5-2015

Geometric Separation and Packing Problems

Ivo Vigan
Graduate Center, City University of New York

GEOMETRIC SEPARATION AND PACKING PROBLEMS

BY

IVO VIGAN

A dissertation submitted to the Graduate Faculty in Computer Science in partial fulfillment of the requirements for the degree of Doctor of Philosophy, The City University of New York

2015

This manuscript has been read and accepted for the
Graduate Faculty in Computer Science in satisfaction for the
dissertation requirement for the degree of Doctor of Philosophy.

Peter Braß

Date — Chair of Examining Committee

Robert Haralick

Date — Executive Officer

Amotz Bar-Noy

Jianting Zhang

Jonathan Lenchner

Supervision Committee

THE CITY UNIVERSITY OF NEW YORK

Abstract

GEOMETRIC SEPARATION AND PACKING PROBLEMS

by

Ivo Vigan

Advisor: Professor Peter Braß

The first part of this thesis investigates combinatorial and algorithmic aspects of geometric separation problems in the plane. In such a setting one is given a set of points and a set of separators such as lines, line segments or disks. The goal is to select a small subset of those separators such that every path between any two points is intersected by at least one separator. We first look at several problems which arise when one is given a set of points and a set of unit disks embedded in the plane and the goal is to separate the points using a small subset of the given disks. Next, we focus on a separation problem involving only one region: Given a region in the plane, bounded by a piecewise linear closed curve, such as a fence, place few guards inside the fenced region such that wherever an intruder cuts through the fence, the closest guard is at most a distance one away. Restricting the separating objects to be lines, we investigate combinatorial aspects which arise when we use them to pairwise separate a set of points in the plane; hereafter we generalize the notion of separability to arbitrary sets and present several enumeration results. Lastly, we investigate a packing problem with a non-convex shape in $\mathbb{R}^3$. We show that $\mathbb{R}^3$ can be packed at a density of 0.222 with tori of major radius one and minor radius going to zero. Furthermore, we show that the same torus arrangement yields the asymptotically optimal number of pairwise linked tori.

Acknowledgments

I would first like to thank my advisor Peter Braß for introducing me to so many different problems and elegant techniques, for guiding me to interesting research directions, for generously supporting me financially, and for arranging countless research and workshop visits. I would like to thank Amotz Bar-Noy for providing many interesting perspectives, both inside and outside of mathematics, and for letting me work on his project while supporting me financially. I would also like to thank Robert Haralick for letting me be part of his data visualization project during my first year of studies. Furthermore, I would like to thank all my committee members for their time.

The following people hosted me for research visits at their universities and research institutes. I am deeply grateful to all of them for spending their time working on problems with me: Christian Knauer at University of Bayreuth, Mark de Berg at University of Eindhoven, Stephen Kobourov at University of Arizona, Hee-Kap Ahn at POSTECH University, Csaba D. Tóth at University of Calgary, Rolf Klein at University of Bonn, Heiko Harborth at Technical University of Braunschweig, Alexander Wolff at University of Wuerzburg, Hyeon-Suk Na at Soongsil University and Marco Laumanns at IBM Research Zürich. I thank Takahisa Toda for sparking my interest in the combinatorics of separation problems and for being a very pleasant and competent co-author.

I would like to thank Ning Xu, Maleni Romero, Zuska Kepplova, Cagil Tasdemir, Konstantios Pouliasis and Deniz Sariöz for many good discussions and for all the fun we had together. Very special thanks goes to George Rabanca; he is not only a roommate, startup co-founder, co-author and a person always interested in hearing about a problem, but also very good friend.

Above all, I would like to thank my parents for their love and continuous support, far beyond what I can express within these lines.

to my parents

Contents

List of Figures

Chapter 1

Introduction and Motivation

In the first part of this thesis we focus on geometric separation problems in the plane. In such a situation one is given a set of points in the plane and a set of separating objects such as lines, line segments or disks and the goal is to select a small subset of those separating objects such that every path between any two points is intersected by at least one separating object.

The motivation for studying these types of problems comes from sensor networks. Historically, sensor networks tried to achieve *full coverage* of a region, where each point of the region has to be within sensing radius of at least one sensor. In [67], an alternative type of coverage, called *Barrier Coverage*, was introduced, where the goal is to place few sensors or guards to detect any intruder into a given region. Barrier coverage received considerable attention (see for example [14],[24],[25],[66],[68],[83],[87]) and the separation problems studied here should be seen as extending this line of research.

The main results of this thesis are split into four chapters.

Chapter 2 considers the problem of given a set of points in the plane and a set of unit disks, separating the points, the goal is to select a minimum size subset of the disks such that any path between any pair of these points is intercepted by some disk. We present a $(9+\epsilon)$-approximation algorithm for this problem which results in a 6-approximation if none of the disks contains a point. Next, we show that this problem is NP-complete, even if no point is contained in any disk. Using a similar reduction, we further show that removing a minimum subset of a given collection of unit disks, such that the such that the plane minus the arrangement of the remaining disks consists of a single connected region is also NP-complete. Lastly, we show that the Multiterminal Cut Problem remains NP-complete on unit disk graphs.

In Chapter 3, we are given a region bounded by a piecewise linear closed curve, such as a fence and we want to place few guards inside the fenced region, such that wherever an intruder cuts through the fence, the closest guard is at most a distance one away. Another way of looking at this problem is from an Art Gallery perspective (see for example [73]), where the polygon represents a gallery and, regardless where on the wall a painting is hanged, the closest guard is at most a distance one away. More formally, we define a geodesic unit disk as all the points in a polygon whose shortest path distance to the disk center is at most one. We present an $O(n \log^2 n + k)$ time 2-approximation algorithm finding a collection of geodesic unit disks covering the boundary of a simple polygon on n vertices, with k denoting the number of disks found by the algorithm. The algorithm then returns a set containing the disk centers. We end this chapter by presenting a simple linear time algorithm which achieves an asymptotically optimal approximation ration when the polygon perimeter is much larger than n.

In Chapter 4 we first enumerate the number of ways a set of points in convex position in

the plane can be pairwise separated by a set of lines and we later generalize this notion and the enumeration results to arbitrary sets. For a set S of n points in the plane, a linear bipartition of S is a set $\{U, S \setminus U\}$ consisting of two disjoint nonempty subsets of S which respectively are fully contained in the two open half-planes bounded by some line. A set $\mathcal{P}$ of linear bipartitions is called a linear separating family for S if every two elements in S are separated by some bipartition in $\mathcal{P}$. In addition, $\mathcal{P}$ is called minimal, if no proper subset of $\mathcal{P}$ is a separating family for S. For the case where the points in S are in convex position, we present a bijection between the set of all minimal linear separating families and a restricted class of edge covers. Using this bijection we enumerate minimal separating families of minimum size and maximum size. We then generalize this result to arbitrary sets, by defining families of bipartitions, as set partitions consisting of at most two components. A family of bipartitions is a separating family for a set if every two elements in the set are separated by some bipartition. We enumerate separating families of arbitrary size. We furthermore enumerate inclusion-wise minimal separating families of minimum and maximum sizes.

In Chapter 5 we change our focus to a packing problem in $\mathbb{R}^3$. Such geometric packing problems received huge attention over many decades (see [9],[18],[39],[40],[41],[42],[49],[50], [74],[82],[85],[101] for books on packings). Still, the sphere is the only body which does not tile $\mathbb{R}^3$ and for which we know the exact packing density [53]. For other bodies such as Platonic solids [26],[93] and ellipsoids [16],[35] dense packing constructions are known, but no proof of optimality exists and a vast amount of related questions remain open (see the books [19] and [31]). On the other hand, there is only a very limited amount of literature studying packings involving non-convex objects, such as the work of Jiao et al. [57],[94],[95]. We thus extend this line of research by considering packings with

the possibly simplest non-convex shape, the torus and show that $\mathbb{R}^3$ can be packed at a density of 0.222... with tori whose minor radius goes to zero. Furthermore, we show that the same torus arrangement yields the asymptotically optimal number of pairwise linked tori.

Chapter 2

On Isolating Points Using Unit Disks

2.1 Introduction

Wireless sensors are being extensively used in applications to provide barriers as a defense mechanism against intruders at important buildings, estates, national borders etc. Monitoring the area of interest by this type of coverage is called *barrier* coverage [67]. Such sensors are also being used to detect and track moving objects such as animals in national parks, enemies in a battlefield, forest fires, crop diseases etc. In such applications it might be prohibitively expensive to attain full coverage, where each point of a given region is within sensing distance of at least one sensor. It suffices to ensure that the object under consideration cannot travel too far before it is detected. Such coverage is called *trap* coverage [10, 84]. Inspired by such applications, we consider the problem of isolating a set of points by a minimum-size subset of a given set of unit radius disks.

A unit disk crudely models the region sensed by a sensor, and the algorithm presented in the next section readily generalizes to disks of arbitrary, different radii.

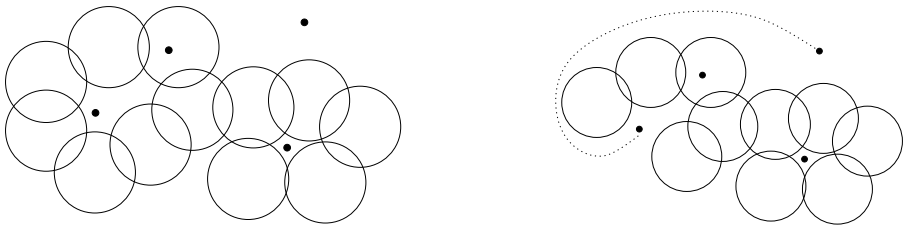

(a) (b)

FIGURE 2.1: (a) A set of disks separating four points, since every path connecting any two points intersects a disk. (b) A set of disks that does not separate the points.

Definition 2.1. A set $\mathcal{D}$ of disks embedded in the plane, referred to as an *arrangement*, *separates* a set P of points, if for any two points $p, q \in P$, every path between p and q intersects at least one disk in $\mathcal{D}$, as shown in Figure 2.1.

Problem 2.1.1 (Point Isolation)**.** *Given a set P of k points and a set $\mathcal{D}$ of n unit disks in the plane, separating P, find a minimum cardinality subset of $\mathcal{D}$ that still separates P.*

For several variants of the geometric set cover problem, approximation algorithms have been designed [7, 28, 69] that improve upon the best guarantees for the combinatorial set cover problem. For the problem of covering points by the smallest subset of a given set of unit disks, there exist a $(9+\epsilon)$-approximation algorithm [1], with $0 < \epsilon \leq 6$, and a PTAS [69]. The PTAS can even be used for disks of arbitrary different radii. Our problem can be viewed as a set cover problem where the elements that need to be covered are not points, but paths. However, known results only imply a trivial $O(n)$-approximation when viewed through this set cover lens.

Contribution and Organization. In Section 2.2, we present a polynomial time algorithm that guarantees a $(9+\epsilon)$-approximation for the Point Isolation problem. It relies on the two-point separation algorithm of [20], which solves the Point Isolation problem optimally when $k = 2$. Our algorithm recursively applies this two-point separation algorithm to find the smallest subset B of $\mathcal{D}$ that separates some pair of points in P. Observe that the arrangement of the disks in B induces exactly two faces in the plane (one bounded and one unbounded) and it thus partitions P into P_1 and P_2, such that all points in P_1 are separated from the points in P_2. The algorithm then recursively finds a separator for P_1 and for P_2, and returns the union of these separators and B. In Section 2.3, we prove that the Point Isolation problem is NP-complete, even if no disk contains any points. We believe that our hardness construction can be used to show hardness for a variety of unit disk problems and we show two of them: in Section 2.4, we show that the NP-complete Multiterminal Cut problem [32] (defined below) remains NP-complete on unit disk graphs. Lastly, in Section 2.5, we show that removing a minimum subset of a given collection of unit disks, such that the plane minus the arrangement of the remaining disks consists of a single connected region is also NP-complete.

Problem 2.1.2 (Multiterminal Cut [32])**.** *Given a graph $G = (V, E)$ and set $S \subseteq V$ of k terminals, find a minimum cardinality set $E' \subseteq E$ of edges such that in $G' = (V, E \setminus E')$ there is no path between any two terminals in S.*

Related Work. Sankararaman et al. [84] investigate a notion of coverage which they call *weak coverage*. Given a region $\mathcal{R}$ of interest (which they take to be a square in the plane) and a set $\mathcal{D}$ of unit disks (sensors), the region is said to be k-weakly covered if each connected component of $\mathcal{R} - \bigcup_{D\in\mathcal{D}} D$ has diameter at most k. They consider the situation when a given set $\mathcal{D}$ of unit disks *completely* covers $\mathcal{R}$, and address the problem

of partitioning $\mathcal{D}$ into as many subsets as possible so that $\mathcal{R}$ is k-weakly covered by every subset.

Given two regions of an arrangement of sensors, the *resilience* [67] with respect to these two regions is the minimum number of sensors that need to be deactivated so that there is a path between the two regions, not intersecting any sensors. In [13], a 3-approximation algorithm for computing the resilience for unit disk sensors is presented; the computational complexity of this problem remains open. In [65], it is shown that computing the resilience for certain types of fat sensors, such as axis aligned rectangles is NP-hard. In [4], it is shown that computing the resilience of a set of line segment sensors is NP-hard. This was extended in [97] to hold even for unit length line segments. The reductions of both [65] of [97] have some resemblance to our NP-hardness reduction, since they both reduce optimization problems on graphs to the resilience problem in a sensor network modeling the graph instances. In the context of separation, [4] addresses the two-point separation problem for a set of line segments. They show that the problem in fact admits a polynomial-time exact algorithm. This work was later extended in [20] to include an exact $O(n^3)$ time algorithm for solving the two-point separation problem on unit disks. Furthermore, they present a hardness result similar to ours, which holds for unit circles but cannot be applied to unit disks.

2.2 Approximation Algorithm

We will refer to the standard notions of vertices, arcs, and faces in arrangements of disks [2]. In particular, for a set $\mathcal{D}$ of n disks, we are interested in the faces in the complement of the union of the disks in $\mathcal{D}$. These are the connected regions in the arrangement of $\mathbb{R}^2 \setminus \bigcup_{D \in \mathcal{D}} D$, which, by slightly abusing notation, we will write as $\mathbb{R}^2 \setminus \bigcup \mathcal{D}$.

Definition 2.2. For any disks in an arrangement of disks $\mathcal{D}$, we refer to the maximal connected subset of its disk boundary which is not contained in any other disk as a *boundary arc*. We denote the collection of all boundary arcs of $\mathcal{D}$ by $\mathcal{B}(\mathcal{D})$.

Theorem 2.3 (Theorem 3.1 in [60], (see also [75]))**.** *If $|\mathcal{D}| \geq 3$, it holds that $|\mathcal{B}(\mathcal{D})| \leq 6|\mathcal{D}| - 12$.*

Observation 2.2.1. *Let $\mathcal{D}$ be a set of at least 3 disks in the plane, and Q a set of points so that (a) no point from Q is contained in any disk from $\mathcal{D}$, and (b) no face in the complement of the union of the disks in $\mathcal{D}$ contains more than one point of Q. It then holds that $|Q| \leq 2|\mathcal{D}| - 4$.*

Proof. Since each connected region in $\mathbb{R}^2 \setminus \bigcup \mathcal{D}$ is bounded by at least three boundary arcs and no boundary arc appears in two regions, it follows from Theorem 2.3 that $|Q| \leq 2|\mathcal{D}| - 4$. □

Covering vs. Separating. The input to our problem is a set $\mathcal{D}$ of n unit disks, and a set P of k points such that $\mathcal{D}$ separates P. Let $P_c \subseteq P$ denote those points contained in some disk of $\mathcal{D}$ and let P_s denote the remaining points. Note that it follows from our definition of separation that any point contained in a disk is separated from all other points. Computing a minimum set cover of $\mathcal{D}$ for P_c is thus equivalent to solving the Point Isolation problem for the points in P_c. In order to obtain an $(\alpha+\beta)$-approximation for the Point Isolation problem, we first compute an α-approximation for the smallest subset of $\mathcal{D}$ that covers P_c. We then compute a β-approximation for the smallest subset of $\mathcal{D}$ that separates P_s. We claim that the combination of the two solutions is an $(\alpha+\beta)$-approximation to the Point Isolation problem. To see this, let $OPT \subseteq \mathcal{D}$ denote an optimal subset that separates P. Suppose that OPT covers k_1 of the points in P_c

and let $k_2 = |P_c| - k_1$. By Observation 2.2.1, it holds that $k_2 \leq 2|OPT|$. Picking one disk to cover each of the k_2 points of P_c not covered by OPT, we see that there exists a cover for P_c of size at most $|OPT| + k_2 \leq 3|OPT|$. Thus, an α-approximate solution to this set cover problem has size at most $3\alpha|OPT|$. Since OPT also separates P_s, a β-approximation for P_s uses at most $\beta|OPT|$ disks. Since the α-approximate set cover algorithm separates each point in P_c from all the points in P and the β-approximate algorithm separates each point in P_s from all the points in P, the two algorithm combined provide a solution to the Point Isolation problem that has size $(3\alpha + \beta)|OPT|$.

Using the PTAS of [69] for the points in P_c results in $\alpha = 1 + \epsilon$, with $\epsilon > 0$; its running time is $O(n^{2(\frac{8\sqrt{2}}{\epsilon})^2+1}|P_c|)$, for $0 < \epsilon \leq 2$, (see [1]). In the rest of this section, we assume that no point of P is contained in any disk of $\mathcal{D}$ and we present a 6-approximation algorithm for the Point Isolation problem in this setting, i.e., we show that $\beta = 6$. Thus combining the two algorithms yields a $(9 + \epsilon)$-approximation algorithm.

Algorithm For a set Q of points separated by a set $\mathcal{D}$ of unit disks, such that no point of Q is contained in any disk of $\mathcal{D}$, we will show that the following algorithm yields a 6-approximation for the Point Isolation problem in this setting.

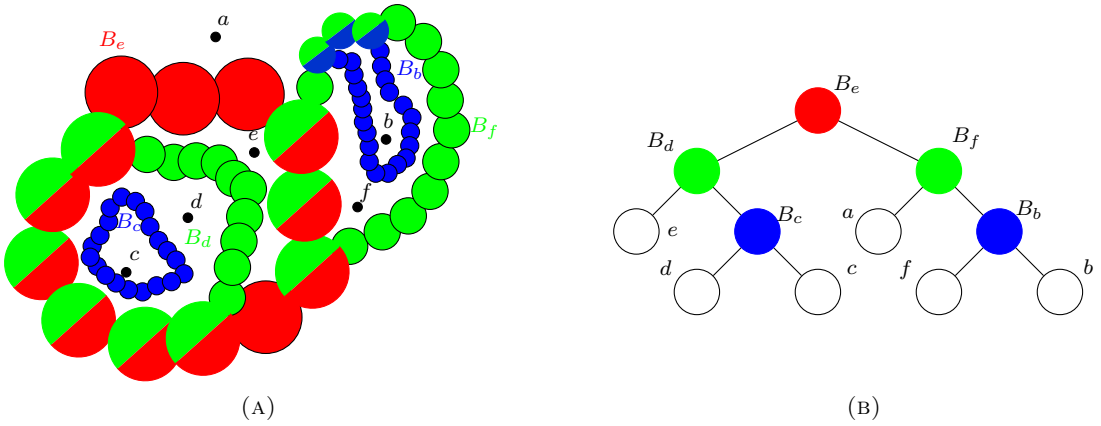

(A) (B)

FIGURE 2.2: Illustration of the charging scheme. (a) A separating arrangement computed by **recSep**, where the colors of a disk encode the separators it is contained in. Disks of two colors appear in two separators (For example, the two-colored disks in the lower left appear in both B_e and B_d). (b) The execution tree of **recSep** and the corresponding charging of the separators to the leaves.

recSep(Q, $\mathcal{D}$)

1. If $|Q| \leq 1$, return $\emptyset$.
2. For every pair of points $s, t \in Q$, invoke the algorithm of [20] to find a minimum cardinality subset $B_{s,t} \subseteq \mathcal{D}$ such that $B_{s,t}$ separates s and t.
3. Let B denote a minimum size separator $B_{s,t}$ over all pairs $s, t \in Q$.
4. Let Q_1 and Q_2 be the partition of Q into two subsets such that each subset corresponds to points in the same face induced by the arrangement of disks in B.

   return $B \cup$ **recSep**(Q_1, $\mathcal{D}$) $\cup$ **recSep**(Q_2, $\mathcal{D}$)

Since in any recursive call, both Q_1 and Q_2 contain fewer points than Q, the **recSep**(P, $\mathcal{D}$) algorithm indeed yields a separator for P. Computing a single separator $B_{s,t}$ takes time $O(n^3)$ (see [20]), and thus computing B takes time $O(k^2n^3)$. Since there are at most k separation steps, the total running time of **recSep**(P, $\mathcal{D}$) is $O(k^3n^3)$.

In each recursive step, Q is partitioned into two sets Q_1 and Q_2, thus the execution-tree

of **recSep** is a binary tree where the leaves correspond to the points in Q and the internal nodes correspond to the computed separators. It is easy to see that a tree, where each internal node has at least 2 children, has more leaves than internal nodes. After the execution of the algorithm, we charge each separator B_p to a point $p \in P$ appearing as a leave in the subtree rooted at B_p, such that each point in P gets charged at most once (see Figure 2.2). In the optimal separating arrangement OPT, we let F_p be the disks contributing to the boundary of the face containing only $p \in P$. For any $p \in P$, it holds that $|B_p| \leq |F_p|$, since otherwise the smallest separator in the corresponding recursive call would be F_p not B_p, contradicting optimality of B_p. Letting $\mathcal{B}(OPT)$ denote the set of boundary arcs of the disks from OPT, and $P_B \subseteq P$ be the set of points to which the separators were charged, we bound the number of disks returned by **recSep** using the following inequalities

$$|\textbf{recSep}(P, \mathcal{D})| \leq \sum_{p \in P_B} |B_p| \leq \sum_{p \in P} |F_p| \leq |\mathcal{B}(OPT)| \leq 6|OPT|.$$

The first inequality holds because each disk used by the algorithm appears in at least one separator, the second inequality was argued above. To see the third inequality, observe that if a disk of OPT appears in $F_{p_1}, \ldots, F_{p_i}$, then its boundary appears in at least i different regions induced by the disk arrangement of OPT and thus by definition, this disk contributes at least i boundary arcs to $\mathcal{B}(OPT)$. The last inequality holds, by Theorem 2.3.

This proves the following theorem.

Theorem 2.4. *The Point Isolation problem can be approximated within a factor of* $9+\epsilon$ *in general and within a factor of* 6 *if no point is contained in a disk of* $\mathcal{D}$.

2.3 NP-Completeness of the Point Isolation Problem

In order to show NP-completeness of the Point Isolation problem, we are going to reduce the Planar Subdivision problem (defined below) to it. This problem will be shown to be NP-complete in Proposition 2.6, by a reduction from the Planar Multiterminal Cut problem which is NP-hard according to Theorem 2.5.

