COUNTING PROBLEMS FOR GEOMETRIC STRUCTURES: RECTANGULATIONS, FLOORPLANS, AND QUASI-PLANAR GRAPHS

EYAL ACKERMAN

COUNTING PROBLEMS FOR GEOMETRIC STRUCTURES: RECTANGULATIONS, FLOORPLANS, AND QUASI-PLANAR GRAPHS

RESEARCH THESIS

SUBMITTED IN PARTIAL FULFILLMENT OF THE

REQUIREMENTS

FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

EYAL ACKERMAN

SUBMITTED TO THE SENATE OF THE TECHNION — ISRAEL INSTITUTE OF TECHNOLOGY

ELUL, 5766 HAIFA SEPTEMBER, 2006

THIS RESEARCH THESIS WAS SUPERVISED BY PROF. GILL BAREQUET AND PROF. RON Y. PINTER UNDER THE AUSPICES OF THE COMPUTER SCIENCE DEPARTMENT

ACKNOWLEDGMENT

I am grateful to my advisors, Gill Barequet and Ron Pinter, for taking me as an M.Sc. student, for suggesting a problem that evolved into the main part of this thesis, and for their guidance throughout my studies. I also thank them for encouraging me when needed, for believing in me, sometimes more than I believed in myself, and for the always friendly atmosphere in our meetings. I thank Gábor Tardos and Dan Romik who are co-authors of works presented in this thesis.

I am in debt to Rom Pinchasi for so many things, I don't know where to begin. His course on combinatorial geometry introduced me to this exciting field of research and to some of its interesting problems and beautiful solutions. Joint discussions we had on quasi-planar graphs have lead to the main result of this thesis. I also thank Rom for sponsoring a visit to the states, and János Pach for being my host in New-York during this visit. I hope to keep collaborating with Rom and consider him to be a friend and a role-model for a researcher.

During my studies I spent three great months, visiting the theoretical computer science group at the Free University in Berlin. I thank the people of the group for their welcoming and friendly working environment. Special thanks to Christian Knauer, Günter Rote, and Kevin Buchin, with whom I enjoyed working and hope to collaborate in the future, and to Kevin and Maike Buchin for their kind hospitality and our joint trips. I also thank the Marie Curie program for financing this visit.

I thank the members of the Center for Graphics and Geometric Computing in the Technion, especially Debbie Miller for her administrative and linguistic help during my studies. I am grateful to Oren Ben-Zwi for many discussions and for being the good friend he is. I will miss the interesting coffee break talks I had with Dudu Amzallag. Finally, I wish to express my love and gratitude to Ofrit for her unconditional love and support.

THE GENEROUS FINANCIAL HELP OF THE TECHNION IS GRATEFULLY ACKNOWLEDGED

Contents

List of Figures

List of Tables

Abstract

The work presented in this thesis lies within the field of *Combinatorial Geometry*. This field, roughly speaking, deals with questions of enumeration, existence, and properties of discrete geometric structures involving geometric objects such as points, lines, polygons, etc.

Given a well-defined set of combinatorial structures, a natural question to ask, is: "how large is this set?" We first study this problem for *rectangulations* of a set of points in the plane. Given a set P of n points within a rectangle R in the plane, a *rectangulation* of P is a subdivision of R into smaller rectangles by axis-parallel segments such that every point is on a different segment. We show that when the relative order of the points forms a *separable* permutation, the number of rectangulations is exactly the number of Baxter permutations on $n+1$ elements. We also show that regardless of the order of the points, the number of *guillotine* rectangulations is always the nth Schröder number, and the total number of rectangulations is $O(18^n/n^4)$.

A related combinatorial structure we have investigated are *floorplans* — rectangular partitions with no point constraints. A *floorplan* represents the relative location of modules on an integrated circuit. It is known that the number of floorplans equals the number of Baxter permutations.We present a simple and efficient *bijection* between Baxter permutations and floorplans, which also suggests an enumeration of floorplans

according to various structural parameters.

We also consider a generalization of floorplans to three and higher dimensions. The special case of *guillotine* partitions is analyzed, and we present both an exact summation formula for their number, as well as an analysis of its asymptotic behaviour. Such partitions play an important role in many research areas and application domains, e.g., computational geometry, computer graphics, and solid modeling, to mention just a few.

Finally, we study the following problem in *extremal* combinatorial geometry. A *topological graph* is called *k-quasi-planar* if it does not contain k pairwise-crossing edges. It is conjectured that for every fixed k, the maximum number of edges in a k-quasi-planar graph on n vertices is $O(n)$. We provide, for the first time, an affirmative answer for the case $k = 4$. We also give the simplest proof and the best upper bound known for the maximum number of edges in 3-quasi-planar graphs on n vertices. Moreover, we show a tight upper bound for 3-quasi-planar graphs in which every pair of edges meet at most once.

List of Symbols

$\mathrm{B}(n)$	the number of Baxter permutations on n elements
C_n	the nth Catalan number
$cr(G)$	the crossing number of a graph G
$f_k(n)$	the maximum number of edges in a k-quasi-planar graph on n vertices
$g_d(n)$	the number of guillotine partitions by n hyperplanes in d dimensions
$P(n)$	the number of 3-dimensional partitions by n rectangles
r_n	the nth Schröder number
S_n	the symmetric group (the group of all permutations on n symbols)
$\Gamma(n)$	the number of guillotine rectangulations of a set of n points
$\Xi(P)$	the set of rectangulations of a set of points P

Chapter 1

Introduction

In this thesis we study several counting problems related to geometric structures. Such problems are interesting both from a theoretic point of view, and from practical reasons, as they sometimes play a role in evaluating the performances of algorithms and data structures.

The first problem, discussed in Chapter 2, concerns the number of different *floorplans*. These structures represent the relative positions between modules of an integrated circuit. During the design of an integrated circuit, one determines the shape, size, and position on chip of every module. The shape of the chip and of each of the modules (blocks) is usually a rectangle. A floorplan describes the relative positions of the blocks, thus, it is often represented by a partition (dissection) of a rectangle by non-intersecting segments. A special kind of floorplans are *slicing* (or *guillotine*) floorplans, in which the subdivision to rectangles can be obtained by recursively cutting a rectangle either vertically or horizontally.

It is known [64] that the number of floorplans and slicing floorplans with n blocks equals the number of *Baxter* and *separable* permutations on n elements, respectively.

These permutations can be characterized by the absence of certain subpermutations. Our contribution is a simple and direct bijection between floorplans and Baxter permutations. This bijection also maps slicing floorplans to separable permutations, and—combined with known results about Baxter permutations—suggests an enumeration of floorplans according to various parameters. The results presented in Chapter 2 also appear in [2].

In Chapter 3 we study a generalization of floorplans to three and higher dimensions, namely, partitions of a 3-dimensional (or d-dimensional) box into smaller boxes. Guillotine partitions, a generalization of guillotine (slicing) floorplans, are a special type of 3-dimensional box partitions. The hierarchical structure of guillotine partitions is useful in many areas, such as integrated circuit layout [34] (where $d = 2$) and approximation algorithms in computational geometry [21, 25, 27, 40]. Guillotine partitions are also the underlying structure of orthogonal *binary space partitions* (BSPs) which are widely used in computer graphics (e.g., for hidden-surface removal [23] and shadow generation [18]), solid modeling [60], motion planning [12], etc. By analyzing the combinatorial properties of guillotine partitions we show that $g_d(n)$, the number of guillotine partitions in d dimensions by n hyperplanes, is $\frac{1}{n}\sum_{k=0}^{n-1}\binom{n}{k}\binom{n}{k+1}(d-1)^k d^{n-k}$. Then, we analyze the asymptotic behavior of $g_d(n)$ and prove that it is $\Theta\Big(\big(2d-1+2\sqrt{d(d-1)}\big)^n/n^{3/2}\Big)$ for a fixed value of d. In fact, our analysis provides a rather accurate estimate of $g_d(n)$. These results appear in [4]. We also consider partitions of a 3-dimensional box into smaller boxes by rectangular "walls." Based on some structural properties of such partitions we describe a way of enumerating them.

In Chapter 4 we consider another constrained variant of rectangular partitions in the plane: Given a set P of n points in general position within a rectangle R,

a *rectangulation* of P is a subdivision of R into smaller rectangles by axis-parallel segments such that every point is on a different segment. Finding a rectangulation that minimizes the sum of lengths of the segments has attracted some attention in the literature [15, 17, 21, 26, 27, 28, 35, 36]. We present efficient methods for generating and enumerating the set of distinct rectangulations of a pair (R, P). We also show that when the relative order of the points in P forms a separable permutation, the number of rectangulations is exactly the number of Baxter permutations on $n + 1$ elements, which is $\Theta(8^n/n^4)$. For point sets in an arbitrary permutation we prove an upper bound of $O(18^n/n^4)$ on the number of rectangulations, using a novel technique of Sharir and Welzl [55]. Some of the results presented in this chapter appear in [3].

In the next chapter, Chapter 5, we study a problem from *extremal* combinatorial geomety. A *geometric graph* is a graph drawn in the plane, such that its vertices are distinct points and its edges are straight-line segments connecting the corresponding points and containing no other vertex. A *topological graph* is defined in a similar way with the exception that the edges can take the form of any simple (i.e., non-self-intersecting) curve (*Jordan arc*). Crossings are allowed, but we assume that each pair of edges has a finite number of common internal points and they properly cross at each such point. (However, these usual restrictions are not important in this work.) A topological graph is *simple* if every pair of its edges intersect at most once (either at a vertex or at a crossing point). Clearly, every geometric graph is a simple topological graph.

In a *Turán-type* problem on topological graphs, one asks for the maximum number of edges in a topological graph on n vertices avoiding a certain pattern. For example, it is easy to prove, using Euler's formula, that a topological graph avoiding a pair of crossing edges (i.e., a *planar* graph) contains at most $3n - 6$ edges. Many other

Turán-type problems on topological graphs were considered in the literature (see, e.g., [38, 47, 48, 50, 58, 59], and [46] for a survey).

A natural generalization of planar graphs are *k-quasi-planar* graphs. A topological graph is k-quasi-planar if it does not contain a set of k pairwise-crossing edges. According to a rather famous conjecture from the early 1990s, for every fixed k there is a constant C_k such that any k-quasi-planar graph on n vertices has at most $C_k \cdot n$ edges. Since 2-quasi-planar graphs are planar graphs, the conjecture trivially holds for $k = 2$ with $C_2 = 3$. The case $k = 3$ was settled for simple topological graphs by Agarwal, Aronov, Pach, Pollack, and Sharir [6]. Later, Pach, Radoičić, and Tóth [48] simplified the proof of Agarwal et al. [6], generalized it to non-simple graphs, and showed that 3-quasi-planar graphs on n vertices have at most $65n$ edges. The best upper bounds on the number of edges in k-quasi-planar graphs, $k > 3$, were $O(n \log n)$ for geometric graphs [62], $O(n \log^{2k-6} n)$ for simple topological graphs [6], and $O(n \log^{4k-12} n)$ for general topological graphs [48]. We provide the first proof that 4-quasi-planar graphs have a linear number of edges. We also give the simplest proof and the best upper bound known for the maximum number of edges in 3-quasi-planar graphs on n vertices. Moreover, we show a tight upper bound of $6.5\, n - O(1)$ for simple 3-quasi-planar graphs. The results presented in this chapter also appear in [1, 5].

We end every chapter with a discussion on the results presented in that chapter and related open problems. More concise conclusions can be found in Chapter 6.

Chapter 2

Floorplans

2.1 Introduction

During the physical design process of an integrated circuit, one determines the shape, size, and position on chip of every module. The shape of the chip and that of each of the modules (blocks) is usually a rectangle. A *floorplan* describes the relative positions of the blocks, thus it is often represented by a partition (dissection) of a rectangle by non-intersecting segments into m rectangles (rooms) such that there is a one-to-one mapping from the n ($\leq m$) blocks to the rooms. In a *mosaic* floorplan there are no empty rooms, that is, $n = m$. A special kind of mosaic floorplans are *slicing* floorplans (here we follow the definition in [56]) in which the subdivision to rectangles can be obtained by recursively cutting either vertically or horizontally a rectangle into two smaller rectangles. Slicing floorplans can also be characterized as mosaic floorplans that do not contain a 'pin-wheel' structure. See Figure 2.1 for examples of general, mosaic, and slicing floorplans.

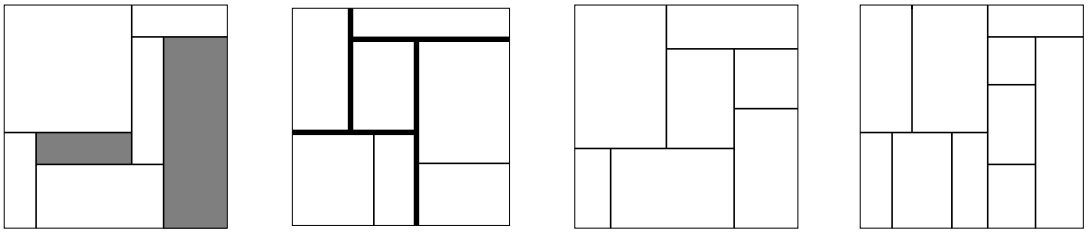

(a) General (empty rooms are shaded)

(b) Mosaic (a clockwise pin-wheel structure is shown in bold)

(c) Mosaic

(d) Slicing

Figure 2.1: Floorplans (b and c are equivalent)
איור 2.1: מפות-תכנון למעגלים מוכללים

Separable and *Baxter* permutations are classes of permutations that can be defined in terms of forbidden subsequences. A separable permutation can be defined as a permutation that does not contain a subsequence of four elements with the same pairwise comparison as 2413 or 3142 (an alternative definition is given in Section 2.2.2). A Baxter permutation has a similar forbidden condition, but it can contain such a subsequence if the absolute difference between the first and last element in the subsequence is greater than one (a more formal definition appears in Section 2.2.1). Thus, separable permutations are a subclass of Baxter permutations.

Sakanushi et al. [51] were the first to consider the number of distinct mosaic floorplans. They found a recursive formula for this number, but did not recognize it to be the same formula suggested by Chung et al. [19] in their analysis of the number of Baxter permutations. Yao et al. [64] showed a bijection between mosaic floorplans and *twin binary trees* whose number is known [22] to be the number of Baxter permutations. They have also shown that the number of slicing floorplans

containing n blocks is the nth Schröder number.

A connection between floorplans and permutations was first presented by Murata et al. [42], who suggested representing floorplans as a pair of permutations (*sequence-pair*). In a later work, Murata et al. [43] described a mapping from sequence-pairs to floorplans. From this mapping one can deduce a mapping from Baxter permutation to mosaic floorplans. Recently and independently, Kajitani [32] has suggested representing a floorplan by a permutation (*single-sequence* in his terminology) and explored, along with others, the properties and advantages of this simple representation [66, 67, 68, 69]. Among other things they showed mappings between (mosaic) floorplans and (Baxter) permutations.

In this chapter we present another bijection between Baxter permutations and mosaic floorplans. This bijection is direct, as opposed to the bijection that can be deduced from the work of Yao et al. [64] and the work of Dulucq and Guibert [22]. The mapping from permutations to floorplans we suggest is much simpler and more efficient than the mapping described in [42]. Comparing with the mappings suggested recently in [66, 67], our mapping is as efficient (has a linear time and space complexity) and at least as simple. Furthermore, the mapping algorithm can easily find the direct neighbors of every block, with performances matching that of the algorithm suggested in [69]. This information is useful for the actual placement of the blocks. The bijection we describe has the following additional merits: First, it maps separable permutations to slicing floorplans. Second, by combining it with known results about Baxter permutations, we obtain enumerations of mosaic floorplans according to various structural parameters, such as the number of vertical segments in the partition and the number of blocks on the boundary of the floorplan.

The rest of this chapter is organized as follows. In Section 2.2 we give some background on Baxter and separable permutations and define an equivalence relation on (mosaic) floorplans. Then we show in Section 2.3 the bijection between Baxter permutations and (equivalence classes of) mosaic floorplans and discuss its applications. In Section 2.4 we explore the enumeration of mosaic floorplans according to various parameters. Finally, we discuss our results in Section 2.5.

2.2 Preliminaries

In order to distinguish between different (mosaic) floorplans we must first define when two floorplans are considered equivalent. Here we follow the definition of Sakanushi et al. [51]. Given a floorplan f a segment s *supports* a room r in f if s contains one of the edges of r. We say that s and r hold a *top-*, *left-*, *right-*, or *bottom-seg-room* relation if s supports r from the respective direction. Two floorplans are equivalent if there is a labeling of their rectangles and segments such that they hold the same seg-room relations. Thus, for example, the floorplans in Figures 2.1(b) and 2.1(c) are equivalent.

2.2.1 Baxter permutations

A Baxter permutation on $[n] = \{1, 2, \ldots, n\}$ is a permutation $\pi = (\sigma_1\sigma_2 \ldots \sigma_n)$ for which there are no four indices $1 \leq i < j < k < l \leq n$ such that

1. $\sigma_k < \sigma_i + 1 = \sigma_l < \sigma_j$; or

2. $\sigma_j < \sigma_l + 1 = \sigma_i < \sigma_k$.

For example, for $n = 4$, 3142 and 2413 are the only non-Baxter permutations. This class of permutations was introduced by Baxter [13] in the context of fixed points of

the composite of commuting functions. The nth Baxter number, $\mathrm{B}(n)$, is the number of Baxter permutations on $[n]$. Chung et al. [19] showed that

$$\mathrm{B}(n) = \sum_{r=0}^{n-1} \frac{\binom{n+1}{r}\binom{n+1}{r+1}\binom{n+1}{r+2}}{\binom{n+1}{1}\binom{n+1}{2}}.$$

Dulucq and Guibert [22] showed one-to-one correspondences between Baxter permutations, *twin binary trees*, and some type of three non-intersecting paths on a grid. Shen et al. [56] analyzed the asymptotic behavior of the Baxter numbers and proved that $\mathrm{B}(n) = \Theta(8^n/n^4)$. The first Baxter numbers (starting from $n = 0$) are $\{0, 1, 2, 6, 22, 92, 422, 2074, \ldots\}$.

2.2.2 Separable permutations

Let $\pi_1 = (\alpha_1, \alpha_2, \ldots, \alpha_n)$ and $\pi_2 = (\beta_1, \beta_2, \ldots, \beta_m)$ be two permutations on $[n]$ and $[m]$, respectively. We say that $\pi = (\sigma_1, \sigma_2, \ldots, \sigma_{n+m})$ is the result of concatenating π_2 *above* π_1 if $\pi_i = \alpha_i$ for $1 \leq i \leq n$ and $\pi_{n+i} = n + \beta_i$ for $1 \leq i \leq m$. Likewise, we say that $\pi = (\sigma_1, \sigma_2, \ldots, \sigma_{n+m})$ is the result of concatenating π_2 *below* π_1 if $\pi_i = m + \alpha_i$ for $1 \leq i \leq n$ and $\pi_{n+i} = \beta_i$ for $1 \leq i \leq m$.

A permutation π is *separable* if either

1. $\pi = (1)$; or

2. There are two separable permutations π_1 and π_2 such that π is the the concatenation of π_2 above or below π_1.

Bose et al. [14] coined the term separable permutations and showed a polynomial-time algorithm for finding a given sub-permutation P within a permutation T, where P is separable. A similar definition was suggested by Shapiro and Stephens [54] in their analysis of permutation-matrices that eventually fill up under *bootstrap percolation*. They have also shown that the number of separable permutations on $[n]$ is the

$(n-1)$st Schröder number[1]. Avis and Newborn [11] showed that separable permutations are exactly the permutations that can be sorted by an unbounded sequence of pop-stacks (in a pop-stack the pop operation unloads the entire stack).

Another characterization of separable permutations is in terms of *forbidden subsequences*. A permutation $\pi = (\sigma_1, \sigma_2, \sigma_3, \ldots, \sigma_n) \in S_n$ *avoids* a certain sub-permutation $\tau \in S_k$ (for $k \leq n$) if it does not contain a subsequence $(\sigma_{i_1}, \sigma_{i_2}, \ldots, \sigma_{i_k})$ with the same pairwise comparisons as τ. The set of permutations on $[n]$ avoiding τ is denoted by $S_n(\tau)$. It can be shown [14] that the set of separable permutations is equal to $S_n(3142, 2413)$, suggesting an alternative proof [63] that their number is the $(n-1)$st Schröder number.

2.3 The Bijection

In this section we show a direct and simple bijection between Baxter permutations and mosaic floorplans. In Section 2.3.1 we describe a mapping from mosaic floorplans to Baxter permutations, while in Section 2.3.2 we present a mapping in the other direction and thus show that these two mappings define a bijection.

2.3.1 Mapping mosaic floorplans to Baxter permutations

In this section we describe a mapping from mosaic floorplans to permutations. It is essentially the same mapping presented implicitly in [51] and explicitly in [67],

[1]The Schröder numbers arise in numerous other enumerative combinatorial problems [57, pp. 239–240]. One example is the number of paths on an orthogonal grid from $(0,0)$ to (n,n) that do not go above the line $y = x$ and use only the steps $(1,0)$, $(0,1)$, and $(1,1)$. When denoting by r_n the nth Schröder number, we have $r_n = \sum_{k=0}^{n} \binom{2n-k}{k} C_{n-k}$, where C_n is the nth Catalan number. It can be shown (see, e.g., [56]) that $r_n = \Theta\left((3+\sqrt{8})^n / n^{1.5}\right)$. The first Schröder numbers (starting from $n = 0$) are $\{0, 1, 2, 6, 22, 90, 394, 1086, 8558, \ldots\}$.

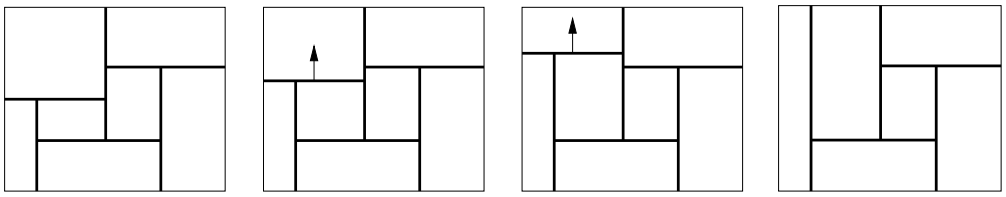

Figure 2.2: Block deletion
איור 2.2: מחיקת בלוק

however, we describe it here for completeness and prove that it always produces a Baxter permutation.

Given a mosaic floorplan of n blocks (rectangles) we can obtain a mosaic floorplan of $n-1$ blocks by using the *block deletion* operation introduced by Hong et al. [30].

Definition 2.3.1 (block deletion) *Let f be a mosaic floorplan with $n > 1$ blocks and let b be the top-left block in f. If the bottom-right corner of b is a '$\dashv$'- (resp., '$\perp$'-) junction, then one can delete b from f by shifting its bottom (resp., right) edge upwards (resp., leftwards), while pulling the T-junctions attached to it until the edge hits the bounding rectangle.*

See Figure 2.2 for an example of the block-deletion operation. Note that we can delete in a similar manner a block from any corner of a floorplan. Using the block-deletion operation we now define a mapping from mosaic floorplans to Baxter permutations.

For example, the permutation that corresponds to the floorplan in Figure 2.3 is 521463. Before we show that the output of Algorithm `FP2BP` is always a Baxter permutation, we need the following definition and observations.

Input: A mosaic floorplan f with n blocks.
Output: A (Baxter) permutation on $[n]$.

Label the blocks of f according to their deletion order from the top-left corner;
Return the permutation of labels obtained by deleting the blocks of f from the bottom-left corner.

Algorithm 1: Algorithm `FP2BP`

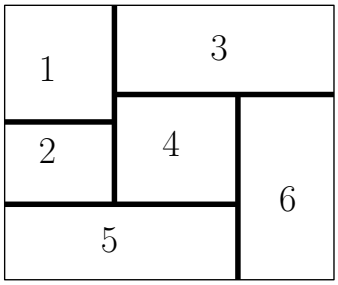

Figure 2.3: Applying Algorithm `FP2BP` on this floorplan yields the permutation 521463

איור 2.3: הפעלת אלגוריתם `FP2BP` על מפת-תכנון זו יוצרת את התמורה 521463

Definition 2.3.2 *Let f be a mosaic floorplan and let b_1 and b_2 be two blocks in f. We say that b_1 is* left of *(resp.,* above*) b_2 if there is either 1. a segment which contains the right (resp., lower) edge of b_1 and the left (resp., upper) edge of b_2; or 2. a block b' such that b_1 is left of (resp., above) b' and b' is left of (resp., above) b_2. If a block b_1 is left-of (resp., above) block b_2 by the first rule, then b_1 is* directly *left-of (resp., above) b_2.*

Observation 2.3.3 ([43, Property 5]) *Let f be a mosaic floorplan and let b_1 and b_2 be two blocks in f. Then exactly one of the following relations holds: b_1 is left of b_2, b_1 is above b_2, b_2 is left of b_1, or b_2 is above b_1.*

Observation 2.3.4 *If a block b_1 precedes a block b_2 according to the top-left corner-deletion order and b_2 precedes b_1 according to the bottom-left corner-deletion order,*

then b_1 is above b_2. Similarly, if b_1 precedes b_2 according to both orders, then b_1 is left of b_2.

Proof: Notice that when a block is deleted from the top-left corner it is to the left or above any other block in the floorplan. Additionally, the relation between any two blocks remains the same after applying the deletion operation. Therefore, if a block b_1 precedes a block b_2 according to the top-left (resp., bottom-left) corner-deletion order, then b_1 is to the left of or above (resp., below) b_2. Hence the claim follows. □

The following observation is easy.

Observation 2.3.5 *If a block b_1 follows immediately a block b_2 according to one of the orders, then there is a segment that contains edges of both b_1 and b_2.*

Next, we prove that Algorithm `FP2BP` always produces a Baxter permutation.

Lemma 2.3.6 *Given a mosaic floorplan f with n blocks, the permutation π obtained by applying Algorithm `FP2BP` on f is a Baxter permutation on $[n]$. Moreover, if f is a slicing floorplan, then π is a separable permutation.*

Proof: Suppose $\pi = (\sigma_1\sigma_2 \dots \sigma_n)$ is not a Baxter permutation. Then there are four indices $1 \leq i < j < k < l \leq n$ such that either 1. $\sigma_k < \sigma_i + 1 = \sigma_l < \sigma_j$; or 2. $\sigma_j < \sigma_l + 1 = \sigma_i < \sigma_k$. Assume that the first case holds, and choose j and k such that $k = j + 1$. According to Observations 2.3.4 and 2.3.5, block σ_i is left of block σ_l, and some segment s_1 supports both blocks. Similarly, block σ_j is below block σ_k, and some segment s_2 supports both blocks. According to Observation 2.3.4, block σ_k is to the left of block σ_l and above block σ_i. Similarly, block σ_j is to the right of block σ_i and below block σ_l. Thus, s_1 and s_2 must intersect (see Figure 2.4). The proof in the second case is similar and is thus omitted.

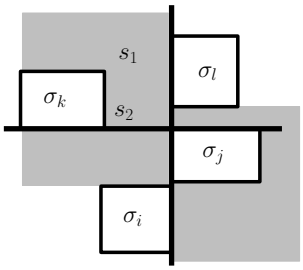

Figure 2.4: An illustration for the proof of Lemma 2.3.6
איור 2.4: הוכחת למה 2.3.6

Now suppose f is a slicing floorplan, and let s be the segment that cuts the bounding rectangle of f into two. Suppose s is horizontal, and denote by f_1 the m blocks above s and by f_2 the $n-m$ blocks below s. Then, the blocks in f_1 precede the blocks of f_2 according to the top-left deletion order, and follow them according to the bottom-left deletion order. By induction, the blocks in f_1 form a separable permutation on $1,\ldots,m$, and the blocks in f_2 form a separable permutation on $m+1,\ldots,n$. Thus, by definition, π is a separable permutation. The proof for the case in which s is vertical is similar. □

Next we show that the mapping defined by Algorithm FP2BP is one-to-one.

Lemma 2.3.7 *Let f_1 and f_2 be two mosaic floorplans, each containing n blocks, and let π_1 and π_2 be the permutations produced by Algorithm FP2BP when applied to f_1 and f_2, respectively. Then, if $f_1 \neq f_2$ then $\pi_1 \neq \pi_2$.*

Proof: We prove by induction on n that if $\pi_1 = \pi_2$ then $f_1 = f_2$. Let b_1 (resp., b_2) be the first block which is removed from the top-left corner of f_1 (resp., f_2), and let s_1 (resp., s_2)be the segment that is shifted in the course of this action. Then s_1 and s_2

must have the same orientation, otherwise the numbers 1 and 2 would have different orders in π_1 and π_2. Let f_1' and f_2' be the resulting floorplans after the deletion. The permutation that corresponds to f_1' and f_2' is the permutation obtained from π_1 by deleting the number 1 and decreasing every remaining number by 1. Thus, by the induction hypothesis, $f_1' = f_2'$. It remains to verify that when we reverse the deletion operation, then the same number of blocks are "pushed" by s_1 and s_2. Indeed, if this number is different then there is a block x which is pushed in one floorplan, say f_1, but not on the other floorplan. Thus, x is to the left of block 1 in f_1 while it is below block 1 in f_2. It follows that 1 and x will have different orders in π_1 and π_2. □

2.3.2 Mapping Baxter permutations to mosaic floorplans

Input: A Baxter permutation $\pi = (\sigma_1\sigma_2 \dots \sigma_n)$.
Output: A mosaic floorplan with n blocks.

```
Draw a block and name it σ_1;
Construct an n × n grid within the block;
for i = 2 to n do
  if σ_i < σ_{i−1} then
    Slice the top-right block by a horizontal segment at the ith level of the grid;
    Name the new block σ_i;
    while the block σ' to the left of σ_i has a label smaller than σ_i do
      Extend block σ_i leftwards (at the expense of σ');
    end while
  else
    Slice the top-right block by a vertical segment at the ith level of the grid;
    Name the new block σ_i;
    while the block σ' below σ_i has a label greater than σ_i do
      Extend block σ_i downwards (at the expense of σ');
    end while
  end if
end for
```

Algorithm 2: Algorithm `BP2FP`

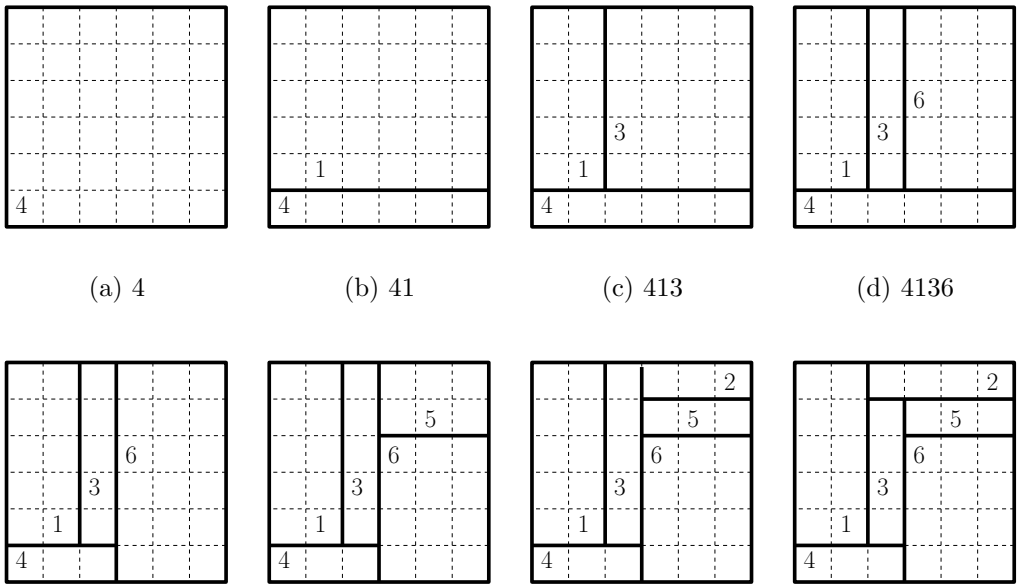

(a) 4 (b) 41 (c) 413 (d) 4136

(e) 4136 (f) 41365 (g) 413652 (h) 413652

Figure 2.5: Applying Algorithm BP2FP to the permutation 413652
איור 2.5: הפעלת אלגוריתם BP2FP על התמורה 413652

Given a Baxter permutation on $[n]$, Algorithm BP2FP constructs a mosaic floorplan with n blocks. See Figure 2.5 for an example. The algorithm simply inserts blocks one by one into the top-right corner of the floorplan. The current block is created by slicing the previous block into two, and is labeled according to the current element in the permutation. If the previous element is smaller (resp., greater) than the current element, then we slice the block vertically (resp., horizontally). The horizontal (resp., vertical) slicing segment is extended leftwards (resp., downwards) as long as the block to the left of it (resp., below) has a smaller (resp., greater) label than the current block.

