk with radius $d_{\hat{\mathcal{C}}}(p, q)$ that has p and q on its boundary and contains at most k points of S different from p and q. From the definition of k-$GG_{\mathcal{C}}(S)$ and k-$DG_{\mathcal{C}}(S)$ it follows that k-$GG_{\mathcal{C}}(S) \subseteq k$-$DG_{\mathcal{C}}(S)$, and it can be a proper subgraph. See Figure 1.8a for an example. Further, $\hat{\mathcal{C}}$ always contains $\mathcal{C}$ in its interior. However, for some non point-symmetric convex $\mathcal{C}$ it is not true that $GG_{\hat{\mathcal{C}}} \subseteq GG_{\mathcal{C}}$; see Figure 1.8b for an example.

For simplicity, denote by $\mathcal{C}_r(a, b)$ a $\mathcal{C}$-disk of radius r with the points a and b on its boundary. For the special case of a *diametral disk*, i.e., when $r = d_{\hat{\mathcal{C}}}(a, b)$,

[3]The Minkowski sum of two sets A and B is defined as $A \oplus B = \{a + b : a \in A, b \in B\}$.

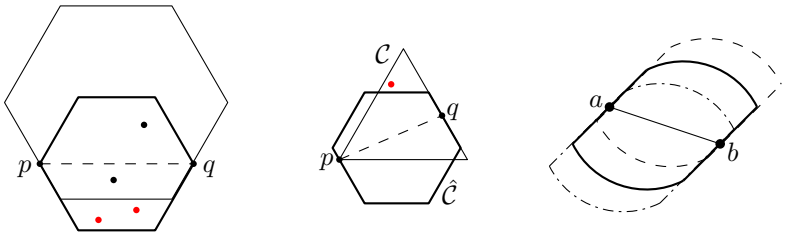

Figure 1.8: (a) $\mathcal{C}$ is a regular hexagon. Edge pq is in 2-$DG_{\mathcal{C}}(S)$ but it is not in 2-$GG_{\mathcal{C}}(S)$. (b) Edge pq is in $GG_{\hat{\mathcal{C}}}(S)$ but it is not in $GG_{\mathcal{C}}(S)$. (c) Many $\mathcal{C}$-disks $\mathcal{C}(a,b)$ may exist for a and b.

we denote it by $\mathcal{C}(a,b)$. Note that $\mathcal{C}(a,b)$ may not be unique when the boundary $\partial\mathcal{C}$ is not smooth, see Figure 1.8c. In addition, we denote by $D_{\mathcal{C}}(c,r)$ the $\mathcal{C}$-disk centered at point c with radius r.

1.5 Outline

In Chapter 2 we consider the flip graph of k-Delaunay triangulations of a point set S. The *flip graph* of S has one vertex for each possible triangulation of S, and an edge connecting two vertices when the two corresponding triangulations can be transformed into each other by a *flip* (i.e., exchanging the diagonal of a convex quadrilateral by the other one). The flip graph is an essential structure in the study of triangulations, but until now it had been barely studied for order-k Delaunay triangulations. In this work we show that, even though the order-k flip graph can be disconnected for $k \geq 3$, any order-k triangulation can be transformed into some other order-k triangulation by at most $k-1$ flips, such that the intermediate triangulations are of order at most $2k-2$, in the following settings: (1) for any $k \geq 0$ when S is in convex position, and (2) for any $k \leq 5$ and any point set S. Our results have several implications on the flip distance between order-k triangulations, as well as on their efficient algorithmic enumeration.

This chapter is based on the following publications:

[12] Elena Arseneva, Prosenjit Bose, Pilar Cano, and Rodrigo I. Silveira. Flips in higher order Delaunay triangulations. In *Proceedings of the 14th Latin American Theoretical Informatics Symposium (LATIN)*, to appear, 2020

[11] Elena Arseneva, Prosenjit Bose, Pilar Cano, and Rodrigo I. Silveira. Flips in higher order Delaunay triangulations. In *36th European Workshop on*

Computational Geometry: extended abstracts (EuroCG), 2020

In Chapter 3 we look at constrained $\mathcal{C}$-Delaunay triangulations and $d_{\mathcal{C}}$-minimum spanning trees, for an arbitrary compact convex set $\mathcal{C}$. We first extend a known result about standard Delaunay triangulations to $\mathcal{C}$-Delaunay graphs. Let $\mathcal{ST}(S)$ be the set of all plane spanning trees of a planar point set S of size n. We prove that for each element T of $\mathcal{ST}(S)$, there exists a weight-decreasing sequence of trees $T_0, \ldots, T_k$ in the $d_{\mathcal{C}}$-distance such that $T_0 = T, T_k = MST_{\mathcal{C}}(S)$ and T_i is the $d_{\mathcal{C}}$-minimum spanning tree of $DG_{\mathcal{C}}(T_{i-1})$. Hence, T_i does not cross T_{i-1} for all $i = 1, \ldots, k$. Later, we look at the $\square$-distance —also known as the L_∞-metric— and give an $\Omega(\log n)$ lower bound for the length of the sequence. Finally, we show an upper bound of $\log n$ on the length of the sequence for a case in which the constraints are crossing the edges of the $MST_{\square}(S)$ in a specific manner.

This chapter is based on the following publication:

[31] Prosenjit Bose, Pilar Cano, and Rodrigo I. Silveira. Sequences of spanning trees for L_∞-Delaunay triangulations. In *34th European Workshop on Computational Geometry: extended abstracts (EuroCG)*, 2018

In Chapter 4 we study Hamiltonicity in the $\mathcal{C}$-Delaunay graphs. More specifically, we provide upper bounds on the minimum value of k for which k-$GG_{\mathcal{C}}(S)$ is Hamiltonian. Since k-$GG_{\mathcal{C}}(S) \subseteq k$-$DG_{\mathcal{C}}(S)$, all results carry over to k-$DG_{\mathcal{C}}(S)$. In particular, we give upper bounds of 24 for every $\mathcal{C}$ and 15 for every point-symmetric $\mathcal{C}$. We also improve these bounds to 7 for squares, 11 for regular hexagons, 12 for regular octagons, and 11 for even-sided regular t-gons (for $t \geq 10$). These constitute the first general results on Hamiltonicity for convex shape Delaunay and Gabriel graphs.

In addition, we show lower bounds of $k = 3$ and $k = 6$ on the existence of a bottleneck Hamiltonian cycle in the k-order Gabriel graph for squares and hexagons, respectively. Finally, we construct a point set such that for an infinite family of regular polygons $\mathcal{P}_t$, the Delaunay graph $DG_{\mathcal{P}_t}$ does not contain a Hamiltonian cycle.

This chapter is based on the following publications:

[30] Prosenjit Bose, Pilar Cano, Maria Saumell, and Rodrigo I. Silveira. Hamiltonicity for convex shape Delaunay and Gabriel graphs. *Computational Geometry*, to appear, 2020

[28] Prosenjit Bose, Pilar Cano, Maria Saumell, and Rodrigo I. Silveira. Hamiltonicity for convex shape Delaunay and Gabriel graphs. In *Proceedings of*

the 16th Algorithms and Data Structures Symposium (WADS), pages 196–210. Springer, 2019

[29] Prosenjit Bose, Pilar Cano, Maria Saumell, and Rodrigo I. Silveira. Hamiltonicity for convex shape Delaunay and Gabriel graphs. In *35th European Workshop on Computational Geometry: extended abstracts (EuroCG)*, 2019

In Chapter 5 we revisit an affine invariant triangulation defined by Gregory M. Nielson [86], that uses the inverse of the covariance matrix of S to define an affine invariant norm, denoted A_S, and an affine invariant triangulation, denoted $DT_{A_S}(S)$. The A_S-norm is a special kind of $\mathcal{C}$-distance where $\mathcal{C}$ is replaced by a type of ellipse. We revisit the A_S-norm from a geometric perspective, and show that $DT_{A_S}(S)$ can be seen as a standard Delaunay triangulation of a transformed point set based on S. We prove that it retains all of its well-known properties. In addition, motivated by this norm and the problem of finding affine geometric methods, we provide different affine invariant sorting methods of a point set S and of the vertices of a polygon P that can be combined with well-known algorithms in order to obtain other affine invariant methods.

This chapter is based on the following publication:

[32] Prosenjit Bose, Pilar Cano, and Rodrigo I. Silveira. Affine invariant triangulations. In *31st Canadian Conference in Computational Geometry (CCCG)*, pages 250–256, 2019

Finally, in the concluding remarks of this thesis we introduce some open problems that were not studied in this thesis but are related to some of our results and generalized Delaunay triangulations.

Higher order Delaunay triangulations and flips

2

A fundamental operation to work with triangulations is the *edge flip*. It consists in removing an edge, shared by two triangles that form a convex quadrilateral, and inserting the other diagonal of the quadrilateral. A flip transforms a triangulation T into another triangulation T' that differs by exactly one edge and two triangles. The flip operation leads naturally to the definition of the *flip graph* of S. Each triangulation of S is represented by a vertex in this graph, and two vertices are adjacent if their corresponding triangulations differ by exactly one flip.

The importance of flips in triangulations comes from the fact, first proved by Lawson [72], that the flip graph is connected. Moreover, the sequence of edge flips connecting any two triangulations has length $O(n^2)$. In fact, it was later shown [73, 96] that any triangulation of S can be converted into $DT(S)$ by performing $O(n^2)$ flips. Each flip in this transformation also results in an increase of the *angle vector* of the triangulation, i.e., the vector of all angles of each of its triangles in increasing order. It is also known that the quadratic upper bound on the diameter of the flip graph is tight in general [58, 66], although it goes down to $\Theta(n)$ if the points in S are in convex position [99]. In general, computing the distance in the flip graph between two given triangulations is a difficult problem, whose complexity was open until recently, when it was shown to be APX-hard [78, 88]. This has drawn considerable attention to the study of certain subgraphs of the flip graph, which define the flip graph of certain classes of triangulations, for instance, bounded degree triangulations [7] or triangulations with perfect matchings [65]. We refer to [36] for a survey.

In this chapter we study the flip graph of higher order Delaunay triangulations. Most previous work on such triangulations focused on algorithmic questions related to finding order-k triangulations that are optimal with respect

to extra criteria [104, 97, 98], or on evaluating their effectiveness in practical settings [20, 45]. One of the few theoretical aspects studied is the asymptotic number of order-k triangulations [82]. In that work, the authors showed that for points drawn uniformly at random, already for $k = 1$ one can expect an exponential number of different order-k triangulations. However, almost nothing is known about the flip graph of order-k triangulations, except that it is connected only for $k \leq 2$ [1].

Abellanas et al. [2] studied the flip graph of triangulations that consists of only order-k edges. In their work an *order-k (Delaunay) edge* e is defined as an edge for which there exists a circle through the endpoints of e that encloses at most k points of S. All edges in an order-k triangulation are order-k. However, the converse is not true: a triangle composed of three order-k edges can have order greater than k. In fact, the lowest order triangulation containing an order-k edge can have order up to $2k-2$ [63]. Similarly as for order-k triangulations, Abellanas et al. [2] showed that the flip graph of triangulations of point sets with edges of order k is connected for $k \leq 1$, but might be disconnected for $k \geq 2$. On the other hand, they proved that for point sets in convex position the flip graph is always connected [2]. However, their proof implies an exponential bound in the diameter of the flip graph. The only previous work on the flip graph of order-k triangulations is by Abe and Okamoto [1], in the context of enumeration algorithms. They observed that for $k \leq 2$, the fact that the flip graph is connected implies that the reverse enumeration framework by Avis and Fukuda [16] can be applied to enumerate all order-k triangulations, spending polynomial time on each of them.

Our countributions. We present several structural properties of the flip graph of order-k triangulations. For points in convex position, we show that for any $k > 2$ there exist a point set in convex position for which the flip graph is not connected. However, we prove that no order-k triangulation is too far from all the other order-k triangulations, in the sense that for any order-k triangulation T there exists another order-k triangulation T' at distance at most $k-1$ in the flip graph of order-$(2k-2)$ triangulations. It is noteworthy that each flip on the path from T to T' increases the angle vector of the triangulation. The bottom line is that while order-k triangulations are not connected via the flip operation, they become connected if a slightly relaxed condition is considered. For points in generic (non-convex) position, we prove the same result for up to $k \leq 5$, although we conjecture that it holds for all k. Our results have several implications on the flip distance between order-k triangulations, as well as on their efficient algorithmic enumeration.

2.1 Preliminaries and general observations

In this section we give some definitions and observations that will be useful for the rest of the chapter.

Let S be a point set in the plane. Throughout this chapter we assume that set S is in *general position*, i.e., that no three points of S lie on a line and no four points of S lie on a circle. Let T be a triangulation of S, and let $\triangle uyv$ be a triangle in T with vertices u,y,v. We will denote by $\bigcirc uyv$ the open disk defined by the enclosed area of the *circumcircle* of $\triangle uyv$ (i.e., the unique circle through u, y, and v). Thus, $\partial\bigcirc uyv$ denotes the circumcircle of $\triangle uyv$. Triangle $\triangle uyv$ is a triangle *of order* k, also called an *order-k* triangle, if $\bigcirc uyv$ contains at most k points of S. A triangulation where all triangles are order-k is an *order-k (Delaunay) triangulation*. Hence, a triangulation T is not order-k if $\bigcirc uvy$ contains more than k points of S for some $\triangle uyv$ in T. The set of all order-k triangulations of S will be denoted $\mathcal{T}_k(S)$.

Let $e = uv$ be an edge in T. Edge e is *flippable* if e is incident to two triangles $\triangle uxv$ and $\triangle uyv$ of T and $uxvy$ is a convex quadrilateral. The edge e is called *illegal* if $\bigcirc uxv$ contains y. Note that this happens if and only if $\bigcirc uyv$ contains x. See Fig. 2.1.b. Otherwise, the edge uv is called *legal*. It is easy to see that an illegal edge is flippable. The *angle vector* $\alpha(T)$ of a triangulation T is the vector whose components are the angles of each triangle in T ordered in increasing order. Let $T' \neq T$ be another triangulation of S. We say that $\alpha(T) > \alpha(T')$ if $\alpha(T) > \alpha(T')$ in lexicographic order. It is well-known that if T' is the triangulation obtained by flipping an illegal edge of T, then $\alpha(T') > \alpha(T)$ [58]. Moreover, since $DT(S)$ maximizes the minimum angle, it follows that $DT(S)$ is the only triangulation where all the edges are legal [96]. This property leads to a simple algorithm for computing the $DT(S)$: start from

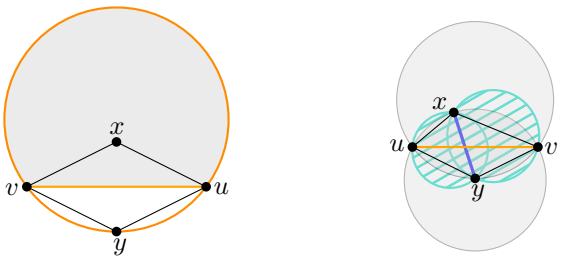

Figure 2.1: (a) An illegal edge uv, with region $\circlearrowleft_y^{uv}$ in gray. (b) The union of the dashed filled disks $\bigcirc uxy$ and $\bigcirc xvy$ is contained in the union of the gray disks $\bigcirc uxv$ and $\bigcirc uyv$.

an arbitrary triangulation of S, and flip illegal edges until none is left. This also implies that the flip graph of all triangulations of S is connected, since any triangulation can be transformed into the $DT(S)$ by a finite number of flips. Let $G(\mathcal{T}_k(S))$ denote the flip graph of $\mathcal{T}_k(S)$.

Next we present several important facts. We start with a well-known observation, illustrated in Fig. 2.1.b.

Observation 2.1.1. Let $\triangle uxv$ and $\triangle uyv$ be two adjacent triangles in a triangulation of S. If edge uv is illegal, then $(\bigcirc uxy \cup \bigcirc xyv) \subset (\bigcirc uxv \cup \bigcirc uvy)$.

In the context of order-k triangulations, Obs. 2.1.1 implies the following.

Lemma 2.1.2 (Abe and Okamoto [1])**.** Let T be a triangulation of S, let uv be an illegal edge of T, and let $\triangle uvx$ and $\triangle uyv$ be the triangles incident to uv in T. If $\triangle uxv$ is of order k, and $\triangle uyv$ is of order l, then triangles $\triangle uxy$ and $\triangle xyv$ are of order k' and l', respectively, for some k', l' with $k' + l' \leq k + l - 2$.

Now we need an extra piece of notation, which we will use extensively. For a triangle $\triangle uyv$, we let $\circledcirc_y^{uv}$ denote the open region bounded by edge uv and the arc of $\partial \bigcirc uyv$ that does not contain y. See Fig 2.1.a.

In what remains of this section we will consider a triangulation T of order $k \geq 3$, and an illegal edge uv adjacent to triangles $\triangle uxv$ and $\triangle uyv$.

Throughout this chapter we will often refer to points of S that are contained—or not—in a certain region. For brevity, we will sometimes omit "of S", and simply refer to points in a certain region, as we do in next observation.

Observation 2.1.3. If $\circledcirc_y^{uv}$ does not contain points, then $(\bigcirc uxy \cap \bigcirc xyv \cap \circledcirc_y^{uv}) \subset (\bigcirc uxv \cap \circledcirc_y^{uv})$ does not contain points.

Proof. Consider the intersection $\partial \bigcirc uxy \cap \partial \bigcirc xyv = \{x, y\}$. Thus, $\bigcirc uxy \cap \bigcirc xyv$ is defined by the area between the arcs from x to y of the circles $\partial \bigcirc uxy$ and $\partial \bigcirc xyv$ that do not contain u and v, respectively. Since y lies in $\bigcirc uxv$ and x lies in $\bigcirc uyv$, the intersection $\bigcirc uxy \cap \bigcirc xyv$ is contained in $\bigcirc uxv \cap \bigcirc uyv$, which does not contain any point in $\circledcirc_y^{uv}$. □

Consider the triangulation T' resulting from flipping uv in T such that $\triangle uxy$ is not of order k. Thus, T' is not of order k. For the sake of simplicity, for any region R in the plane, we denote by $|R|$ the number of points of S in the interior of R.

Using that $|\bigcirc uxv \setminus \{y\}| \leq k-1, |\bigcirc uyv \setminus \{x\}| \leq k-1$ and $|\bigcirc uxy| \geq k+1$, a rather simple counting argument implies the following.

Observation 2.1.4. Each of $\circledcirc_y^{ux} \setminus \bigcirc uxv$ and $\circledcirc_x^{uy} \setminus \bigcirc uyv$ contains at least two points.

Proof. Since $\triangle uyv$ is of order k, $|(\mathbin{\textcircled{}}_y^{ux} \setminus \bigcirc uxv) \cup (\bigcirc uxv \cap \bigcirc uyv)| = |\mathbin{\textcircled{}}_y^{ux} \setminus \bigcirc uxv| + |\bigcirc uxv \cap \bigcirc uyv| \leq k-1$ (1) since x is on $\partial \bigcirc uxv$. Similarly, since $\triangle uxv$ is of order k, $|(\mathbin{\textcircled{}}_x^{uy} \setminus \bigcirc uyv) \cup (\bigcirc uxv \cap \bigcirc uyv)| = |\mathbin{\textcircled{}}_x^{uy} \setminus \bigcirc uyv| + |\bigcirc uxv \cap \bigcirc uyv| \leq k-1$ (2). On the other hand, since $\triangle uxy$ is not of order k, $|\mathbin{\textcircled{}}_y^{ux} \setminus \bigcirc uxv| + |\mathbin{\textcircled{}}_x^{uy} \setminus \bigcirc uyv| + |\bigcirc uxv \cap \bigcirc uyv| \geq k+1$ (3). Combining inequalities (2) and (3) we obtain that $|\mathbin{\textcircled{}}_y^{ux} \setminus \bigcirc uxv| + k - 1 \geq k+1$, thus $|\mathbin{\textcircled{}}_y^{ux} \setminus \bigcirc uxv| \geq 2$. Analogously, combining inequalities (1) and (3), we obtain that $|\mathbin{\textcircled{}}_x^{uy} \setminus \bigcirc uyv| + k - 1 \geq k+1$. Hence, $|\mathbin{\textcircled{}}_x^{uy} \setminus \bigcirc uyv| \geq 2$. □

The next two observations concern the case where the region $\bigcirc uxy \setminus \mathbin{\textcircled{}}_y^{ux}$ contains the maximum possible number of points, i.e., $k-1$. As shown next, in this case the intersection $\mathbin{\textcircled{}}_y^{ux} \cap \bigcirc uxv$ does not contain points of S.

Observation 2.1.5. If $\bigcirc uxy \setminus \mathbin{\textcircled{}}_y^{ux}$ contains $k-1$ points, then $\mathbin{\textcircled{}}_y^{ux} \cap \bigcirc uxv$ does not contain any point.

Proof. Since y is in $\bigcirc uxv$, the arc of $\partial \bigcirc uxy$ from u to v that contains y is contained in $\bigcirc uxv$. Hence, $\bigcirc uxy \setminus \mathbin{\textcircled{}}_y^{ux}$ is contained in $\bigcirc uxv$. Thus, the interior of $\bigcirc uxv \setminus \mathbin{\textcircled{}}_v^{ux}$ contains k points, the $k-1$ points in $\bigcirc uxy \setminus \mathbin{\textcircled{}}_y^{ux}$ and y. Since $\triangle uxv$ is of order k, $\mathbin{\textcircled{}}_y^{ux} \cap \bigcirc uxv$ does not contain any point. □

Let $p_1 \neq y$ in S be such that $\triangle up_1x$ is in T. The next lemma implies that ux is an illegal edge. Furthermore, $\bigcirc up_1x$ cannot contain points in $\mathbin{\textcircled{}}_y^{ux}$.

Lemma 2.1.6. If $\bigcirc uxy \setminus \mathbin{\textcircled{}}_y^{ux}$ contains $k-1$ points, then $\bigcirc up_1x$ contains point y, but does not contain any point of $\mathbin{\textcircled{}}_y^{ux}$.

Proof. From Obs. 2.1.4 there are at least two points in $\mathbin{\textcircled{}}_y^{ux} \setminus \bigcirc uxv$. Since there are points in $\mathbin{\textcircled{}}_y^{ux}$, from Obs. 2.1.5 such points are only in $\mathbin{\textcircled{}}_y^{ux} \setminus \bigcirc uxv$. Let p_1 be the point in $\mathbin{\textcircled{}}_y^{ux}$ such that $\bigcirc up_1x$ does not contain any point of $\mathbin{\textcircled{}}_y^{ux}$. Then, the arc of $\partial \bigcirc up_1x$ from u to x that contains p_1 lies in $\mathbin{\textcircled{}}_y^{ux}$. Thus, the other arc of $\partial \bigcirc up_1x$ from x to u lies outside $\bigcirc uxy$. Therefore, the region $\bigcirc uxy \setminus \mathbin{\textcircled{}}_y^{ux}$ is contained in $\bigcirc up_1x$. Then, $\bigcirc up_1x$ contains k points (the $k-1$ points in $\bigcirc uxy \setminus \mathbin{\textcircled{}}_y^{ux}$ and y). By the same arguments, for any $p' \in \mathbin{\textcircled{}}_y^{ux} \setminus \{p_1\}$, $\bigcirc up'x$ contain at least $k+1$ points (the points in $\bigcirc up_1x$ and p_1). Thus, for any $p' \in \mathbin{\textcircled{}}_y^{ux} \setminus \{p_1\}$, $\triangle up'x$ is not in T. It remains to show that $\triangle up_1x$ is in T. For the sake of a contradiction suppose that $\triangle up_1x \notin T$. Thus, u and x are not adjacent to any point in $\mathbin{\textcircled{}}_y^{ux}$ in T. Then, there exists at least one $p' \in \mathbin{\textcircled{}}_y^{ux}$ such that p' is adjacent to the endpoints of an edge ab crossing twice $\partial \bigcirc uxy$ and is blocking x or u from p' (note that a can be u and b can be x but $ab \neq ux$). Hence, at least one of a or b lies outside $\bigcirc uxy$. Consider the disk $\bigcirc ap'b$. Since ab crosses $\partial \mathbin{\textcircled{}}_y^{ux}$, the arc of $\partial \bigcirc ap'b$ from a to b that contains p' crosses twice the

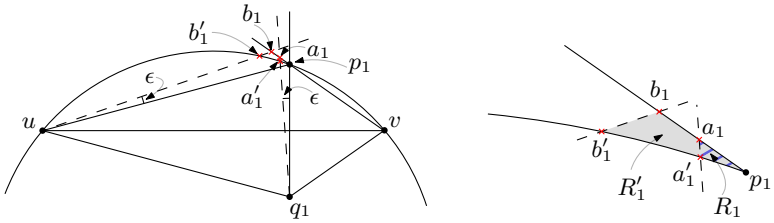

(a) (b)

Figure 2.2: For $k = 4$ we have $\epsilon = 4°$: (a) Shows the intersection points a_1, a_1', b_1 and b_1' given in the proof of Theorem 2.2.1. (b) Shows an example of regions R_1 and R_1' given in the proof of Theorem 2.2.1, represented by the blue raising tilling pattern and the grey filled area, respectively; in this case $R_1 \cap R_1' = R_1$

arc of $\partial \bigcirc uxy$ from u to x that defines $\circledcirc_y^{ux}$. Hence, $\bigcirc uxy \setminus \circledcirc_y^{ux}$ is contained in $\bigcirc ap'b$. Therefore, $\bigcirc ap'b$ contains at least $k+1$ points of S (the $k-1$ points in $\bigcirc uxy \setminus \circledcirc_y^{ux}$, y and either u or x), which is a contradiction. □

2.2 Points in convex position

In this section we show that $k-1$ flips are sufficient to transform any order-k triangulation of a convex point set into some other order-k triangulation, such that all the intermediate triangulations are of order $2k-2$.

Before that, we show that our result is tight in how large the flip distance between two order-k triangulations can be.

Theorem 2.2.1. For any $k > 2$ there is a set S_k of $2k+2$ points in convex position such that $G(\mathcal{T}_k(S_k))$ is not connected. Moreover, there is a triangulation T_k in $\mathcal{T}_k(S_k)$ such that in order to transform T_k into any other triangulation in $\mathcal{T}_k(S_k)$ one needs to perform at least $k-1$ flips.

Proof. In this construction, all angles will be smaller than 180°, so by $\angle xyz$ we denote the smaller of the two angles formed by the rays yx and yz. By angle between two lines we mean the smaller or the two such angles.

We start with a horizontal line segment uv and a point p_1 above it, such that $\angle(vup_1) = 15°$ and $\angle(p_1vu) = 35°$. Let q_1 be the reflection of p_1 with respect to the line through uv. Then, $\angle(up_1v) = \angle(vq_1u) = 130°$. Thus, $\angle(up_1q_1) = \angle(p_1q_1u) = 75°$ and $\angle(vq_1p_1) = \angle(q_1p_1v) = 55°$. Let $\epsilon > 0$ be such that $(205° - 2(k-1)\epsilon°) > 180°$. Let $\bigcirc_1$ be $\bigcirc up_1v$ and ℓ_1 be the line through p_1v. See Fig. 2.3.

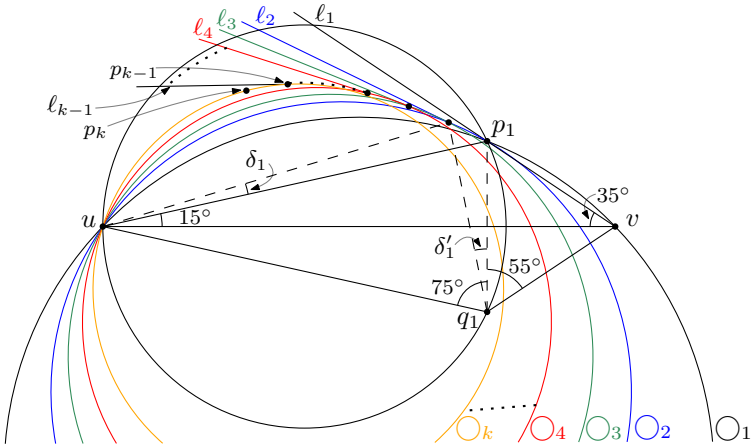

Figure 2.3: Each disk $\bigcirc_i$ contains q_1 and $\circledcirc_{p_1}^{uq_1}$. The colored rays represent the supporting line of the edge $e_i = p_ip_{i-1}$ for all $i > 1$ and respecting $\bigcirc_i$.

Let a_1 be the intersection point of ℓ_1 and the line obtained by rotating the line through q_1p_1 about q_1 counterclockwise by $\epsilon°$. Thus, $\angle a_1q_1v = 55° + \epsilon°$. Let a_1' be the point defined as $q_1a_1 \cap (\partial\bigcirc_1)$. Let b_1 be the point defined as the intersecting of ℓ_1 and the line obtained by rotating counterclockwise $\epsilon°$ about u the line through up_1. Hence, $\angle q_1ub_1 = 35° + \epsilon°$. Let b_1' be the point defined as $ub_1 \cap (\partial \bigcirc_1 \setminus u)$. Notice that by construction, the points a_1 and b_1 do not lie in $\bigcirc_1$. Moreover, there is a 2-dimensional region R_1 outside $\bigcirc_1$ defined by the edges a_1a_1', p_1a_1 and the arc of $\partial\bigcirc_1$ with endpoints p_1 and a_1', that does not contain u. Similarly, there is a 2-dimensional region R_1' outside $\bigcirc_1$ defined by the edges b_1b_1', p_1b_1 and the arc of $\partial\bigcirc_1$ with endpoints p_1 and b_1', that does not contain u. See Fig. 2.2.a. Notice, that since a_1 and b_1 are outside $\bigcirc_1$, then $R_1 \cap R_1'$ defines a non-empty 2-dimensional region. Let p_2 be a point in the interior of $R_1 \cap R_1'$. See Fig. 2.2.b. Let δ_1 be the angle $\angle(p_1up_2)$ and let δ_1' be the angle $\angle(p_1q_1p_2)$. By definition of p_2, $\delta_1 < \epsilon$ and $\delta_1' < \epsilon$. Then,
$\angle(up_2p_1) = \angle(up_2q_1) + \angle(q_1p_2p_1) \geq \angle(ub_1p_1) + \angle(q_1a_1v) \geq \angle(up_1q_1) + \delta_1' - \delta_1 + \angle(q_1p_1v) - \delta_1' = 130 - \delta_1$.
Thus, $\angle(up_2p_1)+\angle(p_1q_1u) > \angle(up_1v)-\delta_1+\angle(p_1q_1u)-\delta_1' = 130-\delta_1+75-\delta_1' \geq 205 - 2\epsilon > 180$. Therefore, $\bigcirc_2 = \bigcirc up_2p_1$ contains q_1. Moreover, $\bigcirc_2$ contains $\circledcirc_{p_1}^{uq_1}$.

Recursively, for $i \in \{1, \ldots, k-1\}$ we find a point p_{i+1} with the following properties: (1)p_{i+1} lies in R_i and R_i' between the supporting line ℓ_i of edge p_ip_{i-1} and outside $\bigcirc_i = \bigcirc_{up_ip_{i-1}}$, where $p_0 = v$. (2) $\angle(up_{i+1}q_1)+\angle(p_{i+1}q_1u) > 130-\delta_i+75-\delta_i'-2(i-1)\epsilon \geq 205-2i\epsilon > 180$. For all $i > 2$, note that by construction p_j does not lie in $\bigcirc_i$ for all $j < i-1$. Also, note that $\circledcirc_{p_1}^{uq_1}$ is contained in $\bigcirc_i$.

See Fig. 2.3. Symmetrically, we define a point set $\{q_1, q_2, \ldots, q_k\}$, such that q_i is the reflection of p_i with respect to the horizontal line passing through u and v.

Let $S' := \{p_1, p_2, \ldots, p_k\}$ and $S'' := \{q_1, q_2, \ldots, q_k\}$. We define $S_k := \{u, v\} \cup S' \cup S''$. Consider the triangulation $T = \{\triangle up_1v, \triangle uq_1v\} \cup \{\triangle up_{i+1}p_i, \triangle uq_{i+1}q_i : \forall i \in \{1, \ldots, k-1\}\}$. See Fig. 2.4. By construction of S_k, each of the triangles is of order k, which is, each triangle formed by points in $S \setminus S'$ contains exactly the k points of S' in the interior of its circumcircle and symmetrically each triangle formed by points in $S \setminus S''$ contains in the interior of its circumcircle exactly the k points of S''. Moreover, the triangulations $T \setminus S'$ and $T \setminus S''$ are the Delaunay triangulations of the point sets $S \setminus S'$ and $S \setminus S''$, respectively.

Let us show that for any order-k triangulation T' different from T, the difference between T and T' consists of a set of edges with one endpoint in S' and the other endpoint in S''. In other words, we will show that any edge p_ip_j or q_iq_j with $i - 1 > j$ is in a triangle that is not order k. Note, that if this is true, then flipping any edge of $T \setminus S'$ or $T \setminus S''$ will result in a triangulation that is not of order k. In addition, the only edge of T that can be transformed by one flip flipped to an edge with one endpoint in S' and the other in S'' is uv. By construction, the triangle $\triangle up_1q_1$ is of order $2k - 2$. This means that flipping any edge of T results in a triangulation that is not order-k. Therefore $G(\mathcal{T}_k(S_k))$ is not connected.

Let us show that for $k - 1 \geq i - 1 > j$, any triangulation of S_k containing edge p_ip_j is not order-k, the other case for q_iq_j is symmetric. To show this, we consider triangle $\triangle p_ip_tp_j$ for any $k \geq i > t > j \geq 0$. Note that there exists a t with $i > t > j$, since p_ip_j is a diagonal of the convex hull of S_k where one of the two triangles containing p_ip_j has a vertex p_t with $i > t > j$.

By construction, points p_i and p_j are not in $\bigcirc up_tp_{t-1}$. Then the arc A of $\partial \bigcirc p_ip_tp_j$ with endpoints p_i and p_j that contains p_t crosses $\partial \bigcirc up_tp_{t-1}$ twice, that is, once in p_t and one more time in one of the two arcs: (1) the arc with endpoints p_i and p_t that does not contain p_j, or (2) the arc with endpoints p_t and p_j that does not contain p_i. Arc A has a portion that lies in $\bigcirc up_tp_{t-1} \cap \mathbin{\rotatebox{180}{$\bigcirc$}}_{q_1}^{up_1}$, since the points p_i, p_t, p_{t-1}, p_j lie in $\mathbin{\rotatebox{180}{$\bigcirc$}}_{q_1}^{up_1}$ In addition, since p_i and p_j do not lie in $\bigcirc up_tp_{t-1}$, $\bigcirc p_ip_tp_j$ contains all the points in $S'' \cup \{u\}$ whose total number is $k + 1$. Thus, flipping any edge of T (including uv as mentioned before) results in a triangulation that is not order-k. Therefore, two things have been showed: (1) the flip graph $G(\mathcal{T}_k(S_k))$ is not connected, and (2) any order-k triangulation differs from T by only edges with one endpoint in S' and the other in S''.

It remains to show that $k-1$ flips are necessary to transform T into another order-k triangulation T'. We have shown above, that each edge of T' that is not in T must have one endpoint in S' and one in S''. Thus, edge uv has to be flipped at some point. Note that if there is an edge p_tq_r, then there is a

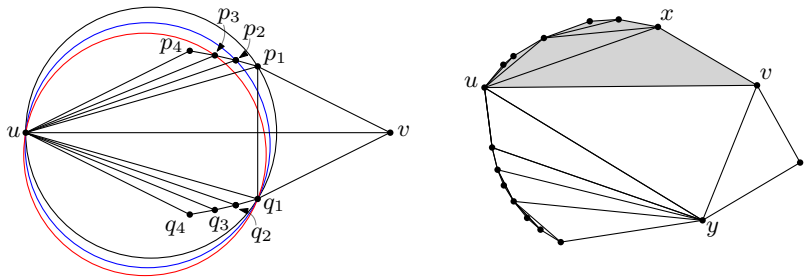

(a) (b)

Figure 2.4: (a) An order-k triangulation at distance at least $k-1$ from any other order-k triangulation (k=4). (b) The gray area corresponds to T^x_{uv}.

triangle $\triangle up_iq_j$ with $i \geq t$ and $j \geq r$: since u is not adjacent to v and u and it is adjacent to p_k and q_k, the minimum i and minimum j such that u is adjacent to p_i and q_j creates a triangle $\triangle up_iq_j$. For $k-1 \geq i \geq 1$ and $k-1 \geq j \geq 1$, let us show that $\bigcirc up_iq_j$ contains all points p_t with $t > i$ and all points q_r with $r > j$, i.e., triangle $\triangle up_iq_j$ is not of order $2k-i-j-1$. This would end the proof, since any edge p_iq_j with $2k-i-j \leq k$ crosses at least $k-1$ edges from T, i.e., edge p_iq_j crosses the edges uv, up_t and uq_r for all $t < i$ and $r < j$. This would conclude that $k-1$ flips are needed in order to transform T into another triangulation of order k.

Notice that disk $\partial \bigcirc up_iq_j$ crosses disk $\partial \bigcirc up_ip_{i-1}$ twice in u and in p_i. See Fig. 2.4.a. Since q_j is in $\bigcirc up_ip_{i-1}$, the arc of $\partial \bigcirc up_iq_j$ between u and p_i that contains q_j lies in $\bigcirc up_ip_{i-1}$. In addition, since the point q_j lies in $\bigcirc up_tp_{t-1}$ for all $t > i$, the arc from u to q_j of $\partial \bigcirc up_iq_j$ that does not contain p_i lies in $\bigcirc up_tp_{t-1}$ for all $t > i$. Also, since p_i is not in $\bigcirc up_tp_{t-1}$, it follows that the arc between u and p_i that contains q_j crosses $\partial \bigcirc up_tp_{t-1}$ twice, and a portion of this arc lies in $\mathbin{\rotatebox{180}{$\bigcirc$}}^{uq_{t-1}}_{p_t}$. Thus, the arc of $\partial \bigcirc up_tp_{t-1}$ between u and p_t that does not contain p_{t-1} lies in $\bigcirc up_iq_j$. Hence, p_t lies in$\bigcirc up_iq_j$. Similarly, q_r is in $\bigcirc up_iq_j$ for all $r > j$. Therefore, since for any order-k triangulation different from T differs by only edges with one endpoint in S' and the other in S'', in order to get some triangle $\triangle up_iq_j$ of order k, we have to flip $k-1$: the edge uv and $k-2$ edges of the form up_t and uq_r with $2k-i-j \leq k, t < i$ and $r < j$. □

Let S be a point set in convex position. Let T be an order-k triangulation of S. We say that T is *minimal* if flipping any illegal edge in T results in a triangulation that is not order-k. Let uv be a diagonal in T and let $\triangle uxv$ and $\triangle uyv$ be the triangles incident to it. Since S is in convex position, the diagonal uv partitions T into two sub-triangulations that only share edge uv. Let T^x_{uv} (respectively, T^y_{uv}) denote the sub-triangulation that contains triangle

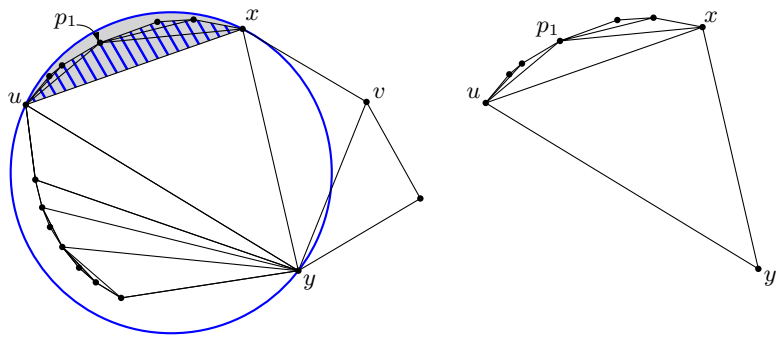

(a) (b)

Figure 2.5: (a) The $\bigcirc uxy$ contains 12 points, i.e, $\triangle$ is not order 11. But the rest of the triangles are of order 8. The gray area is $\textcircled{}_y^{ux}$ and the dashed area corresponds to $T_{ux}^{p_1}$, which in this case is contained in $\textcircled{}_y^{ux}$. (b) Triangulation $T_{ux}^{p_1} \cup \{\triangle uxy\}$ obtained from the one in (a).

$\triangle uxv$ (respectively, $\triangle uyv$). See Fig. 2.4b.

Recall that $|\textcircled{}_y^{ux}|$ referes to the number of points of S in the interior of $\textcircled{}_y^{ux}$. The following lemma will be important for showing the main theorem of this section.

Lemma 2.2.2. Let T be a triangulation with the following properties:

1. Exactly one triangle $\triangle uxy$ in T is not of order $k > 2$,
2. $\textcircled{}_y^{ux}$ contains at least two points,
3. $T_{ux}^{p_1}$ consists only of legal edges, where $\triangle up_1x$ is the triangle in T adjacent to $\triangle uxy$ through edge ux. See Fig 2.5a.

Then, there exists a triangulation T' in $\mathcal{T}_k(S)$ which can be reached from T by flipping at most $|\textcircled{}_y^{ux}| - 1$ illegal edges of $T_{ux}^{p_1}$. Moreover, each intermediate triangulation is of order $|\textcircled{}_y^{ux}| + k - 1$.

Proof. Consider the triangulation $T_{ux}^{p_1} \cup \{\triangle uxy\}$, which is a sub-triangulation of T. See Fig. 2.5b. Since $\textcircled{}_y^{ux}$ contains at least two points in $T_{ux}^{p_1}$, it follows that ux is an illegal edge in $T_{ux}^{p_1} \cup \{\triangle uxy\}$. All other edges in $T_{ux}^{p_1} \cup \{\triangle uxy\}$ are legal since $T_{ux}^{p_1}$ consists of legal edges and uy, xy are convex hull edges by construction.

Let us prove the lemma by induction on $|\textcircled{}_y^{ux}|$.