Problem 2.3.1 (Planar Multiterminal Cut [32])**.** *Given a simple planar graph $G = (V, E)$ and a set $S \subseteq V$ of k terminals, find a minimum cardinality set $E' \subseteq E$ such that in $G' = (V, E \setminus E')$, there is no path between any two nodes in S.*

Theorem 2.5 ([32])**.** *If k is not fixed, the Planar Multiterminal Cut problem is NP-hard.*

Problem 2.3.2 (Planar Subdivision)**.** *Given a simple planar graph $G = (V, E)$ embedded in the plane and a set S of k faces of G, find a minimum cardinality set $E' \subseteq E$ such that in the embedding of $G' = (V, E')$, any path between any two points in two faces of S gets intersected by at least one embedded edge of E'.*

Proposition 2.6. *The Planar Subdivision problem is NP-complete, even on connected graphs.*

Proof. Given an instance $I_1 = (G_1, S_1)$ of the Planar Multiterminal Cut problem, with $G_1 = (V_1, E_1)$, we embed G_1 in the plane (using for example the linear time algorithm of [86]) and we build an instance $I_2 = (G_2, S_2)$ of the Planar Subdivision problem with $G_2 = (V_2, E_2)$ as follows: We let $\overline{G}_2 = (\overline{V}_2, \overline{E}_2)$ be the geometric dual multigraph of the embedded graph G_1. Then we create a simple graph G_2 from $\overline{G}_2$ by subdividing each edge $\{u, v\} \in \overline{E}_2$ into $\{u, x\}$ and $\{x, v\}$ by adding a new vertex x to V_2. We embed G_2 in the plane and let S_2 be the set of faces of G_2 whose dual vertices are in S_1. Observe

that an optimal solution $OPT \subseteq E_2$ for I_2 uses an edge $\{u, x\}$ if and only if it uses edge $\{x, v\}$, with $\{u, v\} \in \overline{E}_2$, since only using one of the two subdivision edges does not change the partition of the plane, with respect to the embedding of G_2. Let $OPT_\cup \subseteq \overline{E}_2$ denote the edges of $\overline{G}_2$ obtained from OPT by merging any subdivision edges $\{u, x\}$, $\{x, v\} \in OPT$ into $\{u, v\} \in OPT_\cup$, thus $|OPT| = 2|OPT_\cup|$. We denote by $OPT^*_\cup \subseteq E_1$ the duals of the edges in $OPT_\cup$. We claim that OPT is an optimal solution for I_2 if and only if $OPT^*_\cup$ is an optimal solution for I_1. To see this, let E'_1 be an arbitrary subset of E_1, let $\overline{E}'_2 \subseteq \overline{E}_2$ be the dual edges of E'_1 and let $E'_2 \subseteq E_2$ be the subdivision edges in G_2 corresponding to $\overline{E}'_2$. Two vertices $u, v \in S_1$ are connected by a path $(u, v_1, \ldots, v_l, v)$ in $G'_1 = (V_1, E_1 \setminus E'_1)$ if and only if there is a sequence $u^*, v^*_1, \ldots, v^*_l, v^*$ of adjacent faces in G_2, which, in $G'_2 = (V_2, E'_2)$, are merged into one face. Therefore, I_2 has a solution of size $2M$ if and only if I_1 has a solution of size M; the factor 2 stemming from subdividing each edge of $\overline{G}_2$. Since the Planar Multiterminal Cut problem is NP-complete on connected graphs and the dual of a connected graph is connected, the Planar Subdivision problem is NP-complete, even on connected graphs. □

Corollary 2.7. *Solving the Point Isolation problem for line segments instead of disks is* NP*-complete.*

In [86], a linear time algorithm is presented which constructs a drawing of a planar graph on n vertices, crossing free using straight line segments and having its vertices lie on an $n \times n$ grid[1]; we call such an embedding a *grid embedding.*

Lemma 2.8. *For any jordan arc embedding of a Planar Subdivision instance (G, S), with $G = (V, E)$ being a connected graph on n vertices, there exists a grid embedding of*

[1] In [86] it is actually shown that this is even possible on an $(n-2) \times (n-2)$ grid, but for our purposes an $n \times n$ grid suffices

G, such that every solution in the original embedding is a solution in the grid embedding and vice versa.

Proof. We say that two embeddings of a graph in the plane are *equivalent*, if one embedding can be continuously transformed into the other. It is well known that every maximal plane graph on at least four vertices is three-connected (Corollary 4.4.7 of [34]) and that every three-connected graph has a unique embedding according to Whitney's Theorem (modulo the choice of the outer face). Given a jordan arc embedding Σ of G, we make G maximally plane by adding $3|V| - 6 - |E|$ edges to G, obtaining a new graph $\overline{G}$ having a corresponding embedding $\overline{\Sigma}$. Embedding $\overline{G}$ on an $n \times n$ grid and removing the $3|V| - 6 - |E|$ additional edges thus results in a grid embedding of G which is equivalent to Σ. Given two equivalent embeddings Σ_1 and Σ_2 of G, we let Σ_1' and Σ_2' be the embeddings of a subgraph G' of G, which are obtained from Σ_1 and Σ_2 by deleting the corresponding edges and vertices of G not contained in G'. Since Σ_1' and Σ_2' are again two equivalent embeddings of G', it holds that every solution for a Planar Subdivision instance in the original embedding is a solution in the grid embedding and vice versa.

□

We now present two lemmas which will be useful for arguing that in a grid embedding, each edge can be replaced by a path of disks such that no disks of different paths intersect.

Lemma 2.9. *In an $n \times n$ grid, the minimum distance between any line l through two grid points and any grid point not on l is $(2n^2 - 2n + 1)^{-\frac{1}{2}}$.*

Proof. W.l.o.g., we fix one point on l to $(0,0)$. Denoting the second point on l by (a,b), we get a line equation of $bx - ay = 0$. Thus, the distance from a point (c,d) to l is $\frac{|bc-ad|}{\sqrt{b^2+a^2}}$. Furthermore, we can assume that $gcd(a,b) = 1$ since otherwise we can divide both coordinates by $gcd(a,b)$. Thus, setting $a = n$ and $b = n-1$ maximizes a^2+b^2, given $gcd(a,b) = 1$ and the minimum non-zero distance[2] is thus at least $(2n^2-2n+1)^{-\frac{1}{2}}$. □

Lemma 2.10. *In an $n \times n$ grid, for any grid point p, the minimum angle between any two distinct lines, each going through p and at least one other grid point respectively, is larger than* $2 \arctan 1/(6n^2)$.

Proof. Let g and h denote two lines through p and (a,b) and (c,d) respectively, having minimum angle, and let the slope of g be larger than the slope of h. First, observe that the slopes of g and h have the same sign, since otherwise their angle is not minimal. Therefore, it is easy to see that the minimum angle between g and h is obtained when $p = (0,0)$. Due to symmetry we can further restrict g and h to be contained in the lower right triangle portion $\{(i,j) \mid 0 \leq j \leq i \leq n\}$ of the grid. Since $b/a > d/c$, due to monotonicity of arctan, it holds that $\arctan b/a - \arctan d/c = \arctan \frac{bc-ad}{1-(bd)/(ac)} \geq \arctan \frac{b/a-d/c}{2} \geq \arctan \frac{1}{2n^2}$. The last inequality holds since all coordinates are integers. Thus, $bc - ad \geq 1$ and therefore $b/a - d/c = \frac{bc-ad}{ac} \geq 1/n^2$. The Lemma then follows from the fact that $\arctan(x) > 2\arctan(x/3)$ holds for all $0 < x < \sqrt{3}$. □

In the remainder of this section, we are going to prove the following theorem.

Theorem 2.11. *The Point Isolation problem is* NP*-complete, even if no point is contained in a disk.*

[2] Observing that $|bc - ad| \geq 1$, yields that the minimum value is achieved at point $(c,d) = (1,1)$.

In order to prove Theorem 2.11, we reduce an instance $I_2 = (G_2, S_2)$ of the Planar Subdivision problem, with G_2 being a connected embedded graph in polynomial time to an instance $I_1 = (\mathcal{D}, S_1)$ of the Point Isolation problem; note that solving I_2 is NP-hard according to Proposition 2.6. We do this by first transforming the embedding of G_2 to an equivalent straight line embedding on an $n \times n$ integer grid as argued in Lemma 2.8.

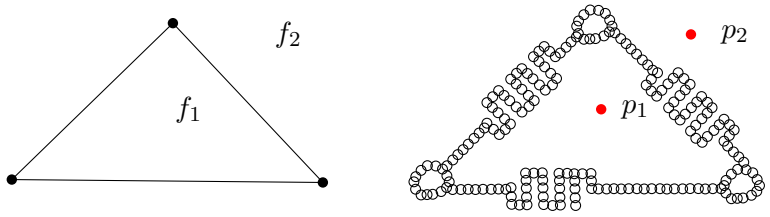

FIGURE 2.3: (a) An instance $I_2 = (G_2, S_2)$ of the Planar Subdivision problem, with $S_2 = \{f_1, f_2\}$. (b) The corresponding instance $I_1 = (\mathcal{D}, S_1)$ of the Point Isolation problem using the Vertex- and Edge gadgets defined in this section, with $S_1 = \{p_1, p_2\}$.

We then replace each edge in the embedding by an edge gadget of Definition 2.12. Such an edge gadget is depicted in Figure 2.4(a) and consists of a path of disks constructed in such a way that every edge gadget contains the same amount of disks, regardless of the length of the corresponding embedded edge. Furthermore, the dimensions of each edge gadget is carefully chosen so that no two disks of different edge gadgets intersect. Having replaced each edge by an edge gadget, we replace each vertex by a vertex gadget of Definition 2.13, shown in Figure 2.4(b). A vertex gadget for a vertex v consists of a cycle of disks which are circularly arranged around v. An edge gadget for an edge $\{u, v\}$ will have non-empty intersection with the vertex gadgets for both u and v, but does not intersect any other vertex or edge gadgets. Furthermore, all vertex gadgets are pairwise disjoint. If we denote the collection of all disks contained in the vertex- and edge gadgets by $\mathcal{D}$, then each face in the embedding of G_2 has a corresponding connected region in the arrangement $\mathbb{R}^2 \setminus \bigcup \mathcal{D}$. We put a point s_1 into the region in $\mathbb{R}^2 \setminus \bigcup \mathcal{D}$ corresponding to the face $s_2 \in S_2$ in the embedding of G_2 and add s_1 to S_1; we obtain an instance

$I_1 = (\mathcal{D}, S_1)$ of the Point Isolation problem. This whole reduction process is illustrated in Figure 2.3. The main task of the reduction is to choose the radius of the disks and the dimension of the gadgets such that every edge gadget consists of the same amount of disks and all edge gadgets are disjoint. Thus, on I_1, removing any disk D from an edge gadget merges the two adjacent regions in $\mathbb{R}^2 \setminus (\bigcup \mathcal{D} \setminus \{D\})$. Note that this is not true for vertex gadgets, i.e., removing disks from vertex gadgets does not necessarily merge any regions containing points. In order to infer the solution size for I_2 from the solution size of I_1, i.e., retrieve the number of edges removed from G_2 from the number of disks removed in the solution for I_1, we have to choose the number of disks contained in a single edge gadget to be larger than the number of disks contained in all the vertex gadgets together. We will make this argument formal in Lemma 2.14.

Definition 2.12. An edge gadget (see Figure 2.4(a)) for an embedded edge $e = \{u, v\}$ of length $2s + 2a + b$ is a path of disks which can be thought of as being placed in an elongated octagon of height h and length $2a + b$. Every edge gadget consists of a path of C_E many radius r disks which are arranged as a straight path at the beginning and the end of the gadget, which we call *hallways*, and as an up-down path in the middle part which we call a *cabin*. While C_E, a, h and s are constant for all edge gadgets, b may vary from $1 - (2s + 2a)$ to $\sqrt{2}n - (2s + 2a)$, depending on the length of the embedded edge e.

The constants a, h and s will be fixed later in this section in such a way that no two disks in two different edge gadgets intersect and that at least one disk of an edge gadget intersects the incident vertex gadgets respectively. We define $C_E = \lceil \frac{\sqrt{2}n - 2s}{2r} \rceil$, since this amounts to the number of disks of radius r needed to represent the longest edge in an $n \times n$ grid embedding as a straight line path of touching disks. Note that we take r as the unit measure for the disks. Since each edge gadget needs to contain an equal amount

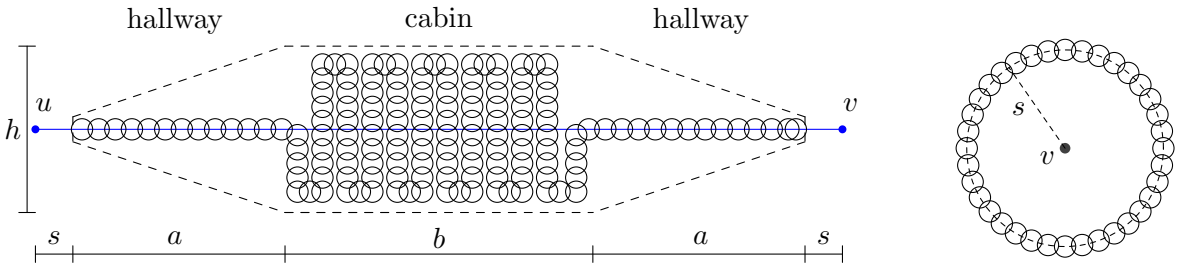

(A) (B)

FIGURE 2.4: (a) An illustration of an edge gadget replacing the edge $\{u, v\}$, (b) an illustration of a vertex gadget.

of disks, we have to place $C_E - a/r$ disks into the cabin of any edge gadget. Arranging the disks in cabins (for edges of length $< \sqrt{2}n$) as an up-down path as described in Definition 2.12 (and shown in Figure 2.4(a)) allows us to put a path consisting of up to $\left\lfloor \frac{h}{2r} \right\rfloor \cdot \left\lfloor \frac{b}{2r} - 1 \right\rfloor$ many disks into the cabin.

Definition 2.13. A vertex gadget, shown in Figure 2.4(b), for an embedded vertex v consists of a cycle of $C_V = \lceil \pi s/r \rceil$ many intersecting disks of radius r which are centered on a circle of radius s centered at v.

Since the first and last disks of an edge gadget for an edge $e = \{u, v\}$ have a point at distance s from u and v respectively, it is easy to see that if $C_V \geq 4$, these disks have a non-empty intersection with the disks of the vertex gadgets for u and v respectively.

We now choose the radius r of the disks and the height h of the edge gadgets such that the seven constraints below hold. These constraints ensure that no two vertex gadgets intersect (1), no two edges gadgets intersect (3, 4, 5), no edge gadget intersects any non-incident vertex gadget (2) and all edge gadgets have an equal amount of disks (6), and a single edge gadget contains more disks than all the vertex gadgets combined (7). If we set $a = 1/4$ it follows that b is at least $1/2 - 2s$ in every edge gadget. We furthermore

fix s to $r(\sqrt{36n^4+1}-1)$. Choosing the radius r to be $\frac{1}{40n^4}$ and the height h to be $\frac{1}{12n^2}$ it is easy to see that the following seven constraints are satisfied for all $n \geq 2$.

1. No two vertex gadgets intersect: $2(r+s) < 1$
2. No edge gadget intersects any vertex gadget other than the ones at its two endpoints: $r + s + h/2 < (2n^2 - 2n + 1)^{-\frac{1}{2}}$, using Lemma 2.9
3. No edge gadget intersects another edge gadget if the corresponding edges do not share a common vertex: $2\frac{h}{2} < (2n^2 - 2n + 1)^{-\frac{1}{2}}$
4. No disk placed in the cabin of an edge gadget intersects any disk in any incident edge gadget: $h/2 < \frac{s+a}{6n^2}$, using Lemma 2.10
5. No two disks contained in the hallways of two incident edge gadgets intersect each other: $s \geq \frac{r}{\sin(2\arctan(1/(6n^2))/2)} - r = r(\sqrt{36n^4+1}-1)$, using Lemma 2.10 and the fact that the minimum distance of two disjoint lines segments occurs at an endpoint of one
6. The cabin of every edge gadget is big enough so that the whole gadget contains a path of C_E many disks: $\left\lceil \frac{\sqrt{2}n-2s}{2r} \right\rceil - 2\left\lceil \frac{a}{2r} \right\rceil \leq \left\lfloor \frac{h}{2r} \right\rfloor \cdot \left\lfloor \frac{1-2(s+a)}{2r} - 1 \right\rfloor$
7. An edge gadget contains more disks than all the vertex gadgets combined: $C_E > nC_V$

Plugging the calculated values $r = \frac{1}{40n^4}$ and $s = r(\sqrt{36n^4+1}-1)$ into C_E, which we earlier defined to be $\lceil \frac{\sqrt{2}n-2s}{2r} \rceil$ yields that an edge gadget consists of $\left\lceil (20\sqrt{2}n^5 - \sqrt{36n^4+1}) + 1 \right\rceil$ disks. Analogously, plugging the calculated values into C_V, which is defined as $\lceil \pi s/r \rceil$, yields that a vertex gadget consists of

$\left\lceil \pi(\sqrt{36n^4+1}-1) \right\rceil$ many disks. We can conclude that the above construction can be computed in polynomial time.

In order to finish the proof of Theorem 2.11, we need to show how to retrieve a solution for an instance I_2 of the Planar Subdivision problem from the solution for I_1, where I_1 is built using the construction described above.

Lemma 2.14. *An instance of $I_2 = (G_2, S_2)$ of the Planar Subdivision problem has a solution of size at most k_2 if and only if I_1 of the Point Isolation problem has a solution of size at most $C_E(k_2+1)-1$, where I_1 is built out of I_2 using the construction described above.*

Proof. Removing any disk D from an edge gadget merges the two adjacent regions in $\mathbb{R}^2 \setminus (\bigcup \mathcal{D} \setminus \{D\})$. Thus, I_2 has a solution consisting of the corresponding k_2 edges, if and only if I_1 uses all the disks of k_2 edge gadgets. On the other hand, removing a disk from a vertex gadget does not necessarily merge two regions. Thus, a solution for I_1 may contain no disks or all disks of any vertex gadget. Since according to constraint 7, it holds that $nC_V < C_E$, it follows that if I_2 has a solution of size at most k_2, then I_1 has a solution of size at most $C_E k_2 + nC_V \leq C_E(k_2+1) - 1$. On the other hand, if I_1 has a solution consisting of at most $k_1 = C_E(k_2+1) - 1$ disks, if follows that I_2 uses at most $k_1/C_E < k_2 + 1$ edges and thus the lemma follows. □

Since the two-point separation algorithm of [20] can be used to test whether all points of P are separated in a potential solution $\mathcal{D}' \subseteq \mathcal{D}$ of the Point Isolation problem, it follows that its decision version is indeed contained in NP.

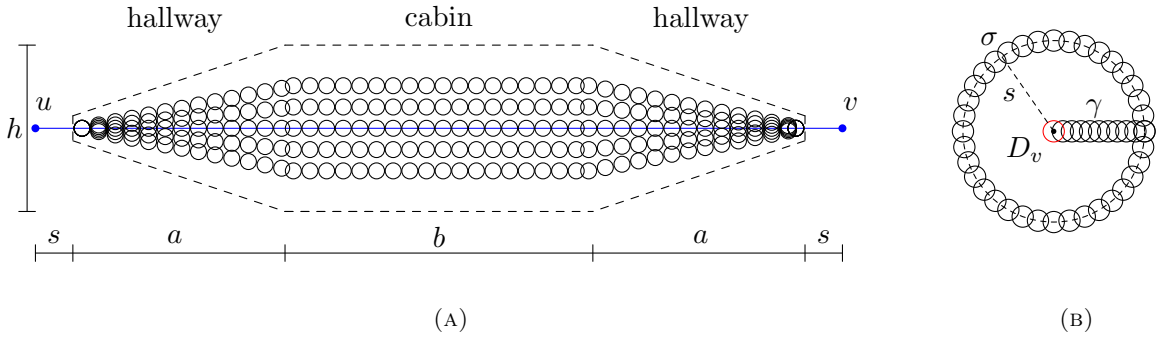

(A) (B)

FIGURE 2.5: (a) An example of an Edge Gadget for an edge of weight 5 in the proof of Theorem 2.15. (b) The vertex gadget for vertex v, where γ is a path of unit disks, where each disks shown corresponds to 16 perturbed copies of a single disk and σ is a cycle where each disks shown corresponds to 16 perturbed copies of a single disks; the red disk is the centroid disk D_v.

2.4 Multiterminal Cut Problem on Unit Disk Graphs

In this section we are going to prove the following theorem about unit disk graphs, i.e., intersection graphs of a collection of unit disks in the plane.

Theorem 2.15. *The (Unweighted) Multiterminal Cut problem remains* NP*-complete on unit disk graphs, if k is not fixed.*

Proof. We make a reduction from the following weighted version of the planar Multiterminal Cut problem, which was proven to be NP-complete in [32].

Problem 2.4.1 ([32])**.** *Given an edge-weighted planar graph $G = (V, E)$, where each edge has a weight in $\{1, \ldots, 5\}$ and each vertex has degree at most 3, and given set $S \subseteq V$ of k terminals, find a minimum weight set $E' \subseteq E$ of edges such that in $G' = (V, E \setminus E')$ there is no path between any two terminals in S.*

Let $I_2 = (G_2, S_2)$ be an instance of the restricted version of the planar Multiterminal Cut problem described in Problem 2.4.1. We embed $G_2 = (V_2, E_2)$ crossing free into an $n \times n$ grid, replace each edge by an edge gadget of Definition 2.16 (see Figure 2.5(a)) and

each vertex by a vertex gadget of Definition 2.17 (see Figure 2.5(b)). We then construct the embedded unit disk graph $G_1 = (V_1, E_1)$ of the disks in all vertex and edge gadgets.

Definition 2.16 (Edge Gadget)**.** An edge gadget for an embedded edge $e = \{u, v\}$ of weight $w \in \{1, \ldots, 5\}$ consists of w paths of radius r disks, which are pairwise disjoint in the cabin of the gadget but start and end from the same point respectively, as shown in Figure 2.5(a). Both, start and end point are w slightly perturbed copies of a single disk touching the boundary of the gadget.

It thus holds that inside an edge gadget for an edge $\{u, v\}$ of weight w, removing w many edges (from the cabin) in G_1 disconnects u from v inside the subgraph of G_1 corresponding to this edge gadget, but removing fewer than w edges leaves u and v connected in this subgraph.

Definition 2.17 (Vertex Gadget)**.** For a vertex v, the vertex gadget consists of 16 slightly perturbed copies of a cycle of $C_V = \lceil \pi s/r \rceil$ disks of radius r which are centered on a circle of radius s around v as shown in Figure 2.5(b). We denote the arrangement of those 16 copies by σ. Furthermore, we place a *centroid disk* D_v at the coordinates of v and connect it to σ by 16 slightly perturbed copies of a path of radius r disks which we denote by γ in Figure 2.5(b).

We build the set $S_1 \subseteq V_1$, by putting the vertex corresponding to the centroid disk D_v into S_1 for every vertex gadget, corresponding to a vertex $v \in S_2 \subseteq V_2$. We thus obtain an instance I_1 of the Unweighted Multiterminal Cut problem on unit disk graphs.

Since both edge and vertex gadgets have the same dimension as the gadgets of the previous section, it holds that all edge gadgets are pairwise disjoint, all vertex gadgets are pairwise disjoint and no edge gadget intersects any vertex gadget other than the

ones of its end vertices. There, each of the w copies of the last disks intersect at least one 16-disk cluster of the vertex gadget. Note that in the vertex gadget, each vertex in the corresponding subgraph of G_1 has degree at least 16. Thus removing less than 16 edges in this subgraph does not disconnect any two vertices in this subgraph; therefore, an optimal solution for I_1 does not remove any edge from the subgraphs of the vertex gadgets. Furthermore, it is easy to see that $E_2' \subseteq E_2$ is a solution of weight k for I_2 if and only if S_1 gets disconnected in I_1, by removing a total of k edges in the edge gadgets corresponding to the edges in E_2'. Since the reduction from I_2 to I_1 can be done in polynomial time, Theorem 2.15 follows. □

2.5 All-Cells-Connection Problem for Unit Disks

While the Point Isolation problem can be interpreted as asking for the minimum number of sensors which have to be turned *on* to detect any transition between the input points, the All-Cells-Connection problem, as defined below, can be interpreted as asking for the minimum number of sensors which have to be turned *off* so that an intruder can transition between *any* connected region induced by the sensors (see Figure 2.6). In the context of line segment sensors, the All-Cells-Connection problem was shown to be NP-complete in [4] and we reuse their reduction idea in this section for unit disk sensors.

Problem 2.5.1 (All-Cells-Connection for unit disks)**.** *Given a set $\mathcal{D}$ of unit disks embedded in the plane, remove a minimum cardinality subset $\mathcal{D}'$ of $\mathcal{D}$ such that the plane minus the arrangement of the remaining disks, i.e., $\mathbb{R}^2 \setminus \bigcup(\mathcal{D} \setminus \mathcal{D}')$, consists of a single connected region.*

Theorem 2.18. *The All-Cells-Connection problem for unit disks is* NP*-complete.*

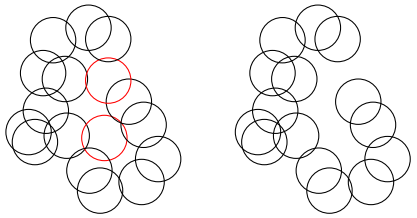

FIGURE 2.6: Left: A set $\mathcal{D}$ of unit disks with its *merging* subset $\mathcal{D}'$ in red. Right: Illustration of $\mathbb{R}^2 \setminus \bigcup(\mathcal{D} \setminus \mathcal{D}')$, consisting of a single connected region.