The output of Algorithm BP2FP is clearly a mosaic floorplan. We show next that

Algorithms BP2FP and FP2BP define a one-to-one correspondence (bijection) between Baxter permutations and mosaic floorplans.

Theorem 2.3.8 *There is a bijection between Baxter permutations on* $[n]$ *and mosaic floorplans with* n *blocks. Moreover, it remains a bijection when restricted to separable permutations on* $[n]$ *and slicing floorplans with* n *blocks.*

Proof: Let π be a Baxter permutation on $[n]$, and let f be the output of Algorithm BP2FP when it is applied on π. Clearly, f is a valid mosaic floorplan containing n blocks. Let π' be the output of Algorithm FP2BP applied to f. To prove the theorem it is enough to show that $\pi' = \pi$. It is easy to see that during the computation of π', the blocks are deleted from the bottom-left corner of f in the same order they were inserted to the top-right corner of f in the course of Algorithm BP2FP. Therefore, it is sufficient to prove that the order in which the blocks of f are deleted from the top-left corner is $1, 2, \ldots, n$. It is clear that the block labeled 1 is the first removed block (no other block is above or to the left of it). Assume that for every $1 \leq i \leq k$ the block labeled i is the ith removed block from the top-left corner. We now show that the next deleted block is the one labeled $k+1$. Suppose that $k+1$ precedes k in π, that is, $\pi = (\ldots, k+1, A, B, k, \ldots)$, where A is a (possibly empty) sequence of integers that are greater than $k+1$ and B is a (possibly empty) sequence of integers that are smaller than k. (There are no other options since π is a Baxter permutation.) Figure 2.6(a) shows the floorplan after k was inserted in the course of Algorithm BP2FP. According to the induction hypothesis, all the blocks in B are removed before block k, so when k is removed (from the top-left corner) the left edge of the block labeled $k+1$ is also on the boundary. The bottom-right corner of k is either a '$\vdash$' or '$\perp$' junction. In the first case $k+1$ is clearly the next block to be deleted. For the

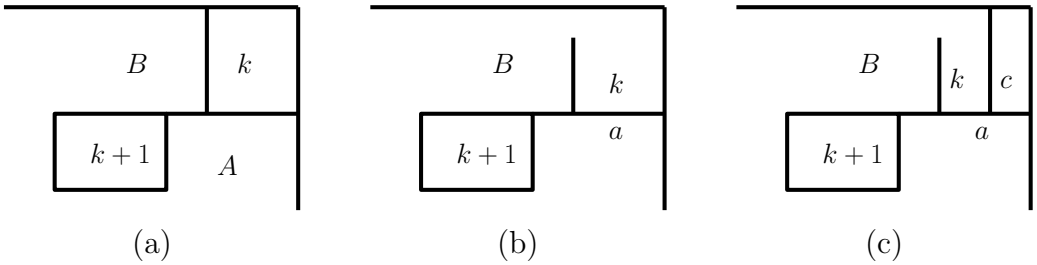

Figure 2.6: Illustration for the proof of Theorem 2.3.8

איור 2.6: הוכחת משפט 2.3.8

second case, note that a '⊥' junction can be formed only when the first block with a label greater than k and to the right of k in π (denote this block by c) has a smaller label than the block below k and sharing the same segment (as c) as a right edge (denote this block by a). Figures 2.6(b,c) illustrate the situation before and after the insertion of c. Note that a is the last of the elements of A and $k < c < a$. If A is empty, then $a = k + 1$; thus, there cannot be such a block c. Otherwise, there must be an integer i such that $k+1 \leq i \leq c-1$, i is to the left of a in π, and $i+1$ is either c or to the right of c. Therefore, $i, a, k, i+1$ form a forbidden subsequence, and π is not a Baxter permutation. The proof for the case k precedes $k+1$ in π is similar and is thus omitted. □

Definition 2.3.9 *Given a floorplan f and a block b in f, the* direct relation set *(DRS) of b is the set of blocks that are directly left of, right of, above, or below b.*

The DRS is important for the actual placement of the blocks on the chip, once their dimensions are set [69].

Theorem 2.3.10 *Let π be a Baxter permutation on $[n]$ and let f be its corresponding floorplan. Then f and the DRS of every block in f can be computed in $O(n)$ time.*

Proof: Algorithm `BP2FP` inserts n blocks one after the other. It is easy to update the DRS of the currently inserted block and of its neighbors while the block is being extended at their expense. For a certain block there could be many update operations, but every block can be chopped at most once from above and at most once from right. Hence, the total number of update operations is $O(n)$. □

Some of the problems which are hard to solve for general or even mosaic floorplans are easier for slicing floorplans, due to their simple structure. Therefore, sometimes one wishes to determine whether a given floorplan is slicing [70]. When a floorplan f is represented as a permutation π, by Theorem 2.3.8 f is slicing if and only if π is separable. Thus, it can be determined in a linear time whether f is a slicing floorplan using the algorithm suggested in [14, pp. 282] to test if a permutation is separable.

2.4 Enumeration of Mosaic Floorplans According to Various Parameters

Baxter permutations are known to be enumerated according to various parameters [22, 37]. Algorithm `BP2FP` along with those results suggest enumerations of mosaic floorplans according to various parameters such as the number of vertical segments, and the number of blocks on the boundary of the floorplan.

Definition 2.4.1 (rise) *Given a permutation $\pi = (\sigma_1\sigma_2\ldots\sigma_n)$, a* rise *(resp.,* descent*) in π is sequence of two consecutive elements $\sigma_i\sigma_{i+1}$ such that $\sigma_i < \sigma_{i+1}$ (resp., $\sigma_i > \sigma_{i+1}$).*

According to Algorithm `BP2FP`, every rise in the input permutation is mapped to a vertical segment in the output floorplan. Mallows [37] considered the enumeration

of Baxter permutations according to the number of rises. The next corollary follows from his result.

Corollary 2.4.2 *The number of mosaic floorplans with n blocks and r vertical segments is*

$$\frac{\binom{n+1}{r}\binom{n+1}{r+1}\binom{n+1}{r+2}}{\binom{n+1}{1}\binom{n+1}{2}}.$$

Definition 2.4.3 (left-to-right minimum/maximum) *Let $\pi = (\sigma_1\sigma_2\ldots\sigma_n)$ be a permutation on $[n]$. An element σ_k is a* left-to-right minimum *(resp.,* maximum*) if $\sigma_k < \sigma_i$ (resp., $\sigma_k > \sigma_i$) for every $1 \leq i < k$.*

Algorithm `BP2FP` maps every left-to-right minimum to a block touching the left edge of the boundary of the output floorplan f. Similarly, every left-to-right maximum in π is mapped to a block touching the bottom edge of the boundary of f. Thus, according to a result of Dulucq and Guibert [22, Theorem 1] we have:

Corollary 2.4.4 *The number of mosaic floorplans with n blocks, r vertical segments, i blocks touching the left edge of the boundary of the floorplan, and s blocks touching the bottom edge of the boundary of the floorplan is*

$$\binom{n+1}{r+1}\frac{si}{n(n+1)}\left(\binom{n-s-1}{n-r-2}\binom{n-i-1}{r-1}-\binom{n-s-1}{n-r-1}\binom{n-i-1}{r}\right).$$

Dulucq and Guibert have considered the enumeration of Baxter permutations according to two other parameters.

Definition 2.4.5 [22, Definition 4] *Given a permutation $\pi = (\sigma_1\sigma_2\ldots\sigma_n)$,*

- *$rd(\pi)$ is the number of rises in $\sigma_i\sigma_{i+1}\ldots\sigma_n$, where $i = \max\{j|\exists k \geq 2 : \sigma_j < \sigma_{j+k} < \sigma_{j+1} < \ldots < \sigma_{j+k-1}\}$ with $\sigma_0 = -1$ and $\sigma_{n+1} = 0$.*

- *$dd(\pi)$ is the number of descents in $\sigma_i\sigma_{i+1}\dots\sigma_n$, where $i = \max\{j|\exists k \geq 2 : \sigma_j > \sigma_{j+k} > \sigma_{j+1} > \dots > \sigma_{j+k-1}\}$ with $\sigma_0 = n+2$ and $\sigma_{n+1} = n+1$.*

We define below the corresponding parameters for a mosaic floorplan f.

Definition 2.4.6 *Let f be a mosaic floorplan and let t be a '$\perp$'- (resp., '$\vdash$'-) junction in f. We say that t is the* last *'$\perp$'- (resp., '$\vdash$'-) junction in f if it is the last '$\perp$'- (resp., '$\vdash$'-) junction which is deleted when the blocks of f are removed from the bottom-left corner. Given a horizontal (resp., vertical) segment s in f, we say that s is* above *(resp.,* left of*) t if s is above the horizontal (resp., vertical) segment*[2] *of t.*

A careful look at the definitions of $rd(\pi)$ and $dd(\pi)$, and the way Algorithm `BP2FP` works leads to the following observation.

Observation 2.4.7 *Let π be a Baxter permutation on $[n]$ and let f be the floorplan produced by the application of Algorithm `BP2FP` to π. Denote by t the last '$\perp$'- (resp., '$\vdash$'-) junction in f. Then, the number of horizontal (resp., vertical) segments in f above (resp., left of) t is $dd(\pi)$ (resp., $rd(\pi)$).*

It follows from this observation and from Theorem 5 in [22] that:

Corollary 2.4.8 *The number of mosaic floorplans with n blocks, r vertical segments, i blocks touching the left edge of the boundary of the floorplan, s blocks touching the bottom edge of the boundary of the floorplan, p vertical segments after the last*

[2] *A segment s_1 is* above *a segment s_2 in a floorplan f if: 1. s_1 is vertical and its lower endpoint is on s_2; or 2. s_2 is vertical and its upper endpoint is on s_1; or 3. s_1 and s_2 are both horizontal and contain opposite edges of a block (rectangle) in f. The relation* left of *is defined in a similar manner.*

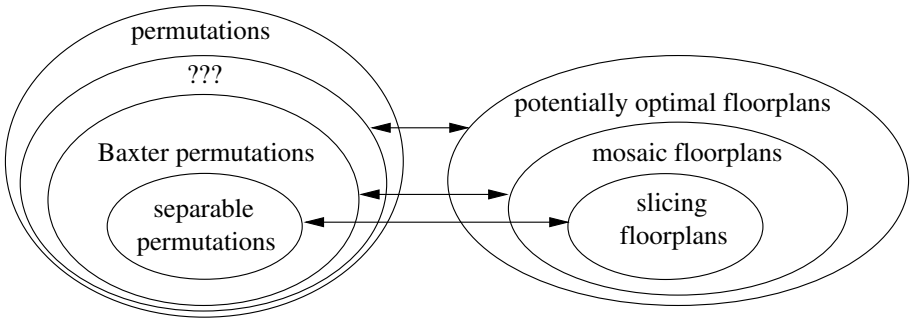

Figure 2.7: The hierarchy of bijections between permutations and floorplans
איור 2.7: הירכית ההתאמות בין תמורות למפות-תכנון

'$\perp$'-junction, and q horizontal segments after the last '$\vdash$'-junction is

$$\begin{vmatrix} \binom{n-1-i-p}{r-p} & \binom{n-1-p}{r-p} & \binom{n-1-s-p}{r-s-p} \\ \binom{n-1-i}{r} & \binom{n-1}{r} & \binom{n-1-s}{r-s} \\ \binom{n-1-i-q}{r} & \binom{n-1-q}{r} & \binom{n-1-s-q}{r-s} \end{vmatrix}.$$

2.5 Discussion

We have presented a bijection between Baxter permutations and mosaic floorplans. Moreover, from this bijection we have also deduced a similar correspondence between separable permutations and slicing floorplans (see Figure 2.7), and have suggested enumerations of mosaic floorplans according to various parameters, such as the number of vertical segments and the number of blocks on certain edges of the boundary of the floorplan. The algorithm we use to map Baxter permutations to mosaic floorplans has applications in integrated circuit (IC) design: it can be used for an easy and efficient construction of a floorplan from the permutation representing it.

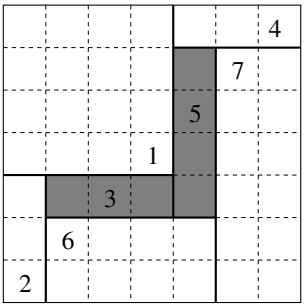

Figure 2.8: The floorplan with empty rooms that corresponds to the non-Baxter permutation 24153

איור 2.8: מפת-תכנון עם חדרים ריקים המתאימה לתמורה 24153

Given a non-Baxter permutation, it is possible to convert it to a Baxter permutation by inserting *dummy* elements, as suggested by Murata et al. [43]. The new permutation can then be mapped to a floorplan containing *empty* rooms (which is a general, non-mosaic floorplan) using Algorithm BP2FP. Given a permutation $\pi = (\sigma_1\sigma_2 \dots \sigma_n)$, the elimination of a forbidden subpattern of the form $\sigma_i \dots \sigma_j\sigma_{j+1} \dots \sigma_k$, such that $\sigma_{j+1} < \sigma_i + 1 = \sigma_k < \sigma_j$ (resp., $\sigma_j < \sigma_k + 1 = \sigma_i < \sigma_{j+1}$), is done by inserting the dummy element σ_k (resp., σ_i) between σ_j and σ_{j+1} and increasing by 1 each of the old elements greater than or equal to σ_k (resp., σ_i). For example, the permutation 2413 is converted to 25**3**14. In the mosaic floorplan that matches the new permutation we mark every block that corresponds to a dummy element as an empty room. Each of these empty rooms is the center of a 'pin-wheel' structure. Figure 2.8 shows an example with two pin-wheels and their corresponding empty rooms.

Finding a floorplan that minimizes criteria such as area or wire-length is a major problem in IC design. It is well-known that an optimal floorplan might contain empty

rooms. However, Young et al. [65] showed that when searching for an optimal floorplan, it is enough to consider floorplans in which every empty room (if such exists) is at the center of pin-wheel structure and has no *room-room* neighbor (that is, a touching room) which is an empty room. Characterizing and enumerating permutations that are mapped to such floorplans is an interesting open problem (as indicated in Figure 2.7).

Chapter 3

Rectangular Partitions in Three and Higher Dimensions

3.1 Introduction

In this chapter we consider the number of partitions of a 3-dimensional box into smaller 3-dimensional boxes. Formally, let B be an axis-aligned 3-dimensional box in $\mathbb{R}^3$. A *partition* (or *subdivision*) of B is a set of interior-disjoint axis-aligned boxes $b_1, b_2, \ldots, b_k$ such that $\bigcup_{i=1}^{k} b_i = B$. As we consider a generalization of (mosaic) floorplans, we do not allow the intersection of two boxes to be a line segment. (This is equivalent to requiring that the intersection of two rectangles in a floorplan not be a single point.) However, the intersection of two boxes can be a single point. Considering the plane portions that induce a partition, we require that each of them be a maximal connected 2-dimensional region of $\bigcup_{i=1}^{k} \partial b_i$.

We classify the partitions into three classes based on the shapes of the plane segments induced by (or inducing) a partition. A *partition-by-rectangles* is a partition

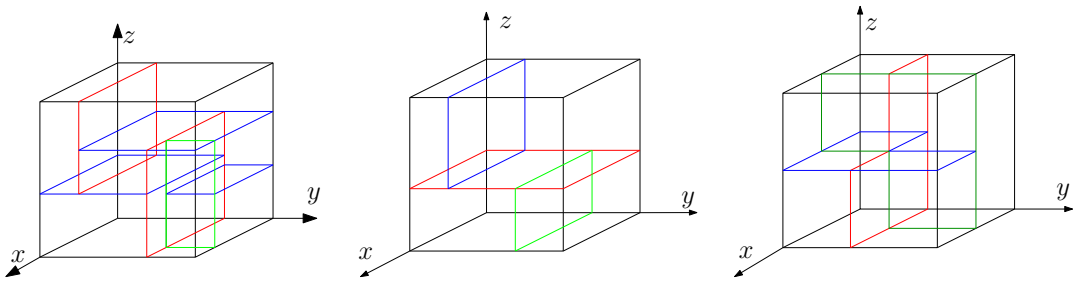

(a) A partition-by-rectangles

(b) A guillotine partition

(c) A general partition

Figure 3.1: 3-dimensional partitions
איור 3.1: חלוקות תלת-מימדיות

in which the plane segments are rectangles (see Figure 3.1(a) for an example). *Guillotine partitions* are a subclass of partitions-by-rectangles, in which the partition can be obtained by recursively cutting a box using a plane into two smaller boxes (see Figure 3.1(b) for an example). Guillotine partitions can be naturally generalized to d dimensions. We call partitions in which some plane segment is not a rectangle *general* partitions (see Figure 3.1(c) for an example).

In Section 3.2 we present an exact summation formula for the number of structurally-different guillotine partitions in d dimensions by n hyperplane segments, and then show that it is $\Theta\left(\left(2d-1+2\sqrt{d(d-1)}\right)^n/n^{3/2}\right)$.[1] In Section 3.3 we prove some properties of partitions-by-rectangles and provide a method to enumerate such partitions. So far we were unable to obtain any meaningful observations about general 3-dimensional partitions.

[1] These results were published in [4].

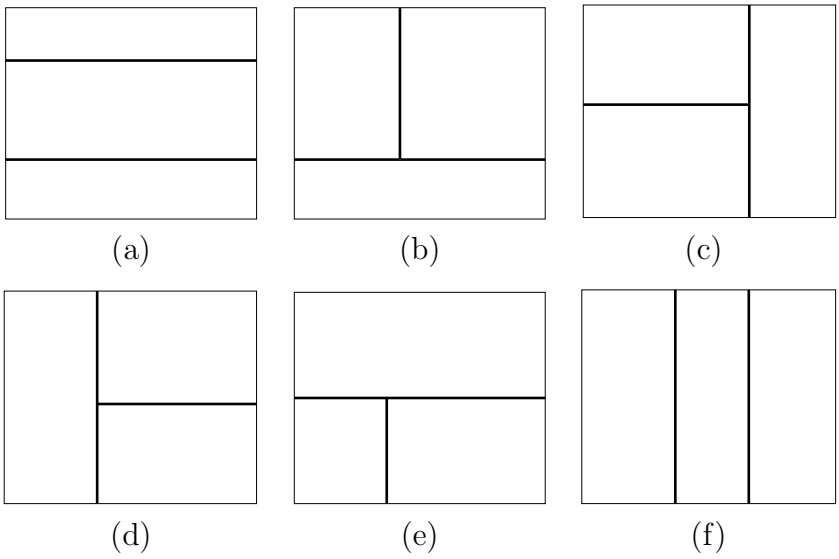

Figure 3.2: Guillotine partitions by two lines in the plane
איור 3.2: חלוקות גליוטינה ע"י שני קווים במישור

3.2 Guillotine Partitions

Given a d-dimensional box B in $\mathbb{R}^d$, a *guillotine partition* of B is a subdivision of B into smaller d-dimensional boxes obtained by first cutting B into two d-boxes by a hyperplane, then recursively cutting these boxes in the same manner (changing the direction of the cut at will). Clearly, there are infinitely-many guillotine partitions with a given number of hyperplanes; however, if we look at the *structure* of the hierarchy that is formed by the hyperplanes, i.e., we care merely about the directions of the cuts rather than their exact positions, then we can count the finitely-many (structurally) different guillotine partitions by n hyperplanes in d dimensions. In this section we are interested in finding this number, denoted $g_d(n)$; to the best of our knowledge, this problem for $d > 2$ has not been investigated to date (see [64] for a discussion of the case $d = 2$). For example, Figure 3.2 lists the $g_2(2) = 6$ possible guillotine partitions by $n = 2$ lines in the plane ($d = 2$). Table 3.1 provides the values

n	$d=2$	$d=3$	$d=4$
0	1	1	1
1	2	3	4
2	6	15	28
3	22	93	244
4	90	645	2,380
5	394	4,791	24,868
6	1,806	37,275	272,188
7	8,558	299,865	3,080,596
8	41,586	2,474,025	35,758,828
9	206,098	20,819,307	423,373,636
10	1,037,718	178,003,815	5,092,965,724
11	5,293,446	1,541,918,901	62,071,299,892
12	27,297,738	13,503,125,805	764,811,509,644
13	142,078,746	119,352,115,551	9,511,373,563,492
14	745,387,038	1,063,366,539,315	119,231,457,692,284
15	3,937,603,038	9,539,785,668,657	1,505,021,128,450,516
16	20,927,156,706	86,104,685,123,025	19,112,961,439,180,588
17	111,818,026,018	781,343,125,570,515	244,028,820,862,442,116
18	600,318,853,926	7,124,072,211,203,775	3,130,592,301,487,969,948
19	3,236,724,317,174	65,233,526,296,899,981	40,333,745,806,536,135,028
20	17,518,619,320,890	599,633,539,433,039,445	521,655,330,655,122,923,980

Table 3.1: First values of $g_d(n)$ for $d = 2, 3, 4$ and $n \leq 20$. The series $g_2(n)$ is the sequence of Schröder numbers [64].

טבלה 3.1: הערכים הראשונים של $g_d(n)$ עבור $d = 2, 3, 4$ ו-$n \leq 20$

of $g_d(n)$ for $d = 2, 3, 4$ and $n \leq 20$.

The hierarchical structure of guillotine partitions is useful in many areas, such as integrated circuit layout [34] (where $d = 2$) and approximation algorithms in computational geometry [21, 25, 27, 40]. Guillotine partitions are also the underlying structure of orthogonal *binary space partitions* (BSPs) which are widely used in computer graphics (e.g., for hidden-surface removal [23] and shadow generation [18]), solid modeling [60], motion planning [12], etc. Two-dimensional guillotine partitions are equivalent to what is known in integrated circuit layout as *slicing floorplans* (see

Chapter 2).

Recall that $g_d(n)$ denotes the number of guillotine partitions of a d-dimensional box in $\mathbb{R}^d$ by n hyperplanes. By analyzing the combinatorial properties of guillotine partitions, we show that $g_d(n) = \frac{1}{n}\sum_{k=0}^{n-1}\binom{n}{k}\binom{n}{k+1}(d-1)^k d^{n-k}$. Then, we analyze the asymptotic behavior of $g_d(n)$ and prove that it is $\Theta\Big(\big(2d-1+2\sqrt{d(d-1)}\big)^n/n^{3/2}\Big)$ for a fixed value of d. In fact, our analysis provides a rather accurate estimate of $g_d(n)$.

The rest of this section is organized as follows. In Subsection 3.2.1 we compute a binomial-sum formula for $g_d(n)$ and provide its generating function, while in Subsection 3.2.2 we obtain an asymptotic value for $g_d(n)$ directly as well as by using the generating function.

3.2.1 The exact number of guillotine partitions

In this section we derive a recursive formula, a binomial-sum formula, and a generating function for $g_d(n)$, the number of guillotine partitions by n hyperplanes in $\mathbb{R}^d$. We first define formally the notion of *structurally-different* guillotine partitions.

Let B be an axis-parallel d-dimensional box in $\mathbb{R}^d$. A *partition* (or *subdivision*) of B is a set S of $k > 0$ interior-disjoint axis-parallel boxes $b_1, b_2, \ldots, b_k$ whose union equals B. A partition $S = \{b_1, b_2, \ldots, b_k\}$ of B is *guillotine* if $k = 1$ or there are a hyperplane h and two disjoint non-empty subsets $S_1, S_2 \subset S$ such that:

1. h splits B into two interior-disjoint boxes B_1 and B_2;
2. S_1 is a guillotine partition of B_1; and
3. S_2 is a guillotine partition of B_2.

Clearly, the hyperplane h must be orthogonal to some axis x_i. We will assume without loss of generality that the interior of B_1 is below h with respect to x_i, and the interior

of B_2 is above h. The definition of a guillotine partition implies a method by which one can obtain the partition of the box B into the small boxes $b_1, b_2, \ldots, b_k$: if $k = 1$ then do nothing, otherwise cut B into two boxes B_1 and B_2 according to h, then continue by cutting recursively B_1 and B_2. One way to describe this cutting procedure is by a binary tree: The tree is a single node (and no edges) in case the d-dimensional box is not partitioned. Otherwise, the root of the tree contains the details of the hyperplane h; the left child of the root is the binary tree that corresponds to the guillotine partition of B_1; and the right child of the root is the binary tree that corresponds to the guillotine partition of B_2.

A guillotine partition can be represented by several trees: For example, an $n \times 1$ two-dimensional box (a rectangle) that is partitioned into n 1×1 boxes (squares) can be represented by C_{n-1} different trees (where C_n is the nth Catalan number). Another example is the two different trees representing a 2×2 two-dimensional box (square) partitioned into four 1×1 boxes (squares). For a *canonical* tree representation of a guillotine partition S of a box B, consider all the hyperplanes h that split B into two boxes B_1 and B_2 with the respective subpartitions $S_1, S_2 \subset S$ such that S_1 (respectively S_2) is a guillotine partition of B_1 (respectively B_2). Among these hyperplanes consider only those that are orthogonal to the axis x_i with the smallest index i, and among them choose the one which is above all the others with respect to x_i. The canonical tree will have this hyperplane as its root and the canonical trees representing guillotine partitions of the resulting two sub-boxes as children.

The hierarchy structure of a guillotine partition is then the canonical tree representing this partition, in which each node that corresponds to a hyperplane h only records the index i such that h is orthogonal to the axis x_i. We refer to this tree as the *structure-tree* of the guillotine partition.

Definition 3.2.1 *Two guillotine partitions by n hyperplanes of a d-dimensional box in $\mathbb{R}^d$ are* structurally equivalent *if they have the same structure-tree representation.*

Clearly, the number of different guillotine partitions of a box $B \subset \mathbb{R}^d$ by n hyperplane does not depend on the dimensions of B. Hence, the notation $g_d(n)$ is used for this number.

Observation 3.2.2 *Let $\mathcal{T}_d(n)$ be the set of binary trees with n nodes, such that every node has a label $\ell \in \{1, 2, \ldots, d\}$ and the label of every right child is different from the label of its parent node. Then $|\mathcal{T}_d(n)| = g_d(n)$.*

Proof: Clearly, two (structurally) different guillotine partitions are represented by two different binary trees such that each node in a tree has a label from $\{1, 2, \ldots, d\}$. Additionally, it follows from the canonical tree representation of a guillotine partition that a structure-tree of a guillotine partition cannot have a node whose right child has the same label as its parent. Therefore, $g_d(n) \leq |\mathcal{T}_d(n)|$. On the other hand, given $t \in \mathcal{T}_d(n)$, one can easily construct a guillotine partition of a d-dimensional box B whose structure-tree is t: Let r be the root of t and let i be the label of r. Cut B into two equal halves by a hyperplane h orthogonal to x_i. Then, cut recursively the half of B below h according to the subtree whose root is the left child of r, and similarly cut the half of B above h according to the subtree whose root is the right child of r. Thus, $|\mathcal{T}_d(n)| \leq g_d(n)$. □

From this observation we can deduce that

$$g_d(n) = \sum_{k=0}^{n-1} N_{n,k}(d-1)^k d^{n-k},$$

where $N_{n,k}$ is the number of binary trees with n nodes, k of which are right children. ($N_{n,k} = \frac{1}{n}\binom{n}{k}\binom{n}{k+1}$ are the well-known *Narayana* numbers; see, for example, [24,

pp. 1033–1034].) This follows easily from the fact that every node that is a right child has one of $d-1$ possible labels, while every other node has one of d possible labels. Thus, we obtain the following binomial-sum formula for the number of guillotine partitions in d dimensions:

Theorem 3.2.3

$$g_d(n) = \frac{1}{n}\sum_{k=0}^{n-1}\binom{n}{k}\binom{n}{k+1}(d-1)^k d^{n-k}.$$

Observation 3.2.2 also yields a recursive formula for $g_d(n)$:

$$g_d(n) = d\cdot\left(g_d(n-1) + \sum_{k=1}^{n-1}\frac{d-1}{d}g_d(k)g_d(n-1-k)\right), \qquad g_d(0)=1, \tag{3.1}$$

where k represents the number of nodes in the right subtree of the root, and the number of these subtrees is $\frac{d-1}{d}\,g_d(k)$ since the root of such a subtree cannot have the same label as its parent node. By rearranging terms is Equation (3.1), we have:

$$g_d(n) = g_d(n-1) + (d-1)\sum_{k=0}^{n-1} g_d(k)g_d(n-1-k), \qquad g_d(0)=1.$$

From this recursive formula we can also easily compute $f_d(z)$, the generating function of $g_d(n)$:

$$\begin{aligned} f_d(z) &= \sum_{n=0}^{\infty} g_d(n)z^n \\ &= 1 + \sum_{n=1}^{\infty}\left(g_d(n-1) + (d-1)\sum_{k=0}^{n-1} g_d(k)g_d(n-1-k)\right)z^n \\ &= 1 + \sum_{n=1}^{\infty} g_d(n-1)z^n + (d-1)\sum_{n=1}^{\infty}\sum_{k=0}^{n-1} g_d(k)g_d(n-1-k)z^n. \end{aligned} \tag{3.2}$$

It is readily seen that

$$\sum_{n=1}^{\infty} g_d(n-1)z^n = z\sum_{n=1}^{\infty} g_d(n-1)z^{n-1} = zf_d(z) \tag{3.3}$$

and

$$\sum_{n=1}^{\infty}\sum_{k=0}^{n-1} g_d(k)g_d(n-1-k)z^n = z\sum_{n=0}^{\infty}\sum_{k=0}^{n} g_d(k)g_d(n-k)z^n = z(f_d(z))^2. \tag{3.4}$$

By substituting Equations (3.3) and (3.4) in Equation (3.2), we obtain

$$f_d(z) = zf_d(z) + (d-1)z(f_d(z))^2 + 1. \tag{3.5}$$

One solution of Equation (3.5) is spurious, and, thus, we get

$$f_d(z) = \frac{1 - z - \sqrt{z^2 - 2(2d-1)z + 1}}{2(d-1)z}.$$

Aside from the intrinsic interest of this generating function, knowing the generating function of a combinatorial sequence is useful for obtaining information on the asymptotics of the sequence by the standard technique (which has been implemented in software) of analyzing the singularities of the function—see Remark 2 in Section 3.2.2. Generating functions are also a useful tool for manipulating combinatorial sequences, and for showing that they satisfy various identities, recurrence equations, etc.