Consider the base case: $|\textcircled{}_y^{ux}| = 2$. Since ux is an illegal edge, $p_1 \in \textcircled{}_y^{ux}$. Consider the triangulation $T_1 = (T \setminus \{ux\}) \cup \{p_1y\}$. Since ux is an illegal edge, by Obs. 2.1.1 $(\bigcirc up_1y \cup \bigcirc p_1xy) \subset (\bigcirc uxy \cup \bigcirc up_1x)$. Since $T_{ux}^{p_1}$ consists of legal edges, $T_{ux}^{p_1}$ is a Delaunay triangulation. Thus, $\bigcirc up_1x$ does not contain any

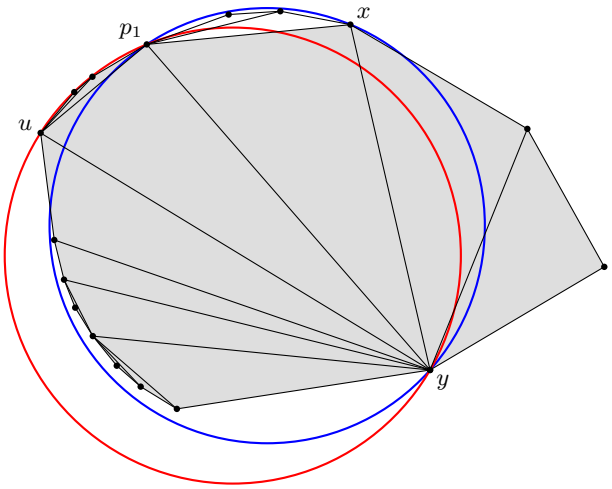

Figure 2.6: The resulted triangulation after flipping ux in T from Fig 2.5.a. Both new triangles $\triangle up_1y$ and $\triangle p_1xy$ are not of order 8. The sub-triangulation in the gray area contains exactly one triangle, $\triangle up_1y$, that is not of order 8.

point from $T^{p_1}_{ux}$. Hence, $\bigcirc up_1y$ contains at most one point in $\circledcirc^{up_1}_y \subset \circledcirc^{ux}_y$ and at most $k-1$ points in $\bigcirc up_1y \setminus \circledcirc^{up_1}_y$ (the points from $\bigcirc up_1x \setminus \{y\}$). Similarly, for $\bigcirc p_1xy$. Thus, $\triangle up_1y$ and $\triangle p_1xy$ are of order k.

Inductive hypothesis: Assume that lemma holds when $|\circledcirc^{ux}_y| \leq m-1$ for $m \geq 3$.

Let $|\circledcirc^{ux}_y| = m$. Since ux is an illegal edge, $p_1 \in \circledcirc^{ux}_y$. Consider the triangulation $T_1 = (T \setminus \{ux\}) \cup \{p_1y\}$. Then, $(\bigcirc up_1y \cup \bigcirc p_1xy) \subset (\bigcirc uxy \cup \bigcirc up_1x)$. In addition, since $T^{p_1}_{ux}$ consists of legal edges, $\bigcirc up_1x$ does not contain any point from $T^{p_1}_{ux}$. Hence, $\bigcirc up_1y$ contains at most $m-1$ points in $\circledcirc^{up_1}_y \subset \circledcirc^{ux}_y$ (the m points in $\circledcirc^{ux}_y$ minus p_1) and at most $k-1$ points in $\bigcirc up_1y \setminus \circledcirc^{up_1}_y$ (the points from $\bigcirc up_1x \setminus \{y\}$). A similar argument holds for $\bigcirc p_1xy$. On the other hand, by Obs. 2.1.3 the disks $\bigcirc up_1y$ and $\bigcirc p_1xy$ do not share points in $\circledcirc^{ux}_y$. Hence, $|\circledcirc^{up_1}_y| \leq |\circledcirc^{ux}_y| - |\circledcirc^{p_1x}_y| - 1 \leq m-1$ and $|\circledcirc^{p_1x}_y| \leq |\circledcirc^{ux}_y| - |\circledcirc^{up_1}_y| - 1 \leq m-1$. Thus, the order of triangles $\triangle up_1y$ and $\triangle p_1xy$ is $|\circledcirc^{ux}_y| + k - 2$. If the new triangles $\triangle up_1y$ and $\triangle p_1xy$ are of order k, then the lemma follows. If exactly one of $\triangle up_1y$ and $\triangle p_1xy$ is not of order k, then by the inductive hypothesis the lemma follows. Assume that both $\triangle up_1y$ and $\triangle p_1xy$ are not of order k. See Fig. 2.6.

Since $\bigcirc up_1y \setminus \circledcirc^{up_1}_y$ contains at most $k-1$ points, $|\circledcirc^{up_1}_y| \geq 2$. Similarly, $|\circledcirc^{p_1x}_y| \geq 2$. Consider p_2 and p_3 in S such that $\triangle up_2p_1$ is adjacent to $\triangle up_1y$, and $\triangle p_1p_3x$ is adjacent to $\triangle p_1xy$. Then, the points in $\circledcirc^{up_1}_y$ are vertices of $T^{p_2}_{up_1}$

and the points in $\bigcirc_y^{p_1x}$ are vertices of $T_{p_1x}^{p_3}$.

Consider the triangulation $T_2 = (T_1 \setminus T_{p_1x}^{p_3}) \cup \{p_1x\}$, which consists of the triangulation T_1 of S minus the points in $T_{p_1x}^{p_3} \setminus \{p_1, x\}$. See the triangulation of the gray area in Fig. 2.6. Note that $\triangle up_1y$ is the only triangle that is not of order k in T_2. Indeed, since the only points in $\bigcirc p_1xy$ in T_2 are in $\bigcirc p_1xy \setminus \bigcirc_y^{p_1x}$, the triangle $\triangle p_1xy$ is of order $k-1$ in T_2. Thus, $\triangle up_1x$ is the only triangle that is not of order k in T_2. Since $T_{up_1}^{p_2} \subset T_{ux}^{p_1}$, $T_{up_1}^{p_2}$ consists of legal edges. Thus, the inductive hypothesis holds for T_2. That is, T_2 is at distance at most $|\bigcirc_y^{up_1}| - 1$ from some triangulation T_2' of order k such that the only edges flipped are in $T_{up_1}^{p_2}$. Moreover, the triangulations between T_2 and T_2' are of order $|\bigcirc_y^{up_1}| + k - 1$. Consider the triangulation $T_3 = (T_1 \setminus T_2) \cup T_2'$. Since T_3 consists of edges in T_1 and at most $|\bigcirc_y^{up_1}| - 1$ flipped edges from $T_{up_1}^{p_2}$, the triangulation T_3 is at distance at most $|\bigcirc_y^{up_1}| - 1$ from T_1 by flipping only illegal edges in $T_{up_1}^{p_2}$. In addition, the triangles between T_1 and T_3 are of order $\max\{|\bigcirc_y^{up_1}|, |\bigcirc_y^{p_1x}|\} + k - 1 \leq |\bigcirc_y^{ux}| + k - 1$. Note that $\triangle p_1xy \in T_3$. By construction, $\triangle p_1xy$ is the only triangle that is not of order k in T_3 since T_2' is of order k and the only that are not order-k in T_1 were $\triangle up_1y$ and $\triangle p_1xy$. Thus, T_3 is of order $|\bigcirc_y^{p_1x}| + k - 1$. Finally, since $T_{p_1x}^{p_3} \subset T_{ux}^{p_1}$, $T_{p_1x}^{p_3}$ consists of legal edges. Thus, by inductive hypothesis T_3 is at distance at most $|\bigcirc_y^{p_1x}| - 1$ from a triangulation T_3' of order k by flipping only illegal edges in $T_{p_1x}^{p_3}$. Also, the triangulations between T_3 and T_3' are of order $|\bigcirc_y^{p_1x}| + k - 1$. Thus, T_1 is at distance at most $|\bigcirc_y^{up_1}| + |\bigcirc_y^{p_1x}| - 2 \leq |\bigcirc_y^{ux}| - 2 = m - 2$ to a k-$DT(S)$ by flipping only illegal edges in $(T_{up_1}^{p_2} \cup T_{p_1x}^{p_3}) \subset T_{ux}^{p_1}$. Therefore, there exists a triangulation T' of order k which resulted after flipping at most $|\bigcirc_y^{ux}| - 1$ illegal edges from $T_{ux}^{p_1}$ in T. Moreover, the triangulations between T and T' are of order $|\bigcirc_y^{ux}| + k - 1$. □

Using Lemma 2.2.2 we now proceed to prove the upper bound.

Theorem 2.2.3. For a point set S in convex position and $k \geq 2$, let $T \neq DT(S)$ be a triangulation in $\mathcal{T}_k(S)$. Then, there exists T' in $\mathcal{T}_k(S)$ such that there is a path between T and T' in $G(\mathcal{T}_{2k-2}(S))$ of length at most $k-1$, where each edge of the path corresponds to flipping an illegal edge.

Proof. Note that for $k = 2$, follow trivially since $G(\mathcal{T}_2(S))$ is connected. Assume $k \geq 3$. If T is not minimal, then T contains an illegal edge e such that flipping e results in an order-k triangulation. Thus, we assume that triangulation T is minimal. Observe first that there must be an illegal edge uv in T incident to triangle $\triangle uxv$ such that all edges of T_{uv}^{x} are legal. Indeed, since T is not an order-0 triangulation, T contains an illegal edge. Any triangulation of S has at least two ears, i. e., triangles with two edges in the convex hull of S, and for any ear in T all three of its edges are legal, otherwise T is not minimal.

Let uv be an illegal edge incident to triangle $\triangle uxv$ such that all edges in T^x_{uv} are legal. Let $\triangle uyv$ be the other triangle in T incident to uv. Consider triangulation $T_1 = (T \setminus \{uv\}) \cup \{xy\}$. Since T is minimal, T_1 is not order-k. The only triangles in T_1 that could be not order-k are the new triangles $\triangle uxy$ and $\triangle xyv$. Without loss of generality assume that $\triangle uxy$ is not order-k. By Lemma 2.1.2, it follows that $\triangle uxy$ is the only one that is not order-k. In addition, $\triangle uxy$ is order-$(2k-2)$. By Obs. 2.1.4 it follows that $2 \leq |\mathcal{O}^{ux}_y| \leq k-1$.

By Lemma 2.2.2 it follows that T_1 can be transformed into an order-k triangulation T' by flipping at most $k-2$ illegal edges. Moreover, in the sequence of triangulations from T_1 to T' every triangulation is of order $2k-2$. The claim follows. □

Observe that due to the known property of illegal edges, each flip in the path from T to T' increases the angle vector of the triangulation. In addition, if T is an order-k triangulation, then the edges of T ar of order k. Independently Abellanas et al. [2] and Gudmundsson et al. [62] showed that there are $O(kn)$ order-k edges. It follows from Theorem 2.2.3 that T can be transformed into $DT(S)$ by a sequence of at most $O((2k-2)n) = O(kn)$ flips, since all the flipped edges are illegal and of order at most $2k-2$, which implies that no order-$(2k-2)$ edge is flipped twice.

Corollary 2.2.4. Let S be a point set in convex position. Any triangulation in $T_k(S)$ can be transformed into $DT(S)$ by a sequence of at most $O(kn)$ triangulations of order at most $2k-2$.

Proof. Let $T \in \mathcal{T}_k$. Let us show that there is a sequence of triangulations $T_0 = T, T_1, \ldots, T_m = DT(S)$. By Sibson [96], T can be transformed into $DT(S)$ by a finite number of illegal flips. For $i \geq 0$, if T_i is not minimal then there exists an illegal edge u_iv_i in two triangles $\triangle u_ix_iv_i$ and $\triangle u_iy_iv_i$ such that $T_{i+1} = T_i \setminus \{u_iv_i\} \cup \{x_iy_i\}$ is an order-k triangulation. If T_i is minimal, then flipping any illegal edge in T_i will result into a triangulation that is not order-k. It follows from Theorem 2.3.1 that there exists an order-k triangulation T_j such that T_i can be transformed into T_j by flipping at most $k-1$ only illegal edges and the intermediate triangulations are order-$(2k-2)$. So, add the sequence of triangulations between T_i and T_j: $T_i, T_i+1, \ldots, T_j$ where $j < i+k$. Since all of the flipped edges (including the ones of order-$(2k-2)$) are illegal, the sequence of order-$(2k-2)$ triangulations is strictly increasing with respect to the angle vector and $DT(S)$ is the only maximum element, at most all of the order-$(2k-2)$ edges are flipped in order to reach $DT(S)$. By Abellanas et al. [2] and Gudmundsson et al. [62], there are $O((2k-2)n) = O(kn)$ flipped edges. □

Theorem 2.2.3 also implies an enumeration algorithm for all order-k triangulations. However, using the preprocessing method given by Silveira and van

Kreveld [97] where all order-k triangles and order-k edges in all order-k triangulation of a point set in general position can be computed in $O(k^2 n \log k + kn \log n)$ expected time, implies a faster algorithm.

Lemma 2.2.5. For any point set S in convex position, there exists an algorithm that enumerates all order-k triangulations of S in polynomial time per triangulation.

The lemma follows since we can store all the order-k possible triangles in order-k triangulations given by Silveira and van Kreveld [97] in a hash table. Then, for each order-k triangle we subdivide the problem into three subsets of S and count all possible triangulations of each subset without repeating triangles already visited.

2.3 General point sets

In this section we consider a general point set S. We show that a triangulation of order $k = 3, 4$ or 5 of S can be transformed into some other order-k triangulation of S by flipping at most $k-1$ illegal edges, and that the intermediate triangulations are of order $2k-2$. Moreover, since we flip only illegal edges, after each flip that transforms a triangulation T to T', we have that $\alpha(T') > \alpha(T)$. Thus, if we keep applying this procedure, we will eventually reach $DT(S)$.

Theorem 2.3.1. Let S be a point set in general position and let T be a triangulation in $\mathcal{T}_k(S)$ for $2 \leq k \leq 5$. There exists T' in $\mathcal{T}_k(S)$ such that there is a path from T to T' in $G(\mathcal{T}_{2k-2}(S))$ of length at most $k-1$, where each edge of the path corresponds to flipping an illegal edge.

In order to prove Theorem 2.3.1, we consider whether T is minimal. If not, the statement follows trivially. If T is minimal, then $k > 2$. Also, for any illegal edge uv in T, flipping uv produces a new and unique triangle $\triangle uxy$ that is not order-k. Next, we show that the statement holds for such two cases and with

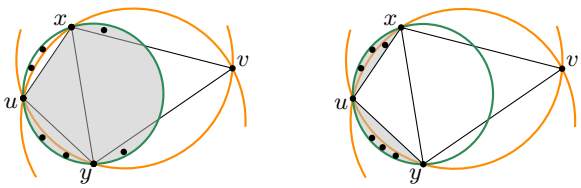

Figure 2.7: In both cases, $k = 5$. (a) There are four points in the gray region $\bigcirc uxy \setminus \circledcirc_y^{ux}$. (b) There are exactly three points in $\circledcirc_y^{ux}$ and exactly three points in $\circledcirc_x^{uy}$.

that we conclude with the theorem.

Lemma 2.3.2. For $k = 3, 4, 5$, the following are the only two possible cases: (a) There are exactly $k-1$ points in $\bigcirc uxy \setminus \circledcirc_y^{ux}$ or exactly $k-1$ points in $\bigcirc uxy \setminus \circledcirc_x^{uy}$. See Fig 2.7.a. (b) There are exactly $k-2$ points in $\circledcirc_y^{ux}$ and exactly $k-2$ points in $\circledcirc_x^{uy}$.

Proof. By Obs. 2.1.1, $\bigcirc uxy$ has at most $2k-2$ points of S. First, we consider $k = 3, 4$ and show that in order for the $\bigcirc uxy$ to have more than k points, then only case (a) can happen. Recall that since $\bigcirc uxy$ contains more than k points, by Obs. 2.1.4 both regions $\circledcirc_y^{ux}$ and $\circledcirc_x^{uy}$ have at least 2 points. In addition, each of $\bigcirc uxy \setminus \circledcirc_y^{ux} \subset \bigcirc uxv$ and $\bigcirc uxy \setminus \circledcirc_x^{uy} \subset \bigcirc uyv$ can contain at most $k-1$ points.

Consider $k = 3$. Since by Obs. 2.1.1 $\bigcirc uxy$ has at most $2k-2 = 2(3)-2 = 4$ points, it follows from Obs. 2.1.4 that there are exactly 2 points in $\circledcirc_y^{ux}$ and exactly 2 in $\circledcirc_x^{uy}$. Thus, both regions $\bigcirc uxy \setminus \circledcirc_y^{ux}$ and $\bigcirc uxy \setminus \circledcirc_x^{uy}$ have exactly $k-1 = 3-1 = 2$ points.

Now, consider $k = 4$. Since there are at least 2 points in each of $\circledcirc_y^{ux}$ and $\circledcirc_x^{uy}$, and $\bigcirc uxy$ contains at least $k+1 = 5$ points, it follows that at least one of the regions $\bigcirc uxy \setminus \circledcirc_y^{ux}$ and $\bigcirc uxy \setminus \circledcirc_x^{uy}$ contains 3 (thus $k-1$) points.

Therefore, we can conclude that for $k = 3, 4$ in order for $\bigcirc uxy$ to contain more that k points, there most be exactly $k-1$ points in $\bigcirc uxy \setminus \circledcirc_y^{ux}$ or exactly $k-1$ points in $\bigcirc uxy \setminus \circledcirc_x^{uy}$.

When $k = 5$, there are two cases to consider. The first one is when there are exactly $k-1$ points in $\bigcirc uxy \setminus \circledcirc_y^{ux}$ or in $\bigcirc uxy \setminus \circledcirc_x^{uy}$. This case can happen, since $\bigcirc uxy \setminus \circledcirc_y^{ux}$ is contained in $\bigcirc uxv$, which can have up to $k-1$ points. Similarly, $\bigcirc uxy \setminus \circledcirc_x^{uy}$ is contained in $\bigcirc uyv$, which can have up to $k-1$ points different from x.

The only other case for $k = 5$ is when there are $k-2 = 3$ points in $\circledcirc_y^{ux}$ and $k-2 = 3$ points in $\circledcirc_x^{uy}$. To observe this, note that if there are fewer than $k-1$ points in $\bigcirc uxy \setminus \circledcirc_y^{ux}$ and $\bigcirc uxy \setminus \circledcirc_x^{uy}$, then there are at most $k-2 = 3$ points in both regions $\bigcirc uxy \setminus \circledcirc_y^{ux}$ or $\bigcirc uxy \setminus \circledcirc_x^{uy}$. On the other hand, since $k = 5$ and $\bigcirc uxy$ contains at least $k+1 = 6$ points, $|\bigcirc uxy \setminus \circledcirc_y^{ux}| + |\bigcirc uxy \setminus \circledcirc_x^{uy}| \geq 6$. Thus, both $\bigcirc uxy \setminus \circledcirc_y^{ux}$ and $\bigcirc uxy \setminus \circledcirc_x^{uy}$ contain at least $k-2 = 3$ points, otherwise one of $\bigcirc uxy \setminus \circledcirc_y^{ux}$ and $\bigcirc uxy \setminus \circledcirc_x^{uy}$ contains at least $4 = k-1$ points, falling into the previous case. It remains to show that the 3 points in $\bigcirc uxy \setminus \circledcirc_y^{ux}$ and $\bigcirc uxy \setminus \circledcirc_x^{uy}$ are actually in $\circledcirc_x^{uy}$ and $\circledcirc_x^{uy}$, respectively. Note that if $\circledcirc_u^{xy}$ does not contain points, then it follows that there are 3 points in both $\circledcirc_y^{ux}$ and $\circledcirc_x^{uy}$. For the sake of a contradiction suppose that there are points in $\circledcirc_u^{xy}$. Since $\bigcirc uxy \setminus \circledcirc_y^{ux}$ contains exactly 3 points and $\circledcirc_x^{uy}$ contains at

least 2 points, it follows that $\bigcirc_u^{xy}$ contains exactly one point and $\bigcirc_x^{uy}$ contains exactly 2 points. Similarly, since $\bigcirc uxy \setminus \bigcirc_x^{uy}$ contains exactly 3 points and $\bigcirc_y^{ux}$ contains at least 2 points, it follows that $\bigcirc_u^{xy}$ contains exactly 1 point and $\bigcirc_y^{ux}$ contains exactly 2 points. Thus, $\bigcirc uxy$ has $|\bigcirc_y^{ux}|+|\bigcirc_x^{uy}|+|\bigcirc_u^{xy}| = 2+1+2 = 5$ points, which contradicts the assumption that $\bigcirc uxy$ contains more than $k = 5$ points. □

Therefore, in order to show that with at most $k-1$ flips the triangulation T can be transformed into some other triangulation T' of order k such that only illegal edges are flipped, for $k = 3, 4, 5$, we consider two cases:

A) [$k = 3, 4, 5$] There are exactly $k-1$ points in $\bigcirc uxy \setminus \bigcirc_y^{ux}$ or exactly $k-1$ points in $\bigcirc uxy \setminus \bigcirc_x^{uy}$. See Fig 2.7.a.

B) [$k = 5$] There are exactly $k-2 = 3$ points in $\bigcirc_y^{ux}$ and exactly $k-2 = 3$ points in $\bigcirc_x^{uy}$. See Fig 2.7.b.

Case A

For $k = 3, 4, 5$; there are exactly $k-1$ points in $\bigcirc uxy \setminus \bigcirc_y^{ux}$ or exactly $k-1$ points in $\bigcirc uxy \setminus \bigcirc_x^{uy}$. Since the case when there are exactly $k-1$ points in $\bigcirc uxy \setminus \bigcirc_y^{ux}$ is symmetric to the one in $\bigcirc uxy \setminus \bigcirc_x^{uy}$, without loss of generality we assume that there are exactly $k-1$ points in $\bigcirc uxy \setminus \bigcirc_y^{ux}$.

We show that T can be transformed into a different order-k triangulation T' by a sequence of of at most $k-1$ flips of only illegal edges. Moreover, the intermediate triangulations are of order $2k-2$.

Lemma 2.3.3. Let S be a point set in general position and let T be a minimal triangulation in $\mathcal{T}_k(S)$, for $3 \leq k \leq 5$. Let uv be an illegal edge of T with adjacent triangles $\triangle uxv$ and $\triangle uyv$. Consider the triangulation $T_1 = (T \setminus \{uv\}) \cup \{ux\}$ such that:

1. $\triangle uxy$ is not of order k,
2. there are $k-1$ points in $\bigcirc uxy \setminus \bigcirc_y^{ux}$.

Then, there exists a triangulation T' in $\mathcal{T}_k(S)$ which can be reach from T by flipping at most $k-1$ only illegal edges. Moreover, the intermediate triangulations are of order $2k-2$.

Proof. Recall by Lemma 2.1.2 that $\bigcirc uxy$ contains at most $2k-2$ points. Thus, $T_1 \in \mathcal{T}_{2k-2}(S)$.

From Lemma 2.1.6 there exists a point $p_1 \in \bigcirc_y^{ux} \setminus \bigcirc uxv$ such that $\triangle up_1x$ is in T and $\bigcirc up_1x$ does not contain a point from $\bigcirc_y^{ux}$. Then, the edge ux is an illegal edge. Moreover, since the arc from u to x in $\partial \bigcirc up_1x$ that contains p_1 is contained in $\bigcirc_y^{ux} \subset \bigcirc uxy$, $\bigcirc up_1x$ contains all of the points in $\bigcirc uxy \setminus \bigcirc_y^{ux}$

including y. Thus, $\bigcirc up_1x$ contains exactly k points, which are the points in $\bigcirc uxy \setminus \circlearrowleft_y^{ux}$ including y.

Consider the triangulation T_2 obtained from T_1 after flipping ux: $T_2 = (T_1 \setminus \{ux\}) \cup \{p_1y\} = (T \setminus \{uv, ux\}) \cup \{xy, p_1y\}$. Since ux is an illegal edge, by Obs. 2.1.1 $(\bigcirc up_1y \cup \bigcirc p_1xy) \subset (\bigcirc uxy \cup \bigcirc up_1x)$. Hence, $\bigcirc up_1y$ and $\bigcirc p_1xy$ contain at most $2k-3$ points (the at most $2k-2$ points in $\bigcirc uxy \cup \bigcirc up_1x$ minus p_1). Thus, $T_2 \in \mathcal{T}_{2k-3}(S) \subset \mathcal{T}_{2k-2}(S)$. Note that for $k = 3$ the lemma already follows, since $2k-3 = 6-3 = 3$. For the remaining of this proof, we assume that $k = 4, 5$. If $T_2 \in \mathcal{T}_k(S)$, then the lemma follows. Assume that T_2 is not of order k.

Thus, there is one triangle $\blacktriangle$ (either $\triangle up_1y$ or $\triangle p_1xy$) that is not order-k. Since $p_1 \in \circlearrowleft_y^{ux}$, the regions $\circlearrowleft_y^{up_1}$ and $\circlearrowleft_y^{p_1x}$ are in $\circlearrowleft_y^{ux}$. Therefore, by Obs. 2.1.4 the interior of circumcircle $\bigcirc\blacktriangle$ of $\blacktriangle$ contains at least 2 points from $\circlearrowleft_y^{ux} \setminus \{p_1\}$. From Obs. 2.1.3 it follows that $\bigcirc up_1y$ and $\bigcirc p_1xy$ do not share a point from $\circlearrowleft_y^{ux} \setminus \{p_1\}$. Since $k = 4, 5$, there are at most $3 = 5-2$ points in $\circlearrowleft_y^{ux} \setminus \{p_1\}$. Then, exactly one triangle $\blacktriangle$ of $\triangle up_1y$ and $\triangle p_1xy$ is not order-k.

Now, we consider the following two cases depending how many points are contained in $\bigcirc\blacktriangle \setminus \circlearrowleft_y^{ux}$: $k-1$ or fewer.

Case 1) The disk $\bigcirc\blacktriangle$ contains $k-1$ points in $\bigcirc\blacktriangle \setminus \circlearrowleft_y^{ux}$. Recall that the points contained in $\bigcirc up_1x$ are the $k-1$ points in $\bigcirc uxy \setminus \circlearrowleft_y^{ux}$ and y. Hence, the $k-1$ points in $\bigcirc\blacktriangle \setminus \circlearrowleft_y^{ux}$ are the $k-1$ points in $\bigcirc uxy \setminus \circlearrowleft_y^{ux}$. Thus, for either $\blacktriangle = \triangle up_1y$ or $\blacktriangle = \triangle p_1xy$, this is an instance of case A again, but now with at most $k-2$ points in either of $\circlearrowleft_y^{up_1}$ or $\circlearrowleft_y^{p_1x}$, respectively.

Assume $\blacktriangle = \triangle up_1y$ (the case $\blacktriangle = \triangle p_2p_1y$ is analogous). Using the same arguments as in the beginning of this proof, there exists a point p_2 in $\circlearrowleft_y^{up_1} \subset \circlearrowleft_y^{ux}$ such that $\triangle up_2p_1 \in T$ and $\bigcirc up_2p_1$ does not contain points in $\circlearrowleft_y^{up_1}$. Then, the edge up_1 is an illegal edge in T_2. Consider $T_3 = (T_2 \setminus \{up_1\}) \cup \{p_2y\}$. Analogously, as before we show that T_3 has order $2k-3-1 = 2k-4$, since $(\bigcirc up_2y \cup \bigcirc p_2p_1y) \subset (\bigcirc up_1y \cup \bigcirc up_2p_1)$. Therefore, $T_3 \in \mathcal{T}_{2k-4}(S) \subset \mathcal{T}_{2k-2}(S)$. Note that for $k = 4$, $\mathcal{T}_{2k-4}(S) = \mathcal{T}_4$. Thus, for $k = 4$ the lemma follows.

Consider $k = 5$. If $T_3 \in \mathcal{T}_5$, then the lemma follows. Otherwise, consider $T_3 \notin \mathcal{T}_5$. Hence, $T_3 \in \mathcal{T}_6(S)$ since $2k-4 = 6$. Thus, one triangle $\hat{\blacktriangle}$ of the new triangles $\triangle up_2y$ or $\triangle p_2p_1y$ is of order 6. Since $p_2 \in \circlearrowleft_y^{up_1} \subset \circlearrowleft_y^{ux} \setminus \{p_1\}$, we have that $\circlearrowleft_y^{up_2} \subset \circlearrowleft_y^{ux} \setminus \{p_1\}$ and $\circlearrowleft_y^{p_2p_1} \subset \circlearrowleft_y^{ux} \setminus \{p_1\}$. Then, by Obs. 2.1.4, $\bigcirc\hat{\blacktriangle}$ contains the 2 points in $\circlearrowleft_y^{ux} \setminus \{p_1, p_2\}$, namely, in either the regions $\circlearrowleft_y^{up_2} \subset \circlearrowleft_y^{ux} \setminus \{p_1, p_2\}$ or $\circlearrowleft_y^{p_2p_1} \subset \circlearrowleft_y^{ux} \setminus \{p_1, p_2\}$. In addition, by Obs. 2.1.3 it follows that $\circlearrowleft_y^{up_2} \cap \circlearrowleft_y^{p_2p_1} = \emptyset$. Hence, exactly one of $\triangle up_2y$ and $\triangle p_2p_1y$ can have order 6. Also, since $\bigcirc\hat{\blacktriangle}$ has exactly 2 points in $\circlearrowleft_y^{ux}$, $\bigcirc\hat{\blacktriangle}$ contains the $k-1 = 4$ points in $\bigcirc up_1y \setminus \circlearrowleft_y^{up_1} \subset \bigcirc up_2p_1$. Therefore, $\bigcirc\hat{\blacktriangle} \setminus \circlearrowleft_y^{ux}$ contains $k-1$ points,

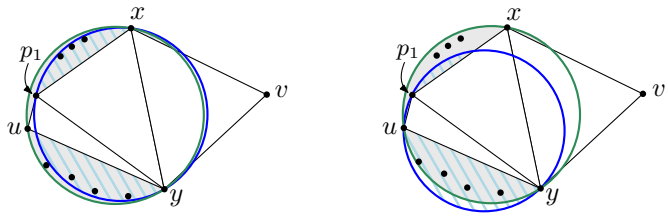

Figure 2.8: In both cases: $k = 5$, the gray filled disk is $\bigcirc uxy$ and both $\textcircled{\scriptsize }_y^{ux}$ and $\textcircled{\scriptsize }_x^{uy}$ contain 4 points. (a) Disk $\bigcirc p_1xy$ does not contain the 4 points in $\bigcirc uxy \setminus \textcircled{\scriptsize }_y^{ux}$. (b) The blue disk $\bigcirc up_1y$ contains the 4 points in $\bigcirc uxy \setminus \textcircled{\scriptsize }_y^{ux}$.

which is again an instance of Case A. Analogously, there is an illegal edge e (either up_2 or p_2p_1) such that the triangulation $T_4 = T_3 \setminus \{e\} \cup \{e'\}$ resulting after flipping e by e' in T_3 is of order at most $2k - 4 - 1 = 5$. Therefore, the lemma follows for this case.

Case 2) The disk $\bigcirc\blacktriangle$ contains fewer than $k - 1$ points in $\bigcirc\blacktriangle \setminus \textcircled{\scriptsize }_y^{ux}$. In addition, since $\textcircled{\scriptsize }_y^{ux} \subset \bigcirc uyv$ contains at most $k - 1$ points, the region of $\bigcirc\blacktriangle$ contained in $\textcircled{\scriptsize }_y^{ux}$ contains at most $k - 2$ points (the points that are in $\textcircled{\scriptsize }_y^{ux}$ minus p_1). Hence, note that this is a special instance of Case B. Thus, this case can only happen for $k = 5$. Moreover, $\bigcirc\blacktriangle$ contains exactly 3 points in $\textcircled{\scriptsize }_x^{uy}$ (either the 3 points in $\textcircled{\scriptsize }_{p_1}^{uy}$ or the 3 points in $\textcircled{\scriptsize }_x^{p_1y}$) and exactly 3 points in $\textcircled{\scriptsize }_y^{ux}$ (either the 3 in $\textcircled{\scriptsize }_y^{up_1}$ or the 3 in $\textcircled{\scriptsize }_y^{p_1x}$). In addition, $\textcircled{\scriptsize }_y^{ux} \subset \bigcirc uyv$ must contain $k - 1 = 4$ points in order for either $\textcircled{\scriptsize }_y^{up_1}$ or $\textcircled{\scriptsize }_y^{p_1x}$ to have exactly 3 points. Since $\bigcirc uxy \setminus \textcircled{\scriptsize }_y^{ux}$ contains $k-1$ points, these must lie in $\textcircled{\scriptsize }_x^{uy} \subset \bigcirc uxv$, since $\textcircled{\scriptsize }_u^{xy} \subset \bigcirc uxv \cap \bigcirc uyv$. But $\bigcirc uyv \setminus \textcircled{\scriptsize }_u^{xy}$ already contains k points: $k - 1$ in $\textcircled{\scriptsize }_y^{ux}$ and x. Notice that $\bigcirc up_1y$ contains $\textcircled{\scriptsize }_x^{uy}$, since the arc from u to y of $\partial \bigcirc up_1y$ that contains p_1 is in $\bigcirc uxy$. Thus, $\textcircled{\scriptsize }_{p_1}^{uy}$ contains exactly $k - 1 = 4$ points. Therefore we have $\blacktriangle = \triangle p_1xy$. See Fig. 2.8.

Let us show that the edge p_1x is an illegal edge. Suppose for the sake of a contradiction that p_1x is not illegal. Then, none of the points in $\textcircled{\scriptsize }_y^{p_1x}$ is in a triangle with both x and p_1. Hence, there exists a point p' in $\textcircled{\scriptsize }_y^{p_1x}$ that is in a triangle $\triangle ap'b \in T$ such that the edge ab is blocking either p_1 or x from all the points in $\textcircled{\scriptsize }_y^{p_1x}$. Thus, either both endpoints are different from x and p_1 or either $p_1 = a$ or $x = b$ but $p_1x \neq ab$. Since $\textcircled{\scriptsize }_y^{ux}$ contains 4 points and $\textcircled{\scriptsize }_y^{p_1x}$ contains 3 points, at least one of the endpoints of ab is not in $\bigcirc uxv \cup \bigcirc uyv$. If ab is blocking p_1 from p' then the arc from a to b of $\partial \bigcirc ap'b$ that contains p' crosses twice $\textcircled{\scriptsize }_y^{ux}$ and has a portion in $\textcircled{\scriptsize }_y^{p_1x} \subset \textcircled{\scriptsize }_y^{ux}$. Thus, $\bigcirc ap'b$ contains y, p_1, u and the 4 points in $\textcircled{\scriptsize }_x^{uy}$, contradicting that $\triangle ap'b$ is in T. If ab is blocking x

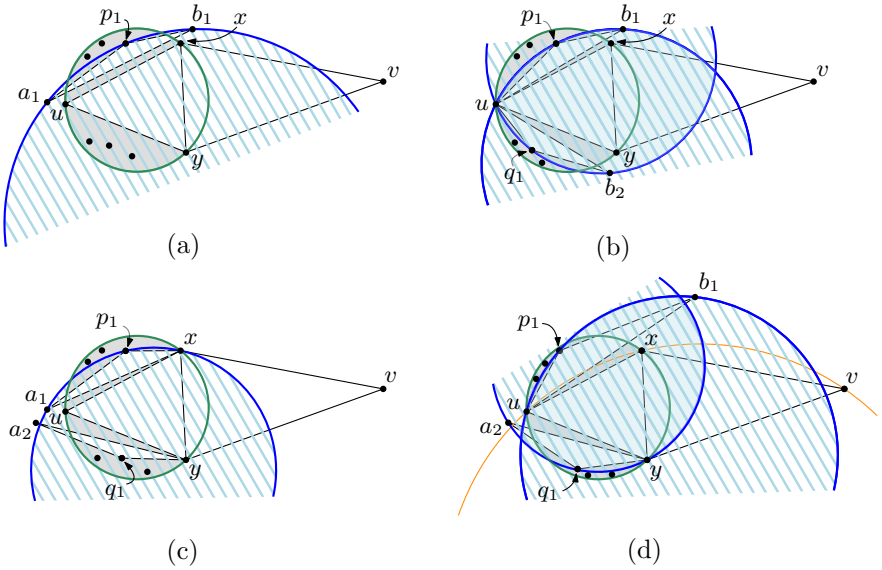

Figure 2.9: (a) Triangle $\triangle a_1p_1b_1$ is of order 6. (b) Triangle $\triangle up_1b_2$ has order 6. (c) Triangle $\triangle a_2q_1y$ is of order 5. (d) Triangles $\triangle up_1b_1$ and $\triangle a_2q_1y$ are of order 5.

from p', then the arc from a to b of $\partial \bigcirc ap'b$ that contains p' has a portion in $\mathcal{D}_y^{p_1x} \subset \mathcal{D}_y^{uv}$. Since at least one of the endpoints of ab is not in $\bigcirc uyv$, $\bigcirc ap'b$ contains y, x, v and the 3 points in $\mathcal{D}_x^{p_1y}$, contradicting that $\triangle ap'b$ is in T.

Therefore, there exists a point p_2 in $\mathcal{D}_y^{p_1x}$ such that $\triangle p_1p_2x$ is in T, and p_1x is illegal. Consider the triangulation $T_3 = T_2 \setminus \{p_1x\} \cup \{p_2y\}$.

Case 2a) Disk $\bigcirc p_1p_2x$ contains at most 3 points from $\mathcal{D}_x^{uy}$. Since the disks $\bigcirc p_2xy$ and $\bigcirc p_1p_2y$ are in the union of $\bigcirc p_1xy \cup \bigcirc p_1p_2x$, which contains at most 7 points including p_2 and y, the order of the triangles $\triangle p_2xy$ and $\triangle p_1p_2y$ is 5. Thus, the lemma follows.
Case 2b) Disk $\bigcirc p_1p_2x$ contains exactly 4 points from $\mathcal{D}_x^{uy}$. Hence, T_3 has order $|\mathcal{D}_x^{uy}| + |\mathcal{D}_y^{p_1x} \setminus \{p_2\}| = 4 + 2 = 6$. Therefore, $T_3 \in \mathcal{T}_6(S) \subset \mathcal{T}_{2k-2}(S)$. If $T_3 \in \mathcal{T}_5(S)$, then the lemma follows. Otherwise, consider the case when T_3 is of order 6. There is at most one triangle $\hat{\blacktriangle}$ from $\triangle p_1p_2y$ and $\triangle p_2yx$ of order 6. Thus the circumcircle of $\hat{\blacktriangle}$ contains the 4 points in $\mathcal{D}_x^{uy}$, implying that there are 4 points in $\bigcirc\hat{\blacktriangle} \setminus \mathcal{D}_y^{p_1x}$, which is an instance of Case A again. Using similar arguments as in case 1, the lemma follows. □

Case B

We consider the resulting triangulation T_1 after flipping uv such that Case B occurs: (1) $\triangle uxy$ is not order-5, (2) there are exactly three points in both $\textcircled{\scriptsize)}_y^{ux}$ and $\textcircled{\scriptsize)}_x^{uy}$.

We show that there exists an order-5 triangulation T' that can be reached from T in the flip graph of $\mathcal{T}_{2k-2}(S)$ by flipping at most 4 illegal edges.

Lemma 2.3.4. Let S be a point set in general position and let T be a minimal triangulation in $\mathcal{T}_5(S)$. Let uv be an illegal edge of T with adjacent triangles $\triangle uxv$ and $\triangle uyv$. Consider the triangulation $T_1 = (T \setminus \{uv\}) \cup \{ux\}$ such that:

1. $\triangle uxy$ is not order-5,
2. there are exactly three points in both $\textcircled{\scriptsize)}_y^{ux}$ and $\textcircled{\scriptsize)}_x^{uy}$.

Then, there exists a triangulation T' in $\mathcal{T}_5(S)$ at flip distance at most 3 from T in $G(\mathcal{T}_8(S))$ such that all of the flipped edges are illegal.

Proof. Since edge uv is an illegal edge in T, Since $\triangle uxy$ is of order 6, it follows from Lemma 2.1.2 that $\triangle xvy$ is of order 2, since $2k - 2 = 8$. Thus, $T_1 \in \mathcal{T}_{2k-2}(S)$.

Next we consider two cases, depending on whether one of ux or uy is illegal or not.

Case 1) One of ux or uy is an illegal edge. Assume ux is an illegal edge (the case when uy is illegal is symmetric). There exists a point p_1 in $\textcircled{\scriptsize)}_y^{ux}$ such that $\triangle up_1x$ is in T. Consider the triangulation $T_2 = T_1 \setminus \{ux\} \cup \{p_1y\}$. If $T_2 \in \mathcal{T}_5(S)$ then the lemma follows. Assume that T_2 is not of order 5.

Note that the disk $\bigcirc up_1y$ is contained in $\bigcirc uxv \cup \bigcirc uyv$ since the arc from u to y in $\partial \bigcirc up_1y$ that contains p_1 is in $\bigcirc uyv$ and the other arc from y to u is contained in $\bigcirc uxv$. Similarly, $\bigcirc p_1xy$ is in $\bigcirc uxv \cup \bigcirc uyv$ since the arc from p_1 to y of $\partial \bigcirc p_1xy$ that contains x is contained in $\bigcirc uyv$ and the other arc from y to p_1 is contained in $\bigcirc uyv$ with one portion in $\bigcirc uxv$. Moreover, since $p_1 \in \textcircled{\scriptsize)}_y^{ux}$, $\bigcirc up_1x$ can contain an extra point in $\bigcirc uxv \setminus (\bigcirc uyv \cup \bigcirc uxy)$. Thus, T_2 is of order 6, since each of $\bigcirc up_1y$ and $\bigcirc p_1xy$ can contain at most 6 points (the 5 points in $\bigcirc uxy \setminus \{p_1\}$ plus the extra point in $\bigcirc up_1x$). Since $|\bigcirc up_1x| + |\bigcirc uxy| = 5 + 6 - 2 = 9$, by Lemma 2.1.2 it follows that at most one of the triangles up_1y and p_1xy is of order 6, say $\triangle up_1y$. Hence, there are 4 points in $\bigcirc up_1y \setminus \textcircled{\scriptsize)}_y^{up_1}$ which is an instance of Case A. So, analogously as in the proof of Lemma 2.3.3, the lemma follows. Anagously, we prove it when $\triangle p_1xy$ is of order 6.

Case 2) None of the edges ux, uy is illegal. Since ux is not illegal, the triangle $\triangle utx$ of T incident to ux must have its third vertex t outside $\bigcirc uxy$, so t cannot be any of the three points in $\textcircled{\scriptsize)}_y^{ux}$. Hence, there are edges blocking

u or x from all of the vertices in $\circledcirc_y^{ux}$. Thus, there exists a point p_1 in $\circledcirc_y^{ux}$ that is incident to one of the edges that is blocking u or x from all the points in $\circledcirc_y^{ux}$.

Let a_1b_1 be the blocking edge such that $\triangle a_1p_1b_1 \in T$. See Fig. 2.9. For point a_1 two cases are possible: $a_1 = u$, or a_1 is outside $\bigcirc uxy$. Similarly, either $b_1 = x$, or b_1 is outside $\bigcirc uxy$. Observe that a_1b_1 intersects $\partial \bigcirc uxy$ twice, and both intersection points are on the closed arc from u to x that defines the boundary of $\circledcirc_y^{ux}$. Now, consider the cases whether a_1 and b_1 are in $\bigcirc uxy$.

If both a_1 and b_1 were outside $\bigcirc uxy$, then $\bigcirc a_1p_1b_1$ would contain at least six points: u, x, y and the three points that are in $\circledcirc_x^{uy}$, see Fig. 2.9a. This would make $\bigcirc a_1p_1b_1$ a triangle of T of order 6, a contradiction. Thus one of a_1, b_1 must be respectively u or x, and the other one must be outside $\bigcirc uxy$. Symmetrically, define triangle $\triangle a_2q_1b_2$ such that $q_1 \in \circledcirc_x^{uy}$ and a_2b_2 is blocking u or y from all the points in $\circledcirc_x^{uy}$. By the same argument as above, one of a_2, b_2 must be respectively u or y, and the other one must be outside $\bigcirc uxy$.