Proof. For a given graph, the Feedback Vertex Set problem (FVS) asks for the minimum cardinality subset of vertices to be removed such that the remaining subgraph is acyclic. In order to prove Theorem 2.18, we are going to reduce a restricted version of FVS to the All-Cells Connection problem for unit disks. In [89] (see also [90]), it is shown that FVS is NP-complete in planar graphs of maximum degree 4. Given such an instance $G = (V, E)$, we embed it into an $n \times n$ grid and replace each edge with the edge gadget[3] of Definition 2.12. We further replace each vertex with a vertex gadget of Definition 2.19 and we refer to the collection of all disks used in the vertex and edge gadgets as $\mathcal{D}$.

The idea of the vertex gadget for a vertex v is to center a radius r disk D_v at v and connect the, at most 4, incident edge gadgets to D_v, using simple paths which are pairwise disjoint; in Definition 2.19, we will argue that this is always possible. It then follows that removing disk D_v merges the, up to 4 regions, adjacent to D_v. Removing a disk in an edge gadget for an edge $\{u, v\}$ never merges more regions than removing D_v (or D_u). Therefore, G has a FVS of size k if and only if the corresponding All-Cells-Connection Problem for $\mathcal{D}$ has a solution of size k.

Definition 2.19. We first provide a conceptual description of the vertex gadget and then we will replace its conceptual paths by pairwise disjoint paths of intersecting disks of radius r. For an embedded vertex v, which has $0 \leq i \leq 4$ incident edges $e_1, \dots e_i$,

[3]The fact that all edge gadgets contain the same amount of disks is irrelevant for the reduction. The only relevant property is that the disks in each edge gadget form a simple path.

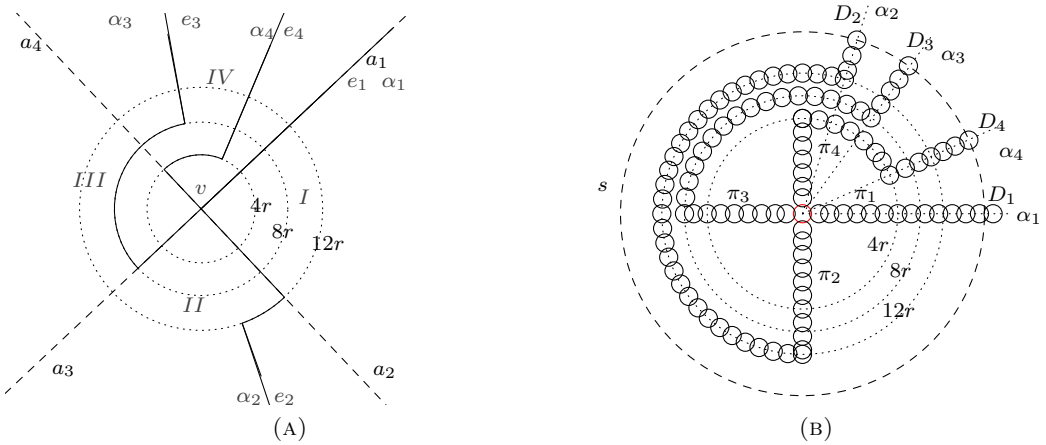

FIGURE 2.7: (a) A conceptual illustration of the Vertex gadget. The cross, satisfying the four constraints of Definition 2.19, is depicted as dashed rays $a_1, \ldots, a_4$. Edge e_1 lies on a_1. Axis $a_2, \ldots, a_4$ are connected to the edges $e_2, \ldots, e_4$, which are truncated by $12r, 8r, 4r$ from v respectively, by arcs of radius $12r, 8r, 4r$ until reaching $\alpha_2, \alpha_3, \alpha_4$ and thereby $e_2, \ldots, e_4$. (b) An illustration of an actual realization of a vertex gadget, having three edges in quadrant IV. The red disk D_v is centered at v. The dashed circle has radius s and the dotted circles have radii $12r, 8r, 4r$ respectively. D_v is connected by the paths $\pi_1, \ldots, \pi_4$ to the last disks $D_1, \ldots, D_4$ of the edge gadget for the edges $e_1, \ldots, e_4$ respectively, which have an angle of $\alpha_1, \ldots, \alpha_4$ with respect to the x-axis.

where in the grid embedding edge e_j has an angle α_j with the x-axis. We center a *cross* at v, which consist of four perpendicular rays $a_1, \ldots, a_4$ introducing four quadrants I, II, III, IV in clockwise order, as shown in Figure 2.7(a). The cross is oriented such that (1) at least one of the incident edges lies on ray a_1, (2) its clockwise quadrant I does not contain any edge incident to v, (3) quadrant II contains at most one edge incident to v and (4) quadrant II and III together contain at most two edges incident to v. A moment of thought shows that it is indeed always possible to center a cross at v satisfying these four constraints. Next, we truncate edges $e_2, \ldots, e_i$ to end at a distance of $12r, 8r, 4r$ respectively from v. We then connect the truncated edges $e_2, \ldots, e_i$ to the rays $a_2, \ldots, a_i$ respectively, by extending an arc of radius $12r, 8r, 4r$ from $a_2, \ldots, a_i$ until we reach $\alpha_2, \ldots \alpha_i$ respectively and thereby orthogonally connecting all the edges $e_1, \ldots e_i$ to v, as shown in Figure 2.7(a). Denoting by π_j the path connecting edge e_j to v, we realize the construction of the vertex gadget by first centering a radius r disk D_v at v. We then center radius r disks on the paths $\pi_1, \ldots, \pi_i$ such that consecutive disks

intersect. Note that any point in the last disk of any incident edge gadget has distance $s - r = r(\sqrt{36n^4 + 1} - 2)$ to v, which is greater than $13r$ for any $n \geq 2$, and thus the whole vertex gadget is contained inside the circle of radius s around v. Furthermore, it is easy to see that all paths of disks $\pi_1, \ldots, \pi_i$ are pairwise disjoint.

Since this reduction can be done in polynomial time, Theorem 2.18 follows.

□

Corollary 2.20. *Given a set P of k points and a set $\mathcal{D}$ of unit disks in the plane separating P, it is* NP*-complete to find a minimum cardinality subset $\mathcal{D}'$ of $\mathcal{D}$ such that there is a path between any two point in P not intersecting any disk in $\mathcal{D} \setminus \mathcal{D}'$.*

Note that the problem of this corollary is a generalization of the Barrier Resilience problem from 2 to k points. The computational complexity of the Barrier Resilience problem for 2 points is still open in the unit disk setting [13].

Chapter 3

Covering the Boundary of a Simple Polygon

3.1 Introduction and Main Results

For two points u and v in a simple polygon P, the *geodesic distance*, denoted by $d(u, v)$, is the length of the shortest path between u and v inside P. A *geodesic unit disk* $D(v)$ centered at a point $v \in P$ is the set of all points in P whose geodesic distance to v is at most 1.

The *boundary* of $D(v)$, denoted by $\partial D(v)$, consists of all points in P which are either exactly at distance 1 from v or at distance at most 1 from v but contained in the polygon boundary ∂P. The *interior* of $D(v)$, denoted by $int(D(v))$, consists of all the points of $D(v)$ not contained on the boundary of $D(v)$, i.e., $int(D(v)) = D(v) \setminus \partial D(v)$, as shown in Fig. 3.1.

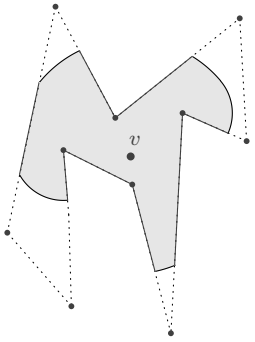

FIGURE 3.1: A polygon (dotted) containing a geodesic disk centered at v, whose interior is depicted in gray and its boundary is drawn in black.

A collection of geodesic disks *covers* the polygon boundary ∂P, if each point of ∂P is contained in at least one disk. In this chapter we present an $O(n \log^2 n + k)$ time 2-approximation algorithm which finds a collection of geodesic unit disks covering the boundary of a simple polygon on n vertices, with k denoting the number of disks found by the algorithm. The algorithm then returns the centers of the disks. We consider the setting where the centers can be placed anywhere inside the polygon, but the algorithm can be easily modified to restrict the centers to lie on ∂P. Furthermore, the *number* of disks can be computed in time $O(n \log^2 n)$.

While it follows from Theorem 7 of [98] that our problem is NP-hard in polygons with holes, its complexity remains open in simple polygons.

The main motivation for studying this problem comes from sensor networks, where *Barrier Coverage* problems have been studied extensively (see for example [14],[24],[25],[66], [68],[83],[87]). In a Barrier Coverage problem the goal is to place few sensors or guards to detect any intruder into a given region. The algorithm in this chapter can be applied to this context: given a region, bounded by a piecewise linear closed border, such as a fence, place few guards inside the fenced region, such that wherever an intruder cuts through the fence, the closest guard is at most distance one away. Another way of looking at this problem is from an Art Gallery perspective (see for example [73]), where the

polygon represents a gallery and, regardless where on the wall a painting is hanged, the closest guard is at most a distance one away.

3.1.1 Related Work

Several papers ([46],[56],[61],[63],[88],[100]) study full coverage of geometric regions with Euclidean disks. For an overview of optimal coverings of squares and triangles with disks see Chapter 1.7 of [19].

In the context of Barrier Coverage, [21] presents a polynomial time algorithm which for two points in the plane and a set of Euclidean disks selects a minimal subset of the disks which separates the two points. Extending the problem to k points, an $O(1)$-approximation algorithm was presented in [47] and NP-hardness was shown in [78]. The same two point separation problem was studied in [5] when segments instead of disks are given.

Covering a simple polygon with a single geodesic disk of minimum radius has been studied in [79] and a linear-time algorithm is presented in [12]. An output sensitive algorithm for computing geodesic disks for a given set of centers and a fixed radius is presented in [17].

Chapter Organization

This chapter is organized as follows. In Section 3.2 we present the algorithm and show that it runs in time $O(n \log^2 n + k)$. In Section 3.3 we prove that the number of centers placed by the algorithm is at most twice the minimum number of centers needed to cover the polygon boundary. In Section 3.4 we show that a simple linear time algorithm

achieves an asymptotically optimal approximation ratio when the polygon perimeter is much larger than n.

3.2 The Algorithm and Its Running Time

Our algorithm makes use of several properties of geodesic Voronoi diagrams which we review below.

3.2.1 Geodesic Voronoi diagrams

A *furthest-site geodesic Voronoi diagram* of k sites in a simple polygon P on n vertices is a decomposition of P into cells such that all points in a cell have the same site *furthest* away from them (in the geodesic metric). As shown in [8], it has combinatorial complexity $O(n+k)$ and can be constructed in time $O((n+k)\log(n+k))$. In Section 2.8 of [8] it is shown that these combinatorial and time complexities are with respect to a *refinement* (also called a shortest path partition) of the Voronoi edges. For all points on a refined edge it holds that their shortest paths to each of the two furthest sites are combinatorially equivalent, i.e., they consist of the same sequence of polygon vertices respectively. Furthermore, Section 3.3 of [8] defines for each of the $O(n+k)$ refined edges, and for each of the two furthest sites s_1 and s_2 defining a Voronoi edge e, the *anchor points* $a_e(s_1)$, $a_e(s_2)$ which are the last points on the shortest path from s_1, s_2 respectively to any point on e. Those anchors can be computed in total $O(n+k)$ time and each time we compute a furthest-site geodesic Voronoi diagram we store the anchors as well as the distance to its site at the refined Voronoi edges. An additional property of this Voronoi diagram is that its edges form a tree, rooted at the *geodesic center* of the k sites, which is defined as the point that minimizes the maximum distance to any

of the sites (see Corollary 2.9.3 of [8]). Therefore, the geodesic center of the sites can be obtained within the same time bound.

The second data structure we use is the *closest-site geodesic Voronoi diagram* which, for k sites in a simple polygon P on n vertices, is a decomposition of P into cells such that all points in a cell have the same site *closest* to them (in the geodesic metric). It has combinatorial complexity $O(n+k)$ and it can be constructed in time $O((n+k)\log(n+k))$ (see [6]).

3.2.2 The ContiguousGreedy Algorithm

In this section we describe a greedy 2-approximation algorithm which finds a collection of geodesic unit disks which cover the boundary of P and returns the set of disk centers. It starts at vertex v_1 of the vertices $v_1, \ldots, v_n$ of P and iteratively extends a contiguous cover Γ of ∂P (in clockwise order) by the maximum amount that can be covered with a single geodesic disk. We denote the clockwise endpoint of Γ by c, thus initially $\Gamma = \{v_1\}$ and $c = v_1$.

We cover segment portions longer than 2 in time linear in the minimum number of disks needed to cover them. With v_u denoting the first uncovered vertex in the current iteration, we partially cover $\overline{cv_u}$ by adding $\lceil d(c, v_u)/2 \rceil - 1$ centers sequentially on $\overline{cv_u}$. By this, we assure that none of those disks contains v_u and, since each disk contains a boundary portion of length 2, the disks placed are indeed optimal with respect to the greedy contiguous extension criterion.

Definition 3.1. For a polygonal chain C, we denote by $\|C\|$ the sum of the lengths of its line segments and we refer to the number of vertices of C by $|C|$.

Definition 3.2. For two points $u, v \in \partial P$, we denote the portion of ∂P in clockwise orientation between u and v by $\partial P[u, v]$.

If c does not lie on a long segment, we compute the next endpoint c' which extends Γ in clockwise order by a maximum length boundary portion which can be covered by a single geodesic unit disk. We do this by finding the first vertex v_u (in clockwise order) such that $\partial P[c, v_u]$ cannot be contained in a single geodesic unit disk. This test is done by calling the TestCover(c, v) procedure discussed below, which, for a boundary point c and a vertex v tests whether $\partial P[c, v]$ can be covered with a single geodesic unit disk. If v_i is the first vertex in clockwise order after c, we find v_u by first using exponential search with the TestCover predicate with c fixed and v set to $v_{i+1}, v_{i+2}, v_{i+4}, ...v_{i+2^k}, ...$ respectively in consecutive steps until TestCover returns $false$ or $i + 2^k > n$. This defines an index-interval containing the index u which can then be found using a simple binary search.

After finding v_u and thereby fully determining the sequence of vertices covered in the current iteration, we use the AugmentShort procedure – discussed below – to compute the new endpoint c' of Γ as well as the center of the next disk.

ContiguousGreedy

$c \leftarrow v_1,\ v_u \leftarrow v_2$

$S \leftarrow \emptyset$

while ∂P not covered:

1. If $\overline{cv_u}$ is longer than 2

   compute centers on $\overline{cv_u}$ at steps of 2; add them to S; update c

2. Update v_u to the first vertex s.t. $\partial P[c, v_u]$ cannot be covered by a single disk, using Exponential and Binary Search with predicate TestCover

3. Use AugmentShort to cover the vertices between c and v_u, and a maximal portion of the edge $\overline{v_{u-1}v_u}$; add new center to S and update c

end while

return S

Definition 3.3. For two points u, v in a simple polygon P, we denote the shortest path in P between u and v by $\pi(u, v)$. We denote the number of its vertices by $|\pi(u, v)|$.

Definition 3.4 ([96])**.** A set Q inside a simple polygon P is called *geodesic convex*, if for any two points $u, v \in Q$, the shortest path $\pi(u, v)$ is contained in Q.

TestCover(c, v)**.** This procedure tests for a boundary point c and a polygon vertex v whether $\partial P[c, v]$ can be covered with a single geodesic unit disk. Observe that if a geodesic unit disk can cover a set of points, then a geodesic unit disk centered at the *geodesic center* of those points obviously also covers them. Let $U = U(c, v)$ denote the sequence of point c and all polygon vertices up to (and including) v in clockwise order. TestCover computes the geodesic center of U and returns true iff it has distance at

most one to all points in U.

Implementation details. We compute the geodesic center of U in a smaller polygon Q containing U. We let Q be $\partial P[c, v] \circ \pi(v, c)$, with $\circ$ denoting the concatenation of two polygonal chains sharing two endpoints. Note that Q may have touching sides, but it is not self-intersecting. Such polygons are referred to as *weakly simple* polygons ([38]) and the geodesic distance within them is well defined. Since Q is the concatenation of a boundary part of P and a shortest path in P it follows that Q is geodesic convex in P, thus implying that the geodesic center of U in Q is the same point as the geodesic center of U in P. We find this geodesic center point by computing the furthest-site geodesic Voronoi diagram $\mathcal{VP}_Q(U)$ of the sites U in Q, traversing the (oriented) Voronoi edges to the root and thereby obtain the geodesic center of U (see Section 3.2.1). Then, for each site in U we test whether the distance to the geodesic center is at most one.

Computational complexity. Computing $\pi(v, c)$ takes time $O(|\pi(v, c)| \log n)$ after $O(n)$ global pre-processing time, using the algorithm of [52]; concatenating two polygonal chains to construct Q takes constant time. Computing $\mathcal{VP}_Q(U)$ takes $O(|Q| \log(|Q|))$ time and the geodesic center can be obtained from $\mathcal{VP}_Q(U)$ in the same time bound. Computing the distance from the geodesic center to all sites in U can be done in time $O(|Q|)$ (see [51]), by building the shortest path tree rooted at the geodesic center. Therefore, the procedure has an overall time complexity of $O(|Q| \log n)$.

Knowing the first vertex v_u such that $\partial P[c, v_u]$ cannot be covered with a single geodesic unit disk, we compute the center of the next disk and compute the new endpoint c' of Γ the following AugmentShort procedure.

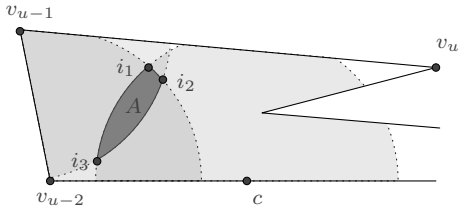

FIGURE 3.2: Illustration of A, i.e., the intersection of the geodesic unit disks centered at points in $\overline{U} = \{c, v_{u-2}, v_{u-1}\}$ as well as the disk-disk intersection points $I = \{i_1, i_2, i_3\}$.

AugmentShort(c, v_u)**.** For the new endpoint c' of Γ it needs to hold that $\partial P[c, c']$ can be covered with one geodesic unit disk, and for any $c'' \in \partial P$, with $\|\partial P[c, c'']\| > \|\partial P[c, c']\|$ it is not possible to cover $\partial P[c, c'']$ with a single geodesic unit disk. Let $\overline{U} = \overline{U}(c, v_{u-1})$ denote the clockwise sequence of point c and all vertices up to (and including) v_{u-1}. We construct $Q = \partial P[c, v_u] \circ \pi(v_u, c)$ and denote by A the intersection of the geodesic unit disks centered at the points $\overline{U}$ in Q. This intersection is non-empty by construction and the center of the next disk lies in A. We denote by I the set of all disk-disk intersection points on ∂A as shown in Fig. 3.2. Lemma 3.5 below justifies the steps taken to find c'.

Lemma 3.5. *Given a simple polygon P, let A be the non-empty intersection of a collection of geodesic unit disks in P and let $\overline{\alpha\beta}$ be a line-segment in P, such that for all $a \in A$, $d(a, \alpha) \leq 1$ and $d(a, \beta) > 1$. For any point $c \in \overline{\alpha\beta}$ and any disk center q furthest away from c, $d(c, A) = 1$ if and only if either:*

a) $d(c, q) = 2$ and $\pi(c, q) \cap \partial D(q) \in A$, or

b) $d(c, I) = 1$ and $\pi(c, q) \cap \partial D(q) \notin A$,

with I denoting the disk-disk intersection points on ∂A and $d(c, Y) = \min_{s \in Y} d(c, s)$, for a point set Y in P.

Proof. (Lemma 3.5) Let p be the point in A closest to the point c and by q a center farthest from c. Notice that since A is geodesic convex, p is unique and it lies on ∂A. We prove Lemma 3.5 with the help of the following two observations.

Observation 3.2.1. *If $p \in \partial D(q')$ for some center q' and $\pi(q', c) \cap \partial D(q') \notin A$ then $p \in I$.*

Proof. Let $p' = \pi(q', c) \cap \partial D(q')$ and assume that $p \notin I$, thus p is in the interior of all the disks defining A, other than $D(q')$. Furthermore, since $D(q')$ is geodesic convex, $\pi(p, p') \cap A$ contains a point a, with $a \neq p$. Since $d(p', q') = d(p, q') = 1$ and p' is on the shortest path from q' to c, by uniqueness of the shortest path, $d(c, p') < d(c, p)$. Then, by Lemma 3.6, $d(c, a) < \max\{d(c, p), d(c, p')\} = d(c, p)$, contradicting that p is the point in A closest to c. □

Observation 3.2.2. *If $p \in A \setminus I$ then $p = \pi(c, q) \cap \partial D(q)$.*

Proof. Since $p \notin I$ there is a unique center q' such that $p \in \partial D(q')$. As shown in Observation 3.2.1, if $p \in D(q')$ and $\pi(c, q') \cap D(q') \notin A$ then $p \in I$. Therefore $\pi(c, q') \cap D(q')$ is contained in A and we denote this point by p'.

Observe that p' is contained in $\pi(c, q')$ and is at distance 1 from q', thus $d(q', c) = d(q', p') + d(p', c) = 1 + d(p', c)$. Clearly, $d(q', c) \leq d(q', p) + d(p, c) = 1 + d(p, c)$ and since p is the closest point in A to c, it follows that $p = p'$. Now observe that, since $d(q', c) = 1 + d(p, c)$ and the distance between p and any other center is less than 1 (because $p \notin I$), q' is the farthest center, i.e., $q' = q$. Therefore $p = \pi(c, q) \cap D(q)$ holds as claimed. □

We now prove Lemma 3.5 :

"⇒". We distinguish two cases based on whether $p \in A \setminus I$ or $p \in I$. If $p \in A \setminus I$, then $p = \pi(c,q) \cap D(q)$ as shown in Observation 3.2.2. Therefore, $\pi(c,q) \cap D(q) \in A$ and $d(c,q) = d(c,p) + d(p,q) = 2$, and thus condition a) holds. If $p \in I$ and $\pi(c,q) \cap D(q) \notin A$ condition b) holds. Otherwise if $p \in I$ and $\pi(c,q) \cap D(q) \in A$, let $p' = \pi(c,q) \cap D(q)$. We can write the distance $d(c,q)$ as $d(c,p') + d(p',q)$. Since p' lies on $\partial D(q)$, $d(p',q) = 1$ and since $d(c,A) = 1$ this implies that $d(c,p') \geq 1$. Therefore, distance $d(c,q) \geq 2$.

By the triangle inequality it also holds that $d(c,q) \leq d(c,p) + d(p,q)$. Since p is the closest point in A to c, $d(c,p) = 1$ by hypothesis. Since p lies on $\partial D(q)$, $d(p,q) = 1$. Therefore, $d(c,q) \leq 2$ and combining this with $d(c,q) \geq 2$ from above, $d(c,q) = 2$ and again condition a) holds.

"⇐" a) Let $p' = \pi(c,q) \cap \partial D(q)$. Then $d(c,A) \leq d(c,p') = d(c,q) - d(q,p') = 1$. If $d(c,A) < 1$, by definition $d(c,p) < 1$. Since $p \in A$, $d(p,q) \leq 1$ and by the triangle inequality, $d(c,q) \leq d(c,p) + d(p,q) < 2$ which contradicts $d(c,q) = 2$.

b) Since $d(c,I) \leq 1$ and $I \subseteq A$, obviously $d(c,A) \leq 1$. For $p \in A$ the closest point to c in A, assume that $d(c,p) < 1$. Since $d(c,I) = 1$, $p \notin I$. Therefore, by Observation 3.2.2, $p = \pi(c,q') \cap \partial D(q')$, and thus this intersection is in A contradicting the hypothesis.

□

We use the following steps to determine c' on $e = \overline{v_{u-1}v_u}$.

Step 1) Find the point x_1 on e closest to v_u, whose distance to its furthest point q in $\overline{U}$ is exactly 2 and $\pi(x_1,q) \cap \partial D(q) \in A$, if such a point x_1 exists.

Step 2) Find the point x_2 on e closest to v_u, whose distance to its closest point in I is exactly 1.