3.2.2 The asymptotic number of guillotine partitions

In this section we compute the asymptotic number of guillotine partitions.

Theorem 3.2.4 *For every $d \in \mathbb{N}$, as $n \to \infty$*

$$g_d(n) = (1+o(1))\frac{\sqrt{2d(d-1) + (2d-1)\sqrt{d(d-1)}}}{2(d-1)\sqrt{\pi}} \cdot \frac{\left(2d-1+2\sqrt{d(d-1)}\right)^n}{n^{3/2}}. \tag{3.6}$$

Proof: Let $g_d(n) = \sum_{k=0}^{n-1} b_d(n,k)$, where

$$b_d(n,k) = \frac{d^n (n!)^2 ((d-1)/d)^k (n-k)}{(k!)^2((n-k)!)^2 n(k+1)}.$$

By Stirling's formula, $m! = (1 + O(m^{-1}))\sqrt{2\pi m}(m/e)^m$. Therefore,

$$b_d(n,k) = (1 + O(k^{-1} + (n-k)^{-1}))\frac{n}{2\pi k^2}\exp(nQ_d(k/n)), \tag{3.7}$$

where $Q_d : [0,1] \to [0,\infty)$ is the function defined by

$$Q_d(t) = \log d + t\log\frac{d-1}{d} - 2t\log t - 2(1-t)\log(1-t).$$

Let us analyze the properties of Q_d. Since

$$Q_d'(t) = \log\frac{d-1}{d} + 2\log\frac{1-t}{t} = \log\frac{(d-1)(1-t)^2}{dt^2},$$

it follows easily that Q_d has a unique maximum at $t = t_d := \sqrt{d(d-1)} - (d-1)$, with the value

$$Q_d(t_d) = \log(2d - 1 + 2\sqrt{d(d-1)})$$

and second derivative

$$v_d := Q_d''(t_d) = \frac{-2}{1-t_d} - \frac{2}{t_d} = -2 \cdot \frac{2d(d-1) + (2d-1)\sqrt{d(d-1)}}{d(d-1)}.$$

In other words, one has the Taylor expansion

$$Q_d(t) = \log(2d - 1 + 2\sqrt{d(d-1)}) - \tfrac{2d(d-1)+(2d-1)\sqrt{d(d-1)}}{d(d-1)}(t-t_d)^2 + O((t-t_d)^3),$$

which is valid in some neighborhood of t_d.

Let $\varepsilon_n = n^{-5/12}$, and write

$$g_d(n) = \sum_{k=0}^{n-1} b_d(n,k) = \sigma_1 + \sigma_2,$$

where

$$\sigma_1 = \sum_{|k - t_d n| > n\varepsilon_n} b_d(n,k), \qquad \sigma_2 = \sum_{|k - t_d n| \le n\varepsilon_n} b_d(n,k).$$

First, we show that the contribution of σ_1 to $g_d(n)$ is negligible. Indeed, if $k \approx tn$, then

$$\frac{b_d(n,k-1)}{b_d(n,k)} = \frac{d(n-k-1)k^3}{(d-1)(n-k)(k+1)(n-k+1)^2} \approx \frac{dt^2}{(d-1)(1-t)^2}.$$

Since t_d has precisely the property that $dt^2/((d-1)(1-t)^2) < 1$ for $t < t_d$ and $dt^2/((d-1)(1-t)^2) > 1$ for $t > t_d$, this implies in particular that $b_d(n,k-1) < b_d(n,k)$ for all $k < (t_d - \varepsilon_n)n$, and $b_d(n,k-1) > b_d(n,k)$ for all $k > (t_d + \varepsilon_n)n$, and, therefore,

$$\begin{aligned}
\sigma_1 &\leq n \cdot \max\Big(b_d(n, \lfloor (t_d - \varepsilon_n)n \rfloor), b_d(n, \lceil (t_d + \varepsilon_n)n \rceil)\Big) \\
&\leq O\Big[\exp(n \max(Q_d(t_d - \varepsilon_n), Q_d(t_d + \varepsilon_n)))\Big] \\
&= (2d - 1 + 2\sqrt{d(d-1)})^n O\Big[\exp\Big(-v_d n \varepsilon_n^2 + O(n\varepsilon_n^3)\Big)\Big] \\
&= o\Big((2d - 1 + \sqrt{d(d-1)})^n / n^{3/2}\Big).
\end{aligned}$$

Second, turn to σ_2, which contributes the bulk of the value of $g_d(n)$. Use Equation (3.7), noting that in this range of values of k, $O(k^{-1} + (n-k)^{-1}) = O(n^{-1})$, uniformly, and denoting $k = t_d n + u\sqrt{n}$:

$$\begin{aligned}
\sigma_2 &= (1 + O(n^{-1})) \sum_{|k - t_d n| < n\varepsilon_n} \frac{1}{2\pi(t_d + u/\sqrt{n})^2 n} \exp(nQ_d(t_d + u/\sqrt{n})) \\
&= (1 + O(n^{-1} + \varepsilon_n)) \sum_{|k - t_d n| < n\varepsilon_n} \frac{\big(2d - 1 + 2\sqrt{d(d-1)}\big)^n}{2\pi t_d^2 n} \cdot \exp(-v_d u^2 + O(n\varepsilon_n^3)) \\
&= (1 + O(n^{-1} + n^{-5/12} + n^{-1/4})) \frac{\big(2d - 1 + 2\sqrt{d(d-1)}\big)^n}{2\pi t_d^2 n^{3/2}} \int_{-\infty}^{\infty} e^{-v_d u^2} du \\
&= (1 + o(1)) \frac{\big(2d - 1 + 2\sqrt{d(d-1)}\big)^n}{2\pi t_d^2 n^{3/2}} \cdot \frac{\sqrt{\pi}}{\sqrt{v_d}}.
\end{aligned}$$

This, upon some simple manipulations, gives Equation (3.6). □

Remarks:

1. More careful estimates can be used to improve the $(1+o(1))$ term to $(1+O(n^{-1}))$. A complete asymptotic expansion in powers of n^{-1} can also be obtained.

2. Relation (3.6) can be obtained from the generating function $f_d(z)$ using standard saddle-point techniques, as described, e.g., in [44]. The Maple package `gdev` [52] produces such asymptotic estimates automatically. Upon loading the package and typing the command

```
> equivalent(1-z-sqrt(z^2-2*(2*d-1)*z+1))/(2*(d-1)*z),z,n);
```

   one obtains an output, which after some reformatting and simplification, is seen to be equivalent to Equation (3.6).

When d is considered constant, Equation (3.6) readily yields:

Corollary 3.2.5 $g_d(n) = \Theta\left(\frac{\left(2d-1+2\sqrt{d(d-1)}\right)^n}{n^{3/2}}\right)$.

For example, there are $\Theta((3+2\sqrt{2})^n/n^{3/2}) \approx \Theta(5.828^n/n^{3/2})$ (resp., $\Theta((5+2\sqrt{6})^n/n^{3/2})$ $\approx \Theta(9.900^n/n^{3/2})$ structurally-different guillotine partitions with n lines (resp., planes) in the plane (resp., in 3-space).

3.3 Partitions by Rectangles

In this section we consider partitions of 3-dimensional boxes into smaller boxes such that the "walls" inducing the partition are rectangles. We call such partitions *partitions-by-rectangles*. We explore some of the properties of these partitions and suggest a method for enumerating them.

In order to count the number of different partitions-by-rectangles we must define when two such partitions are equivalent. We use a definition similar to the definition of equivalent floorplans (see Section 2.2). Let b be a box in a partition-by-rectangles and let r be a rectangle induced by the partition such that r is parallel to the xy plane. We say that r and b have an *xy-before rectangle-box relation* if r supports b from below. In a similar way, we define *xy-after*, *yz-before*, *yz-after*, *xz-before*, and *xz-after* rectangle-box relations. Using these definitions, we define two partitions-by-rectangles as *equivalent* if they hold the same rectangle-box relations up to renaming. Furthermore, by slightly "sliding" r parallel to itself while extending or shrinking accordingly the boxes supported by r, we maintain an equivalent partition-by-rectangles. Thus, we can assume, without loss of generality, that there are no two rectangles in a partition-by-rectangles that lie in the same plane.

3.3.1 Properties of partitions-by-rectangles

We prove most of the properties of partitions-by-rectangles by induction on the number of boxes in the partition, so let us begin by describing a *box deletion* operation (which is similar to the block-deletion operation for mosaic floorplans described in Section 2.3.1). Let P be a partition-by-rectangles into n boxes, and let b be a box touching one of the eight corners of the bounding box of P. Suppose that one of the rectangles defining P is also a face of b, and denote this rectangle by r. If we 'slide' r parallel to itself towards the face of the bounding box, while extending every other box supported by r, we maintain a valid partition-by-rectangles. If this process is continued until r coincides with a face of the bounding box, the result is a partition-by-rectangles into $n - 1$ boxes (since b was deleted). Next, we show that such a process is always possible. To this aim we need the following easy observation.

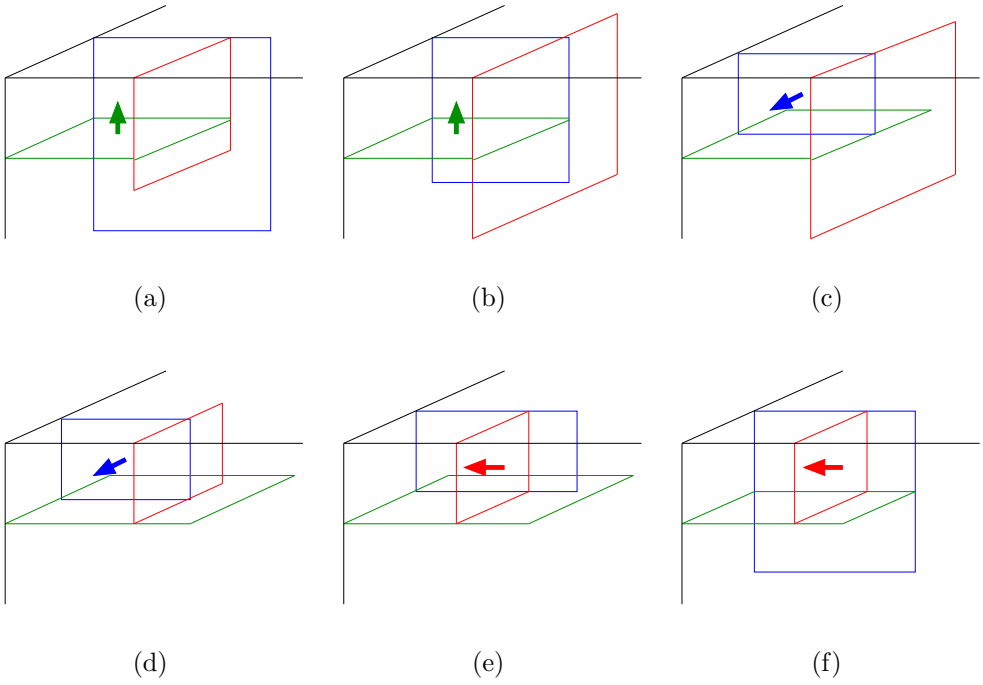

Figure 3.3: The possible cases for the proof of Observation 3.3.2
איור 3.3: המקרים האפשריים בהוכחת הבחנה 3.3.2

Observation 3.3.1 *Every edge of every rectilinear polygon in a partition is contained in one of the other polygons defining the partition, or in a face of the bounding box.*

Observation 3.3.2 *Let P be a partition-by-rectangles of a box B. Then, the box-deletion operation can be applied to any of the boxes at the corners of B.*

Proof: The proof follows from Observation 3.3.1 and from a simple case analysis (see Figure 3.3). □

Corollary 3.3.3 *The number of rectangles defining a partition-by-rectangles into n boxes is $n + 1$.*

The box-deletion operation is clearly reversible. We call the inverse operation *box insertion*.

Corollary 3.3.4 *A partition-by-rectangles can be constructed by a series of box insertions at a fixed corner of the bounding box.*

Recall that a box partition might contain two boxes intersecting at a single point (see Figure 3.1(c)). Next, we prove a different characterization of partitions-by-rectangles.

Lemma 3.3.5 *A box partition does not contain two boxes intersecting in a single point if and only if it is a partition-by-rectangles.*

Proof: Proving that a partition-by-rectangles does not contain two boxes sharing only a vertex can be done easily by induction on n and using the box-insertion operation.

Conversely, let v be a concave vertex of a rectilinear polygon r which is one of the polygons defining a partition P. Assume, without loss of generality, that v is as shown in Figure 3.4(a). Look at the boxes in P that are supported by r from below and contain v on their boundary. Exactly one of these boxes must have v as one of its vertices, as in Figure 3.4(b); having more than one would imply that there are two polygons contained in the same plane, as in Figure 3.4(c). Suppose, without loss of generality, that this box is the right one (refer to Figure 3.4(b)), and denote it by b_1. Now, look at the boxes in P that are supported by r from above and contain v on their boundary. Again, there are two such boxes and exactly one of them has v as a vertex; denote it by b_2. If b_2 is the right box of the two, then there are two polygons contained in the same plane (see Figure 3.4(d)), thus, b_2 must be the left box, and, therefore, b_1 and b_2 intersect at a single point (see Figure 3.4(e)). □

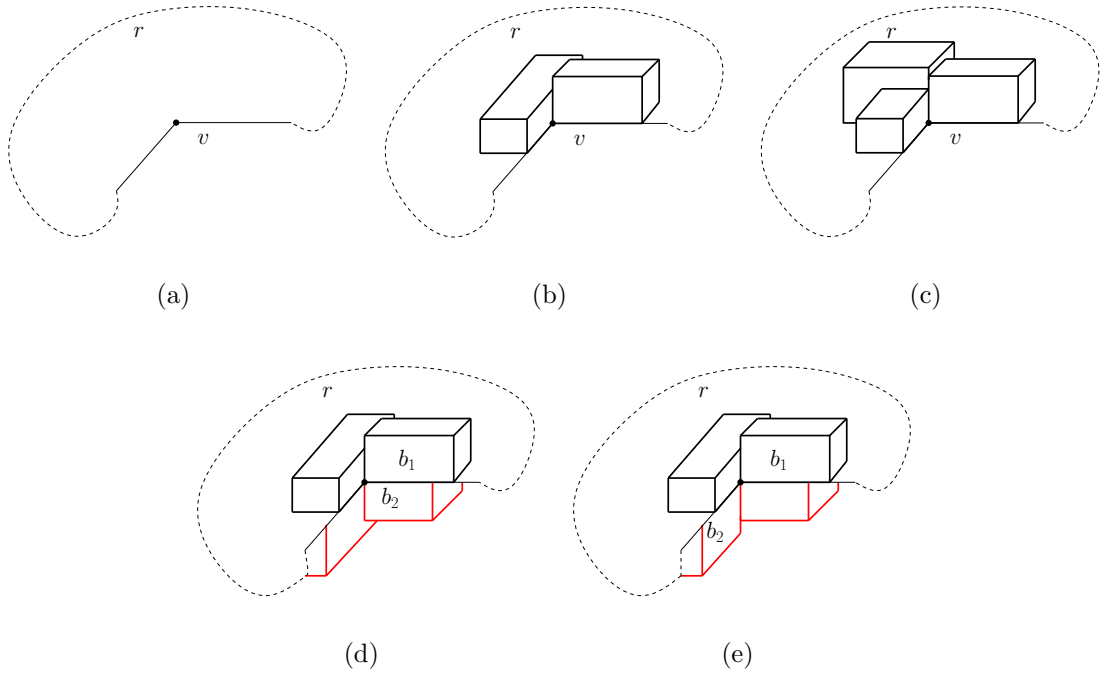

Figure 3.4: Illustrations for the proof of Lemma 3.3.5
איור 3.4: הוכחת למה 3.3.5

3.3.2 Computing the number of partitions-by-rectangles

In this section we show a method of computing the number of distinct partitions-by-rectangles using a more precise characterization of such partitions. First, we observe that partitions-by-rectangles have a certain hierarchal structure.

Definition 3.3.6 *Given a partition-by-rectangles P, a set of boxes in P whose union is a box is called a* block.

Definition 3.3.7 *A partition P is an* extrusion *if its boxes can be grouped into blocks such that the eight vertices of every block are on faces of the bounding box. That is, the blocks form an extrusion of a 2-dimensional rectangular partition. A partition is*

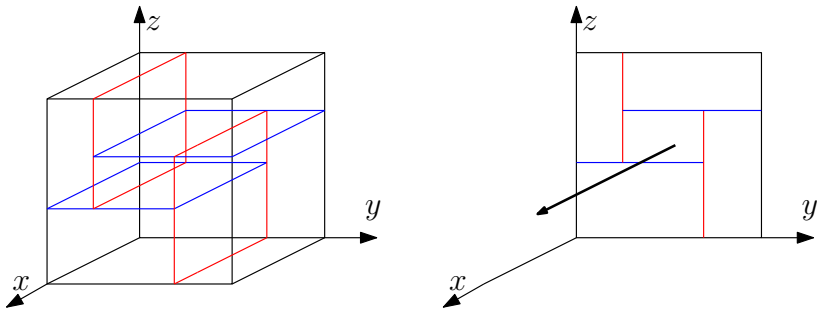

(a) An x-extrusion partition

(b) The extruded 2D partition

Figure 3.5: An extruded partition
איור 3.5: חלוקה "נמתחת"

an x-extrusion (resp., y- or z-extrusion) if all vertices of the blocks lie on two planes parallel to the yz-plane (resp., xz- or xy-plane).

See Figure 3.5 for an example of an extruded partition. It follows from the definition that an extruded partition is x-extrusion, y-extrusion, or z-extrusion (however, note that the 'or' is not exclusive, e.g., a partition might be both x- and y-extrusion).

Observation 3.3.8 *Let P be an x-extrusion partition and let $B_1, B_2, \ldots, B_k$ be blocks into which the boxes of P are grouped, such that k is as small as possible. Let F be a face of the bounding box which is parallel to the yz plane, and let $b_1, b_2, \ldots, b_m$ be a set of boxes in P, such that $b_1, b_2, \ldots, b_m$ touch F, and their projection on F is a rectangle (see Figure 3.6(a)). Then, there is a single block in $\{B_1, B_2, \ldots, B_k\}$ that contains $b_1, b_2, \ldots, b_m$.*

Proof: Denote by r the rectangle that is the union of the projection of $b_1, b_2, \ldots, b_m$ on F. The projection of $B_1, B_2, \ldots, B_k$ on F is a 2D partition into the rectangles

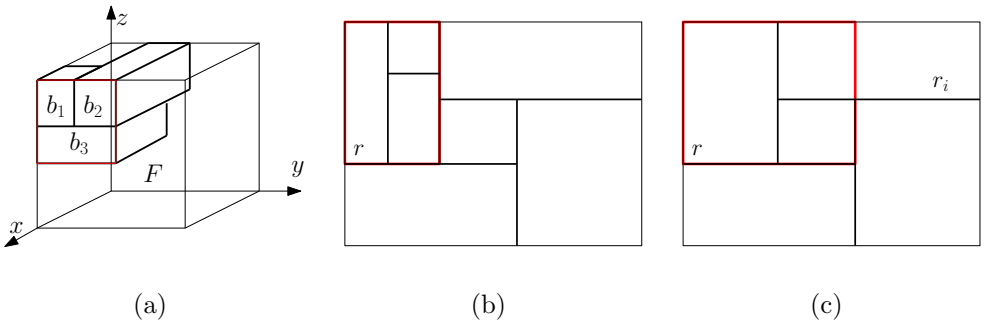

Figure 3.6: Illustrations for the proof of Observation 3.3.8
איור 3.6: הוכחת הבחנה 3.3.8

$r_1, r_2, \ldots, r_k$. If r is contained in one of those rectangles then the claim holds. If r is the union of some of the rectangles (see Figure 3.6(b)), then k is not minimal since the blocks that correspond to those rectangles can be grouped into a single block. Thus, there must be a rectangle r_i (for $1 \leq i \leq k$) that intersects r but does not contain r nor is contained in r (see Figure 3.6(c)). However, in this case it must be that two of the rectangles defining P lie in the same plane. □

Observation 3.3.9 *Let P be a partition-by-rectangles and let F be a face of the bounding box parallel to the yz-plane. Suppose that r is one of the rectangles defining P such that one of the edges of r is on F and the vertices of this edge are on edges of the bounding box (refer to Figure 3.7). Then, the boxes of P can be grouped into blocks $B_1, B_2, \ldots, B_k$ forming a z-extrusion, such that r is one of the rectangles defining the blocks $B_1, B_2, \ldots, B_k$.*

Proof: Let a be the edge of r on F and let b be the opposite edge. If b is on the bounding box, then we are done. Otherwise, it follows from Observation 3.3.1 that b is contained in another rectangle, denoted by r_1, whose height is the height of the

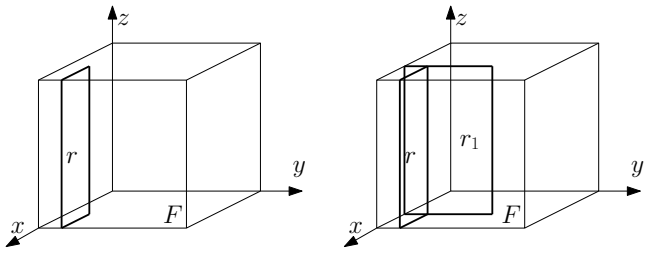

(a) (b)

Figure 3.7: Observation 3.3.9
איור 3.7: הבחנה 3.3.9

bounding box and is parallel to the yz-plane (see Figure 3.7(b)). We can now repeat the same arguments for each of the edges of r_1 that are orthogonal to the xy plane: either such an edge is contained in the bounding box, or it is contained in a rectangle r_j whose height is the height of the bounding box. Since the number of rectangles is finite, by repeating the same arguments for each such "new" rectangle we must end up with a set of rectangles that define a partition of P into blocks, such that P is a z-extrusion. □

Lemma 3.3.10 *Let P be a partition-by-rectangles into $n > 1$ boxes. Then, the boxes of P can be grouped into $k \geq 2$ blocks that form an extruded partition.*

Proof: By induction on n. For $n = 2$ the claim clearly holds. Let P be a partition-by-rectangles into $n > 2$ boxes. Delete one box, denoted by b_n, from a corner of the bounding box, w.l.o.g., by sliding it along the x-axis. We obtain a partition-by-rectangles P' into $n - 1$ boxes. Thus, by the induction hypothesis, the boxes of P' can be grouped into k blocks $B_1, B_2, \ldots, B_k$, $k \geq 2$, that form an extruded partition. Next, we consider the possible cases according to which type of extrusion P' is.

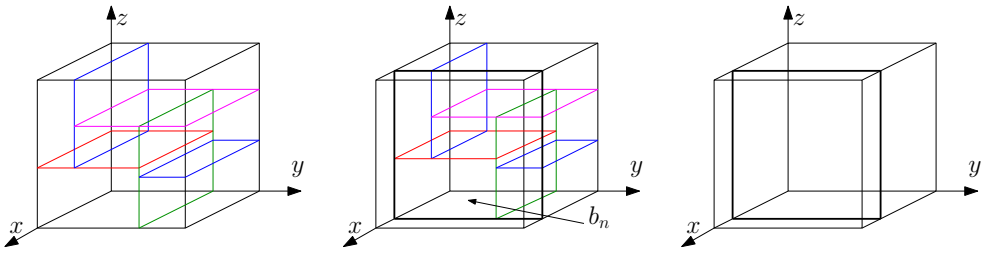

(a) Before inserting b_n

(b) After inserting b_n

(c) Blocks forming a y- and z-extrusion

Figure 3.8: First case in the proof of Lemma 3.3.10

איור 3.8: מקרה ראשון בהוכחת למה 3.3.10

P' is an x-extrusion. Assume that k is as small as possible. Let $b_{i_1}, b_{i_2}, \ldots, b_{i_\ell}$ be the boxes "pushed" away by inserting b_n into P'. Denote by F the face of the bounding box through which b_n is inserted. Then, the union of the faces of $b_{i_1}, b_{i_2}, \ldots, b_{i_\ell}$ on F forms a rectangle r. If r covers the whole face of the bounding box, then P is a y-extrusion (as well as a z-extrusion), b_n being one block and the union of all the other boxes being the other block (see Figure 3.8). Otherwise, P must be an x-extrusion: simply add b_n to the one block (which is unique, according to Observation 3.3.8) containing $b_{i_1}, b_{i_2}, \ldots, b_{i_\ell}$ (see Figure 3.9).

P' is a z-extrusion. In this case the intersection of the rectangles defining the blocks $B_1, B_2, \ldots, B_k$ with F is a set of guillotine cuts of F, which partition F into a set of rectangles $r_1, r_2, \ldots, r_m$ (see Figure 3.10(a)). If r is contained in r_1, then P is a z-extrusion (its boxes can be grouped into the blocks $B_1 \cup b_n, B_2, \ldots, B_k$, assuming that B_1 is the block corresponding to r_1). If r is not contained in r_1, then its height must be the height of the bounding box, otherwise there will be two

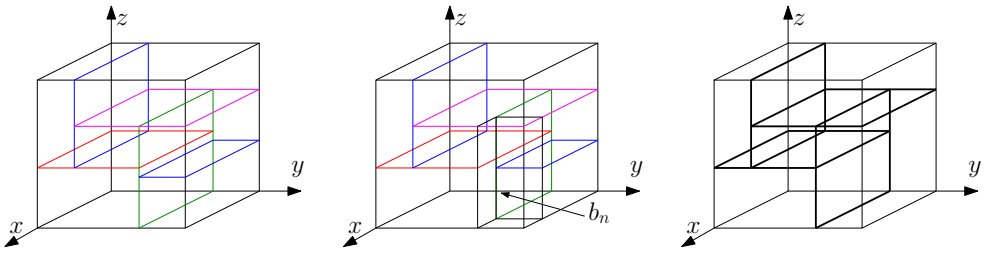

(a) Before inserting b_n

(b) After inserting b_n

(c) Blocks forming an x-extrusion

Figure 3.9: Second case in the proof of Lemma 3.3.10
איור 3.9: מקרה שני בהוכחת למה 3.3.10

rectangles defining P that are contained in the same plane (see Figure 3.10(b)). Let h be the left edge of r (refer to Figure 3.10(c)) and let a be the rectangle in the set of rectangles defining P whose edge is h. According to Observation 3.3.9, the boxes of P can be grouped into the blocks $B'_1, B'_2, \ldots, B'_m$ forming a z-extrusion, such that a is one of the rectangles defining the blocks $B'_1, B'_2, \ldots, B'_m$. Thus, P is also a z-extrusion by the blocks $b_n, B'_1, B'_2, \ldots, B'_m$.

P' is a y-extrusion. This case is similar to the case in which P' is a z-extrusion, and is thus omitted. □

By using their hierarchical structure, we can compute the number of partitions-by-rectangles. We use the following notations:

- $P(n)$: The number of partitions by n rectangles.
- $P_x(n)$: The number of partitions by n rectangles that are x-extrusions ($P_y(n)$ and $P_z(n)$ have similar meanings).

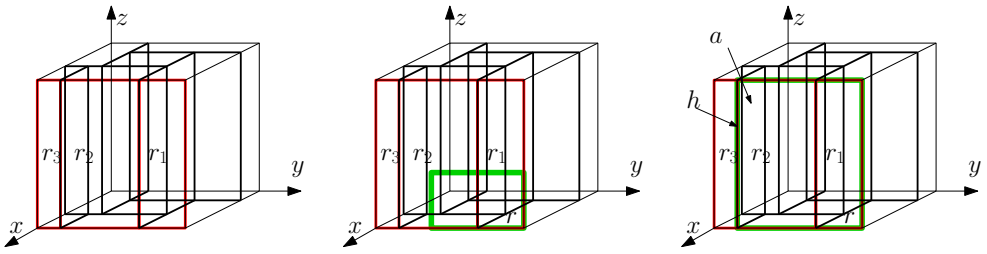

Figure 3.10: Third case in the proof of Lemma 3.3.10
איור 3.10: מקרה שלישי בהוכחת למה 3.3.10

- $P_{\overline{x}}(n)$: The number of partitions by n rectangles that are not x-extrusions ($P_{\overline{y}}(n)$ and $P_{\overline{z}}(n)$ have similar meanings).

- $P_{xy}(n)$: The number of partitions by n rectangles that are both x-extrusions and y-extrusions ($P_{xz}(n)$ and $P_{yz}(n)$ have similar meanings).

- $P_{\overline{xy}}(n)$: The number of partitions by n rectangles that are not both x-extrusions and y-extrusions ($P_{\overline{xz}}(n)$ and $P_{\overline{yz}}(n)$ have similar meanings).