Consider $a_1 = a_2 = u$. See Fig. 2.9b. Since both $\bigcirc a_1p_1b_1$ and $\bigcirc a_2q_1b_2$ contain x and y, at least one of the following holds: $b_1 \in \bigcirc a_2q_1b_2$ or $b_2 \in \bigcirc a_1p_1b_1$. This again contradicts our assumption that T is of order 5.

Consider $b_1 = x$ and $b_2 = y$. See Fig. 2.9c. Since both $\bigcirc a_1p_1b_1$ and $\bigcirc a_2q_1b_2$ contain u, x and y, at least one of the following holds: $a_1 \in \bigcirc a_2q_1b_2$ or $a_2 \in \bigcirc a_1p_1b_1$. This also contradicts our assumption that T is of order 5.

Finally, consider $a_1 = u$ and $b_2 = y$ (the remaining case is symmetric, $a_2 = u$ and $b_1 = x$). See Fig. 2.9d. Notice that b_1 has to lie in the interior of $\bigcirc uyv$. Otherwise, $\bigcirc up_1b_1$ contains at least 6 points (the three points in $\circledcirc_x^{uy}$, and $\{x, y, v\}$), which is a contradiction. On the other hand, notice that triangulation T must contain the triangle $\triangle uxb_1$. Otherwise, b_1 form a triangle with an edge ub_3, such that b_3 is outside $\bigcirc uyv$, since $\bigcirc uyv$ contains already 5 points (the three points in $\circledcirc_y^{ux}$, and $\{x, b_1\}$). Hence, $\triangle ub_1b_3$ is of order 6, since part of the arc of $\partial \bigcirc ub_1b_3$ from u to b_3 that contains b_1 is in $\bigcirc uxy$ and $\bigcirc uyv$, the disk $\bigcirc ub_1b_3$ must contain the three points in $\circledcirc_x^{uy}$, and $\{x, y, v\}$, which is a contradiction. Therefore, $\triangle uxb_1$ is in T. Moreover, $\bigcirc uxb_1$ contains the three points in $\circledcirc_y^{ux}$ since the arc from u to x of $\partial \bigcirc uxb_1$ that does not contain b_1 is contained in $\bigcirc uxy \setminus \circledcirc_v^{ux}$. Hence, the edge ub_1 is an illegal edge in T_1.

Now, consider the triangulation $T_2 = T_1 \setminus \{ub_1\} \cup \{p_1x\}$.

Let us show that $\triangle p_1b_1x$ is of order 4. First, notice that if a point is in the intersection $\bigcirc uxv \cap \bigcirc uyv$, then such point has to be in $\circledcirc_y^{ux}$, otherwise $\bigcirc uyv$ contains more than 5 points (the three points in $\circledcirc_y^{ux}$, $\{x, b_1\}$ plus the extra point in $(\bigcirc uyv \cap \bigcirc uxv) \setminus \circledcirc_y^{ux}$). Notice that the arc of $\partial \bigcirc p_1b_1x$ from b_1 to p_1 that contains x is in $\bigcirc uyv$. Hence, if $\partial(\bigcirc p_1b_1x)$ intersects twice $\partial(\bigcirc uxv)$ (otherwise $\bigcirc p_1b_1x$ and $\bigcirc uxv$ are tangent), $(\bigcirc p_1b_1x \cap \bigcirc uxv) \subset$

($\bigcirc uxv \cap \bigcirc uyv$). Thus, $\bigcirc p_1b_1x$ can only contain points in $\mathcal{O}_y^{uv} \setminus \{p_1, x, b_1\}$ and the points in $\bigcirc uxb_1 \setminus (\bigcirc uxv \cup \{p_1\})$. Since $\bigcirc uxb_1$ contains the points in $\mathcal{O}_y^{uv} \setminus \{x, b_1\}$, $\bigcirc uxb_1$ contains at most 4 points (the at most 5 points in $\bigcirc uxb_1$ minus p_1).

Now, let us show that triangle $\triangle up_1x$ is of order 6. Since T is minimal, then $\triangle up_1x$ is not of order 5. Notice that p_1 lies in $\mathcal{O}_y^{ux} \setminus \bigcirc uxv$, otherwise $\bigcirc up_1b_1$ contains more than 5 points: Since the arc from u to b_1 of $\partial \bigcirc up_1b_1$ that contains p_1 crosses twice $\partial \bigcirc uxv$ and a portion of such arc is in $\mathcal{O}_v^{ux}$, $\bigcirc up_1b_1$ contains the 3 points in $\mathcal{O}_x^{uy}$, and $\{x, y, v\}$, which is a contradiction. Thus, disk $\bigcirc up_1x$ must contain the points in $(\bigcirc up_1b_1 \cap \bigcirc uxv) \setminus \{x\}$ which are the 3 points in $\mathcal{O}_x^{uy}$ and y. In addition, since the arc of $\partial \bigcirc up_1x$ from u to x containing p_1 is in $\mathcal{O}_y^{ux}$, the disk $\bigcirc up_1x$ can also contain the 2 points in $\mathcal{O}_y^{ux} \setminus \{p_1\}$. Hence, $\triangle up_1x$ is of order 6. Thus, $T_2 \in \mathcal{T}_6(S) \subset \mathcal{T}_8(S)$. Since p_1 is in $\mathcal{O}_y^{ux}$, the edge ux is an illegal edge in T_2.

Consider the triangulation $T_3 = T_2 \setminus \{ux\} \cup \{p_1y\}$. Notice that the two new triangles $\triangle up_1y$ and $\triangle p_1xy$ are of order 5 since the union of $\bigcirc up_1x \cup \bigcirc uxy$ only contains the points in $\bigcirc uxy$. Therefore, T_3 is of order 5. □

Finally, Theorem 2.3.1 follows.

proof of Theorem 2.3.1. Note that for $k = 2$, the theorem follows since $G(\mathcal{T}_2(S))$ is connected. If T is not minimal, then there exists an illegal edge e such that the resulting triangulation T' after flipping e is of order k. Assume that T is a minimal triangulation. Hence, by Lemmas 2.3.3 and 2.3.4 there exists an order-k triangulation T' that is flip distance at most $k-1$ in $G(\mathcal{T}_{2k-2}(S))$ such that all the flipped edges are illegal. □

Finally, since there are $O(kn)$ order-k edges (see [2, 62]), it follows from Theorem 2.3.1 that for $k \leq 5$, any order-k triangulation can be transformed into $DT(S)$ by a sequence of at most $O(kn)$ triangulations of order at most $2k-2$.

Corollary 2.3.5. Let S be a point set in general position. For $k \leq 5$, any triangulation in $T_k(S)$ can be transformed into $DT(S)$ by a sequence of at most $O(kn)$ triangulations of order $2k-2$.

Proof. Let $T \in \mathcal{T}_k$. Let us show that there is a sequence of triangulations $T_0 = T, T_1, \ldots, T_m = DT(S)$. By Sibson [96], T can be transformed into $DT(S)$ by a finite number of illegal flips. For $i \geq 0$, if T_i is not minimal then there exists an illegal edge u_iv_i in two triangles $\triangle u_ix_iv_i$ and $\triangle u_iy_iv_i$ such that $T_{i+1} = T_i \setminus \{u_iv_i\} \cup \{x_iy_i\}$ is an order-k triangulation. If T_i is minimal, then flipping any illegal edge in T_i will result into a triangulation that is not order-k. It follows from Theorem 2.3.1 that there exists an order-k triangulation T_j at flip distance

at most $k-1$ in $G(\mathcal{T}_{2k-2}(S))$ from T_i such that all the flipped edges are illegal. So, add the sequence of triangulations between T_i and T_j: $T_i, T_i{+}1, \ldots, T_j$ where $j < i+k$. Since all of the flipped edges (including the ones of order-$(2k-2)$) are illegal, the sequence of order-$(2k-2)$ triangulations is strictly increasing with respect to the angle vector and $DT(S)$ is the only maximum element, at most all of the order-$(2k-2)$ edges are flipped in order to reach $DT(S)$. By Abellanas et al. [2] and Gudmundsson et al. [62], there are $O((2k-2)n) = O(kn)$ flipped edges. □

Using Theorem 2.3.1, the reverse search framework of Avis and Fukuda [16] and pre-processing the order-k triangles of a points set S in general position, we can enumerate all of the triangulations of order-k in expected polynomial time per triangulation.

Corollary 2.3.6. For $k \leq 5$, there exists an algorithm that enumerates all k-order Delaunay triangulations of any point set in general position in expected $O(k^4n^2)$ time per triangulation.

Proof. Let S be a point set in general position. First, we compute the Delaunay triangulation of S which takes $O(n \log n)$ time. Also, we use the pre-processing algorithm that computes all the triangles and edges that are in triangulations of order k of S defined by Silveira and van Kreveld [97], which takes $O(k^2 n \log k + kn \log n)$ expected time. We place each triangle and corresponding edges in a hash table. Now, in order to enumerate all of the k-order Delaunay triangulations of S for $k \leq 5$, we do a reverse search method as the one define by Avis and Fukuda [16] where $DT(S)$ defines the root element of order-k triangulations of S. For a triangulation T in $\mathcal{T}_k(S)$ in order to find its child T' in $\mathcal{T}_k(S)$ and the parent of T' of a rooted spanning tree of $\mathcal{T}_k(S)$, we do the following:

For each vertex v in T, we look at any pair of incident edges $e_1 = (v, a)$ and $e_2 = (v, b)$ to v such that $ab \notin T$ and there are at most $k-1$ edges in T crossing ab. Then, in expected $O(1)$ time, we can check in the hash table if $\triangle avb$ is of order k. Since we do this for each vertex, in at most $\sum_{v \in S} \frac{deg_T(v)(deg_T(v)-1)}{2} = O(n^2)$ we can check which are the possible order-k triangles that T can reach after $O(k-1)$ edges being flipped, in order to reach another triangulation of order k. By [62] there are $O(kn)$ edges of order k that belong to a triangulation of order k. Hence, there are $O(kn)$ possible triangles of order k that are not in T where for each triangle $\triangle avb$ exactly one of its edges ab is not in T and is crossed by at most $k-1$ edges of T. Thus, we have to flip at most $k-1$ edges from T in order to get a triangulation that contains $\triangle avb$. In order to know which triangulations of order k can be obtained after flipping at most $k-1$ edges of T, we do the following. Per each possible triangle $\triangle avb$ of order k with

exactly one edge ab not in T, we find the vertex p_1 of edge vp_1 that crosses ab in T with the following properties: 1) $\triangle ap_1b$ is of order k, and 2) $\bigcirc ap_1b$ does not contain point p such that $vp \in T$ and crosses ab. This takes $O(k)$ expected time. We do the same for each point p adjacent to v and edge vp crosses either ap_1 or p_1b. Since there are at most $k-1$ edges crossing ab, this process takes $O(k^2)$ expected time. Thus, in at most $O(k^3n)$ time we can find possible children of T. Per each possible child T' of T, we can check in $O(kn)$ time the triangulation T'' of order k with maximum angle vector with $\alpha(T'') > \alpha(T)$. Therefore, there exists an algorithm that enumerates the elements of $\mathcal{T}_k(S)$ in $O(k^4n^2)$ expected time per triangulation. □

2.4 Conclusions

In this chapter we presented the first general results on the flip graph of order-k Delaunay triangulations. We showed that already for points in convex position, the flip graph may not be connected. This is in contrast to the flip graph of triangulations that consist of order-k *edges*, for which the flip graph is always connected [2]. Our main result is that $k-1$ flips are sometimes necessary and always sufficient to transform an order-k triangulation into some other order-k triangulation, for any $k \geq 2$ if the points are in convex position. Moreover, we proved that these $k-1$ flips go through triangulations of order $2k-2$. This is a noteworthy result, and one of the first results on order-k Delaunay triangulations proven for any value of k. In the setting of general point sets, we also showed that for $k = 3, 4, 5$, the order-k triangulations are at flip distance at most $k-1$ from some other order-k triangulation within the flip graph of order-$(2k-2)$ triangulations. This result also implies that the flip distance between any two order-k triangulations is $O(kn)$, which is consistent with the fact that the diameter of the flip graph of all triangulations is $\Theta(n^2)$.

Our results imply an enumeration algorithm using the Avis and Fukuda framework [16], generalizing the results that Abe and Okamoto obtained for $k \leq 2$ [1]. For the case of convex position this is not of practical importance, as in that case one can obtain a more efficient method by first pre-computing all order-k triangles, and then recursively enumerating all order-k triangulations. However, our results imply the first non-trivial enumeration results for points in generic position for $3 \leq k \leq 5$. It should be mentioned that small values of k are the most important ones in practice, since for small orders the triangle shape is still close to Delaunay, but at the same time they are enough to obtain significantly better triangulations [91].

Clearly, the main question left open is what happens in general when $k \geq 6$. For larger orders our techniques present issues due to a large increase in the

number of cases that need to be considered. However, we conjecture that the same results obtained for convex position hold in general, and in particular, that any order-k triangulation is at flip distance at most $k-1$ from another order-k triangulation.

Convex shape Delaunay graphs and trees

3

Flips in triangulations are an instance of a more general question: Given a class of spanning graphs $\mathcal{G}$ of a point set S, which local operations can transform one element of $\mathcal{G}$ into another? Several operators have been introduced for different classes of spanning graphs that are not triangulations, such as perfect matchings, plane Hamiltonian cycles and plane spanning trees [6, 9, 36]. One class of graphs that has received much attention in these type of operators, is the set of plane spanning trees of a point set.

Consider the set of all *plane* spanning trees [1] of S, denoted $\mathcal{ST}(S)$. Avis and Fukuda [16] considered a graph with vertex set $\mathcal{ST}(S)$ where two vertices are connected if the corresponding trees can be transformed into each other by an *edge move*, i.e., the corresponding two trees differ by exactly one edge. They showed that this graph is connected with diameter $O(n)$. In other words, Avis and Fukuda showed that any two distinct plane spanning trees can be transformed into each other by performing a linear number of edge moves. Observe that the edge move is a very general operation, since once we remove one edge of a tree, there might be a quadratic number of possible replacements.

A more local operation for transforming two different plane trees is the *edge slide*, which consists of replacing the edge uv from $T \in \mathcal{ST}(S)$ by an edge vw such that w is adjacent to u in T. See Figure 3.1. Consider the graph $G(\mathcal{ST}(S))$ with vertex set $\mathcal{ST}(S)$ and a pair of vertices are joined if the corresponding trees can be transformed into each other by a slide edge operation. Aichholzer et al. [5] showed that $G(\mathcal{ST}(S))$ is connected. Later, it was shown that the diameter of $G(\mathcal{ST}(S))$ is $O(n^2)$ [8].

[1] A planar graph G is called *plane* when $V(G)$ is mapped to a set of points in the plane and edges are mapped to segments whose interiors do not intersect. A plane graph refers to the planar embedding of a planar graph.

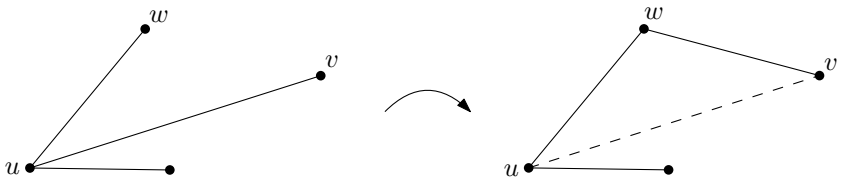

Figure 3.1: Example of slide edge operation in a tree.

In order to prove that $G(\mathcal{ST}(S))$ is connected, Aichholzer et al. [5] used another operation that transforms one plane spanning tree into a complementary plane spanning tree. This result is interesting on its own, since they showed that any plane spanning tree of S can be transformed into the Euclidean minimum spanning tree of S by repeatedly computing Euclidean minimum spanning trees of the constrained Delaunay triangulations. More precisely, the following iterative procedure converges to the Euclidean minimum spanning tree of S. Start with an arbitrary spanning tree T_0 of S. Compute the Delaunay triangulation of S taking the edges of T_0 as constraints. Next, compute the Euclidean minimum spanning tree T_1 of this constrained triangulation, and repeat. Aichholzer et al. [5] also showed a tight bound of $\Theta(\log n)$ on the length of this sequence.

A natural question with respect to $\mathcal{C}$-distances is whether there exists a sequence that transforms any plane spanning tree of S into the $d_{\mathcal{C}}$-minimum spanning tree of S such that two consecutive plane trees are contained in the constrained $\mathcal{C}$-Delaunay graph.

Our contributions. In this chapter we first show that the result by Aichholzer et al. [5] can be generalized to arbitrary $\mathcal{C}$-distance functions, i.e., there exists a sequence of trees $T_0, T_1, \ldots T_k$ such that $T_i = MST_{\mathcal{C}}(DT_{\mathcal{C}}(T_{i-1}))$ and this procedure converges in $O(n)$ steps to the $d_{\mathcal{C}}$-minimum spanning tree of S, where $|S| = n$. The main ingredient of our result is, as in [5], a fixed tree theorem (Theorem 3.2.1) that states that once one iteration of the above procedure does not produce a change, then it must have reached the $d_{\mathcal{C}}$-minimum spanning tree of S. We note, however, that our proofs are different from those used for the Euclidean metric, since several key lemmas in [5] rely on properties of circles. Later, we show a lower bound of $\Omega(\log n)$ on the sequence of trees that converges to the $d_{\square}$-minimum spanning tree. Finally, for $\square$-distance we prove that for a special case of plane spanning trees, the length of this sequence is $O(\log n)$.

3.1 General observations

In this section we give different results that will be essential for the proof of the fixed tree theorem.

Let S be a point set in the plane and let $\mathcal{C}$ be a compact convex set that contains the origin in it interior. Throughout this chapter we will assume that all pairwise $\mathcal{C}$-distances between points in S are different. Let G be a planar graph with vertex set S and edges $E(G)$. Let p and q be two points in S. We say that a path is a *pq-path* in G if it is contained in G and it starts from p and ends in q.

The following property is a known characterization of the edges in the minimum spanning tree.

Property 3.1.1. For any graph G, an edge $e \in G$ is not present in $MST_{\mathcal{C}}(G)$ if and only if there is a path in G between the endpoints of e that solely consists of edges shorter than e.

Next we show that for each pair of points p and q that are visible to each other there exists a path that connects them in $DG_{\mathcal{C}}(G)$ fully contained in $\mathcal{C}_r(p,q)$ for all $r \geq d_{\mathcal{C}}(p,q)$.

Lemma 3.1.2. Let p, q be two visible vertices such that pq in G. Then, for all $r \geq d_{\mathcal{C}}(p,q)$ the $\mathcal{C}$-disk $\mathcal{C}_r(p,q)$ contains a pq-path in $DG_{\mathcal{C}}(G)$.

Proof. Let $r \geq d_{\mathcal{C}}(p,q)$ and let C be a $\mathcal{C}_r(p,q)$. The proof is by induction on the number of points of S contained in the interior of the disk C. If C does not contain points in its interior, then by definition, there exists the edge pq in $DG_{\mathcal{C}}(S)$.

Now, suppose that the result holds for $k \geq 0$ points contained in $\mathcal{C}_r(p,q)$. Assume that C contains $k+1$ vertices of S in its interior.

Note that if none of the points in the interior of C is visible to both of p and q, then by definition the edge $pq \in DG_{\mathcal{C}}(G)$. Consider the case when at least one point $v \in S$ contained in C is visible to both points. Since v is in the interior of C, there exist $C' = \mathcal{C}_{r'}(p,v)$ and $C'' = \mathcal{C}_{r''}(v,q)$ that are fully contained in C. Hence, C' and C'' contain at most k points in their interior. Then, by inductive hypothesis, there exist a pv-path P' and a vq-path P'' that are contained in $C' \subset C$ and $C'' \subset C$, respectively. Therefore, there is a pq-path in $P' \cup P''$ that is contained in C (note that $P' \cup P''$ might repeat edges, i.e., it is not necessarily a pq-path but it contains one). □

Observe that for any edge $uv \in MST_{\mathcal{C}}(S)$ it holds that any $\mathcal{C}(u,v)$ does not contain any point of S in its interior, otherwise, from Lemma 3.1.2 there exists an uv-path with edges solely shorter than uv, contradicting Property 3.1.1.

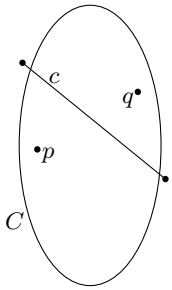

Figure 3.2: The points p and q are separated in C by a constraint c.

We say that p and q are *separated* in a $\mathcal{C}$-disk C if p and q are in C and there exists a constraint $c \in E(G)$ crossing C such that c divides C into two different sets such that p and q are in different sets, we refer to Figure 3.2. Using Lemma 3.1.2 we can show that if two points are contained in a homothet C of $\mathcal{C}$ where no edge is separating them (note that, despite this, the points might not be visible to each other), then there exists a path between the two points in $DG_{\mathcal{C}}(G)$ that is contained in C.

Lemma 3.1.3. Let C be a $\mathcal{C}_r(p,q)$ for some $r \geq d_{\hat{\mathcal{C}}}(p,q)$ such that there is no constraint separating p and q in C, then there exists a pq-path of $DG_{\mathcal{C}}(G)$ contained in C.

Proof. Since there is no constraint separating p and q, there exists a path $P = p_0, p_1, \ldots, p_m$ in the visibility graph in C, where $p_0 = p$ and $p_m = q$. The vertices p_i and p_{i-1} can see each other, since the edge $p_{i-1}p_i$ is in the visibility graph for all $i \in \{1, \ldots, n\}$. Also, since p_{i-1} and p_i are in C, there is an r_i such that $C_i = \mathcal{C}_{r_i}(p_{i-1}, p_i)$ is contained in C for each i. Thus by Lemma 3.1.2 there exists a $p_{i-1}p_i$-path P_i in C_i for all i. Therefore, there is a pq-path in $P_1 \cup P_2 \cup \ldots \cup P_m$ that is contained in C (note that $P_1 \cup P_2 \cup \ldots \cup P_m$ might repeat edges, i.e., it is not necessarily a pq-path but it contains one). □

A characteristic of circles is that two different circles intersect either in exactly one point or exactly two points where one arc of a circle is enclosed by the other and vice versa. Something similar is true for convex curves.

Proposition 3.1.4 (Ge Xia [108])**.** Let $\partial\mathcal{C}$ be the boundary of $\mathcal{C}$, which is a closed convex curve in the plane. The intersection of two distinct homothets of $\partial\mathcal{C}$ is the union of two sets, each of which is either a segment, a single point, or empty.

We know that if two points z and w lie in the interior of a $\mathcal{C}(p,q)$, then $d_{\mathcal{C}}(z,w) \leq d_{\mathcal{C}}(p,q)$ and that there exists a $\mathcal{C}_r(z,w)$ contained in $\mathcal{C}(p,q)$. How-

ever, it is possible that $\mathcal{C}(z,w)$ is not fully contained in $\mathcal{C}(p,q)$. See Figure 3.3a.

We observe that if the points z and w lie in the line segment pq then there exists a $\mathcal{C}(z,w)$ in the interior of $\mathcal{C}(p,q)$.

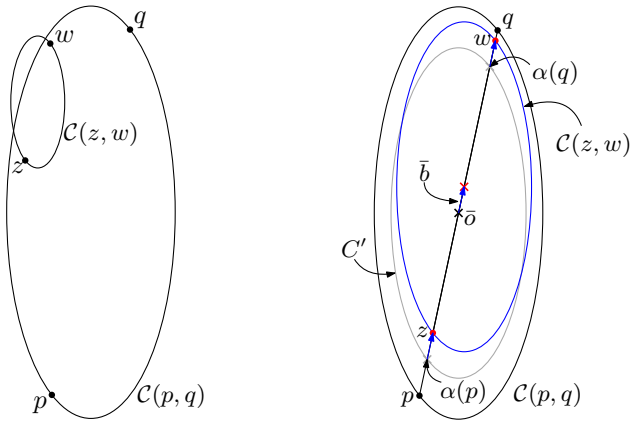

Figure 3.3: (a) The unique $\mathcal{C}(z,w)$ is not fully contained in $\mathcal{C}(p,q)$. (b) The line segment zw lies in the line segment pq. The origin corresponds to the midpoint of pq. The gray ellipse is the homothet C' of $\mathcal{C}(p,q)$ scaled by $\frac{d(z,w)}{d(p,q)} = \frac{1}{\lambda}$, and the ellipse $\mathcal{C}(z,w)$ is the homothet of C' translated by a vector $\bar{b}$.

Observation 3.1.5. Let z and w be two points that lie in the line segment pq, then for each $C = \mathcal{C}(p,q)$ there exists a $\mathcal{C}(z,w)$ contained in C, and vice versa, for each $C' = \mathcal{C}(z,w)$ there exists a $\mathcal{C}(p,q)$ that contains C'.

Proof. Since scaling of $\mathcal{C}$ does not depend on where the center of $\mathcal{C}$ is, we consider without loss of generality a convex set $\mathcal{C}'$, homothet of $\mathcal{C}$, such that the origin $\bar{o}$ is in the midpoint of pq. Since $d(z,w) = \lambda d(p,q)$ we can scale C by $1/\lambda$ with a linear transformation α and translate the new scale C' by a vector $\bar{b}$, denoted $C' + \bar{b}$, that maps $\bar{o}$ to the midpoint b of zw, then $\alpha(p)\alpha(q) + \bar{b} = zw$. See Figure 3.3b. By construction, $C' + \bar{b}$ is a $\mathcal{C}_r(z,w)$. In addition, $C' + \bar{b}$ is a $\mathcal{C}(z,w)$, otherwise there exists a smaller radius (scaling factor) r such that $\mathcal{C}_r(z,w)$ is smaller than C', then scaling it by λ we obtained a smaller $\mathcal{C}(p,q)$ than C, which is a contradiction. Note that if we show that $d_{\mathcal{C}'}(\bar{o},u) \leq d_{\mathcal{C}'}(\bar{o},p)$ for all $u \in (C' + \bar{b})$, then $C' + \bar{b}$, that is a $\mathcal{C}(z,w)$, is contained in C. Let $u \in (C' + \bar{b})$, then $d_{\mathcal{C}'}(\bar{o},u) \leq d_{\mathcal{C}'}(\bar{o},b) + d_{\mathcal{C}'}(b,u) \leq d_{\mathcal{C}'}(\bar{o},b) + d_{\mathcal{C}'}(b,z) \leq d_{\mathcal{C}'}(\bar{o},b) + d_{\mathcal{C}'}(b,p) = d_{\mathcal{C}'}(\bar{o},p)$. The second part of the observation is proved analogously. □

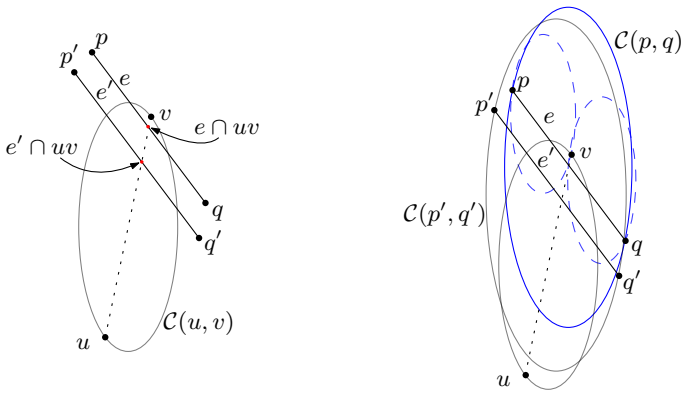

(a) (b)

Figure 3.4: Edges e and e' are edges crossing $uv \in MST_{\mathcal{C}}$: (a) The edge e is closer to v than e', and e' is closer to u than e. (b) None of $\mathcal{C}(p,q)$ and $\mathcal{C}(p',q')$ contain u, but both contain v.

Let $t = uv$ be an edge of the $MST_{\mathcal{C}}(S)$ crossed by a set of constraints. Let C_t be a $\mathcal{C}(u,v)$. Let $e = pq$ and $e' = pq$ be two edges in $DG_{\mathcal{C}}(G)$ that cross t. We say that e is closer than e' to v if $d(e \cap t, v) \leq d(e' \cap t, v)$. See Figure 3.4a.

Lemma 3.1.6. Let $e = pq$ and $e' = p'q'$ be edges of $DG_{\mathcal{C}}(T)$ crossing uv such that u is not contained in any $\mathcal{C}(p',q')$ and, e is closer to v than e'. Then, there exists a $\mathcal{C}(p,q)$ that contains v. See Figure 3.4b.

Proof. Let C_t be a $\mathcal{C}(u,v)$. Since $uv \in MST_{\mathcal{C}}(S)$, p and q lie outside C_t. Consider the intersection points z', w' of $e' \cap \partial C_t$. Let $C_{e'}$ be a $\mathcal{C}(p',q')$. By Observation 3.1.5 there exists a $C_{z'w'} = \mathcal{C}(z',w')$ contained in $C_{e'}$. Since z' and w' lie on ∂C_t, $z'w'$ partitions C_t into two convex subsets of C_t. Thus, by Proposition 3.1.4 one of the two convex subsets defined by $z'w'$ in C_t is totally contained in $C_{z'w'}$: Since $C_{e'}$ does not contain u, $C_{z'w'}$ does not contain u neither. Hence, $C_{z'w'}$ contains the subset of C_t defined by $z'w'$ that contains v.

Since e crosses t closer to v than e', the intersection points z and w of $e \cap \partial C_t$ are contained in $C_{z'w'}$. Hence, $d_{\hat{\mathcal{C}}}(z,w) \leq d_{\hat{\mathcal{C}}}(z',w')$. Note that if there exists a $\mathcal{C}(z,w)$ that contains v in its interior, then, by Observation 3.1.5 there exists a $\mathcal{C}(p,q)$ that contains v in its interior. For the sake of a contradiction suppose that none of the $\mathcal{C}(z,w)$ contains v. Let C_{zw} be a $\mathcal{C}(z,w)$. Since z and w are on $\partial\mathcal{C}(u,v)$ and $v \notin C_{zw}$, by Proposition 3.1.4 it follows that C_{zw} contains the subset of C_t that contains u. Since e' is closer to u than e, C_{zw} contains z' and w'. Then, $d_{\hat{\mathcal{C}}}(z,w) = d_{\hat{\mathcal{C}}}(z',w')$. Thus, C_{zw} is a $\mathcal{C}(z',w')$ that contains u. Therefore, by Observation 3.1.5 it follows that there exists a $\mathcal{C}(p',q')$ that

contains u, a contradiction. □

3.2 Fixed tree theorem

Let $\mathcal{ST}(S)$ be the set of all crossing-free spanning trees of S. For each element $T \in \mathcal{ST}(S)$, we define the sequence $T_0, T_1, T_2, \ldots$ where $T_0 = T$, $T_i = MST_{\mathcal{C}}(DG_{\mathcal{C}}(T_{i-1}))$ for all $i > 0$. In this section we will prove the convergence of this sequence to the canonical element $MST_{\mathcal{C}}(S)$ of $\mathcal{ST}(S)$. In other words, we show that any spanning tree of S can be transformed into the $MST_{\mathcal{C}}(S)$ by a finite sequence of crossing-free spanning trees.

The following theorem shows that the fixed tree in the sequence of $T_0, T_1 \ldots$ is unique and it is the $MST_{\mathcal{C}}(S)$.

Theorem 3.2.1 (Fixed tree theorem)**.** Let $T \in \mathcal{ST}_S$. $T = MST_{\mathcal{C}}(DG_{\mathcal{C}}(T))$ if and only if $T = MST_{\mathcal{C}}(S)$.

Proof. The "if" part is trivial by definition of $MST_{\mathcal{C}}(S)$. Let us prove the "only if" part. Let $T = MST_{\mathcal{C}}(DT_{\mathcal{C}}(T))$, and assume for the sake of a contradiction that $T \neq MST_{\mathcal{C}}(S)$. Then, there exists an edge $t = uv \in MST_{\mathcal{C}}(S)$ that does not belong to T. Note that there is at least one edge in T that crosses t: otherwise, $t \in DG_{\mathcal{C}}(T)$. Thus, t must be in $MST_{\mathcal{C}}(DG_{\mathcal{C}}(T))$, a contradiction. Hence, there is at least one edge in T that crosses t.

Let C_t be a $\mathcal{C}(u, v)$. Each edge in $DG_{\mathcal{C}}(T)$ crossing t crosses C_t and has its endpoints outside C_t. Let $c = ab \in T$ be the constraint closest to v. We consider whether there exists a $\mathcal{C}(a, b)$ that contains v.

If there exists a $C_c = \mathcal{C}(a, b)$ that contains v, then none of the edges of T separates v from a and b in C_c. Otherwise, such edge crosses t and is closer to v than c, a contradiction. Hence, by Lemma 3.1.3 there is an av-path, P_a, and a vb-path, P_b, contained in C_c. Therefore there is an ab-path in $P_a \cup P_b$ that solely consists of edges shorter than c, contradicting our hypothesis by Property 3.1.1.

If there does not exist a $\mathcal{C}(a, b)$ that contains v, then by Proposition 3.1.4, any $\mathcal{C}(a, b)$ must contain u. Consider the closest constraint $c' = a'b'$ to u. Thus, by Lemma 3.1.6 there exists a $C_{c'} = \mathcal{C}(a', b')$ that contains u. Again, there is no edge in $C_{c'}$ that separates u from a' and b'. Thence, by Lemma 3.1.3 there is an $a'u$-path, $P_{a'}$, and an ub'-path, $P_{b'}$, contained in $C_{c'}$. Therefore there is an $a'b'$-path in $P_{a'} \cup P_{b'}$ that solely consists of edges shorter than c, contradicting our hypothesis by Property 3.1.1. Therefore $T = MST_{\mathcal{C}}(S)$. □

Consider again an arbitrary tree $T \in \mathcal{ST}(S)$ and a sequence $T_0, T_1, \ldots$, such that $T_0 = T, T_i = MST_{\mathcal{C}}(DG_{\mathcal{C}}(T_{i-1}))$ for all $i \geq 1$. Notice that this is a $d_{\mathcal{C}}$-length-decreasing sequence, since T_i has smaller $d_{\mathcal{C}}$-weight than T_{i-1}, unless both are identical trees. Also, by definition of T_i, the trees T_i and T_{i-1} do not

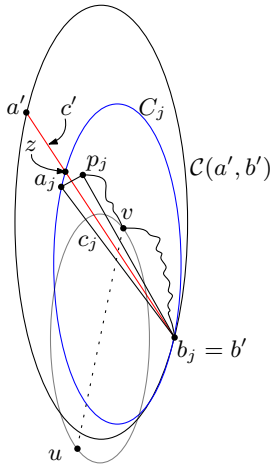

Figure 3.5: Edge $c_j = a_j b_j$ is the closest edge to v in T_j with $j = n - 3$ and $\triangle_j = \triangle a_j p_j b_j$ is in $DT_{\mathcal{C}}(T_j)$. Edge c' is in T_{j-1} and it is blocking the visibility from a_j to p_j.

cross, since they belong to the same plane graph, namely $DG_{\mathcal{C}}(T_{i-1})$, for all $i > 0$. Since the size of S is finite, there are only a finite number of crossing-free spanning trees of S. Thus, as a consequence of Theorem 3.2.1 we obtain a $d_{\mathcal{C}}$-length-decreasing sequence of trees in $\mathcal{ST}(S)$ which reaches a fixed point $T_k = MST_{\mathcal{C}}(S)$ in a finite number of steps. We will call this sequence the *canonical sequence* of T.

Theorem 3.2.2. For any $T \in \mathcal{ST}(S)$ there exists a finite sequence $T_0, T_1, \ldots, T_k$ in $\mathcal{ST}(S)$, such that $T_0 = T$ and $T_k = MST_{\mathcal{C}}(S)$.

Figure 3.6 shows a canonical sequence that converges in 4 steps where $\mathcal{C}$ is an axis-aligned square.

3.2.1 Upper bound

The next question to answer is what the length of the canonical sequence is. In this section we give a linear upper bound for the length of such sequence.

First, we show that for any $i \geq 0$, if an edge is in the constrained Delaunay graph of T_i and does not appear in T_{i+1}, then such edge cannot appear in any of the following trees of the canonical sequence.

Lemma 3.2.3. Let $e = pq$ be an edge in $DG_{\mathcal{C}}(T_i)$ that is not in tree T_{i+1}, with $i \geq 0$. Then $d \notin T_{i+j}$ for all $j > 1$.

Proof. Recall that $T_{i+1} = MST_{\mathcal{C}}(DG_{\mathcal{C}}(T_i))$. Since $e \notin T_{i+1}$, then by Property 3.1.1 there exist a pq-path in $DG_{\mathcal{C}}(T_i)$ that solely consists of edges of

$d_{\mathcal{C}}$-length smaller than $d_{\hat{\mathcal{C}}}(p,q)$. In particular, T_{i+1} contains a pq-path with this property, such a path is in $DG_{\mathcal{C}}(T_{i+1})$, hence $e \notin MST_{\mathcal{C}}(DG_{\mathcal{C}}(T_{i+1})) = T_{i+2}$. By induction, e can never reappear in the following trees. □

Next, we show that the canonical sequence converges in $O(n)$ steps, where n is the size of S. The proof is similar to the one in Aichholzer et al. [5] for circles, yet we make use of special properties of arbitrary convex shapes.

Theorem 3.2.4. The canonical sequence has length $O(n)$, where $|S| = n$.

Proof. Let $k \geq 2$ and $j = k-2$. Then, there exists an edge $t = uv \in MST_{\mathcal{C}}(S)$ that is not in T_j. Thus, by Proposition 3.1.4 there exists a constraint $c = ab$ in T_j such that $\mathcal{C}(a,b)$ contains either u or v, without loss of generality say v. Hence, by Lemma 3.1.6 the edge $c_j = a_jb_j$ closest to v has the property that $C_j = \mathcal{C}(a_j,b_j)$ contains v. Then, v is not separated by an edge in T from a_j and b_j in C_j. Thus, by Lemma 3.1.3 there exists an a_jb_j-path that goes through v contained in C_j. Hence, there exists a point p_j in C_j that is visible to both a_j and b_j such that $\triangle_j = \triangle a_jp_jb_j \in DG_{\mathcal{C}}(T_j)$. Since $a_jp_jb_j$ is an a_jb_j-path that consists of edges solely shorter than c_j, the point p_j had to be separated by an edge $c' = a'b'$ of T_{j-1} for every $\mathcal{C}(a_j,b_j)$. But c' cannot cross c_j, since $c_j \in DG_{\mathcal{C}}(T_{j-1})$. See Figure 3.5. Note that there exists a $\mathcal{C}(a',b')$ that contains p_j in its interior: For the sake of a contradiction suppose that p_j is not contained in any $\mathcal{C}(a',b')$. We note that c' partitions C_j into two parts, one of which, by Proposition 3.1.4, is contained in any $\mathcal{C}(a',b')$, that is the one that does not contain p_j. Since, c' separates p_j from either a_j of b_j, c_j is contained in any $\mathcal{C}(a',b')$. Now, consider the intersecting points z and w in $c' \cap \partial C_j$ (note that at most one of the following can happen: $a_j = z = a'$ or $b_j = w = b'$). Thus, $d_{\hat{\mathcal{C}}}(z,w) \leq d_{\hat{\mathcal{C}}}(a_j,b_j)$. Also, by Observation 3.1.5, for each $C = \mathcal{C}(a',b')$ there is a $C_{zw} = \mathcal{C}(z,w)$ contained in C. Then, by Proposition 3.1.4, C_{zw} contains c_j. Therefore, $d_{\hat{\mathcal{C}}}(a_j,b_j) = d_{\hat{\mathcal{C}}}(z,w)$. Thus, C_j is a $\mathcal{C}(z,w)$. Hence, $\mathcal{C}(z,w)$ contains p_j, a contradiction.

Consider the closest edge $c_{j-1} = a_{j-1}b_{j-1}$ to p_j that separates p_j from a_j or b_j in C_{j-1}. Then, there exists a $\mathcal{C}(a_{j-1},b_{j-1})$ that contains p_j. Moreover, there is no edge in T_{j-1} that can separate p_j from a_{j-1} and b_{j-1} in C_{j-1}, otherwise, c_j wouldn't be the closest edge to p_j crossing p_ja_j and C_j in T_{j-1}. Therefore, there exists a p_{j-1} in C_{j-1}, such that $\triangle_{j-1} = \triangle a_{j-1}p_{j-1}b_{j-1}$ is in $DT_{\mathcal{C}}(T_{j-1})$. Consequently, we obtained a sequence $\triangle_j, \triangle_{j-1}, \ldots \triangle_0$ of empty triangles.

Observe that for each point p_i the edge c_j blocks the visibility to each $c_{i+1}, \ldots c_j$. Thus, we can map each index i in $\{0,\ldots,j-1\}$ to an endpoint of edge c_{i+1} that is different from the one in c_i (they can share at most one endpoint). Since p_0 and c_0 do not participate, then $j \leq n-3$. Therefore, $k \leq n-1$. □

3.3 Bounds for square spanning trees

In this section we study the canonical sequence when $\mathcal{C}$ is an axis-aligned square.

We denote the x- and y-coordinate of a point p by p_x and p_y, respectively. An important property of the $\square$-distance is that it is the same as the L_∞-metric, i.e., for any pair $p, q \in S$, $d_\square(p,q) = \max\{|p_x - q_x|, |p_y - q_y|\}$. In particular, any $\mathcal{C}(p,q)$ contains p and q on opposite sides of $\partial\mathcal{C}(p,q)$.

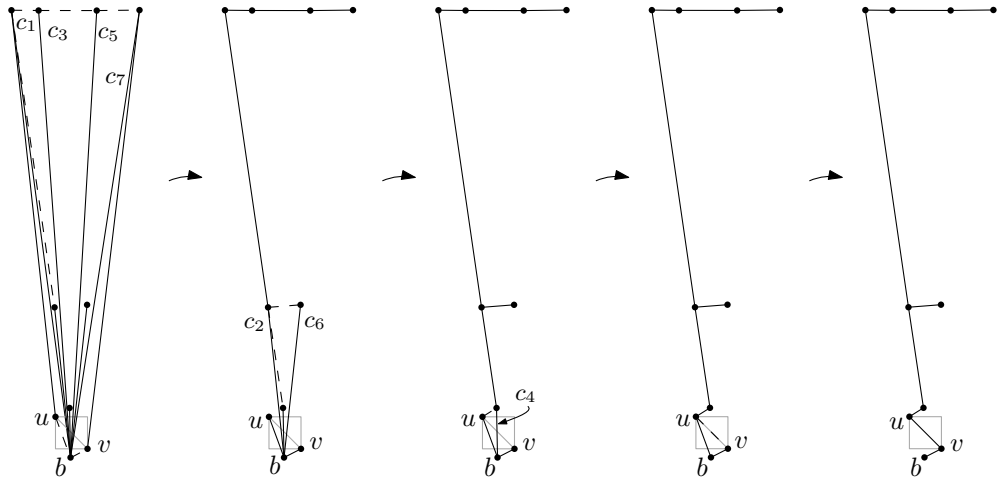

Figure 3.6: Example of a sequence with a spanning tree of a 10-point set that converges to $MST_\square(S)$ in 4 steps. The dashed edges represent the appearing edges at stage $i+1$.