Step 3) Set $c' \leftarrow x_2$ if x_1 does not exist or $d(x_2,v_u) < d(x_1,v_u)$. In this case we add the

point in I closest to c' as the new disk center to the set S of centers. Otherwise $c' \leftarrow x_1$ and the point $\pi(x_1, q) \cap \partial D(q)$ is the new disk center which gets added to S.

Note that since v_{u-1} will be covered in this iteration and v_u won't be covered, $d(v_{u-1}, A) \leq 1 < d(v_u, A)$. By continuity of the geodesic distance, there is a point c' on e, with $d(c', A) = 1$ and thus by Lemma 3.5 either x_1 or x_2 exists.

In Step 1, to find x_1 if it exists, we construct the (refined) furthest-site geodesic Voronoi diagram of the sites $\overline{U}$ in Q and traverse the Voronoi vertices $\gamma_1, \ldots, \gamma_m$ on e, ordered in the direction from v_u to v_{u-1} and set $\gamma_{m+1} = v_{u-1}$. For each such vertex we check in $O(\log |Q|)$ time whether the distance to (one of) its furthest site(s) is at most 2, using an $O(\log |Q|)$ time shortest path query ([52]) after pre-processing Q in $O(|Q|)$ time. Once we find the first γ_j with distance at most 2, if it exists, this determines a sub-segment $\overline{\gamma_j \gamma_{j-1}}$ on e containing a point x at distance exactly 2 from its furthest site q. Note that since the shortest paths to the furthest site q have the same combinatorial structure for all points on the refined Voronoi edge $\overline{\gamma_j \gamma_{j-1}}$, we find the point at distance 2 to q in constant time since we stored the anchor point $a_{\overline{\gamma_j \gamma_{j-1}}}(q)$ at the edge $\overline{\gamma_j \gamma_{j-1}}$ (see Section 2.1). We check if $\pi(x, q) \cap \partial D(q) \in A$, by computing $D(q)$ in time $O(|Q|)$ using [51] and finding in $O(\log |Q|)$ time the arc α of $D(q)$ separating q from x. We traverse the edges of $\pi(x, q)$ and for each edge we test in $O(1)$ time if it intersects α. Denoting the intersection point by p, we check if $p \in A$, by computing the shortest path tree to the sites in U and test if the distance to all sites is at most 1 in time $O(|Q|)$. If this intersection is in A, we set x_1 to x.

Claim 3.2.3. *There can be at most two points on e that have distance exactly 2 from their respective furthest site; if there are two such points, one of them must be v_{u-1}.*

We prove this claim using the following lemma.

Lemma 3.6 (Lemma 1 [79]; see also Lemma 2.2.1 [8])**.** *Given three points a, b, c in a simple polygon, for $x \in \pi(b, c)$, the distance $d(a, x)$ is a convex function on $\pi(b, c)$, with $d(a, x) < \max\{d(a, b), d(a, c)\}$.*

Proof. Assume that there are two points p_1 and p_2 on $e \setminus \{v_{u-1}\}$ that are at distance 2 from their respective furthest sites, with p_1 closer to v_{u-1} than p_2, thus $p_1 \in \overline{v_{u-1}p_2} \setminus \{v_{u-1}, p_2\}$. Let q_1 be a center furthest away from p_1. Clearly $d(q_1, v_{u-1}) \leq 2$ since both q_1 and v_{u-1} are at distance at most 1 from any point in A. Since $d(q_1, p_2) \leq 2$ and $p_1 \in \overline{v_{u-1}p_2} \setminus \{v_{u-1}, p_2\}$, by Lemma 3.6 $d(q_1, p_1) < 2$, contradicting the assumption that $d(q_1, p_1) = 2$. □

According to the above claim, the only other candidate for x_1 is v_{u-1}. Thus, if $\pi(x_1, q) \cap \partial D(q) \notin A$ we check in $O(\log |Q|)$ time if the point v_{u-1} is at distance exactly 2 from its furthest site and if so, we set x_1 to v_{u-1}. If x_1 exists $\partial P[c, x_1]$ can be covered with one geodesic unit disk, because the point $\pi(x_1, q) \cap \partial D(q)$ has distance exactly 1 to x_1 and lies in A.

In Step 2, to find x_2, we first construct the set I of the disk-disk intersection points of A; we do this without explicitly computing A. To construct I, we look at the furthest-site geodesic Voronoi diagram of the sites $\overline{U}$ in Q constructed in the Step 1. Since any point in I has two points in $\overline{U}$ at distance 1, every point in I lies on a Voronoi edge. For every site $s \in \overline{U}$ we look at the refined edges of $\sigma(s)$ and for such edge e we access its

anchor point $a_e(s)$ as well as the distance from s to the endpoints of e, in constant time. We test if there is a point on e having distance 1 to s, again in $O(1)$ time. If such a point exists then this is a disk-disk intersection point and we add it to I. Since we need constant time for each refined Voronoi edge, I can be computed in total time $O(|Q|)$.

Having computed I, we construct the closest-site geodesic Voronoi diagram of the sites I in Q. We traverse the Voronoi vertices $\gamma_1, \ldots, \gamma_m$ on e, ordered in the direction from v_u to v_{u-1} and set $\gamma_{m+1} = v_{u-1}$. For each such vertex we check whether the distance to (one of) its closest site(s) is at most 1 again by an $O(\log n)$ time shortest path distance query. Once we find the first such vertex γ_j on $e = \overline{v_{u-1}v_u}$, if it exists, we have determined a sub-segment $\overline{\gamma_j\gamma_{j-1}}$ on e where x_2 lies. Letting $i \in I$ be the corresponding closest site, by Lemma 3.6, we find the point in $\overline{\gamma_j\gamma_{j-1}}$ at distance 1 from i by computing the intersection point of a geodesic unit disk centered at i with $\overline{\gamma_j\gamma_{j-1}}$, in time $O(|Q|)$, using the funnel algorithm of [51].

There can be at most two points on e that have distance exactly 1 from i; if there are two such points, one of them must be v_{u-1}. This can be seen directly from the fact that $d(i, v_{u-1}) \leq 1$, and Lemma 3.6. We set x_2 to the one closer to v_u. It is easy to see that x_2 is feasible, i.e., $\partial P[c, x_2]$ can be covered with one geodesic unit disk, because $d(i, \overline{U}) \leq 1$ and $d(i, x_2) = 1$.

In Step 3, $c' \leftarrow x_2$, if either x_1 does not exist or $d(x_2, v_u) < d(x_1, v_u)$, thus c' indeed extends Γ maximally because x_2 is the point on e closest to v_u having distance exactly 1 to the closest point in I, i.e., to the center of the geodesic unit disk placed in this iteration. Otherwise $c' \leftarrow x_1$ and x_1 is the point on e closest to v_u having distance exactly 2 to the furthest center in $\overline{U}$; any point on e closer to v_u has distance larger than

2 from that center and is thus infeasible.

Computational Complexity / Summary. Constructing Q takes time $O(|Q| \log n)$ as argued in the TestCover(c, v) paragraph before. Step 1 needs $O(|Q| \log |Q|)$ time to construct the geodesic furthest-site Voronoi diagram of $\overline{U}$ in Q and $O(|Q| \log |Q|)$ time to find a sub-segment of the edge e possibly containing x_1, since there are only $O(|Q|)$ Voronoi vertices in total and we spend $O(\log |Q|)$ on them for finding the sub-segment. The last step is to test if $\pi(x, q) \cap \partial D(q) \in A$, which takes time $O(|Q|)$ as argued above.

In Step 2, we spend $O(|Q|)$ time to construct the set I and $O(|Q| \log |Q|)$ time to construct the geodesic closest-site Voronoi diagram of the sites I. We then traverse edge e in $O(|Q| \log |Q|)$ time to find a sub-segment of the edge e possibly containing x_2, and determine x_2 on this sub-segment in $O(|Q|)$ time.

Thus the overall time spent in AugmentShort is $O(|Q| \log n)$.

Total Running Time.

Let $\mathcal{Q}$ be the set of all polygons constructed throughout the whole execution of ContiguousGreedy. In each polygon $Q \in \mathcal{Q}$ we spend $O(|Q| \log n)$ time in TestCover and possibly $O(|Q| \log n)$ time in AugmentShort as argued above. Since in each iteration of ContiguousGreedy, Γ is extended to cover at least one new polygon vertex, there are at most n iterations of the main *while* loop. Furthermore, covering long segments of ∂P takes total time $O(k)$. Since according to Lemma 3.7, $\sum_{Q \in \mathcal{Q}} |Q| = O(n \log n)$, the running time of ContiguousGreedy is $O(n \log^2 n + k)$.

Lemma 3.7. $\sum_{Q \in \mathcal{Q}} |Q| = O(n \log n)$.

Proof. Each polygon of $\mathcal{Q}$ constructed in the ContiguousCover algorithm has the form $Q = \partial P[c, w] \circ \pi(w, c)$, with c an arbitrary point on ∂P and w a vertex of P. We call $\partial P[c, w]$ the ∂-portion and $\pi(c, w)$ the π-portion of the polygon Q. Notice that every polygon constructed in AugmentShort was also constructed in a TestCover call and thus it suffices to bound the number of polygons constructed in all TestCover calls.

Observe that $|\mathcal{Q}| = O(n \log n)$, since in each iteration, Γ is extended to cover at least one new vertex, thus there are at most n iteration, and in each iteration we construct $O(\log n)$ polygons during Exponential and Binary Search. Observe that if every vertex of P is contained in $O(\log n)$ polygons of $\mathcal{Q}$ then $\sum_{Q \in \mathcal{Q}} |Q| = O(n \log n)$. This holds because for each $Q \in \mathcal{Q}$ there is at most one vertex of Q which is not a vertex in P, namely the point c.

Since v_1 is covered both in the first and last iteration of the algorithm, we are pessimistically bounding the number of polygons containing v_1 by $|\mathcal{Q}| = O(n \log n)$. To then prove the lemma it is enough to show that every vertex of P except v_1 is contained in $O(\log n)$ polygons of $\mathcal{Q}$. For that we fix a vertex v_k, with $1 < k \leq n$, and show that v_k appears in the ∂-portion of $O(\log n)$ polygons and v_k appears in the π-portion of $O(\log n)$ polygons of $\mathcal{Q}$.

To bound the number of appearances of v_k on the ∂-portion of a polygon we fix the unique iteration i^* in which v_k is first covered. Since TestCover is used as a predicate in Exponential and Binary search, in iteration i^* it is called $O(\log n)$ times and thus v_k appears in $O(\log n)$ polygons during this iteration. Observe, that in subsequent iterations, when $i > i^*$, vertex v_k is not part of the ∂-portion of any constructed polygon. For an iteration $i < i^*$, let v_{u_i} be the first uncoverable vertex (denoted by v_u in the algorithm) found in iteration i, thus $u_i \leq k$; let q_i be the number of polygons in which

v_k appears on the ∂-portion during this iteration i. Also observe that u_{i-1} is the index of the first vertex of P covered in iteration i. We claim that

$$k - u_i \leq \frac{k - u_{i-1}}{2^{q_i - 1}}, \text{for any } 1 \leq i < i^* \tag{3.1}$$

and defining $u_0 = 1$, implies that $\sum_{i=1}^{i^*-1} q_i \leq \log k \leq \log n$.

For $q_i = 0$, inequality (3.1) holds trivially. Otherwise, since v_k is not covered during this iteration, Exponential Search stops after the first time v_k appears on the ∂-portion of a constructed polygon. This leaves a search interval of size at most $k - u_{i-1}$. During Binary Search, there are exactly $q_i - 1$ search intervals which contain both v_u and v_k. Since the interval size is halved at each step and all search intervals containing both v_{u_i} and v_k have size at least $k - u_i$, inequality (3.1) follows.

So far we have shown that v_k appears on the ∂-portion of $O(\log k)$ polygons in $\mathcal{Q}$ before iteration i^*, $O(\log n)$ times during iteration i^* and does not appear in subsequent iterations. Therefore, all together, v_k appears on ∂-portions of $O(\log n)$ polygons in $\mathcal{Q}$.

To bound the number of appearances of v_k on the π-portion of a polygon, let $\mathcal{Q}_k \subseteq \mathcal{Q}$ be the set of polygons containing v_k on their π-portion but not on the ∂-portion. By Observation 3.2.4 below, any two polygons in $\mathcal{Q}_k$ intersect on their ∂-portion because they both contain v_k on their π-portion. Since by construction the ∂-portion of each polygon Q ends with a vertex, any two polygons in $\mathcal{Q}_k$ have a vertex in common on their ∂-portion. This is true because the ∂-portion of those polygons are subsequences of $(v_1, ..., v_n, v_1)$ and it is easy to see that there is a vertex $v_{k'}$ that belongs to the ∂-portion of all $Q \in \mathcal{Q}_k$. Since $v_{k'}$ appears on ∂-portions of $O(\log n)$ polygons, $|\mathcal{Q}_k| = O(\log n)$.

□

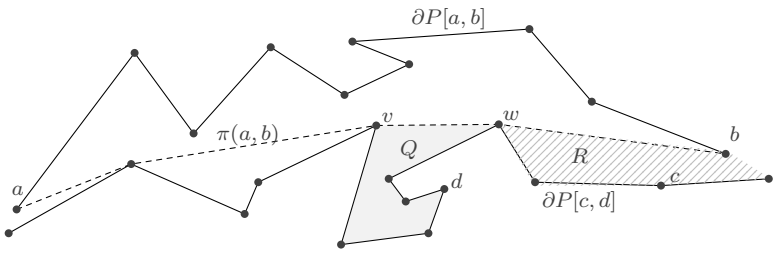

FIGURE 3.3: Illustration of the proof of Observation 3.2.4.

Observation 3.2.4. *For a, b, c, d four distinct points on ∂P, if $\pi(a,b) \cap \pi(c,d)$ contains a polygon vertex not contained in $\partial P[a,b] \cup \partial P[c,d]$, then $\partial P[a,b] \cap \partial P[c,d] \neq \emptyset$.*

Proof. Let v be a vertex contained in $\pi(a,b) \cap \pi(c,d)$ not in $\partial P[a,b] \cup \partial P[c,d]$ and assume for contradiction that $\partial P[a,b]$ and $\partial P[c,d]$ are disjoint. Then either $\partial P[c,d] \subset \partial P[a,v]$ or $\partial P[c,d] \subset \partial P[v,b]$. W.l.o.g. assume $\partial P[c,d] \subset \partial P[v,b]$. For w the successor vertex of v in $\pi(a,b)$, let Q be the simple polygon bounded by $\overline{vw} \circ \partial P[w,v]$. If both c, d are contained in Q, meaning $w \notin \partial P[c,d]$, since v is a convex vertex in Q, it holds that $v \notin \pi(c,d)$, a contradiction. Otherwise, let R be the geodesic convex set bounded by $\pi(w,b) \circ \partial P[b,w]$. If both c, d are contained in R, then by geodesic convexity, $\pi(c,d) \subseteq R$ and thus $v \notin \pi(c,d)$. Otherwise $d \in Q$ and $c \in R$ as shown in Fig. 3.3. Since in that case $Q \cap R = \{w\}$ and $R \cup Q$ is again a geodesic convex set, $\pi(c,d) = \pi(c,w) \circ \pi(w,d)$. Again, since v is a convex vertex in Q, $v \notin \pi(w,d)$, and thus v not in $\pi(c,d)$, a contradiction.

□

3.3 Approximation Ratio

Let OPT denote a set of geodesic unit disks optimally covering ∂P. In order to prove the 2-approximation we prove the existence of a *coloring* for ∂P using $|OPT|$ distinct colors and introducing at most $\max\{2|OPT| - 2, 1\}$ monochromatic boundary portions.

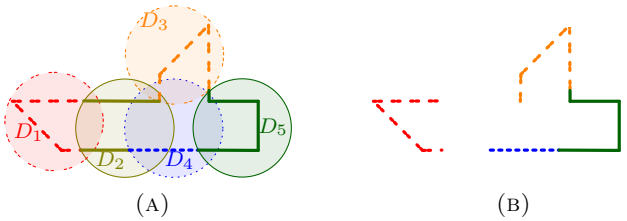

FIGURE 3.4: (a) A crossing disk-free coloring of ∂P. (b) The two pockets induced by color 2.

We then show that CONTIGUOUSGREEDY uses at most one disk per monochromatic boundary portion (plus possibly one additional disk for the unique monochromatic boundary portion containing v_1), which implies the 2-approximation factor of CONTIGUOUSGREEDY.

A *coloring* of ∂P is a function $\gamma : \partial P \to \mathbb{N}$. The number of colors used by γ is defined as the cardinality of the image of γ. A *block* is a connected component of ∂P colored with a single color. We let ∂P_i denote the subset of the polygon boundary colored with color i and we call each connected component of $\partial P \setminus \partial P_i$ a *pocket* of ∂P induced by color i (see Fig. 3.4(b)).

A coloring of ∂P is called *crossing-free* if for any two distinct colors i, j, it holds that ∂P_j is contained in a single pocket induced by color i.

For a collection $\mathcal{D} = \{D_1, \ldots, D_k\}$ of disks covering ∂P, a *disk-coloring* of ∂P w.r.t. $\mathcal{D}$ is a function $\gamma_{\mathcal{D}} : \partial P \to \{1, \ldots, k\}$, such that $\gamma(x) = i \Rightarrow x \in D_i$, i.e., a point on ∂P can only be colored with one of the indices of the disks covering it (see Fig. 3.4(a)).

Definition 3.8. For a coloring γ, two of its colors r and b *cross* each other, if there are two pockets induced by color r containing blocks of color b.

Observe that if two colors r and b cross each other, there are at least two blocks B_r^1, B_r^2 of color r and two blocks B_b^1, B_b^2 of color b such that sequence of blocks $B_r^1, B_b^1, B_r^2, B_b^2$

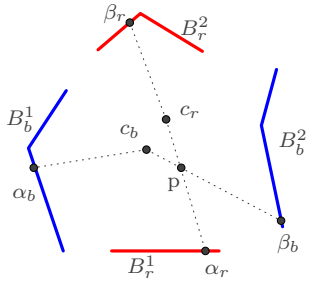

FIGURE 3.5: Illustration of the four alternating blocks $B_r^1, B_b^1, B_r^2, B_b^2$ and the corresponding points α_r, β_r and α_b, β_b; the disk centers c_r and c_b, as well as the intersection point p of $\pi(c_r, \alpha_r)$ and $\pi(c_b, \beta_b)$.

occurs in clockwise order on ∂P as shown in Fig. 3.5.

Lemma 3.9. *In any disk-coloring, if two colors r and b cross each other, one of the following holds: 1) There exists a pocket induced by color r which contains blocks of color b and all these blocks can be re-colored with color r, s.t. the resulting coloring is still a disk-coloring. 2) There exists a pocket induced by color b which contains blocks of color r and all these blocks can be re-colored with color b, s.t. the resulting coloring is still a disk-coloring.*

Proof. Suppose this is not possible. Since neither B_r^1 nor B_r^2 can be colored with b, there are points $\alpha_r \in B_r^1$ and $\beta_r \in B_r^2$ which lie outside of disk D_b. If we denote the center of D_b by c_b, it therefore holds that $d(c_b, \alpha_r) > 1$ and $d(c_b, \beta_r) > 1$ (see Fig. 3.5). Analogously, there are two points $\alpha_b \in B_b^1$ and $\beta_b \in B_b^2$, s.t. α_b and β_b can not be colored with color r. This again implies that both points lie outside of disk D_r centered at c_r and thus $d(c_r, \alpha_b) > 1$ and $d(c_r, \beta_b) > 1$.

Lemma 3.10. *For any collection of disks covering ∂P, there exists a crossing free disk-coloring of ∂P.*

Proof. Consider the four paths $\pi(c_r, \alpha_r)$, $\pi(c_r, \beta_r)$, $\pi(c_b, \alpha_b)$ and $\pi(c_b, \beta_b)$. Due to the alternating arrangement of the four blocks $B_r^1, B_b^1, B_r^2, B_b^2$ – and therefore of $\alpha_r, \alpha_b, \beta_r, \beta_b$,

on the polygon boundary, one of the paths from c_r must intersect with one of the paths from c_b. Assume w.l.o.g. that $\pi(c_r, \alpha_r)$ intersects $\pi(c_b, \beta_b)$ and let p be an intersection point. Again, w.l.o.g., assume that $d(c_r, p) \leq d(c_b, p)$. Then, by the triangle inequality $d(c_r, \beta_b) \leq d(c_b, \beta_b) \leq 1$ contradicting our assumption that $d(c_r, \beta_b) > 1$.

□

We are now going to prove Lemma 3.10, which states that for any collection of disks covering ∂P, there exists a crossing free disk-coloring of ∂P.

For a given disk-coloring w.r.t. a collection of disks $\mathcal{D}$, we let l_{ij} be the number of pockets induced by color i which contain blocks of color j; observe that $l_{ij} = l_{ji}$. We refer to

$$\sum_{1 \leq i < j \leq |\mathcal{D}|} (l_{ij} - 1)$$

as the *crossing number* of the disk-coloring. By definition it holds that the crossing number of a coloring is zero if and only if the coloring is crossing free. We now let $\gamma = \gamma_{\mathcal{D}}$ be a disk-coloring of ∂P w.r.t. disks $\mathcal{D}$, having minimum crossing number (over all disk-colorings w.r.t. $\mathcal{D}$). Assume for contradiction that the crossing number of γ is not zero and let r and b be two colors of γ which cross each other, i.e., $l_{rb} = l_{br} \geq 2$. Then, w.l.o.g., according to Lemma 3.9, there exists a pocket $\mathcal{P}_b$ induced by color b, in which all blocks of color r can be colored with b and the coloring remains a valid disk-coloring w.r.t. $\mathcal{D}$. We refer to the resulting disk-coloring as $\hat{\gamma}$ and by $\hat{l}_{ij}$ to the number of pockets induced by color i of $\hat{\gamma}$ which contain blocks of color j (again in $\hat{\gamma}$). Lastly we denote by $\hat{\mathcal{P}}_r$ the pocket induced by color r in $\hat{\gamma}$, fully containing $\mathcal{P}_b$ as shown in Figure 3.6.

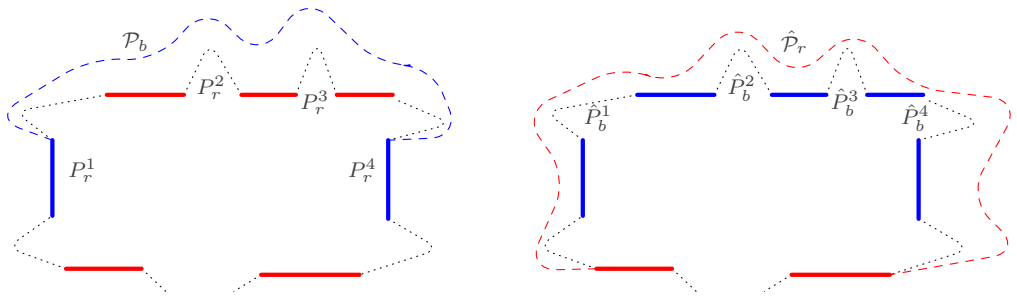

(A) (B)

FIGURE 3.6: (a) Illustration of $\mathcal{P}_b$ in the disk-coloring γ. (b) Illustration of $\hat{\mathcal{P}}_r$ in the disk-coloring $\hat{\gamma}$.

We are going to show that $\hat{\gamma}$ has a smaller crossing number than γ, thus contradicting the assumption that γ is the disk-coloring with minimum crossing number. For this, we extend the definition of l_{ij} to parts of the polygon boundary: for a contiguous subset ∂Q of ∂P, we denote by $l_{ij}[\partial Q]$ the number of pockets induced by color i which are fully contained in ∂Q and which contain blocks of color j.