Due to symmetry, $P_x(n) = P_y(n) = P_z(n)$, $P_{\overline{x}}(n) = P_{\overline{y}}(n) = P_{\overline{z}}(n)$, $P_{xy}(n) = P_{xz}(n) = P_{yz}(n)$, and $P_{\overline{xy}}(n) = P_{\overline{xz}}(n) = P_{\overline{yz}}(n)$. Using the inclusion-exclusion principle we can write:

$$\begin{aligned} P(n) &= P_x(n) + P_y(n) + P_z(n) - P_{xy}(n) - P_{xz}(n) - P_{yz}(n) \\ &= 3(P_x(n) - P_{xy}(n)) \\ P_x(n) &= \sum_{k=1}^{n} \mathrm{B}(k+1) \cdot \sum_{i_1+i_2+\ldots+i_{k+1}=n-k} \prod_{j=1}^{k+1} P_{\overline{x}}(i_j), \end{aligned}$$

where $P_x(0) = 1$ and $\mathrm{B}(n)$ denotes the nth Baxter number, which is also the number of mosaic floorplans [64]. Each mosaic floorplan by k segments can be extruded into an x-extrusion with k blocks, each of which can be further partitions by some of the $n-k$ remaining rectangles (provided that this partition is not an x-extrusion).

$$\begin{aligned}
P_{\overline{x}}(n) &= P_y(n) - P_{xy}(n) + P_z(n) - P_{xz}(n) - P_{yz}(n) \\
&= 2P_x(n) - 3P_{xy}(n), \quad \text{where } P_{\overline{x}}(0) = 1 \\
P_{xy}(n) &= \sum_{k=1}^{n} \sum_{i_1+i_2+\ldots+i_{k+1}=n-k} \prod_{j=1}^{k+1} P_{\overline{xy}}(i_j), \quad \text{where } P_{xy}(0) = 1 \\
P_{\overline{xy}}(n) &= P_x(n) + P_y(n) - 2P_{xy}(n) + P_z(n) - P_{xz}(n) - P_{yz}(n) \\
&= 3P_x(n) - 4P_{xy}(n), \quad \text{where } P_{\overline{xy}}(0) = 1
\end{aligned}$$

The equations for $P_{\overline{x}}(n)$ and $P_{\overline{xy}}(n)$ are easy to justify. Considering the equation for $P_{xy}(n)$, notice that k rectangles can form an x-extrusion that is also a y-extrusion in exactly one way (k guillotine cuts that are parallel to the xy-plane). The $k+1$ blocks formed by these k rectangles can be further subdivided by the remaining $n-k$ rectangles, provided that they do not subdivide a block into an x-extrusion which is also a y-extrusion (to prevent double counting of the same partition). The first values of $P(n)$ appear in Table 3.2. Note that $P(n)$ can be computed in this way just for small values of n, since it involves computing all the possible partitions (of an integer) for every integer $1, 2, \ldots, n-1$.

3.4 Conclusions and Open Problems

In this chapter we considered the number of partitions of a 3-dimensional box into smaller 3-dimensional boxes. For guillotine partitions we were able to give a tight

n	$g_3(n)$	$P(n)$
0	1	1
1	3	3
2	15	15
3	93	93
4	645	651
5	4,791	4,917
6	37,275	39,111
7	299,865	322,941
8	2,474,025	2,742,753
9	20,819,307	23,812,341
10	178,003,815	210,414,489
11	1,541,918,901	1,886,358,789
12	13,503,125,805	17,116,221,531
13	119,352,115,551	156,900,657,561
14	1,063,366,539,315	1,450,922,198,319
15	9,539,785,668,657	13,519,506,382,953

Table 3.2: The number of partitions by n rectangles: guillotine and overall
טבלה 3.2: מספר החלוקות באמצעות n מלבנים

asymptotic estimation of the number of partitions in any dimension using any number of hyperplanes. This provides context to an enumeration of all guillotine partitions, for example, for finding the one that optimizes a given measure.

An interesting related open problem arises when we also care about the *relative order* between guillotine cuts on opposite sides of their parent cut; for example, we distinguish between the partitions shown in Figures 3.11(a) and 3.11(b). This gives rise to new modeling and enumeration problems which have some bearing on floorplanning of integrated circuits.

Considering partitions-by-rectangles in three dimensions, we showed some properties of such partitions, most important of which characterizes partitions-by-rectangles as partitions that can be recursively obtained by extracting 2-dimensional partitions

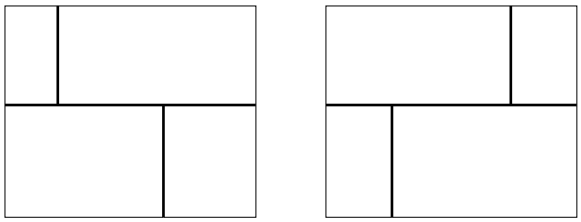

Figure 3.11: What is the number of guillotine partitions when these are considered different partitions?

איור 3.11: מה מספר חלוקות הגליוטינה כאשר חלוקות אלה שונות זו מזו?

(mosaic floorplans). This characterization provides a way of computing the number of partitions-by-rectangles. It is an open problem to come up with simpler expressions for the number of such partitions.

Chapter 4

Enumerating Rectangulations

4.1 Introduction

Given a set P of n points within a rectangle R, a *rectangulation* (or *rectangular partition*) of (R, P) is a subdivision of R into rectangles by non-intersecting axis-parallel segments, such that every point in P lies on a segment.

The problem of finding a rectangulation that minimizes the sum of lengths of the segments (known as RGP [26], or RPP [15]) has attracted some attention. First, it was introduced by Lingas et al. [36] as a special case of partitioning a rectilinear polygon containing rectilinear holes into rectangles. The motivation for this partitioning problem comes from the design of integrated circuits. This problem, as well as RGP, were shown to be NP-hard [36]. Later, several approximation algorithms for RGP were suggested (see, e.g., [21, 26, 27, 28, 35]), including a polynomial-time approximation scheme [17, 40]. RGP has also applications to stock (or die) cutting in the presence of material defects.

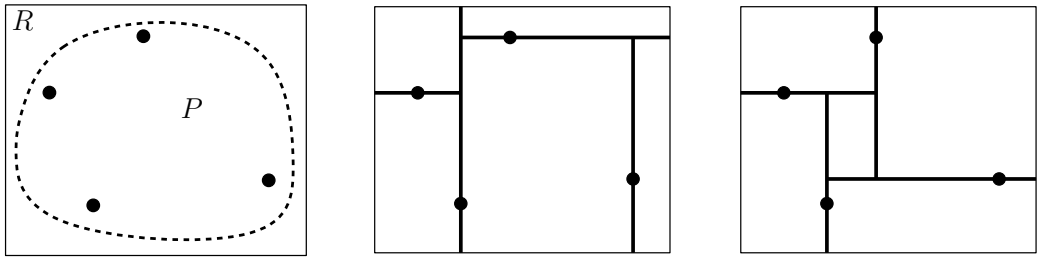

Figure 4.1: Rectangulations of (R, P)
(R, P) איור 4.1: מילבונים של

When the points in P are in general position in the sense that no two points have the same x or y coordinate, i.e., the points are *noncorectilinear*, then the complexity class of the minimization problem (known as RGNLP [26], or NCRPP [15]) is still unknown. However, Calheiros et al. [15] have shown that an optimal solution must comprise exactly n non-intersecting segments.

In this chapter we consider the *number* of such rectangulations, namely:

> Given a set P of n noncorectilinear points in the plane within a rectangle R, how many different ways are there to divide R into smaller rectangles by n (non-intersecting) segments such that every point in P lies on a segment?

See Figure 4.1 for examples of such rectangulations.

It is easy to see that the set of rectangulations of (R, P) does not depend on R. Thus, we use the notation $\Xi(P)$ to denote the set of rectangulations of P. Moreover, we observe that $\Xi(P)$ depends only on the relative order of the points in P. We represent this order by a permutation π on $[n]$ (reflecting the order of y-coordinates with respect to the x-coordinates when listing the points in P from left to right), and

show that if π is a *separable* permutation [14], then the number of rectangulations is the $(n+1)$st Baxter number.

When the permutation of the points is arbitrary, we use a novel technique of Sharir and Welzl [55] to show an upper bound of $O(18^n/n^4)$ for the number of rectangulations. We also show that the number of *guillotine* rectangulations (see Definition 4.4.1) in this case is the nth Schröder number.

We also consider the relation between rectangulations and floorplans and show that given a set P of n points whose permutation is separable, and a mosaic floorplan f with n segments, f can be drawn such that every point in P is on exactly one segment of f. From this result we conclude a stronger version of a result of de Fraysseix, de Mendez and Pach [20] about the embedding of bipartite planar graphs as contact graphs of vertical and horizontal segments in the plane.

The rest of this chapter is organized as follows. In Section 4.2 we describe two methods to enumerate rectangulations. Next, we show the upper bound for the number of rectangulations. In Section 4.4 we discuss guillotine rectangulations. The heart of this chapter (Section 4.5) is an analysis of the exact number of rectangulations. We start by observing that this number depends only on the permutation of the points in P, then we show that for points arranged in an identity permutation the number of rectangulations is $\mathrm{B}(n+1)$. Next we define separable permutations and generalize this result for them. In Section 4.6 we discuss the relation between rectangulations and floorplans, and finally, we conclude in Section 4.7.

4.2 Computer-assisted Enumeration

In this section we present two methods of computing the number of rectangulations. The first generates all the rectangulations using two simple operators; the second method counts the number of rectangulations without actually generating them, and is thus more efficient.

4.2.1 Enumeration by generating all the rectangulations

Following we define two operators that enable us to explore the space of all the rectangulations of a given point set P (within a rectangle R). Given a rectangulation x we can obtain new rectangulations by applying each of the following operators on x.

Definition 4.2.1 (Flip) *Let p be a point in P such that the segment s containing p does not contain any endpoints of other segments. The operator* `Flip`(x, p) *changes the orientation of s from vertical to horizontal or vice-versa.*

Definition 4.2.2 (Rotate) *Let s_1 be a segment that contains one or more endpoints of other segments, and let t be such an endpoint which is extreme on s_1 (closest to one of its endpoints). Denote by s_2 the segment terminated at t. The operator* `Rotate`(x, t) *extends s_2 beyond t until it reaches another segment (or the boundary) and shortens s_1 to t.*

See Figure 4.2 for examples of the `Flip` and `Rotate` operators.

Given a set of (noncorectilinear) points P, we denote by $G(P) = (V, E)$ the graph of rectangulations of P, where $V = \{x : x \text{ is a rectangulation of } P\}$ and $E = \{(x_1, x_2) : x_2 \text{ is reachable from } x_1 \text{ by a single } \texttt{Flip} \text{ or } \texttt{Rotate} \text{ operation}\}$. $G(P)$ is undirected since both operators are clearly reversible.

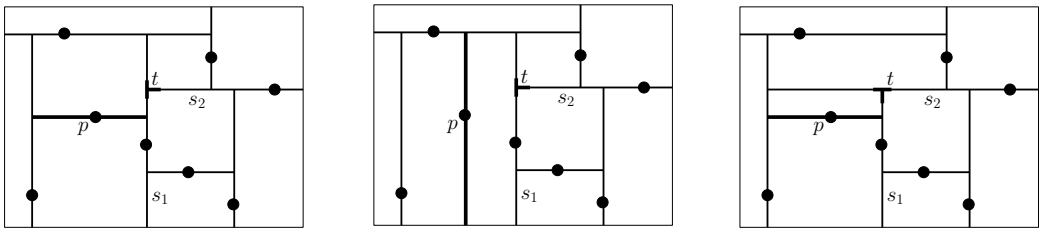

(a) A rectangulation x (b) `Flip`(x, p) (c) `Rotate`(x, t)

Figure 4.2: Applying the `Flip` and `Rotate` operators
`Rotate`-ו `Flip` פעולות :4.2 איור

Lemma 4.2.3 *Let P be a set of noncorectilinear points in the plane and let $G(P)$ be the graph of rectangulations of P. Then $G(P)$ is connected.*

Proof: Let x_1 and x_2 be two different rectangulations, and let x_v be the rectangulation in which all the segments are vertical. The rectangulation x_v can be reached from both x_1 and x_2 by a finite series of `Rotate` and `Flip` operations: Shorten every horizontal segment that contains endpoints of other segments by the `Rotate` operator, then turn it into a vertical segment by the `Flip` operator. Therefore there is a path between x_1 and x_2 (through x_v) in $G(P)$. □

It is easy to see that $O(n^2)$ `Flip` or `Rotate` operations are enough to reach x_v from any given rectangulation. Therefore, the *diameter*[1] of $G(P)$ is $O(n^2)$. However, we did not find any examples for a pair of rectangulations whose distance is super-linear.

Problem 4.2.4 *What is the maximum diameter of the graph of rectangulations of a set of n points?*

It follows from Lemma 4.2.3 that it is possible to generate and iterate over all the rectangulations of P by traversing $G(P)$ by, say, a standard depth-first search. Every

[1]The diameter of a graph is the longest distance between a pair of its vertices.

vertex in $G(P)$ has $\Theta(n)$ neighbors, as every segment can be either flipped or rotated. Since handling every neighbor takes $O(n)$ time, exploring all the rectangulations in such a way takes $O(n^2|\Xi(P)|)$ time and $O(n|\Xi(P)|)$ space. However, one can do better.

Theorem 4.2.5 *Given a set of n points P, it is possible to traverse all the rectangulations of P in $O(|\Xi(P)|\log(n))$ time and $O(n^3)$ space.*

Proof: We use a method of Avis and Fukuda [10] known as *Reverse Search*. The key observation is that in order to visit all the vertices (rectangulations) of a graph, it is enough to traverse a *spanning tree* of the graph. This saves time (we do not explore all the edges) and space (there is no need to keep a record of the already-visited vertices).

Given a set of points P, let x_h be the rectangulation of P in which all the segments are horizontal. For every rectangulation except x_h we designate one of its neighbors to be its "parent" in such a way that every rectangulation is a descendant of x_h. The parent of a rectangulation $x \neq x_h$ is defined as follows: Let s be the leftmost vertical segment in x. If the operator `Flip` can be applied on s, then the result of applying it is the parent of s. Otherwise, we can apply the `Rotate` operator and shorten s. If we can shorten s from below, then the rectangulation we get as a result is the parent of x. Otherwise, the rectangulation we get by shortening s from above is the parent of x. Clearly, every rectangulation (except x_h) has a parent, and every rectangulation is a descendant of x_h.

In order to find the children of a certain rectangulation x we can keep a pointer to the leftmost vertical segment in x, and a sorted list of flippable horizontal segments

that are to the left of it (that is, segments that pass through points which are left of the vertical segment). The children of x are obtained by either:

1. Flipping one of the horizontal segments. In this case the flipped segment becomes the leftmost vertical segment and the list of flippable horizontal segments is the list of flippable horizontal segments to the left of it that do not contain endpoints of the flipped segment.

2. Extending the leftmost vertical segment using a `Rotate` operation (downwards if it is possible, or upwards if it is possible and it is impossible to shorten it from below by a `Rotate` operation). In this case the leftmost vertical segment remains the same. Additionally, at most one segment is added to the list of flippable horizontal segments and at most one segment is removed from this list.

Updating the sorted list of segments can be performed in $O(\log n)$ using a (slightly modified) deterministic skip list [41]. Therefore, the time complexity of enumerating (by generating) all the rectangulations is $O(|\Xi(P)| \log n)$. The depth of the spanning tree is bounded by $O(n^2)$, since, when traversing from parent to child, the leftmost vertical segment is either extended or replaced by a vertical segment to the left of it. Thus, the space complexity is $O(n^3)$. □

4.2.2 Fast enumeration of rectangulations

Let x be a rectangulation of P, a set of n points, and let ℓ be a horizontal line not containing any point from P. The intersection of x and ℓ can be represented by a binary word of length $n + 2$, in which the $(i + 1)$st bit (from left to right) is set if ℓ intersects a vertical segment that passes through the ith point (left-to-right)

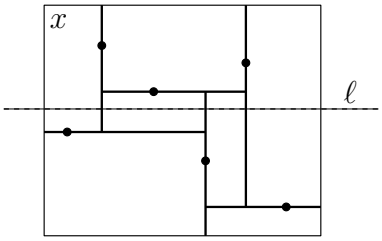

Figure 4.3: 10101101 represents the intersection of x and ℓ
איור 4.3: המחרוזת 10101101 מייצגת את החיתוך של x ו-ℓ

in P. (For convenience, the first and last bits of the word are always set, in order to represent the intersection of the sweeping line with the bounding rectangle.) See Figure 4.3 for an example. If we sweep ℓ from bottom to top (skipping over the points of P) we get a sequence of $n+1$ binary words of length $n+2$ that represents the rectangulation x. For example, the rectangulation in Figure 4.3 is represented by the sequence (10001001, 10001101, 10001101, 10101101, 10100101, 10100101, 10100101). This observation suggests a way of computing the number of rectangulations of P as follows. Define the following directed acyclic graph $G = (V, E)$:

1. Set two distinct vertices v_N and v_S, and $(n+1)2^n$ vertices of the form v_w^j, for every $w \in 1\{0,1\}^n1$ and $1 \leq j \leq n+1$. For $1 \leq j \leq n$, a vertex of the form v_w^j corresponds to an intersection of the sweeping line just below the jth point (from bottom to top), resulting in the sequence w. A vertex of the form v_w^{n+1} corresponds to an intersection of the sweeping line just above the nth point (from bottom to top), resulting in the sequence w.

2. Set edges from v_S to v_w^1 and from v_w^{n+1} to v_N, for every $w \in 1\{0,1\}^n1$.

3. Let $p_2, p_3, \ldots, p_{n+1}$ be the points of P, such that p_{k+1} is the kth point from left

to right. Let p_i be the $(i-1)$st point from left to right, and the jth point from bottom to top. Denote by w_k the kth bit of w. Then, the neighbors of v_w^j are defined by the following rules:

(a) If $w_i = 1$, then v_w^j has only one neighbor, v_w^{j+1}. This case corresponds to a vertical segment through p_i.

(b) Assume that $w_i = 0$. This case corresponds to a horizontal segment through p_i. The neighbors of v_w^j are all the vertices $v_{w'}^{j+1}$ that satisfy:

   i. $w'_i = 0$ (since the segment through p_i is horizontal); and

   ii. there are integers $1 \leq l < i$ and $i < r \leq n+2$ (representing the left and right endpoints of the horizontal segment through p_i) such that:

      A. $w_l = w'_l = w_r = w'_r = 1$;

      B. $w_s = w'_s$ for every $1 \leq s < l$ and $r < s \leq n+2$;

      C. $w'_s = 0$ for every $l < s < r$ such that p_s is below p_i; and

      D. $w_s = 0$ for every $l < s < r$ such that p_s is above p_i.

   See Figure 4.4 for an example of the neighbors of a certain vertex according to this rule.

Consequently, the number of rectangulations is the number of paths in G from v_S to v_N. Next we show that counting the number of rectangulations in this way can be implemented in $O(n^4 2^n)$ time and $O(n^3 2^n)$ space.

Lemma 4.2.6 *Given a set of n points P, let $G = (V, E)$ be the corresponding DAG of rectangulations. Then, $|E| = \Theta(n^3 2^n)$.*

Proof: Let p_{j+1} be the jth point left-to-right and the kth point bottom-to-top, and let e_k be the number of edges of the form $(v_w^k \to v_{w'}^{k+1})$. Then $e_k = 2^{n-1} + (j-1)(n-$

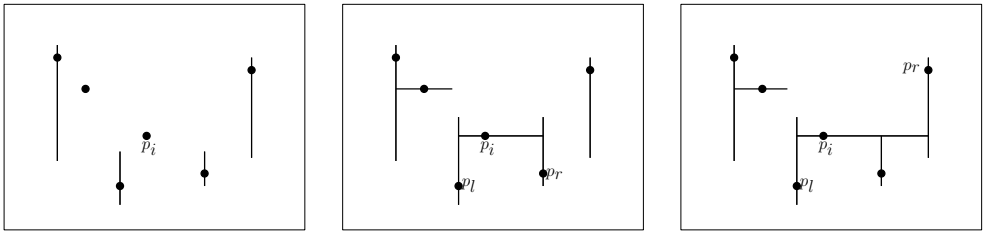

(a) $v^j_{1\ldots101011\ldots1}$ (b) $v^{j+1}_{1\ldots101011\ldots1}$ (c) $v^{j+1}_{1\ldots101001\ldots1}$

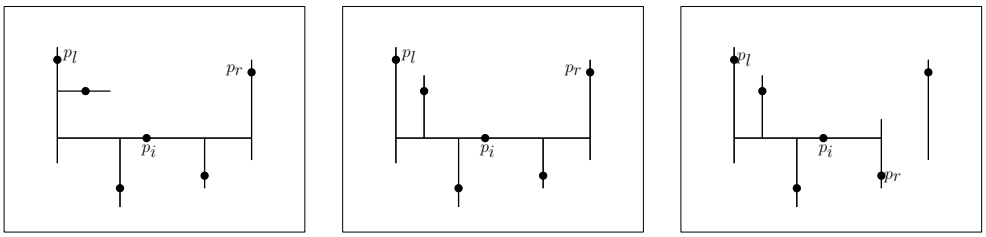

(d) $v^{j+1}_{1\ldots100001\ldots1}$ (e) $v^{j+1}_{1\ldots110001\ldots1}$ (f) $v^{j+1}_{1\ldots110011\ldots1}$

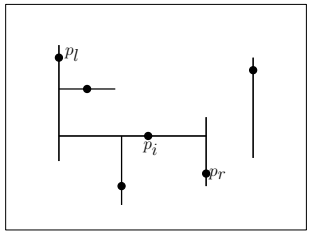

(g) $v^{j+1}_{1\ldots100011\ldots1}$

Figure 4.4: $v^j_{1\ldots101011\ldots1}$ and its neighbors according to rule 3(b)
3(b) ושכניו לפי כלל $v^j_{1\ldots101011\ldots1}$:4.4 איור

$j)2^{n-3} + (j-1)2^{n-2} + (n-j)2^{n-2} + 2^{n-1}$. The first summand stands for all the words w in which the jth bit is set (i.e., there is a vertical segment through p_{j+1}), from which there is only one out-edge. The second summand represents all the cases in which the endpoints of the horizontal segment through p_{j+1} are set by coordinates of other points from P: one to the left of p_{j+1} and the other to the right of p_{j+1}. There are $(j-1)(n-j)$ options to choose such a pair and 2^{n-3} options to set the other bits in w and w' (if a bit in w can be either 0 or 1, then the corresponding bit in w' has only one option, and vice versa). The rest of the summands are for the cases in which one or two of the endpoints of the horizontal segment through p_{j+1} are on the bounding rectangle. Therefore,

$$|E| = 2 \cdot 2^n + \sum_{k=1}^{n} e_k = \Theta(n^3 2^n).$$

□

Constructing G takes $O(n^2 2^n + n|E|)$ time since computing the neighbors of every vertex v_w^j can be performed in $O(n + d_{\text{out}}(v_w^j))$ time. Computing the number of paths from a source vertex to a sink vertex in a DAG takes $O(|E|)$ time. Since every vertex is represented by an n-bit word,[2] the time complexity of this enumeration algorithm is $O(n^4 2^n)$. Considering the space complexity, note that the DAG is composed of $n+3$ "levels" and the edges are only between consecutive levels. Thus, it is enough to hold in memory only two consecutive levels and therefore the space complexity is $O(n^3 2^n)$.

[2]Likewise, a factor of $\log n$ should be added in the analysis of the previous algorithm; however, we follow the common practice and omit this factor.

4.3 An Upper Bound on the Number of Rectangulations

In this section we show that the number of rectangulations of a set of n points is $O(18^n/n^4)$. Our proof is based on the proof technique of Sharir and Welzl [55] for the currently best upper bound on the number of *triangulations* of a planar point set. Their method is an extension of work by Santos and Seidel [53] who showed the previously best upper bound on the number of such triangulations. Santos and Seidel's method can be used to show an upper bound of $O(20^n/n^4)$ on the number of rectangulations [3]. However, we do not see any straightforward relation between the number of rectangulations and the number of triangulations of a set of points.

Theorem 4.3.1 *Let $f(n)$ be the maximum number of rectangulations of n noncorectilinear points by n segments. Then $f(n) \leq 18^n/\binom{n+4}{4}$.*

Proof: By induction on n. For $n = 0, 1, 2$ the claim holds trivially: $f(0) = 1 = 18^0/\binom{4}{4}$, $f(1) = 2 < 3.6 = 18^1/\binom{5}{4}$, and $f(2) = 6 < 21.6 = 18^2/\binom{6}{4}$. Let x be a rectangulation of a set of points P within a rectangle R, such that $|P| = n \geq 3$ and $\Xi(P) = f(n)$. A *T-junction* is an endpoint of a segment on another segment, or on R's boundary. The *degree* of a point $p \in P$ in x, $d(p, x)$, is the number of T-junctions on the segment that contains p. We denote by $n_i(x)$ the number of points with degree i in x. Clearly, $\sum_i n_i(x) = n$. Every segment is bounded by two T-junctions, thus every segment s contributes at most 4 to the total sum of degrees: 2 to the point it contains, and 1 to every point that is contained in a segment bounding s (if it is not a boundary segment). Thus, the total sum of degrees is $4n - b_x$, where b_x is the number of T-junctions on the boundary of R in x.

Observation 4.3.2 *Let* $n_2(P) = \sum_{x \in \Xi(P)} n_2(x)$ *be the total number of degree-2 points in all the rectangulations of* P*. Then,* $n_2(P) = 2 \sum_{p \in P} |\Xi(P \setminus \{p\})|$.

Proof: Consider a specific point $p \in P$. Every rectangulation x of $P \setminus \{p\}$ can be extended to two rectangulations of P by splitting the rectangle of x in which p lies into two rectangles by either a vertical or a horizontal segment through p. Moreover, these are exactly the rectangulations of P in which the degree of p is 2. □

Next, we assign a *charge* to every point in every rectangulation x of P: $ch(p, x) := 5 - d(p, x)$. Let $ch(x) = \sum_{p \in P} ch(p, x)$ be the charge of a rectangulation x, and let $ch(\Xi(P))$ be the total charge over all the rectangulations of P.

Observation 4.3.3 $ch(\Xi(P)) \geq (n + 4)|\Xi(P)|$.

Proof: Let x be a rectangulation of P. Then $ch(x) = \sum_i (5 - i)n_i(x) = 5 \sum_i n_i(x) - \sum_i i \cdot n_i(x)$. Recall that $\sum_i n_i(x) = n$ and that the total sum of degrees is $\sum_i i \cdot n_i(x) = 4n - b_x$, where b_x is the number of T-junctions on R in x. Thus, $ch(x) = n + b_x$. One can easily see that any rectangulation of a set of at least three points has at least four T-junctions on R. Hence, $ch(x) \geq n + 4$ and the claim follows. □

Given a point p in a rectangulation x', let x be the rectangulation obtained from x' by continuously applying the `Rotate` operation on the segment through p until the degree of p becomes two. We say that (p, x) is the *successor* of (p, x') and that (p, x') is a *predecessor* of (p, x). If $d(p, x') = 2$, then (p, x') is the successor of itself. Note that every point in a rectangulation has exactly one successor, however, a pair (p, x) might have many predecessors. We proceed by distributing the charges, such that every pair of a point p in a rectangulation x' sends its charge to its successor. Let $ch'(\cdot)$ be the charge of a point in a rectangulation after this step.

Observation 4.3.4 *Let p be a point in a rectangulation x. Then, if $d(p,x) = 2$, then $ch'(p,x) \leq 9$. Otherwise, $ch'(p,x) = 0$.*

Proof: If $d(p,x) > 2$, then (p,x) sends all of its charge, thus, $ch'(p,x) = 0$. Assume now that $d(p,x) = 2$ and let $ch_i(p,x)$ be the total charge sent to (p,x) by predecessors (p,x') such that $d(p,x') = i$. Then, $ch_2(p,x) = 3$, since the only predecessor (p,x') of (p,x), such that $d(p,x') = 2$ is (p,x) itself. Let s be the segment through p in x. Since the `Rotate` operation is reversible, every predecessor (p,x') of (p,x), such that $d(p,x') = i$, can be reached by applying $i-2$ `Rotate` operations on s. Moreover, the predecessor is uniquely determined once we decide how many of the $i-2$ `Rotate` operations are applied to each of the endpoints of s. Therefore, (p,x) has at most $i-1$ predecessors (p,x'), such that $d(p,x') = i$, and $ch_i(p,x) \leq (5-i)(i-1)$. In particular, $ch_3(p,x) \leq 4$ and $ch_4(p,x) \leq 3$.

Note that if $ch_4(p,x) = 3$, then it is possible to extend s by using four `Rotate` operations, two on each of its endpoints. Thus, (p,x) has a predecessor (p,x'), such that $d(p,x') = 6$ and $ch_6(p,x) \leq -1$. Hence, $ch'(p,x) = \sum_i ch_i(p,x) \leq ch_2(p,x) + ch_3(p,x) + ch_4(p,x) + ch_6(p,x) \leq 9$. □

Combining Observations 4.3.2, 4.3.3, and 4.3.4 we have:

$$(n+4)f(n) = (n+4)|\Xi(P)| \leq ch(\Xi(P)) \leq 9n_2(P) = 18\sum_{p\in P}|\Xi(P\backslash\{p\})| \leq 18n\cdot f(n-1).$$

By the induction hypothesis, $f(n-1) \leq 18^{n-1}\binom{n+3}{4}$, and the claim follows. □

4.4 Guillotine Rectangulations

In this section we consider a special class of rectangulations: guillotine rectangulations.

Definition 4.4.1 (Guillotine rectangulation) *In a* guillotine *rectangulation the segments can be ordered so that when the partition is executed according to that order, the current segment always partitions a rectangle into two rectangles.*

For example, the rectangulation in Figure 4.1(b) is guillotine, whereas the rectangulation in Figure 4.1(c) is not. In this section we consider the number of guillotine rectangulations. It is easy to see that this number depends only on the number of points in P. Let $\Gamma(n)$ be the number of guillotine rectangulations when $|P| = n$. Clearly, $\Gamma(n)/2$ guillotine rectangulations contain a vertical segment cutting the bounding rectangle into two rectangles, while the remaining $\Gamma(n)/2$ rectangulations contain a horizontal segment cutting the bounding rectangle into two. Considering only the first set and denoting by k the first point left to right, through which passes a vertical segment cutting the bounding rectangle into two, we derive the following recursive formula for $\Gamma(n)$:

$$\Gamma(n)/2 = \Gamma(n-1) + \sum_{k=2}^{n} \left(\frac{1}{2}\Gamma(k-1) \right) \Gamma(n-k),$$

where $\Gamma(0) = 1$. The formula holds since for $k = 1$ there are $\Gamma(n-1)$ guillotine rectangulations, while for $k > 1$ the segment through the kth point splits the bounding rectangle into two rectangles: the right one has $\Gamma(n-k)$ guillotine rectangulations, while the left one has $\Gamma(k-1)/2$ guillotine rectangulation as it must be cut into two by a horizontal segment.

This formula is equivalent to a recursive formula of the nth (large[3]) Schröder number:

$$r_n = r_{n-1} + \sum_{k=0}^{n-1} r_k r_{n-1-k}, \qquad r_0 = 1.$$

[3]The nth small Schröder number counts the number of possible bracketing on a word of n letters. For $n > 1$ it is exactly half of the nth large Schröder number.

Thus, we have:

Theorem 4.4.2 *Given a rectangle R which encloses a set P of n noncorectilinear points, the number of guillotine rectangulations of (R, P) is the nth Schröder number.*

4.5 The Exact Number of Rectangulations

In this section we investigate $|\Xi(P)|$—the exact number of rectangulations (guillotine and non-guillotine) of a set P of n noncorectilinear points within a rectangle R. We start by observing that $|\Xi(P)|$ depends only on the permutation of the points in P. Next, we show that for identity permutations the number of rectangulations equals the $(n+1)$st Baxter number. Finally, we generalize this result for the class of *separable* permutations.