3.3.1 A lower bound

In this section we give a lower bound on the length of the canonical sequence when $\mathcal{C}$ is a square.

A natural question, once we know that this sequence converges, is how fast it reaches the $MST_\square(S)$. As a first step for answering this question we give a lower bound based on a construction shown in Figure 3.6, similar to the one given in [5], which has length $\Theta(\log n)$.

Consider points u and b such that the edge $t = uv$ is a diagonal of the unique square $Q_t = \mathcal{C}(u,v)$. Let $n = 2^m + 2$ for some $m \in \mathbb{N}$. Let b be a point below Q_t. Consider points $a_1, a_2, \ldots a_{n-3}$ above Q_t such that each a_j has the following properties: (1) the line segment $c_j = a_j b$ crosses the top and bottom sides of Q_t, (b) c_j is closer to u than c_s for all $j < s$, and (c) if j is odd, then c_j has length ℓ, otherwise, c_j has length $\frac{\ell}{3^{i-1}}$ where $j \equiv 2i$ modulo 2^{i+1} for

some $i \in \{2, \ldots, m\}$. Refer to Figure 3.6. Let $S := \{u, v, b, a_1 \ldots, a_{n-3}\}$. By construction, $t \in MST_{\square}(S)$.

Now, consider the spanning tree T of S with edges $c_1, \ldots c_{n-3}, ub$ and vb. Then, t is crossed by all the constraints $c_1, \ldots, c_{n-3}$. Notice that at step $i-1$, the $d_{\square}$-longest constraints of T_{i-1} are the constraints with length $\frac{\ell}{3^{i-1}}$ and these constraints are the only ones of T_{i-1} not present at T_i for $1 \leq i < m$. Indeed, let $c_k = ab$ be a constraint crossing t in T_{i-1} where $k \not\equiv 2i$ modulo 2^{i+1}, then $d_{\square}(a_i, b) \leq \frac{\ell}{3^{i-1}}$, where $c_i = a_i b$ for all $1 \leq i < k$. Thus, any $a_i b$-path different from edge $a_i b$ would contain an edge with length $\frac{2\ell}{3^{i-1}}$, refer to Figure 3.7. Also, the constraints c_j where $j \equiv 2i$ modulo 2^{i+1} disappear at stage $i-1$, since there exists a path between c_j's endpoints with edges shorter than c_j, refer to Figure 3.7. Hence the tree undergoes $m + 1 = \log_2(n-2) + 1$ iterations.

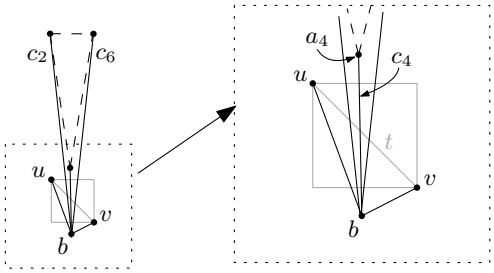

Figure 3.7: The dashed edges have $d_{\square}$-length $\frac{2\ell}{3^{i-1}}$, while each of the $a_4 b$-paths different from $a_4 b$ passes through one of the endpoints of c_2 or c_6 that are different from b.

Theorem 3.3.1. There exists a point set S and $T \in \mathcal{ST}(S)$, such that its canonical sequence has length $\Theta(\log n)$.

3.3.2 An upper bound

In this section we show a special case when the upper bound of $O(\log n)$ on the length of the canonical sequence when $\mathcal{C}$ is a square.

Note that if we show that an arbitrary edge of $MST_{\square}(S)$ appears in $O(\log n)$ steps, then the sequence of $d_{\square}$-spanning trees has length $O(\log n)$. In this section we will show that an arbitrary edge uv of the $MST_{\square}(S)$ that is only crossed by constraints that cross opposite sides of a square $\partial\mathcal{C}(u, v)$, then uv appears in a tree of the canonical sequence in at most $O(\log n)$ steps.

Let $t = uv$ be an edge of the $MST_{\square}(S)$ crossed by a set of constraints. Let Q_t be a $\mathcal{C}(u, v)$. Without loss of generality we assume that $d_{\square}(u, v) = |u_x - v_x|$ and that u and v lie on the right and left side of Q_t, respectively. Let e be an edge crossing t. We say that e is a *diagonal* when e crosses t in consecutive sides

of ∂Q_t, or call it *vertical* and *horizontal* if e crosses opposite sides that are top- and bottom-sides or left- and right-sides of ∂Q_t, respectively. In this section, we show that if an edge $t \in MST_{\square}(S)$ is only crossed by vertical constraints, then t appears in $O(\log n)$ steps. The case when the constraints are horizontal is symmetric.

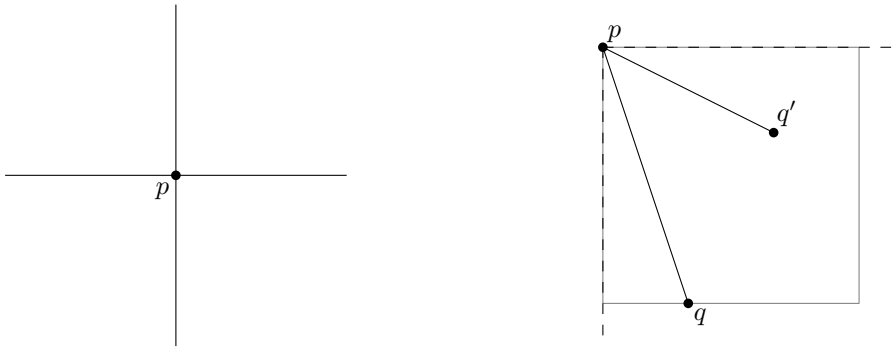

Figure 3.8: Left: quadrants at p. Right: both q and q' are in the same quadrant at p. Point q' is in $\square(p, q)$.

Recall that the quadrants at a point p of S are given by the four quadrants defined by the lines $x = p_x$ and $y = q_y$. See Figure 3.8(left). Consider two edges pq and pq' such that $d_{\square}(p, q') < d_{\square}(p, q)$. Observe that if q and q' are in the same quadrant at p, then q' lies in a $\square(p, q)$: since the $Q = \square(p, q)$ with vertex p lies in the same quadrant at p as q and q'. Thus, since $d_{\square}(p, q') < d_{\square}(p, q)$, it follows that q' lies in Q. Therefore, $d_{\square}(q, q') < d_{\square}(p, q)$. See Fig. 3.8(right).

Observation 3.3.2. Consider edges pq and pq' incident to p. If q and q' are in the same quadrant at p, then $d_{\square}(q, q') < \max\{d_{\square}(p, q), d_{\square}(p, q')\}$.

Consider a vertical edge $e = pq$ and the intersection points z, w in $e \cap \partial Q_t$. We have that $d_{\square}(z, w) = |z_y - w_y| = d_{\square}(u, v)$. Thus, the rectangle R_{zw} with z and w as opposite vertices has height greater or equal than its width. The rectangle R_{zw} is a homothet of the rectangle with opposite vertices p and q. Therefore, $d_{\square}(p, q) = |p_y - q_y|$.

Observation 3.3.3. If pq is a vertical edge in $DG_{\square}(T)$, then $d_{\square}(p, q)=|p_y-q_y|$.

Moreover, there exists a $\mathcal{C}(z, w)$ that contains both u and v. Hence, by Observation 3.1.5 there exists a $\mathcal{C}(p, q)$ that contains both u and v. Thus, the following observation holds.

Observation 3.3.4. If $e = pq$ is a vertical edge crossing $t = uv \in MST_{\square}(S)$, then $d_{\square}(p, v) < d_{\square}(p, q)$ and $d_{\square}(v, q) < d_{\square}(p, q)$.

Next, we observe that if two edges e and e' are vertical edges crossing t, then

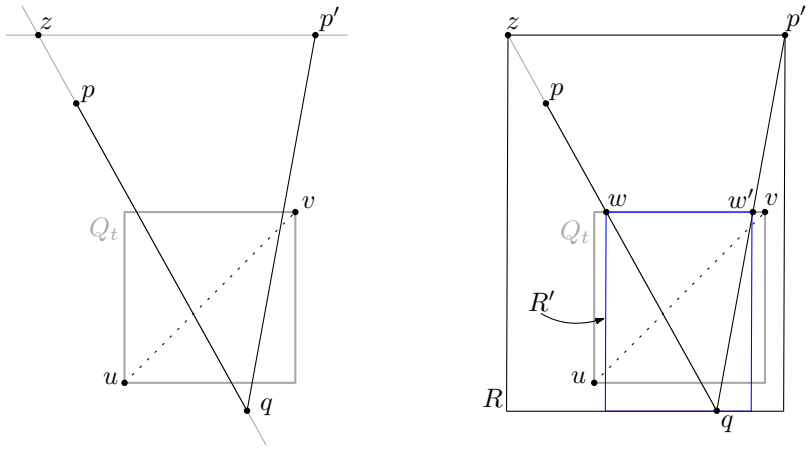

Figure 3.9: (a) Point z lies in the supporting line of pq and $z_y{=}p'_y$. (b) Rectangle R has vertices z and p' on its top-side and q lies on its boundary.

the $d_\square$ distance of the two endpoints below Q_t is shorter than the $d_\square$-length of the largest of e and e'.

Observation 3.3.5. Let $e = pq$ and $e' = p'q'$ be two vertical edges crossing Q_t with the next properties: (a) $p_y > q_y$ and $p'_y > q'_y$, and, (b) $d_\square(p, q) < d_\square(p', q')$. Then, the following hold:

1. If $p \neq p'$, then $d_\square(p, p') < d_\square(p', q')$.
2. If $q \neq q'$, then $d_\square(q, q') < d_\square(p', q')$.

Proof. Since $d_\square(p, q) < d_\square(p', q')$, at least one of $p \neq p'$ and $q \neq q'$ happens. Since both cases are symmetric, without loss of generality we assume that $p \neq p'$. Note that the edges $p'q$ and pq' are also vertical edges, since they are either edges (if $q = q'$) or diagonals of the polygon $pqq'p'$. Also, note that one of $p'q$ or pq' has smaller $d_\square$-length than $p'q'$, since either $q_y > q'_y$ or $p_y < p'_y$, otherwise $p'q'$ has smaller $d_\square$-length than pq, which is a contradiction. Consider the edge e'' from $\{p'q, pq'\}$ that has smaller $d_\square$-length than $p'q'$. Without loss of generality assume that $e'' = p'q$ and $p_y < p'_y$. Observe that if p and p' belong to the same quadrant at q, then by Observation 3.3.2 the statement holds. Assume that p and p' are in different quadrants at q. Consider the point z on the supporting line of e with $z_y = z'_y$. Let R rectangle be the rectangle that contains q on its boundary and the points z and p' are vertices of R. See Figure 3.9. Rectangle R has height $|p'_y - q_y| = |z'_y - q_y|$ and width $|p'_x - z_x| \geq |p'_x - p_x|$. Consider the intersection points w of the top side of Q_t and pq, and the intersection point w' of top side of Q_t and $p'q$. Consider the rectangle R' that has w and w' as vertices and contains

q on its bottom side. See Fig. 3.9b. Then, R' has width $|w_x - w'_x| \leq d_\square(u, v)$ and height greater than $d_\square(u, v)$. Thus, $|q_y - w'_x| < |w_x - w'_x|$. The rectangle R' is a homothet of R. Hence, the height of R is greater than its width. Therefore, $d_\square(p, p') < |p'_y - q_y| = d_\square(p', q) \leq d_\square(p', q')$. □

Let $C_i := \{c_1^i, c_2^i, \dots c_m^i\}$ be the set of vertical constraints in T_i crossing t sorted in the following fashion: $d(c_j^i \cap t, u) < d(c_{j+1}^i \cap t, u)$ for all $1 \leq j < m$.

Observe that c_1^i and c_m^i are the closest edges to u and v, respectively. In addition, $m \leq n - 3$, since the edges in T_i with endpoints u and v do not cross uv and T_i has $n - 1$ edges. Let $c_j^i \in C_i$, we denote $c_j^i = a_j b_j$ where $a_{j_y} > b_{j_y}$ for all j, i.e., a_j is above b_j. Let s_t, s_b, s_r, s_l denote the top-side, bottom-side, right-side and left-side of ∂Q_t, respectively.

Consider two consecutive constraints c_j^i and c_{j+1}^i. Let z be the intersection point $c_j^i \cap s_t$ and z' be the intersection point $c_{j+1}^i \cap s_t$. Let P be the polygon $a_j z z' a_{j+1}$. Since the line segments $a_j z, zz'$ and $z' a_{j+1}$ are contained in c_i^j, s_t and c_i^{j+1}, respectively, the constraints crossing P are only crossing the segment $a_j a_{j+1}$, we refer to Figure 3.10a. Consider the convex hull CH_{a_j} of all the points of S contained in P. Notice that if $a_j \neq a_{j+1}$, then $a_j a_{j+1}$ is an edge of CH_{a_j}. If CH_{a_j} is different than a point (i.e., $CH_{a_j} \neq a_j$ when $a_j = a_{j+1}$), then we denote by ch_{a_j} the longest $a_j a_{j+1}$-path in CH_{a_j} (note that CH_{a_j} can be an edge $a_j a_{j+1}$). Similarly, we define a chain between b_j and b_{j+1} denoted ch_{b_j}. See Figure 3.10a.

Consider the intersection points $c_j^i \cap t$ and $c_{j+1}^i \cap t$, we say that an edge e is between c_j^i and c_{j+1}^i if e crosses t in the line segment $(c_j^i \cap t)(c_{j+1}^i \cap t)$. See Figure 3.10b.

The following lemma shows that for every two consecutive constraints c_j^i and c_{j+1}^i constraints in C_i, then at most one vertical constraint appears in T_{i+1} between c_j^i and c_{j+1}^i (possibly one of c_j^i and c_{j+1}^i).

Lemma 3.3.6. Consider two consecutive vertical constraints c_j^i and c_{j+1}^i crossing t and consider the chains ch_{a_j} and ch_{b_j}. Let a be the bottom-most point of ch_{a_j} and b the top-most point of ch_{b_j}. If there is a constraint $c \in C_{i+1}$ between c_j^i and c_{j+1}^i, then such a constraint is the only constraint of C_{i+1} that appears between c_i^j and c_i^{j+1}. Moreover, $c = ab$.

Proof. By Observation 3.3.3 the length of any vertical constraint is given by the y-coordinates of its endpoints, then ab is the shortest possible edge between c_j^i and c_{j+1}^i. Let us show that the edge $c = ab$ crossing t is the only possible constraint of T_{i+1} between c_j^i and c_{j+1}^i. For the sake of a contradiction suppose that there exists a constraint $c' = a'b' \in T_{i+1}$ with $d_\square(a', b') > d_\square(a, b)$ that is between c_j^i and c_{j+1}^i. Consider squares $Q_a = \mathcal{C}(a, a')$ and $Q_b = \mathcal{C}(b, b')$. Note

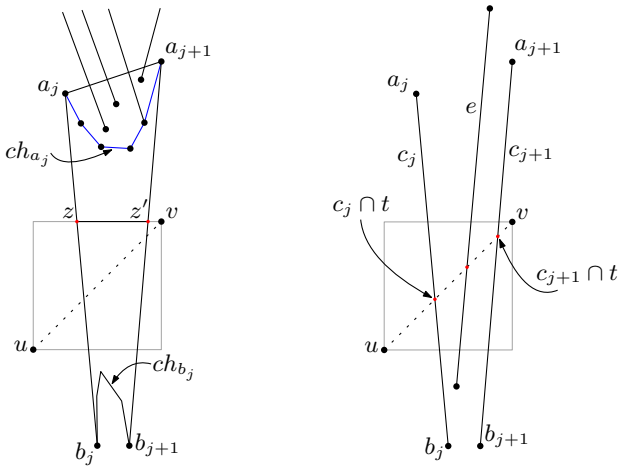

Figure 3.10: (a) The polygon $a_j z z' a_{j+1}$ is only crossed by constraints in the line segment $a_j a_{j+1}$. The chain ch_{a_j} is defined by the $a_j a_{j+1}$-path different from $a_j a_{j+1}$ in the convex hull of the points of S that lie in the polygon $a_j z z' a_{j+1}$. (b) Edge e is between c_j and c_{j+1}.

that the bottom side of ∂Q_a lies at most at the height of a_y and the top side of ∂Q_b lies at least at the height of b_y. In addition, by Observation 3.3.5 the square Q_a has side-length less than $\mathcal{C}(a', b')$. Similarly, Q_b has side-length less than $\mathcal{C}(a', b')$. Since a defines the bottom-most point of the chain ch_{a_j}, there is no edge in T_i that separates a from a' in Q_a. Similarly, since b is the top-most point in the chain ch_{b_j}, there does not exist a constraint in T_i that separates b from b' in Q_b. Thus, by Lemma 3.1.3 there exists an $a'a$-path P_a and a bb'-path P_b in $DG_{\square}(T_i)$ contained in Q_a and Q_b, respectively. Therefore, there is an $a'b'$-path in $P_a \cup \{ab\} \cup P_b$ that solely consists of edges shorter than $a'b'$, contradicting Property 3.1.1. □

In the following lemma we show that if there appears a constraint c of C_{i+1} between two consecutive constraints c^i_j and c^i_{j+1} in C_i such that c is neither c^i_j or c^i_{j+1}, or adjacent to both c^i_j and c^i_{j+1}, then there must have been two constraints in C_{i-1} between c^i_j and c^i_{j+1}. See Figure 3.11.

Lemma 3.3.7. Let $i \geq 1$. If a constraint $c = ab$ in C_{i+1} is between $c^i_j = a_j b_j$ and $c^i_{j+1} = a_{j+1} b_{j+1}$, and at least one of a and b is not in $\{a_j, a_{j+1}, b_j, b_{j+1}\}$, then there are at least two constraints in C_{i-1} that are between c^i_j and c^i_{j+1}.

Proof. Assume without loss of generality that a is not in $\{a_j, a_{j+1}, b_j, b_{j+1}\}$.

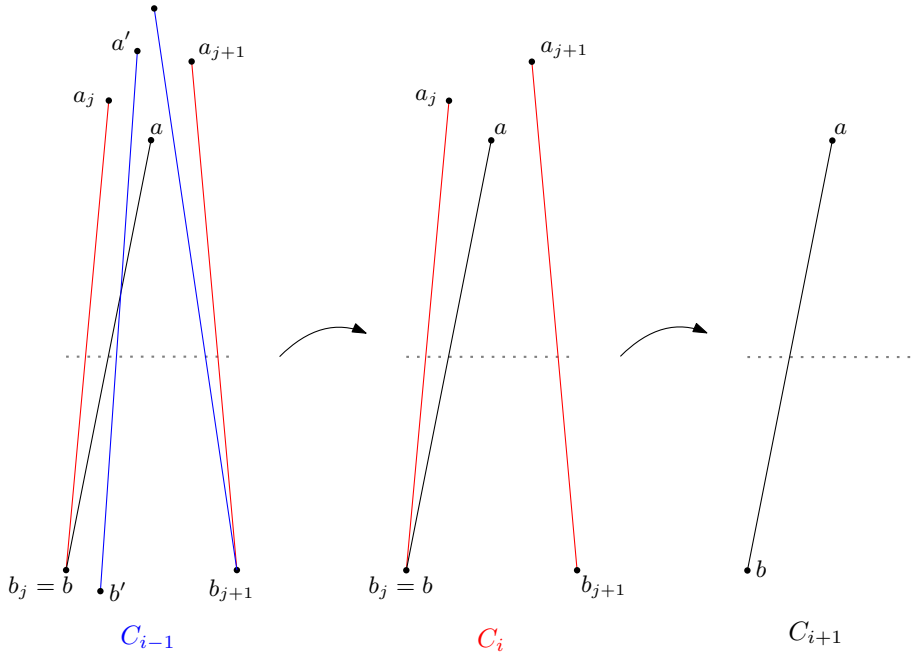

Figure 3.11: There are two blue vertical constraints in T_{i-1} that are between two consecutive red constraints $c^i_j = a_j b_j$ and $c^i_{j+1} = a_{j+1} b_{j+1}$ in T_i. The constraint $c = ab$ is in C_{i+1} and $a \notin \{a_j, a_{j+1}\}$.

Then, there must be an edge $c' = a'b'$ blocking a from b, otherwise ab is visible in T_{i-1}. Thus by Property 3.1.1 there exists an ab-path P that solely consists of edges shorter than ab in T_i. Hence, P is in $DG_\square(T_i)$, by Property 3.1.1 we get a contradiction.

In addition, c' has to be in C_{i-1}, otherwise c' has an endpoint with either y-coordinate less than a_y or greater than b_y, contradicting Lemma 3.3.6. See Figure 3.12a.

Observe that there exists a constraint c'' in T_{i-1} that blocks the visibility from a to either a_j and b_j, and a different constraint c''' in T_{i-1} that blocks the visibility from a to either a_{j+1} and b_{j+1}: suppose for the sake of a contradiction that a is visible to either b_j or a_j. Consider the consecutive constraints $c^{i-1}_{j'}$ and $c^{i-1}_{j'+1}$ in C_{i-1} such that c^i_j appears in between $c^{i-1}_{j'}$ and $c^{i-1}_{j'+1}$. We consider whether a_j is visible to a. If a_j is visible to a, then $a_j = a_{j'+1}$, otherwise there exists a point p in $ch_{a_{j'}}$ with $p_y < a_{j_y}$, which contradicts Lemma 3.3.6. Then, a must be blocked from b_j by a constraint in C_{i-1}, otherwise there is a square $\mathcal{C}(a, b_j)$ with side-length less than $d_\square(a_j, b_j)$ such that a and b_j are not separated by a constraint in T_{i-1}. Hence, by Lemma. 3.1.3 there exists an $a_j a$-path P_{a_j}

and an ab_j-path P_{b_j} in $DG_\square(T_{i-1})$ such that $P_{a_j} \cup P_{b_j}$ solely consists of shorter edges than c^i_j, contradicting Property 3.1.1. Analogously we prove that there exists another constraint blocking a_l from a_{j+1} or b_{j+1}. Note that one of the constraints c'' and c''' can be the constraint c' that is blocking the visibility from a_l to b_l. □

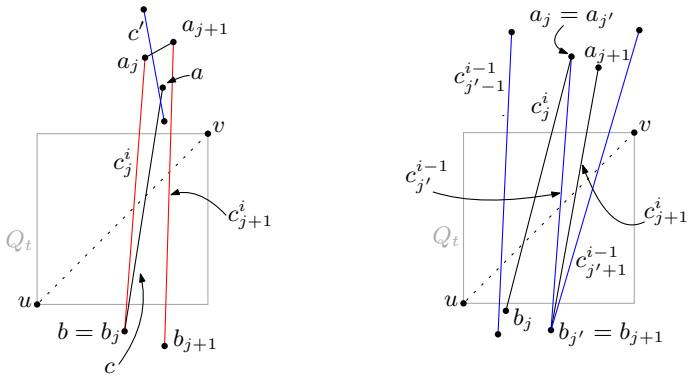

Figure 3.12: Forbidden configurations: (a) Edge c' is in T_{i-1} and $c \in C_{i+1}$. Edge c' blocks the visibility from a to $b = b_{j'}$ and one of the endpoints of c' has lower y-coordinate than a. (b) The consecutive constraints c^i_j and c^i_{j+1} are not incident but both edges are incident to $c^{i-1}_{j'}$. Constraint $c^{i-1}_{j'}$ is the only constraint of C_{i-1} between c^i_j and c^i_{j+1}.

We show that if there is exactly one constraint $c^{i-1}_{j'}$ in C_{i-1} between two consecutive constraints c^i_j and c^i_{j+1} in C_i and $c^{i-1}_{j'}$ is incident to both c^i_j and c^i_{j+1}, then c^i_j and c^i_{j+1} are incident as well.

Proposition 3.3.8. Let c^i_j and c^i_{j+1} consecutive constraints in C_i such that exactly one edge $c^{i-1}_{j'}$ of C_{i-1} is in between c^i_j and c^i_{j+1}. If c^i_j and c^i_{j+1} are both incident to $c^i_{j'}$, then c^i_j and c^i_{j+1} share a vertex.

Proof. Without loss of generality assume that $d_\square(a_j, b_j) < d_\square(a_{j+1}, b_{j+1})$. Assume for the sake of a contradiction that $a_{j'} = a_j$ and $b_{j'} = b_j$. See Figure 3.12b. Since there is exactly one edge of C_{i-1} between c^i_j and c^i_{j+1}, the edge c^i_j appears between $c^{i-1}_{j'-1}$ and $c^{i-1}_{j'}$, and the edge c^i_{j+1} appears between $c^{i-1}_{j'}$ and $c^{i-1}_{j'+1}$. On the other hand, by Observation 3.3.5 $d_\square(a_j, a_{j+1}) < d_\square(a_{j+1}, b_{j+1})$. Consider $Q_a = \mathcal{C}(a_j, a_{j+1})$ and $Q_b = \mathcal{C}(b_j, b_{j+1})$. By Lemma 3.3.7 a_{j+1} has the lowest y-coordinate in $ch_{a_{j'}}$ and b_j has the largest y-coordinate in $ch_{b_{j'-1}}$. Hence, there is no edge separating a_j from a_{j+1} in Q_a and there is no edge separating

b_j from b_{j+1} in Q_b. Then, by Lemma 3.1.3, there exists an a_ja_{j+1}-path P_a of $DG_{\square}(T_{i-1})$ in Q_a and a b_jb_{j+1}-path P_b of $DG_{\square}(T_{i-1})$ in Q_b. Therefore, there is an $a_{j+1}b_{j+1}$-path in $P_a \cup a_jb_j \cup P_b$ that solely consists of edges shorter than c^i_{j+1}, which contradicts Property. 3.1.1. □

Note that if for each pair of consecutive constraints c^{i+1}_j and c^{i+1}_{j+1} in C_{i+1} there are at least two constraints in C_i between c^{i+1}_j and c^{i+1}_{j+1}, then we obtain at most $\frac{|C_i|}{2}$ constraints in C_{i+1} that appear between consecutive pairs of constraints in C_i. Hence, $|C_{i+1}| \leq \frac{|C_i|}{2}$.

Observation 3.3.9. If for each pair of consecutive constraints in C_{i+1} there are at least two constraints C_i in between them, then $|C_{i+1}| \leq \frac{|C_i|}{2}$.

Also note that there cannot appear any constraint $c = ab$ in T_{i+1} crossing t such that c crosses t in either the line segment $u(c^i_1 \cap t)$ or $(c^i_m \cap t)v$ where $|C_i| = m$. Otherwise, without loss of generality we assume c crosses the line segment $u(c^i_1 \cap t)$. Then, by similar arguments to those in Theorem 3.2.1, there exist an au-path P_a and an ub-path P_b such that $P_a \cup P_b$ contains an ab-path that solely consists of shorter edges than c, which contradicts Property 3.1.1. Therefore, we only need to focus on the edges that appear between pairs of consecutive constraints in C_i.

Proposition 3.3.10. If for each pair of consecutive constraints in C_i there appears exactly one edge in C_{i+1} between them such that $c^{i+1}_1 \cup c^{i+1}_2 \cup \ldots \cup c^{i+1}_{m-1}$ defines a path P in T_{i+1} with $|P| \geq 2$. Then, $|C_{i+2}| \leq \frac{|C_i|}{2}$.

Proof. Since C_{i+1} consists of edges in a path P and exactly one edge appear per pair of consecutive constraints in C_i, it follows that $|C_{i+1}| = |C_i| - 1$. Since for each consecutive pair of constraints c^{i+1}_j and c^{i+1}_{j+1} there was exactly one constraint in C_i between c^{i+1}_j and c^{i+1}_{j+1}, then if there appears one edge of C_{i+2} in between them, from Lemma 3.3.7 it follows that such edge has to have both endpoints in $\{a_j, b_j, a_{j+1}, b_{j+1}\}$. Since c^{i+1}_j and c^{i+1}_{j+1} are incident, then the new edge has to be exactly one of c^{i+1}_j and c^{i+1}_{j+1}. Thus, $|C_{i+2}| \leq \frac{|C_i|}{2}$. □

We say that a set of consecutive pairs of constraints in C_i is of *type 1* if it satisfies the hypothesis from Observation 3.3.9. We say that a set of consecutive pairs of constraints in C_i is of *type 2* if it satisfies the hypothesis from Proposition 3.3.10.

Theorem 3.3.11. $|C_{i+3}| \leq \frac{2}{3}|C_i|$.

Proof. Let k_1 be the number of consecutive pairs of constraints in C_i in a set of type 1. Let k_2 be the number of consecutive pairs of constraints in C_i in a

set of type 2. Let k_3 the number of remaining consecutive pairs of constraints in C_i. Notice that $|C_i| = k_3 + k_1 + k_2 + 1$.

Thus, by Obs 3.3.9 and Proposition 3.3.10 $|C_{i+2}| \leq k_3 + \frac{k_1+k_2}{2}$.

Now, let us count the remaining constraints. So, we get k_3 different constraints of C_{i+1} in each of the remaining consecutive pairs of constraints in C_i that are not in sets of type 1 and 2. Let c_j^{i+1} and c_{j+1}^{i+1} be a pair of these types of pairs. Since there was exactly one edge from C_i in between c_j^{i+1} and c_{j+1}^{i+1}, by Lemma 3.3.7 if there is one constraint c from C_{i+2} in between c_j^{i+1} and c_{j+1}^{i+1}, then either one of these happens: (1) c is one of c_j^{i+1} and c_{j+1}^{i+1} or (2) c shares a vertex with both c_j^{i+1} and c_{j+1}^{i+1}. If a pair of consecutive constraints in C_{i+2} are disjoint such that both are in case (1), we will call these pairs *type 3-1*, and if one of the pair of consecutive constraints is of the case (1) and the other is of case (2), then we will call these pairs *type 3-2*. See Figure 3.13 for a complete example.

Hence, by Proposition 3.3.8 we get a set of paths and disjoint constraints in C_{i+2}. So, we have a new set of type 1 and 2 consecutive pairs of constraints in C_{i+2} and a set of consecutive pairs of constraints that are disjoint of type 3-1 and 3-2 in C_{i+2}. Let k_3'' denote the number of consecutive constraints of type 1 and 2 in C_{i+1} and let k'_{3-1} and k'_{3-2} the number of pairs of consecutive constraints in C_{i+2} of type $3-1$ and $3-2$, respectively. Note that for each pair of type $3-1$ there is a partition of a subset of pairs in C_{i+1} where 4 consecutive pairs of constraints contain exactly one pair of consecutive constraints from C_{i+2} (the pair of type 3-1). Refer to Figure 3.13. Also, for each pair of type 3-2 there is a partition of the pairs of consecutive constraints of C_{i+1} where 3 consecutive pairs contain exactly one pair of consecutive constraints from C_{i+2} (the pair of type 3-2 in C_{i+2}). See Figure 3.13. Thus, we obtained that $k_3 = 2k'_{3-1} + \frac{3}{2}k'_{3-2} + k''_3$.

Therefore, by Proposition 3.3.10 and Observation 3.3.9 this forms at most $k'_{3-1} + k'_{3-2} + \frac{1}{2}k''_3 \leq \frac{2}{3}k_3$ new constraints in C_{i+3} from the k_3 remaining pairs to count. This gives us,

$$|C_{i+3}| \leq \frac{2}{3}(k_3) + \frac{1}{2}(k_1 + k_2) \leq \frac{2|C_i|}{3} - 1$$

□

As a consequence of Theorem 3.3.11 we obtain that there are a constant number of edges crossing t after $i = O(\log n)$ steps. From Lemma 3.3.6 we have that $|C_{i-1}| \leq |C_i| - 1$. Thus, we obtain the following corollary.

Corollary 3.3.12. Let $|S| = n$. Let $T \in \mathcal{ST}(S)$ such that T crosses the edges from $MST_{\square}(S)$ by only vertical constraints. Then, the length of the canonical sequence of T is $O(\log n)$.

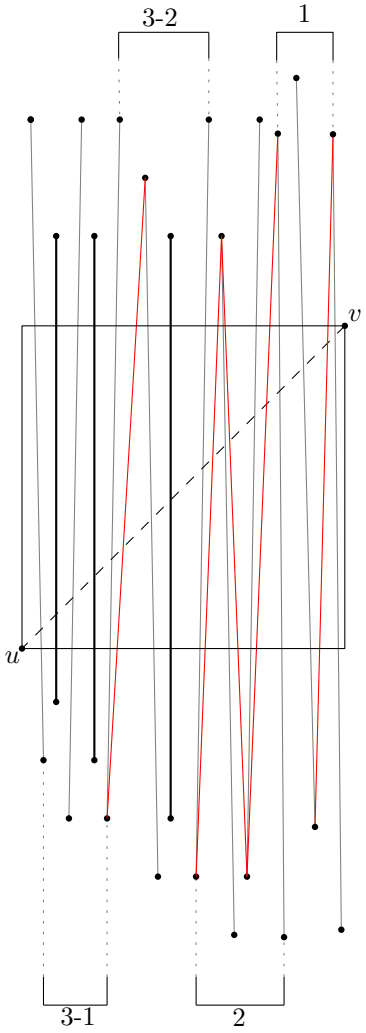

Figure 3.13: The gray edges represent the constraints in C_{i+1} that do not appear in C_{i+2} and the black edges are constraints in $C_{i+1} \cap C_{i+2}$. The red edges represent the constraints in C_{i+1} that share vertices with two different consecutive constraints from C_{i+1}. Here, we make a partition of pairs of two consecutive constraints in C_{i+1} by pair of consecutive constraints of type 1, 2, 3-1 and 3-2. Note that the number of edges of type 2 drops by half in the next iteration, C_{i+3}.

3.4 Conclusions

For an arbitrary convex $\mathcal{C}$, we have extended the convergence of the canonical sequence of crossing-free spanning trees for the Euclidean metric [5] to the $\mathcal{C}$-distance. We have also given a linear upper bound for the length of the canonical sequence. For the particular case when the convex is a square, we give a $\Omega(\log n)$ and an upper bound of $O(\log n)$ where for each edge of the $d_{\square}$-minimum spanning tree, the constraints in the initial tree only cross opposite sides of at least one of the smallest squares containing such edge.

In order to show this for any tree in the $d_{\square}$-distance we considered both cases, when the constraints cross opposite sides of a smallest square containing an edge from $MST_{\square}(S)$ or the constraints cross consecutive sides of such square. When dealing with constraints crossing consecutive sides, most of the ingredients for the proof of Theorem 3.3.11 can be generalized with the exception of Lemma 3.3.7. So, the main open question is how to close the gap between $O(n)$ and $\Omega(\log n)$ for $d_{\square}$-distances. We believe that the right answer should be $O(\log n)$. However, it seems that a different approach is needed. Perhaps something using the constrained $RNG_{\square}$, similar to the upper bound of $O(\log n)$ for Euclidean distance given in [5], yet it is not clear how to use their approach since they use a lot of properties from circles that do not hold for squares.

An even more general open problem is whether there exists a convex shape where the length of the canonical sequence is actually linear.

Convex shape Delaunay graphs and Hamiltonicity

4

Shamos [93] conjectured that the Delaunay triangulation contains a Hamiltonian cycle. This conjecture sparked a flurry of research activity. Although Dillencourt [48] disproved this conjecture, he showed that Delaunay triangulations are *almost* Hamiltonian [49], that is, they are 1-tough.[1] Focus then shifted on determining how much to loosen the definition of the Delaunay triangulation to achieve Hamiltonicity. Chang et al. [39] showed that 19-RNG is Hamiltonian.[2] Abellanas et al. [2] proved that 15-GG is Hamiltonian. Currently, the lowest known upper bound is by Kaiser et al. [67] who showed that 10-GG is Hamiltonian. All of these results are obtained by studying properties of *bottleneck Hamiltonian cycles*. Given a planar point set, a bottleneck Hamiltonian cycle is a Hamiltonian cycle whose maximum edge length is minimum among all Hamiltonian cycles of the point set. Biniaz et al. [22] showed that there exist point sets such that its 7-GG does not contain a bottleneck Hamiltonian cycle, implying that this approach cannot yield an upper bound lower than 8. Despite this, it is conjectured that 1-DG is Hamiltonian [2].

As for Hamiltonicity in convex shape Delaunay graphs, not much is known. Bonichon et al. [24] proved that every plane triangulation is Delaunay-realizable where homothets of a triangle act as the empty convex shape. This implies that there exist *triangle-DG* graphs that do not contain Hamiltonian paths or cycles. Biniaz et al. [21] showed that 7-$DG_{\triangle}$ contains a bottleneck Hamiltonian cycle and that there exist points sets where 5-$DG_{\triangle}$ does not contain a bottleneck Hamiltonian cycle. Ábrego et al. [3] showed that the $DG_{\square}$ admits a Hamiltonian path, while Saumell [92] showed that the $DG_{\square}$ is not necessarily 1-tough, and

[1] Recall that a graph G is *1-tough* if removing any k vertices from G results in $\leq k$ connected components.

[2] According to the definition of k-RNG in [39], they showed Hamiltonicity for 20-RNG.

Type of shape $\mathcal{C}$	$k \leq$	$k \geq$	Bottleneck-$k \geq$
Circles	10 [67]	1 [48]	8 [22]
Equilateral triangles	7 [21]	1 [21]	6 [21]
Squares	7 [Thm. 4.3.3]	1 [92]	3 [Lemma 4.4.1]
Regular hexagons	11 [Thm. 4.3.6]	1 [Lemma 4.5.1]	6 [Lemma 4.4.2]
Regular octagons	12 [Thm. 4.3.8]	1 [Lemma 4.5.1]	-
Regular t-gons (t even, $t \geq 10$)	11 [Thm. 4.3.7]	-	-
Regular t-gons ($t = 3m$ with m odd, $m \geq 3$)	24 [Thm. 4.1.7]	1[Thm. 4.5.2]	-
Point-symmetric convex	15 [Thm. 4.2.4]	-	-
Arbitrary convex	24 [Thm. 4.1.7]	-	-

Table 4.1: Bounds on the minimum k for which k-$DG_{\mathcal{C}}(S)$ is Hamiltonian and for which k-$GG_{\mathcal{C}}(S)$ contains a $d_{\mathcal{C}}-$bottleneck Hamiltonian cycle.

therefore does not necessarily contain a Hamiltonian cycle.

Our contributions. We generalize the above results by replacing the disk with an arbitrary convex shape $\mathcal{C}$. We show that the k-Gabriel graph, and hence also the k-Delaunay graph, is Hamiltonian for any convex shape $\mathcal{C}$ when $k \geq 24$. Furthermore, we give improved bounds for point-symmetric shapes, as well as for even-sided regular polygons. Table 4.1 summarizes the bounds obtained. Finally, we provide some lower bounds on the existence of a Hamiltonian cycle for an infinite family of regular polygons, and bottleneck Hamiltonian cycles for the particular cases of hexagons and squares. Together with the results of Bose et al. [33], our results are the first results on graph-theoretic properties of generalized Delaunay graphs that apply to arbitrary convex shapes.

Our results rely on the use of normed metrics and packing lemmas. In fact, in contrast to previous work on Hamiltonicity for generalized Delaunay graphs, our results are the first to use properties of normed metrics to obtain simple proofs for various convex shape Delaunay graphs.

4.1 Hamiltonicity for general convex shapes

In this section we show that the 24-order $\mathcal{C}$-Gabriel graph is Hamiltonian for any point set S in general position.

In this chapter we let S be a set of points in the plane satisfying the following general position assumption: For each pair $p, q \in S$, any minimum homothet of $\mathcal{C}$ having p and q on its boundary does not contain any other point of S on its boundary.

Let $\mathcal{H}$ be the set of all Hamiltoninan cycles of the point set S. Define the $d_{\hat{\mathcal{C}}}$-length sequence of $h \in \mathcal{H}$, denoted $ds_{\mathcal{C}}(h)$, as a sequence of edges of h

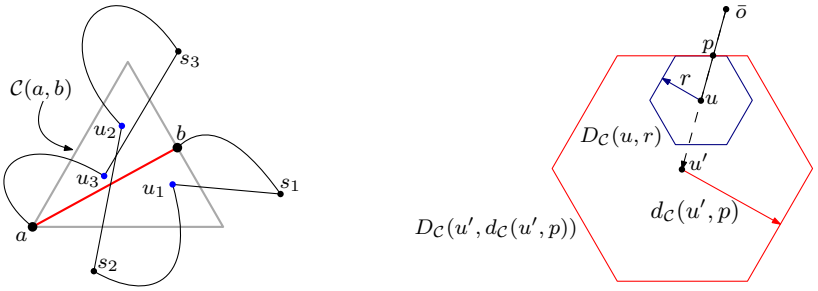

Figure 4.1: (a) Example of U in $\mathcal{C}(a,b)$. (b) $D_{\mathcal{C}}(u,r)$ is contained in $D_{\mathcal{C}}(u', d_{\mathcal{C}}(u',p))$ where $u' = \lambda u$ with $\lambda > 1$.

sorted in decreasing order with respect to the $d_{\hat{\mathcal{C}}}$-metric. Sort the elements of $\mathcal{H}$ in lexicographic order with respect to their $d_{\hat{\mathcal{C}}}$-length sequence, breaking ties arbitrarily. This order is strict. For $h_1, h_2 \in \mathcal{H}$, if h_1 is smaller than h_2 in this order, we write $h_1 \prec h_2$.

Let h be the minimum element in $\mathcal{H}$, often called *bottleneck Hamiltonian cycle*. The approach we follow to prove our bounds, which is similar to the approach in [2, 39, 67], is to show that h is contained in k-$GG_{\mathcal{C}}(S)$ for a small value of k. The strategy for proving that h is contained in 24-$GG_{\mathcal{C}}(S)$ is to show that for every edge $ab \in h$ there are at most 24 points in the interior of any $\mathcal{C}(a,b)$. In order to do this, we associate each point in the interior of an arbitrary fixed $\mathcal{C}(a,b)$ to another point. Later, we show that the $d_{\hat{\mathcal{C}}}$-distances between such associated points and a is at least $d_{\hat{\mathcal{C}}}(a,b)$. Finally, we use a packing argument to show that there are at most 24 associated points, which leads to a maximum of 24 points contained in $\mathcal{C}(a,b)$.

Let $ab \in h$; we assume without loss of generality that $d_{\hat{\mathcal{C}}}(a,b) = 1$. Let $U = \{u_1, u_2, \ldots, u_k\}$ be the set of points in S different from a and b that are in the interior of an arbitrary fixed $\mathcal{C}(a,b)$.[3] When traversing h from b to a, we visit the points of U in the order $u_1, \ldots, u_k$. For each point u_i, define s_i to be the point preceding u_i in h. See Figure 4.1a.