For the rest of the proof, let k be an arbitrary color of γ (and thus also of $\hat{\gamma}$). Since every pocket induced by color r in γ (and in $\hat{\gamma}$) is either contained in $\hat{\mathcal{P}}_r$ or in $\partial P \setminus \hat{\mathcal{P}}_r$, it holds that

$$l_{rk} = l_{rk}[\hat{\mathcal{P}}_r] + l_{rk}[\partial P \setminus \hat{\mathcal{P}}_r] \quad \text{and} \quad \hat{l}_{rk} = \hat{l}_{rk}[\hat{\mathcal{P}}_r] + \hat{l}_{rk}[\partial P \setminus \hat{\mathcal{P}}_r]. \tag{3.2}$$

Similarly, it holds that

$$l_{bk} = l_{bk}[\mathcal{P}_b] + l_{bk}[\partial P \setminus \mathcal{P}_b] \quad \text{and} \quad \hat{l}_{bk} = \hat{l}_{rb}[\mathcal{P}_b] + \hat{l}_{bk}[\partial P \setminus \mathcal{P}_b]. \tag{3.3}$$

Furthermore, since $\hat{\gamma}$ does not differ from γ in $\partial P \setminus \mathcal{P}_b$, it holds that

$$\hat{l}_{bk}[\partial P \setminus \mathcal{P}_b] = l_{bk}[\partial P \setminus \mathcal{P}_b] \tag{3.4}$$

and analogously, since $(\partial P \setminus \hat{\mathcal{P}}_r) \subseteq (\partial P \setminus \mathcal{P}_b)$, it holds that

$$\hat{l}_{rk}[\partial P \setminus \hat{\mathcal{P}}_r] = l_{rk}[\partial P \setminus \hat{\mathcal{P}}_r]. \tag{3.5}$$

Next, we are going to show that

$$\hat{l}_{rk}[\hat{\mathcal{P}}_r] + \hat{l}_{bk}[\mathcal{P}_b] \leq l_{rk}[\hat{\mathcal{P}}_r] + l_{bk}[\mathcal{P}_b]. \tag{3.6}$$

We are going to prove this by distinguishing two cases: 1) $\hat{l}_{rk}[\hat{\mathcal{P}}_r] > l_{bk}[\mathcal{P}_b]$. Since in $\hat{\gamma}$, by definition $\hat{\mathcal{P}}_r$ is a single pocket induced by r, it follows that $\hat{l}_{rk}[\hat{\mathcal{P}}_r] = 1$ and thus $l_{bk}[\mathcal{P}_b] = 0$. Observe that $l_{bk}[\mathcal{P}_b] = 0$ means that no block of color k was present in $\mathcal{P}_b$ in the γ coloring, and this implies that $\hat{l}_{bk}[\mathcal{P}_b] = 0$. Furthermore, since a block of color k appears inside $\hat{\mathcal{P}}_r$ in the coloring $\hat{\gamma}$, a block of color k appeared inside $\hat{\mathcal{P}}_r$ in the coloring γ. Thus it holds that $l_{rk}[\hat{\mathcal{P}}_r] \geq 1$ which together establishes Eq. (3.6). 2) $\hat{l}_{rk}[\hat{\mathcal{P}}_r] \leq l_{bk}[\mathcal{P}_b]$. We only need to show that $\hat{l}_{bk}[\mathcal{P}_b] \leq l_{rk}[\hat{\mathcal{P}}_r]$. To see this, let $P_r^1, \ldots, P_r^t$ be the pockets induced by color r in γ, which are contained in $\hat{\mathcal{P}}_r$ (ordered clockwise). Observe that since in $\hat{\gamma}$ each block in $\mathcal{P}_b$ which was of color r in γ gets colored with color b, there are again exactly t such pockets $\hat{P}_b^1, \ldots, \hat{P}_b^t$ induced by color b in $\hat{\gamma}$ which are fully contained in $\mathcal{P}_b$. Next, observe that for any $1 \leq p \leq t$ it holds that $\hat{P}_b^p \subseteq P_r^p$. Thus if in $\hat{\gamma}$ a block of color k is contained in a pocket $\hat{P}_b^p$ then this block was contained in pocket P_r^p in γ. This indeed implies that $\hat{l}_{bk}[\mathcal{P}_b] \leq l_{rk}[\hat{\mathcal{P}}_r]$ proving Eq. (3.6) for this second case.

Using Eq. (3.2) - (3.6), we obtain

$$\begin{aligned}
\hat{l}_{rk} + \hat{l}_{bk} &\overset{(3.2),(3.3)}{=} \hat{l}_{rk}[\hat{\mathcal{P}}_r] + \hat{l}_{rk}[\partial P \setminus \hat{\mathcal{P}}_r] + \hat{l}_{bk}[\mathcal{P}_b] + \hat{l}_{bk}[\partial P \setminus \mathcal{P}_b] \\
&\overset{(3.4),(3.5)}{=} \hat{l}_{rk}[\hat{\mathcal{P}}_r] + l_{rk}[\partial P \setminus \hat{\mathcal{P}}_r] + \hat{l}_{bk}[\mathcal{P}_b] + l_{bk}[\partial P \setminus \mathcal{P}_b] \\
&\overset{(3.6)}{\leq} l_{rk}[\hat{\mathcal{P}}_r] + l_{rk}[\partial P \setminus \hat{\mathcal{P}}_r] + l_{bk}[\mathcal{P}_b] + l_{bk}[\partial P \setminus \mathcal{P}_b] \\
&\overset{(3.2),(3.3)}{=} l_{rk} + l_{bk}.
\end{aligned}$$

Furthermore, since color k was chosen arbitrarily, it holds that

$$\sum_{\substack{i \in \{1,\ldots,|\mathcal{D}|\} \\ \setminus \{r,b\}}} (\hat{l}_{ri} + \hat{l}_{bi}) \leq \sum_{\substack{i \in \{1,\ldots,|\mathcal{D}|\} \\ \setminus \{r,b\}}} (l_{ri} + l_{bi}).$$

Because we colored all blocks of color r in $\mathcal{P}_b$ by color b, it follows that $\hat{l}_{rb} = l_{rb} - 1$ and it thus indeed holds that

$$\sum_{1 \leq i < j \leq |\mathcal{D}|} \hat{l}_{ij} < \sum_{1 \leq i < j \leq |\mathcal{D}|} l_{ij},$$

contradicting the assumption that γ has the smallest crossing number.

□

Lemma 3.11. *For a crossing-free coloring γ using κ colors, let Π_γ be the set of blocks induced by γ. If $\kappa > 1$ then $|\Pi_\gamma| \leq 2(\kappa - 1)$.*

Proof. We prove the lemma by induction on the number of colors. For $\kappa = 2$, since γ is crossing-free it is easy to see that $|\Pi_\gamma| \leq 2$ and thus the lemma holds. Assuming the lemma holds for $\kappa - 1$ colors, we show it also holds for κ colors. For any color i used by

γ, let $\mathcal{B}_i \subseteq \Pi_\gamma$ be the set of blocks of color i. If for all i, $|\mathcal{B}_i| \leq 1$, the lemma trivially holds. Otherwise fix i to be a color with $|\mathcal{B}_i| \geq 2$ and observe that the number of pockets induced by color i is $|\mathcal{B}_i|$.

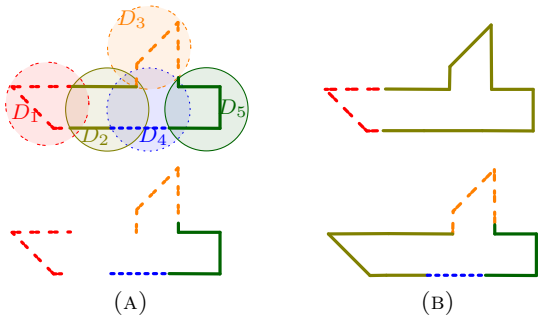

FIGURE 3.7: (a) A disk-coloring example w.r.t. disks $D_1, \ldots, D_5$; in the lower part the two pockets induced by color 2 are shown. (b) shows the two polygon colorings in the induction step for color 2 in the proof of Lemma 3.11.

Let $\mathcal{P}_1, ..., \mathcal{P}_{|\mathcal{B}_i|}$ be the pockets induced by color i. For each such pocket $\mathcal{P}_j$ we create a new coloring γ_j of ∂P, with

$$\gamma_j(x) = \begin{cases} \gamma(x) & \text{if } x \in \mathcal{P}_j \\ i & \text{otherwise,} \end{cases}$$

as illustrated in Fig. 3.7(b).

Since γ is crossing-free it is easy to see that for any pocket $\mathcal{P}_j$, the coloring γ_j is also a crossing-free. Denoting the number of colors of γ_j by κ_j, it holds that $1 < \kappa_j < \kappa$, for all $1 \leq j \leq |\mathcal{B}_i|$. Letting Π_{γ_j} be the set of blocks induced by the coloring γ_j, by induction hypothesis $|\Pi_{\gamma_j}| \leq 2(\kappa_j - 1)$.

Observe that each Π_{γ_j} contains exactly one block not in Π_γ. Also, the blocks in $\mathcal{B}_i$ are exactly the blocks not appearing in any of the Π_{γ_j}. Therefore, since the number of pockets induced by color i equals the number of blocks in $\mathcal{B}_i$, it holds that $|\Pi_\gamma| = \sum_{j=1}^{|\mathcal{B}_i|} |\Pi_{\gamma_j}|$. Thus we obtain

$$|\Pi_\gamma| = \sum_{j=1}^{|\mathcal{B}_i|} |\Pi_{\gamma_j}| \leq 2 \sum_{j=1}^{|\mathcal{B}_i|} (\kappa_j - 1) \tag{3.7}$$

and because each of the colorings $\gamma_1, \ldots, \gamma_{|\mathcal{B}_i|}$ is crossing-free,

$$1 + \sum_{j=1}^{|\mathcal{B}_i|} (\kappa_j - 1) = \kappa, \tag{3.8}$$

where $(\kappa_j - 1)$ is the number of colors the coloring γ (and also γ_j) uses for the pocket $\mathcal{P}_j$. The addition of 1 on the left hand side of (3) attributes for color i, which was not counted in any of the pockets. Plugging (3) into (2), the lemma follows.

□

Theorem 3.12. *The number of disk centers placed by* ContiguousGreedy *is at most* $2|OPT| - 1$.

Proof. If $|OPT| = 1$ then, by its greedy nature, ContiguousGreedy also uses only one disk. If $|OPT| > 1$, let γ_{OPT} be a crossing free disk-coloring of ∂P w.r.t. OPT, whose existence is guaranteed by Lemma 3.10. We let $(B_1, B_2, ..., B_m)$ be the collection of blocks induced by γ_{OPT} ordered as they appear on ∂P in clockwise order, with B_1 the block containing v_1. We split B_1 at v_1 into two blocks B_l and B_r, with B_l being the portion of B_1 counterclockwise from v_1, and $B_r = B_1 \setminus B_l$.

Now observe that by the greedy nature, every disk D computed by ContiguousGreedy extends Γ so that $\Gamma \cup D$ fully covers at least one new block in the sequence $(B_r, B_2, ..., B_m, B_l)$. Therefore, after computing at most $m + 1$ disks, $\Gamma = \partial P$. By Lemma 3.11, it holds that $m \leq 2(|OPT| - 1)$ and the theorem follows.

□

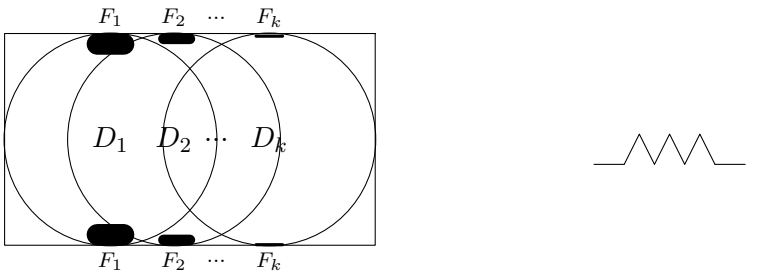

(A) (B)

FIGURE 3.8: (a) Illustration of the polygon containing foldings $F_1, \ldots, F_k$ on the boundary. A global greedy algorithm starts covering the two F_1 foldings on opposite sides by the disk D_1, the two F_2 foldings by a disk D_2 and so on, while OPT still only uses a constant number of disks to cover ∂P. (b) Illustration of a folding.

3.3.1 Tightness of Analysis

The analysis for the 2-approximation ratio of CONTIGUOUSGREEDY is almost tight, even for convex polygons, as can be seen by a rectangle of length n and height $\epsilon > 0$. It can be covered with $n/(2\sqrt{1-\epsilon^2/4})$ many geodesic unit disks (by centering them on the median line at height $\epsilon/2$). On the other hand, CONTIGUOUSGREEDY centers disks in steps of 2 on the boundary, thus after finishing one side of the rectangle, each disk introduced a small uncovered hole on the other side. CONTIGUOUSGREEDY covers those holes by placing another $n/2$ disks contiguously on the other side of the polygon, resulting in a total of n disks needed.

Another natural greedy approach is to cover the largest amount of uncovered boundary at each step. This algorithm results in an approximation ratio of $\Omega(\log n)$, i.e., it is unbounded with respect to $|OPT|$. An example where this greedy rule performs badly is illustrated in Fig. 3.8(a). The parts of the boundary denoted by $F_1, \ldots, F_k$ are dense *foldings* as shown in Fig. 3.8(b) where the boundary length of F_1 is twice that of F_2, four times that of F_3, and so on. The global greedy algorithm first covers the two F_1 sections on opposite sides of the boundary (illustrated by D_1 in Fig. 3.8(a)), then the two F_2 sections continuing in this way until the two F_k sections are covered, thereby

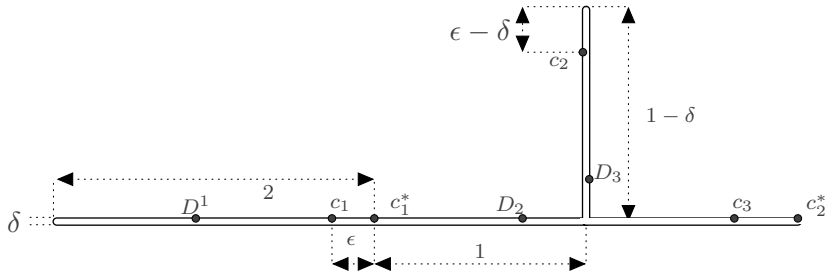

FIGURE 3.9: Illustration of the δ-thin polygon where an ϵ-approximate contiguous extension algorithm results in an approximation ratio larger than 2.

having used k disks to cover the foldings, (plus some constant number of disks to cover the rest of ∂P). Notice that when the height of the polygon is arbitrary close to 2, the number of foldings can be made arbitrary large, while OPT only uses a constant number of disks to cover ∂P.

It is worth noting that it is crucial that ContiguousGreedy *exactly* computes the maximum extension of the contiguous boundary covered by a single geodesic unit disk in each iteration. Only approximately (even with ϵ precision) extending the contiguously covered part by a single geodesic unit disk results in an approximation factor of at least 4 (instead of 2). To see this, we refer to Fig. 3.9, where c_1 is the endpoint of the ϵ-approximate contiguous greedy extension in the first step and c_1^* is the corresponding exact endpoint (obtained from ContiguousGreedy). The ϵ-approximate algorithm continues by centering a disk at D_2 which covers the boundary from c_1 up to c_2. At this point, an exact extension could cover the boundary from c_2 up to c_2^*. However, the approximate algorithm may only cover up to c_3, by, for example, centering the third disk at D_3. ContiguousGreedy covers up to c_2^* using only two disks. Copying the polygon-section between c_1^* and c_2^*, shows that an ϵ-approximate algorithm performs at least twice as bad as ContiguousGreedy.

3.4 Covering Large Perimeters

In this section we show that if the polygon perimeter L is significantly larger than n, i.e., $L \geq n^{1+\delta}$, with $\delta > 0$, a simple linear time algorithm achieves an approximation ratio which goes to one as L/n goes to infinity. For this, we decompose ∂P into long and short portions, based on the length of the corresponding *medial axis*. The medial axis is the set of points in P which have more than one closest point on ∂P. It forms a tree whose edges are either line segments or parabolic arcs and it can be computed in linear time [27]. For a line segment edge, the closest points to the boundary are a subset of two polygon edges; for a parabolic edge, the closest boundary points are a polygon vertex and a subset of a polygon edge. The idea of the algorithm is to identify long edges of the medial axis (of length at least some constant $c > 2$), and to cover the corresponding polygon boundary section (referred to as *corridors*) almost optimally using only a constant number of disks more than OPT uses to cover the corridor. It is easy to see that each corridor stemming from a parabolic arc can be covered with at most two more disks than OPT uses, by centering disks at distance 2 from each other on the corresponding polygon boundary segment and one disk on the corresponding polygon vertex. Each corridor consisting of a pair of polygon boundary segments can be covered by greedily centering disks on the corresponding medial axis as long as each disk contains corridor portions of length more than two; if the length becomes two or less, greedily center the disks on corridor segments in steps of two. Observe that also in this case, the number of disks needed to cover a corridor is at most two more than OPT uses and their centers can be computed in time linear in their number. This holds since there is at most one point where the covering changes from centering disks on the medial axis to centering disks on ∂P. The rest of the polygon, i.e., the short portions, can be covered greedily by centering $O(n)$ disks on ∂P.

Let $\mathcal{D}$ be the set of all disks placed by the algorithm, $\mathcal{D}_L \subseteq \mathcal{D}$ the disks covering the corridors and $\mathcal{D}_S \subseteq \mathcal{D}$ the $O(n)$ disks covering the short portion of ∂P. Since the number of edges in the medial axis is $O(n)$ (see [27]) and the procedure for covering the long corridors uses at most two more disks than OPT for each corridor, $|\mathcal{D}_L| \leq |OPT| + O(n)$. It therefore holds that $|\mathcal{D}| = |\mathcal{D}_L| + |\mathcal{D}_S| \leq |OPT| + O(n)$. It is easy to see that the disks of OPT which contain a polygon vertex cover at most an $O(n)$ portion of ∂P implying that $|OPT| = \Omega(L)$. Therefore, the approximation ratio can be written as

$$\frac{|\mathcal{D}|}{|OPT|} \leq 1 + \frac{O(n)}{|OPT|} = 1 + \frac{O(n)}{\Omega(L)} = 1 + O\left(\frac{n}{L}\right) = 1 + o\left(1\right).$$

Chapter 4

Combinatorial Separation Results

4.1 Separating Convex Points with Lines

4.1.1 Introduction

Separating objects in the plane using lines in such a way that each object is contained in its own cell was studied in the computational geometry community because of its potential applications to manufacturing, constructive solid geometry and statistical classification (see [45]). Several papers ([45],[22],[70]) focus on algorithmic and complexity-theoretical aspects of separating points in the plane with lines. On the other hand, the only combinatorial aspect studied is the minimum number of hyperplanes required to separate points, in both the general and convex cases (see [44]).

We would thus like to initiate combinatorial studies in this setting by considering the number of ways to separate points in convex position with lines. As we will see, this provides a connection between noncrossing trees (see for example [33],[72],[76],[77]) and linked diagrams studied in [11],[71],[91],[92].

Since there are generally uncountably many lines introducing the same separation, we define the following notion of linear bipartitions. For a set S of points in the plane, a *linear bipartition* of S is a set $\{U, S \setminus U\}$ consisting of two disjoint nonempty subsets of S which respectively are fully contained in the two open half-planes bounded by some line.

Definition 1. A set $\mathcal{P}$ of linear bipartitions is called a *linear separating family* for S if for every distinct elements $p, q \in S$ there is a linear bipartition $\{U, S \setminus U\}$ in $\mathcal{P}$ such that $p \in U$ and $q \in S \setminus U$. Furthermore, $\mathcal{P}$ is called *minimal*, if no proper subfamily of $\mathcal{P}$ separates S.

We say that a linear separating family for S is of *maximum* (*minimum*) size if its cardinality is the largest (smallest) among all linear separating families for S.

Example 1. Given a set $S = \{p_1, \ldots, p_6\} \subseteq \mathbb{R}^2$ of points in convex position, the set $\mathcal{P} = \{P_1, \ldots, P_4\}$ defined below and illustrated in Figure 4.1 (left) is an example of a minimal linear separating family for S.

$$P_1 = \{\{p_1, p_2, p_3, p_4\}, \{p_5, p_6\}\}, \qquad P_2 = \{\{p_1, p_2, p_3, p_6\}, \{p_4, p_5\}\}$$

$$P_3 = \{\{p_1, p_3, p_4, p_5, p_6\}, \{p_2\}\}, \qquad P_4 = \{\{p_1, p_4, p_5, p_6\}, \{p_2, p_3\}\}$$

4.1.2 Two Bijections

In this section we introduce a graph representation for linear separating families from which we obtain two bijections.

We call elements in $[n] = \{1, \ldots, n\}$ *vertices* and unordered pairs of vertices are called *edges*. A collection H of edges on $[n]$ is called an *edge cover*, if every vertex in $[n]$ is incident to some edge in H. Furthermore, two edges $\{a, b\}$ and $\{c, d\}$ are called *crossing* if $a < c < b < d$ or $c < a < d < b$.

Definition 2. A collection of edges H is said to be *crossingly connected* if every two edges e, e' in H can be connected by a sequence $e, e_1, \ldots, e_k, e'$ of edges in H such that every two consecutive edges in the sequence are either indicent or crossing.

This notion extends the usual notion of connectivity of edges on $[n]$. Every connected edge set on $[n]$ in the usual sense is crossingly connected, but the converse is clearly not true in general.

Theorem 4.1.1. *There exists a bijection between the set of all linear separating families for a set of n points in convex position in the plane and the set of all crossingly connected edge covers on* $[n]$.

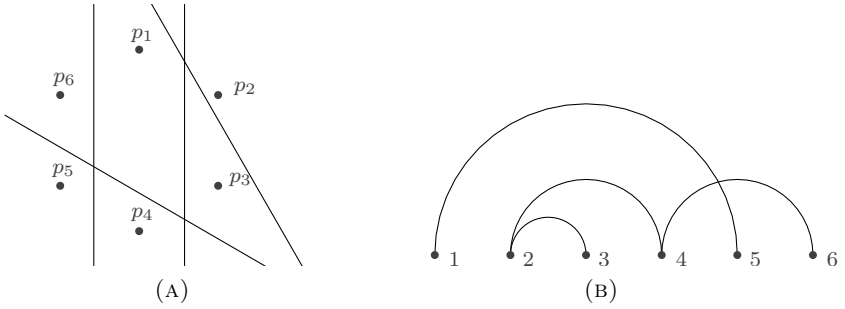

FIGURE 4.1: Illustration of the minimal linear separating family of Example 1 and its corresponding edge cover.

Proof. Let $S = \{p_1, \ldots, p_n\}$ be a set of n points in convex position in the plane, and suppose that $p_1, \ldots, p_n$ are arranged in clockwise order. For each component of a linear bipartition P of S, we select the first point when traversing it in clockwise order. If p_i and p_j are two such points, we associate P with the set $\{i, j\}$ of indices. As an example,

$\{\{p_1, p_2, p_3, p_6\}, \{p_4, p_5\}\}$ gets associated with the indices $\{4, 6\}$ as shown in Figure 4.1 (right). This correspondence clearly establishes a bijection between the set of all linear bipartitions of S and the set of all edges on $[n]$.

Let $\mathcal{C}$ be a linear separating family for S and let $H_{\mathcal{C}}$ be the edge set obtained by the above correspondence. Then $H_{\mathcal{C}}$ is an edge cover on $[n]$, since if a vertex $i \in [n]$ is not incident to any edge in $H_{\mathcal{C}}$, then no bipartitions in $\mathcal{C}$ separate p_i from its counterclockwise neighbor.

In order to show that $H_{\mathcal{C}}$ is crossingly connected, assume the contrary and let I, J be two of its crossingly connected components. We define U_I and U_J to be the sets of all points in S whose indices are incident to edges in I and J, respectively.

Since no edge from I is adjacent to any edge in J, the two point sets U_I and U_J are disjoint. Let $\underline{p_I}$ be the first point in U_I such that no point of U_J appears when traversing S from $\underline{p_I}$ in clockwise order and let $\overline{p_I}$ be the last such point in U_I; $\underline{p_J}$ and $\overline{p_J}$ are defined analogously (see Figure 4.2).

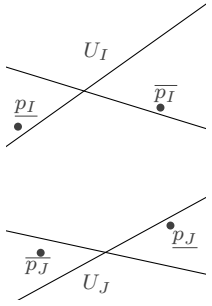

FIGURE 4.2: Illustration of U_I and U_J.