4.5.1 Rectangulations and permutations

Definition 4.5.1 *Given a set P of noncorectilinear points, we refer to the relative order of the points in P as the permutation of P and denote it by $\pi(P)$.*

Representing the relative order of the points by a permutation $\pi = (\sigma_1\sigma_2 \ldots \sigma_n)$ is feasible since the points are noncorectilinear. By $\sigma_i = j$ we mean that the ith point along the x-axis is the jth point along the y-axis. It is easy to see that given two sets of points, P_1 and P_2, such that $\pi(P_1) = \pi(P_2)$, we always have $\Xi(P_1) = \Xi(P_2)$. In other words, the number of rectangulations does not depend on the actual point coordinates, it depends only on the permutation of points. Therefore, we will also use the notation $\Xi(\pi)$. However, computational enumerations we have performed showed

that when $\pi(P_1) \neq \pi(P_2)$ it is possible to have $|\Xi(\pi(P_1))| \neq |\Xi(\pi(P_2))|$. For example, $|\Xi(1234)| = 92$, whereas $|\Xi(3142)| = 93$.

4.5.2 The number of rectangulations of identity permutations

Lemma 4.5.2 *Let $\mathcal{I}_n$ be the identity permutation on $[n]$. Then $|\Xi(\mathcal{I}_n)| = B(n+1)$.*

Proof: Given a rectangulation x we denote by *bottom*(x) (resp., *top*(x)) the set of vertical segments touching the bottom (resp., top) edge of the bounding rectangle R. Similarly, *left*(x) (resp., *right*(x)) denotes the set of horizontal segments touching the left (resp., right) edge of R. Let $T_n(i,j)$ be the number of different rectangulations x of n points with the identity permutation, such that $|\text{top}(x)| = i$ and $|\text{right}(x)| = j$. Then we can write the following recurrence relation for $n > 0$:

$$T\left(n+1, i+1, j+1\right) = \sum_{k=1}^{\infty} \left(T\left(n, i, j+k\right) + T\left(n, i+k, j\right)\right), \tag{4.1}$$

where $T_0(0,0) = 1$ and $T_n(i,j) = 0$ for $n < 0$. To understand why this relation holds, note that we can create a rectangulation x of $n+1$ points such that $|\text{top}(x)| = i+1$ and $|\text{right}(x)| = j+1$ from a rectangulation x' of n points, such that $|\text{top}(x')| = i$ and $|\text{right}(x')| = j+k$ (for $k \geq 1$), by:

1. Adding an additional point p to the right and above all the points of x';
2. Setting a vertical segment s through p; and
3. Extending s downwards using `Rotate` operations until $k-1$ segments are removed from right(x).

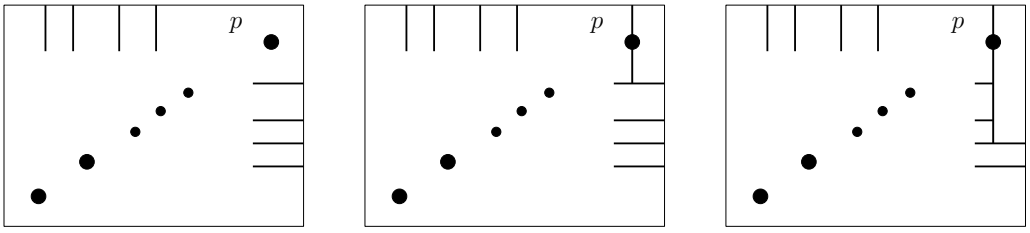

(a) Adding an additional point

(b) Setting a vertical segment

(c) Extending the vertical segment

Figure 4.5: From $T_n(i, j+k)$ to $T_{n+1}(i+1, j+1)$
איור 4.5: המעבר מ-$T_n(i, j+k)$ ל-$T_{n+1}(i+1, j+1)$

Figure 4.5 shows these steps. We can create in a similar way a rectangulation x of $n+1$ points, for which $|\text{top}(x)| = i+1$ and $|\text{right}(x)| = j+1$, from a rectangulation x' of n points, such that $|\text{top}(x')| = i+k$ (for $k \geq 1$) and $|\text{right}(x')| = j$, by passing a horizontal segment through a new point p. Clearly, every rectangulation x of $n+1$ points can be created from a rectangulation x' of n points as described above, and there are no two different rectangulations x'_1, x'_2 of n points that lead to the same rectangulation of $n+1$ points. Therefore,

$$|\Xi(\mathcal{I}_n)| = \sum_{i,j \geq 0} T_n(i,j), \tag{4.2}$$

which is exactly $\text{B}(n+1)$ by [19]. □

4.5.3 Separable permutations and their number of rectangulations

In this section we show that $|\Xi(\pi)| = \text{B}(n+1)$ if π is a separable permutation. Recall that separable permutations were introduced in Section 2.2.2.

Let x be a rectangulation. The *interface* of x, denoted by $\mathcal{F}(x)$, is an ordered quadruple (l, t, r, b), such that $l = |\text{left}(x)|$, $t = |\text{top}(x)|$, $r = |\text{right}(x)|$, and $b = |\text{bottom}(x)|$. We denote by $\Xi(\pi, \mathcal{F})$ the set of rectangulations with permutation π and interface $\mathcal{F}$.

Lemma 4.5.3 *For every n, l, t, r, b, $|\Xi(\mathcal{I}_n, (l, t, r, b))| = |\Xi(\mathcal{I}_n, (l, b, r, t))|$.*

Proof: Let x be a rectangulation of a point set P that lies within a rectangle R. We denote by $\mathcal{H}(x)$ (resp.,$\mathcal{V}(x)$) the set of segments in x touching both vertical (resp., horizontal) edges of R. Clearly, $\mathcal{H}(x) \neq \emptyset$ implies $\mathcal{V}(x) = \emptyset$ and $\mathcal{V}(x) \neq \emptyset$ implies $\mathcal{H}(x) = \emptyset$.

Given a rectangulation x, we call a pair of segments $s_1 \in \text{top}(x)$ and $s_2 \in \text{bottom}(x)$, such that s_1 is to the left of s_2, *$\natural$-segments*. If, in addition, there is no other segment $s \in \text{top}(x) \cup \text{bottom}(x)$ between s_1 and s_2, we say that s_1 and s_2 are *adjacent* $\natural$-segments. The next observation will be useful in the sequel.

Proposition 4.5.4 *Given a rectangulation x of (R, P), such that $\pi(P) = \mathcal{I}_n$ and $\mathcal{H}(x) = \mathcal{V}(x) = \emptyset$:*

1. *There is a pair of $\natural$-segments in x; or*

2. *There are segments $s_1 \in$ left(x) and $s_2 \in$ right(x) such that s_1 is above s_2.*

Proof: $\mathcal{H}(x) = \emptyset$ implies that $\text{top}(x) \neq \emptyset$ and $\text{bottom}(x) \neq \emptyset$. Suppose that there is no pair of $\natural$-segments in x. That is, all the segments in $\text{top}(x)$ are to the right of all the segments in $\text{bottom}(x)$. Let a be the rightmost segment in $\text{bottom}(x)$, and let b be the leftmost segment in $\text{top}(x)$. Let c and d be the horizontal segments terminating a and b, respectively (there must be such segments since $\mathcal{V}(x) = \emptyset$). Suppose further

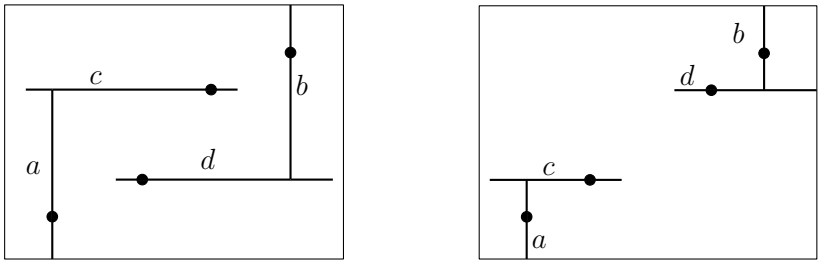

(a) c is not lower than d (b) c is lower than d

Figure 4.6: Illustrations for the proof of Proposition 4.5.4
איור 4.6: הוכחת טענה 4.5.4

that the height of c is at least the height of d (see Figure 4.6(a)). Then, as we now show, there must be a horizontal segment in left(x) whose height is at least the height of c. We traverse c to the left. If we reach the left edge of R, then c is the sought segment. Otherwise, we reach a vertical segment e that terminates c. It must be that $e \notin \text{top}(x)$, since all the segments in top(x) are to the right of a. Therefore, there is a horizontal segment f that terminates e from above. We proceed this way leftward and upward until we reach the left edge of R. Thus, there is a segment $s_1 \in \text{left}(x)$ which is not lower than c. Using the same arguments one can show that there exists a segment $s_2 \in \text{right}(x)$ which is not higher than d. Thus, s_1 and s_2 are the segments we seek.

The other case in which c is lower than d, (see Figure 4.6(b)) is handled in an similar manner. The claim follows. □

Proposition 4.5.5 *Let X be the set of all the rectangulations of (R, P) when $\pi(P) = \mathcal{I}_n$. Then there is a mapping $\psi : X \to X$ such that for every rectangulation $x \in X$:*

1. $|\mathcal{H}(x)| = |\mathcal{H}(\psi(x))|$ *and* $|\mathcal{V}(x)| = |\mathcal{V}(\psi(x))|$*; and*
2. *if* $\mathcal{F}(x) = (l, t, r, b)$ *then* $\mathcal{F}(\psi(x)) = (l, b, r, t)$.

According to the these properties, $\psi(x)$ has the same number of segments crossing from left to right and from bottom to top as x, and the same interface as x except the numbers of top-touching and bottom-touching segments which are interchanged. Note that these properties are not trivial and do not follow from simple symmetry arguments.

Proof: We will build such a mapping by induction on n. When $n = 1$ there are only two possible rectangulations, each one corresponding to itself. Assume that such a mapping ψ exists for all the rectangulations of $n' < n$ points. Let x be a rectangulation of n points arranged in the identity permutation, such that $\mathcal{F}(x) = (l, t, r, b)$. There are three cases:

1. $\mathcal{V} \neq \emptyset$;
2. $\mathcal{H} \neq \emptyset$; or
3. $\mathcal{V} = \mathcal{H} = \emptyset$.

We now describe $\psi(x)$ in each of these cases.

1. $\mathcal{V}(x) \neq \emptyset$. Let s be the leftmost segment in $\mathcal{V}(x)$. We find the corresponding rectangulations for the points to the left and to the right of s, and concatenate them to create $y = \psi(x)$ (see Figure 4.7). Clearly, $\mathcal{F}(y) = (l, b, r, t)$, $|\mathcal{H}(y)| = |\mathcal{H}(x)|$, and $|\mathcal{V}(y)| = |\mathcal{V}(x)|$.

2. $\mathcal{H}(x) \neq \emptyset$. Let s be the lowest segment in $\mathcal{H}(x)$. Let x' be the rectangulation we get when we reflect x with respect to the primary diagonal (along the points). The points of x' are arranged in the identity permutation, $\mathcal{V}(x') \neq \emptyset$ and $\mathcal{F}(x') = (b, r, t, l)$, thus x' qualifies for the previous case. Let y be the rectangulation

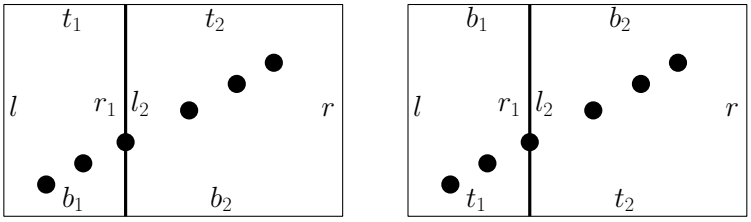

(a) x (b) $\psi(x)$

Figure 4.7: $\psi(x)$ when $\mathcal{V}(x) \neq \emptyset$
$\mathcal{V}(x) \neq \emptyset$ כאשר $\psi(x)$:4.7 איור

we get when reflecting $\psi(x')$ with respect to the secondary diagonal. Clearly, $\mathcal{F}(y) = (l, b, r, t)$, $|\mathcal{H}(y)| = |\mathcal{H}(x)|$, and $|\mathcal{V}(y)| = |\mathcal{V}(x)|$. See Figure 4.8 for an illustration of these steps.

3. $\mathcal{H}(x) = \mathcal{V}(x) = \emptyset$. In this case there are two subcases:

   (a) There is a pair of $\natural$-segments in x; or

   (b) There is no such pair of segments.

   Proposition 4.5.4 guarantees that in the second subcase there are segments $s_1 \in \text{left}(x)$ and $s_2 \in \text{right}(x)$, such that s_1 is higher than s_2. By following the same series of steps described above (see Figure 4.8), we can reduce this subcase to the first subcase.

   Let us, then, consider the first subcase of the current case. Let (a, b) be the leftmost pair of adjacent $\natural$-segments. Let $x(-a)$ be the rectangulation induced by the points to the left of a, and let $x(a-b)$ and $x(b-)$ be the rectangulations induced by the points between a and b, and the points to the right of b, respectively. We construct $\psi(x)$ by concatenating $\psi(x(-a))$, $x(a-b)$, and $\psi(x(b-))$. However, since a and b do not cut R we need to "combine" $\text{right}(\psi(x(-a)))$

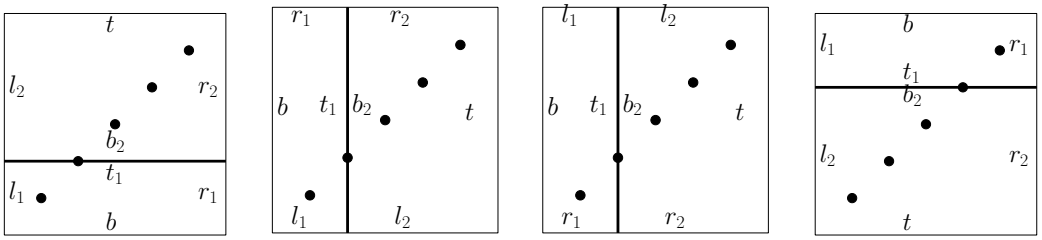

(a) x

(b) $x' := x$ reflected w.r.t. the primary diagonal

(c) $\psi(x')$

(d) $\psi(x) := \psi(x')$ reflected w.r.t. the secondary diagonal

Figure 4.8: $\psi(x)$ when $\mathcal{H}(x) \neq \emptyset$
$\mathcal{H}(x) \neq \emptyset$ **באשר** $\psi(x)$ **:4.8 איור**

with bottom$(x(a-b)) \cup \{a, b\}$, and top$(x(a-b)) \cup \{a, b\}$ with left$(\psi(x(b-)))$ in order to create a valid rectangulation. Here are the details of this combination: Suppose the ith (bottom to top) segment in right$(x(-a))$ is terminated by the jth (left to right) segment in bottom$(x(a-b)) \cup \{a, b\}$. We stretch the ith segment in right$(\psi(x(-a)))$ until the jth segment in bottom$(x(a-b)) \cup \{a, b\}$, and vice versa. We do the same in order to combine top$(x(a-b)) \cup \{a, b\}$ with left$(\psi(x(b-)))$. The result is a rectangulation y such that $\mathcal{F}(y) = (l, b, r, t)$ and $\mathcal{V}(y) = \mathcal{H}(y) = \emptyset$. Figure 4.9 shows an example of the steps in this case.

□

It is not hard to prove the next property of ψ (e.g., by induction on the number of points).

Observation 4.5.6 *ψ preserves pairs of adjacent $\natural$-segments (although their dimensions might change) and does not introduce such new pairs.*

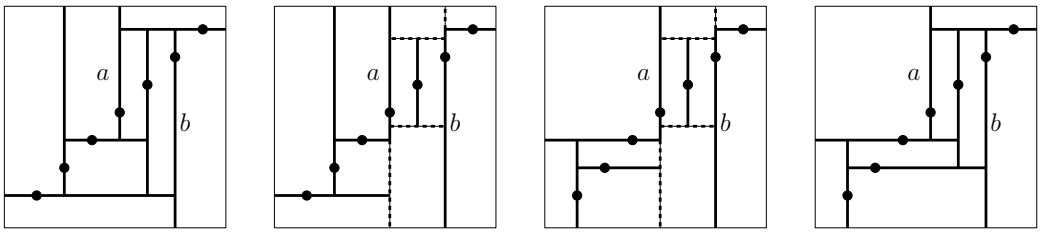

(a) x

(b) Divide x into two sub-problems

(c) Solve the sub-problems

(d) Combine the partial rectangulations

Figure 4.9: $\psi(x)$ when x contains ⊧-segments
איור 4.9: $\psi(x)$ כאשר x מכיל מקטעי-⊧

Proposition 4.5.7 *ψ is one-to-one.*

Proof: We show the claim by induction on n. For $n = 1$, ψ is one-to-one. Let us assume that ψ is one-to-one for every $n' < n$. Let x_1 and x_2 be two different rectangulations of n points, and let $y_1 = \psi(x_1)$ and $y_2 = \psi(x_2)$. We consider the different cases as in the definition of ψ.

1. $\mathcal{V}(x_1) \neq \emptyset$. If $\mathcal{V}(x_2) = \emptyset$ then clearly $y_1 \neq y_2$ since by Proposition 4.5.5 we have $\mathcal{V}(y_1) \neq \emptyset$ and $\mathcal{V}(y_2) = \emptyset$. Otherwise, if $\mathcal{V}(x_2) \neq \emptyset$, then if the leftmost vertical segment in $\mathcal{V}(x_2)$ is different from the leftmost vertical segment in $\mathcal{V}(x_1)$, then $y_1 \neq y_2$ since applying ψ on a rectangulation x does not change the leftmost vertical segment in $\mathcal{V}(x)$. If the same segment is the leftmost segment both in $\mathcal{V}(x_1)$ and $\mathcal{V}(x_2)$, then we can conclude by the induction hypothesis.

2. $\mathcal{H}(x_1) \neq \emptyset$. This case is similar to the previous case, and is thus omitted.

3. $\mathcal{H}(x_1) = \mathcal{V}(x_1) = \emptyset$. As in the definition of ψ, in this case we consider two subcases:

   (a) There is a pair of $\natural$-segments in x; or

   (b) There is no such pair of segments.

   According to Observation 4.5.6, if x_1 contains $\natural$-segments and x_2 does not, then $y_1 \neq y_2$. Assume that (a_1, b_1) and (a_2, b_2) are the leftmost pairs of adjacent $\natural$-segments in x_1 and x_2, respectively, and let p_1, q_1, p_2, q_2 be the points through which a_1, b_1, a_2, b_2 pass, respectively. If $p_1 \neq p_2$ or $q_1 \neq q_2$, then by Observation 4.5.6 $y_1 \neq y_2$. Otherwise, one of the induced rectangulations in x_1 must be different from its corresponding induced rectangulation in x_2, or x_1 and x_2 are different in the way the induced rectangulations are "combined." In the first case, it follows from the induction hypothesis and the definition of ψ that $y_1 \neq y_2$. In the second case, after applying ψ on the induced rectangulations of x_1 and x_2, they are "combined" in a similar way as in x_1 and x_2, therefore again we have $y_1 \neq y_2$.

   The second subcase is handled in a similar manner, and is thus omitted.

□

We are now able to conclude that $|\Xi(\mathcal{I}_n, (l, t, r, b))| = |\Xi(\mathcal{I}_n, (l, b, r, t))|$, for every n, l, t, r, b. □

Corollary 4.5.8 *Let $\overline{\mathcal{I}}_n$ be the reverse identity permutation on $[n]$ $(n, n-1, \ldots, 1)$, then for every n, l, t, r, b $|\Xi(\mathcal{I}_n, (l, t, r, b))| = |\Xi\left(\overline{\mathcal{I}}_n, (l, t, r, b)\right)|$.*

Proof: Let x be a rectangulation of n points in the identity permutation, such that $\mathcal{F}(x) = (l, t, r, b)$. When x is reflected with respect to the x-axis we get a rectangulation x' of n points in the reverse identity permutation, such that $\mathcal{F}(x') = (l, b, r, t)$. The corollary follows directly from this fact and from Lemma 4.5.3. □

Lemma 4.5.9 *Let π be a separable permutation of n points. Then for every interface $\mathcal{F}$, $|\Xi(\pi, \mathcal{F})| = |\Xi(\mathcal{I}_n, \mathcal{F})|$.*

Proof: By induction on n. For $n = 1$ a permutation of one point is both the identity permutation and a separable permutation. Assume the claim is true for every separable permutation of $n' < n$ points, and let π be a separable permutation of n points. π may be a concatenation-above or a concatenation-below of two separable permutations. Suppose that π is the result of concatenating a separable permutation $\pi_2 \in S_{n-k}$ above another separable permutation $\pi_1 \in S_k$. Then all the rectangulations of π can be created by considering every pair of a rectangulation of π_1 and a rectangulation of π_2, and by combining every such pair in all the possible combinations (see Figure 4.10). Note that given x_1 and x_2, rectangulations of π_1 and π_2, respectively, the number of rectangulations of π that are created by combining x_1 and x_2 in all the possible combinations depends only on $\mathcal{F}(x_1)$ and $\mathcal{F}(x_2)$. Moreover, the interface of every such combined rectangulation also depends only on $\mathcal{F}(x_1)$ and $\mathcal{F}(x_2)$ and the way they were combined.

According to the induction hypothesis, for every pair of interfaces $\mathcal{F}_1$ and $\mathcal{F}_2$ we have $|\Xi(\pi_1, \mathcal{F}_1)| = |\Xi(\mathcal{I}_k, \mathcal{F}_1)|$ and $|\Xi(\pi_2, \mathcal{F}_2)| = |\Xi(\mathcal{I}_{n-k}, \mathcal{F}_2)|$. All the rectangulations of $\mathcal{I}_n$ can be created by combining all the pairs of a rectangulation of $\mathcal{I}_k$ and a rectangulation of $\mathcal{I}_{n-k}$ in all possible combinations. Again, the number of combinations and the interface of every such combined rectangulation depends only on

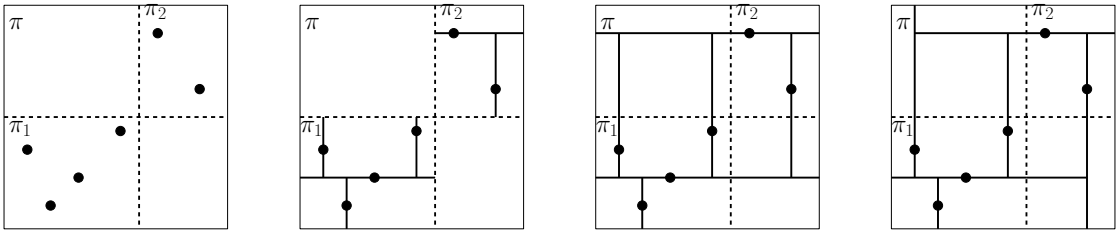

(a) π_2 concatenated above π_1

(b) Rectangulations of π_1 and π_2

(c) One combination of the rectangulations of π_1 and π_2

(d) Another combination of the rectangulations of π_1 and π_2

Figure 4.10: Rectangulations of a separable permutation
איור 4.10: מילבונים של תמורה בת-הפרדה

the interfaces of the rectangulations of $\mathcal{I}_k$ and $\mathcal{I}_{n-k}$, and on the way they were combined. Thus, for every concatenation-above separable permutation π and interface $\mathcal{F}$, $|\Xi(\pi, \mathcal{F})| = |\Xi(\mathcal{I}_n, \mathcal{F})|$.

Suppose now that π is the result of concatenating a separable permutation $\pi_2 \in S_{n-k}$ below another separable permutation $\pi_1 \in S_k$. It follows from Corollary 4.5.8 that for every pair of interfaces $\mathcal{F}_1$ and $\mathcal{F}_2$, $|\Xi(\mathcal{I}_k, \mathcal{F}_1)| = |\Xi(\overline{\mathcal{I}}_k, \mathcal{F}_1)|$ and $|\Xi(\mathcal{I}_{n-k}, \mathcal{F}_2)| = |\Xi(\overline{\mathcal{I}}_{n-k}, \mathcal{F}_2)|$. Using the induction hypothesis we conclude that for every pair of two interfaces $\mathcal{F}_1$ and $\mathcal{F}_2$, $|\Xi(\pi_1, \mathcal{F}_1)| = |\Xi(\overline{\mathcal{I}}_k, \mathcal{F}_1)|$ and $|\Xi(\pi_2, \mathcal{F}_2)| = |\Xi(\overline{\mathcal{I}}_{n-k}, \mathcal{F}_2)|$. Then, according to the combination arguments given above and by using Corollary 4.5.8, for every concatenation-below separable permutation π and interface $\mathcal{F}$, $|\Xi(\pi, \mathcal{F})| = |\Xi(\overline{\mathcal{I}}_n, \mathcal{F})| = |\Xi(\mathcal{I}_n, \mathcal{F})|$.

In conclusion, the claim holds for all separable permutations. □

Theorem 4.5.10 *Given a rectangle R which encloses a set P of n noncorectilinear points, such that $\pi(P)$ is a separable permutation on $[n]$, $|\Xi(R,P)| = B(n+1)$.*

Proof: The claim follows from Lemmata 4.5.2 and 4.5.9. □

4.6 Rectangulations and Floorplans

Recall that "point-free" rectangulations, that is, a subdivisions of a rectangle into smaller rectangles by n non-intersecting axis-parallel segments are equivalent to mosaic floorplans whose number is known to be $\mathrm{B}(n+1)$ (see Chapter 2). In this section we prove that given a set of points P in a separable permutation and a mosaic floorplan f by n segments, there is a unique way of "combining" P and f into a rectangulation.

Theorem 4.6.1 *Given a mosaic floorplan f with n segments and a set P of n points arranged in a separable permutation π, there is a unique rectangulation of P, x, such that $x \setminus P$ is equivalent to f.*

Proof: We will show that it is possible to create a rectangulation of a set of points whose permutation is π and its underlying mosaic floorplan is f. It then follows that an equivalent rectangulation can be created for P. Since by Theorem 4.5.10 the number of rectangulations of a set of n points in a separable permutation is $\mathrm{B}(n+1)$ and this is also the number of mosaic floorplans with n segments [64], it follows that the combination is unique.

We now prove by induction on n that for every mosaic floorplan f by n segments and a separable permutation $\pi \in S_n$ it is possible to create a rectangulation x of a set of n points whose permutation is π such that the underlying floorplan of x is f.

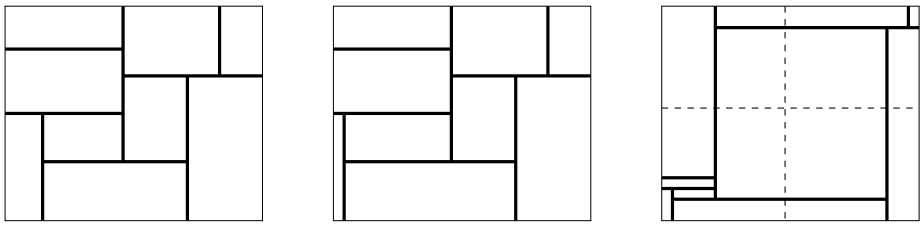

(a) (b) (c)

Figure 4.11: Illustrations for the proof of Theorem 4.6.1
איור 4.11: הוכחת משפט 4.6.1

Examining the bottom-left rectangle in f note that its top-right corner is either of the form $\dashv$ or T. In the first case by 'sliding' the horizontal segment creating the '$\dashv$'-junction downwards (resp., upwards) while 'stretching' (resp., 'shrinking') the vertical segments attached to it (if such exist) we create a floorplan which is equivalent to f. In the second case one can slide the vertical segment of the 'T'-junction leftwards or rightwards. Note that if a segment is shifted until it hits the boundary then we obtain a mosaic floorplan of $n-1$ segments.

Now suppose π can be formed by concatenating the permutation $\pi_2 \in S_{n-k}$ above the permutation $\pi_1 \in S_k$. We create two sets of the segments of f in the following manner. We start by shrinking the bottom-left rectangle in f by sliding one of its edges as described above. We stop sliding this edge when it is in a small distance $\epsilon > 0$ from the boundary (see Figures 4.11(a) and 4.11(b)), but we notice that if this edge vanishes in the boundary then it is possible to slide another segment in a similar manner.

We continue by sliding this segment until it is within a distance of 2ϵ from the boundary. Likewise, we slide each of the next $k-2$ segments: the ith segment is shifted either leftwards or downwards until it is within $i\epsilon$ distance from the boundary.

This ensures we maintain a valid floorplan equivalent to f.

In a similar way we can 'group' the other $n-k$ segments near the top-right corner of f. See Figure 6 for illustrations of this process. Now divide f into four parts by drawing a vertical and a horizontal line through its center. Every segment in the top-left and the bottom-right parts is partly contained in either the bottom-left or the top-right parts as well. Additionally, the bottom-left part is actually a floorplan with k segments whereas the top-right part is a floorplan with $n-k$ segments. By induction it is possible to embed a set of k points whose permutation is π_1 into the first floorplan, and a set of $n-k$ points whose permutation is π_2 into the second floorplan. Therefore it is possible to embed a set of n points whose permutation is π into f.

The case in which π is concatenation-below permutation is handled in a similar manner (this time the segments are grouped at the top-left and bottom-right corners). □

4.7 Conclusions

We showed that the number of rectangulations (by n segments) of a set P of n noncorectilinear points depends only on the permutation in which the points are arranged. For any arrangement, the number of guillotine rectangulations is always the nth Schröder number and the total number of rectangulation is $O(18^n/n^4)$.

For point sets in a separable permutation we proved that the number of rectangulations is the $(n+1)$st Baxter number. Moreover, for every mosaic floorplan f with n segments there is a unique way to embed a set of n points, arranged in a separable permutation, in f. This strengthens a result of de Fraysseix et al. [20]: they showed

n	B$(n+1)$	Minimum number of rectangulations	Maximum number of rectangulations
4	92	93	93
5	422	424	428
6	2,074	2,080	2,122
7	10,754	10,776	11,092
8	58,202	58,290	60,524
9	326,240	326,608	342,938

Table 4.1: Empirical results of the number of rectangulations for non-separable permutations

איור 4.11: תוצאות אמפיריות עבור מספר המילבונים של תמורות שאינן בנות-הפרדה

that every bipartite planar graph can be represented as the contact graph[4] of a set of non-intersecting vertical and horizontal segments in the plane. It follows from our results that given a set P of n noncorectilinear points in the plane, arranged in a separable permutation, and a planar bipartite graph $G = (V, E)$ such that $|V| = n$, it is possible to represent G as a contact graph of a set S of n vertical and horizontal segments such that every segment in S contains a single point from P.