Note that if a point p is in the interior of $\mathcal{C}(a,b)$, then for any q on the boundary of $\mathcal{C}(a,b)$ there exists a $\mathcal{C}$-disk (not necessarily diametral) through p and q contained in $\mathcal{C}(a,b)$. Moreover, any diametral disk through p and q has size smaller than or equal to the size of this $\mathcal{C}$-disk. Therefore, $d_{\hat{\mathcal{C}}}(a,u_i) < 1$ and $d_{\hat{\mathcal{C}}}(b,u_i) < 1$ for any $i \in \{1,\ldots,k\}$. Furthermore, we have the following:

Claim 4.1.1. Let $1 \leq i \leq k$. Then $d_{\hat{\mathcal{C}}}(a,s_i) \geq \max\{d_{\hat{\mathcal{C}}}(s_i,u_i), 1\}$

[3] Since S is in general position, only a and b can lie on the boundary of $\mathcal{C}(a,b)$.

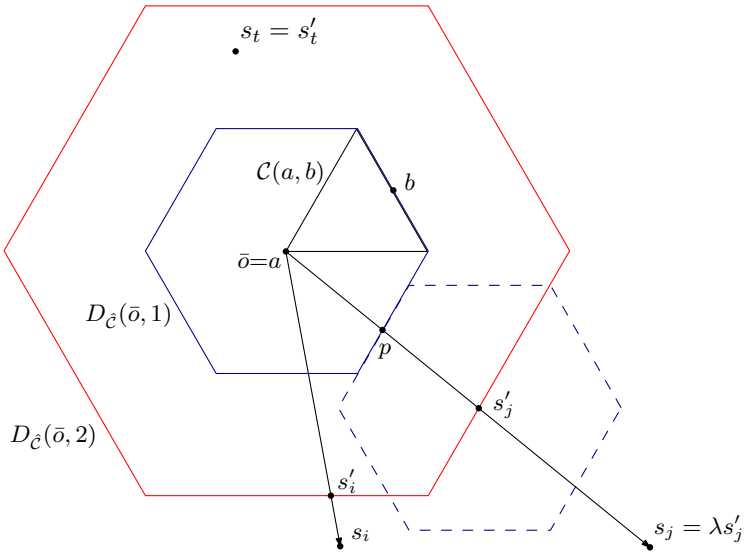

Figure 4.2: $\mathcal{C}(a,b)$ has radius 1, $D_{\hat{\mathcal{C}}}(\bar{o},1)$ is the $\hat{\mathcal{C}}$-disk with radius 1 centered at $\bar{o}$ and $D_{\hat{\mathcal{C}}}(\bar{o},2)$ is the $\hat{\mathcal{C}}$-disk with radius 2 centered at $\bar{o}$. The points s'_i and s'_j are projections of s_i and s_j on $\partial D_{\hat{\mathcal{C}}}(\bar{o},2)$, respectively. The dashed $\hat{\mathcal{C}}$-disk is centered at s'_j and has radius 1.

Proof. If $s_1 = b$, then $d_{\hat{\mathcal{C}}}(a,s_1) = 1$ and $d_{\hat{\mathcal{C}}}(s_1,u_1) < 1$. Otherwise, define $h'{=}(h \setminus \{ab, s_iu_i\}) \cup \{as_i, u_ib\}$. For sake of a contradiction suppose that $d_{\hat{\mathcal{C}}}(a,s_i){<}\max\{d_{\hat{\mathcal{C}}}(s_i,u_i),1\}$. It holds that $d_{\hat{\mathcal{C}}}(a,s_i) < \max\{d_{\hat{\mathcal{C}}}(s_i,u_i), d_{\hat{\mathcal{C}}}(a,b)\}$ since $d_{\hat{\mathcal{C}}}(a,b) = 1$. Also, $d_{\hat{\mathcal{C}}}(u_i,b) < 1$ since $u_i \in \mathcal{C}(a,b)$. Thus, $\max\{d_{\hat{\mathcal{C}}}(a,s_i), d_{\hat{\mathcal{C}}}(u_i,b)\} < \max\{d_{\hat{\mathcal{C}}}(s_i,u_i), d_{\hat{\mathcal{C}}}(a,b)\}$. Therefore $h' \prec h$, which contradicts the definition of h. □

Claim 4.1.1 implies that, for each $i \in \{1,\ldots,k\}$, s_i is not in the interior of $\mathcal{C}(a,b)$.

Claim 4.1.2. Let $1 \le i < j \le k$. Then $d_{\hat{\mathcal{C}}}(s_i,s_j) \ge \max\{d_{\hat{\mathcal{C}}}(s_i,u_i), d_{\hat{\mathcal{C}}}(s_j,u_j), 1\}$.

Proof. For sake of a contradiction suppose that $d_{\hat{\mathcal{C}}}(s_i,s_j) < \max\{d_{\hat{\mathcal{C}}}(s_i,u_i), d_{\hat{\mathcal{C}}}(s_j,u_j), 1\}$. Consider the Hamiltonian cycle $h' = h \setminus \{(a,b),(s_i,u_i),(s_j,u_j)\} \cup \{(s_i,s_j),(u_i,a),(u_j,b)\}$. As in Claim 4.1.1 we have that $d_{\hat{\mathcal{C}}}(u_i,a) < 1$ and $d_{\hat{\mathcal{C}}}(u_j,b) < 1$. So, $\max\{d_{\hat{\mathcal{C}}}(s_i,s_j), d_{\hat{\mathcal{C}}}(u_i,a), d_{\hat{\mathcal{C}}}(u_j,b)\} < \max\{d_{\hat{\mathcal{C}}}(s_i,u_i), d_{\hat{\mathcal{C}}}(s_j,u_j), d_{\hat{\mathcal{C}}}(a,b)\}$. Therefore, $h' \prec h$ which contradicts the minimality of h. □

The *$d_{\mathcal{C}}$-distance from a point* v to a region C is given by the minimum $d_{\mathcal{C}}$-distance from v to any point u in C. Notice that if $v \notin C$, then the $d_{\mathcal{C}}$-distance from v to C is defined by a point on ∂C. This can be seen by taking an $\varepsilon \in \mathbb{R}^+$

small enough such that $D_{\mathcal{C}}(v, \varepsilon)$ does not intersect C, and making ε grow until $D_{\mathcal{C}}(v, \varepsilon)$ hits C.

Observation 4.1.3. Let $u \notin D_{\mathcal{C}}(\bar{o}, r)$ for some $r \in \mathbb{R}^+$ and let p be the intersection point of $\partial D_{\mathcal{C}}(\bar{o}, r))$ and line segment $\bar{o}u$. Then, the $d_{\mathcal{C}}$-distance from u to $D_{\mathcal{C}}(\bar{o}, r)$ is $d_{\mathcal{C}}(u, p)$.

Proof. Since p is in $\partial D_{\mathcal{C}}(\bar{o}, r)$ and $u \notin D_{\mathcal{C}}(\bar{o}, r)$, the vector $u = \lambda p$ for some $\lambda > 1 \in \mathbb{R}$. In addition, the $d_{\mathcal{C}}$-distance from u to $D_{\mathcal{C}}(\bar{o}, r)$ is al least $d_{\mathcal{C}}(u, p)$. For the sake of a contradiction suppose that the $d_{\mathcal{C}}$-distance from u to $D_{\mathcal{C}}(\bar{o}, r)$ is less than $d_{\mathcal{C}}(u, p)$. Thus, there exists a point $v \in \partial D_{\hat{\mathcal{C}}}(\bar{o}, r)$ such that $d_{\hat{\mathcal{C}}}(u, v) < d_{\hat{\mathcal{C}}}(u, p)$, and $r\lambda = d_{\hat{\mathcal{C}}}(\bar{o}, u) \leq d_{\hat{\mathcal{C}}}(\bar{o}, v) + d_{\hat{\mathcal{C}}}(v, u) < d_{\hat{\mathcal{C}}}(\bar{o}, v) + d_{\hat{\mathcal{C}}}(p, u) = r + r\lambda - r = r\lambda$, which is a contradiction. □

Without loss of generality assume that a is the origin $\bar{o}$. Since for any point u in $\mathcal{C}(a, b)$, $d_{\hat{\mathcal{C}}}(\bar{o}, u) = d_{\hat{\mathcal{C}}}(a, u) \leq 1$, we have that $D_{\hat{\mathcal{C}}}(\bar{o}, 1)$ contains $\mathcal{C}(a, b)$. Also, from Claim 4.1.1, we have that s_i is not in the interior of $D_{\hat{\mathcal{C}}}(\bar{o}, 1)$ for all $i \in \{1, \ldots, k\}$. Let $D_{\hat{\mathcal{C}}}(\bar{o}, 2)$ be the $\hat{\mathcal{C}}$-disk centered at $\bar{o} = a$ with radius 2. For each $s_i \notin D_{\hat{\mathcal{C}}}(\bar{o}, 2)$, define s'_i as the intersection of $\partial D_{\hat{\mathcal{C}}}(\bar{o}, 2)$ with the ray $\overrightarrow{\bar{o}s_i}$. Let $s'_i = s_i$ when s_i is inside $D_{\hat{\mathcal{C}}}(\bar{o}, 2)$. See Figure 4.2.

Observation 4.1.4. If $s_j \notin D_{\hat{\mathcal{C}}}(\bar{o}, 2)$ (with $1 \leq j \leq k$), the $d_{\hat{\mathcal{C}}}$-distance from s'_j to $D_{\hat{\mathcal{C}}}(\bar{o}, 1)$ is 1.

Proof. Since $s_j \notin D_{\hat{\mathcal{C}}}(\bar{o}, 2)$, s'_j is on the boundary of $D_{\hat{\mathcal{C}}}(\bar{o}, 2)$ and $d_{\hat{\mathcal{C}}}(\bar{o}, s'_j) = 2$. Let p be the intersection point of $\partial D_{\hat{\mathcal{C}}}(\bar{o}, 1)$ and $\bar{o}s_j$. Then $d_{\hat{\mathcal{C}}}(\bar{o}, p) = 1$. By Observation 4.1.3 the $d_{\hat{\mathcal{C}}}$-distance from s'_j to $D_{\hat{\mathcal{C}}}(\bar{o}, 1)$ is $d_{\hat{\mathcal{C}}}(s'_j, p) = d_{\hat{\mathcal{C}}}(p, s'_j) = d_{\hat{\mathcal{C}}}(\bar{o}, s'_j) - d_{\hat{\mathcal{C}}}(\bar{o}, p) = 2 - 1 = 1$. □

The following claim is needed to prove our key lemma. Intuitively, this claim shows that if there is a point-symmetric $\mathcal{C}$-disk C of radius r centered at a point u such that $r \leq d_{\mathcal{C}}(u, \bar{o})$, then C is contained in any $\mathcal{C}$-disk with $\partial C \cap \overrightarrow{\bar{o}u}$ on its boundary such that its center u' lies on the ray $\bar{o}u$ and is farther to $\bar{o}$ than u. For an example, see Figure 4.1b.

Claim 4.1.5. Let $\mathcal{C}$ be a point-symmetric convex shape. Let u be a point in the plane different from the origin $\bar{o}$. Let $r < d_{\mathcal{C}}(u, \bar{o})$. Let p be the intersection point of $\partial D_{\mathcal{C}}(u, r)$ and line segment $\bar{o}u$. Let $u' = \lambda u$, with $\lambda > 1 \in \mathbb{R}$, be a point defined by vector u scaled by a factor of λ. Then $D_{\mathcal{C}}(u, r) \subset D_{\mathcal{C}}(u', d_{\mathcal{C}}(u', p))$. (See Figure 4.1b.)

Proof. Let $q \in D_{\mathcal{C}}(u, r)$; then $d_{\mathcal{C}}(u, q) \leq d_{\mathcal{C}}(u, p)$. Since u is on the line segment $u'p$, we have that $d_{\mathcal{C}}(u', p) = d_{\mathcal{C}}(u', u) + d_{\mathcal{C}}(u, p)$. Hence $d_{\mathcal{C}}(u', q) \leq d_{\mathcal{C}}(u', u) +$

$d_{\mathcal{C}}(u,q) \leq d_{\mathcal{C}}(u',u) + d_{\mathcal{C}}(u,p) = d_{\mathcal{C}}(u',p)$. Therefore, $D_{\mathcal{C}}(u,r)$ is contained in $D_{\mathcal{C}}(u', d_{\mathcal{C}}(u',p))$. □

Using the previous claims we can prove a key lemma stating that for every pair of points s'_i and s'_j, we have that $d_{\hat{\mathcal{C}}}(s'_i, s'_j) \geq 1$. From this lemma we can conclude that any pair of $\hat{\mathcal{C}}$-disks with radius $\frac{1}{2}$ centered at s'_i and s'_j are internally disjoint, which allows us to bound $|U|$ via a packing argument.

Lemma 4.1.6. For any pair s_i and s_j with $i \neq j$, we have that $d_{\hat{\mathcal{C}}}(s'_i, s'_j) \geq 1$.

Proof. If both s_i and s_j are in $D_{\hat{\mathcal{C}}}(\bar{o}, 2)$, then from Claim 4.1.2 we have that $d_{\hat{\mathcal{C}}}(s'_i, s'_j) = d_{\hat{\mathcal{C}}}(s_i, s_j) \geq 1$. Otherwise, we assume, without loss of generality, that $d_{\hat{\mathcal{C}}}(\bar{o}, s_j) \geq d_{\hat{\mathcal{C}}}(\bar{o}, s_i)$. Then, $s_j \notin D_{\hat{\mathcal{C}}}(\bar{o}, 2)$. Since s'_j is on the line segment $\bar{o}s_j$, we have $s_j = \lambda s'_j$ for some $\lambda > 1$. Let p be the intersection point of $\partial D_{\hat{\mathcal{C}}}(\bar{o}, 1)$ and $\bar{o}s_j$. Since $d_{\hat{\mathcal{C}}}$ defines a norm, we have $d_{\hat{\mathcal{C}}}(\lambda s'_j, \bar{o}) = \lambda d_{\hat{\mathcal{C}}}(s'_j, \bar{o})$. By Observation 4.1.4 we have that $d_{\hat{\mathcal{C}}}(s_j, p) = d_{\hat{\mathcal{C}}}(s_j, \bar{o}) - d_{\hat{\mathcal{C}}}(p, \bar{o}) = \lambda d_{\hat{\mathcal{C}}}(s'_j, \bar{o}) - 1 = 2\lambda - 1$. From Observation 4.1.3 it follows that the $d_{\hat{\mathcal{C}}}$-distance from s_j to $D_{\hat{\mathcal{C}}}(\bar{o}, 1)$ is equal to $d_{\hat{\mathcal{C}}}(s_j, p)$. Further, $d_{\hat{\mathcal{C}}}(s_j, s'_j) = d_{\hat{\mathcal{C}}}(s_j, \bar{o}) - d_{\hat{\mathcal{C}}}(s'_j, \bar{o}) = 2\lambda - 2$. Let us prove that $d_{\hat{\mathcal{C}}}(s'_i, s'_j) \geq 1$. For sake of a contradiction suppose that $d_{\hat{\mathcal{C}}}(s'_i, s'_j) < 1$. Let $D_{s'_j} = D_{\hat{\mathcal{C}}}(s'_j, 1)$. By Observation 4.1.4, $d_{\hat{\mathcal{C}}}(s'_j, p) = 1$. Therefore, p is on $\partial D_{s'_j}$. Now, we consider two cases:

Case 1) $s_i \in D_{\hat{\mathcal{C}}}(\bar{o}, 2)$. Then $d_{\hat{\mathcal{C}}}(\bar{o}, s_i) \leq 2$. Since $d_{\hat{\mathcal{C}}}(s'_i, s'_j) < 1$, we have $s_i \in D_{s'_j}$. From Claim 4.1.5 it follows that $D_{s'_j}$ is contained in $D_{\hat{\mathcal{C}}}(s_j, d_{\hat{\mathcal{C}}}(s_j, p))$. Thus, $s'_i \in D_{\hat{\mathcal{C}}}(s_j, d_{\hat{\mathcal{C}}}(s_j, p))$ and $d_{\hat{\mathcal{C}}}(s_j, s'_i) = d_{\hat{\mathcal{C}}}(s_j, s_i) \leq d_{\hat{\mathcal{C}}}(s_j, p)$. Since S is in general position, u_j is in the interior of $D_{\hat{\mathcal{C}}}(\bar{o}, 1)$. Hence, $d_{\hat{\mathcal{C}}}(s_j, s_i) \leq d_{\hat{\mathcal{C}}}(s_j, p) < d_{\hat{\mathcal{C}}}(s_j, u_j)$, which contradicts Claim 4.1.2.

Case 2) $s_i \notin D_{\hat{\mathcal{C}}}(\bar{o}, 2)$. Then $d_{\hat{\mathcal{C}}}(\bar{o}, s_i) > 2$. Thus, $s_i = \delta s'_i$ for some $\delta > 1 \in \mathbb{R}$. Moreover, since $d_{\hat{\mathcal{C}}}(\bar{o}, s_j) \geq d_{\hat{\mathcal{C}}}(\bar{o}, s_i)$ and s'_i, s'_j are on $\partial D_{\hat{\mathcal{C}}}(\bar{o}, 2)$, $\delta \leq \lambda$. Hence, s_i is on the line segment $s'_i(\lambda s'_i)$. Let $D_{s_j} = D_{\hat{\mathcal{C}}}(s_j, 2\lambda - 1)$. Note that $\lambda < 2\lambda - 1$ because $\lambda > 1$. Since $d_{\hat{\mathcal{C}}}$ defines a norm, $d_{\hat{\mathcal{C}}}(s_j, \lambda s'_i) = d_{\hat{\mathcal{C}}}(\lambda s'_j, \lambda s'_i) = \lambda d_{\hat{\mathcal{C}}}(s'_j, s'_i) < \lambda < 2\lambda - 1$. Hence, $\lambda s'_i \in D_{s_j}$. In addition, since $d_{\hat{\mathcal{C}}}(s_j, p) = 2\lambda - 1$, from Claim 4.1.5 it follows that $D_{s'_j} \subseteq D_{s_j}$. Therefore, $s'_i \in D_{s_j}$. Thus, the line segment $s'_i(\lambda s'_i)$ is contained in D_{s_j}. Hence, $s_i \in D_{s_j}$. Then, $d_{\hat{\mathcal{C}}}(s_j, s_i) \leq 2\lambda - 1 = d_{\hat{\mathcal{C}}}(s_j, p) < d_{\hat{\mathcal{C}}}(s_j, u_j)$ which contradicts Claim 4.1.2. □

Theorem 4.1.7. For any set S of points in general position and convex shape $\mathcal{C}$, the graph 24-$GG_{\mathcal{C}}(S)$ is Hamiltonian.

Proof. For each s_i we define the $\hat{\mathcal{C}}$-disk $D_i = D_{\hat{\mathcal{C}}}(s'_i, \frac{1}{2})$. We also set $D_0 := D_{\hat{\mathcal{C}}}(\bar{o}, \frac{1}{2})$ (recall that we can assume without loss of generality that $a = \bar{o}$). By Lemma 4.1.6, each pair of $\hat{\mathcal{C}}$-disks D_i and D_j $(0 < i < j \leq k)$ are internally disjoint. Note that, if s'_i is on $\partial D_{\hat{\mathcal{C}}}(\bar{o}, 2)$, then D_0 and D_i are internally disjoint.

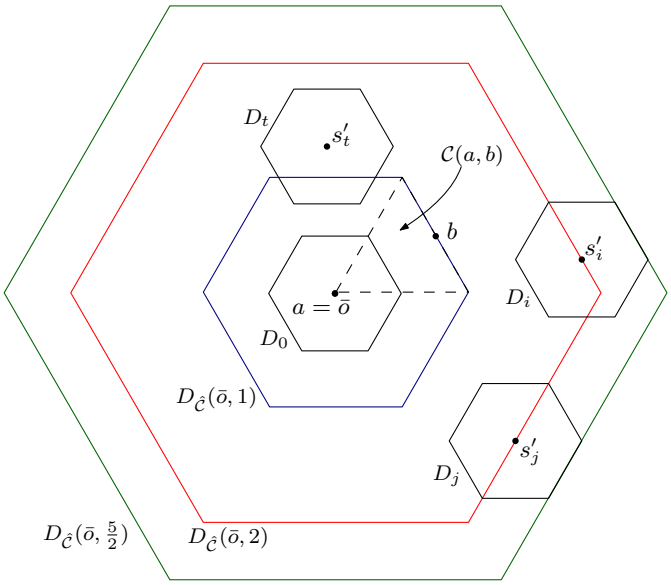

Figure 4.3: The $\hat{\mathcal{C}}$-disks D_0, D_i, D_j and D_t of radius $\frac{1}{2}$ are centered at a, s'_i, s'_j and s'_t, respectively. Such $\hat{\mathcal{C}}$-disks are contained in the $\hat{\mathcal{C}}$-disk $D_{\hat{\mathcal{C}}}(\bar{o}, \frac{5}{2})$.

On the other hand, if s'_i is in the interior of $D_{\hat{\mathcal{C}}}(\bar{o}, 2)$, then by definition $s'_i = s_i$. Thus, by Claim 4.1.1 D_0 is internally disjoint from D_i. See Figure 4.3. Since $s'_i \in D_{\hat{\mathcal{C}}}(\bar{o}, 2)$ for all i, each disk D_i is inside $D_{\hat{\mathcal{C}}}(\bar{o}, \frac{5}{2})$. In $D_{\hat{\mathcal{C}}}(\bar{o}, \frac{5}{2})$, there can be at most $\frac{Area(D_{\hat{\mathcal{C}}}(\bar{o}, \frac{5}{2}))}{Area(D_0)} = \frac{(\frac{5}{2})^2 Area(\hat{\mathcal{C}})}{(\frac{1}{2})^2 Area(\hat{\mathcal{C}})} = 25$ internally disjoint disks of type D_i. Thus, since D_0 is centered at a, there are at most 24 points s'_i in $D_{\hat{\mathcal{C}}}(\bar{o}, 1)$. As a consequence, there are at most 24 points of S in the interior of $\mathcal{C}(a, b)$, and the bottleneck Hamiltonian cycle of S is contained in 24-$GG_{\mathcal{C}}(S)$. □

4.2 Hamiltonicity for point-symmetric convex shapes

In this section we improve Theorem 4.1.7 for the case where $\mathcal{C}$ is convex and point-symmetric. We use similar arguments to those in Section 4.1.

Consider h defined as before, i.e., h is the minimum Hamiltonian cycle in $\mathcal{H}$. Let ab be an edge in h and consider an arbitrary fixed $\mathcal{C}(a, b)$. In this section it will be more convenient to assume without loss of generality that $d_{\mathcal{C}}(a, b) = 2$ and that $\mathcal{C}(a, b)$ is centered at the origin $\bar{o}$. Thus, $\mathcal{C}(a, b) = D_{\mathcal{C}}(\bar{o}, 1)$, see Figure 4.4. Consider again the set $U = \{u_1, \ldots, u_k\}$ defined as in Section 4.1, and let s_i be the predecessor of u_i in h.

Using that $d_{\mathcal{C}}(a, b) = 2d_{\hat{\mathcal{C}}}(a, b)$ when $\mathcal{C}$ is point-symmetric, we can prove the

following claims.

Claim 4.2.1. $d_{\mathcal{C}}(s_i, a) \geq \max\{d_{\mathcal{C}}(s_i, u_i), 2\}$.

Proof. By Claim 4.1.1 we have that $d_{\mathcal{C}}(s_i, a) = 2d_{\hat{\mathcal{C}}}(s_i, a) \geq 2\max\{d_{\hat{\mathcal{C}}}(s_i, u_i), 1\} = \max\{2d_{\hat{\mathcal{C}}}(s_i, u_i),\ 2\} = \max\{d_{\mathcal{C}}(s_i, u_i), 2\}$. □

Claim 4.2.2. Let $1 \leq i < j \leq k$, then $d_{\mathcal{C}}(s_i, s_j) \geq \max\{d_{\mathcal{C}}(s_i, u_i), d_{\mathcal{C}}(s_j, u_j), 2\}$.

Proof. By Claim 4.1.2 we have that $d_{\mathcal{C}}(s_i, s_j) = 2d_{\hat{\mathcal{C}}}(s_i, s_j) \geq 2\max\{d_{\hat{\mathcal{C}}}(s_i, u_i), d_{\hat{\mathcal{C}}}(s_j, u_j), 1\} = \max\{2d_{\hat{\mathcal{C}}}(s_i, u_i), 2d_{\hat{\mathcal{C}}}(s_j, u_j), 2\} = \max\{d_{\mathcal{C}}(s_i, u_i), d_{\mathcal{C}}(s_j, u_j), 2\}$. □

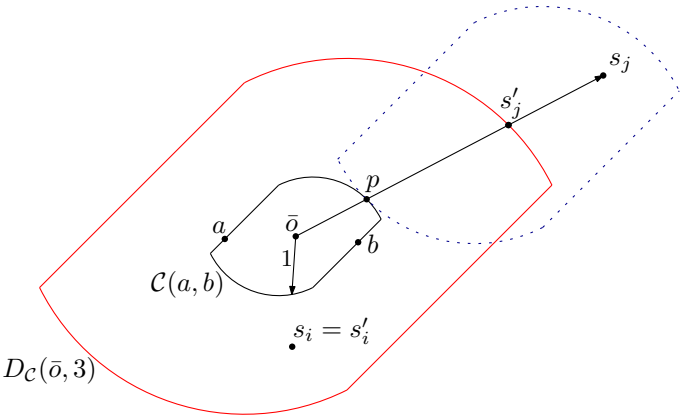

Figure 4.4: $\mathcal{C}(a, b)$ has radius 1 and it is centered at $\bar{o}$. The point s_j is not in $D_{\mathcal{C}}(\bar{o}, 3)$, so s'_j is the intersection point of $\bar{o}s_j \cap \partial D_{\mathcal{C}}(\bar{o}, 3)$. The dotted $\mathcal{C}$-disk is centered at s'_j and has radius 2.

From Claim 4.2.1, we have that s_i is not in the interior of $D_{\mathcal{C}}(\bar{o}, 1) = \mathcal{C}(a, b)$ for all $i \in \{1, \ldots, k\}$. Let $D_{\mathcal{C}}(\bar{o}, 3)$ be the $\mathcal{C}$-disk centered at $\bar{o}$ with radius 3. For each $s_i \notin D_{\mathcal{C}}(\bar{o}, 3)$, define s'_i as the intersection of $\partial D_{\mathcal{C}}(\bar{o}, 3)$ with the ray $\overrightarrow{\bar{o}s_i}$. We let $s'_i = s_i$ when s_i is inside $D_{\mathcal{C}}(\bar{o}, 3)$. See Figure 4.4.

The following lemma is similar to Lemma 4.1.6. We show that every pair s'_i and s'_j are at $d_{\mathcal{C}}$-distance at least 2. This lemma allows us again to reduce our problem to a packing problem.

Lemma 4.2.3. For any pair s_i and s_j with $i \neq j$, we have that $d_{\mathcal{C}}(s'_i, s'_j) \geq 2$. Moreover, if at least one of s_i and s_j is not in $D_{\mathcal{C}}(\bar{o}, 3)$, then $d_{\mathcal{C}}(s'_i, s'_j) > 2$.

Proof. If both s_i and s_j are in $D_{\mathcal{C}}(\bar{o}, 3)$, then from Claim 4.2.2 we have that $d_{\mathcal{C}}(s'_i, s'_j) = d_{\mathcal{C}}(s_i, s_j) \geq 2$. Otherwise, assume without loss of generality that $d_{\mathcal{C}}(\bar{o}, s_j) \geq d_{\mathcal{C}}(\bar{o}, s_i)$. Then $s_j \notin D_{\mathcal{C}}(\bar{o}, 3)$ and s'_j is on the line segment $\bar{o}s_j$.

Thus, $s_j = \lambda s'_j$ for some $\lambda > 1$. Let p be the intersection point of $\partial\mathcal{C}(a,b)$ and $\bar{o}s_j$. By Observation 4.1.3 the $d_\mathcal{C}$-distance from s'_j to $\mathcal{C}(a,b)$ is $d_\mathcal{C}(s'_j,p) = d_\mathcal{C}(s'_j,\bar{o}) - d_\mathcal{C}(p,\bar{o}) = 2$. Since $d_\mathcal{C}$ defines a norm, $d_\mathcal{C}(s_j,p) = d_\mathcal{C}(s_j,\bar{o}) - d_\mathcal{C}(p,\bar{o}) = \lambda d_\mathcal{C}(s'_j,\bar{o}) - 1 = 3\lambda - 1$, and this corresponds to the $d_\mathcal{C}$-distance from s_j to $\mathcal{C}(a,b)$. Further, $d_\mathcal{C}(s_j,s'_j) = d_\mathcal{C}(s_j,\bar{o}) - d_\mathcal{C}(s'_j,\bar{o}) = 3\lambda - 3$. For sake of contradiction we suppose that $d_\mathcal{C}(s'_i,s'_j) \leq 2$. Thus, s'_i is in $D_\mathcal{C}(s'_j,2)$. We consider the following two cases.

Case 1) $s_i \in D_\mathcal{C}(\bar{o},3)$. Then $d_\mathcal{C}(\bar{o},s_i) \leq 3$. Since $d_\mathcal{C}(s'_i,s'_j) \leq 2$, $s_i = s'_i \in D_\mathcal{C}(s'_j,2)$. From Claim 4.1.5 follows that $D_\mathcal{C}(s'_j,2) \subset D_\mathcal{C}(s_j,3\lambda-1)$. Thus, $s_i \in D_\mathcal{C}(s_j,3\lambda-1)$. Hence, $d_\mathcal{C}(s_j,s'_i) = d_\mathcal{C}(s_j,s_i) \leq d_\mathcal{C}(s_j,p)$. Since S is in general position, u_j is in the interior of $\mathcal{C}(a,b)$. Therefore, $d_\mathcal{C}(s_j,s_i) \leq d_\mathcal{C}(s_j,p) < d_\mathcal{C}(s_j,u_j)$, which contradicts Claim 4.2.2.

Case 2) $s_i \notin D_\mathcal{C}(\bar{o},3)$. Then $s'_i \in \partial D_\mathcal{C}(\bar{o},3)$ and $s_i = \delta s'_i$ for some $\delta > 1$. Moreover, since $d_\mathcal{C}(\bar{o},s_j) \geq d_\mathcal{C}(\bar{o},s_i)$ and s'_i, s'_j are on the boundary of $D_\mathcal{C}(\bar{o},3)$, $\delta \leq \lambda$. Hence, s_i is on the line segment $s'_i(\lambda s'_i)$. Note that $2\lambda < 3\lambda - 1$ because $\lambda > 1$. Since $d_\mathcal{C}$ defines a norm, $d_\mathcal{C}(s_j,\lambda s'_i) = \lambda d_\mathcal{C}(s'_j,s'_i) \leq 2\lambda < 3\lambda - 1$. Hence, $\lambda s'_i \in D_\mathcal{C}(s_j,3\lambda-1)$. In addition, from Claim 4.1.5 it follows that $D_\mathcal{C}(s'_j,2) \subseteq D_\mathcal{C}(s_j,3\lambda-1)$. Thus, $s'_i \in D_\mathcal{C}(s_j,3\lambda-1)$ and the line segment $s'_i(\lambda s'_i)$ is contained in $D_\mathcal{C}(s_j,3\lambda-1)$. Then, $s_i \in D_\mathcal{C}(s_j,3\lambda-1)$ and $d_\mathcal{C}(s_j,s_i) \leq 3\lambda - 1 = d_\mathcal{C}(s_j,p) < d_\mathcal{C}(s_j,u_j)$, which contradicts Claim 4.2.2. □

Theorem 4.2.4. For any set S of points in general position and point-symmetric convex shape $\mathcal{C}$, the graph 15-$GG_\mathcal{C}(S)$ is Hamiltonian.

Proof. For each $s_i \in S$ we define the $\mathcal{C}$-disk $D_i = D_\mathcal{C}(s'_i,1)$. We also set $D_0 := D_\mathcal{C}(a,1)$. From Lemma 4.2.3, each pair of $\mathcal{C}$-disks D_i and D_j are internally disjoint, for $0 < i < j \leq k$. Note that, if s'_i is on $\partial D_\mathcal{C}(\bar{o},3)$, then D_0 and D_i are internally disjoint. On the other hand, if s'_i is in the interior of $D_\mathcal{C}(\bar{o},3)$, then by definition $s'_i = s_i$. Thus, by Claim 4.2.1 D_0 is internally disjoint from D_i. Consider $D_\mathcal{C}(\bar{o},4)$. Since, $s'_i \in D_\mathcal{C}(\bar{o},3)$ for all $i \in \{1,\dots k\}$, then each disk D_i is inside $D_\mathcal{C}(\bar{o},4)$. Hence, in $D_\mathcal{C}(\bar{o},4)$ there can be at most $\frac{Area(D_\mathcal{C}(\bar{o},4))}{Area(\mathcal{C})} = \frac{4^2 Area(\mathcal{C})}{Area(\mathcal{C})} = 16$ internally disjoint disks of type D_i. Since D_0 is centered at a, there are at most 15 points s'_i in $D_\mathcal{C}(\bar{o},3)$. Therefore, there are at most 15 points of S in $\mathcal{C}(a,b)$, and the bottleneck Hamiltonian cycle of S is contained in 15-$GG_\mathcal{C}(S)$. □

4.3 Hamiltonicity for regular polygons

An important family of point-symmetric convex shapes is that of regular even-sided polygons. When $\mathcal{C}$ is a regular polygon $\mathcal{P}_t$ with t sides, for t even, we

can improve the previous bound by analyzing the properties of the shape for different values of t.

4.3.1 Hamiltonicity for squares

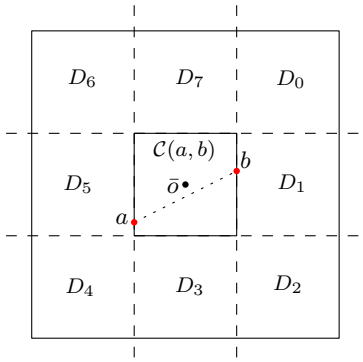

Figure 4.5: Lines $x = -1, x = 1, y = -1$, and $y = 1$ split $D_\square(3, \bar{o})$ into nine unit squares: $\mathcal{C}(a, b), D_0, \ldots, D_7$.

First, we consider the case when the polygon is a square. In this case, we divide $D_\square(\bar{o}, 3)$ into 9 disjoint squares of radius 1 and show that there can be at most one point of $\{a, s'_1, \ldots, s'_k\}$ in each such square. We use lines $x = -1, x = 1, y = -1$, and $y = 1$ to split $D_\square(\bar{o}, 3)$ into 9 squares of radius 1. Refer to Figure 4.5. Let $D_0, D_1, \ldots, D_7$ be the squares of radius 1 in $D_\square(\bar{o}, 3)$ different from $C(a, b)$, ordered clockwise, and where D_0 is the top-right corner square. In the following lemma we prove that there is at most one point of $\{a, s'_1, \ldots, s'_k\}$ in each D_i. Let indices be taken modulo 8. Note that each D_i shares a side with D_{i-1}, and for each odd i, D_i shares a side with $\mathcal{C}(a, b)$. Moreover, there exists a D_i that contains a on its boundary. We will associate any point in $D_\square(\bar{o}, 3)$ (not in the interior of $\mathcal{C}(a, b)$) to a unique square D_i in the following way: Let p be a point in D_i. If p does not lie on the shared boundary of D_i and some other D_j, then p is associated to D_i. If i is odd and p is the intersection point $D_i \cap D_{i-1} \cap D_{i-2}$, then p is associated to D_{i-2} (p can be a or b). Otherwise, if p is on the edge $D_i \cap D_{i-1}$, then p is associated to D_{i-1}.

Observation 4.3.1. Any two points at $d_\square$-distance 2 in a unit square must be on opposite sides of the square.

Lemma 4.3.2. There is at most one s'_j associated to each D_i. Moreover, the D_i containing the point a on its boundary has no s'_j associated to it.

Proof. Suppose that there are two points s'_j and s'_m associated to D_i. From Lemma 4.2.3 we have that $d_\square(s'_j, s'_m) \geq 2$. Also, since D_i is a unit square,

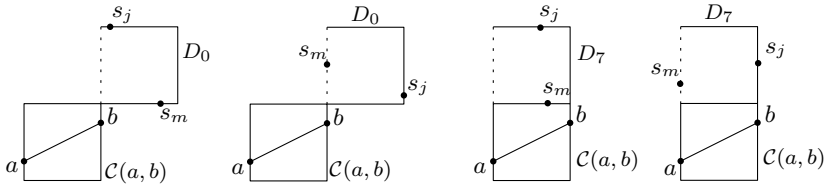

Figure 4.6: Cases (a) and (b) contradict Claim 4.2.2. Case (c) contradicts our general position assumption. In case (d) only one point, s_j, is associated to D_7.

$d_\square(s'_j, s'_m) \leq 2$. Therefore, $d_\square(s'_j, s'_m) = 2$. Then Lemma 4.2.3 implies that s_j and s_m must be inside $D_\mathcal{C}(\bar{o}, 3)$. In addition, by Observation 4.3.1, the points s_j and s_m are on opposite sides of the boundary of D_i. For simplicity we will assume that $d_\square(\bar{o}, s_j) \geq d_\square(\bar{o}, s_m)$. If i is even, then the $d_\square$-distance of s_j to $\mathcal{C}(a,b)$ is exactly 2. We refer to Figure 4.6a and 4.6b. Recall that, by our general position assumption, u_j is in the interior of $\mathcal{C}(a,b)$. Thus, the $d_\square$-distance from s_j to $\mathcal{C}(a,b)$ is less than $d_\square(s_j, u_j)$, i.e., $d_\square(s_j, u_j) > 2$. Hence, $d_\square(s_j, s_m) = 2 < d_\square(s_j, u_j)$ which contradicts Claim 4.2.2. Therefore, if i is even, there is at most one point in D_i, which is associated to it.
If i is odd, then s_j is either on $D_i \cap D_{i-1}$ or $D_i \cap D_{i+1}$, or on $D_i \cap \partial D_\mathcal{C}(\bar{o}, 3)$ (see Figure 4.6c and 4.6d). If s_j is on $D_i \cap D_{i-1}$ or $D_i \cap D_{i+1}$, then only one of s_j and s_m is associated to D_i. If s_j is on $D_i \cap \partial D_\mathcal{C}(\bar{o}, 3)$, then by Observation 4.3.1, s_m is on $\mathcal{C}(a,b)$, which by our general position assumption implies that $s_m = b$, since $s_m \neq a$. Thus, $d_\square(s_j, u_j) > 2 = d_\square(s_j, s_m)$, which contradicts Claim 4.2.2. Therefore, there is only one point associated to D_i.

Finally, if D_i contains a, then there is no point s'_j in D_i. Indeed, assume for sake of a contradiction that $s'_j \in D_i$. Then, s_j is not in D_i, otherwise, $d_\square(a, s_j) < d_\square(s_j, u_j)$, contradicting Claim 4.2.1. Thus, s'_j is on $D_i \cap \partial D_\square(\bar{o}, 3)$ and $s_j = \lambda s'_j$ for some $\lambda > 1$. Hence, $d_\square(s'_j, a) = 2$, which means that $a \in D_\square(s'_j, 2)$. Let p be the point $\bar{o}s'_j \cap \partial\mathcal{C}(a,b)$. By Claim 4.1.5, $D_\square(s'_j, 2) \subset D_\square(s_j, d_\square(s_j, p))$. So, $a \in D_\square(s_j, d_\square(s_j, p))$ and $d_\square(s_j, a) < d_\square(s_j, u_j)$, contradicting Claim 4.2.1. □

Theorem 4.3.3. For any set S of points in general position, the graph 7-$GG_\square(S)$ is Hamiltonian.

Proof. From Lemma 4.3.2 we have that for each $0 \leq i \leq 7$ there is at most one point of $\{a, s'_1, \ldots, s'_k\}$ associated to D_i, and any square containing a has no s'_i associated to it. Since there is at least one D_i containing a, there are at most 7 points s'_j in $D_\square(\bar{o}, 3)$. Therefore, there are at most 7 points of S in the interior of $\mathcal{C}(a,b)$, and the bottleneck Hamiltonian cycle of S is contained in

7-$GG_{\square}(S)$. □

4.3.2 Hamiltonicity for regular hexagons

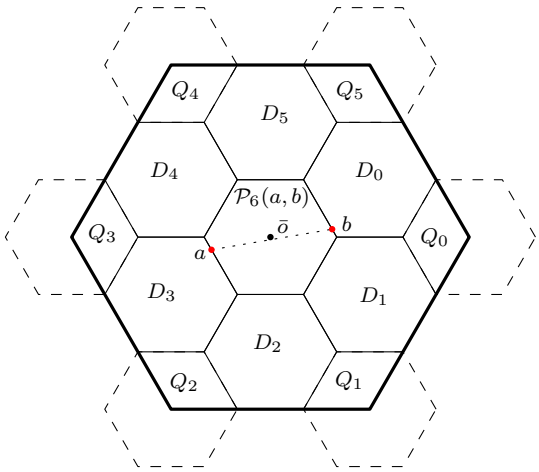

Figure 4.7: The bold hexagon is the boundary of $D_{\mathcal{P}_6}(\bar{o}, 3)$. Such hexagon is divided into 13 interior-disjoint regions: 6 quadrangles—a third of a unit $\mathcal{P}_6$-disk—and 7 unit $\mathcal{P}_6$-disks.

The analysis for the case of hexagons is similar to the previous one. First we divide the hexagon $D_{\mathcal{P}_6}(\bar{o}, 3)$ into 13 different regions $\mathcal{C}(a,b), D_0, \ldots, D_5$, $Q_0, \ldots, Q_5$, shown in Figure 4.7. Let indices be taken modulo 6. We will associate a point in $D_{\mathcal{P}_6}(\bar{o}, 3)$ (not in the interior of $\mathcal{C}(a,b)$) to a region D_i or Q_i in the following fashion. If a point is in the interior of D_i or Q_i we say that such point is associated to D_i or Q_i, respectively. If a point is on the edge $D_i \cap D_{i-1}$ or edge $D_i \cap Q_{i-1}$, then such point is associated to D_{i-1} and Q_{i-1}, respectively. In the case when a point is the vertex $D_i \cap D_{i-1} \cap Q_{i-1}$, we say that such point is associated to D_{i-1}. When a point is on the edge $D_i \cap Q_i$ then we associate it with D_i. See Figure 4.8.

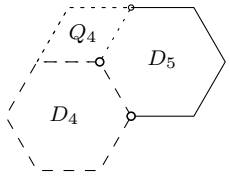

Figure 4.8: The dashed boundary of D_5 is associated to D_4 and the dotted one is associated to Q_4. The rest of D_5 is associated to D_5.

Observation 4.3.4. Any two points at $d_{\mathcal{P}_6}$-distance 2 in a unit hexagon D must be on opposite sides of D.

In the following lemma we show that the hexagon $D_{\mathcal{P}_6}(\bar{o}, 3)$ contains at most 11 points $s'_1, \ldots, s'_k$.

Lemma 4.3.5. There is at most one point s'_j associated to each region of type D_i or Q_i. Moreover, there is no point s'_j in the hexagon D_i that contains a.

Proof. If a point is in the interior of D_i or Q_i then by Observation 4.3.4 there is no other point in the same region.