Since we selected I and J arbitrarily, we can assume without loss of generality that $\underline{p_J}$ is the clockwise neighbor of $\overline{p_I}$. Let Q be a bipartition in $\mathcal{C}$ which separates $\overline{p_I}$ from $\overline{p_J}$ and let $\{a, b\}$ be the edge corresponding to Q. It then follows that when traversing S from $\underline{p_J}$ to $\overline{p_J}$ in clockwise order p_a appears but not p_b or vice versa. Thus $\{a, b\}$ has

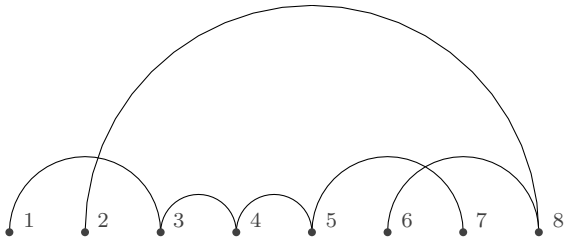

FIGURE 4.3: An example of a crossingly connected cycle formed by the path $(2, 8, 6)$ and the path $(1, 3, 4, 5, 7)$. After removing either edge $\{3, 4\}$ or $\{4, 5\}$ the graph is still a crossingly connected edge cover.

one vertex, say a, in J implying that $\{a, b\}$ belongs to J. On the other hand, since Q separates $\overline{p_I}$ from $\overline{p_J}$, p_b appears after $\overline{p_J}$ in clockwise order. This contradicts that $\overline{p_J}$ is the last point in U_J and therefore we obtained that $H_{\mathcal{C}}$ is a crossingly connected edge cover on $[n]$.

From the argument above, it is straightforward to verify that any crossingly connected edge cover on $[n]$ corresponds to some linear separating family for S. □

We refer to a minimal sequence of crossingly connected edges connecting two vertices as a *crossingly connected path*. Furthermore, we call a set H of edges a *crossingly connected cycle* if for every $e \in H$, the edge set $H - e$ contains exactly one crossingly connected path for all pairs of vertices in H (see Figure 4.3).

Theorem 4.1.2. *There exists a bijection between the set of all minimal linear separating families for n points in convex position in the plane and the set of all crossingly connected edge covers on $[n]$ such that no crossingly connected cycle contains a path of length 3.*

Proof. Let $\mathcal{C}$ be a linear separating family for n points in convex position in the plane and let $H_{\mathcal{C}}$ be the edge set on $[n]$ corresponding to $\mathcal{C}$. Using Theorem 4.1.1 it suffices to show that $\mathcal{C}$ is minimal if and only if $H_{\mathcal{C}}$ has no crossingly connected cycle which contains a path of length 3.

Assume that $H_{\mathcal{C}}$ has a crossingly connected cycle σ which contains a path P of length 3. Observe that if an edge in $H_{\mathcal{C}}$ crosses or is adjacent to the inner edge e of P, then it either crosses or is adjacent to some other edge in σ. Thus, $H_{\mathcal{C}} - e$ is still a crossingly connected edge cover, implying that $\mathcal{C}$ is not minimal.

Conversely, assume that $\mathcal{C}$ is not minimal and let e be an edge in $H_{\mathcal{C}}$ such that $H_{\mathcal{C}} - e$ is still a crossingly connected edge cover on $[n]$. Since the vertices of e are both incident to edges in $H_{\mathcal{C}} - e$, the edge e has to be the inner edge of a path of length 3 contained in a crossingly connected cycle in $H_{\mathcal{C}}$. □

A tree on $[n]$ is called *noncrossing*, if no two edges are crossing.

Corollary 4.1.3. *An edge cover of size k on $[n]$ satisfying the latter condition in Theorem 4.1.2 consists of $n - k$ noncrossing trees.*

4.1.3 Enumeration Results

In this section we enumerate minimal linear separating families for points in convex position in the plane.

Proposition 4.1.4. *The maximum and minimum sizes of a minimal linear separating family of a set of n points in convex position in the plane are $n-1$ and $\lceil n/2 \rceil$, respectively.*

Proof. As shown in [45], for a linear separating family to be minimal it consists of at most $n - 1$ linear bipartitions. Furthermore, it is easy to construct a minimal linear separating family of size $n - 1$.

On the other hand, since any edge cover on n vertices consists of at least $\lceil n/2 \rceil$ edges, this number gives a lower bound for the size of any linear separating family. Furthermore,

$\lceil n/2 \rceil$ linear bipartitions are sufficient for separation, since $H = \{\{i, i + \lfloor n/2 \rfloor\} : 1 \leq i \leq \lceil n/2 \rceil\}$ is a crossingly connected edge cover on $[n]$. □

Thus from Corollary 4.1.3 it follows that any minimal linear separating family of maximum size corresponds to a single noncrossing tree. From [36] the following theorem immediately follows (see also [72]).

Theorem 4.1.5. *The number of noncrossing trees on $[n]$ is equal to the number of ternary trees with $n-1$ internal vertices.*

It is known (see [54]) that there are $\frac{1}{2n-1}\binom{3n-3}{n-1}$ ternary trees having $n-1$ internal vertices. We thus obtain the following result.

Theorem 4.1.6. *The number of minimal linear separating families of maximum size for a convex n-point set in the plane is $\frac{1}{2n-1}\binom{3n-3}{n-1}$.*

To enumerate minimal linear separating families of minimum size, we have to distinguish the cases where n is even and n is odd. For the even case, we define a *matching* as an edge cover where all edges are pairwise non-incident. Then according to Corollary 4.1.3, minimal linear separating families of minimum size correspond to crossingly connected matchings.

For the odd case, we define a *near-matching* to be an edge cover such that all but two edges are pairwise non-incident. Minimal linear separating families of minimum but odd size then correspond to crossingly connected near-matchings.

If we denote by e_k the number of crossingly connected matchings on $[2k]$ then, as proven in [62], the following theorem holds.

Theorem 4.1.7. *The ordinary generating function* $E = \sum_{k\geq 1} e_k x^k = x + x^2 + 4x^3 + 27x^4 + \cdots$ *satisfies the differential equation*

$$E' = \frac{E^2 + E - x}{2xE}.$$

As pointed out in [62] and derived in [80], the recursion

$$e_k = (k-1) \sum_{i=1}^{k-1} e_i \cdot e_{k-i} \tag{4.1}$$

holds for all $k \geq 2$, with $e_1 = 1$.

As for the asymptotic behavior of e_k, it was shown in [92] that

$$\lim_{k\to\infty} e_k/p_k = \mathrm{e}^{-1}, \tag{4.2}$$

where p_k denotes the number of matchings on $[2k]$, which is well known to be $(2k-1)!! \equiv (2k-1)(2k-3)\cdots 3 \cdot 1$.

In order to enumerate minimal linear separating families of minimum size for point sets of odd cardinality, let f_k denote the number of crossingly connected near-matchings on $[2k+1]$.

Theorem 4.1.8. *The ordinary generating function* $F = \sum_{k\geq 1} f_k x^k = 3x + 15x^2 + 126x^3 + 1395x^4 + \cdots$ *is related to* E *by the equation*

$$F = \frac{E^3 + 2E^2 + (1 - 3x - 6x^2)E - x}{2xE}.$$

Proof. We embed the vertices of $[2k+1]$ on a straight line from left to right in increasing order. Let H be a crossingly connected near-matching on $[2k+1]$ and observe that H contains a unique path γ of length 2. For the moment assume that γ contains vertices $1, a, b$ $(a < b)$ with vertex 1 having degree 2 as shown in Figure 4.4. Let c_k denote the number of crossingly connected near-matchings on $[2k+1]$ in which vertex 1 has degree 2. Removing γ decomposes H into a collection $\mathcal{M} = \{M_1, \ldots, M_t\}$ of crossingly connected components with $k_1 + \cdots + k_t = k - 1$, where $k_i = |M_i|$.

Observe that every crossingly connected component M_i of $\mathcal{M}$ either surrounds only a, only b or both a and b, where we say that M_i surrounds a vertex x if $s < x < t$ for some pair of vertices s, t incident to edges in M_i. We partition $\mathcal{M}$ into three (possibly empty) classes $\mathcal{M}_1 = \{M_1, \ldots, M_{t_1}\}$, $\mathcal{M}_2 = \{M_{t_1+1}, \ldots, M_{t_2}\}$ and $\mathcal{M}_3 = \{M_{t_2+1}, \ldots, M_t\}$ corresponding to the three cases above (see Figure 4.4).

Conversely, the near-matching H can be constructed out of a sequence $\overline{\mathcal{M}} = (M_1, \ldots, M_t)$ of smaller matchings, by interpreting each M_i as a matching on $[2k_i]$ which has one of the $2k_i - 1$ inner segments $[j, j+1]$ in the line embedding of $[2k_i]$ marked $(1 \leq j < 2k_i)$. For the last matching M_t we mark an additional segment which possibly coincides with the already marked segment.

We divide $\overline{\mathcal{M}}$ into three, possibly empty, parts $\overline{\mathcal{M}}_1 = (M_1, \ldots, M_{t_1})$, $\overline{\mathcal{M}}_2 = (M_{t_1+1}, \ldots, M_{t_2})$ and $\overline{\mathcal{M}}_3 = (M_{t_2+1}, \ldots, M_t)$. The number of ways for such divisions is $\binom{t+1}{2}$.

A marked segment in $M_i \in \overline{\mathcal{M}}_3$ will be expanded to contain M_{i+1}. In the case $i = t$, where we have two marked segments $[j, j+1]$ and $[l, l+1]$ with $j \leq l$, we expand $[j, j+1]$ to contain M_1 and $[l, l+1]$ to contain M_{t_1+1} and continue the segment expansion process in both $\overline{\mathcal{M}}_1$ and $\overline{\mathcal{M}}_2$ until M_{t_1} and M_{t_2} are expanded to contain the vertex a and b respectively.

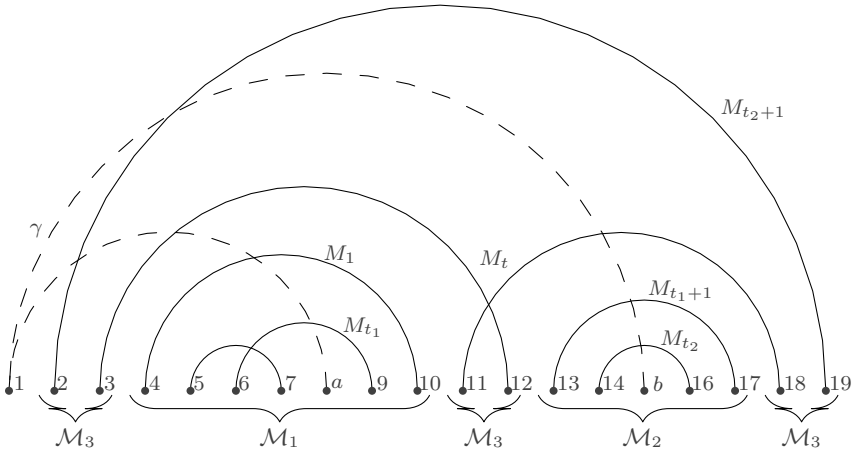

FIGURE 4.4: Illustration of the decomposition of a crossingly connected near-matching into $\mathcal{M}_1$, $\mathcal{M}_2$ and $\mathcal{M}_3$. In this example, M_1, M_{t_1+1}, M_{t_2} and M_{t_2+1} all have segment $[1, 2]$ marked, M_{t_1} has segment $[3, 4]$ marked and M_t has segments $[1, 2]$ and $[3, 4]$ marked.

Therefore, for $k > 1$, c_k can be expressed as

$$c_k = \sum_{t>0} \sum_{k_1+\cdots+k_t=k-1} g_{k_1} \cdots g_{k_t} \left\{ \binom{t+1}{1} + \binom{t+1}{2} k_t \right\},$$

where $g_k = (2k-1)e_k$. The expression in curly brackets accounts for the partitioning of $\overline{\mathcal{M}}$ in both the cases when $\mathcal{M}_3$ is empty and when there are two marked segments in $M_t \in \overline{\mathcal{M}}_3$.

Letting $G = \sum_{k\geq 1} g_k x^k$ denote the generating function of g_k, we can write the generating function $C = \sum_{k\geq 1} c_k x^k$ of c_k as

$$\begin{aligned} C &= \sum_{t>0} xG^{t-1}\left\{\binom{t+1}{2}xG' + (t+1)G\right\} \\ &= \frac{x^2G'}{(1-G)^3} + x\left\{\sum_{t\geq 0}(t+1)G^t - 1\right\} \\ &= \frac{x^2G'}{(1-G)^3} + \frac{x}{(1-G)^2} - x. \end{aligned}$$

Combining the fact that $G = 2xE' - E$ with $E' = \frac{E^2+E-x}{2xE}$ from Theorem 4.1.7, yields $\frac{E}{x} = \frac{1}{1-G}$. Therefore, C can be expressed as

$$\begin{aligned} C &= \frac{E^3G'}{x} + \frac{E^2}{x} - x \\ &= \frac{E^3(E' + 2xE'')}{x} + \frac{E^2}{x} - x \\ xC &= E^3E' + 2xE^3E'' + E^2 - x^2. \end{aligned}$$

where for the second to last equation the fact that $G' = E' + 2xE''$ is used.

Furthermore, using

$$E'' = -\frac{E^4 + E^3 - xE + x^2}{4x^2E^3}$$

yields

$$2xC = E^2 + E - 2x^2 - x. \tag{4.3}$$

Since c_k denotes the number of crossingly connected near-matchings on $[2k+1]$ in which vertex 1 has degree 2, one can obtain f_k by applying a cyclic permutation on each near-matching enumerated by c_k. Since there are $2k+1$ cyclic permutations on $[2k+1]$ we conclude that $f_k = (2k+1)c_k$. In terms of generating functions this yields

$$F = 2xC' + C$$

and therefore

$$F = \frac{E^3 + 2E^2 + (1 - 3x - 6x^2)E - x}{2xE}.$$

□

Combining equations (4.1) and (4.3) one obtains the following corollary.

Corollary 4.1.9. *The number of crossingly connected near-matchings on* $[2k+1]$ *is*

$$f_k = e_{k+1}\frac{(2k+1)(k+1)}{2k}$$

for all $k \geq 1$.

Theorem 4.1.10. *The number* s_n *of minimal linear separating families of minimum size for* n *points in convex position in the plane is*

$$s_n = \begin{cases} (k-1)\sum_{i=1}^{k-1} s_{2i}s_{2(k-i)} & \text{if } n = 2k \\ s_{n+1}\frac{(2k+1)(k+1)}{2k} & \text{if } n = 2k+1 \end{cases}$$

for all $n \geq 3$, *with* $s_2 = 1$.

Denoting by q_k the number of near-matchings on $[2k+1]$ the following proposition holds.

Proposition 4.1.11.

$$\lim_{k\to\infty} f_k/q_k = \mathrm{e}^{-1}$$

Proof. From equation (4.2) it holds that

$$\lim_{k\to\infty} e_{k+1}\frac{2^{k+1}(k+1)!}{(2k+2)!} = \mathrm{e}^{-1}$$

and it is easy to show that $q_k = (2k+1)!/(2^k(k-1)!)$.

Thus it follows that

$$\lim_{k\to\infty} f_k/q_k = \lim_{k\to\infty} e_{k+1}\frac{2^{k+1}(k+1)!}{(2k+2)!}\frac{(2k+2)(2k+1)}{4k^2} = \mathrm{e}^{-1}.$$

□

4.2 Separating arbitrary Sets

In this section we generalize the enumeration results from the previous section to arbitrary sets.

4.2.1 Introduction

A *bipartition* of a set S is $\{S\}$ or an unordered pair $\{U, V\}$ of nonempty subsets of S such that $U \cap V = \emptyset$ and $U \cup V = S$. Note that we allow $\{S\}$ as a bipartition, because it corresponds to the case where the ground set S is divided into S and $\emptyset$. A collection of bipartitions of S is a *separating family* for S if every two elements in S are separated by some bipartition in the collection , that is, they are contained in different components of some bipartition. A separating family for S is *minimal* if no proper subfamily is a separating family for S.

Example 2. Let $S = \{1, 2, 3, 4\}$ and let P_1, P_2, Q_1, Q_2, Q_3 be the bipartitions given as

$$P_1 = \{\{1, 2\}, \{3, 4\}\}, \qquad Q_1 = \{\{1\}, \{2, 3, 4\}\},$$
$$P_2 = \{\{1, 3\}, \{2, 4\}\}, \qquad Q_2 = \{\{1, 2\}, \{3, 4\}\},$$
$$Q_3 = \{\{1, 2, 3\}, \{4\}\}.$$

The family of bipartitions $\{P_1, P_2\}$ is a minimal separating family of minimum size for S, while $\{Q_1, Q_2, Q_3\}$ is a minimal separating family of maximum size. Here the size of a separating family denotes its cardinality.

The concept of separating families appears in the following search problem. Suppose that we are given a finite set S and a collection $\{P_1, \ldots, P_m\}$ of bipartitions of S. For an unknown element x in S, we choose a bipartition P_i and we are allowed to ask which component of P_i contains x, thereby narrowing down the range containing x. The goal is to locate the unknown element x by asking a series of such questions. One can easily observe that for every element in S there exists a series of questions which leads to finding it if and only if $\{P_1, \ldots, P_m\}$ is a separating family for S. Rényi [81] initiated the

study of the search problem described above, although he didn't employ bipartitions but subsets of S as questions. Since then, many authors have studied combinatorial problems related to finding the minimum size of a separating family under various constraints (see [3] [58] [59] for a survey). In this chapter we deal with bipartitions since we do not want to distinguish subsets of S from their complements.

Motivation and Contribution: As mentioned above, separating families can be considered as sets of questions which allow to locate an arbitrary unknown element in the ground set. It is natural to ask how many such sets of questions exist for a given ground set. Since one is often interested in inclusion-wise minimal sets, we first concentrate on separating families which are minimal and enumerate those of maximum size. This is done by obtaining a bijection from the set of minimal separating families of maximum size for a set S to the set of spanning trees on S. This result partially answers our enumeration question in the minimal case, but the intermediate cases are left open. In a next step we enumerate separating families of arbitrary size, which need not be minimal. This result includes minimal separating families of minimum size as a special case.

The remainder of thsi chapter is organized as follows. In Section 4.2.2 we enumerate all minimal separating families of maximum size for a finite set. In Section 4.2.3 we extend this analysis and enumerate all separating families of arbitrary size, which need not be minimal.

4.2.2 Minimal separating families of maximum size and spanning trees

In this section we enumerate all minimal separating families of maximum size for a finite set S by obtaining a bijection onto the set of all spanning trees on S.

Proposition 4.2.1. *The maximum size of a minimal separating family for an n-element set is $n-1$.*

Proof. Let S be an n-element set and let a be a fixed element in S. It is easy to see that the collection of all bipartitions of the form $\{\{x\}, S \setminus \{x\}\}$ for any x other than a is a minimal separating family of size $n-1$, for $n > 2$. To see that the maximum size is at most $n-1$, let $\mathcal{P}$ be any minimal separating family. Let us remove members from $\mathcal{P}$ one by one, as shown below, where Q_1, Q_2, Q_3 are the bipartitions given in Example 2.

$$\{1\},\{2\},\{3\},\{4\} \overset{Q_1}{\to} \{1,2\},\{3\},\{4\} \overset{Q_2}{\to} \{1,2,3\},\{4\} \overset{Q_3}{\to} \{1,2,3,4\}$$

The first step shows that $\{1\}$ and $\{2\}$ are merged by removing Q_1, because neither Q_2 nor Q_3 separate them. From the minimality of $\mathcal{P}$ it follows that in each step some isolated components get merged; in other words, the number of isolated components decreases by at least one. Since there are n isolated components in the initial state, the size of $\mathcal{P}$ is at most $n-1$.

□

Definition 3. Let S be a nonempty finite set. For each minimal separating family $\mathcal{P}$ for S we associate a graph denoted by $\Phi_S(\mathcal{P})$ such that the vertices of the graph correspond to the elements in S and two vertices are adjacent if the corresponding elements in S are separated by exactly one bipartition in $\mathcal{P}$.

The definition above induces a correspondence between the bipartitions in $\mathcal{P}$ and the edges in $\Phi_S(\mathcal{P})$. One can easily observe that each edge of $\Phi_S(\mathcal{P})$ corresponds to a unique bipartition in $\mathcal{P}$, while each bipartition may correspond to multiple edges (see the left graph in Figure 4.5).

FIGURE 4.5: The graphs induced by the minimal separating families of the bipartitions $\{P_1, P_2\}$ and $\{Q_1, Q_2, Q_3\}$ given in Example 2.

Theorem 4.2.2. *Let S be a nonempty finite set. The mapping $\Phi_S \colon \mathcal{P} \mapsto \Phi_S(\mathcal{P})$ is a bijection from the set of all minimal separating families of maximum size for S to the set of all spanning trees on S.*

Proof. Proposition 4.2.1 states that the maximum size of a minimal separating family for an n-element set S is $n-1$. Let G be a maximal subgraph of $\Phi_S(\mathcal{P})$ such that the vertex set of G is S and no two edges in G are separated by the same bipartition in $\mathcal{P}$. Here an edge is *separated* by a bipartition if the end-vertices of the edge are separated by this bipartition.

Such a subgraph G exists, because the empty graph on S trivially satisfies the two conditions above, thus a desired maximal subgraph can be obtained by adding edges of $\Phi_S(\mathcal{P})$ as long as the second condition holds. By the definition of $\Phi_S(\mathcal{P})$, each edge in G corresponds to a unique bipartition in $\mathcal{P}$, and by the second condition no two edges in G correspond to the same bipartition. On the other hand each bipartition in $\mathcal{P}$ corresponds to some edge in G, because otherwise, we could further extend G by adding one of the edges in $\Phi_S(\mathcal{P})$ corresponding to the bipartition: note that each bipartition in $\mathcal{P}$ corresponds to at least one edge in $\Phi_S(\mathcal{P})$, since $\mathcal{P}$ is minimal. Therefore, the edges in G are in one-to-one correspondence with the bipartitions in $\mathcal{P}$.

Next suppose for contradiction that G contains a cycle. Consider any bipartition $P \in \mathcal{P}$ separating some edge in a cycle. As argued above P only separates its own edge e_P of G. This is however impossible because there are two ways to connect the end-vertices of e_P in the cycle and thus P separates not only e_P but also some edge in the other path connecting the end-vertices of e_P. Therefore, we conclude that G is a forest. Since the vertex set of G is of size n and the edge set of G is of size $n-1$, it follows that G is a spanning tree on S.

To prove that $\Phi_S(\mathcal{P})$ is a spanning tree on S, suppose $G \neq \Phi_S(\mathcal{P})$. Thus there is an edge f in $\Phi_S(\mathcal{P})$ which is not in G, and adding f to G yields a cycle.

Let e and e' be two distinct edges in the cycle other than f. The second condition implies that the bipartitions P_e and $P_{e'}$ corresponding to e and e' do not separate any edges in G other than e and e', respectively. Thus both P_e and $P_{e'}$ must separate f. Since each edge of $\Phi_S(\mathcal{P})$ corresponds to a unique bipartition, we obtain $P_e = P_{e'}$, which is a contradiction. Therefore, we obtain $G = \Phi_S(\mathcal{P})$ and $\Phi_S(P)$ is indeed a spanning tree on S. It is straightforward to see that Φ_S is a one-to-one and onto mapping.

□

Since Cayley's formula [23] states that the number of spanning trees in the complete graph on n labeled vertices is n^{n-2}, we immediately obtain the following enumeration result.

Theorem 4.2.3. *The number of minimal separating families of maximum size for an n-element set is n^{n-2}.*

4.2.3 Enumerating separating families of arbitrary size

In this section we shift our attention to separating families which need not be minimal and enumerate separating families of arbitrary size by representing them as matrices whose entries are 0 or 1. This result includes the enumeration of minimal separating families of minimum size as a special case.

Throughout this section we assume without loss of generality that $S = \{1, \ldots, n\}$.

Definition 4. For each bipartition P of S, we define $b(P)$ to be the vector of length n whose i-th coordinate is given by

$$b_i(P) = \begin{cases} 1 & \text{if } P \text{ separates 1 and } i \\ 0 & \text{otherwise.} \end{cases}$$

Any k-tuple $(P_1, \ldots, P_k)$ of bipartitions can then be encoded as an $n \times k$ matrix $M_{(P_1,\ldots,P_k)}$ whose j-th column vector is $b(P_j)$.

We distinguish between a family of bipartitions and a tuple of bipartitions by denoting the former as $\mathcal{P}$, $\mathcal{Q}$, *etc* and the latter as P, Q, *etc.* We call vectors and matrices whose entries are 0 or 1 $(0,1)$-*vectors* and $(0,1)$-*matrices*, respectively.