Counting the number of rectangulations for non-separable permutations is still an open question. Our computations have led us to the following conjecture:

Conjecture 4.7.1 *Given a set P of n noncorectilinear points, such that $\pi(P)$ is a non-separable permutation on $[n]$, $|\Xi(P)| > B(n+1)$. Moreover, there is at least one way of embedding P in any mosaic floorplan containing n segments.*

For example, when $n = 4$ there are two non-separable permutations (3142 and 2413), and for each of them (not surprisingly, since one is the reverse of the other) the number of rectangulations is 93 (as opposed to B$(5) = 92$ for separable permutations). For $n = 5$ the number of rectangulations varies from 424 to 428 (as opposed to

[4]In a contact graph there is an edge between every two touching elements.

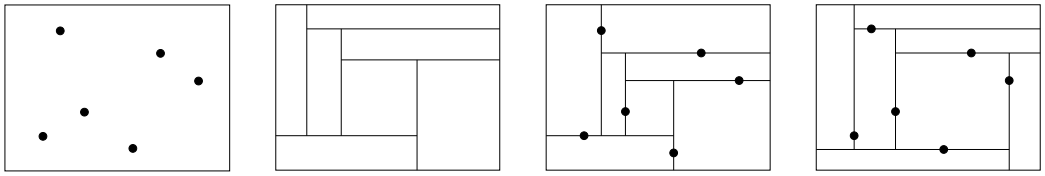

(a) A set of points P in a non-separable permutation $\pi = 263154$

(b) A mosaic floorplan f

(c) A possible embedding of P into f

(d) A different embedding of P into f

Figure 4.12: Two possible embeddings for a non-separable permutation
איור 4.12: שני שיבוצים אפשריים עבור תמורה שאינה בת-הפרדה

B(6) = 422 for separable permutations), but some values appear and some do not. Our empirical results are listed in Table 4.1.

Perhaps the extra number of rectangulations for non-separable permutations can be computed by counting the number of different ways in which they can be embedded in some mosaic floorplans. Figure 4.12 shows, for example, two possible ways of embedding a set of points in a non-separable permutation into a certain mosaic floorplan.

Other questions of interest are:

1. Improve the upper bound of $O(18^n/n^4)$, perhaps to $O(16^n/n^4)$ by showing that for every mosaic floorplan and a set of points, once the orientations of the segments through every point are set, there is at most one way of embedding the points into the floorplan.

2. What is the number of rectangulations when the problem is generalized to higher dimensions?

3. What is the complexity of the original minimum edge-length partitioning problem (RGNLP)? Furthermore, what is its computational complexity when restricted to separable (or even monotone) permutations?

Chapter 5

Quasi-Planar Graphs

5.1 Introduction

5.1.1 Graph-theory preliminaries

A *graph*, G, is a pair (V, E), where V is a set of *vertices* and E is a set of *edges*, such that each edge $e \in E$ is associated with two vertices $u, v \in V$. Two edges are *parallel* if they are associated with the same pair of vertices. An edge $e = (u, v)$ is a *self-loop* if $u = v$. Unless stated otherwise, we consider only *simple* graphs, that is, graphs with no parallel edges or self-loops. If every pair of distinct vertices are associated by an edge, then the graph is *complete*. The complete graph on n vertices is denoted by K_n. Two graphs, $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, are *isomorphic* if they have the same number of vertices and edges, and there is a mapping $\ell : V_1 \to V_2$ such that $(v_1, u_1) \in E_1$ if and only if $(\ell(v_1), \ell(u_1)) \in E_2$. A graph $G' = (V', E')$ is a *subgraph* of a graph $G = (V, E)$ if $V' \subseteq V$ and $E' \subseteq E \cap (V' \times V')$. If a graph G' is isomorphic to a subgraph of G, we say that G *contains* G'.

Extremal Graph Theory studies graphs that are extremal, in some sense, among the graphs that satisfy certain properties. For example, in [61] Turán answered the following question: What is the maximum number of edges that a graph with n vertices that does not contain K_k can have? Obviously, one can ask similar questions by replacing K_k with any fixed graph H. Such questions are known as *Turán-type* problems.

5.1.2 Geometric graph-theory preliminaries

Geometric Graph Theory studies graphs drawn in the plane. A *geometric graph* is a graph drawn in the plane such that its vertices are distinct points and its edges are straight-line segments connecting the corresponding points and containing no other vertex. A *topological graph* is defined in a similar way with the exception that the edges can take the form of any simple (i.e., non-self-intersecting) curve (*a Jordan arc*). Crossings are allowed, but we assume that each pair of edges has a finite number of common internal points and they properly cross at these points (these usual restrictions, however, are not important in this work). A topological graph is *simple* if every pair of its edges intersect at most once (either at a vertex or at a crossing point). Clearly, every geometric graph is a simple topological graph. Throughout this work we make no distinction between a topological graph and its underlying abstract graph.

Turán-type problems can naturally be asked for topological graphs: what is the maximum number of edges in a topological graph on n vertices that avoids a certain *geometric* pattern? Perhaps the simplest pattern one can consider is a pair of crossing edges. Clearly, a topological graph with no pair of crossing edges is a planar graph, and it follows from Euler's Polyhedral Formula that such a graph has at most $3n - 6$

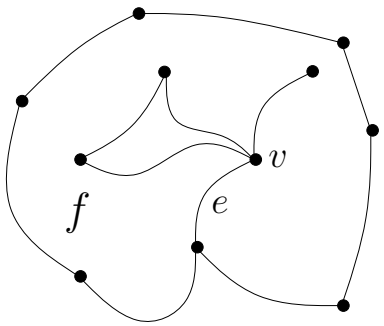

Figure 5.1: The edge e (resp., vertex v) appears twice (resp., three times) along the boundary of f

איור 5.1: צמתים וקשתות המופיעים יותר מפעם אחת לאורך פאה

edges, assuming $n \geq 3$. Let us recall the proof: Let G be a topological graph with no pair of crossing edges. Denote by V, E, and F, the vertex, edge, and face[1] sets of G, respectively. Let $|f|$ be the size of a face f, that is, the number of edges along the boundary of f (note that an edge might appear twice along the boundary of a face; see Figure 5.1 for an example). Every edge appears either along the boundary of exactly two faces, or twice along the boundary of a single face, thus we have: $\sum_{f \in F} |f| = 2|E|$. Since we consider graphs with no self-loops or parallel edges and $|V| \geq 3$, the size of every face is at least three. Thus, $3|F| \leq 2|E|$. According to Euler's formula $|V|+|F|-|E| \geq 2$, where equality holds if G is connected. Combining these two inequalities we have $|E| \leq 3|V|-6$. If G is a *triangulation*, that is, a planar graph in which the size of every face is three, then the inequalities become equalities, and therefore, the upper bound is tight.

Notice that the proof and the resulting bound still hold even if G is not a simple (abstract) graph, as long as there are no faces of size one or two. We will get back to this issue when we consider generalizations of quasi-planar graphs in Section 5.2.3.

[1] A *face* of a plane graph G is a connected region of $\mathbb{R}^2 \setminus G$.

Quasi-planar graphs are a natural generalization of planar graphs:

Definition 5.1.1 *A topological graph is* k-quasi-planar *if it does not contain* k *pairwise crossing edges (i.e., a set of* k *edges in which every pair is crossing).*

We denote by $f_k(n)$ the maximum number of edges in a k-quasi-planar graph on n vertices. Since a 2-quasi-planar graph is simply a planar graph, we have $f_2(n) = 3n - 6$. It is conjectured that for every fixed k, the maximum number of edges in a k-quasi-planar graph is linear in the number of vertices. Namely,

Conjecture 5.1.2 *For every fixed integer* $k > 1$*, there is a constant* C_k*, such that every* k*-quasi-planar graph on* n *vertices has at most* $C_k n$ *edges.*

5.1.3 Previous work

Conjecture 5.1.2 was first stated by Pach [45, Problem 3.3] as a problem suggested by Bernd Gärtner. In this paper Pach showed that $f_k(n) = O(n^{2-(1/2^{k-2})})$. Pach, Shahrokhi, and Szegedy [49] proved a better upper bound of $O(n \log^{2k-4} n)$ for simple topological graphs. Conjecture 5.1.2 was settled for simple 3-quasi-planar graphs by Agarwal, Aronov, Pach, Pollack, and Sharir [6]. Their result also implied an improvement by a factor of $\Theta(\log^2 n)$ for the upper bound for k-quasi-planar graphs. Later, Pach, Radoičić, and Tóth [48] simplified the proof of Agarwal et al. [6], and generalized it to nonsimple graphs as well. They showed that $f_3(n) \leq 65n$ and $f_k(n) = O(n \log^{4k-12} n)$. For *x-monotone* k-quasi-planar graphs, that is, graphs for which every vertical line crosses every edge at most once, Valtr [62] obtained an upper bound of $O(n \log n)$ for every fixed k.

The special case of *convex* geometric k-quasi-planar graphs (that is, graphs whose vertices are in convex position) was considered by Capoyleas and Pach [16], who

showed that in this case

$$f_k(n) = \begin{cases} \binom{n}{2} & \text{if } n \leq 2k-1, \\ 2(k-1)n - \binom{2k-1}{2} & \text{if } n \geq 2k-1. \end{cases}$$

Many other Turán-type problems on topological graphs avoiding certain geometric patterns were considered in the literature (see, e.g., [38, 47, 48, 50, 58, 59], and [46] for a survey).

In our proofs we use the *discharging method*. In this technique one usually assigns "charges" (or weights) to elements of the input (typically vertices or faces of a graph). The total charge is then computed twice: once after assigning the charges, and once more after charges have been re-distributed (*discharging* phase). Perhaps the most famous example of using the discharging method is in the proof of *The Four Color Theorem* [9].

5.2 3-Quasi-Planar Graphs

In this section we consider 3-quasi-planar graphs. Recall that it is already known that $f_3(n) = O(n)$ [6, 48]. In Section 5.2.1 we give a simple proof for this fact that yields the better upper bound of $8n$ edges for n vertices (recall that the previous best upper bound was $65n$ [48]). Our best construction, described in Section 5.2.2, with $7n - O(1)$ edges comes very close to this bound. In Section 5.2.3 we show matching upper and lower bounds for several relaxations and restrictions of the problem of the maximum number of edges in 3-quasi-planar graphs. In particular, we show that the maximum number of edges of a simple 3-quasi-planar topological graph is $6.5n - O(1)$.

5.2.1 The upper bound for $f_3(n)$

Theorem 5.2.1 *The maximum number of edges in a 3-quasi-planar graph on $n \geq 3$ vertices is at most $8n - 20$.*

Proof: Consider a 3-quasi-planar topological graph G on n vertices. We assume, without loss of generality, that among the 3-quasi-planar graphs with the same underlying abstract graph, G is chosen to have the minimum overall number of intersections. We can also assume that G is connected, otherwise the statement easily follows by induction. Let $X(G)$ denote the set of points where the edges of G intersect, and let us obtain G' by adding the vertices of $X(G)$ to the graph G and subdividing the edges of G accordingly. Note that G' is a crossing-free topological graph, i.e., a plane drawing of a planar graph. We refer to the vertices in $V(G)$ as *original* vertices, while the vertices in $X(G)$ are the *new* vertices. As no three edges of G cross at a single vertex, all new vertices are degree-4 vertices of G'. Denote by $F(G')$ the set of faces of G', and let $|f|$ be the number of edges along the boundary of a face $f \in F(G')$. (Note that it is possible for an edge of G' to appear twice along the boundary of a face, in which case it is counted with multiplicity.) Given a face f, we denote by $v(f)$ the number of original vertices that appear along the boundary of f (note that a vertex can appear more than once along the boundary of a face; again, in this case it is counted with multiplicity). We will use the terms *triangles*, *quadrilaterals*, and *pentagons* to refer to faces of size 3, 4, and 5, respectively. An integer prefix of a name of a face f denotes the number of original vertices $v(f)$ along the boundary of the face f. For example, a 2-pentagon is a face of size 5 that has two original vertices along its boundary.

We proceed by assigning charges to the faces of G' such that the charge of a face f, $ch(f)$, is $|f|+v(f)-4$. Summing the total charges over all the faces of G' we have:

$$\sum_{f\in F(G')} ch(f) = \sum_{f\in F(G')} \Big(|f|+v(f)-4\Big) = 2|E(G')| + \sum_{f\in F(G')} v(f) - 4|F(G')| = 4n-8, \tag{5.1}$$

where the last equality follows from Euler's formula and from the next equalities:

$$\sum_{f\in F(G')} v(f) = \sum_{u\in V(G)} d(u) = \sum_{u\in V(G')} d(u) - \sum_{u\in X(G)} d(u) = 2|E(G')| - 4\Big(|V(G')| - |V(G)|\Big),$$

where $d(u)$ denotes the degree of the vertex u.

Since G is drawn with the least possible number of crossings, G' does not contain faces of size one or two. We cannot have 0-triangles, as they would yield three pairwise crossing edges. Thus, every face in G' has a non-negative charge.

Next, we redistribute the charges without affecting the total charge found in (5.1). We make sure that the new charge $ch'(f)$ of a face f satisfies $ch'(f) \geq v(f)/5$, as follows. Clearly, the only faces that do not satisfy this bound with the original charge ch are 1-triangles. Let f be a 1-triangle, and let u be the original vertex of f. Let e'_1 and e'_2 be the two sides of f incident to u. Let e_1 and e_2 be the edges of G of which e'_1 and e'_2 are segments. Let us examine the faces of G' touching e_1 from the same side as f. Along e_1 we find f followed by one or more additional faces $f_1, f_2, \ldots$. Let f_i be the first of these faces that is not a 0-quadrilateral (see Figure 5.2). Clearly, such a face exists as the last f_j has at least one original vertex on its boundary. Moreover, one can easily see that we obtain the same faces $f_1, \ldots, f_i$ if we switch the roles of e_1 and e_2; therefore, f_i is well defined despite the ambiguity about e_1 and e_2. Note that f_i is not a triangle, as we do not have 0-triangles and e_1 and e_2 are not connecting the same pair of vertices. Let us shift $1/5$ unit of charge from f_i to f. We say that

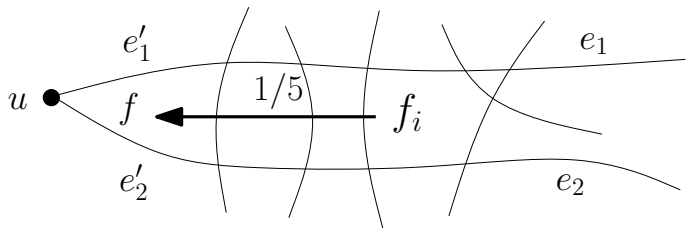

Figure 5.2: The first round of charge redistribution
איור 5.2: השלב הראשון בפיזור המטענים

f_i loses this charge *through* the edge it is sharing with f_{i-1} (or f if $i = 1$). Let ch' be the charge obtained from ch after doing this shift for all 1-triangles f.

Notice that $ch'(f) = 1/5 = v(f)/5$ for every 1-triangle f and $ch'(f) = 0 = v(f)/5$ for every 0-quadrilateral f. Let f be a face of G' that is neither a 0-quadrilateral nor a 1-triangle. We have $ch(f) = |f| + v(f) - 4 \geq 1$ and $ch'(f) = ch(f) - x_f$, where x_f is the total charge f lost in the redistribution. As f loses at most $1/5$ unit of charge through each of its edges and only if both endpoints of the edge are new, we have $x_f \leq (|f| - v(f))/5$. Thus, $ch'(f) \geq (2/5)v(f) + (4/5)(ch(f) - 1) \geq 2v(f)/5$ in this case.

Therefore, we have $ch'(f) \geq v(f)/5$ for all faces f of G'. As a second round of redistribution we collect all the extra charge around an original vertex. For every face f with $v(f) > 0$ we compute the extra charge $ch'(f) - v(f)/5$ and distribute it evenly among the $v(f)$ original vertices along the boundary of f. Each such original vertex receives $(ch'(f) - v(f)/5)/v(f)$ units of charge from f (as always, if the same vertex appears several times along the boundary it receives this charge several times). Let $ch''(f)$ be the remaining charge at a face f, and let $ch''(u)$ be the total charge

accumulated at an original vertex u. By the construction of ch'' we have

$$ch''(f) \geq v(f)/5, \tag{5.2}$$

for every face f, with equality unless $v(f) = 0$. From equation (5.1) we have

$$\sum_{f \in F(G')} ch''(f) + \sum_{u \in V(G)} ch''(u) = 4n - 8. \tag{5.3}$$

It remains to prove a lower bound on $ch''(u)$ for original vertices u. We claim

$$ch''(u) \geq 4/5 \tag{5.4}$$

To prove this claim let us consider the graph G_u obtained from G by removing the vertex u and all its incident edges. Let G'_u be the corresponding plane drawing, i.e., we introduce all crossing points as new vertices and subdivide the edges accordingly. Let f_u be the face of G'_u containing u. The size of f_u is at least 2. Here we assume G_u is not empty (if G is a star, then the statement of Theorem 5.2.1 clearly holds) and no edge is self-crossing (otherwise G does not have the minimal number of overall crossings). Let w be a vertex of f_u. There is at least one face f of G' that touches both u and w. If $|f| \geq 5$, then f alone contributes at least 1 unit of charge to u, so the claim holds. The same is true if f is a 4-quadrilateral, a 3-quadrilateral, or a 2-quadrilateral with the two original vertices not being neighbors. As f cannot be a 1-triangle, the remaining cases are a 2-triangle contributing 3/10, a 3-triangle contributing 7/15, a 1-quadrilateral contributing at least 2/5, and a 2-quadrilateral (with neighboring original vertices) contributing at least 7/10 (see Figure 5.3). One can finish the proof by a case analysis noting that (a) except for the 1-quadrilateral case, we find an edge of G incident to u that is not involved in any crossing, and therefore, the faces on both sides of this edge contribute charge to u, and (b) there

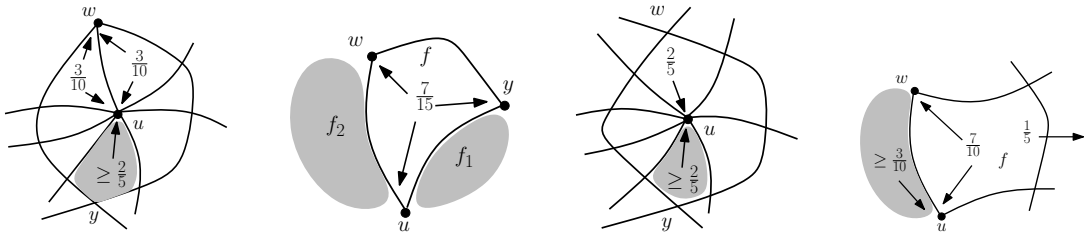

(a) (b) (c) (d)

Figure 5.3: The second round of charge redistribution
איור 5.3: השלב השני בפיזור המטענים

is at least one other vertex w' of f_u besides w. The minimal value of $ch''(u) = 4/5$ is only possible if f_u is subdivided into two 1-quadrilaterals and a few 1-triangles in G'.

Combining (5.2), (5.3), and (5.4) we have:

$$4n - 8 \geq \sum_{f \in F(G')} v(f)/5 + \sum_{u \in V(G)} \frac{4}{5} \geq \frac{2}{5}|E(G)| + \frac{4}{5}n.$$

Therefore, $|E(G)| \leq 8n - 20$. □

Remarks. Note that one can stop after the first round of charge redistribution and use Equation (5.1) together with $ch'(f) \geq v(f)/5$ to obtain $|E(G)| \leq 10n - 20$. An even simpler way to prove a linear bound on the number of edges in a 3-quasi-planar graph is the following: Let G be a 3-quasi-planar graph on n vertices. We claim $|E(G)| \leq 19n$. For the proof one can assume G is connected and $d(u) \geq 20$ for every $u \in V(G)$; otherwise we conclude by induction. Now set charges as follows: for every $u \in V(G')$ set $ch(u) := d(u) - 4$, and for every $f \in F(G')$ set $ch(f) := |f| - 4$. Then by Euler's formula the total charge is -8. However, we can distribute the charges such that every element has a non-negative charge as follows. First, every original vertex contributes $\frac{4}{5}$ units of charge to each of the faces incident to it. The only

elements with a negative charge after this step are 1-triangles. Each 1-triangle now obtains another $\frac{1}{5}$ unit of charge in the way described in the proof of Theorem 5.2.1 and ends up with a zero charge.

5.2.2 The lower bound for $f_3(n)$

Agarwal et al. [6] argued that for sufficiently large n one can always construct a 3-quasi-planar graph with n vertices and roughly $6n$ edges, by considering two (almost) edge-disjoint triangulations of the same set of n points. Next we show that $f_3(n) \geq 7n - O(1)$.

Theorem 5.2.2 *For every sufficiently large integer n, there is a 3-quasi-planar graph on n vertices with $7n - O(1)$ edges.*

Proof: Figure 5.4 shows a construction that yields a parametric family of n-vertex 3-quasi-planar graphs with $7n-O(1)$ edges, as follows: First, consider a hexagonal grid such that for each hexagon we draw all the diagonals but one as straight line edges. Then we add a curved edge for the missing diagonal as in Figure 5.4(a). Finally, we add longer edges, two from every vertex as in Figure 5.4(b). This 3-quasi-planar graph already has $7n - O(\sqrt{n})$ edges. To improve the error term we wrap the graph around a cylinder so that we have three hexagons around the cylinder and draw five more edges on each of the top and bottom faces. Considering m layers of hexagons we have $n = 6m+6$ vertices and $7n-29$ edges. For m not divisible by 6, the constant term gets slightly worse. □

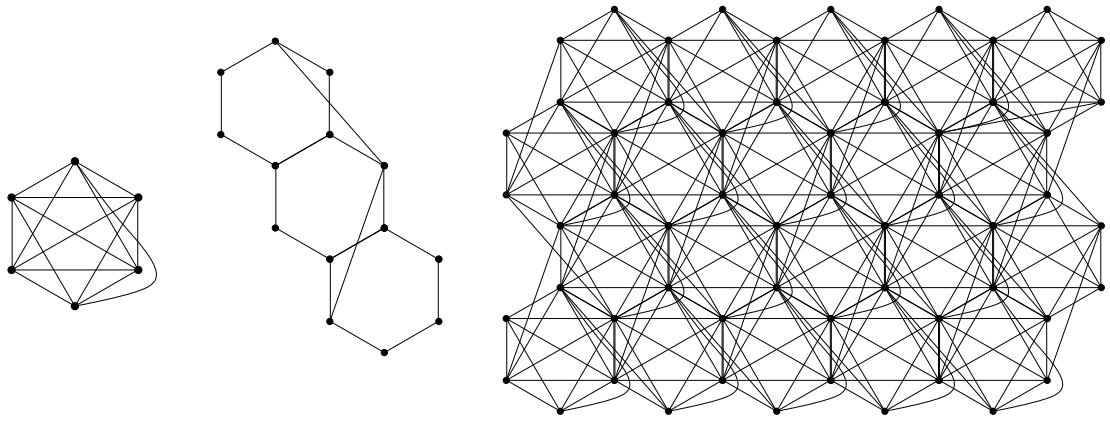

(a) (b) (c)

Figure 5.4: A construction for a 3-quasi-planar graph with n vertices and $7n - O(1)$ edges

איור 5.4: בנייה של גרף 3-מעין-מישורי עם n צמתים ו-$7n - O(1)$ קשתות

5.2.3 Generalizations and restrictions

Observe that the proof of Theorem 5.2.1 did not use our assumptions defining 3-quasi-planar graphs in their full generality. We assumed that they—as topological graphs—do not contain several edges connecting the same pairs of vertices. This assumption may seem to be crucial as one can have an arbitrary number of parallel edges even in a planar graph. However, recall that Euler's formula proves the $3n - 6$ bound on the number of edges in any plane drawing as long as there is no face bounded by two edges (see Section 5.1.2). In other words, we can allow for parallel edges as long as they are drawn with at least one vertex in any 2-gon they determine. This generalization of the bound on the number edges of a planar graph is often useful. We can generalize Theorem 5.2.1 the same way and the proof presented in the previous section will still apply. For the generalization of topological graphs where parallel

edges are allowed, we use the term *generalized topological graph*.

We have not used the 3-quasi-planarity assumption in full generality either. All we used is the assumption that the drawing does not yield 0-triangles, i.e., three pairwise crossing edges determining an empty triangle with no vertex inside. Here we consider three edges crossing each other at a single point an "infinitesimal" 0-triangle, and so we do not allow this either. We, therefore, have the following generalization of Theorem 5.2.1:

Theorem 5.2.3 *Consider a generalized topological graph G on $n \geq 3$ vertices. Assume that*

- *G has no self-loops;*
- *any 2-gon formed by two (possibly intersecting) parallel edges contains a vertex inside; and*
- *any triangle formed by three pairwise intersecting edges has a vertex inside.*

Then G has at most $8n - 20$ edges.

This generalization of Theorem 5.2.1 is surprisingly tight.

Theorem 5.2.4 *There exist generalized topological graphs satisfying the conditions of Theorem 5.2.3 with n vertices and $8n - O(1)$ edges.*

Proof: We start with the construction proving Theorem 5.2.2 and add more edges. For each of the curved edges we add a parallel edge that is drawn as the central reflection of the original one (see Figure 5.5(a)). Each pair of parallel edges yields a 2-gon with two vertices inside. Then we further add the (straight) edges as shown in Figure 5.5(b). Each of these new edges will create two sets of three pairwise crossing

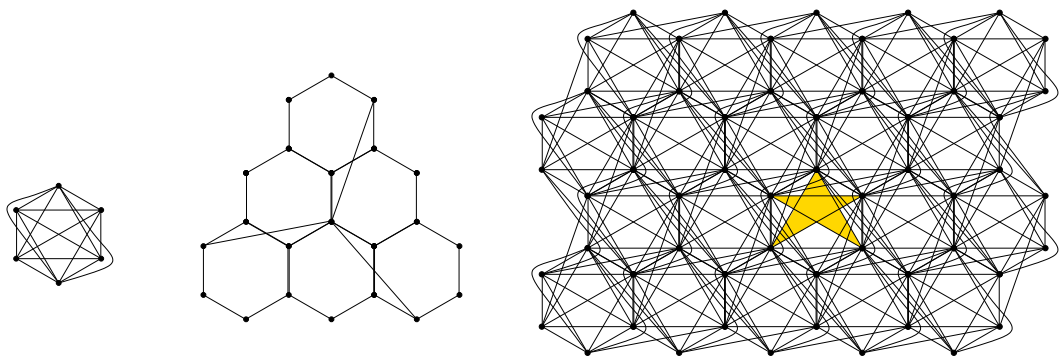

(a) (b) (c)

Figure 5.5: A generalized topological graph satisfying the conditions of Theorem 5.2.3

איור 5.5: גרף טופולוגי-מוכלל המקיים את התנאים במשפט 5.2.3

edges with exactly one vertex in either triangle. See Figure 5.5(c) for the resulting graph. The total number of edges is $8n - O(1)$. □

Remarks. It is interesting to see how the charges are distributed in the graph constructed in the proof of Theorem 5.2.4. All the faces f end up with a charge of $v(f)/5$, so in particular all 0-faces are 0-quadrilaterals or 0-pentagons. Moreover, the original edges involved in any 0-pentagon form a 5-star with the 0-pentagon in the middle and each arm consisting of a few 0-quadrilaterals and a 1-triangle (see Figure 5.5(c)). All original vertices u (except for a few close to the boundary) are surrounded by a 2-gon split up into many 1-triangles and two 1-quadrilaterals, so they end up with the minimal charge of 4/5.

We may allow for only one of (a) parallel edges with at least one vertex in the 2-gon or (b) three pairwise crossing edges if the resulting triangle has a vertex inside. For both resulting generalizations of the 3-quasi-planar graphs we still have that

$8n - 20$ is an upper bound on the number of edges but the obvious modification of the construction in the proof of Theorem 5.2.4 gives only examples with $7.5n - O(1)$ edges.

An alternate approach is to consider natural restrictions of the original problem. Recall that a topological graph is *simple* if every pair of its edges meet at most once (either at an endpoint or at a crossing point).

Theorem 5.2.5 *The maximum number of edges in a simple* 3*-quasi-planar graph on* $n \geq 4$ *vertices is at most* $6.5n - 20$*, and this bound is tight up to some additive constant.*

Proof: The proof is very similar to the proof of Theorem 5.2.1, so we omit most of the details. In this case we can derive a stronger bound since it can be shown that for every vertex $u \in V(G)$ we have $ch''(u) \geq \frac{7}{5}$. The main observation is that the face f_u of G'_u containing u cannot be a 2-gon because G is a simple topological graph, and it can be a triangle, but only if one of the vertices of the triangle is an original vertex (otherwise we would have three pairwise crossing edges). In the extremal configuration yielding charge $ch''(u) = \frac{7}{5}$, u is surrounded by two 2-triangles, two 1-quadrilaterals and any number of 1-triangles.

The matching lower bound is obtained from the same construction shown in Figure 5.4, without adding the curved edges. □

Note that this bound is tight also for geometric graphs. To see this, one has to make the construction consist of straight line edges. It naturally does consist of straight line edges before we wrap it around a cylinder. The wrapping can also be managed with straight line edges if we use five or more hexagons in every layer instead of three.

The pair of Theorems 5.2.2 and 5.2.5 seems to give the first clear distinction within the extremal function of a forbidden configuration among topological graphs and simple topological graphs.

As a combination of generalization and relaxation, we state the following result:

Theorem 5.2.6 *Let G be a simple topological graph on $n \geq 4$ vertices and with no three pairwise crossing edges that form an empty triangular face. Then the number of edges in G is at most $7n - 20$, and this bound is tight up to some additive constant.*

Proof: Once again the upper bound is proven following the proof of Theorem 5.2.1, and noticing that for every vertex $u \in V(G)$, after re-distributing the charges for the second time, we have $ch''(u) \geq \frac{6}{5}$. Now the minimal contribution to an original vertex comes from three 1-quadrilateral faces.

For the lower bound we use the construction described in Figure 5.5, this time using no curved edges. □

Finally, we mention that the construction for Theorem 5.2.5 has many edges that are not crossed by any other edge. This is no coincidence, since the minimal vertex-charge of 7/5 can only be achieved with such edges. If one only considers simple 3-quasi-planar graphs with *all edges crossed*, then one can prove an upper bound of $6n - O(1)$ on the number of edges. This is tight as witnessed by the construction presented in the proof of Theorem 5.2.5 with the non-crossed edges removed.

5.3 4-Quasi-Planar Graphs

The best upper bounds for the number of edges in a 4-quasi-planar graph on n vertices were: $O(n \log n)$ for geometric graphs [62]; $O(n \log^2 n)$ for simple topological

graphs [6]; and $O(n \log^4 n)$ for topological graphs. In this section we show that $f_4(n) = O(n)$, thus settling Conjecture 5.1.2 for the case $k = 4$.