Note that if Q_i contains two points at $d_{\mathcal{P}_6}$-distance 2, then by Observation 4.3.4 such points are exactly $D_i \cap Q_i \cap \partial D_{\mathcal{P}_6}(\bar{o}, 3)$ and $D_{i+1} \cap Q_i \cap \partial D_{\mathcal{P}_6}(\bar{o}, 3)$. Since the points on $D_i \cap Q_i$ are associated to D_i, the intersection point $D_i \cap Q_i \cap \partial D_{\mathcal{P}_6}(\bar{o}, 3)$ is not associated to Q_i. Thus, there is at most one point associated to Q_i.

If D_i contains a point s'_j that is on $D_i \cap \partial D_{\mathcal{P}_6}(\bar{o}, 3)$, then there cannot be another $s'_m \in D_i$: Otherwise, by Observation 4.3.4, s'_m would be on the boundary of $\mathcal{P}_6(a, b)$, in which case $s'_m = b$ due to our general position assumption. Since, $d_{\mathcal{P}_6}(s'_j, s'_m) = 2$, it follows from Lemma 4.2.3 that s_j would be in $D_{\mathcal{P}}(\bar{o}, 3)$. Thus, $d_{\mathcal{P}_6}(s_j, s_m) = 2 < d_{\mathcal{P}_6}(s_j, u_j)$ which contradicts Claim 4.2.2. Consequently, if D_i contains two points s'_j and s'_m then by Observation 4.3.4 either: 1) one is on the edge $D_i \cap Q_i$ and the other is on the edge $D_i \cap D_{i-1}$ (see Figure 4.9a); or 2) one is on the edge $D_i \cap D_{i+1}$ and the other is on the edge $D_i \cap Q_{i-1}$ (see Figure 4.9b). In either case, just one point is associated to D_i.

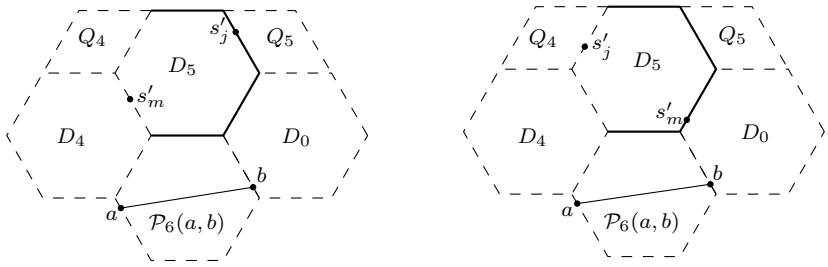

Figure 4.9: In both, (a) and (b), D_5 contains exactly two points at $d_{\mathcal{P}_6}$-distance 2.

Finally, if D_i contains a, then there is no point s'_j in D_i. Indeed, suppose for sake of contradiction that $s'_j \in D_i$. Then, s_j is not in D_i because, by Claim 4.2.1, $d_{\mathcal{P}_6}(a, s_j) \geq d_{\mathcal{P}_6}(s_j, u_j)$. Thus, s'_j is on $D_i \cap \partial D_{\mathcal{P}_6}(\bar{o}, 3)$ and $s_j = \lambda s'_j$ for some $\lambda > 1$. Hence, $d_{\mathcal{P}_6}(s'_j, a) = 2$ and $a \in D_{\mathcal{P}_6}(s'_j, 2)$. Let p be the intersection point $\bar{o}s'_j \cap \partial \mathcal{P}_6(a, b)$. By Claim 4.1.5, $D_{\mathcal{P}_6}(s'_j, 2) \subset D_{\mathcal{P}_6}(s_j, d_{\mathcal{P}_6}(s_j, p))$ and,

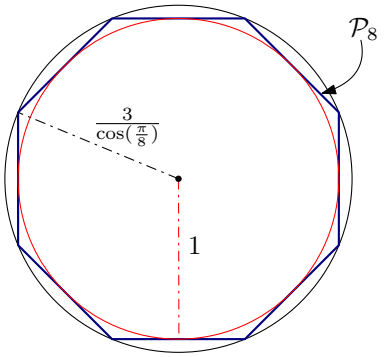

Figure 4.10: The incircle of the octagon $\mathcal{P}_8$ has Euclidean radius 1. The octagon $\mathcal{P}_8$ is inscribed in a circle of Euclidean radius $\frac{3}{\cos(\frac{\pi}{8})}$; such circle is also known as the circumcircle of $\mathcal{P}_8$.

thus, $a \in D_{\mathcal{P}_6}(s_j, d_{\mathcal{P}_6}(s_j, p))$. We obtain that $d_{\mathcal{P}_6}(s_j, a) < d_{\mathcal{P}_6}(s_j, u_j)$, which again contradicts Claim 4.2.1. □

The following theorem holds.

Theorem 4.3.6. For any set S of points in general position, the graph 11-$GG_{\mathcal{P}_6}(S)$ is Hamiltonian.

4.3.3 Hamiltonicity for regular even-sided t-gons where $t \geq 8$

For the remaining regular polygons with an even number of sides, we use the circumcircle [4] of $D_{\mathcal{P}_t}(\bar{o}, 3)$ in order to give an upper bound on the number of points in $D_{\mathcal{P}_t}(\bar{o}, 3)$ at pairwise Euclidean distance at least 2. Without loss of generality we assume that the incircle [5] of the unit $\mathcal{P}_t$-disk has Euclidean radius 1. See Figure 4.10.

In this section we will first treat the case $t \geq 10$, and afterwards the case $t = 8$.

Theorem 4.3.7. For any set S of points in general position and regular polygon $\mathcal{P}_t$ with even $t \geq 10$, the graph 11-$GG_{\mathcal{P}_t}(S)$ is Hamiltonian.

Proof. Let $\mathcal{P}_t$ be a polygon with $t \geq 10$ sides and t even. Then $D_{\mathcal{P}_t}(\bar{o}, 3)$ is inscribed in a circle of radius $r = \frac{3}{\cos(\frac{\pi}{t})}$. Since the function $\cos(\frac{\pi}{t})$ is an increasing function for $t \geq 2$, we have that $r \leq \frac{3}{\cos(\frac{\pi}{10})}$. Therefore, $D_{\mathcal{P}_t}(\bar{o}, 3)$ is

[4] The *circumcircle* of a polygon $\mathcal{P}$ is the smallest circle that contains $\mathcal{P}$.

[5] The *incircle* of a polygon $\mathcal{P}$ is the largest circle in the interior of $\mathcal{P}$ that is tangent to each side of $\mathcal{P}$.

inside the circumcircle of a decagon with incircle of radius 3. In addition, from Lemma 4.2.3 we know that for any pair of points s'_i, s'_j in $D_{\mathcal{P}_t}(\bar{o}, 3)$, $d_{\mathcal{P}_t}(s'_i, s'_j) \geq 2$. Since the incircle of the 2-unit $\mathcal{P}_t$-disk has Euclidean radius 2, we have that $d(s'_i, s'_j) \geq 2$. Hence, it suffices to show that there are at most 12 points in $D_{\mathcal{P}_t}(\bar{o}, 3)$ at pairwise Euclidean distance at least 2. Fodor [56] proved that the minimum radius R of a circle having 13 points at pairwise Euclidean distance at least 2 is $R \approx 3.236$, which is greater than $\frac{3}{\cos(\frac{\pi}{10})} \approx 3.154$. Thus, $D_{\mathcal{P}_t}(\bar{o}, 3)$ contains at most 12 points at pairwise $d_{\mathcal{P}_t}$-distance at least 2. Since a is also at $d_{\mathcal{P}_t}$-distance at least 2 from all s'_i's, there are at most 11 points of S inside $\mathcal{P}_t(a, b)$. □

For the case of octagons, the radius of the circumcircle of $D_{\mathcal{P}_8}(\bar{o}, 3)$ is greater than 3.236, so we cannot use the result in [56]. However, we can use a similar result from Fodor [57] to prove an analogous theorem:

Theorem 4.3.8. For any set S of points in general position, the graph 12-$GG_{\mathcal{P}_8}(S)$ is Hamiltonian.

Proof. From Lemma 4.2.3 we know that, for any pair of points s'_i, s'_j in $D_{\mathcal{P}_8}(\bar{o}, 3)$, $d_{\mathcal{P}_8}(s'_i, s'_j) \geq 2$. Since the incircle of the 2-unit $\mathcal{P}_8$-disk has Euclidean radius 2, we have that $d(s'_i, s'_j) \geq 2$. Hence, it suffices to show that there are at most 13 points in $D_{\mathcal{P}_8}(\bar{o}, 3)$ at pairwise Euclidean distance at least 2. The regular octagon $D_{\mathcal{P}_8}(\bar{o}, 3)$ is inscribed in a circle of radius $r = \frac{3}{\cos(\frac{\pi}{8})} \approx 3.247$. By a result of Fodor [57], the smallest radius R of a circle containing 14 points at pairwise Euclidean distance at least 2 is $R \approx 3.328$. Hence, $D_{\mathcal{P}_8}(\bar{o}, 3)$ contains at most 13 points at pairwise Euclidean distance at least 2. Since a is also at $d_{\mathcal{P}_8}$-distance at least 2 from all s'_i's, there are at most 12 points of S inside $\mathcal{P}_8(a, b)$. □

4.4 Bottleneck Hamiltonian cycles in k-$GG_{\square}$ and k-$GG_{\mathcal{P}_6}$

In this section we give lower bounds on the minimum values of k for which the graphs k-$GG_{\square}$ and k-$GG_{\mathcal{P}_6}$ contain a bottleneck Hamiltonian cycle. This is useful to understand to what extent we can use the bottleneck Hamiltonian cycle for showing Hamiltonicity in a k-$GG_{\mathcal{C}}$ in order to improve the known upper bounds on k. The proofs are very similar to those in [22, 21, 67].

Lemma 4.4.1. There exists a point set S with $n \geq 17$ points such that 2-$GG_{\square}(S)$ does not contain any $d_{\square}$-bottleneck Hamiltonian cycle of S.

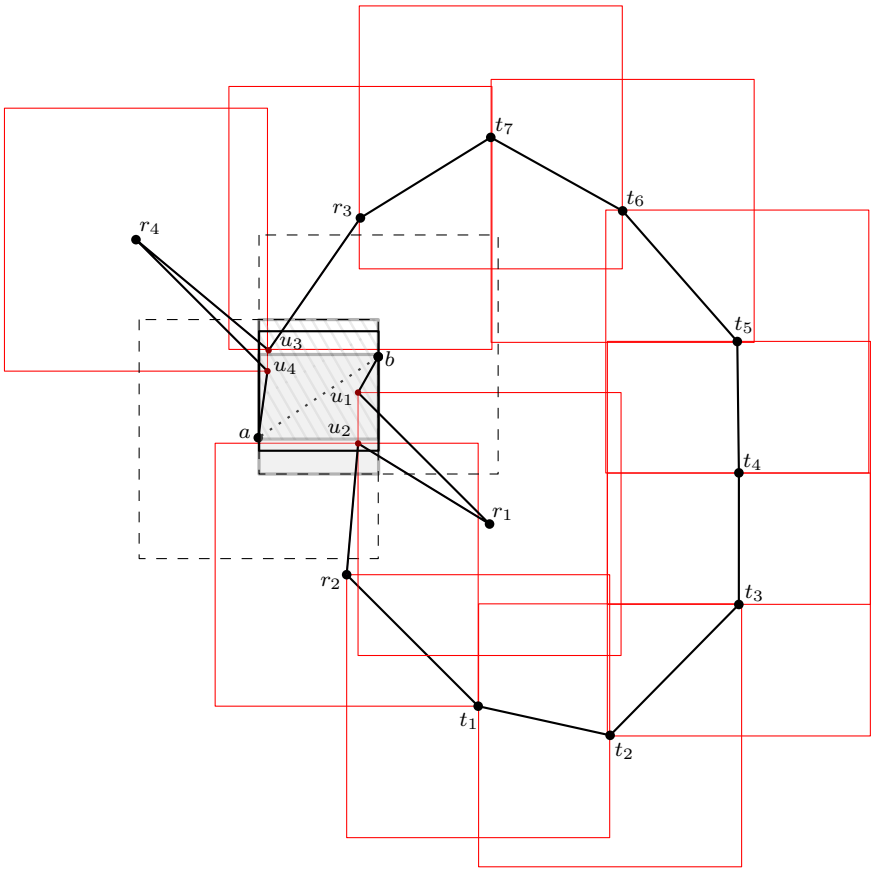

Figure 4.11: The diagonal-pattern square is a $\mathcal{C}(a,b)$ with a as a vertex, and the gray-filled square is a $\mathcal{C}(a,b)$ with b as vertex. The union of both squares contains all the possible $\mathcal{C}(a,b)$. The bold edges belong to h in the proof of Lemma 4.4.1.

Proof. Consider the point set S in Figure 4.11. The length of edge ab is $d_{\square}(a,b) = 1$, and the two dashed squares have radius 1 and are centered at a and b. Notice that any $\mathcal{C}(a,b)$ contains at least 3 points from $U = \{u_1, u_2, u_3, u_4\}$, so $ab \notin 2\text{-}GG_{\square}(S)$.

Let $R = \{r_1, r_2, r_3, r_4, t_1, \ldots t_7\}$. For each point in R there is a red square centered at such point with radius $1+\varepsilon$, where ε is a small positive value. Thus, $d_{\square}(r_i, u_i) = 1 + \varepsilon$, $d_{\square}(r_i, a) > 1 + \varepsilon, d_{\square}(r_i, b) > 1 + \varepsilon$ and $d_{\square}(r_i, r_j) > 1 + \varepsilon$, for $i \neq j$. The cycle h=$(a, b, u_1, r_1, u_2, r_2, t_1, t_2, t_3, t_4, t_5, t_6, t_7, r_3, u_3, r_4, u_4, a)$ is Hamiltonian and the maximum length of its edges in the $d_{\square}$-distance is $1 + \varepsilon$. Hence, any $d_{\square}$-bottleneck Hamiltonian cycle of S has at most $1 + \varepsilon$ maximum edge $d_{\square}$-length.

We will show that the edge ab is in every $d_{\square}$-bottleneck Hamiltonian cycle of S. Let h' be a $d_{\square}$-bottleneck Hamiltonian cycle. Since the $d_{\square}$-distance from a and b to any point in R is greater than $1+\varepsilon$, in h', a and b can only be connected between each other or to the points in U. Note that u_2 has to be connected to r_1 and r_2 in h', since otherwise r_1 or r_2 would be adjacent to an edge whose $d_{\square}$-length is greater than $1+\varepsilon$. Similarly, u_3 has to be connected to r_3 and r_4 in h', since otherwise r_3 or r_4 would be adjacent to an edge of $d_{\square}$-length greater than $1+\varepsilon$. Finally, a and b have to be connected to each other, since otherwise both would be adjacent to u_1 and u_4, which does not produce a Hamiltonian cycle.

In summary, ab is included in any $d_{\square}$-bottleneck Hamiltonian cycle, and since $ab \notin 2\text{-}GG_{\square}(S)$, the lemma holds. □

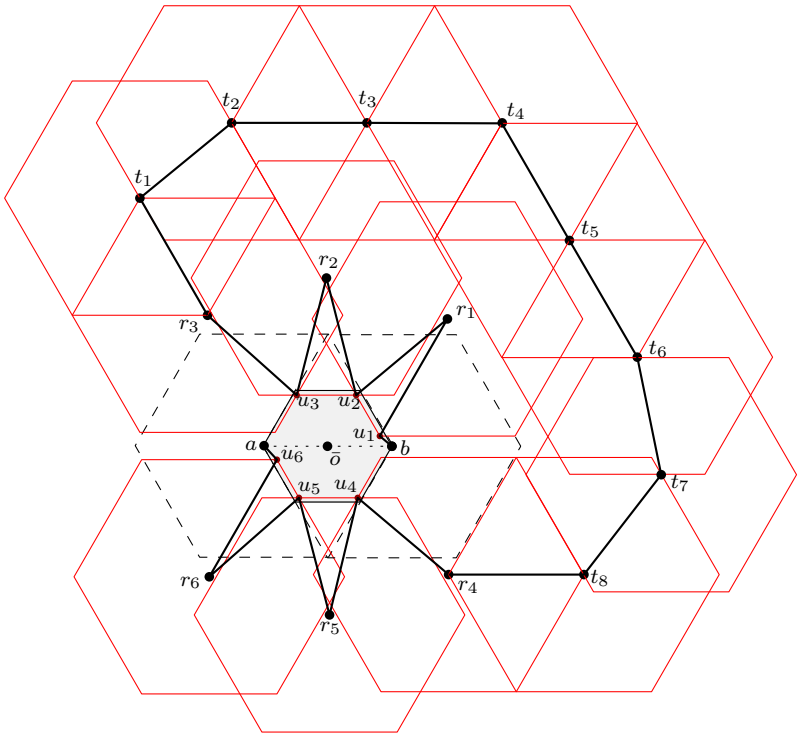

Figure 4.12: The gray hexagon is the unique $\mathcal{C}(a,b)$, and it contains 6 points of S. The bold edges belong to h in the proof of Lemma 4.4.2.

Lemma 4.4.2. There exists a point set S with $n \geq 22$ points such that 5-$GG_{\mathcal{P}_6}(S)$ does not contain any $d_{\mathcal{P}_6}$-bottleneck Hamiltonian cycle of S.

Proof. We proceed in the same fashion as in the previous proof. Consider the

point set S in Figure 4.12. The length of edge ab is $d_{\mathcal{P}_6}(a,b) = 1$, and the dashed hexagons have radius 1 and are centered at a and b. Notice that there is exactly one $\mathcal{C}(a,b)$, and it contains all points from $U = \{u_1, \ldots, u_6\}$. Therefore, $ab \notin 5\text{-}GG_{\mathcal{P}_6}(S)$. Let $R = \{r_1, \ldots, r_6, t_1, \ldots t_8\}$. For each point in R there is a red regular hexagon centered at such point with radius $1+\varepsilon$, where ε is a small positive value. Thus, $d_{\mathcal{P}_6}(r_i,u_i)=1+\varepsilon$, $d_{\mathcal{P}_6}(r_i,a)>1+\varepsilon, d_{\mathcal{P}_6}(r_i,b)>1+\varepsilon$ and $d_{\mathcal{P}_6}(r_i,r_j)>1+\varepsilon$, with $i \neq j$. The cycle $h=(a,b,u_1,r_1,\ u_2,r_2,u_3,r_3,t_1,t_2,\ldots,$ $t_8,r_4,u_4,r_5,u_5,r_6,u_6,a)$ is Hamiltonian, and the maximum length of its edges in the $d_{\mathcal{P}_6}$-distance is $1+\varepsilon$.

Let h' be a $d_{\mathcal{P}_6}$-bottleneck Hamiltonian cycle. Let us show that $ab \in h'$. Since the $d_{\mathcal{P}_6}$-distance from a and b to any point in R is greater than $1+\varepsilon$, in h', a and b can only be connected between each other or to the points in U. Note that u_3 has to be adjacent to r_2 and r_3; otherwise, r_2 or r_3 would be adjacent to an edge of $d_{\mathcal{P}_6}$-length greater than $1+\varepsilon$. Similarly, u_2, u_4, u_5 have to be adjacent to r_2 and r_3, to r_4 and r_5, and to r_5 and r_6, respectively. Finally, a and b have to be connected to each other, otherwise both would be adjacent to u_1 and u_6 which does not produce a Hamiltonian cycle. Therefore, ab is included in any $d_{\square}$-bottleneck Hamiltonian cycle, and since $ab \notin 5\text{-}GG_{\mathcal{P}_6}(S)$, the lemma holds. □

4.5 Non-Hamiltonicity for regular polygons

Until now we have discussed upper and lower bounds for k, so that $k\text{-}GG_{\mathcal{C}}$ contains a bottleneck Hamiltonian cycle. As mentioned in Section 1.4, $k\text{-}GG_{\mathcal{C}} \subseteq k\text{-}DG_{\mathcal{C}}$, thus all upper bounds given in the previous sections hold for k-order $\mathcal{C}$-Delaunay graphs as well, but not the lower bounds. In this section we present point sets for which $DG_{\mathcal{P}_t}$ is not Hamiltonian. For $t=4$, Saumell [92] showed that for any $n \geq 9$ there exists a point set S such that $DG_{\square}(S)$ is non-Hamiltonian, so we focus on $t \geq 5$.

First, we present particular cases of $t \geq 5$ for which $DG_{\mathcal{P}_t}$ is non-Hamiltonian. Later on, we present a generalization of these point sets and show the non-Hamiltonicity for an infinite family of $DG_{\mathcal{P}_t}$.

4.5.1 Non-Hamiltonicity for regular polygons with small number of sides

In this section we prove that $DG_{\mathcal{P}_t}$ fails to be Hamiltonian for every point set when $t = 5, 6, \ldots, 11$ (see Figure 4.13).

Lemma 4.5.1. For any $n \geq 7$ and any $t \in \{5, 6, \ldots, 11\}$, there exists an n-point set S such that $DG_{\mathcal{P}_t}(S)$ is non-Hamiltonian.

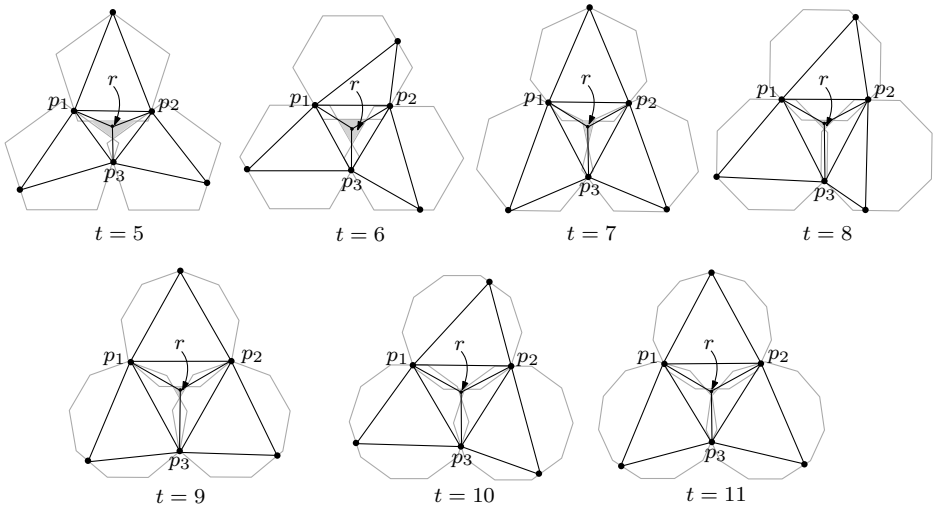

Figure 4.13: For each $t \in \{5, 6, 7, 8, 9, 10, 11\}$ the graph $DG_{\mathcal{P}_t}(S)$ is non-Hamiltonian.

Proof. Let $t \in \{5, 6, \ldots, 11\}$. Consider the graph $DG_{\mathcal{P}_t}(S)$ in Figure 4.13 for such t. Note that such graph is indeed a $\mathcal{P}_t$-Delaunay graph, since for each edge there exists a $\mathcal{P}_t$-disk that contains its vertices on its boundary and is empty of other points of S. Also, note that some edges from the convex hull of S do not appear in such graphs. Finally, notice that there exists an area r that is not contained in any of the $\mathcal{P}_t$-disks associated to the edges of the outer face or the triangle $\triangle p_1p_2p_3$. Such area can have an arbitrary number of points in its interior, say $n - 6$. Now, let $G' = DG_{\mathcal{P}_t}(S) \setminus \{p_1, p_2, p_3\}$. The graph G' consists of 4 connected components, so $DG_{\mathcal{P}_t}(S)$ is not 1-tough. Since every Hamiltonian graph is 1-tough, $DG_{\mathcal{P}_t}(S)$ is non-Hamiltonian. □

4.5.2 An infinite family of regular polygons such that $DG_{\mathcal{P}_t}$ is non-Hamiltonian

Based on the point sets given in the previous section, we construct an n-point set S, with $n \geq 7$, such that the following theorem holds.

Theorem 4.5.2. Let $\mathcal{P}_t$ be a regular t-gon, where $t > 3$ and t is an odd number and multiple of three. For any $n \geq 7$, there exists an n-point set S such that $DG_{\mathcal{P}_t}(S)$ is non-Hamiltonian.

Our construction is a generalization of the ones in the previous section. However, in order to be able to prove that $DG_{\mathcal{P}_t}(S)$ has the desired structure for arbitrary large values of t, we have to define it in a very precise way.

Before we proceed to prove Theorem 4.5.2, we need some new definitions and a few auxiliary claims.

Let $\mathcal{P}_t$ be a regular t-gon, where $t = 3(2m+1)$ for some positive integer m. Without loss of generality, we assume that $\mathcal{P}_t$ is oriented so that its bottom side is horizontal. We also assume that its vertices are given in counterclockwise order.

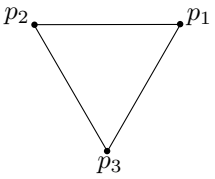

Figure 4.14: An equilateral triangle pointing downwards.

Consider three points p_1, p_2 and p_3 in the plane that define an equilateral triangle T as in Figure 4.14. Let c be the circumcenter of the triangle T. Let C_1, C_2 and C_3 be three circles circumscribing the triangles $\triangle p_1p_2c$, $\triangle p_2p_3c$ and $\triangle p_3p_1c$, respectively. These three circles are *Johnson circles*,[6] they have the same radius r, and they intersect at c. Let c_1, c_2 and c_3 be the centers of C_1, C_2 and C_3, respectively. Notice that the line segments p_2c_2 and c_3p_1 are vertical, and that $\angle cc_ip_{i+1} = \frac{\pi}{3}$ and $\angle p_ic_ic = \frac{\pi}{3}$, for all i modulo 3. See Figure 4.15.

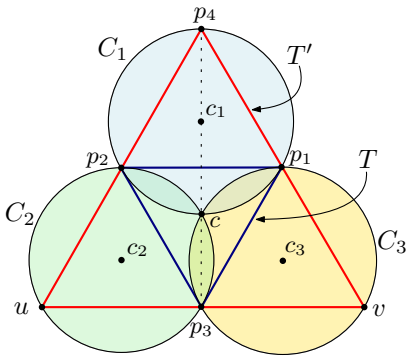

Figure 4.15: The circles C_1, C_2, C_3 contain triangles $\triangle p_1p_2c$, $\triangle p_2p_3c$ and $\triangle p_3p_1c$, respectively. The big triangle $T' = \triangle p_4uv$ is the anticomplementary triangle of T.

Consider the *anticomplementary triangle*[7] $T' = \triangle p_4uv$ of T defined as in

[6] A set of *Johnson circles* is a set of three circles of the same size that mutually intersect each other in a single point. For a survey about the properties of Johnson circles, we refer the reader to the *Johnson Theorem* [80].

[7] Let C be the circle with center $C_1 \cap C_2 \cap C_3$ and radius $2r$. The *anticomplementary*

Figure 4.15. Let $\mathcal{P}_t^1, \mathcal{P}_t^2$ and $\mathcal{P}_t^3$ be the three t-gons inscribed in C_1, C_2 and C_3, respectively. See Figure 4.16.

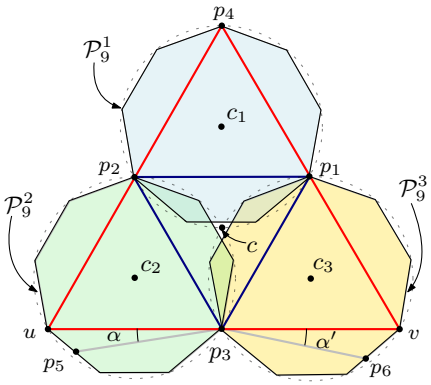

Figure 4.16: Inscribed t-gons for $t = 9$. The angles α and α' are less than $\frac{\pi}{9}$.

Claim 4.5.3. The points p_1, p_2 and p_3 are on $\partial\mathcal{P}_t^3 \cap \partial\mathcal{P}_t^1, \partial\mathcal{P}_t^1 \cap \partial\mathcal{P}_t^2$ and $\partial\mathcal{P}_t^2 \cap \partial\mathcal{P}_t^3$, respectively.

Proof. Recall that $t = 3(2m+1)$ with $m > 0$. Let a_1b_1 be the bottom side of $\partial\mathcal{P}_t^1$. Since the line segment c_1c is vertical, c_1c bisects a_1b_1. Thus, the angle formed by c_1c and the i-th vertex of $\partial\mathcal{P}_t^1$ is given by $\frac{\pi}{t} + \frac{2(i-1)\pi}{t}$. In particular, for $i = m+1$ we obtain $\frac{2m\pi}{t} + \frac{\pi}{t} = \frac{(2m+1)\pi}{3(2m+1)} = \frac{\pi}{3}$, which is precisely $\angle cc_1p_2$. Hence, $p_2 \in \partial\mathcal{P}_t^1$. The proof for $p_1 \in \partial\mathcal{P}_t^1$ is symmetric.

Since the bottom sides of $\partial\mathcal{P}_t^2$ and $\partial\mathcal{P}_t^3$ are horizontal, the top-most vertices of $\partial\mathcal{P}_t^2$ and $\partial\mathcal{P}_t^3$ are p_2 and p_1, respectively. Therefore $p_1 \in \partial\mathcal{P}_t^1 \cap \partial\mathcal{P}_t^3$ and $p_2 \in \partial\mathcal{P}_t^1 \cap \partial\mathcal{P}_t^2$.

On the other hand, since the top-most point of $\partial\mathcal{P}_t^2$ is p_2, the angle formed by p_2c_2 and the i-th vertex of $\partial\mathcal{P}_t^2$ is given by $\frac{2i\pi}{t}$. In particular, for $i = 2m+1$ we obtain $\frac{2(2m+1)\pi}{3(2m+1)} = \frac{2\pi}{3}$, which is precisely $\angle p_2c_2p_3$. Thus, $p_3 \in \partial\mathcal{P}_t^2$. Similarly, we can show $p_3 \in \partial\mathcal{P}_t^3$. □

Given points a and b, we next show how to define a polygon which we call the *$\mathcal{P}_t$-of-influence of a and b*. Recall that the vertices $v_1, \ldots, v_t$ of $\mathcal{P}_t$ are oriented counterclockwise, where v_1 is the top-most one. The i-th oriented edge of $\mathcal{P}_t$ is defined by $e_i = \overrightarrow{v_iv_{i+1}}$. We define the oriented line ℓ_i as the supporting line of the edge e_i with the same orientation as e_i. For each ℓ_i, we consider two oriented lines parallel to ℓ_i, one passing through a and another through b. Among all these lines, we only take those having a and b on its left or on the line. Now,

triangle of T has as vertices the three tangent points of C with the Johnson circles.

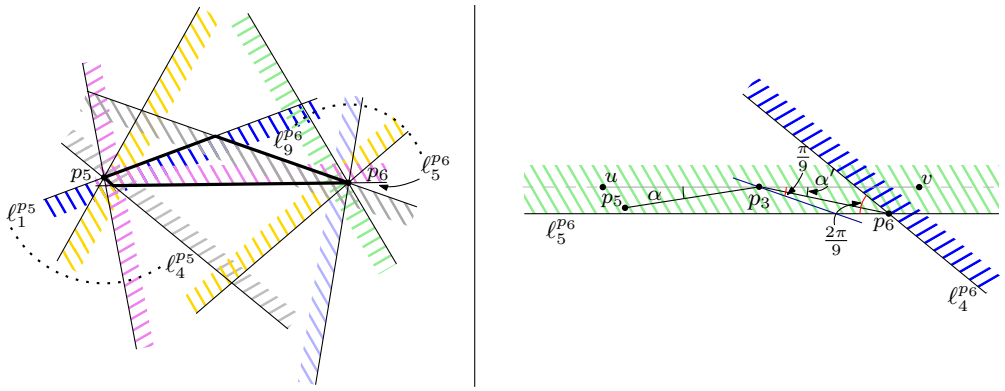

(a) (b)

Figure 4.17: The dashed area next to each line represents the half-plane with points on the left of the line. (a) The bold polygon is the $\mathcal{P}_9$-of-influence of p_5 and p_6. (b) The points p_5 and p_3 are on the left of $\ell_5^{p_6}$ and on the right of $\ell_4^{p_6}$.

consider the left half-planes defined by such oriented lines; the intersection of these half-planes defines the $\mathcal{P}_t$-of-influence of a and b. See Figure 4.17a. Since a point p is in a $\mathcal{P}_t$-disk if p is on the left of each supporting line ℓ_i or on ℓ_i, any $\mathcal{P}_t$-disk containing a and b contains their $\mathcal{P}_t$-of-influence.

Let p_5 and p_6 be two points on the boundary of $\mathcal{P}_t^2$ and $\mathcal{P}_t^3$, respectively, such that $\alpha = \angle p_5 p_3 u < \frac{\pi}{t}$ and $\alpha' = \angle v p_3 p_6 < \frac{\pi}{t}$. See Figure 4.16.

Claim 4.5.4. Any $\mathcal{P}_t$-disk containing p_5 and p_6 on its boundary contains p_3 in its interior.

Proof. Note that if p_3 is in the interior of the $\mathcal{P}_t$-of-influence of p_6 and p_5 then the claim follows. Let us show first that p_3 is in the $\mathcal{P}_t$-of-influence of p_6 and p_5 (but not necessarily in its interior). We denote by ℓ_i^p the parallel line to ℓ_i passing through point p. Without loss of generality assume that p_5 is above the horizontal line passing through p_6. Note that α' is equal to the inner angle at p_6 formed by the horizontal line passing through p_6 and edge p_6p_3, and this angle is less than $\frac{\pi}{t}$. Also, note that for $h = \frac{t-1}{2} + 1$, ℓ_h is horizontal. Finally, observe that the angle formed by consecutive ℓ_i^p and ℓ_{i+1}^p is $\frac{2\pi}{t}$. Then, p_3 is contained in the wedge defined by $\ell_h^{p_6}$ and $\ell_{h-1}^{p_6}$ with inner angle $\frac{2\pi}{t}$ that lies above $\ell_h^{p_6}$. Refer to Figure 4.17b. Since $\alpha' < \frac{\pi}{t}$, this wedge contains the wedge defined by edge p_6p_3 and $\ell_h^{p_6}$, which contains p_5. Thus, p_5 is on the left of $\ell_h^{p_6}$ and on the right of $\ell_{h-1}^{p_6}$. The lines of the form $\ell_i^{p_6}$ that have p_5 on its left are the ones encountered when rotating $\ell_h^{p_6}$ along p_6 counterclockwise until it hits p_5; the total angle of rotation is π minus the inner angle formed by p_5p_6 and

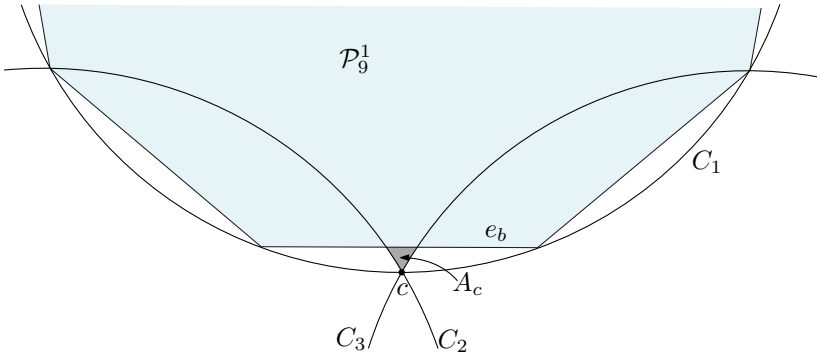

Figure 4.18: The gray area A_c is contained in the interior of $C_1 \setminus (C_2 \cup C_3)$ and is not contained in $\mathcal{P}_9^1$.

$\ell_h^{p_6}$. Therefore, these lines are $\ell_h^{p_6}, \ell_{h+1}^{p_6}, \ldots, \ell_t^{p_6}$. Since the wedge defined by $\ell_h^{p_6}$ and $\ell_t^{p_6}$ containing p_5 has angle $\frac{\pi}{t}$, p_3 also lies on such wedge and p_3 is on the left of $\ell_t^{p_6}$. Moreover, the angle of the cone containing p_5 formed by $\ell_h^{p_6}$ and $\ell_i^{p_6}$, for any $i \in \{h+1, \ldots, t\}$, is at least $\frac{\pi}{t}$. Hence, p_3 lies on the left of $\ell_i^{p_6}$ for all $i \in \{h, \ldots, t\}$. Similarly, we show that $\ell_i^{p_5}$ has p_6 on its left if and only if $i = \{1, \ldots, \frac{t-1}{2}\}$, and these lines also have p_3 on its left. Thus, p_3 is in the $\mathcal{P}_t$-of-influence of p_5 and p_6. Moreover, since p_3 is strictly on the left of all the mentioned relevant lines, p_3 is in the interior of the $\mathcal{P}_t$-of-influence of p_5 and p_6. Therefore, p_3 is in the interior of any $\mathcal{P}_t$-disk containing p_5 and p_6. □

Now, we proceed to prove Theorem 4.5.2.

Proof of Theorem 4.5.2. Since the bottom side e_b of $\mathcal{P}_t^1$ is horizontal and the intersection of the three circles C_1, C_2, C_3 is only the point c, there is an empty space in C_1 bounded by e_b and the circular arcs of C_2 and C_3 with endpoints c and the intersection points $e_b \cap C_2$ and $e_b \cap C_3$. Let us call such area A_c. See Figure 4.18. Let S' be a set of $n-6$ points in general position contained in A_c. Let $S = \{p_1, p_2, p_3, p_4, p_5, p_6\} \cup S'$.

Since for $i = 1, 2, 3$ the $\mathcal{P}_t$-disk $\mathcal{P}_t^i$ contains no point of S in its interior, from Claim 4.5.3 it follows that the edges p_1p_2, p_2p_3, p_3p_1 are in $DG_{\mathcal{P}_t}(S)$.[8] Also, since for each of the edges p_5p_2, p_2p_4, p_4p_1, p_1p_6, p_6p_3 and p_3p_5, its endpoints lie on $\partial\mathcal{P}_t^i$ for some fixed $i \in \{1, 2, 3\}$, such edges are in $DG_{\mathcal{P}_t}(S)$. By Claim 4.5.4, $p_5p_6 \notin DG_{\mathcal{P}_t}(S)$. Hence, the outerface of $DG_{\mathcal{P}_t}(S)$ is given by the edges $p_5p_2, p_2p_4, p_4p_1, p_1p_6, p_6p_3$ and p_3p_5.

[8]Notice that each of the $\mathcal{P}_t^i$ satisfies the following property: For any two of the three points of S on its boundary, the $\mathcal{P}_t$-disk can be slightly perturbed so that the two chosen points remain on its boundary and the third point lies in the exterior of the $\mathcal{P}_t$-disk.

The graph $DG_{\mathcal{P}_t}(S)$ is not 1-tough because $DG_{\mathcal{P}_t}(S) \setminus \{p_1, p_2, p_3\}$ consists of four connected components, namely, $\{p_4\}, \{p_5\}, \{p_6\}$ and $DG_{\mathcal{P}_t}(S')$. Therefore, $DG_{\mathcal{P}_t}(S)$ is not Hamiltonian. □

4.6 Conclusions

In this chapter we have presented the first general results on Hamiltonicity for higher-order convex-shape Delaunay and Gabriel graphs. By combining properties of metrics and packings, we have achieved general bounds for any convex shape, and improved bounds for point-symmetric shapes, as well as for even-sided regular polygons. For future research, we point out that our results are based on bottleneck Hamiltonian cycles, in the same way as all previously obtained bounds [2, 39, 67]. However, in several cases, as we show in Section 4.4, this technique is reaching its limit. Therefore a major challenge to effectively close the existing gaps will be to devise a different approach to prove Hamiltonicity of Delaunay graphs.

An Affine Delaunay triangulation and more

5

Affine invariance of various geometric structures is an important property in areas like graphics and computer aided geometric design. In the context of triangulations, consider a triangulation algorithm T, which given a point set S computes a triangulation $T(S)$. We say that T is *affine invariant* if and only if for any invertible *affine transformation* α (see Section 5.1 for a formal definition), the triangulations $\alpha(T(S))$ and $T(\alpha(S))$ are the same; i.e., triangle $\triangle pqr$ is in $T(S)$ if and only if triangle $\triangle\alpha(p)\alpha(q)\alpha(r)$ is in $T(\alpha(S))$. Note that α is not known to the triangulation algorithm.

Besides the Delaunay triangulation, another famous triangulation is the *minimum weight triangulation*, denoted MWT, which minimizes the sum of the Euclidean-length of its edges. The Delaunay triangulation may fail to be a minimum weight triangulation by a factor of $\Theta(n)$ where n is the size of S [69]. It is easy to see that neither the Delaunay triangulation nor the minimum weight triangulation is affine invariant in general (see Figure 5.1). This is because non-uniform stretching can make a point previously outside of a circumcircle become inside, or edge lengths can increase non-uniformly.

Affine invariance is also important in the analysis and visualization of data, to guarantee for instance that different units of measurement do not influence the geometric structure, such as a triangulation that is computed. For this reason, Nielson [85] defined an affine invariant normed metric A_S of a point set S, denoted A_S-norm, where for each point $v \in S$ and any affine transformation α, $A_S(x) = A_{\alpha(S)}(\alpha(x))$. The A_S-norm produces ellipses (see Figure 5.2) as the boundary of the A_S-norm disk and using this notion Nielson [86] defined an A_S-Delaunay triangulation that is affine invariant. Nielson's approach does not distinguish if the point set is rotated or reflected. While this is not an issue to obtain an affine invariant Delaunay triangulation, it makes the method

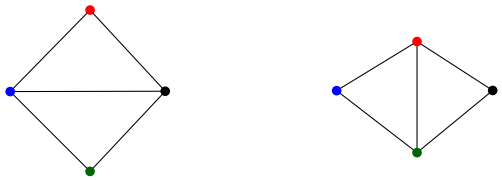

Figure 5.1: The points on the left correspond to an affine transformation of the points on the right, where each point is mapped to the point with the same color. The DT and MWT (equal in this case) differ between the left and right points, hence they are not affine invariant.

unsuitable to construct other triangulations or geometric objects, like the ones discussed in Section 5.3. Surprisingly the whole topic of affine invariant geometric constructions has gone virtually unnoticed in the computational geometry literature. Note that the use of the title "Affine invariant triangulations" have also been used by Haesevoets et al. [64], but with different meaning, to study affine invariant triangulation methods of the union of given triangles at different times (two triangles can overlap in different regions).

Our contributions. We revisit the A_S-norm and explain the geometry behind it in order to understand how the A_S-Delaunay triangulation behaves. We show that such triangulations have a spanning ratio related to the spanning ratio of the standard Delaunay triangulation, and that the hierarchy of subgraphs of the Delaunay triangulation, such as the minimum spanning tree [102] or the relative neighborhood graph [102] is also affine invariant. In addition, we show how to use the A_S-norm to compute different affine invariant orderings of a point set (radial order, sweep line ordering, and a polygon traversal ordering). Using these affine invariant orderings as subroutines, we can adapt standard geometric algorithms for computing a triangulation of a point set or a polygon to become affine invariant.