Example 3. The matrix-representations of the tuples (P_1, P_2) and (Q_1, Q_2, Q_3) whose entries are given in Example 2 are

$$M_{(P_1,P_2)} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \\ 1 & 0 \\ 1 & 1 \end{pmatrix}, \qquad M_{(Q_1,Q_2,Q_3)} = \begin{pmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \end{pmatrix}.$$

The following lemma is straightforward to verify.

Lemma 4.2.4. *The mapping $P \mapsto M_P$ is a bijection from the set of all k-tuples of bipartitions of S to the set of all $(0,1)$-matrices of size $n \times k$ such that the entries in the first row are all 0.*

Note that column vectors whose entries are all 0 correspond to the bipartition $\{S\}$.

Lemma 4.2.5. *Let $\mathcal{P}$ be a family of bipartitions of S, and let P be a tuple obtained from $\mathcal{P}$ by ordering its members. It holds that $\mathcal{P}$ is a separating family for S if and only if every two row vectors of M_P are distinct.*

Proof. It is easy to see that for every two elements $i, j \in S$ there is a bipartition in $\mathcal{P}$ separating them if and only if the i-th row vector and the j-th row vector of M_P are distinct. From this observation this lemma immediately follows.

□

The following result is known (see for example Lemma 1 in [37] and Theorem 4 in [59]), but for the completeness we prove it.

Proposition 4.2.6. *The minimum size of a separating family for an n-element set is $\lceil \log_2 n \rceil$.*

Proof. Suppose for contradiction that the minimum size m of a separating family is less than $\lceil \log_2 n \rceil$. We have $2^m < n$ and since the number of all $(0,1)$-vectors of length m is 2^m, at least two row vectors in an $n \times m$ matrix would coincide. This is a contradiction according to Lemma 4.2.5.

To see that the minimum size is at most $\lceil \log_2 n \rceil$, we construct a tuple P of bipartitions in such a way that the row vectors of the corresponding matrix M_P are all distinct.

To achieve this, a row length of $\lceil \log_2 n \rceil$ is sufficient, and by Lemma 4.2.5 the family consisting of the entries of P is a separating family.

□

On the other hand, the maximum size of a separating family for an n-element set is 2^{n-1} because a separating family of maximum size contains all possible bipartitions.

Let us denote by $\left\{ {k \atop i} \right\}$ the number of partitions of a k-element set into i nonempty subsets. This number is known as a *Stirling number of the second kind* (see [48, §6.1]).

Example 4. As shown below, there are 6 ways to partition the 4-element set $\{1, 2, 3, 4\}$ into three nonempty subsets, yielding $\left\{ {4 \atop 3} \right\} = 6$.

$$\{\{1\}, \{2\}, \{3, 4\}\} \quad \{\{1\}, \{2, 3\}, \{4\}\} \quad \{\{1\}, \{2, 4\}, \{3\}\}$$
$$\{\{1, 2\}, \{3\}, \{4\}\} \quad \{\{1, 3\}, \{2\}, \{4\}\} \quad \{\{1, 4\}, \{2\}, \{3\}\}$$

Lemma 4.2.7. *The number of sequences of length k containing each element from a set of i symbols at least once is $i! \left\{ {k \atop i} \right\}$.*

Proof. When we identify distinct positions in a sequence of length k with k distinct objects, there are $\left\{ {k \atop i} \right\}$ many ways to distribute k distinct objects into i indistinguishable boxes, with no box left empty. If the boxes are distinguishable, then there are $i! \left\{ {k \atop i} \right\}$ many possibilities.

□

Let us denote by $\tau_{n,i}$ the number of separating families of size i for an n-element set.

Example 5. The following diagrams illustrate $\tau_{3,3} = 4$, where in the leftmost diagram the circle, the left arc and the right arc represent the bipartitions $\{\{1, 2, 3\}\}$, $\{\{1, 3\}, \{2\}\}$, and $\{\{1, 2\}, \{3\}\}$, respectively.

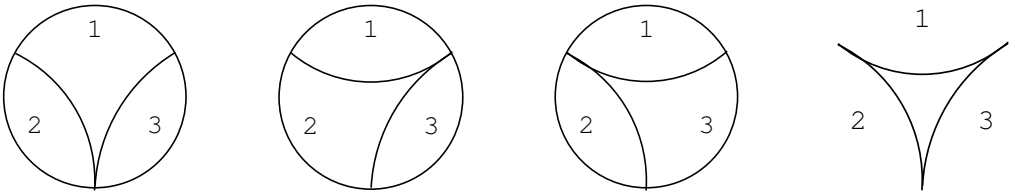

Lemma 4.2.8. *The following equation holds*

$$\sum_{i=1}^{k} i! \left\{ {k \atop i} \right\} \tau_{n,i} = (2^k - 1)(2^k - 2) \cdots (2^k - n + 1),$$

where $2 \leq n$ *and* $1 \leq k \leq 2^{n-1}$.

Proof. It is easy to verify the case $k < \lceil \log_2 n \rceil$. For the case $k \geq \lceil \log_2 n \rceil$, we prove the equation by counting in two ways the number of $(0,1)$-matrices of size $n \times k$ whose entries in the first row are all 0 and for which every two row vectors are distinct.

For the first way to count, observe that the desired matrices can be obtained by arranging distinct nonzero $(0,1)$-vectors of length k from the second row to the n-th row. Since the number of all nonzero $(0,1)$-vectors of length k is $2^k - 1$, the number of such matrices is $\binom{2^k-1}{n-1} (n-1)!$.

For the second way to count, observe that the column vectors of a desired matrix need not be distinct. For each $i \in \{1, \ldots, k\}$ we count the number of all desired $n \times k$ matrices having i distinct column vectors. We can assume $i \geq \lceil \log_2 n \rceil$ because otherwise there are no such matrices. Clearly those i column vectors form a separating family of size i. Conversely, the desired matrices can be obtained from separating families of size i by arranging the column vectors corresponding to the i bipartitions with repetition from the first column to the k-th column in such a way that each vector occurs at least once. Combining this observation with Lemma 4.2.7, it follows that the number of all matrices constructed in this way is $i! \left\{ {k \atop i} \right\} \tau_{n,i}$. By summing over all $i \in \{1, \ldots, k\}$, we obtain the left side of the equation stated in this lemma.

□

Let us denote by $\begin{bmatrix} k \\ i \end{bmatrix}$ the number of permutations of k elements which contain exactly i permutation cycles. This number is known as an *unsigned Stirling number of the first kind* (see [48, §6.1]).

Example 6. As shown below, there are 3 permutations of the 3-element set $\{1, 2, 3\}$ containing two permutation cycles, yielding $\begin{bmatrix} 3 \\ 2 \end{bmatrix} = 3$.

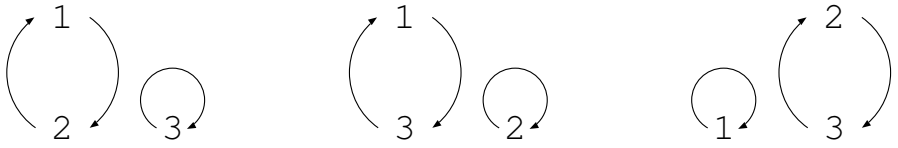

Theorem 4.2.9. *The number $\tau_{n,k}$ of separating families of size k for an n-element set is*

$$\frac{(n-1)!}{k!} \sum_{i=1}^{k} (-1)^{k-i} \begin{bmatrix} k \\ i \end{bmatrix} \binom{2^i - 1}{n-1},$$

where $2 \leq n$ and $1 \leq k \leq 2^{n-1}$.

Proof. The inversion formula for Stirling numbers (see [15, §3.1]) states that if two integer sequences $\{a_i\}_{1 \leq i \leq n_0}$ and $\{b_i\}_{1 \leq i \leq n_0}$ satisfy $b_k = \sum_{i=1}^{k} \left\{ {k \atop i} \right\} a_i$, then they also satisfy $a_k = \sum_{i=1}^{k} (-1)^{k-i} \begin{bmatrix} k \\ i \end{bmatrix} b_i$. From Lemma 4.2.8 we obtain this theorem.

□

Proposition 4.2.10. *The number of minimal separating families of minimum size for an n-element set is*

$$\frac{(n-1)!}{\lceil \log_2 n \rceil !} \binom{2^{\lceil \log_2 n \rceil} - 1}{n-1}.$$

Proof. Since every separating family of minimum size is minimal, it suffices to calculate the number of separating families of minimum size. According to Proposition 4.2.6 their minimum size for an n-element set is $\lceil \log_2 n \rceil$. Substituting $\lceil \log_2 n \rceil$ for k in the formula of Theorem 4.2.9 yields the expression stated in this proposition, since $\binom{2^i-1}{n-1} = 0$ for $i \leq \lceil \log_2 n \rceil$.

□

This number corresponds to the number of state assignments of an n-state machine in switching theory (see [64, chap. 12]).

Let us denote by $\sigma_{n,k}$ the number of separating families of size k for an n-element set such that every bipartition consists of exactly two components, i.e. it does not contain $\{S\}$.

Lemma 4.2.11. *The equation $\sigma_{n,k-1} + \sigma_{n,k} = \tau_{n,k}$ holds for all numbers n and k with $2 \leq n$ and $2 \leq k \leq 2^{n-1}$.*

Proof. Observe in Example 5 that the first three diagrams can be identified with separating families of size 2 by excluding $\{S\}$ and that the rightmost diagram is a unique separating family of size 3 such that every bipartition consists of exactly two components. Thus we obtain $\sigma_{3,2} + \sigma_{3,3} = \tau_{3,3}$. In general one obtains that separating families of size k which contain $\{S\}$ are identified with those of size $k-1$ which do not contain $\{S\}$. Since there are $\sigma_{n,k-1}$ such families, this lemma follows.

□

Lemma 4.2.12.

$$\begin{bmatrix} k+1 \\ i+1 \end{bmatrix} = k! \sum_{j=i}^{k} \frac{1}{j!} \begin{bmatrix} j \\ i \end{bmatrix}$$

Proof. It holds (see [48, §6.1]) that $\begin{bmatrix} k+1 \\ i+1 \end{bmatrix} = k! \sum_{j=0}^{k} \begin{bmatrix} j \\ i \end{bmatrix} / j!$. Since $\begin{bmatrix} j \\ i \end{bmatrix} = 0$ for $0 \leq j < i$, we obtain the equation.

□

Proposition 4.2.13. *The number $\sigma_{n,k}$ of separating families of size k for an n-element set such that every bipartition consists of exactly two components is*

$$\frac{(n-1)!}{k!} \sum_{i=1}^{k} (-1)^{k-i} \begin{bmatrix} k+1 \\ i+1 \end{bmatrix} \binom{2^i - 1}{n-1},$$

where $2 \leq n$ and $1 \leq k < 2^{n-1}$.

Proof. Using Lemma 4.2.11 and 4.2.12, we can calculate $\sigma_{n,k}$ as follows.

$$\begin{aligned}
\sigma_{n,k} &= \sum_{j=1}^{k} (-1)^{k-j} \tau_{n,j} \\
&= (n-1)! \sum_{j=1}^{k} \sum_{i=1}^{j} \frac{(-1)^{k-i}}{j!} \begin{bmatrix} j \\ i \end{bmatrix} \binom{2^i - 1}{n-1} \\
&= (n-1)! \sum_{i=1}^{k} \sum_{j=i}^{k} \frac{(-1)^{k-i}}{j!} \begin{bmatrix} j \\ i \end{bmatrix} \binom{2^i - 1}{n-1} \\
&= (n-1)! \sum_{i=1}^{k} \frac{(-1)^{k-i}}{k!} \begin{bmatrix} k+1 \\ i+1 \end{bmatrix} \binom{2^i - 1}{n-1}
\end{aligned}$$

□

Lemma 4.2.14. *The equation $\sigma_{n,k-1} (k-1)! = \sigma_{k,n-1} (n-1)!$ holds for all numbers n, k such that $2 \leq n$ and $2 \leq k \leq 2^{n-1}$.*

Proof. Consider $(0,1)$-matrices whose entries in the first row and in the first column are all 0 and for which every two row vectors and every two column vectors are distinct. Since both conditions are symmetric with respect to rows and columns, the matrix transpose operation induces a bijection between the set of $(0,1)$-matrices of size $n \times k$ satisfying the conditions and the set of those of size $k \times n$ satisfying the conditions. Observe for example that the following matrices which are transposes of each other satisfy the conditions.

$$M = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 1 & 1 \end{pmatrix} \qquad M^t = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 \end{pmatrix}$$

Furthermore, every $(0,1)$-matrix of size $n \times k$ satisfying the conditions can be obtained by arranging the bipartitions in a separating family of size $k-1$ which does not contain $\{S\}$ from the second column to the k-th column. Thus the number of such matrices is $\sigma_{n,k-1}(k-1)!$. Since the same argument applies to $(0,1)$-matrices of size $k \times n$, this lemma follows.

□

From Proposition 4.2.13 and Lemma 4.2.14, one obtains the following proposition.

Proposition 4.2.15. *The number $\sigma_{n,k}$ of separating families of size k for an n-element set such that every bipartition consists of exactly two components is*

$$\sum_{i=1}^{n-1} (-1)^{n-1-i} \begin{bmatrix} n \\ i+1 \end{bmatrix} \binom{2^i - 1}{k},$$

where $2 \leq n$ *and* $1 \leq k < 2^{n-1}$.

Proof.

$$\begin{aligned}\sigma_{n,k} &= \frac{(n-1)!}{k!}\,\sigma_{k+1,n-1} \\ &= \frac{(n-1)!}{k!}\,\frac{k!}{(n-1)!}\sum_{i=1}^{n-1}(-1)^{n-1-i}\begin{bmatrix} n \\ i+1 \end{bmatrix}\binom{2^i-1}{k}\end{aligned}$$

□

From Lemma 4.2.11 and Proposition 4.2.15, one obtains the following theorem.

Theorem 4.2.16. *The number* $\tau_{n,k}$ *of separating families of size* k *for an* n*-element set is*

$$\sum_{i=1}^{n-1}(-1)^{n-1-i}\begin{bmatrix} n \\ i+1 \end{bmatrix}\binom{2^i}{k},$$

where $2 \leq n$ *and* $2 \leq k < 2^{n-1}$.

Proof.

$$\begin{aligned}\tau_{n,k} &= \sigma_{n,k} + \sigma_{n,k-1} \\ &= \sum_{i=1}^{n-1}(-1)^{n-1-i}\begin{bmatrix} n \\ i+1 \end{bmatrix}\left\{\binom{2^i-1}{k} + \binom{2^i-1}{k-1}\right\} \\ &= \sum_{i=1}^{n-1}(-1)^{n-1-i}\begin{bmatrix} n \\ i+1 \end{bmatrix}\binom{(2^i-1)+1}{k}\end{aligned}$$

□

Proposition 4.2.17. *The minimum size of an element set for which there is a separating family of* k *arbitrary bipartitions is* $\lceil \log_2 k \rceil + 1$.

Chapter 5

Packing $\mathbb{R}^3$ with Thin Tori

5.1 Introduction and Main Result

Geometric packing problems in $\mathbb{R}^3$ received huge attention over many decades (see [9],[18],[39],[40],[41],[42],[49],[50],[74],[82],[85],[101] for books on packings). Still, the sphere is the only body which does not tile $\mathbb{R}^3$ and for which we know the exact packing density [53]. For other bodies such as Platonic solids [26],[93] and ellipsoids [16],[35] dense packing constructions are known, but no proof of optimality exists and a vast amount of related questions remain open (see the books [19] and [31]).

On the other hand, there is only a very limited amount of literature studying packings involving non-convex objects, such as the work of Jiao et al. [57],[94],[95]. We would thus like to extend this line of research by considering packings with the possibly simplest non-convex shape, the torus.

Note that although Conway and Hopcroft showed in [29] that using the axiom of choice it is possible to fill $\mathbb{R}^3$ with unit circles, the problem addressed here is of different type

since we are dealing with solid bodies and not with one dimensional curves.

Throughout this chapter we refer to a torus with major radius 1 and minor radius going to zero as a *thin torus*.

Theorem 5.1. $\mathbb{R}^3$ *can be packed with thin tori at a density of* 0.222....

Theorem 5.2. *It is possible to pairwise link* $\Theta(1/s^2)$ *many tori of major radius* 1 *and minor radius* $s < 1/3$ *and this is asymptotically optimal.*

5.2 Proof of Theorem 1

Given a torus lying in a plane H and a point p on its surface, there are three different types of non-trivial circles through p. One can draw a circle through p which lies in a plane parallel to the plane H and one can draw a circle through p in the plane perpendicular to H. A third type of circles is obtained by cutting the torus open at p along a plane which is bitangential to the torus and then draw two circles on the surface of the torus along the cut as shown in Figure 5.1a. These circles have a radius which correspond to the major radius of the torus and were first observed by Yvon Villarceau; thus they are called *Villarceau circles* [99]. Most prominently, Villarceau circles appear in Topology, when the Hopf fibration [55] of a 3-sphere is stereographically projected into $\mathbb{R}^3$.

In order for the torus arrangement to pack $\mathbb{R}^3$ with positive density, it is necessary that some tori are linked. This holds, since the volume of the bounding box of a torus with minor radius s is $\Theta(s)$, while the volume of the torus is $\Theta(s^2)$. Thus for $s \to 0$, the packing density of an arrangement of unlinked thin tori goes to 0. Therefore, it is not a

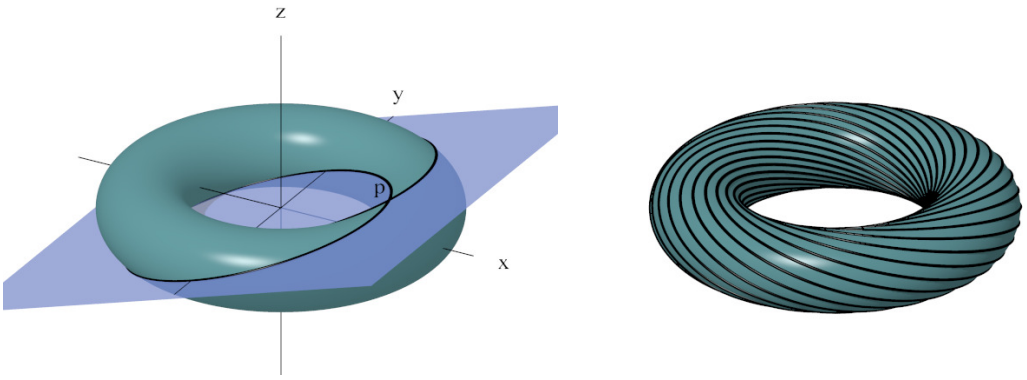

(A) (B)

FIGURE 5.1: (a) An illustration [43] of a bitangential plane and its two Villarceau circles through p. (b) A family of Villarceau circles obtained by rotating the bitangential cutting plane in discrete steps around the z-axis.

priori clear that a packing of thin tori with positive density exists.

Throughout this note we denote by $T(R, r)$ a torus of major radius R and minor radius r. The basic idea of the torus arrangement is to first pack $\mathbb{R}^3$ with an auxiliary lattice packing of fat tori $T_F = T(1, r)$ for some constant r. All such tori are placed parallel to the xy-plane and are centered at points of the lattice, generated by the vectors

$$\begin{pmatrix} 2+\sqrt{3}r \\ 0 \\ r \end{pmatrix}, \begin{pmatrix} 1+r \\ \sqrt{3}(1+r) \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 2r \end{pmatrix}.$$

Since the parallelepiped of this lattice has volume $2\sqrt{3}\left(2+\sqrt{3}r\right)(1+r)\,r$, the lattice packing of $\mathbb{R}^3$ with T_F tori has a density of

$$\delta_L(r) = \frac{\pi^2 r}{\sqrt{3}\left(2+\sqrt{3}r\right)(1+r)}. \tag{5.1}$$

Inside every T_F torus we build a second auxiliary structure by forming a nested sequence of concentric tori $T_{\lfloor \frac{r-s}{2s} \rfloor}, \ldots, T_1$ centered at lattice points, with $T_k = T(1, 2ks)$. Here s denotes the minor radius of the thin tori which we will ultimately use to pack $\mathbb{R}^3$. Next, a $2s$ neighborhood of the surface of each T_k torus gets packed with thin $T(1, s)$ tori as illustrated in Figure 5.2. For this we construct a family of Villarceau circles by rotating the bitangential cutting plane of T_k in discrete steps around the z-axis in such a way that the smallest distance between any two Villarceau circles is at least $2s$ (see Figure 5.1b). Note that we only chose one of the two resulting Villarceau circles associated with each cutting plane. Replacing every Villarceau circle by a $T(1, s)$ torus yields a packing of the $2s$ neighborhood of the surface of T_k. Denoting by V_k the volume of the union of all $T(1, s)$ tori on T_k, we obtain a total volume of

$$V = \sum_{k=1}^{\lfloor \frac{r-s}{2s} \rfloor} V_k$$

for the nested arrangement of thin tori inside T_F and thus a packing density of $\delta_T(r) = V/(2\pi^2 r^2)$ with respect to T_F. Therefore, the packing density of the $T(1, s)$ torus arrangement with respect to $\mathbb{R}^3$ is $\delta(r) = \delta_T(r)\delta_L(\mathrm{r})$.

The remaining part of this section is used to calculate the volume V.

Lemma 5.3. *Given a torus $T(R, r)$, lying in the xy-plane and two Villarceau circles c_0, c_1 lying in the bitangential planes H_0, H_1 respectively. In order for the minimum distance between c_0 and c_1 to be at least d, an angular distance between H_0 and H_1 of $\psi = \arcsin \frac{dR}{r(R-r)}$ around the z-axis suffices, given $r \leq R - d$.*

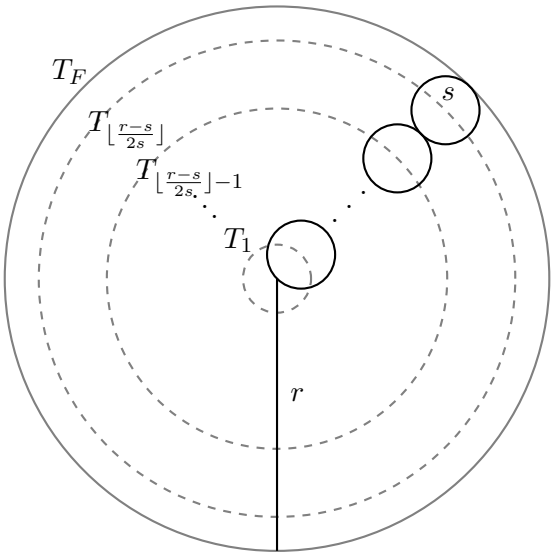

FIGURE 5.2: A vertical cut through T_F, the nested sequence of T_k tori and a schematic representation of the $T(1, s)$ tori through one point on each T_k.

Proof. A Villarceau circle of a torus $T(R, r)$ centered at the origin and lying in the xy-plane can be parameterized (see [30]) as

$$\mathbf{c}(\psi, t) = \begin{pmatrix} \cos\psi & \sin\psi & 0 \\ -\sin\psi & \cos\psi & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \sqrt{R^2 - r^2}\cos t \\ r + R\sin t \\ r\cos t \end{pmatrix}.$$

Here ψ denotes the rotation angle of the bitangential cutting plane around the z-axis and t defines the location on the Villarceau circle. At $\psi = 0$, the derivatives of $\mathbf{c}$ are

$$\frac{\partial \mathbf{c}}{\partial t} = \begin{pmatrix} -\sqrt{R^2 - r^2}\sin t \\ R\cos t \\ -r\sin t \end{pmatrix}$$

and

$$\frac{\partial \mathbf{c}}{\partial \psi} = \begin{pmatrix} r + R\sin t \\ -\sqrt{R^2 - r^2}\cos t \\ 0 \end{pmatrix}.$$

The closest distance between the circle $c_0 = \mathbf{c}(0,t)$ and a circle c_1 rotated around the z-axis appears on c_0 when $\frac{\partial \mathbf{c}}{\partial \psi}$ assumes the smallest value in the direction perpendicular to $\frac{\partial \mathbf{c}}{\partial t}$, i.e. at the point where

$$\left|\frac{\partial \mathbf{c}}{\partial \psi}\right| \sin\gamma$$

attains its minimum, with γ denoting the angle between $\frac{\partial \mathbf{c}}{\partial t}$ and $\frac{\partial \mathbf{c}}{\partial \psi}$.