Theorem 5.3.1 *For any integer $n > 2$, every topological graph on n vertices with no four pairwise crossing edges has at most $36(n-2)$ edges.*

Proof: It is easy to see that $f_4(3) = 3 < 36(3-2)$. Let G be a topological graph on $n > 3$ vertices and without four pairwise crossing edges. We denote by $V(G)$ the vertex set of G, and by $E(G)$ the edge set of G. Given an edge $e \in E(G)$ and two points p and q on e, we will use the notation $e|_{p,q}$ to denote the segment of e between p and q. For a vertex v we denote by $d(v)$ the degree of v. If there is a vertex $v \in V(G)$ such that $d(v) = 1$, then we can conclude the theorem by induction. Hence, we assume the degree of every vertex in G is at least two. Assume, w.l.o.g., that G is drawn with the least possible number of crossings (still, as a 4-quasi-planar graph), and that there are no three edges crossing at the same point. Let e_1 and e_2 be two edges of G that intersect at least twice. A region bounded by segments of e_1 and e_2 that connect two consecutive intersection points is called a *lens*. We observe, as in [48], that G has no *empty* lenses, that is, lenses that do not contain a vertex of G. If there were empty lenses, then the number of crossings in G could be reduced. For the same reason G does not contain self-intersecting edges.

We proceed, as in the proof of Theorem 5.2.1, by considering G', the plane graph induced by G. That is, $V(G') = V(G) \cup X(G)$, where $X(G)$ is the set of crossing points in G; and $e' \in E(G')$ if e' is a segment of an edge of G that connects two vertices in $V(G')$ and contains no other vertex from $V(G')$. We refer to the edges of G' as *p-edges*, in order to distinguish them from the edges of G. Again, we set the charge of every face f to be $|f| + v(f) - 4$, and notice that the overall charge is $4n - 8$.

A *wedge* is a triplet $w = (v, l, r)$, such that $v \in V(G)$, l and r are edges incident to v, and l immediately follows r in a clockwise order of the edges touching v. Our plan is to re-distribute the charges, such that there will be no faces with a negative charge and every wedge will be charged with at least $\frac{1}{18}$ units of charge. Then, the total charge over all the wedges will be $\frac{2|E(G)|}{18} \leq 4n - 8$, and the theorem will follow. (Note that the number of wedges in which a vertex $v \in V(G)$ participates is $d(v)$. Here we use the assumption that the degree of every vertex is at least two.) Since a face of size one yields a self-intersecting edge, and a face of size two yields an empty lens or two parallel edges, the only faces with a negative charge are 0-triangles. We proceed by describing a method to charge these faces. Then, we will show how to charge the wedges of original vertices.

Charging 0-triangles. Each 0-triangle obtains one unit of charge in the same way a 1-triangle obtained 1/5 units of charge in the proof of Theorem 5.2.1. Let t be a 0-triangle, let e_1 be one of its p-edges, and let f_1 be the other face incident to e_1 (see Figure 5.6(a)). It must be that $|f_1| > 3$, for otherwise there would be an empty lens, or two parallel edges. If $v(f_1) > 0$ or $|f_1| > 4$, we move $\frac{1}{3}$ units of charge from f_1 to t, and say that f_1 *contributed* $\frac{1}{3}$ units of charge to t *through* e_1. Otherwise, f_1 must be a 0-quadrilateral. Let e_2 be the opposite p-edge to e_1 in f_1, and let f_2 be the other face incident to e_2. Applying the same arguments as above, we conclude that either f_2 contributes $\frac{1}{3}$ units of charge to t through e_2, or f is also a 0-quadrilateral. In the second case we continue to the next face, that is, to the other face that is incident to the opposite p-edge to e_2 in f_2. However, at some point we must encounter a face that is not a 0-quadrilateral. Denote by f_i this face, by f_{i-1} the face preceding f_i, and by e_i the edge incident to both of these faces. Then f_i will contribute $\frac{1}{3}$ units of

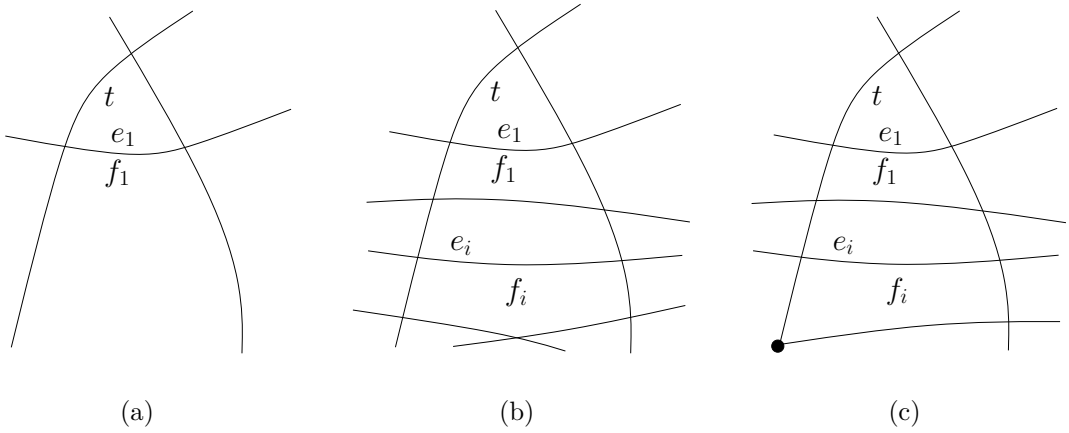

Figure 5.6: Charging a 0-triangle
איור 5.6: טעינת 0-משולש

charge to t through e_i (see Figures 5.6(b,c)).

In a similar way t obtains $\frac{2}{3}$ units of charge from its other p-edges. Thus, after re-distributing charges this way, the charge of every 0-triangle is 0. Note that a face can contribute charge through each of its p-edges at most once. Therefore, every face f such that $|f| + v(f) \geq 6$ still has a non-negative charge. It remains to verify that 1-quadrilaterals and 0-pentagons, which had only one unit of charge to contribute, also have a non-negative charge. Indeed, a 1-quadrilateral contributes to at most two 0-triangles, since the endpoints of a p-edge through which it contributes must be vertices from $X(G)$. A 0-pentagon, on the other hand, can contribute to at most three 0-triangles by the following easy observation.

Observation 5.3.2 *A* 0*-pentagon contributes charge to at most three* 0*-triangles. Moreover, if it contributes to three* 0*-triangles, then the contribution must be done through consecutive p-edges.*

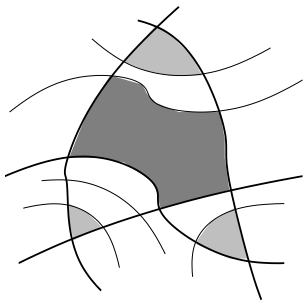

Figure 5.7: A 0-pentagon contributing to three 0-triangles through non-consecutive p-edges implies four pairwise crossing edges
איור 5.7: 0-מחומש התורם לשלושה 0-משולשים דרך צלעות לא עוקבות מעיד על ארבע קשתות נחתכות הדדית

Proof: One can easily check that a contribution to three 0-triangles through non-consecutive p-edges implies four pairwise crossing edges (see Figure 5.7). If a 0-pentagon contributes to more than three 0-triangles, then there must be three non-consecutive p-edges through which it contributes. □

Charging wedges. After the previous step, the faces with a zero charge, apart from 0-triangles and 0-quadrilaterals, are 0-pentagons that contributed to three 0-triangles, and 0-hexagons that contributed to six 0-triangles. We call such faces *bad* faces. Faces that have a positive charge are called *good* faces. Our goal now is to find some extra charge for each wedge. This extra charge will be found next to a *farthest uncut $\mathcal{A}$-crossing* or *$\mathcal{X}$-crossing* of the wedge. At this point we require a few new definitions.

Let $w = (v, l, r)$ be a wedge, and let e be an edge crossing l at p and r at q, such that $e|_{p,q}$ does not cross l or r. We denote by $w|_{e,p,q}$ the area to the left of the closed curve formed by traversing from v on $l|_{v,p}$, $e|_{p,q}$, and $r|_{q,v}$.

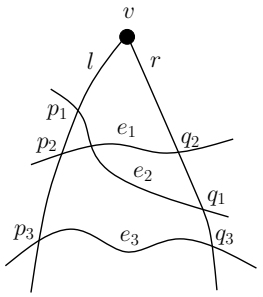

(a) $\mathcal{A}$-crossings of a wedge. (e_3, p_3, q_3) is uncut and is farther than (e_1, p_1, q_1) and (e_2, p_2, q_2).

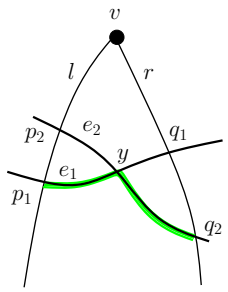

(b) an $\mathcal{X}$-crossing of a wedge. The visible part of the $\mathcal{X}$-crossing is marked.

Figure 5.8: Crossing patterns of a wedge
wedge איור 5.8: חיתוכים של

Definition 5.3.3 (uncut $\mathcal{A}$-crossing) *Let $w = (v, l, r)$ be a wedge. An* $\mathcal{A}$-crossing *of w is a triplet $cr = (e, p, q)$ such that e is an edge crossing l at p and r at q such that: (1) $e|_{p,q}$ does not intersect l or r; and (2) the endpoints of e are not in $w|_{e,p,q}$. We say that cr is an* uncut $\mathcal{A}$*-crossing of w if $e|_{p,q}$ is a p-edge. For an illustration, refer to Figure 5.8(a).*

Let $cr = (e, p, q)$ be an $\mathcal{A}$-crossing of a wedge $w = (v, l, r)$. We use the notation $w|_{cr}$ as an abbreviation for $w|_{e,p,q}$, and say that cr is an *empty* $\mathcal{A}$-crossing of w if there are no original vertices in $w|_{cr}$. Given another $\mathcal{A}$-crossing of w, $cr' = (e', p', q')$, we say that cr is *farther* than cr' (and cr' is *closer* than cr), if $p' \in l|_{v,p}$ and $q' \in r|_{v,q}$. Clearly, not every two $\mathcal{A}$-crossings of a wedge are comparable, but uncut $\mathcal{A}$-crossings are.

Definition 5.3.4 ($\mathcal{X}$-crossing) *Let w be a wedge, and let $cr_1 = (e_1, p_1, q_1)$ and $cr_2 = (e_2, p_2, q_2)$ be two $\mathcal{A}$-crossings of w. Then (cr_1, cr_2) is an* $\mathcal{X}$-crossing *of w if $e_1|_{p_1,q_1}$ and $e_2|_{p_2,q_2}$ intersect exactly once.*

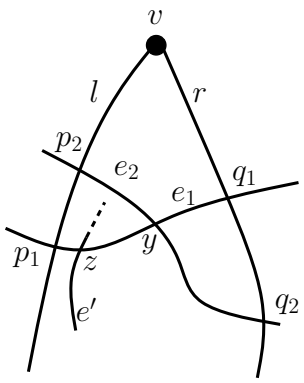

Figure 5.9: An illustration for the proof of Observation 5.3.5
איור 5.9: הוכחת הבחנה 5.3.5

Let $x = \big(cr_1 = (e_1, p_1, q_1), cr_2 = (e_2, p_2, q_2)\big)$ be an $\mathcal{X}$-crossing of a wedge $w = (v, l, r)$. Then the notation $w|_x$ represents the area $w|_{cr_1} \cup w|_{cr_2}$. Assuming $p_2 \in l|_{v,p_1}$ (and therefore, $q_1 \in r|_{v,q_2}$), the boundary of $w|_x$ is the closed curve formed by $l|_{v,p_1}$, $e_1|_{p_2,y}$, $e_2|_{y,q_2}$, and $r_{q_2,v}$, where y is the intersection point of $e_1|_{p_1,q_1}$ and $e_2|_{p_2,q_2}$. We call the curve $e_1|_{p_1,y} \cup e_2|_{y,q_2}$ the *visible* part of $w|_x$ and denote it by $\mathrm{Vis}(x)$, when it is clear to which wedge we refer. We denote by $\mathrm{Vis}(x)_l$ and $\mathrm{Vis}(x)_r$ the two components of $\mathrm{Vis}(x)$, $e_1|_{p_1,y}$ and $e_2|_{y,q_2}$, respectively. See Figure 5.8(b) for an example. The next observation will be useful later on.

Observation 5.3.5 *Let $x = \big(cr_1 = (e_1, p_1, q_1), cr_2 = (e_2, p_2, q_2)\big)$ be an empty $\mathcal{X}$-crossing of a wedge $w = (v, l, r)$, such that $\mathrm{Vis}(x)_l \subset e_1$. Then an edge e' that crosses $\mathrm{Vis}(x)_l$ (resp., $\mathrm{Vis}(x)_r$) must cross l (resp., r) and must not cross e_2 (resp., e_1).*

Proof: Let y be the crossing point of $e_1|_{p_1,q_1}$ and $e_2|_{p_2,q_2}$, and let z be the crossing point of e' and $e_1|_{p_1,y}$ (see Figure 5.9). Since x is an empty $\mathcal{X}$-crossing of w, e' must cross the boundary of $w|_x$ at least one more time. If it crosses $e_1|_{p_1,q_1}$ at any point other than z, we have an empty lens. Therefore, e' cannot cross $\mathrm{Vis}(x)$ at another

point. Otherwise, if e' crosses $r|_{v,q_2}$, then it must also cross $e_2|_{p_2,y}$. This implies four pairwise crossing edges: e', e_1, e_2, and r. Thus e' must cross l. Moreover, it must not cross e_2 since this also yields four pairwise crossing edges. The proof for an edge crossing $\text{Vis}(x)_r$ is similar and is thus omitted. □

Let x and x' be two $\mathcal{X}$-crossings of a wedge w. We say that x is a *farther* $\mathcal{X}$-crossing of w than x' if one $\mathcal{A}$-crossing of x is farther than one $\mathcal{A}$-crossing of x' and the other $\mathcal{A}$-crossing of x' is not farther than the other $\mathcal{A}$-crossing of x. In a similar way, we say that an uncut $\mathcal{A}$-crossing is *farther* (resp., *closer*) than an $\mathcal{X}$-crossing of the same wedge if it is farther (resp., closer) than both $\mathcal{A}$-crossings of the $\mathcal{X}$-crossing.

Given a wedge $w = (v, l, r)$, we look for an empty uncut $\mathcal{A}$-crossing or $\mathcal{X}$-crossing of w, such that there is no empty uncut $\mathcal{A}$-crossing or $\mathcal{X}$-crossing farther than it. If there is no such uncut $\mathcal{A}$-crossing or $\mathcal{X}$-crossing, then the face incident to v,l, and r is not a 1-triangle. Thus, its charge is at least $\frac{1}{3}$ units, from which we use $\frac{1}{18}$ units to charge w. If there is such an empty uncut $\mathcal{A}$-crossing $cr = (e, p, q)$, then let f be the face incident to $e|_{p,q}$ outside $w|_{cr}$. f is not a triangle as this would yield an empty lens or parallel edges, nor can it be a 0-quadrilateral since this would imply an empty uncut $\mathcal{A}$-crossing farther than cr. If f is a bad pentagon, then it follows from Observation 5.3.2 that there is an empty $\mathcal{X}$-crossing farther than cr. Therefore, f must be a good face that will contribute $\frac{1}{18}$ units of charge to w through $e|_{p,q}$. It remains to consider the case in which there is an empty $\mathcal{X}$-crossing x, such that there is no empty uncut $\mathcal{A}$-crossing or $\mathcal{X}$-crossing farther than x.

Denote by $cr_1 = (e_1, p_1, q_1)$ and $cr_2 = (e_2, p_2, q_2)$ the two $\mathcal{A}$-crossings forming x, such that $\text{Vis}(x)_l = e_1|_{p_1,y}$, where y is the intersection point of $e_1|_{p_1,q_1}$ and $e_2|_{p_2,q_2}$. Let f_1 be the face that is incident to y and outside $w|_x$. Suppose f_1 is a 0-triangle and let e_3 be the third edge incident to it (see Figure 5.10(a)). It follows from

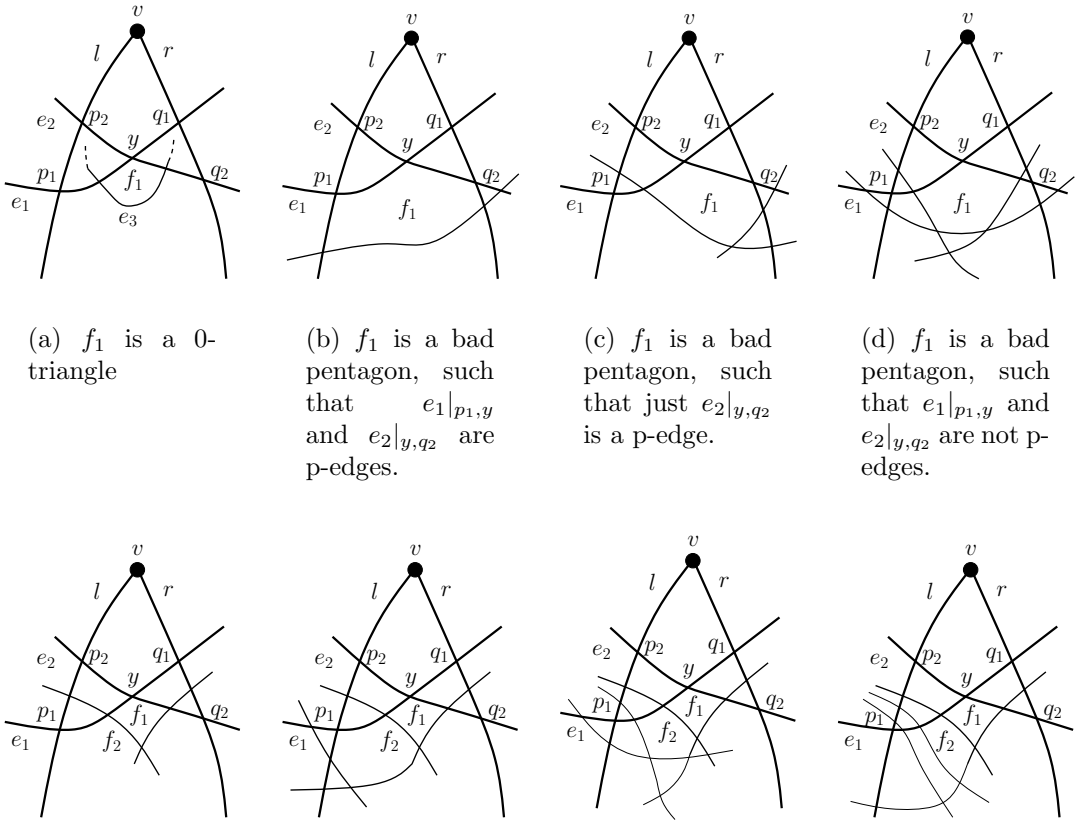

(a) f_1 is a 0-triangle

(b) f_1 is a bad pentagon, such that $e_1|_{p_1,y}$ and $e_2|_{y,q_2}$ are p-edges.

(c) f_1 is a bad pentagon, such that just $e_2|_{y,q_2}$ is a p-edge.

(d) f_1 is a bad pentagon, such that $e_1|_{p_1,y}$ and $e_2|_{y,q_2}$ are not p-edges.

(e) f_2

(f) f_2 is a bad pentagon

(g) f_2 is a bad pentagon

(h) All the faces along $\text{Vis}(x)_l$ are 0-quadrilaterals

Figure 5.10: Obtaining a charge near a farthest $\mathcal{X}$-crossing

איור 5.10: השגת מטען ליד $\mathcal{X}$-crossing רחוק ביותר

Observation 5.3.5 that e_3 must cross l, thus l, e_1, e_2, and e_3 are pairwise crossing. If f_1 is a bad pentagon, then we consider the possible cases according to whether none, one, or both of $e_1|_{p_1,y}$ and $e_2|_{y,q_2}$ are p-edges. In case both of them are p-edges (see Figure 5.10(b)), then there is an empty uncut $\mathcal{A}$-crossing of w that is farther than x. In case one of them, say $e_2|_{y,q_2}$, is a p-edge (see Figure 5.10(c)), then there is an empty $\mathcal{X}$-crossing farther than x. If none of them is a p-edge (see Figure 5.10(d)), then there must be four pairwise crossing edges. In a similar way, if f_1 is a bad hexagon, then there must be four pairwise crossing edges, or an $\mathcal{X}$-crossing of w farther than x.

Therefore, if f_1 is not a good face, then it must be a 0-quadrilateral. Suppose f_1 is a 0-quadrilateral and let f_2 be the face outside of $w|_x$ that shares a p-edge with f_1 and is incident to $\text{Vis}(x)_l$ (see Figure 5.10(e)). If there is no such face, or there is no face outside $w|_x$ that shares a p-edge with f_1 and is incident to $\text{Vis}(x)_r$, then there is an empty $\mathcal{X}$-crossing farther than x. Examining f_2, one can see by inspection that it cannot be a 0-triangle, as this implies four pairwise crossing edges. If f_2 is a bad pentagon (see Figure 5.10(f,g)) or a bad hexagon, then again there must be four pairwise crossing edges or an empty $\mathcal{X}$-crossing farther than x. Thus, either f_2 is a good face or it is a 0-quadrilateral. In the second case, we examine the next face, that is, the face $f_3 \neq f_1$ such that f_3 is outside $w|_x$, shares a p-edge with f_1, and is incident to $\text{Vis}(x)_l$. If there is no such face, then we have an empty $\mathcal{X}$-crossing farther than x. Otherwise, we can apply the same arguments we used for f_2 on f_3, proceed to the next face if f_3 is not a good face, and so on. At some point we must encounter a good face, for otherwise we have an empty $\mathcal{X}$-crossing farther than x (see Figure 5.10(h)). Let f_i be the first good face we encounter along $\text{Vis}(x)_l$, and let e_i be the p-edge of f_i that is contained in $\text{Vis}(x)_l$. Then f_i contributes $\frac{1}{18}$ units of charge to w through e_i.

Next, we prove that after charging every wedge of an original vertex, as above,

there are no faces with a negative charge. For that, we need to show that a face cannot contribute to "too many" wedges. For a (good) face f and one of its p-edges m, we say that f is a *possible $\mathcal{X}$-contributor* to a wedge w through m if there is an empty $\mathcal{X}$-crossing of w, x, such that f is outside $w|_x$ and $m \subset \text{Vis}(x)_l$.

Observation 5.3.6 *Let f be a face and let m be one of its p-edges. Then f is a possible $\mathcal{X}$-contributor through m to at most one wedge.*

Proof: Suppose there is a face f that is a possible $\mathcal{X}$-contributor through one of its p-edges, m, to two wedges, $w_1 = (v_1, l_1, r_1)$ and $w_2 = (v_2, l_2, r_2)$. Let e be the edge containing m; then there are four points, p_1, q_1, p_2, q_2, such that $cr_1 = (e, p_1, q_1)$ is an $\mathcal{A}$-crossing of w_1 and $cr_2 = (e, p_2, q_2)$ is an $\mathcal{A}$-crossing of w_2. Denote by $x_1 = \big((e, p_1, q_1), (e'_1, p'_1, q'_1)\big)$ the empty $\mathcal{X}$-crossing of w_1, such that $m \subset \text{Vis}(x_1)_l$, and by $x_2 = \big((e, p_2, q_2), (e'_2, p'_2, q'_2)\big)$ the empty $\mathcal{X}$-crossing of w_2, such that $m \subset \text{Vis}(x_2)_l$. Suppose we sort p_1, q_1, p_2, q_2 in the order in which they appear when traversing e from one of its endpoints to the other, such that when traversing m the face f is to our right. Then, p_i must precede q_i, for $i = 1, 2$, since f is outside of $w|_{x_i}$. Assume, w.l.o.g., that p_1 precedes p_2. It follows from Observation 5.3.5 that l_2 crosses l_1 (see Figure 5.11(a)). Since $m \subset e|_{p_1,q_1} \cap e|_{p_2,q_2}$, the order of the four points is either p_1, p_2, q_1, q_2 or p_1, p_2, q_2, q_1. Let us consider these cases:

Case 1: Suppose q_1 precedes q_2 (see Figure 5.11(b)). Since $w_1|_{x_1}$ and $w_2|_{x_2}$ do not contain any original vertex, l_2 must cross l_1 (one more time) and r_1 (see Figure 5.11(c)), or r_2 must cross l_1 and r_1 (see Figure 5.11(d)). The first case yields an empty lens. In the second case, note that e'_1 must cross either l_2 (see Figure 5.11(e)) or r_2 (see Figure 5.11(f)), yielding four pairwise crossing edges.

Case 2: Suppose q_1 precedes q_2 (see Figure 5.11(g)). Then the edge e'_2 must cross e

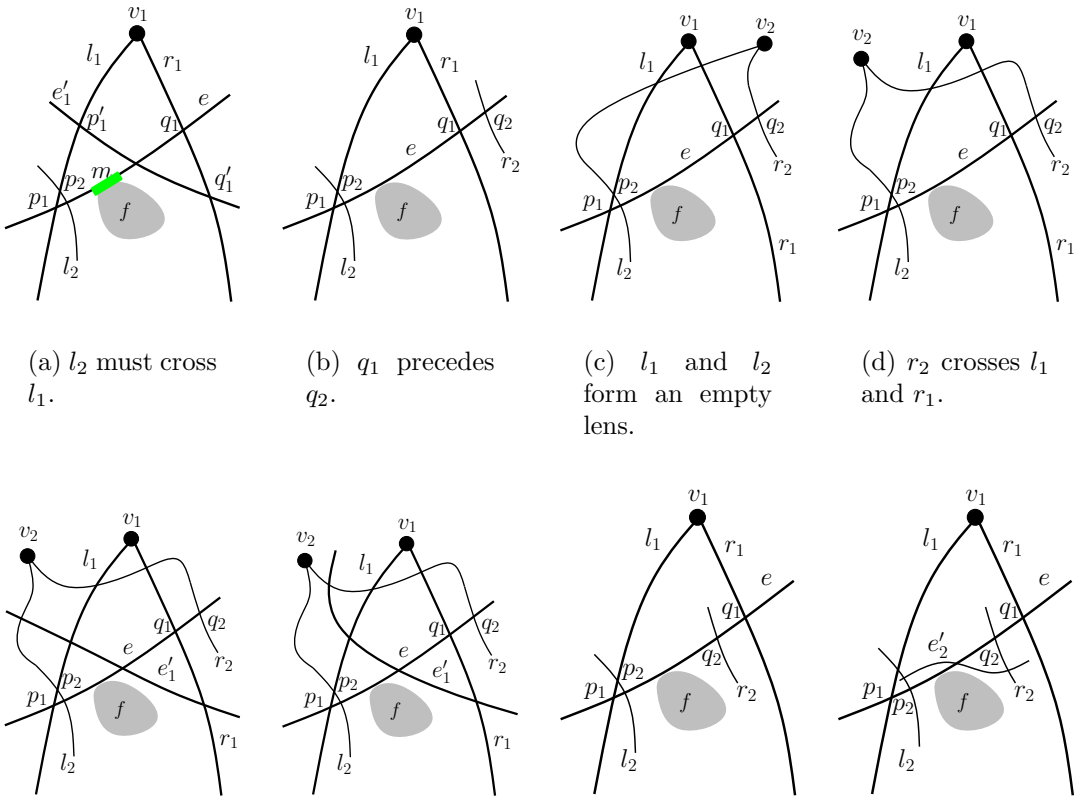

(a) l_2 must cross l_1.

(b) q_1 precedes q_2.

(c) l_1 and l_2 form an empty lens.

(d) r_2 crosses l_1 and r_1.

(e) e'_1 crosses l_2.

(f) e'_1 crosses r_2.

(g) q_2 precedes q_1.

(h) e'_2 must cross l_1.

Figure 5.11: Illustrations for the proof of Observation 5.3.6

איור 5.11: הוכחת הבחנה 5.3.6

twice, creating an empty lens, or cross l_1, yielding four pairwise crossing edges (see Figure 5.11(h)).

Since all the cases imply forbidden configurations (an empty lens or four pairwise crossing edges), we conclude that f cannot be a possible $\mathcal{X}$-contributor through m to more than one wedge. □

It follows from Observation 5.3.6 that a face f, such that $|f| \geq 7$, ends up with a charge of at least $|f| - 4 - |f|(\frac{1}{3} + \frac{1}{18}) > 0$. Likewise, k-quadrilaterals, for $k > 0$, and good pentagons and hexagons end up with a non-negative charge, as their charge after charging the 0-triangles was at least $\frac{1}{3}$. Summing up the charges over all the wedges we have $\frac{2|E(G)|}{18} \leq 4n - 8$; hence $|E(G)| \leq 36n - 72$. □

5.4 Conclusions and Open Problems

We have shown that $7n - O(1) \leq f_3(n) \leq 8n - O(1)$. So far we have been unable to close the small gap between the lower bound and the upper bound. The upper bound on $f_3(n)$ can be used to obtain a lower bound for the number of 3-tuples of 3-pairwise crossing edges in any drawing of a given graph G.

Corollary 5.4.1 *Let G be a graph with n vertices and m edges, and let $cr_3(G)$ be the minimum possible number of 3-tuples of 3-pairwise crossing edges in a drawing of G. Then, if $e > 10n$, then $cr_3(G) \geq \frac{1}{5000}e^5/n^4$.*

Proof: The proof is similar to the proof of the famous *Crossing Number Lemma* [8]. It follows from Theorem 5.2.1 that $cr_3(G) \geq m - 8n$: Assume there is a drawing of G with $cr_3(G)$ 3-tuples of 3-pairwise crossing edges, such that $cr_3(G) < m - 8n$. Then by removing $cr_3(G)$ edges, one from each of the 3-tuples of 3-pairwise crossing

edges, we end up with more than $8n$ edges and no 3-pairwise crossing edges; this is a contradiction.

Now choose every vertex of G with uniform probability p, to be set later. Denote by G' the induced graph, and by n' and m' its number of vertices and edges, respectively. Then, $\mathrm{E}[n'] = pn$, $\mathrm{E}[m'] = p^2 m$, and $\mathrm{E}[cr_3(G')] = p^6 cr_3(G)$.

From the linearity of expectation we have $\mathrm{E}[cr_3(G')] \geq \mathrm{E}[m'] - 8\mathrm{E}[n']$; thus, $cr_3(G) \geq m/p^4 - 8n/p^5$. It is now easy to verify that $p = 10n/m$ maximizes $cr_3(G)$ in this expression, and yields the stated lower bound. □

In a similar way Theorem 5.3.1 implies:

Corollary 5.4.2 *Let G be a graph with n vertices and m edges, and let $cr_4(G)$ be the minimum possible number of 4-tuples of 4-pairwise crossing edges in a drawing of G. Then, if $e > 33.6n$, then $cr_4(G) \geq 9.9 \cdot 10^{-11} e^7/n^6$.*

Our most important result is the first proof that a 4-quasi-planar graph on n vertices has $O(n)$ edges. An interesting open problem is to determine the exact constant hidden in the O-notation. By noticing that it is impossible that all the faces incident to a vertex of G are 1-triangles (as done in the proof of Theorem 5.2.1), one can reduce this constant to 28.8, but this is probably still not tight. The linear upper bound for the maximum number of edges in 4-quasi-planar graphs, combined with the analysis in [48] and [49] also yield the following corollaries.