5.1 Preliminaries

An *affine transformation* $\alpha : X \to Y$ is a function of the form $\alpha(v) = Mv + b$ where X and Y are affine space mapped, M is a linear transformation on each vector in X, and b is a vector in Y. In this chapter we will work in $\mathbb{R}^2$, i.e., $X = Y = \mathbb{R}^2$, M is a matrix in $\mathbb{R}^2 \times \mathbb{R}^2$ and b is a vector in $\mathbb{R}^2$. For the rest of this chapter we will not distinguish a point from a vector unless notation is confusing, and we will assume that α is invertible, i.e., it is a non-degenerate function and $det(M) \neq 0$. Let S be a point set, we denote $\alpha(S)$ the

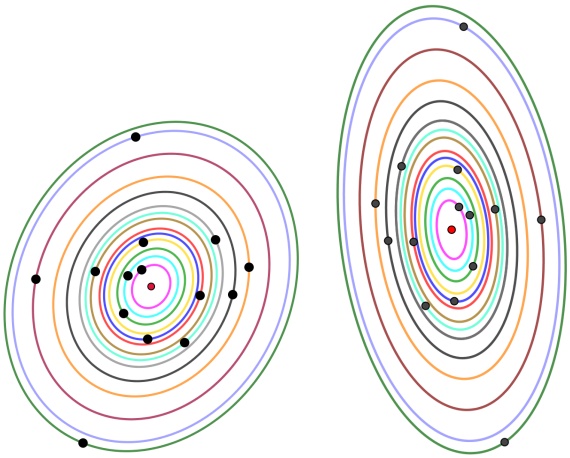

(a) Point set S (b) Point set S'

Figure 5.2: Point set S' is an affine transformation of S. Each color of the ellipses represents the corresponding boundary of the A_S- and $A_{S'}$-disk centered at the red point (mean) and containing the corresponding point of each transformation.

set $\{\alpha(v) : v \in S\}$. The following proposition states some well-known properties of affine transformations.

Proposition 5.1.1. [38] Let $\alpha(v) = Mv + b$ be an affine invertible transformation on the affine space X and let S be a point set in X. Then function α:

1. maps lines (resp., line segments) to lines (resp., line segments),
2. preserves parallelism between lines and line segments,
3. maps a simple n-gon to a simple n-gon,
4. preserves the ratio of lengths of two parallel segments,
5. preserves the ratio of areas of two object, and
6. maps the mean of S to the mean of $\alpha(S)$.

Let $S = \{p_1, p_2, \ldots, p_n\}$ be a point set in the plane of size n. The *order type* of S is a mapping that assigns to each ordered triple $i, j, k \in \{1, \ldots, n\}$ the orientation of p_i, p_j and p_k (either clockwise or counterclockwise). It can be shown that order types are preserved up to a change of sign, by checking for each triple $p_i, p_j, p_k \in S$ the signed area of the triangles $\triangle p_i p_j p_k$ and $\triangle \alpha(p_i)\alpha(p_j)\alpha(p_k)$ given by the following cross product $\alpha(p_i - p_j) \times \alpha(p_k - p_j) = det(M)((p_i - p_j) \times (p_k - p_j))$.

We denote the coordinates of each point p_i in S by (p_{i_x}, p_{i_y}). Nielson [85] defines an affine invariant normed metric, that we call A_S-norm, in the following way.

Let

$$\mu_x = \frac{\sum_{i=1}^{n} p_{i_x}}{n}, \quad \mu_y = \frac{\sum_{i=1}^{n} p_{i_y}}{n}, \quad \sigma_x^2 = \frac{\sum_{i=1}^{n} (p_{i_x} - \mu_x)^2}{n}$$
$$\sigma_y^2 = \frac{\sum_{i=1}^{n} (p_{i_y} - \mu_y)^2}{n}, \quad \sigma_{xy} = \frac{\sum_{i=1}^{n} (p_{i_x} - \mu_x)(p_{i_y} - \mu_y)}{n}.$$

Note that the mean $\mu = (\mu_x, \mu_y)$ is the barycenter of S. The *covariance matrix* of a point set S is defined as $\Sigma = \begin{pmatrix} \sigma_x^2 & \sigma_{xy} \\ \sigma_{xy} & \sigma_y^2 \end{pmatrix}$ and its inverse as $\Sigma^{-1} = \frac{1}{\sigma_x^2\sigma_y^2 - \sigma_{xy}^2} \cdot \begin{pmatrix} \sigma_y^2 & -\sigma_{xy} \\ -\sigma_{xy} & \sigma_x^2 \end{pmatrix}$.

Then, the A_S-norm metric is defined as $A_S(p_x, p_y) = \begin{pmatrix} p_x & p_y \end{pmatrix} \Sigma^{-1} \begin{pmatrix} p_x \\ p_y \end{pmatrix}$.

The matrix Σ^{-1} is also known as the *concentration matrix* [51], which defines a norm with respect to the normal (Gaussian) distribution defined by S. The eigenvectors of Σ and Σ^{-1} are the same and these vectors define the principal orthogonal directions of how spread the point set is with respect to its mean (barycenter) μ. In other words, if we compute the Gaussian manifold defined by the bivariate normal distribution given by the point set S and then cut the Gaussian manifold with a plane parallel to the plane $z = 0$, then we obtain an ellipse. See Figure 5.2. Such an ellipse has principal orthogonal axes defined by the eigenvectors of Σ^{-1}. Thus, the boundary of an A_S-disk will be defined by a homothet of the resulting ellipse. In addition, the boundary of the *unit A_S-disk* will be represented by the ellipse with principal axes being parallel to the eigenvectors of Σ^{-1}, and the magnitude of each principal axis will be given by the square root of the eigenvalue of the corresponding unit eigenvector.

Throughout this chapter, we say that a point set S is in *general position* if no three points are collinear and all points in S are at different A_S-norm distance from the mean μ. Since the boundary of the A_S-disk is an ellipse, Nielson computes the A_S-Delaunay triangulation, denoted $DT_{A_S}(S)$, where an edge uv is in $DT_{A_S}(S)$ provided that it satisfies the A_S-disk empty property, and shows the following.

Theorem 5.1.2 (Nielson [86])**.** The triangulation represented by $DT_{A_S}(S)$ is affine invariant.

5.2 The A_S-Delaunay triangulation revisited

In this section we discuss the connection between the standard Delaunay triangulation and the A_S-Delaunay triangulation.

Let S be a point set in $\mathbb{R}^2$ in general position. Consider the $2 \times n$ matrix N such that for each point p in S there is one column in N represented by the vector $\overline{p-\mu}$. Then, one can check that $\Sigma = \frac{1}{n}NN^T$. If a point set $S' = \alpha(S)$ and $\alpha(\bar{p}) = M\bar{p} + b$ with $p \in S$, is an affine transformation of the point set S, then its mean is given by $\alpha(\mu)$ and the covariance matrix Σ' of S' is given by $\Sigma' = M\Sigma M^T$, recall that M is a 2×2 matrix in $\mathbb{R}^2$ with $det(M) \neq 0$.

Since S is in general position, $det(\Sigma) \neq 0$. Thus, Σ is invertible. Moreover, since Σ is a square symmetric matrix, $\Sigma = Q\Lambda Q^T$ where Q is the matrix of eigenvectors of Σ, Λ is the diagonal matrix of eigenvalues and $Q^{-1} = Q^T$. Recall that if a matrix M is diagonal, then $M = M^T$ and that M^t with $t \in \mathbb{R}^+$ has entries $M^t_{i,j}$. Therefore, we can also rewrite the covariance matrix as $\Sigma = (Q\Lambda^{\frac{1}{2}})(Q\Lambda^{\frac{1}{2}})^T$. Furthermore, Q represents a rotation matrix and $\Lambda^{\frac{1}{2}}$ represents the scaling factor of a point set where the covariance matrix is the identity matrix $\mathbb{I}$. Looking carefully at this representation of Σ and Σ' above, we obtain that $(Q\Lambda^{\frac{1}{2}})^{-1}S$ is an affine transformation of S with $\mathbb{I}$ as its covariance matrix. We refer to the point set $(Q\Lambda^{\frac{1}{2}})^{-1}S$ as the point set S *normalized*. Note that the unit $A_{(Q\Lambda^{\frac{1}{2}})^{-1}S}$-disk is the Euclidean unit disk. This implies that the $A_{(Q\Lambda^{\frac{1}{2}})^{-1}(S)}$-Delaunay triangulation of $(Q\Lambda^{\frac{1}{2}})^{-1}S$ is given by the Delaunay triangulation of $(Q\Lambda^{\frac{1}{2}})^{-1}S$, which together with Theorem 5.1.2 proves the following proposition.

Proposition 5.2.1. Let S be a point set in general position in $\mathbb{R}^2$ and let $\Sigma = Q\Lambda Q^T$ be its covariance matrix where Q is the matrix of eigenvectors of Σ and Λ is the diagonal matrix of eigenvalues. Then, $DT((Q\Lambda^{\frac{1}{2}})^{-1}S) = DT_{A_S}(S)$.

Let $S' = \alpha(S)$ be an affine transformation of the point set S with covariance matrix $\Sigma' = Q'\Lambda'Q'^T$. Notice that even though the point sets $(Q\Lambda^{\frac{1}{2}})^{-1}S$ and $(Q'\Lambda'^{\frac{1}{2}})^{-1}S'$ have the Euclidean metric as the A_S-norm, both point sets can be different, since the unit eigenvectors are the ones mapped to the unit base $\{(1,0),(0,1)\}$. In other words, $(Q\Lambda^{\frac{1}{2}})^{-1}S$ and $(Q'\Lambda'^{\frac{1}{2}})^{-1}S'$ are the same point set *up to rotations and reflections*.

For instance, in Figure 5.3 we have on the left a point set that defines the Euclidean metric as its A_S-norm, in the middle we have the point set resulted from a rotation and scaling of the point set on the left. Finally, the point set on the right is a normalization of the point set in the middle. Note that the two point sets on the right and left are different by a rotation.

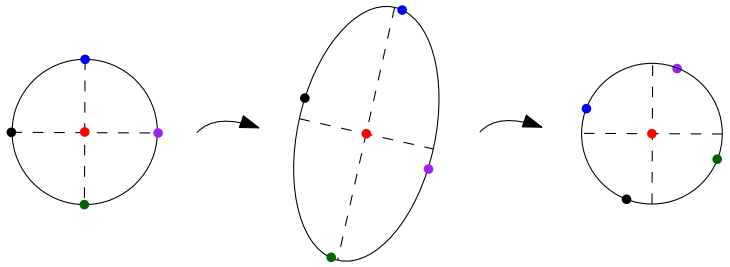

Figure 5.3: (*left*) Point set with A_S-norm the Euclidean metric, (*middle*) an affine transformation of point set in the left, (*right*) a resulting point set when normalizing the point set in the middle.

Another example is depicted in Figure 5.4, where we have two point sets in the middle that can be transformed into each other by an affine transformation – one is a reflection of the other by the x-axis. When we normalize each of the two middle point sets, we obtain the point set on the extreme left and right, respectively. Note that such these two point sets are different. Moreover, the point sets on the left and right have different order type.

On the other hand, a nice implication of Proposition 5.2.1 is that the A_S-Delaunay triangulation behaves in many ways like a standard Delaunay triangulation. For instance, from the results given by Dillencourt [49], it follows that the $DT_{A_S}(S)$ contains a perfect matching when $|S|$ is even, and that it is 1-tough.

It is known that the standard Delaunay triangulation is a spanner, see [50, 68, 108]. The following theorem shows that $DT_{A_S}(S)$ is also a spanner whose spanning ratio is a constant factor times the spanning ratio of the standard Delaunay triangulation. The constant factor depends on the eigenvalues of the co-variance matrix.

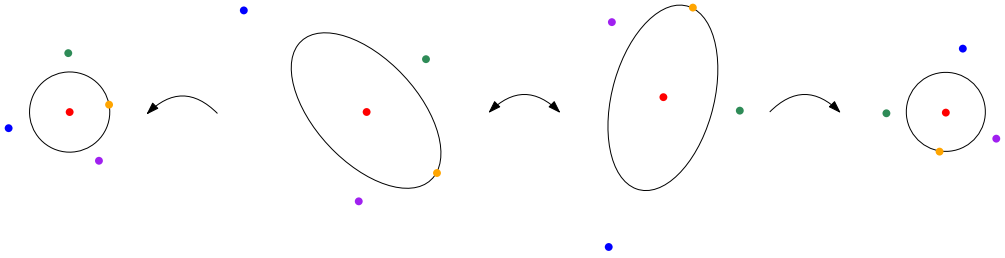

Figure 5.4: The two point sets in the middle are the initial point sets, one is an affine transformation of the other one. The point set on the left is the obtained point set when normalizing the left-middle point set and the point set on the right is the obtained point set when normalizing the right-middle point set.

Theorem 5.2.2. Let S be a point set in general position and let $\Sigma = Q\Lambda Q^T$ be the covariance matrix of S. Let $\lambda_{\max}$ and $\lambda_{\min}$ be the maximum and minimum eigenvalues of Σ, respectively. Then,

$$sr(DT_{A_S}(S)) \leq \left(\frac{\lambda_{\max}}{\lambda_{\min}}\right)^{\frac{1}{2}} \cdot sr(DT((Q\Lambda^{\frac{1}{2}})^{-1}S)).$$

Proof. Let $S' = (Q\Lambda^{\frac{1}{2}})^{-1}S$. The triangulation $DT(S')$ is a standard Delaunay triangulation. For every pair of points $u, v \in S'$ let δ_{uv} be a shortest path from u to v contained in $DT(S')$. By definition the spanning ratio $\frac{\sum_{(p_i,p_j)\in\delta_{uv}} d(p_i,p_j)}{d(u,v)} \leq sr(DT(S'))$.

Note that the only thing that changes the spanning ratio is when the graph $DT(S')$ is stretched with different scaling factors in the x- and y-coordinates. As discussed previously, such scaling is defined by the square root of the eigenvalues of Σ given in the diagonal matrix Λ. Hence, for any $u, v \in S'$ we have $d(\Lambda^{\frac{1}{2}}u, \Lambda^{\frac{1}{2}}v) \leq \lambda_{\max}^{\frac{1}{2}} \cdot d(u,v)$ and $d(\Lambda^{\frac{1}{2}}u, \Lambda^{\frac{1}{2}}v) \geq \lambda_{\min}^{\frac{1}{2}} \cdot d(u,v)$. Therefore, $sr(DT_{A_S}(S)) \leq (\frac{\lambda_{\max}}{\lambda_{\min}})^{\frac{1}{2}} \cdot sr(DT(S'))$. □

5.3 Primitives for other affine invariant geometric constructions

A natural line of study to follow is to consider other geometric objects with affine invariant construction algorithms, such as algorithms for triangulating a point set besides the A_S-Delaunay triangulation, triangulating a simple polygon, or computing a k-angulation of a point set, among others. In this section we identify some ingredients for defining such methods.

At first sight, when such constructions rely on a metric, then the A_S-norm can be used, for instance to compute an affine invariant Delaunay triangulations or a minimum weight triangulation. In general, these types of algorithms can be split into two categories: (1) based on an empty disk property, and (2) based on the rank of the length of the $\binom{n}{2}$ edges of a point set. Later we present different examples of algorithms that result to be affine invariant using the A_S-norm.

However, this does not always work. Many algorithmic techniques rely on the given order of the points, such as a radial order or the order obtained by sweeping the point set in a given direction. If we can use the A_S-norm in order to normalize S by mapping all of its points to a point set S' where the $A_{S'}$-norm is the Euclidean distance, then we can apply the standard algorithms in S' for obtaining an affine invariant radial order or sweep-line order. Yet, as noted in Section 5.2, this is not enough. This only solves the scaling factors on the x- and y-coordinates, but it does not solve rotations or reflections. For instance,

if S is a point set such that A_S-norm is the Euclidean distance and S' is a rotation of S, then the boundary of the $A_{S'}$-disk is also defined by a circle, and there is no transformation needed. However, the resulting order of a sweep-line with respect to the x-coordinates for S and S' will not necessarily be the same. Another simple example is when P is a simple polygon with vertex set S, such that S has as A_S-disk the Euclidean disk. Consider point set S' obtained by reflecting P with respect to the x-axis, and the resulting polygon P'. Again, S' will have as $A_{S'}$-disk the Euclidean disk. However, the clockwise order of S and S' might be different.

Thus, the next question is, what do we need in order to create affine invariant sorting methods? For any sorting method, we always need to define which point is the initial point, so, we have to be able to choose a point $\wp$ in the point set S, such that for any affine transformation α, we always choose $\alpha(\wp)$ as the initial point. Second, for the cases of radial order and sweep-line we need a ray $\overrightarrow{\wp\mu}$ and a line $\ell_\wp$, respectively, such that for any affine transformation α, we always choose $\overrightarrow{\alpha(\wp)\alpha(\mu)}$ and line $\alpha(\ell_\wp)$, respectively. Finally, since there might be reflections, which change clockwise for counterclockwise and left for right of a ray, we need to be able to choose the right orientation (clockwise or counterclockwise, right or left) whether there is a reflection or not. We can solve this problem if we are able to choose a point δ, such that for any affine transformation α, we always choose $\alpha(\delta)$. Then the direction is given depending on which side of the plane $\alpha(\delta)$ lies with respect to $\overrightarrow{\alpha(\wp)\alpha(\mu)}$.

We say that a function f of S is affine invariant if and only if $\alpha(f(S)) = f(\alpha(S))$ for any affine transformation α. For instance, the median $f(S) = \mu$ is an affine invariant function. Hence, note that if we manage to define affine invariant functions that compute a point, a ray, and an oriented line, there is no need for mapping S to a point set S' with $A_{S'}$-norm as the Euclidean metric.

Let S be a point set in general position. Let u, v be two points in S. We say that point u *is on the right of* $\overrightarrow{wz}$ if the signed area of $\triangle zwu$ is positive. Otherwise, u *is on the left of* $\overrightarrow{wz}$. Let $\alpha(x) = Mx+b$ be an affine transformation of $\mathbb{R}^2$. Using that $\alpha(w-z) \times \alpha(w-u) = det(M)((w-z) \times (w-u))$ one can check that if $det(M)<0$, then the orientation of $\alpha(u)$ with respect to $\overrightarrow{\alpha(w)\alpha(z)}$ changes. Otherwise, the orientation remains the same. We say that u and v *lie on the same side of a directed line segment* $\overrightarrow{wz}$ if $sign((w-z) \times (w-u)) = sign((w-z) \times (w-v))$.

Observation 5.3.1. Let $\overrightarrow{wz}$ be a directed line segment in $\mathbb{R}^2$ and let v, u be two points in S. Let α be an affine transformation. The points u and v are on the same side of $\overrightarrow{wz}$ if and only if $\alpha(u)$ and $\alpha(v)$ are on the same side of $\overrightarrow{\alpha(w)\alpha(z)}$.

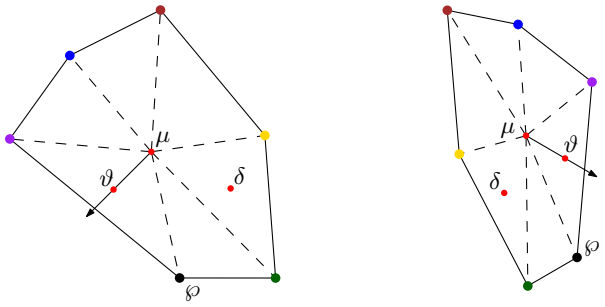

Figure 5.5: The bold polygons define the convex hull of two point sets that can be transform into each other by an affine transformation. Equal color indicate pairs of points and their affine transformations. The triangle defined by the points μ, black and purple is the triangle with largest area. In addition, the triangle defined by the points μ, yellow and green is the triangle with second largest area.

Proof. Since u and v are on the same side of $\overrightarrow{wz}$, $sign((w-z)\times(w-u)) = sign((w-z)\times(w-v))$. Hence,
$sign(\alpha(w)-\alpha(z))\times(\alpha(w)-\alpha(u))) = sign(det(M))sign((w-z)\times(w-u)) =$
$sign(det(M))sign((w-z)\times(w-v)) = sign((\alpha(w)-\alpha(z))\times(\alpha(w)-\alpha(v)))$. □

Note that there may be several ways of implementing affine invariant functions that compute the desired primitives for sorting: a point $\wp \in S$, a ray $\overrightarrow{\wp\mu}$ and a point δ that defines an orientation with respect to $\overrightarrow{\wp\mu}$. In the rest of this section we provide two different procedures for defining the desired primitives.

In the following observation, we use the convex hull and the barycenter of S in order to compute the primitives.

Observation 5.3.2. Let S be a point set in general position in the plane. Let μ be the mean of S and let $CH(S)$ be the convex hull of S. For each edge $p_ip_{i+1} \in CH(S)$, consider the triangle $p_ip_{i+1}\mu$. Assume that the areas of these triangles are pairwise different. Let ϑ and δ be the barycenters of the triangles with largest and second largest area, respectively, denoted $\triangle_{B_1}$ and $\triangle_{B_2}$, respectively. Consider the ray $\overrightarrow{\mu\vartheta}$. If δ is on the left of $\overrightarrow{\mu\vartheta}$, then let $\wp$ be the vertex of $\triangle_{B_1}$ that is on the left of $\overrightarrow{\mu\vartheta}$. Otherwise, let $\wp$ be the vertex of $\triangle_{B_1}$ that is on the right of $\overrightarrow{\mu\vartheta}$. The functions $f_1(S) = \vartheta, f_2(S) = \delta$ and $f_3(S) = \wp$ are affine invariant. See Figure 5.5.

Proof. Let $S' = \alpha(S)$ be an affine transformation of S. Consider the convex hull $CH(S')$ of S' and for each edge $p'_ip'_{i+1}$ of $CH(S')$ we consider the triangles $\triangle p'_ip'_{i+1}\alpha(\mu)$. Let $\triangle'_{B_1}$ and $\triangle'_{B_2}$ be the triangles with largest and second largest

area of the triangles defined by the edges of $CH(S')$ and $\alpha(\mu)$. Let $f_1(S') = \vartheta'$ and $f_2(S') = \delta'$ be the barycenters of $\triangle'_{B_1}$ and $\triangle'_{B_2}$, respectively. If δ' is on the left of $\overrightarrow{\alpha(\mu)\vartheta}$, then let $f_3(S') = \wp'$ be the vertex in $\triangle'_{B_1}$ that is on the left of $\overrightarrow{\alpha(\mu)\vartheta'}$. Otherwise, let $f_3(S') = \wp'$ be the vertex $\triangle'_{B_1}$ that is on the right of $\overrightarrow{\alpha(\mu)\vartheta'}$.

Let us show that $\alpha(f_1(S)) = \alpha(\vartheta) = \vartheta' = f_1(S'), \alpha(f_2(S)) = \alpha(\delta) = \delta' = f_2(S')$ and $\alpha(f_3(S)) = \alpha(\wp) = \wp' = f_3(S')$. From Proposition 5.1.1(6) the mean of S' is $\alpha(\mu)$. From Proposition 5.1.1(3) the convex hull of S is mapped to the convex hull S'. Also, from Proposition 5.1.1(3), for each edge $p_ip_{i+1} \in CH(S)$, the triangle $\triangle p_ip_{i+1}\mu$ is mapped to $\triangle\alpha(p_i)\alpha(p_{i+1})\alpha(\mu)$. Let $\triangle$ be a triangle defined by the endpoints of an edge in $CH(S)$ and μ, that is different from $\triangle_{B_1}$. Since $Area(\triangle_{B_1}) > Area(\triangle)$, $\frac{Area(\triangle_S)}{Area(\triangle)} > 1$. From Proposition 5.1.1(5) $\frac{Area(\alpha(\triangle_{B_1}))}{Area(\alpha(\triangle))} = \frac{Area(\triangle_{B_1})}{Area(\triangle)} > 1$. Thus, $Area(\alpha(\triangle_{B_1})) > Area(\alpha(\triangle))$. Therefore, $\alpha(\triangle_{B_1}) = \triangle'_{B_1}$. By the same arguments, $\alpha(\triangle_{B_2}) = \triangle'_{B_2}$. Hence, ϑ' and δ' are the barycenters of $\alpha(\triangle_{B_1})$ and $\alpha(\triangle_{B_2})$, respectively. From Proposition 5.1.1(6) follows that $\alpha(\vartheta) = \vartheta'$ and $\alpha(\delta) = \delta'$. From Proposition 5.1.1(1) we have that $\overrightarrow{\mu\vartheta}$ is mapped to $\overrightarrow{\alpha(\mu)\vartheta'}$. From Observation 5.3.1 follows that $\alpha(\wp) = \wp'$. □

Under the assumption that the area of the triangles defined by each edge of the $CH(S)$ and μ are pairwise different, we can define affine invariant functions that compute a point $\wp$ of S, a ray $\overrightarrow{\wp\mu}$, a line defined by μ and any of $\wp$ and ϑ; and a direction defined by a ray $\overrightarrow{\wp\mu}$ and point δ. On the other hand, notice that using the area of a triangle formed by the mean μ of S and any pair of points, we obtain an affine invariant order of the edges of the complete graph of S such that each edge has weight equal to the area of the triangle formed by its endpoints and μ. Under the assumption that pairwise triangles have different area, we can also obtain affine invariant algorithms that are based on the rank of the edges of the complete graph.

We will say that a point u is the *A_S-closest* point to μ if u minimizes the A_S-distance to μ. Using the fact that A_S-norm is affine invariant, we observe that there is another method for distinguishing the desired primitives for affine invariant radial and sweep-line ordered.

Observation 5.3.3. Let S be a point set in general position and let μ be the mean of S and let $\alpha(S)$ be an affine transformation of S with mean μ'. The k-th A_S-closest point to μ is mapped to the k-th $A_{\alpha(S)}$-closest point to μ'.

Proof. Since S is in general position, all of the points in S are at different A_S-distance to μ. Hence, we can order the points in S in increasing order with respect to the A_S-distance from each point to μ. Since A_S is an affine

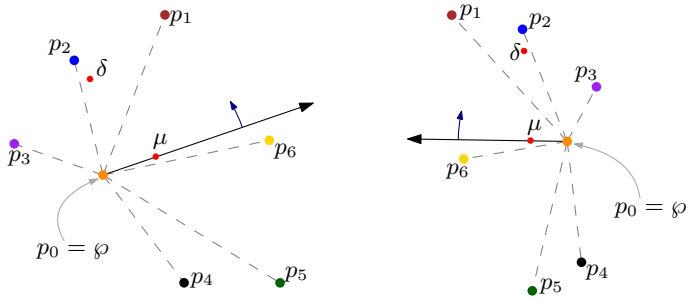

Figure 5.6: An affine invariant ordering of the point sets S (*left*) and its affine transformation $S'{=}\alpha(S)$ (*right*).

invariant norm, the points in $\alpha(S)$ are also at different $A_{\alpha(S)}$-distance to the mean $\alpha(\mu) = \mu'$. This again defines an order of $\alpha(S)$. Since the A_S-distance is invariant under affine transformations, it follows that the increasing order with respect to the A_S-distance from each point to μ is affine invariant. Thus, the k-th A_S-closest point to μ is mapped to the $A_{\alpha(S)}$-closest point to μ'. □

From Observation 5.3.3 it follows that we can also use the A_S-norm in order to define affine invariant functions $f_1(S)$ and $f_2(S)$, such that each function computes a point in the following fashion. Let μ be the mean of S, $f_1(S) = \wp$ be the A_S-closest point to μ. Let $f_2(S) = \delta$ be the second A_S-closest point to μ if it is not on the line defined by μ and $\wp$. Otherwise, let $f_2(S) = \delta$ be the third A_S-closest point to μ.

5.4 Affine invariant sorting algorithms of a point set

In this section we present two different sorting algorithms that together with the affine invariant primitives we defined before result into two different affine invariant sorting methods.

5.4.1 Affine invariant radial ordering

Let μ, δ and $\wp$ be three non-collinear points, such that $\wp$ is in S. Consider the following radial sorting procedure of S.

RadiallySort$(S, \mu, \delta, \wp)$: Sort points radially around $\wp$ in the direction of δ, starting from μ. See Figure 5.6.

Using Observations 5.3.1 and 5.3.3 we can prove that the radial order produced by RadiallySort is affine invariant.

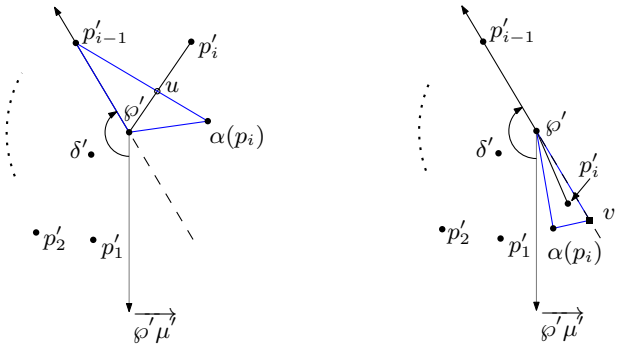

Figure 5.7: (a) p'_i is not in $T = \triangle p'_{i-1}\wp'\alpha(p_i)$ and u is the intersecting point of $\wp' p'_i$ and $\alpha(p_i)p'_{i-1}$. (b) p'_i is in the interior of $T = \triangle v\wp'\alpha(p_i)$ and the Euclidean distance from $\wp'$ to v is the same from $\wp'$ to p'_{i-1}.

Theorem 5.4.1. Let S be a point set in general position and let $f_1(S) = \mu, f_2(S) = \delta$ and $f_3(S) = \wp$ be three affine invariant functions such that $\mu, \delta, \wp$ are three non-collinear points and $\wp \in S$. Then, RadiallySort$(S, \mu, \delta, \wp)$ computes an affine invariant radial order.

Proof. Let α be an affine transformation. Let $p_0 = \wp, p_2, \ldots, p_{n-1}$ be the order of S given by RadiallySort$(S, \mu, \delta, \wp)$ and let $p'_0 = \wp', p'_2, \ldots, p'_{n-1}$ be the order of $\alpha(S)$ given by RadiallySort$(\alpha(S), f_1(\alpha(S)) = \mu', f_2(\alpha(S)) = \delta', f_3(\alpha(S)) = \wp')$, respectively. Then, $\alpha(p_0) = \alpha(\wp) = \wp' = p_0, \alpha(\mu) = \mu'$ and $\alpha(\delta) = \delta'$. It remains to show that $p'_i = \alpha(p_i)$ for all $1 \leq i \leq n-1$. If p_1 is on the ray $\overrightarrow{\wp\mu}$, then $\alpha(p_1)$ is on the ray $\overrightarrow{\wp'\mu'}$. Thus, $p'_1 = \alpha(p_1)$. Assume that p_1 is not on $\overrightarrow{\wp\mu}$.

Assume by induction that the order is affine invariant for the first $i-1$ points and let us show $\alpha(p_i) = p'_i$. Assume for the sake of a contradiction that $\alpha(p_i) \neq p'_i$. First, let us show that p'_i is on the same side as $\alpha(p_i)$ with respect to $\overrightarrow{\wp' p'_{i-1}}$ (we consider the ray $\overrightarrow{\wp'\mu'}$ when $i = 1$). Without loss of generality assume that δ' is to the right of $\overrightarrow{\wp'\mu'}$, since the other case is symmetric. If $\alpha(p_i)$ is to the right of $\overrightarrow{\wp' p'_{i-1}}$, then p'_i has to be to the right of $\overrightarrow{\wp' p'_{i-1}}$ since p'_i is hit first while rotating $\overrightarrow{\wp' p'_{i-1}}$ counterclockwise. If $\alpha(p_i)$ is to the left of $\overrightarrow{\wp' p'_{i-1}}$, then by Observation 5.3.1 p_i is to the left (resp., *to the right*) of $\overrightarrow{\wp p_{i-1}}$ if δ is to the right (resp., *to the left*) of $\overrightarrow{\wp\mu}$. Thus, there is no point of S to the right (resp., to the left) of $\overrightarrow{\wp\mu}$ if δ is to the left (to the right) of $\overrightarrow{\wp\mu}$. Hence, by Observation 5.3.1, there is no point to the left of $\overrightarrow{\wp' p'_{i-1}}$, in particular, p'_i cannot be on that side.

Therefore, $\alpha(p_i)$ and p'_i lie on the same side with respect to $\overrightarrow{\wp' p'_{i-1}}$.

Let ℓ be the line containing $\wp'$ and p'_{i-1} and let v be the point on ℓ on opposite direction of $\overrightarrow{\wp' p'_{i-1}}$ with respect to $\wp'$, such that the length from $\wp'$ to v is the same as the length from $\wp'$ to p'_{i-1}. See Figure 5.7b.

We define a triangle T in the following fashion.

- If $\alpha(p_i)$ and p'_i are to the right of $\overrightarrow{\wp' p'_{i-1}}$, then $T = \triangle p'_{i-1}\wp'\alpha(p_i)$.
- If $\alpha(p_i)$ and p'_i are to the left of $\overrightarrow{\wp' p'_{i-1}}$, then $T = \triangle v\wp'\alpha(p_i)$

Then, either p'_i is in T or not. See Figure 5.7. Again, if p'_i is contained in T then from Proposition 5.1.1(3), $\alpha^{-1}(p'_i)$ is contained in T, which contradicts that p_i appears immediately after p_{i-1} when sorting S around $\wp$. On the other hand, if p'_i is not in T, then line segment $\wp' p'_i$ intersects T on the edge opposite to $\wp'$. Let u be the intersection point of $\wp' p'_i$ and the edge of T opposite to $\wp'$. The point u is in triangle T and in line segment $\wp' p'_1$. Thus, by Properties 1 and 3 of Proposition 5.1.1, $\alpha^{-1}(u)$ is in $\alpha(T)$ and on line segment $\wp\alpha^{-1}(p'_i)$, which again contradicts the fact that p_i appears first while ordering S around $\wp$. Therefore, $p'_i = \alpha(p_i)$. □

5.4.2 Affine invariant sweep-line ordering

Let μ, δ and $\wp$ be three non-collinear points, such that $\wp$ is in S. Consider the following sweep-line order method.

SweepLine$(S, \mu, \delta, \wp)$: Let $\ell^\perp$ be the line containing μ and $\wp$. Let t and b be a farthest point from the line $\ell^\perp$ that lies on the same and opposite side of δ with respect to $\ell^\perp$, respectively. Then, sort the vertices in S by sweeping with lines parallel to $\ell^\perp$ from t towards b. If two vertices u, v lie on the same line parallel to $\ell^\perp$, we say that v appears before u if and only if δ is on the left of $\overrightarrow{\wp\mu}$ and v is on the left of $\overrightarrow{u\mu}$.

Note that a consequence of Observation 5.3.1 is that if two vertices of P lie on the same side with respect to $\ell^\perp$, then the corresponding vertices in $\alpha(P)$ lie again on the same side with respect to $\alpha(\ell^\perp)$.

Observation 5.4.2. Let ℓ be a line in $\mathbb{R}^2$ and let v, u be two points in S. Let α be an affine transformation. If u and v are on the same side of ℓ, then $\alpha(u)$ and $\alpha(v)$ are on the same side of $\alpha(\ell)$.

The following theorem shows that such ordering is affine invariant when the primitives are affine invariant. See Figure 5.8.

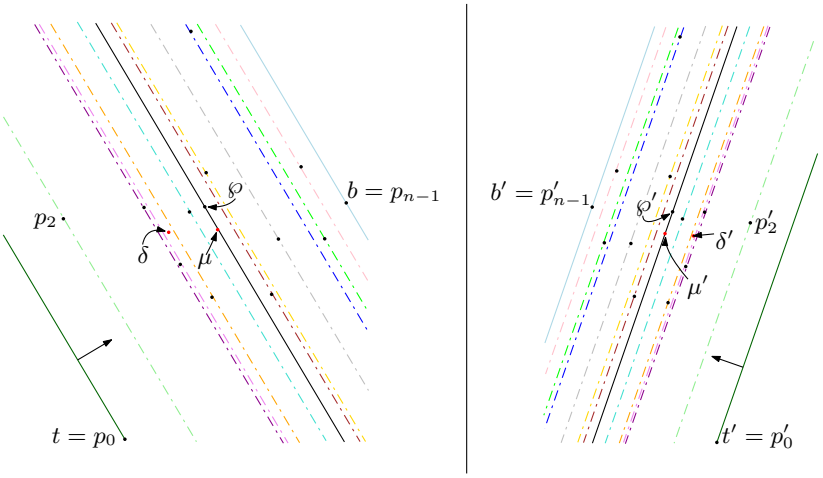

Figure 5.8: Each colored line is a line parallel to the black line containing the corresponding $f_3(S) = \wp$ and $f_1(S) = \mu$. Each color corresponds to the same line in the transformation α. SweepLine$(S, \mu, \delta, \wp)$ is affine invariant.

Theorem 5.4.3. Let S be a point set in general position and let $f_1(S) = \mu, f_2(S) = \delta$ and $f_3(S) = \wp$ be three affine invariant functions such that $\mu, \delta, \wp$ are three non-collinear points and $\wp \in S$. Then, SweepLine$(S, \mu, \delta, \wp)$ computes an affine invariant sweep-line of S.

Proof. Let $\alpha(x) = Mx + b$ be an affine transformation. Let $f_1(\alpha(S)) = \mu'$, $f_2(\alpha(S)) = \delta', f_3(\alpha(S)) = \wp'$. Then, $\alpha(\mu) = \mu', \alpha(\wp) = \wp', \alpha(\delta) = \delta'$. Let $p_0, p_1, \ldots, p_{n-1}$ be the resulting order of SweepLine$(S, \mu, \delta, \wp)$ and let p'_0, $p'_1, \ldots, p'_{n-1}$ be the resulting order of SweepLine$(\alpha(S), \mu', \delta', \wp')$. Let us show that $\alpha(p_i) = p'_i$ for all $i \in \{0, 1, \ldots, n-1\}$.

Let $\ell'^{\perp}$ be the line containing μ' and $\wp'$. From Proposition 5.1.1(1), $\alpha(\ell^{\perp}) = \ell'^{\perp}$. We denote by $\ell'^{\perp}_v$ the line parallel to $\ell'^{\perp}$ containing the point $v \in S$. Let $t' \in \alpha(S)$ be a farthest point from $\ell'^{\perp}$ that lies on the same side of δ' with respect to $\ell'^{\perp}$ and let b' be a farthest point from $\ell'^{\perp}$ that lies on the opposite side of δ' with respect to $\ell'^{\perp}$. Thus, all points of $\alpha(S)$ that are not on $\ell'^{\perp}_{t'}$ lie on the same side of $\ell'^{\perp}_{t'}$. Also, all points of $\alpha(S)$ that are not on $\ell'^{\perp}_{b'}$ lie on the same side of $\ell'^{\perp}_{b'}$.

On the other hand, t and δ are on the same side of $\ell^{\perp}$. Thus, by Observation 5.4.2, $\alpha(t)$ and δ' are on the same side of $\ell'^{\perp}$. Since all points of S that are not on $\ell^{\perp}_t$ lie on the same side of $\ell^{\perp}_t$, by Observation 5.4.2 it follows that all points of $\alpha(S)$ that are not on $\ell'^{\perp}_{\alpha(t)}$ lie on the same side of $\ell'^{\perp}_{\alpha(t)}$. Since both $\alpha(t)$ and t' lie on the same side of δ' with respect to $\ell'^{\perp}$, it follows that both lines $\ell'^{\perp}_{t'}$ and $\ell'^{\perp}_{\alpha(t)}$ lie on the same side with respect to $\ell'^{\perp}$. Finally, since there

is exactly one line parallel to $\ell'^{\perp}$ on the same side of δ' with respect to $\ell'^{\perp}$ that passes through a point of S such that all points of S that are not in such line are on the same side, it follows that $\ell'^{\perp}_{\alpha(t)} = \ell'^{\perp}_{t'}$.

Similarly, since all points of S that are not on $\ell_b^{\perp}$ lie on the same side of $\ell_b^{\perp}$, by Observation 5.4.2 it follows that all points of $\alpha(S)$ that are not on $\ell'^{\perp}_{\alpha(b)}$ lie on the same side of $\ell'^{\perp}_{\alpha(b)}$. In addition, since b is not on $\ell_t^{\perp}$, $\alpha(b)$ is not on $\ell'^{\perp}_{t'}$. Therefore, $\ell'^{\perp}_{\alpha(b)} = \ell'^{\perp}_{b'}$. Thus, the initial and the final lines while sweeping S with lines parallel to $\ell^{\perp}$ are preserved under affine transformations. Assume that the order of the lines parallel to $\ell^{\perp}$ up to the $(i-1)$-th line while sweeping S is preserved under affine transformations. Let $\ell_i^{\perp}$ be the i-th line parallel to $\ell^{\perp}$ while sweeping S. From Observation 5.4.2 it follows that for each point v of S that lies on the same side of $\ell_i^{\perp}$ as t, $\alpha(v)$ lies on the same side of $\alpha(\ell_i^{\perp})$ as t'. Since all points of S that are on the same side of $\ell_i^{\perp}$ as t have been swept by the first $i-1$ lines parallel to $\ell^{\perp}$, it follows that $\alpha(\ell_i^{\perp})$ is the i-th line parallel to $\ell'^{\perp}$ while sweeping $\alpha(S)$.

It remains to prove that if two points v and u of S lie on the same line parallel to $\ell^{\perp}$, then the order given is affine invariant. Let v and u be two points in S that lie on the same line parallel to $\ell^{\perp}$ such that v appears before u while ordering S. From Proposition 5.1.1(1), $\alpha(u)$ and $\alpha(v)$ lie on the same line parallel to $\ell'^{\perp}$. Assume δ is on the left of $\overrightarrow{\wp\mu}$, then v is on the left of $\overrightarrow{u\mu}$. Consider the following cases.

Case 1) If $det(M) > 0$, then δ' is on the left of $\overrightarrow{\wp'\mu'}$ and $\alpha(v)$ is on the left of $\overrightarrow{\alpha(u)\mu'}$. Therefore, $\alpha(v)$ appears before $\alpha(u)$.

Case 2) If $det(M) < 0$, then δ' is on the right of $\overrightarrow{\wp'\mu'}$ and $\alpha(v)$ is on the right of $\overrightarrow{\alpha(u)\mu'}$. Therefore, $\alpha(v)$ appears before $\alpha(u)$.

The case when δ is on the right of $\overrightarrow{\wp\mu}$ is symmetric. □

5.5 Applications to affine invariant geometric objects

In this section we give different application of the A_S-norm and methods given in Section 5.4 that result in affine invariant algorithms.

5.5.1 Affine invariant algorithms based on A_S-norm

Consider a norm $\rho : \mathbb{R}^2 \to \mathbb{R}^+$. In this section we mention different geometric objects of a point set S constructed using the ρ-legth of all possible $\binom{n}{2}$ edges in S. Using the fact that the A_S-norm is affine invariant, we have that the order of the edges between each pair of points in S given by the A_S-norm is also

affine invariant. Therefore, using such ordering gives us a set of affine invariant algorithms for computing geometric objects.