Since

$$\sin\gamma = \sqrt{1 - \left(\frac{\frac{\partial \mathbf{c}}{\partial \psi} \cdot \frac{\partial \mathbf{c}}{\partial t}}{\left|\frac{\partial \mathbf{c}}{\partial \psi}\right| \left|\frac{\partial \mathbf{c}}{\partial t}\right|}\right)^2}$$

we want to minimize

$$\sqrt{\left|\frac{\partial \mathbf{c}}{\partial \psi}\right|^2 - \left(\frac{\frac{\partial \mathbf{c}}{\partial \psi} \cdot \frac{\partial \mathbf{c}}{\partial t}}{\left|\frac{\partial \mathbf{c}}{\partial t}\right|}\right)^2},$$

which simplifies to

$$\frac{r}{R}\sqrt{\left(2\,rR\sin(t) + r^2 + R^2 + ((\sin(t))^2 - 1)r^2\right)}$$

and attains its minimum at $t = -\arcsin R/r$. Since $r < R$ the minimum occurs when $\sin t = -1$. If we denote this point $(0, r-R, 0)$ by q, its minimum distance to the cutting plane H_1 provides a lower bound for the minimum distance d between c_0 and c_1. Since the normal unit vector of H_1 is $\mathbf{n_1} = (r/R\cos\psi, -r/R\sin\psi, \sqrt{1 - r^2/R^2})$,

$$d \geq \mathbf{n_1} \cdot q = \frac{r(R-r)}{R} \sin\psi.$$

Thus an angular distance between H_0 and H_1 around the z-axis of at least

$$\psi = \arcsin \frac{dR}{r(R-r)}$$

results in a minimum distance of at least d between c_0 and c_1.

□ □

It follows from Lemma 5.3 that the $2s$ neighborhood of the surface of a T_k torus can be packed with $\left\lfloor \frac{2\pi}{\arcsin \frac{1}{k(1-2ks)}} \right\rfloor$ many $T(1,s)$ tori. Using the fact that $11x/7 \geq \arcsin x$ for $0 \leq x \leq 1$, the volume of the union of these tori is

$$V_k \geq 2\pi^2 s^2 \left\lfloor \frac{14}{11} \pi k(1-2ks) \right\rfloor.$$

For $s \to 0$, the volume of the union of all $T(1,s)$ tori inside one T_F torus becomes

$$\overline{V} = \lim_{s\to 0} \sum_{k=1}^{\lfloor \frac{r-s}{2s} \rfloor} V_k \geq \frac{7}{66} \pi^3 r^2 (3 - 2r)$$

and thus the packing density $\delta_T(r)$ with respect to T_F is

$$\delta_T(r) = \frac{7}{132} \pi (3 - 2r). \tag{5.2}$$

Combining equation (5.1) and (5.2) the packing density of the thin tori with respect to $\mathbb{R}^3$ evaluates to

$$\delta(r) = \delta_L(r)\delta_T(r) = \frac{7\pi^3 r\,(3-2\,r)}{132\sqrt{3}\left(2+\sqrt{3}r\right)(1+r)}$$

which, at $r = 0.441...$, obtains the maximum value of $0.222....$

5.3 Proof of Theorem 2

In order to prove Theorem 2, we show that in the nested torus construction of the previous section all thin tori are pairwise linked. Let c_0 and c_1 be two Villarceau circles on two nested (auxiliary) unit tori T_0 and T_1 with minor radius r_0 and r_1 respectively and assume $r_0 < r_1$. We prove the claim by showing that c_1 and T_0 are linked, which, since Villarceau circles are fully contained in the surface of a torus, is equivalent to showing that c_1 is linked with the unit circle c defining the axis of revolution of T_0. Furthermore, wlog we may assume that the cutting plane defining c_1 is unrotated (around the z-axis), since otherwise we may just rotate the whole coordinate system around the z-axis. It thus follows that c_1 intersects the y-axis at distances $-1+r_1$ and $1+r_1$. In order to show that c_1 and c are linked, it suffices to note that y-axis intersects c inside the interval $(-1+r_1, 1+r_1)$ exactly once, namely at 1. Since all Villarceau circles on the same torus are linked (see [30]), the claim follows.

We set the minor radius r of T_F, defined in the previous section, equal to 1, i.e. T_F becomes a Horn torus and we note that the outermost torus of the nested structure inside T_F has minor radius $1 - s$. As argued in the previous section, the number of pairwise linked tori of minor radius s in the nested construction is

$$\sum_{k=1}^{\lfloor \frac{1-s}{2s} \rfloor} \left\lfloor \frac{14}{11} \pi k (1 - 2ks) \right\rfloor,$$

which is lower bounded by

$$\frac{1}{132} \frac{(3\,s-1)\left(-7\,\pi - 21\,\pi\,s + 66\,s + 28\,\pi\,s^2\right)}{s^2}. \tag{5.3}$$

It is easy to see that expression (5.3) is contained in $\Omega(1/s^2)$ for $s \in (0, 1/3)$. On the other hand, a simple area argument shows that at most $O(1/s^2)$ tori of major radius 1 and minor radius s can be linked with a single such torus, thus implying that the construction is asymptotically optimal. Obviously, the construction is also asymptotically optimal for linking many tori with one single torus.

Bibliography

[1] Rashmisnata Acharyya, Manjanna Basappa, and Gautam Das. Unit disk cover problem in 2d. In *Computational Science and Its Applications, ICCSA 2013*, volume 7972 of *Lecture Notes in Computer Science*, pages 73–85. Springer Berlin Heidelberg, 2013.

[2] Pankaj K. Agarwal and Micha Sharir. *Davenport-Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, 1998.

[3] M. Aigner. *Combinatorial Search*. John Wiley & Sons, 1988.

[4] Helmut Alt, Sergio Cabello, Panos Giannopoulos, and Christian Knauer. Minimum cell connection and separation in line segment arrangements. *CoRR*, abs/1104.4618, 2011.

[5] Helmut Alt, Sergio Cabello, Panos Giannopoulos, and Christian Knauer. On some connection problems in straight-line segment arrangements. In *EuroCG*, 2011.

[6] Boris Aronov. On the geodesic voronoi diagram of point sites in a simple polygon. *Algorithmica*, 4(1-4):109–140, 1989.

[7] Boris Aronov, Esther Ezra, and Micha Sharir. Small-size epsilon-nets for axis-parallel rectangles and boxes. *SIAM J. Comput*, 39(7):3248–3282, 2010.

[8] Boris Aronov, Steven Fortune, and Gordon Wilfong. The furthest-site geodesic voronoi diagram. *Discrete and Computational Geometry*, 9(1):217–255, 1993.

[9] T. Aste and D. Weaire. *The Pursuit of Perfect Packing.* CRC Press, Bristol, 2008.

[10] Paul Balister, Zizhan Zheng, Santosh Kumar, and Prasun Sinha. Trap coverage: Allowing coverage holes of bounded diameter in wireless sensor networks. In *In Proc. of IEEE INFOCOM, Rio de Janeiro*, 2009.

[11] Dror Bar-Natan. On the vassiliev knot invariants, 1995.

[12] Luis Barba, Prosenjit Bose, Matias Korman, Jean-Lou De Carufel, Hee-Kap Ahn, and Eunjin Oh. A linear-time algorithm for the geodesic center of a simple polygon. SoCG '15, 2015. (accepted).

[13] Sergei Bereg and David G. Kirkpatrick. Approximating barrier resilience in wireless sensor networks. In *Proc. 5th ALGOSENSORS*, pages 29–40, SpringerLNCS Vol. 5804,, Springer, 2009. LNCS Vol. 5804.

[14] Sergey Bereg and David G. Kirkpatrick. Approximating barrier resilience in wireless sensor networks. In *ALGOSENSORS*, volume 5804, pages 29–40. Springer, 2009.

[15] C. Berge. *Principles of Combinatorics.* Academic Press, 1971.

[16] A. Bezdek and W. Kuperberg. Packing Euclidean space with congruent cylinders and with congruent ellipsoids. *Victor Klee Festschrift, ed. P. Gritzmann and B. Sturmfels , American Mathematical Society*, 60:71–80, 1991.

[17] MagdaleneG. Borgelt, Marc Kreveld, and Jun Luo. Geodesic disks and clustering in a simple polygon. In Takeshi Tokuyama, editor, *Algorithms and Computation*,

volume 4835 of *Lecture Notes in Computer Science*, pages 656–667. Springer Berlin Heidelberg, 2007.

[18] K. Böröczky. *Finite Packing and Covering.* Cambridge University Press, Cambridge, 2004.

[19] P. Brass, W. Moser, and J. Pach. *Research Problems in Discrete Geometry.* Springer, New York, 2005.

[20] Sergio Cabello and Panos Giannopoulos. The complexity of separating points in the plane. In *Proceedings of the twenty-ninth annual symposium on Computational geometry*, pages 379–386. ACM, 2013.

[21] Sergio Cabello and Panos Giannopoulos. The complexity of separating points in the plane. In *Proceedings of the Twenty-ninth Annual Symposium on Computational Geometry*, SoCG '13, pages 379–386, New York, NY, USA, 2013. ACM.

[22] Gruia Calinescu and Adrian Dumitrescu. P.j.: Separating points by axisparallel lines. *International Journal of Computational Geometry and Applications*, 15:575–590, 2005.

[23] A. Cayley. A theorem on trees. *Quart. J. Math*, 23:376–378, 1889.

[24] Ai Chen, S. Kumar, and T.H. Lai. Local barrier coverage in wireless sensor networks. *IEEE Transactions on Mobile Computing*, 9(4):491–504, 2010.

[25] Ai Chen, Ten H. Lai, and Dong Xuan. Measuring and guaranteeing quality of barrier-coverage in wireless sensor networks. In *Proceedings of the 9th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, MobiHoc '08, pages 421–430, New York, NY, USA, 2008. ACM.

[26] E. Chen, M. Engel, and S. Glotzer. Dense crystalline dimer packings of regular tetrahedra. *Discrete and Computational Geometry*, 44:253–280, 2010.

[27] Francis Chin, Jack Snoeyink, and CaoAn Wang. Finding the medial axis of a simple polygon in linear time. In John Staples, Peter Eades, Naoki Katoh, and Alistair Moffat, editors, *Algorithms and Computations*, volume 1004 of *Lecture Notes in Computer Science*, pages 382–391. Springer Berlin Heidelberg, 1995.

[28] Kenneth L. Clarkson and Kasturi Varadarajan. Improved approximation algorithms for geometric set cover. In *Proc. Symposium on Computational Geometry*, pages 135–141. SCG '05, 2005.

[29] J. H. Conway and H. T. Croft. Covering a sphere with congruent great-circle arcs. *Proceedings of the Cambridge Philosophical Society*, 60:787, 1964.

[30] H. S. M. Coxeter. *Introduction to Geometry.* Wiley, New York, 1961.

[31] H.T. Croft, K.J. Falconer, and R.K. Guy. *Unsolved Problems in Geometry.* Springer, New York, 1991.

[32] Elias Dahlhaus, David Johnson, Christos Papadimitriou, Paul Seymour, and Mihalis Yannakakis. *The Complexity of Multiterminal Cuts.* SIAM J. Comput, 1994.

[33] Emeric Deutsch and Marc Noy. Statistics on non-crossing trees. *Discrete Mathematics*, 254(1-3):75 – 87, 2002.

[34] Reinhard Diestel. *Graph Theory.* Number 173 in Graduate Texts in Mathematics. Springer, 2010.

[35] J. Donev, F. H. Stillinger, P. M. Chaikin, and S. Torquato. Unusually dense crystal packings of ellipsoids. *Phys. Rev. Lett.*, 92:255506, Jun 2004.

[36] Serge Dulucq and Jean-Guy Penaud. Cordes, arbres et permutations. *Discrete Mathematics*, 117(1 - 3):89 – 105, 1993.

[37] A. Dumitrescu. Efficient algorithms for generation of combinatorial covering suites. In Toshihide Ibaraki, Naoki Katoh, and Hirotaka Ono, editors, *Algorithms and Computation*, volume 2906 of *Lecture Notes in Computer Science*, pages 300–308. Springer Berlin / Heidelberg, 2003.

[38] Adrian Dumitrescu and Csaba D. Tóth. Light orthogonal networks with constant geometric dilation. *Journal of Discrete Algorithms*, 7(1):112 – 129, 2009.

[39] Paul Erdös, P Gruber, and Joseph Hammer. *Lattice Points. Pitman Monographs and Surveys in Pure and Applied Mathematics 39.* Longman Scientific & Technical, Harlow, Harlow, Essex, England, 1989.

[40] G Fejes Tóth and W Kuperberg. Packing and covering with convex sets, chapter 3.3. In P. Gruber and J. Wills, editors, *Vol B of Handbook of Convex Geometry.* North Holland, 1993.

[41] G Fejes Tóth and W. Kuperberg. A survey of recent results in the theory of packing and covering. In J. Pach, editor, *New Trends in Discrete and Computational Geometry*, volume 10 of *Algorithms and Combinatorics*, pages 251–279. Springer, Berlin, New York, 1993.

[42] L Fejes Tóth. *Lagerungen in der Ebene auf der Kugel und im Raum.* Grundlehren der mathematischen Wissenschaften in Einzeldarstellungen mit besonderer Berücksichtigung der Anwendungsgebiete. Springer, Berlin, 1972.

[43] P. Fradin. http://melusine.eu.org/lab/bgr/povray/texgraph/villarceau.bgr, 2009.

[44] R. Freimer. Investigations in geometric subdivisions: linear shattering and cartographic map coloring. *Ph.D. Thesis, Cornell University*, 2000.

[45] Robert Freimer, Joseph S. B. Mitchell, and Christine Piatko. On the complexity of shattering using arrangements. *2nd Canadian Conference on Computational Geometry*, 1990.

[46] Stefan Funke, Alex Kesselman, Fabian Kuhn, Zvi Lotker, and Michael Segal. Improved approximation algorithms for connected sensor cover. *Wirel. Netw.*, 13(2):153–164, April 2007.

[47] Matt Gibson, Gaurav Kanade, and Kasturi Varadarajan. On isolating points using disks. In *Proceedings of the 19th European Conference on Algorithms*, ESA'11, pages 61–69, Berlin, Heidelberg, 2011. Springer-Verlag.

[48] R.L. Graham, D.E. Knuth, and O. Patashnik. *Concrete Mathematics*. Addison-Wesley Professional, 2nd edition, 1994.

[49] P Gritzman and JM Wills. Finite packing and covering, chapter 3.4. In P. Gruber and J. Wills, editors, *Vol B of Handbook of Convex Geometry*. North Holland, Amsterdam, London, 1993.

[50] P. M. Gruber and C. G. Lekkerkerker. *Geometry of Numbers*, volume 37. North-Holland, Amsterdam, 1987.

[51] Leonidas Guibas, John Hershberger, Daniel Leven, Micha Sharir, and RobertE. Tarjan. Linear-time algorithms for visibility and shortest path problems inside triangulated simple polygons. *Algorithmica*, 2(1-4):209–233, 1987.

[52] Leonidas J. Guibas and John Hershberger. Optimal shortest path queries in a simple polygon. *Journal of Computer and System Sciences*, 39(2):126 – 152, 1989.

[53] T. C. Hales. A proof of the Kepler conjecture. *The Annals of Mathematics*, 162(3):1065–1185, 2005.

[54] Peter Hilton and Jean Pedersen. Catalan numbers, their generalization, and their uses. *The Mathematical Intelligencer*, 13:64–75, 1991. 10.1007/BF03024089.

[55] H. Hopf. Über die Abbildungen der dreidimensionalen Sphäre auf die Kugelfläche. *Mathematische Annalen*, 104:637–665, 1931.

[56] Chi-Fu Huang and Yu-Chee Tseng. The coverage problem in a wireless sensor network. In *Proceedings of the 2Nd ACM International Conference on Wireless Sensor Networks and Applications*, WSNA '03, pages 115–121, New York, NY, USA, 2003. ACM.

[57] Y. Jiao, F. H. Stillinger, and S. Torquato. Optimal packings of superballs. *Phys. Rev. E*, 79:041309, Apr 2009.

[58] G.O.H. Katona. Combinatorial search problems. In J.N. Srivastava, F. Harary, C.R. Rao, G.-C. Rota, and S.S. Shrikhande, editors, *A Survey of Combinatorial Theory*, pages 285–308. North-Holland Publishing Company, Amsterdam, Netherlands, 1973.

[59] G.O.H. Katona. Rényi and the combinatorial search problems. *Studia Scientiarum Mathematicarum Hugarica*, 26:363–378, 1991.

[60] Klara Kedem, Ron Livne, János Pach, and Micha Sharir. On the union of jordan regions and collision-free translational motion amidst polygonal obstacles. *Discrete & Computational Geometry*, 1(1):59–71, 1986.

[61] Richard Kershner. The number of circles covering a set. *American Journal of Mathematics*, pages 665–671, 2010.

[62] Martin Klazar. On non p-recursiveness of numbers of matchings (linear chord diagrams) with many crossings. *Adv. Appl. Math*, 30:126–136.

[63] Ren-Song Ko. The complexity of the minimum sensor cover problem with unit-disk sensing regions over a connected monitored region. *IJDSN*, 2012, 2012.

[64] Z. Kohavi and N.K. Jha. *Switching and Finite Automata Theory*. Cambridge University Press, 3rd edition, 2010.

[65] Matias Korman, Maarten Löffler, Rodrigo I. Silveira, and Darren Strash. On the complexity of barrier resilience for fat regions. In Paola Flocchini, Jie Gao, Evangelos Kranakis, and Friedhelm Meyer auf der Heide, editors, *Algorithms for Sensor Systems*, Lecture Notes in Computer Science, pages 201–216. Springer Berlin Heidelberg, 2014.

[66] Santosh Kumar, Ten H. Lai, and Anish Arora. Barrier coverage with wireless sensors. In *Proceedings of the 11th Annual International Conference on Mobile Computing and Networking*, MobiCom '05, pages 284–298, New York, NY, USA, 2005. ACM.

[67] Santosh Kumar, Ten H. Lai, and Anish Arora. Barrier coverage with wireless sensors. in. In *MobiCom '05: Proceedings of the 11th annual international conference on Mobile computing and networking*, pages 284–298. New York, 2005.

[68] Benyuan Liu, Olivier Dousse, Jie Wang, and Anwar Saipulla. Strong barrier coverage of wireless sensor networks. In *Proceedings of the 9th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, MobiHoc '08, pages 411–420, New York, NY, USA, 2008. ACM.

[69] Nabil H. Mustafa and Saurabh Ray. Ptas for geometric hitting set problems via local search. In *Proc. Symposium on Computational Geometry*, pages 17–22. SCG 09, 2009.

[70] Subhas C. Nandy, Tetsuo Asano, and Tomohiro Harayama. Shattering a set of objects in 2d. *Discrete Appl. Math.*, 122:183–194, October 2002.

[71] Albert Nijenhuis and Herbert S Wilf. The enumeration of connected graphs and linked diagrams. *Journal of Combinatorial Theory, Series A*, 27(3):356 – 359, 1979.

[72] Marc Noy. Enumeration of noncrossing trees on a circle. *Discrete Math.*, 180:301–313, February 1998.

[73] Joseph O'Rourke. *Art Gallery Theorems and Algorithms*. Oxford University Press, Inc., New York, NY, USA, 1987.

[74] J. Pach and P. K. Agarwal. *Combinatorial Geometry*. John Wiley & Sons, New York, 1995.

[75] János Pach. On the complexity of the union of geometric objects. In Jin Akiyama, Mikio Kano, and Masatsugu Urabe, editors, *Discrete and Computational Geometry*, volume 2098 of *Lecture Notes in Computer Science*, pages 292–307. Springer Berlin Heidelberg, 2001.

[76] Alios Panholzer and Helmut Prodinger. Bijection for ternary trees and non-crossing trees. *Discrete Math.*, 250:181–195, May 2002.

[77] Alois Panholzer. Non-crossing trees revisited: cutting down and spanning subtrees. In *Proceedings, Discrete Random Walks 2003, Cyril Banderier and*, pages 265–276, 2003.

[78] Rainer Penninger and Ivo Vigan. Point set isolation using unit disks is np-complete. In *Fall Workshop of Computational Geometry*, 2012.

[79] R. Pollack, M. Sharir, and G. Rote. Computing the geodesic center of a simple polygon. *Discrete and Computational Geometry*, 4:611–626, 1989.

[80] Paul R and Stein. On a class of linked diagrams, i. enumeration. *Journal of Combinatorial Theory, Series A*, 24(3):357 – 366, 1978.

[81] A. Rényi. On the theory of random search. *Bull. Amer. Math. Soc.*, 71:809–828, 1965.

[82] C.A. Rogers. *Packing and Covering, Cambridge, 1964.* Cambridge Univ. Press, Cambridge, 1964.

[83] Anwar Saipulla, Cedric Westphal, Benyuan Liu, and Jie Wang. Barrier coverage of line-based deployed wireless sensor networks. In *INFOCOM*, pages 127–135. IEEE, 2009.

[84] Swaminathan Sankararaman, Alon Efrat, Srinivasan Ramasubramanian, and Javad Taheri. Scheduling sensors for guaranteed sparse coverage. *CoRR*, abs/0911.4332, 2009.

[85] G. Scheithauer. *Zuschnitt- und Packungsoptimierung.* Springer, Wiesbaden, 2008.

[86] Walter Schnyder. Embedding planar graphs on the grid. In *Proceedings of the First Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '90, pages 138–148. Society for Industrial and Applied Mathematics, 1990.

[87] Changxiang Shen, Weifang Cheng, Xiangke Liao, and Shaoliang Peng. Barrier coverage with mobile sensors. In *International Symposium on Parallel Architectures, Algorithms, and Networks, 2008. I-SPAN 2008.*, pages 99–104, May 2008.

[88] S. Slijepcevic and M. Potkonjak. Power efficient organization of wireless sensor networks. In *IEEE International Conference on Communications, 2001. ICC 2001.*, volume 2, pages 472–476, 2001.

[89] Ewald Speckenmeyer. *Untersuchungen zum Feedback Vertex Set Problem in ungerichteten Graphen*. PhD thesis, Paderborn, 1983.

[90] Ewald Speckenmeyer. On feedback vertex sets and nonseparating independent sets in cubic graphs. *Journal of Graph Theory*, 12(3), 1988.

[91] Paul R Stein. On a class of linked diagrams, i. enumeration. *Journal of Combinatorial Theory, Series A*, 24(3):357 – 366, 1978.

[92] P.R. Stein and C.J. Everett. On a class of linked diagrams ii. asymptotics. *Discrete Mathematics*, 21(3):309 – 318, 1978.

[93] S. Torquato and Y. Jiao. Dense packings of the platonic and archimedean solids. *Nature*, 460:876–879, 2009.

[94] S. Torquato and Y. Jiao. Exact constructions of a family of dense periodic packings of tetrahedra. *Phys. Rev. E*, 81:041310, Apr 2010.

[95] Salvatore Torquato and Yang Jiao. Organizing principles for dense packings of nonspherical hard particles: Not all shapes are created equal. *Phys. Rev. E*, 86(1):011102, 2012.

[96] G. T. Toussaint. Computing geodesic properties inside a simple polygon. *Rev. Intell. Artific.*, 1989.

[97] Kuan-Chieh Robert Tseng and David Kirkpatrick. On barrier resilience of sensor networks. In Thomas Erlebach, Sotiris Nikoletseas, and Pekka Orponen, editors,

Algorithms for Sensor Systems, volume 7111 of *Lecture Notes in Computer Science*, pages 130–144. Springer Berlin Heidelberg, 2012.

[98] I. Vigan. Packing and Covering a Polygon with Geodesic Disks. Proceedings of the 1st Mexican Conference on Discrete Mathematics and Computational Geometry, MCDMCG13, pages 243–252, 2013.

[99] Y. Villarceau. Théoréme sur le tore. *Nouvelles Annales de Mathématiques*, 7:345–347, 1848.

[100] Bang Wang, Han Xu, Wenyu Liu, and Hui Liang. A novel node placement for long belt coverage in wireless networks. *IEEE Transactions on Computers*, 62(12):2341–2353, 2013.

[101] C. Zong. *Sphere Packings*. Springer, New York, 1999.