The first object is the MWT of S that we already defined, which minimizes the sum of the ρ-length of its edges. The *Greedy triangulation* of S is a triangulation that adds at each step an edge in strict increasing order of the N-length of the edges provided that the new edge cannot intersect in the interior of a previously inserted edge. The *ρ-minimum spanning tree* of S, denoted $MST_\rho(S)$, is a spanning tree of G with minimum total edge length. The *ρ-closest pair* of S computes the edge with lowest ρ-length. The *k-ρ-nearest neighbor graph* of S, denoted k-$NNG_\rho(S)$, is the graph with vertex set S and an edge $uv \in k$-$NNG_\rho(S)$ provided that the ρ-distance between u and v is the k-th smallest ρ-distance from u to any other point in S.

Consider an affine transformation α, since $A_S(u-v) = A_{\alpha(S)}(\alpha(u)-\alpha(v))$, we obtain that if $A_S(u-v) \leq A_S(u'-v')$, then $A_{\alpha(S)}(\alpha(u)-\alpha(v)) \leq A_{\alpha(S)}(\alpha(u')-\alpha(v'))$. Hence, the increasing and decreasing order of the A_S-length of the edges in S and $\alpha(S)$ is the same.

Observation 5.5.1. There exists an affine invariant order of the $\binom{n}{2}$ edges in any point set in general position.

Therefore, any greedy algorithm defined by the A_S-length of the edges of S is affine invariant, i.e., geometric objects that can be obtained by the rank of its edges with respect to the A_S-norm. In particular, the $MST_{A_S}(S)$ is one of them.

Finally, since the A_S-disks are affine invariant, we obtain that the empty A_S-disk property is affine invariant.

Corollary 5.5.2. Let S be a point set in general position. Then, the following objects are affine invariant: $GG_{A_S}(S)$, $RNG_{A_S}(S)$, k-$DG_{A_S}(S)$, k-$GG_{A_S}(S)$, k-$RNG_{A_S}(S)$.

In general, any object constructed based on the empty region property defined by the A_S-disk, will result affine invariant.

Note that, from Aurenhammer and Paulini [15] it follows that $MST_{A_S}(S) \subset GG_{A_S}(S) \subset GG_{A_S}(S) \subset DT_{A_S}(S)$.

5.5.2 An affine invariant Graham triangulation

One of the most popular algorithms for computing the convex hull of a point set S is Graham's scan [60]. A nice property of this algorithm is that a modification of the algorithm can also produce a triangulation, sometimes called the *Graham triangulation* [55]. In this section we present an affine invariant version of Graham's scan using the A_S-norm metric.

Our method is based on the algorithm for Graham triangulation [103]. The

method consists in first choosing a point p and then radially sorting the remaining points around p, say $p_1, \ldots, p_{n-1}$. Once the points are sorted, the algorithm adds edge pp_i for all $i \in \{1, \ldots, n-1\}$. Finally, the algorithm visits each remaining point p_i in order, and adds edge p_ip_j if and only if p_j is visible from p_i, for all $j > i$.

The Graham triangulation can be computed in linear time when S is given radially ordered. Moreover, since the edges of the triangulation are added according to the radial sort of S, it follows that if the radial order is affine invariant, then the triangulation is affine invariant.

The following result is an implication of Theorem 5.4.1.

Corollary 5.5.3. There exists an affine invariant Graham triangulation for any point set S in general position.

5.5.3 An affine invariant Hamiltonian triangulation

When the point set S has at least one point in the interior of its convex hull, then S can be triangulated by the *insertion method*, which consists of computing the convex hull of S and then inserting points from the interior in arbitrary order. Every time a point v is inserted, the edges connecting v with the points defining the face that contains v are added.

Note that since the convex hull of a point set is preserved under affine transformations, it follows that if the points that are in the interior of the convex hull are inserted in an affine invariant order, then the insertion method computes an affine invariant triangulation of S. Hence, applying the insertion method to S such that the interior points in the convex hull of S are inserted in the order given by RadiallySort computes an affine invariant triangulation.

Corollary 5.5.4. There exists an insertion method that computes an affine invariant triangulation for any point set in general position.

These triangulations are Hamiltonian, i.e., their duals [1] contain a Hamiltonian path, and are of interest for fast rendering in computer graphics [10, 55].

Corollary 5.5.5. There exist methods that compute affine invariant triangulations that are Hamiltonian for any point set in general position.

[1] The *dual* of a graph G has as vertices the faces of G and an edge between vertices is added if and only if the two faces share an edge in G.

5.5.4 An affine invariant triangulation of a polygon by ear clipping

In this section we define a traversal of a polygon that is affine invariant.

An *ear* of a polygon is a triangle formed by three consecutive vertices p_1, p_2 and p_3 such that the line segment p_1p_3 is a *diagonal*[2] of the polygon. It is a well-known fact that every simple polygon contains two ears (see Meisters [81]). By recursively locating and chopping an ear, one can triangulate any simple polygon. This method is known as *ear clipping*.

It follows from Properties 1 and 3 of Proposition 5.1.1 that the diagonals of a simple polygon are preserved under affine transformations. Thus, the ears of a simple polygon are also preserved. Hence, if the ear clipping procedure locates an ear by traversing the polygon in an affine invariant order, then such procedure computes an affine invariant triangulation. The traversal of the polygon in an affine invariant order depends only on finding an affine invariant starting point, and on deciding correctly whether to traverse it clockwise or counterclockwise. The latter depends only on whether the affine transformation contains an odd number of reflections.

Let P be a simple polygon with vertex sequence $S := \{v_1, \ldots, v_n\}$, such that S is in general position. Let μ, δ and $\wp$ be three non-collinear points such that $\wp$ is in S. Consider the following traversal of P.

Traversal$(P, \mu, \delta, \wp)$: If δ is on the left of $\overrightarrow{\mu\wp}$, then order S by traversing P from $\wp$ in counterclockwise order. Otherwise, order S by traversing P from $\wp$ in clockwise order. See Figure 5.9.

Using arguments similar to the ones used to prove Theorem 5.4.1, we show that Traversal is affine invariant.

Theorem 5.5.6. Let P be a simple polygon with vertex sequence $S := \{v_1, \ldots, v_n\}$, such that S is in general position. Let $f_1(S) = \mu$, $f_2(S) = \delta$ and $f_3(S) = \wp$ be three affine invariant functions such that $\mu, \delta, \wp$ are three non-collinear points and $\wp \in S$. Then, Traversal$(P, \mu, \delta, \wp)$ traverses P in an affine invariant order.

Proof. Let $\alpha(P)$ be an affine transformation of P. The vertices of $\alpha(P)$ are $\alpha(S)$. Consider Traversal$(\alpha(P), f_1(\alpha(S)) = \mu', f_2(\alpha(S)) = \delta', f_3(\alpha(S)) = \wp')$. By definition, we have that $\alpha(\mu) = \mu', \alpha(\wp) = \wp', \alpha(\delta) = \delta'$.

Let $p_0 = \wp, p_1, \ldots, p_{n-1}$ be the resulting order of Traversal$(P, \mu, \delta, \wp)$ and let $p'_0 = \wp', p'_1, \ldots, p'_{n-1}$ be the resulting order of Traversal$(\alpha(P), \mu', \delta', \wp')$. Let us show that $\alpha(p_i) = p'_i$ for all $i \in \{0, 1, \ldots, n-1\}$.

[2] A *diagonal* of a polygon is a line segment between two non-consecutive vertices that is totally contained inside the polygon.

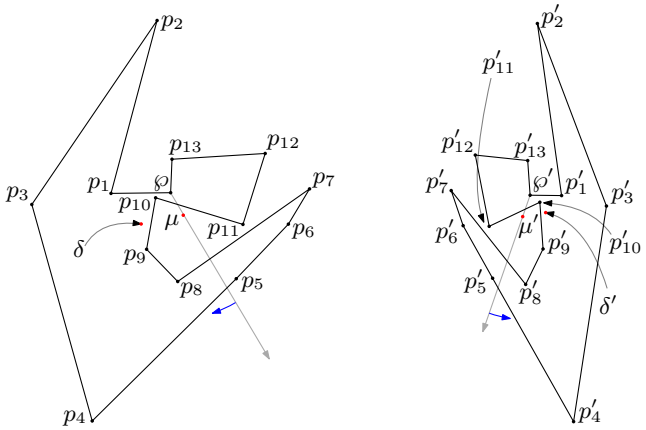

Figure 5.9: An affine invariant traversal of two simple polygons with point set S and $\alpha(S)$ from Figure 5.2.

Since the ordering is defined while traversing the simple polygon, it suffices to show that $\alpha(p_0) = p'_0$ and that the vertices are traversed in the same order. By definition, $p'_0 = \wp' = \alpha(\wp) = \alpha(p_0)$. In addition, the orientation of the points of P changes in $\alpha(P)$ when there is an odd number of reflections in α. Hence, if the orientation of the points of P changes, then δ' is on the opposite direction of $\overrightarrow{\wp'\mu'}$ as δ from $\overrightarrow{\wp\mu}$. Therefore, Traversal$(P, \mu, \delta, \wp)$ traverses P in the opposite direction of Traversal$(\alpha(P), \mu', \delta', \wp')$ traverses P' if the orientation of the points in P is changed in $\alpha(P)$. Thus, the vertices are traversed in the same order if the orientation of the points in P is opposite to the one in $\alpha(P)$. Otherwise, if the orientation does not change, then Traversal$(\alpha(P), \mu', \delta', \wp')$ traverses the points clockwise (resp., counterclockwise) if and only if Traversal$(P, \mu, \delta, \wp)$ traverses the points clockwise (resp., counterclockwise). □

The following result is an implication of Theorem 5.5.6.

Corollary 5.5.7. There exists an affine invariant ear clipping triangulation for any simple P with vertex set in general position.

5.5.5 An affine invariant triangulation of a polygon by sweep-line

In this section we give another affine invariant triangulation algorithm for any simple polygon.

A simple polygon P is *monotone with respect to a line* ℓ if for any line $\ell^\perp$ perpendicular to ℓ, the intersection of P with $\ell^\perp$ is connected. A line ℓ divides the plane into two half planes, and we say that two points *lie on the same side of* ℓ if they lie on the same half plane. Let v be a vertex of P and ℓ_v be the

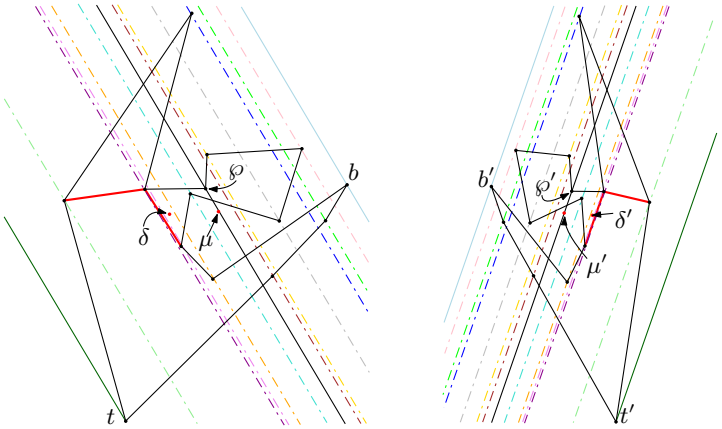

Figure 5.10: Each colored line is a line parallel to the black line. The red diagonals partition the simple polygon into monotone polygons. Each diagonal contains a cusp with respect to the black line. Each color corresponds to the same line in the transformation.

line containing v that is parallel to ℓ. We say that v is an ℓ-*cusp* if v is a reflex vertex of P and its neighbors in P lie on the same side of ℓ_v. A characterization of a monotone polygon is stated in the following property.

Property 5.5.8 (Edelsbrunner [54])**.** A polygon is ℓ-monotone if and only if it does not contain an $\ell^\perp$-cusp.

An ℓ-monotone polygon can be triangulated by sweeping its vertices with line $\ell^\perp$ (see, e.g., Garey et. al. [59]).

If P is not ℓ-monotone, then one can split P into ℓ-monotone subpolygons by adding diagonals in order to break all $\ell^\perp$-cusps. Then each resulting ℓ-monotone polygon can be triangulated independently. The way to add these diagonals is due to Lee et. al. [76] and also uses a sweep line of $\ell^\perp$. For each vertex v in P take the line $\ell_v^\perp$ and partition P into trapezoids defined by the intersection of the $\ell_v^\perp$ lines and the edges of P. If a vertex v is an $\ell^\perp$-cusp, add the diagonal of P to the vertex u that is not in the same side as the neighbors of v with respect to $\ell_v^\perp$ and u is the opposite vertex in a trapezoid containing v. See Figure 5.10.

Using Observation 5.4.2, we show that $\ell^\perp$-cusps are preserved under affine transformations.

Lemma 5.5.9. Let $\alpha(P)$ be an affine transformation of P. Each $\ell^\perp$-cusp of P is mapped to an $\alpha(\ell^\perp)$-cusp in $\alpha(P)$.

Proof. Let v be an $\ell^\perp$-cusp of P. Let us show that $\alpha(v)$ is an $\alpha(\ell^\perp)$-cusp in

$\alpha(P)$. Let v^- and v^+ be the neighbors of v in P. If v^- and v^+ are on the same side of $\ell_v^\perp$ then by Observation 5.4.2 $\alpha(v^+)$ and $\alpha(v^-)$ lie on the same side of $\ell_v^\perp$. Moreover, the triangle $\triangle v^- v v^+$ is not in P. From Proposition 5.1.1(3) follows that the triangle $\triangle\alpha(v^-)\alpha(v)\alpha(v^+)$ is not in $\alpha(P)$. Hence, the interior angle of $\alpha(v)$ in $\alpha(P)$ is greater than π. So, $\alpha(v)$ is an $\alpha(\ell^\perp)$-cusp in $\alpha(P)$. □

Therefore, if P is ℓ-monotone and $\ell^\perp$ is perpendicular to ℓ, then $\alpha(P)$ is ℓ^*-monotone where ℓ^* is perpendicular to $\alpha(\ell^\perp)$.

Thus in order to compute an affine invariant triangulation of P it remains to give an affine invariant sweep-line order of its vertices. Therefore, from Theorem 5.4.3 we obtain the following corollary.

Corollary 5.5.10. There exists an affine invariant algorithm for triangulating a polygon P with vertex set in general position using a SweepLine preprocessing of its vertices.

5.6 Conclusions

In this chapter we initiated the study of affine invariant geometric algorithms, a topic absent in the computational geometry literature. We revisited Nielson's affine invariant norm, whose unit disk represents how spread the point set is with respect to its mean. We also proposed affine invariant point sorting methods, which are necessary for other affine invariant geometric constructions. Our methods heavily rely on being able to distinguish three points. To this end, we gave two methods for distinguishing such points in Section 5.3. However, for this we had to require that the points are in certain general position. Otherwise, the point set becomes highly symmetric, which introduces a problem of indistinguishability. In particular, an interesting open question is to what extent such restriction can be removed, while still being able to distinguish rotations and reflections. Finally, we believe that finding affine invariant methods to construct other geometric objects is a promising direction for future research.

Concluding remarks

In this thesis we presented different contributions in the study of generalized Delaunay graphs. Each chapter presents a different problem in a specific generalization of the Delaunay graph. At the end of each chapter we presented different future lines of study to follow in each particular problem. We devote this last chapter to introduce a few other different open problems related to generalized Delaunay graphs.

In Chapter 4 we have shown upper bounds on the minimum k for which the k-order convex shape Delaunay graph is Hamiltonian. Restricting the triangulation in its connectivity, a long-standing open problem is whether 3-connected Delaunay triangulations are Hamiltonian [48]. By loosening slightly the cycle property, the next question is whether the Delaunay triangulation always contains a Hamiltonian path. This question was answered in the affirmative when the convex distance is the square distance (L_∞-metric) [3]. An interesting problem in this setting, is to determine for which convex shapes this is also true (in particular, circles).

A closely related problem is the one of perfect matchings. Using the fact that Delaunay triangulations are 1-tough, Dillencourt [49] showed that the Delaunay triangulation always contains a perfect matching. When focus shifted to convex shape Delaunay graphs, Ábrego et al. [3] proved that the $\square$-Delaunay graph always contains a perfect matching. However, not all convex shape Delaunay graphs have a perfect matching. At least that is the case for triangles, since any triangulation is triangle-Delaunay-realizable [24]. Biniaz et al. [21] proved that the 2-order $\triangle$-Delaunay graph (i.e., when the triangle is equilateral) always contains a perfect matching. Our results on Hamiltonicity in Chapter 4 imply bounds for the minimum k for which the k-order convex shape Delaunay graph contains a perfect matching. However, we think that it must be possible to generalize the 1-toughness result of Dillencourt for point-symmetric convex shape Delaunay graphs and thus the containment of a perfect matching. In particular, the short proof given recently by Biniaz [19] gives some hope to believe this is true.

If we consider the k-order Delaunay graphs, any of the graph-theoretic properties have been studied by Abellanas et al. [2] and Bose et. al [34]. In particular,

Worst-case possible values for the studied properties. SR, D, κ, χ, and θ_c denote the spanning ratio, diameter, connectivity, chromatic number, and constrained geometric thickness, respectively. The results hold for specific ranges of k; see the complete statements throughout the paper.

Gabriel graphs	Delaunay graphs
$\Theta(\sqrt{\frac{n}{k}}) \leqslant \max_{\lvert S\rvert=n} SR(k\text{-GG}(S)) \in O(\sqrt{n})$	$1.5931 \leqslant \max_{\lvert S\rvert=n} SR(k\text{-DG}(S)) \leqslant 1.998$
$\lceil \frac{n-1}{k+1} \rceil \leqslant \max_{\lvert S\rvert=n} D(k\text{-GG}(S)) \leqslant \frac{n-2}{k+1} + 1$	$\lfloor \frac{n}{2(k+1)} \rfloor \leqslant \max_{\lvert S\rvert=n} D(k\text{-DG}(S)) \leqslant \frac{n-2}{2(k+1)} + 1$
$\min_{\lvert S\rvert=n} \kappa(k\text{-GG}(S)) = k+1$	$\min_{\lvert S\rvert=n} \kappa(k\text{-DG}(S)) = 2(k+1)$
$3(k+1) \leqslant \max_{\lvert S\rvert=n} \chi(k\text{-GG}(S)) \leqslant 6(k+1)$	$4(k+1) \leqslant \max_{\lvert S\rvert=n} \chi(k\text{-DG}(S)) \leqslant 6(k+1)$
$3 \leqslant \max_{\lvert S\rvert=n} \theta_c(1\text{-GG}(S)) \leqslant 4$	$3 \leqslant \max_{\lvert S\rvert=n} \theta_c(1\text{-DG}(S)) \leqslant 4$
$\frac{3k}{2} \leqslant \max_{\lvert S\rvert=n} \theta_c(k\text{-GG}(S)) \leqslant 18k^2 + 3k$	$\frac{3k}{2} \leqslant \max_{\lvert S\rvert=n} \theta_c(k\text{-DG}(S)) \leqslant 18k^2 + 3k$

Figure 5.11: This table appears as Table 1 in Bose et al. [34].

Bose et. al [34] gave several bounds for the spanning ratio, diameter, chromatic number and constrained geometric thickness, in both: k-Delaunay graphs and k-Gabriel graphs. However, most of these bounds have gaps that need to be closed. See Figure 5.11. Thus, the question is which are the correct bounds? or can we get better bounds for each: lower and upper bounds? A different but related line of study concerns the bounds in [34]. Which of these bounds are preserved for k-order $\mathcal{C}$-Delaunay graphs? For instance, we believe that the bounds given for the diameter of the k-$DG(S)$ must hold for k-order $\mathcal{C}$-Delaunay graphs, since their proof does not use a particular property of a circle but properties of homothets of a circle.

Other properties of study are the ones given by Biniaz et al. [21], in which they studied the k-order $\triangle$-Delaunay graphs. More precisely, they showed bounds on the minimum k for which the k-$DG_{\triangle}(S)$ contains a $d_{\triangle}$-bottleneck of the following settings: biconnected spanning graph, Hamiltonian cycle and perfect matching. In Chapter 4 we gave bounds for the minimum k in which the k-order $\mathcal{C}$-Delaunay graph contains a $d_{\mathcal{C}}$-bottleneck Hamiltonian cycle. Yet, besides trying to improve such bounds, there is the line of what happens with $d_{\mathcal{C}}$-bottleneck biconnected spanning graphs and perfect matchings. Particularly, we believe that the bounds given by Biniaz et al. [21] on $d_{\triangle}$-bottleneck biconnected spanning graphs should also easily be generalized to any $\mathcal{C}$-distance.

In summary, generalized Delaunay graphs are of great interest, with a vast sea of open questions. Thus, we invite everyone interested in Delaunay graphs to dive into this sea in order to find some solutions and, most likely, come up with more new questions waiting to be answered.

Bibliography

[1] Yusuke Abe and Yoshio Okamoto. On algorithmic enumeration of higher-order Delaunay triangulations. In *Proceedings of the 11th Japan-Korea Joint Workshop on Algorithms and Computation (WAAC)*, pages 19–20, 2008. – Cited on pages 32, 34, and 54.

[2] Manuel Abellanas, Prosenjit Bose, Jesús García, Ferran Hurtado, Carlos M. Nicolás, and Pedro Ramos. On structural and graph theoretic properties of higher order Delaunay graphs. *International Journal of Computational Geometry and Applications*, 19(06):595–615, 2009. – Cited on pages 32, 43, 52, 53, 54, 79, 81, 102, and 125.

[3] Bernardo M. Ábrego, Esther M. Arkin, Silvia Fernández-Merchant, Ferran Hurtado, Mikio Kano, Joseph S.B. Mitchell, and Jorge Urrutia. Matching points with squares. *Discrete & Computational Geometry*, 41(1):77–95, 2009. – Cited on pages 79 and 125.

[4] Oswin Aichholzer, Franz Aurenhammer, Siu-Wing Cheng, Naoki Katoh, Günter Rote, Michael Taschwer, and Yin-Feng Xu. Triangulations intersect nicely. *Discrete & Computational Geometry*, 16(4):339–359, 1996. – Cited on page 16.

[5] Oswin Aichholzer, Franz Aurenhammer, and Ferran Hurtado. Sequences of spanning trees and a fixed tree theorem. *Computational Geometry*, 21(1):3–20, 2002. – Cited on pages 57, 58, 65, 66, and 77.

[6] Oswin Aichholzer, Franz Aurenhammer, Hannes Krasser, and Peter Brass. Pseudotriangulations from surfaces and a novel type of edge flip. *SIAM Journal on Computing*, 32(6):1621–1653, 2003. – Cited on page 57.

[7] Oswin Aichholzer, Thomas Hackl, David Orden, Pedro Ramos, Günter Rote, André Schulz, and Bettina Speckmann. Flip graphs of bounded degree triangulations. *Graphs and Combinatorics*, 29(6):1577–1593, 2013. – Cited on page 31.

[8] Oswin Aichholzer and Klaus Reinhardt. A quadratic distance bound on sliding between crossing-free spanning trees. *Computational geometry*, 37(3):155–161, 2007. – Cited on page 57.

[9] Greg Aloupis, Luis Barba, Stefan Langerman, and Diane L. Souvaine. Bichromatic compatible matchings. *Computational geometry*, 48(8):622–633, 2015. – Cited on page 57.

[10] Esther M. Arkin, Martin Held, Joseph S.B. Mitchell, and Steven S. Skiena. Hamiltonian triangulations for fast rendering. *The Visual Computer*, 12:429–444, 1996. – Cited on page 119.

[11] Elena Arseneva, Prosenjit Bose, Pilar Cano, and Rodrigo I. Silveira. Flips in higher order Delaunay triangulations. In *36th European Workshop on Computational Geometry: extended abstracts (EuroCG)*, 2020. – Cited on page 27.

[12] Elena Arseneva, Prosenjit Bose, Pilar Cano, and Rodrigo I. Silveira. Flips in higher order Delaunay triangulations. In *Proceedings of the 14th Latin American Theoretical Informatics Symposium (LATIN)*, to appear, 2020. – Cited on page 27.

[13] Tetsuo Asano, Binay Bhattacharya, Mark Keil, and Frances Yao. Clustering algorithms based on minimum and maximum spanning trees. In *Proceedings of the 4th Annual Symposium on Computational geometry(SoCG)*, pages 252–257, 1988. – Cited on page 20.

[14] Franz Aurenhammer, Rolf Klein, and Der-Tsai Lee. *Voronoi diagrams and Delaunay triangulations*. World Scientific Publishing Company, 2013. – Cited on pages 17, 21, 23, and 26.

[15] Franz Aurenhammer and Günter Paulini. On shape Delaunay tessellations. *Information Processing Letters*, 114(10):535–541, 2014. – Cited on pages 24, 26, and 118.

[16] David Avis and Komei Fukuda. Reverse search for enumeration. *Discrete Applied Mathematics*, 65(1-3):21–46, 1996. – Cited on pages 32, 53, 54, and 57.

[17] Marc Benkert, Joachim Gudmundsson, Herman Haverkort, and Alexander Wolff. Constructing interference-minimal networks. In *International Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM)*, pages 166–176. Springer, 2006. – Cited on page 22.

[18] Marshall Bern and David Eppstein. Mesh generation and optimal triangulation. *Computing in Euclidean geometry*, 1:23–90, 1992. – Cited on page 16.

[19] Ahmad Biniaz. A short proof of the toughness of Delaunay triangulations. In *Symposium on Simplicity in Algorithms*, pages 43–46. SIAM, 2020. – Cited on page 125.

[20] Ahmad Biniaz and Gholamhossein Dastghaibyfard. Slope fidelity in terrains with higher-order Delaunay triangulations. In *16th International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision*, pages 17–23. Václav Skala-UNION Agency, 2008. – Cited on page 32.

[21] Ahmad Biniaz, Anil Maheshwari, and Michiel Smid. Higher-order triangular-distance Delaunay graphs: Graph-theoretical properties. *Computational Geometry*, 48(9):646–660, 2015. – Cited on pages 79, 80, 93, 125, and 126.

[22] Ahmad Biniaz, Anil Maheshwari, and Michiel Smid. Bottleneck matchings and Hamiltonian cycles in higher-order Gabriel graphs. In *Proceedings of the 32nd European Workshop on Computational Geometry (EuroCG)*, pages 179–182, 2016. – Cited on pages 79, 80, and 93.

[23] John Adrian Bondy and Uppaluri Siva Ramachandra Murty. *Graph theory with applications*, volume 290. Citeseer, 1976. – Cited on page 18.

[24] Nicolas Bonichon, Cyril Gavoille, Nicolas Hanusse, and David Ilcinkas. Connections between theta-graphs, Delaunay triangulations, and orthogonal surfaces. In *International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 266–278. Springer, 2010. – Cited on pages 79 and 125.

[25] Houman Borouchaki and Paul Louis George. Delaunay triangulation and meshing: Application to finite elements. *Hermes, Paris*, 1998. – Cited on pages 17 and 21.

[26] Otakar Borůvka. O jistém problému minimálním. *Práce Mor. Prırodved. Spol. v Brne (Acta Societ. Scienc. Natur. Moravicae)*, 3(3):37–58, 1926. – Cited on page 20.

[27] Otakar Borůvka. Prıspevek k rešenı otázky ekonomické stavby elektrovodnıch sıtı (contribution to the solution of a problem of economical construction of electrical networks). *Elektronickỳ obzor*, 15:153–154, 1926. – Cited on page 20.

[28] Prosenjit Bose, Pilar Cano, Maria Saumell, and Rodrigo I. Silveira. Hamiltonicity for convex shape Delaunay and Gabriel graphs. In *Proceedings of the 16th Algorithms and Data Structures Symposium (WADS)*, pages 196–210. Springer, 2019. – Cited on page 28.

[29] Prosenjit Bose, Pilar Cano, Maria Saumell, and Rodrigo I. Silveira. Hamiltonicity for convex shape Delaunay and Gabriel graphs. In *35th European Workshop on Computational Geometry: extended abstracts (EuroCG)*, 2019. – Cited on page 29.

[30] Prosenjit Bose, Pilar Cano, Maria Saumell, and Rodrigo I. Silveira. Hamiltonicity for convex shape Delaunay and Gabriel graphs. *Computational Geometry*, to appear, 2020. – Cited on page 28.

[31] Prosenjit Bose, Pilar Cano, and Rodrigo I. Silveira. Sequences of spanning trees for L_∞-Delaunay triangulations. In *34th European Workshop on Computational Geometry: extended abstracts (EuroCG)*, 2018. – Cited on page 28.

[32] Prosenjit Bose, Pilar Cano, and Rodrigo I. Silveira. Affine invariant triangulations. In *31st Canadian Conference in Computational Geometry (CCCG)*, pages 250–256, 2019. – Cited on page 29.

[33] Prosenjit Bose, Paz Carmi, Sébastien Collette, and Michiel Smid. On the stretch factor of convex Delaunay graphs. *Journal of Computational Geometry*, 1(1), 2010. – Cited on pages 23, 24, and 80.

[34] Prosenjit Bose, Sébastien Collette, Ferran Hurtado, Matias Korman, Stefan Langerman, Vera Sacristán, and Maria Saumell. Some properties of k-Delaunay and k-Gabriel graphs. *Computational Geometry*, 46(2):131–139, 2013. – Cited on pages 125 and 126.

[35] Prosenjit Bose, Jean-Lou De Carufel, and André van Renssen. Constrained generalized delaunay graphs are plane spanners. *Computational Geometry*, 74:50–65, 2018. – Cited on page 24.

[36] Prosenjit Bose and Ferran Hurtado. Flips in planar graphs. *Computational Geometry*, 42(1):60–80, 2009. – Cited on pages 31 and 57.

[37] Peter A. Burrough, Rachael McDonnell, Rachael A. McDonnell, and Christopher D. Lloyd. *Principles of geographical information systems*. Oxford university press, 2015. – Cited on page 21.

[38] Owen Byer, Felix Lazebnik, and Deirdre L. Smeltzer. *Methods for Euclidean geometry*, volume 37. MAA, 2010. – Cited on page 105.

[39] Maw-Shang Chang, Chuan Yi Tang, and Richard C.T. Lee. 20-relative neighborhood graphs are Hamiltonian. *Journal of Graph Theory*, 15(5):543–557, 1991. – Cited on pages 79, 81, and 102.

[40] L. Paul Chew. Constrained Delaunay triangulations. *Algorithmica*, 4(1-4):97–108, 1989. – Cited on page 21.

[41] L. Paul Chew. There are planar graphs almost as good as the complete graph. *Journal of Computer System Sciences*, 39(2):205–219, 1989. – Cited on pages 21 and 24.

[42] Nicos Christofides. Worst-case analysis of a new heuristic for the travelling salesman problem. Technical report, Carnegie-Mellon Univ Pittsburgh Pa Management Sciences Research Group, 1976. – Cited on page 20.

[43] Mark de Berg, Marc J. van Kreveld, Otfried Schwarzkopf, and Mark Overmars. *Computational geometry algorithms and applications*. Springer, 2000. – Cited on page 18.

[44] Leila De Floriani. Surface representations based on triangular grids. *The Visual Computer*, 3(1):27–50, 1987. – Cited on pages 16 and 21.

[45] Thierry De Kok, Marc J. van Kreveld, and Maarten Löffler. Generating realistic terrains with higher-order Delaunay triangulations. *Computational Geometry*, 36(1):52–65, 2007. – Cited on page 32.

[46] Boris N. Delaunay. Sur la sphere vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk*, 7(793-800):1–2, 1934. – Cited on page 17.

[47] René Descartes. *Principia Philosophiae*. Ludovicus Elzervirius, Amsterdam, 1644. – Cited on page 16.

[48] Michael B. Dillencourt. A non-Hamiltonian, nondegenerate Delaunay triangulation. *Information Processing Letters*, 25(3):149–151, 1987. – Cited on pages 79, 80, and 125.

[49] Michael B. Dillencourt. Toughness and Delaunay triangulations. *Discrete & Computational Geometry*, 5:575–601, 1990. – Cited on pages 79, 108, and 125.

[50] David P. Dobkin, Steven J. Friedman, and Kenneth J. Supowit. Delaunay graphs are almost as good as complete graphs. *Discrete & Computational Geometry*, 5(4):399–407, 1990. – Cited on page 108.

[51] Yadolah Dodge and Daniel Commenges. *The Oxford dictionary of statistical terms*. Oxford University Press on Demand, 2006. – Cited on page 106.

[52] Adrian Dumitrescu, André Schulz, Adam Sheffer, and Csaba D. Tóth. Bounds on the maximum multiplicity of some common geometric graphs. *SIAM Journal on Discrete Mathematics*, 27(2):802–826, 2013. – Cited on page 16.

[53] R.D. Düppe and H.H. Gottschalk. Automatische interpolation von isolinien bei willkürlichen stützpunkten. *Allgemeine Vermessungsnachrichten*, 77:423–426, 1970. – Cited on page 16.

[54] Herbert Edelsbrunner, Leonidas J. Guibas, and Jorge Stolfi. Optimal point location in a monotone subdivision. *SIAM Journal on Computing*, 15(2):317–340, 1986. – Cited on page 122.

[55] Ruy Fabila-Monroy and Jorge Urrutia. Graham triangulations and triangulations with a center are Hamiltonian. *Information Processing Letters*, 93:295–299, 2005. – Cited on pages 118 and 119.

[56] Ferenc Fodor. The densest packing of 13 congruent circles in a circle. *Contributions to Algebra and Geometry*, 44(2):431–440, 2003. – Cited on page 93.

[57] Ferenc Fodor. Packing of 14 congruent circles in a circle. *Studies of the University of Zilina. Mathematical Series*, 16:25–34, 2003. – Cited on page 93.

[58] Steven Fortune. Voronoi diagrams and Delaunay triangulations. In *Computing in Euclidean geometry*, pages 225–265. World Scientific, 1995. – Cited on pages 31 and 33.

[59] Michael R. Garey, David S. Johnson, Franco P. Preparata, and Robert E. Tarjan. Triangulating a simple polygon. *Information Processing Letters*, 7(4):175–179, 1978. – Cited on page 122.

[60] Ronald L. Graham. An efficient algorithm for determining the convex hull of a finite planar set. *Information Processing Letters*, 1:132–133, 1972. – Cited on page 118.

[61] Ronald L. Graham and Pavol Hell. On the history of the minimum spanning tree problem. *Annals of the History of Computing*, 7(1):43–57, 1985. – Cited on page 20.

[62] Joachim Gudmundsson, Mikael Hammar, and Marc J. van Kreveld. Higher order Delaunay triangulations. *Computational Geometry*, 23(1):85–98, 2002. – Cited on pages 22, 43, 52, and 53.

[63] Joachim Gudmundsson, Herman J. Haverkort, and Marc J. van Kreveld. Constrained higher order Delaunay triangulations. *Computational Geometry*, 30(3):271–277, 2005. – Cited on pages 21 and 32.

[64] Sofie Haesevoets, Bart Kuijpers, and Peter Z. Revesz. Affine-invariant triangulation of spatio-temporal data with an application to image retrieval. *ISPRS International Journal of Geo-Information*, 6(4):100, 2017. – Cited on page 104.

[65] Michael E. Houle, Ferran Hurtado, Marc Noy, and Eduardo Rivera-Campo. Graphs of triangulations and perfect matchings. *Graphs and Combinatorics*, 21(3):325–331, 2005. – Cited on page 31.

[66] Ferran Hurtado, Marc Noy, and Jorge Urrutia. Flipping edges in triangulations. *Discrete & Computational Geometry*, 22(3):333–346, 1999. – Cited on page 31.

[67] Tomáš Kaiser, Maria Saumell, and Nico Van Cleemput. 10-Gabriel graphs are Hamiltonian. *Information Processing Letters*, 115(11):877–881, 2015. – Cited on pages 79, 80, 81, 93, and 102.

[68] J. Mark Keil and Carl A. Gutwin. The Delaunay triangulation closely approximates the complete Euclidean graph. In *Proceedings of the 1st Workshop on Algorithms and Data Structures (WADS)*, pages 47–56. Springer, 1989. – Cited on page 108.

[69] David G. Kirkpatrick. A note on Delaunay and optimal triangulations. *Information Processing Letters*, 10:127–128, 1980. – Cited on page 103.

[70] Joseph B. Kruskal. On the shortest spanning subtree of a graph and the traveling salesman problem. *Proceedings of the American Mathematical society*, 7(1):48–50, 1956. – Cited on page 20.

[71] Philip M. Lankford. Regionalization: theory and alternative algorithms. *Geographical Analysis*, 1(2):196–212, 1969. – Cited on page 20.

[72] Charles L. Lawson. Transforming triangulations. *Discrete Mathematics*, 3(4):365 – 372, 1972. – Cited on page 31.

[73] Charles L. Lawson. Software for C1 surface interpolation. In *Mathematical software*, pages 161–194. Elsevier, 1977. – Cited on pages 16 and 31.

[74] Der-Tsai Lee. Proximity and reachability in the plane. Technical report, Iillinois University at Urbana-Champaign coordinated Science Lab, 1978. – Cited on page 21.

[75] Der-Tsai Lee and Arthur K. Lin. Generalized Delaunay triangulation for planar graphs. *Discrete & Computational Geometry*, 1(3):201–217, 1986. – Cited on page 21.

[76] Der-Tsai Lee and Franco P. Preparata. Location of a point in a planar subdivision and its applications. *SIAM Journal on Computing*, 6(3):594–606, 1977. – Cited on page 122.

[77] Errol Lynn Lloyd. On triangulations of a set of points in the plane. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science (SFCS)*, pages 228–240. IEEE, 1977. – Cited on page 16.

[78] Anna Lubiw and Vinayak Pathak. Flip distance between two triangulations of a point set is np-complete. *Computational Geometry*, 49:17–23, 2015. – Cited on page 31.

[79] Lihong Ma. *Bisectors and Voronoi diagrams for convex distance functions.* PhD thesis, Citeseer, 2000. – Cited on page 23.

[80] Dana N. Mackenzie. Triquetras and porisms. *The College Mathematics Journal*, 23(2):118–131, 1992. – Cited on page 98.

[81] Gary H. Meisters. Polygons have ears. *The American Mathematical Monthly*, 82:648–651, 1975. – Cited on page 120.

[82] Dieter Mitsche, Maria Saumell, and Rodrigo I. Silveira. On the number of higher order Delaunay triangulations. *Theoretical Computer Science*, 412(29):3589–3597, 2011. – Cited on page 32.

[83] Wolfgang Mulzer and Günter Rote. Minimum-weight triangulation is NP-hard. *Journal of the ACM (JACM)*, 55(2):1–29, 2008. – Cited on page 16.

[84] Jaroslav Nešetřil, Eva Milková, and Helena Nešetřilová. Otakar borůvka on minimum spanning tree problem translation of both the 1926 papers, comments, history. *Discrete mathematics*, 233(1-3):3–36, 2001. – Cited on page 20.

[85] Gregory M. Nielson. Coordinate free scattered data interpolation. In *Topics in multivariate approximation*, pages 175–184. Elsevier, 1987. – Cited on pages 103 and 105.

[86] Gregory M. Nielson. A characterization of an affine invariant triangulation. In *Geometric modelling*, pages 191–210. Springer, 1993. – Cited on pages 29, 103, and 106.

[87] Atsuyuki Okabe, Barry Boots, Kokichi Sugihara, and Sung Nok Chiu. *Spatial tessellations: concepts and applications of Voronoi diagrams*, volume 501. John Wiley & Sons, 2009. – Cited on pages 17 and 21.

[88] Alexander Pilz. Flip distance between triangulations of a planar point set is apx-hard. *Computational Geometry*, 47(5):589–604, 2014. – Cited on page 31.

[89] Franco P Preparata and Michael I. Shamos. *Computational geometry: an introduction*. Springer Science & Business Media, 2012. – Cited on page 18.

[90] Robert Clay Prim. Shortest connection networks and some generalizations. *The Bell System Technical Journal*, 36(6):1389–1401, 1957. – Cited on page 20.

[91] Natalia Rodríguez and Rodrigo I. Silveira. Implementing Data-Dependent triangulations with higher order Delaunay triangulations. *ISPRS International Journal of Geo-Information*, 6(12):390, 2017. – Cited on pages 22 and 54.

[92] Maria Saumell Mendiola. *Some problems on proximity graphs*. PhD thesis, Universitat Politècnica de Catalunya, 2011. – Cited on pages 79, 80, and 96.

[93] Michael Ian Shamos. *Computational Geometry*. PhD thesis, Yale University, 1978. – Cited on page 79.

[94] Micha Sharir and Adam Sheffer. Counting triangulations of planar point sets. *The Electronic Journal of Combinatorics*, 18(1), 2011. – Cited on page 16.

[95] Jonathan Shewchuk, Tamal K. Dey, and Siu-Wing Cheng. *Delaunay mesh generation*. Chapman and Hall/CRC, 2016. – Cited on pages 16, 17, and 21.

[96] Robin Sibson. Locally equiangular triangulations. *The Computer Journal*, 21(3):243–245, 1978. – Cited on pages 16, 31, 33, 43, and 52.

[97] Rodrigo I. Silveira and Marc J. van Kreveld. Optimal higher order Delaunay triangulations of polygons. *Computational Geometry*, 42(8):803–813, 2009. – Cited on pages 22, 32, 44, and 53.

[98] Rodrigo I. Silveira and Marc J. van Kreveld. Towards a definition of higher order constrained Delaunay triangulations. *Computational Geometry*, 42(4):322–337, 2009. – Cited on page 32.

[99] Daniel D. Sleator, Robert E. Tarjan, and William P. Thurston. Rotation distance, triangulations, and hyperbolic geometry. *Journal of the American Mathematical Society*, 1(3):647–681, 1988. – Cited on page 31.

[100] Tung-Hsin Su and Ruei-Chuan Chang. The k-Gabriel graphs and their applications. In *International Symposium on Algorithms*, pages 66–75. Springer, 1990. – Cited on page 22.

[101] Alfred H. Thiessen. Precipitation averages for large areas. *Monthly weather review*, 39(7):1082–1089, 1911. – Cited on page 16.

[102] Godfried T. Toussaint. The relative neighbourhood graph of a finite planar set. *Pattern recognition*, 12(4):261–268, 1980. – Cited on pages 20, 21, and 104.

[103] Godfried T. Toussaint and David Avis. On a convex hull algorithm for polygons and its application to triangulation problems. *Pattern Recognition*, 15(1):23–29, 1982. – Cited on page 118.

[104] Marc J. van Kreveld, Maarten Löffler, and Rodrigo I. Silveira. Optimization for first order Delaunay triangulations. *Computational Geometry*, 43(4):377–394, 2010. – Cited on page 32.

[105] Georges Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. deuxième mémoire. recherches sur les parallélloèdres primitifs. *Journal für die reine und angewandte Mathematik*, 134:198–287, 1908. – Cited on page 17.

[106] Eugene Wigner and Frederick Seitz. On the constitution of metallic sodium. *Physical Review*, 43(10):804, 1933. – Cited on pages 16 and 20.

[107] Johnson William and Robert Mehl. Reaction kinetics in processes of nucleation and growth. *Transactions of the American Institute for Mining and Metallurgical Engineering*, 135:416–442, 1939. – Cited on page 16.

[108] Ge Xia. The stretch factor of the Delaunay triangulation is less than 1.998. *SIAM Journal on Computing*, 42(4):1620–1659, 2013. – Cited on pages 17, 60, and 108.

"So long, and thanks for all the fish."