Corollary 5.4.3 *For any fixed integer $k > 4$, a simple topological graph on n vertices with no k pairwise crossing edges has $O(n \log^{2k-8} n)$ edges.*

Corollary 5.4.4 *For any fixed integer $k > 4$, a topological graph on n vertices with no k pairwise crossing edges has $O(n \log^{4k-16} n)$ edges.*

The main open problem remains to settle Conjecture 5.1.2 for $k \geq 5$. An affirmative answer to the following problem from [6], would immediately imply Conjecture 5.1.2.

Problem 5.4.5 *Is it possible to color the edges of a given k-quasi-planar graph with a constant number of colors, such that there is no pair of monochromatic crossing edges?*

It is known [33] that $O(k \log k)$ colors suffice for convex geometric k-quasi-planar graphs. A result of McGuinness [39] implies an affirmative answer also for simple 3-quasi-planar graphs for which there is a closed curve that intersects every edge exactly once.

Our proof technique might also be useful for other Turán-type problems for topological graphs. Given a topological graph G, consider the induced plane graph and charge every face $|f|$ of it with $|f| + v(f) - 4$ units of charge. Then, the overall charge is $4n-8$. Now, if charges can be distributed such that the average charge of a face remains non-negative, and every wedge of the topological graph has some charge $\varepsilon(n)$, then the number of edges in G is at most $(4n-8)/(2\varepsilon(n))$. If one wishes to distribute charges such that every face has a non-negative charge (as is done in the proofs of Theorems 5.2.1 and 5.3.1), then the problematic faces are 0-triangles and 0-pentagons. We now give two other examples where this technique is useful.

Example 1: What is the maximum number of edges in a topological graph on n vertices where no edge is crossed more than once?
The answer to this question was shown to be $4n-8$ by Pach and Tóth [50] in a more general setting. However, obtaining this bound using our technique is easy. Given such a graph G (drawn without empty lenses), every face of the induced plane graph

has a positive charge. Moreover, the first face of a wedge has at least 1/2 units of charge to contribute to the wedge (the worst case is a 2-triangle). Thus, $\varepsilon = 1/2$, and the upper bound follows.

Example 2: The proof of Theorem 5.2.1 shows that a topological graph on n vertices, drawn such that its induced plane graph contains no 0-triangles, has at most $8n - O(1)$ edges. One can ask what is the maximum number of edges in graphs that can be drawn with no 0-k-gons in their induced plane graph.

Corollary 5.4.6 *A simple topological graph that can be drawn such that the plane graph induced by its drawing contains no* 0*-pentagons has at most* $6n - O(1)$ *edges, and this bound is sharp up to some additive constant.*

Proof: Given such a graph G, we assign a charge of $|f| + v(f) - 4$ to every face f in the induced plane graph. Then, each 0-triangle obtains one unit of charge in the way described in the proof of Theorem 5.3.1. Since there are no 0-pentagons, every face remains with a non-negative charge. Next, every wedge obtains 1/3 units of charge and the upper bound follows. Figure 5.12 shows that this upper bound is tight (up to some additive constant). □

Problem 5.4.7 *What is the maximum number of edges in a topological graph that can be drawn such that there are no* 0*-quadrilaterals (resp.,* 0*-*k*-gons,* $k > 5$*) in the plane graph induced by its drawing?*

A lower bound for $f_k(n)$. It would be interesting to generalize the construction in Section 5.2.2 and obtain a non-trivial lower bound for $f_k(n)$, for $k > 3$. A trivial bound can be deduced from the tight bound of $2(k-1)n - \binom{2k-1}{2}$ for convex k-quasi-planar graphs on n vertices [16]. By placing $n - (k-1)$ points in a convex position,

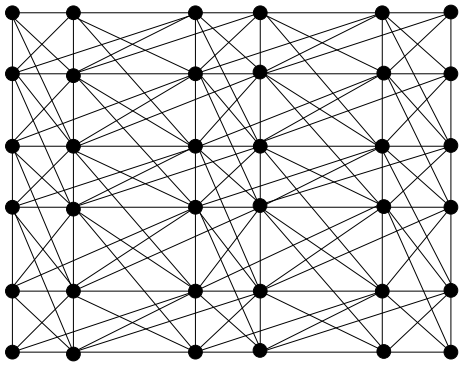

Figure 5.12: A construction for a topological graph on n vertices with no induced 0-pentagons and $6n - O(1)$ edges

איור 5.12: גרף טופולוגי עם n צמתים ו-$6n - O(1)$ קשתות שאינו משרה 0-מחומשים

such that they all "see" another set of $k-1$ points in a convex position, one can obtain a (geometric) k-quasi-planar graph with $2(k-1)(n-k+1) - \binom{2k-1}{2} + (n-k+1)(k-1) + (k-1)^2 = 3(k-1)n - (4k-3)(k-1)$ edges.

Finally, we mention some other problems related to quasi-planar graphs.

Problem 5.4.8 *What is the maximum integer $m(k)$, such that $K_{m(k)}$ is a k-quasi-planar graph?*

According to Figure 5.13, $m(3) \geq 9$. For geometric graphs, this is also the best possible bound [7]. Theorems 5.2.1 and 5.2.5 imply that $m(3) \leq 10$ for simple topological graphs, and $m(3) \leq 14$ for non-simple graphs.

Problem 5.4.9 *Given a graph G and an integer k, determine whether G can be drawn as a k-quasi-planar graph.*

For $k = 2$, the problem becomes a planarity test and can be solved in linear time (see, e.g., [31]). However, we suspect that this problem is NP-hard even for $k = 3$. It

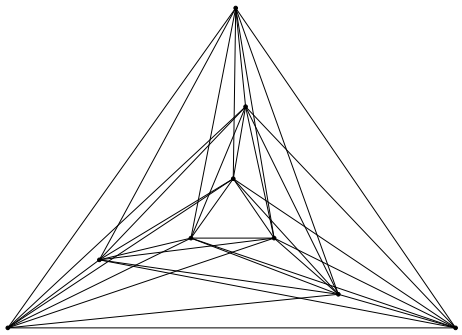

Figure 5.13: K_9 drawn as a geometric 3-quasi-planar graph
איור 5.13: K_9 מצוייר כגרף גיאומטרי 3-מעין מישורי

is known [29] that deciding whether a graph can be drawn with at most one crossing per edge is intractable.

The *crossing number* of a graph G, $cr(G)$, is the minimum possible number of crossings with which G can be drawn. Numerous works have studied different aspects of the crossing number (see [46] for a survey). Let G be a graph that can be drawn as a k-quasi-planar graph, and let $cr_k(G)$ be the minimum possible number of crossings in such a drawing of G. It seems reasonable that for some graphs G, in order to draw them as k-quasi-planar graphs, one has to "pay" by having more crossings than $cr(G)$.

Problem 5.4.10 *Is there a function $f(n,k)$, such that $cr(G) > f(n,k) \cdot cr_k(G)$?*

Chapter 6

Conclusions

In this work we have studied several counting problems in combinatorial geometry. In Chapters 2-4 we considered *rectangular partitions*, with and without point-constraints, in the plane and in higher dimensions. The first problem, discussed in Chapter 2, concerns the number of *mosaic floorplans*—rectangular partitions with no point constraints in the plane. We showed a simple and efficient bijection between floorplans with n blocks (rectangles) and Baxter permutations on $[n]$. From this bijection one can deduce a similar correspondence between separable permutations and slicing floorplans. Moreover, combined with known results on Baxter permutations, the bijection suggests enumerations of floorplans according to various parameters, such as the number of vertical segments. Finding the exact number of general floorplans remains an interesting open problem with applications to integrated circuits design.

In Chapter 3 we discussed a generalization of mosaic floorplans to box partitions in three and higher dimensions. For *guillotine partitions*, a generalization of slicing floorplans, we were able to provide an exact summation formula as well as an asymptotic analysis for the number of structurally-different partitions in d dimensions by

n hyperplanes. We also showed that 3-dimensional partitions by rectangular "walls" are equivalent to extrusions of 2-dimensional mosaic floorplans. Using this property one can count the number of such partitions for small values of n (the number of "walls"). However, finding an exact or even recursive formula for the number of partitions-by-rectangles is an open question.

We call 2-dimensional rectangular partitions subject to point-constraints *rectangulations*, and we discussed their number in Chapter 4. We showed that the number of rectangulations (by n segments) of a set P of n noncorectilinear points depends only on the permutation in which the points are arranged. For any permutation, the number of guillotine rectangulations is always the nth Schröder number and the total number of rectangulation is $O(18^n/n^4)$. For point sets in a separable permutation we proved that the number of rectangulations is the $(n+1)$st Baxter number. We suspect that *any* set of points has at least that many rectangulations.

In Chapter 5 we considered the maximum number of edges in *k-quasi-planar graphs*—graphs that can be drawn with no k pairwise crossing edges. We gave a simple proof that a 3-quasi-planar graph on n vertices has at most $8n - O(1)$ edges, thus improving the previously best upper bound of $65n$ [48]. It is an open question to close the gap between the new upper bound and our best lower bound construction with $7n-O(1)$ edges. For *simple* 3-quasi-planar graphs we were able to show a sharp bound (up to some additive constant) of $6.5n-O(1)$ edges. Our most important result is the first proof that a 4-quasi-planar graph on n vertices has $O(n)$ edges. This is a further step forward in proving a cojecture from the early 1990s saying that for every fixed k, the number of edges in a k-quasi-planar graph is linear in the number of vertices. At the moment, we do not know whether our new proof technique can be extended to 5-quasi-planar graphs and greater values of k. Still, the new upper

bound for 4-quasi-planar graphs results in improved upper bounds for the maximum number of edges in k-quasi-planar graphs, for $k > 4$.

References

[1] E. Ackerman. On the maximum number of edges in topological graphs with no four pairwise crossing edges. In *Proc. 22nd ACM Symp. on Computational Geometry*, pages 259–263, Sedona, AZ, 2006.

[2] E. Ackerman, G. Barequet, and R. Y. Pinter. A bijection between permutations and floorplans, and its applications. *Discrete Applied Mathematics*, 154(12):1674–1684, 2006.

[3] E. Ackerman, G. Barequet, and R. Y. Pinter. On the number of rectangulations of a planar point set. *J. Combinatorial Theory, Ser. A*, 113(6):1072–1091, 2006.

[4] E. Ackerman, G. Barequet, R. Y. Pinter, and D. Romik. The number of guillotine partitions in d dimensions. *Information Processing Letters*, 98(4):162–167, 2006.

[5] E. Ackerman and G. Tardos. On the maximum number of edges in quasi-planar graphs. *J. Combinatorial Theory, Ser. A*. to appear.

[6] P. K. Agarwal, B. Aronov, J. Pach, R. Pollack, and M. Sharir. Quasi-planar graphs have a linear number of edges. *Combinatorica*, 17(1):1–9, 1997.

[7] O. Aichholzer and H. Krasser. The point set order type data base: A collection of applications and results. In *Proc. 13th Canadian Conf. on Computational Geometry*, pages 17–20, Waterloo, Ontario, Canada, Aug. 2001.

[8] M. Ajtai, V. Chvátal, M. Newborn, and E. Szemerédi. Crossing-free subgraphs. *Annals of Discrete Mathematics*, 12:9–12, 1982.

[9] K. Appel and W. Haken. Every planar map is four colorable. Part I. Discharging. *Illinois J. Math.*, 21:429–490, 1977.

[10] D. Avis and K. Fukuda. Reverse search for enumeration. *Discrete Applied Mathematics*, 65:21–46, 1996.

[11] D. Avis and M. Newborn. On pop-stacks in series. *Utilitas Mathematica*, 19:129–140, 1981.

[12] C. Ballieux. Motion planning using binary space partitions. Technical report inf/src/93-25, Utrecht University, The Netherlands, 1993.

[13] G. Baxter. On fixed points of the composite of commuting functions. *Proc. Amer. Math. Soc.*, 15:851–855, 1964.

[14] P. Bose, J. Buss, and A. Lubiw. Pattern matching for permutations. *Information Processing Letters*, 65:277–283, 1998.

[15] F. Calheiros, A. Lucena, and C. de Souza. Optimal rectangular partition. *Networks*, 41(1):51–67, 2003.

[16] V. Capoyleas and J. Pach. A Turán-type problem on chords of a convex polygon. *J. Combinatorial Theory, Ser. B*, 56(1):9–15, 1992.

[17] M. Cardei, X. Cheng, X. Cheng, and D.-Z. Du. A tale on guillotine cut. In P. M. Pardalos and H. Wolkowicz, editors, *Novel Approaches to Hard Discrete Optimization*, volume 37 of *Fields Institute Communications*, pages 41–54. American Mathematical Society, Providence, RI, 2003.

[18] N. Chin and S. Feiner. Near real-time shadow generation using BSP trees. *Computer Graphics*, 23(3):99–106, 1989.

[19] F. R. K. Chung, R. L. Graham, V. E. Hoggatt, and M. Kleiman. The number of Baxter permutations. *J. Combinatorial Theory, Ser. A*, 24(3):382–394, 1978.

[20] H. de Fraysseix, P. de Mendez, and J. Pach. A left-first search algorithm for planar graphs. *Discrete & Computional Geometry*, 13:459–468, 1995.

[21] D. Z. Du, L. Q. Pan, and M. T. Shing. Minimum edge length guillotine rectangular partition. Technical report MSRI 02418-86, University of California, Berkeley, CA, 1986.

[22] S. Dulucq and O. Guibert. Baxter permutations. *Discrete Mathematics*, 180:143–156, 1998.

[23] H. Fuchs, Z. M. Kedem, and B. Naylor. On visible surface generation by a priori tree structures. *Computer Graphics*, 14(3):124–133, 1980.

[24] I. M. Gessel and R. P. Stanley. Algebraic enumeration. In R. L. Graham, M. Grötschel, and L. Lovász, editors, *Handbook of Combinatorics*, volume 2, pages 1021–1061. MIT Press, 1995.

[25] T. F. Gonzalez, M. Razzazi, M.-T. Shing, and S.-Q. Zheng. On optimal guillotine partitions approximating optimal d-box partitions. *Computational Geometry: Theory & Applications*, 4:1–11, 1994.

[26] T. F. Gonzalez and S.-Q. Zheng. Bounds for partitioning rectilinear polygons. In *Proc. 1st ACM Symp. on Computational Geometry*, pages 281–287, Baltimore, MD, 1985.

[27] T. F. Gonzalez and S.-Q. Zheng. Improved bounds for rectangular and guillotine partitions. *J. of Symbolic Computation*, 7:591–610, 1989.

[28] T. F. Gonzalez and S.-Q. Zheng. Approximation algorithms for partitioning a rectangle with interior points. *Algorithmica*, 5:11–42, 1990.

[29] A. Grigoriev and H. L. Bodlaender. Algorithms for graphs embeddable with few crossings per edge. *Algorithmica*. to appear.

[30] X. Hong, G. Huang, Y. Cai, J. Gu, S. Dong, C. K. Cheng, and J. Gu. Corner block list: an effective and efficient topological representation of non-slicing floorplan. In *Proc. IEEE/ACM Int. Conf. on Computer-Aided Design*, pages 8–12, San Jose, CA, 2000.

[31] J. Hopcroft and R. Tarjan. Efficient planarity testing. *J. ACM*, 21(4):549–568, 1974.

[32] Y. Kajitani. The single sequence that unifies placement and floorplanning, 2003. presented at the *1st Conf. for Asia Universities on Semiconductors Design*.

[33] A. Kostochka and J. Kratochvíl. Covering and coloring polygon-circle graphs. *Discrete Mathematics*, 163:299–305, 1997.

[34] T. Lengauer. *Combinatorial algorithms for integrated circuit layout*. John Wiley & Sons, Inc., New York, NY, 1990.

[35] C. Levcopoulos. Fast heuristics for minimum length rectangular partitions of polygons. In *Proc. 2nd ACM Symp. on Computational Geometry*, pages 100–108, Yorktown Heights, NY, 1986.

[36] A. Lingas, R. Y. Pinter, R. L. Rivest, and A. Shamir. Minimum edge length rectilinear decompositions of rectilinear figures. In *Proc. 20th Allerton Conf. on Communication, Control, and Computing*, pages 53–63, Monticello, IL, 1982.

[37] C. L. Mallows. Baxter permutations rise again. *J. Combinatorial Theory, Ser. A*, 27(3):394–396, 1979.

[38] A. Marcus and G. Tardos. Intersection reverse sequences and geometric applications. *J. Combinatorial Theory, Ser. A*, 113(4):675–691, 2006.

[39] S. McGuinness. Colouring arcwise connected sets in the plane I. *Graphs and Combinatorics*, 16:429–439, 2000.

[40] J. S. B. Mitchell. Guillotine subdivisions: Part II—a simple polynomial-time approximation scheme for geometric k-MST, TSP, and related problems. *SIAM J. on Computing*, 28:1298–1309, 1999.

[41] J. I. Munro, T. Papadakis, and R. Sedgewick. Deterministic skip lists. In *Proc. 3rd ACM-SIAM Symp. on Discrete Algorithms*, pages 367–375, Orlando, FL, 1992.

[42] H. Murata, K. Fujiyoshi, S. Nakatake, and Y. Kajitani. VLSI module placement based on rectangle-packing by the sequence-pair. *IEEE Trans. on Computer-Aided Design of Integrated Cirtcuits and Systems*, 15(12):1518–1524, 1996.

[43] H. Murata, K. Fujiyoshi, T. Watanabe, and Y. Kajitani. A mapping from sequence-pair to rectangular dissection. In *Proc. Asia and South Pacific Design Automation Conf.*, pages 625–633, Chiba , Japan, 1997.

[44] A. M. Odlyzko. Asymptotic enumeration methods. In R. L. Graham, M. Grötschel, and L. Lovász, editors, *Handbook of Combinatorics*, volume 2, pages 1063–1229. MIT Press, 1995.

[45] J. Pach. Notes on geometric graph theory. In J. E. Goodman, R. Pollack, and W. Steiger, editors, *Discrete and Computational Geometry: Papers from DIMACS special year*, volume 6 of *DIMACS series*, pages 273–285. American Mathematical Society, Providence, RI, 1991.

[46] J. Pach. Geometric graph theory. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, pages 219–238. CRC Press, Boca Raton, 2nd edition, 2004.

[47] J. Pach, R. Pinchasi, M. Sharir, and G. Tóth. Topological graphs with no large grids. *Graphs and Combinatorics*, 21(3):355–364, 2005.

[48] J. Pach, R. Radoičić, and G. Tóth. Relaxing planarity for topological graphs. In J. Akiyama and M. Kano, editors, *JCDCG*, volume 2866 of *Lecture Notes in Computer Science*, pages 221–232. Springer, 2002.

[49] J. Pach, F. Shahrokhi, and M. Szegedy. Applications of the crossing number. *Algorithmica*, 16(1):111–117, 1996.

[50] J. Pach and G. Tóth. Graphs drawn with few crossings per edge. *Combinatorica*, 17(3):427–439, 1997.

[51] K. Sakanushi, Y. Kajitani, and D. P. Mehta. The quarter-state-sequence floorplan representation. *IEEE Trans. on Circuits and Systems I: Fundamental Theory and Applications*, 50(3):376–386, 2003.

[52] B. Salvy. The `gdev` Maple package (appears in the `algolib` package library). available at `http://algo.inria.fr/libraries`.

[53] F. Santos and R. Seidel. A better upper bound on the number of triangulations of a planar point set. *J. Combinatorial Theory, Ser. A*, 102(1):186–193, 2003.

[54] L. Shapiro and A. B. Stephens. Bootstrap percolation, the Schröder numbers, and the n-Kings problem. *SIAM J. on Discrete Mathematics*, 4:275–280, 1991.

[55] M. Sharir and E. Welzl. Random triangulations of planar point sets. In *Proc. 22nd ACM Symp. on Computational Geometry*, pages 273–281, Sedona, AZ, 2006.

[56] Z. C. Shen and C. C. N. Chu. Bounds on the number of slicing, mosaic, and general floorplans. *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, 22(10):1354–1361, 2003.

[57] R. P. Stanley. *Enumerative Combinatorics*, volume 2. Cambridge, 1999.

[58] G. Tardos. Construction of locally plane graphs. *SIAM J. on Discrete Mathematics*. to appear.

[59] G. Tardos and G. Tóth. Crossing stars in topological graphs. In J. Akiyama, M. Kano, and X. Tan, editors, *JCDCG*, volume 3742 of *Lecture Notes in Computer Science*, pages 184–197. Springer, 2004.

[60] W. C. Thilbault and B. F. Naylor. Set operations on polyhedra using binary space partitioning trees. *Computer Graphics*, 21(4):153–162, 1987.

[61] P. Turán. On an extremal problem in graph theory. *Mathmatikai Lapok*, 48:436–452, 1941.

[62] P. Valtr. Graph drawings with no k pairwise crossing edges. In G. D. Battista, editor, *Graph Drawing*, volume 1353 of *Lecture Notes in Computer Science*, pages 205–218. Springer, 1997.

[63] J. West. Generating trees and the Catalan and Schröder numbers. *Discrete Mathematics*, 146:247–262, 1995.

[64] B. Yao, H. Chen, C. Cheng, and R. Graham. Floorplan representations: Complexity and connections. *ACM Trans. on Design Automation of Electronic Systems*, 8:55–88, 2003.

[65] E. Young, C. Chu, and C. Shen. Twin binary sequences: a non-redundant representation for general non-slicing floorplan. *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, 22(4):457–469, 2003.

[66] X. Zhang and Y. Kajitani. A normalized configuration of floorplans and ABLR-relations. In *Int. Conf. on Communications, Circuits and Systems*, pages 1218–1222, Chengdu, China, 2004.

[67] X. Zhang and Y. Kajitani. Space-planning: placement of modules with controlled empty area by single-sequence. In *Proc. Asia and South Pacific Design Automation Conf.*, pages 25–30, Yokohama , Japan, 2004.

[68] X. Zhang and Y. Kajitani. Theory of T-junction floorplans in terms of single-sequence. In *IEEE Int. Symp. on Circuits and Systems*, pages 341–344, Vancouver, Canada, 2004.

[69] X. Zhu, C. Zhuang, and Y. Kajitani. A general packing algorithm based on single-sequence. In *Int. Conf. on Communications, Circuits and Systems*, pages 1257–1261, Chengdu, China, 2004.

[70] C. Zhuang, X. Zhu, Y. Takashima, S. Nakatake, and Y. Kajitani. An algorithm for checking slicing floorplan based on HPG and its application. In *Int. Conf. on Communications, Circuits and Systems*, pages 1223–1227, Chengdu, China, 2004.

בעיות ספירה במבנים גיאומטריים:
מילבונים, מפות-תכנון למעגלים מוכללים
וגרפים מעין-מישוריים

אייל אקרמן

בעיות ספירה במבנים גיאומטריים:
מילבונים, מפות-תכנון למעגלים מוכללים
וגרפים מעין-מישוריים

חיבור על מחקר

לשם מילוי חלקי של הדרישות לקבלת תואר
דוקטור לפילוסופיה

אייל אקרמן

הוגש לסנט הטכניון — מכון טכנולוגי לישראל

אלול תשס"ו חיפה ספטמבר 2006

חיבור על מחקר שנעשה בהדרכת פרופ׳ גיל ברקת ופרופ׳ רון י. פינטר

הפקולטה למדעי המחשב

הכרת תודה

אני אסיר תודה למנחים שלי, גיל ברקת ורון פינטר, על שהציעו שאלת מחקר שהתפתחה לחלק העיקרי בתזה זו ועל הנחייתם במהלך המחקר. אני מודה להם על אמונם בי, על העידוד והתמיכה בתקופות פחות טובות ועל היחס הידידותי והאווירה הנעימה בפגישות בינינו. תודה גם ל-Gábor Tardos ודן רומיק שהיו שותפים לעבודות הנכללות בתזה זו.

תודה מיוחדת שלוחה לרום פנחסי. נהנתי מאוד מהקורס בנושאים בגיאומטריה קומבינטורית שלימד, ומהעבודה המשותפת איתו מאוחר יותר. רום הציג בפני את בעיית הקיצון הקשורה לגרפים מעין-מישוריים, והדיונים בבעיה זו הובילו לתוצאה המרכזית בעבודה זו. אני מודה לרום על ארגון ותמיכה כספית בנסיעה שערכתי לארה״ב, ומודה ל-János Pach על שאירח אותי בניו-יורק בביקור זה. אני מקווה להמשך שיתוף הפעולה עם רום ורואה בו מודל לחיקוי כחוקר.

במהלך לימודיי התארחתי כשלושה חודשים בקבוצת המחקר בתיאוריה של מדעי המחשב ב-Free University Berlin. אני מודה לחברי הקבוצה על סביבת העבודה המסבירת פנים והידידותית שיצרו. תודה מיוחדת ל-Christian Knauer, Günter Rote ו-Kevin Buchin איתם נהנתי לשתף פעולה ול-Kevin ו-Maike Buchin על האירוח הנדיב והטיולים המשותפים. תודה לתוכנית האיר-ופאית Marie Curie שאיפשרה ביקור זה.

אני מודה לחברי המרכז לחישוב גרפי וגיאומטרי בטכניון ובייחוד לדבי מילר על עזרתה המינהלת-ית והלשונית לאורך לימודי. תודה לאורן בן-צבי על החבר שהינו ועל דיונים משותפים רבים מספור. יחסרו לי השיחות המעניינות עם דודו אמזלג בהפסקות הקפה שלנו. לסיום, אני רוצה להביע את אהבתי ותודתי לבת-זוגי עופרית על תמיכתה ואהבתה.

אני מודה לטכניון על התמיכה הכספית הנדיבה בהשתלמותי

תוכן עניינים

רשימת איורים

רשימת טבלאות

תקציר

נושאי המחקר המוצגים בעבודה זו שייכים לתחום הנקרא גיאומטריה קומבינטורית. תחום זה עוסק בבעיות ספירה, קיום, ותכונות של מבנים גיאומטריים דיסקרטיים, המורכבים מאובייקטים גיאומטריים כגון נקודות, קווים, ומצולעים.

בהינתן קבוצה של מבנים קומבינטורים, אחת השאלות הטבעיות לשאול היא: "עד כמה גדולה קבוצה זו?" אנו חוקרים בעיה זו בעבור מילבונים של קבוצת נקודות במישור. באופן פורמלי, בהינתן קבוצה P של n נקודות במישור הנמצאות בתוך מלבן R, מילבון של P היא חלוקה של המלבן R למלבנים קטנים על ידי קטעים מקבילים לצלעות R, כך שכל נקודה ב-P נמצאת על קטע שונה. מציאת מילבון הממזער את סך אורכי הקטעים היא בעיה שנידונה במספר רב של מאמרים בספרות [15, 17, 21, 26, 27, 28, 35, 36] .

אנו מראים שכאשר הסדר היחסי בין הנקודות יוצר תמורה בת-הפרדה, מספר המילבונים הוא מספר תמורות Baxter על $n+1$ אלמנטים. בנוסף, אנו מראים כי לכל סדר יחסי של הנקודות, מספר מילבוני הגיליוטינה הוא תמיד מספר Schröder ה-n-י, ומספר המילבונים הוא $O(18^n/n^4)$. אנו מציגים גם שיטות למעבר מהיר על כל מרחב המילבונים של קבוצה נתונה של נקודות בעזרת שתי פעולות מקומיות.

מבנים קומבינטוריים נוספים אותם חקרנו הם מפות תכנון למעגלים מוכללים (*floorplans*). במהלך תכנון מעגל מוכלל (integrated circuit) נקבעים הצורה, הגודל והמיקום של כל רכיב במעגל (שבב). לרוב, צורת השבב והרכיבים הינה מלבן. מפת-תכנון (floorplan) מתארת את המיקום היחסי של הרכיבים על השבב, ולכן מיוצגת בדרך כלל באמצעות

חלוקה של מלבן למלבנים קטנים יותר. סוג מיוחד של חלוקות הן חלוקות גליוטינה, בהן חלוקת המלבן מתבצעת באופן רקורסיבי על ידי חיתוך של מלבן לשני מלבנים קטנים יותר באמצעות קטע אנכי או אופקי.

ידוע כי מספר מפות התכנון השונות שווה למספרן של תמורות Baxter. אנו מציגים התאמה פשוטה ויעילה בין תמורות Baxter ומפות תכנון, אשר ממנה ניתן להסיק מהו מספר מפות התכנון בהתאם לפרמטרים מבניים שונים, כגון מספר קטעי החלוקה האנכיים. כמו כן, התאמה זו ממפה חלוקות גיליוטינה לתמורות בנות-הפרדה.

ניתן להכליל חלוקות מלבניות ליותר משני מימדים ולבחון את מספר הדרכים השונות לחלק תיבה תלת-מימדית לתיבות תלת-מימדיות קטנות יותר. חלוקות גיליוטינה מוכללות באופן טבעי לשלושה מימדים ומעלה. אנו מציגים נוסחת סכום מדוייקת בעבור מספר חלוקות הגיליוטינה השונות על ידי n על-מישורים ומראים שבאופן אסימפטוטי מספרן הוא $\Theta\left(\left(2d-1+2\sqrt{d(d-1)}\right)^n/n^{3/2}\right)$ בעבור מימד קבוע d. לחלוקות גליוטינה תפקיד חשוב בתחומי מחקר ויישומים רבים, כגון גיאומטריה חישובית, גרפיקה ממוחשבת, ומידול. אנו דנים גם בחלוקות תלת-מימדיות בהן החלוקה מתבצעת באמצעות "קירות" מלבניים. בהתבסס על מספר תכונות מבניות של חלוקות אלה, אנו מתארים שיטה לספירתן.

בעיות-קיצון בתורת הגרפים (extremal graph theory) דנות בגרפים "קיצוניים", באופן כלשהו, מבין גרפים המקיימים תכונות מסויימות. אנו חוקרים בעית-קיצון בתורת הגרפים הגיאומטריים. גרף גיאומטרי הוא גרף המצויר במישור כך שקודקודיו הם נקודות שונות, וקשתותיו הן קטעים ישרים המחברים את הנקודות המתאימות ואינם מכילים אף קודקוד אחר. גרף טופולוגי מוגדר באופן דומה, אלא שקשתותיו יכולות להיות כל עקום פשוט (כלומר, עקום שאינו חותך את עצמו). חיתוכים בין הקשתות אפשריים, אך אנו מניחים כי כל זוג קשתות נחתך מספר סופי של פעמים. גרף טופולוגי נקרא פשוט אם כל זוג קשתות נפגש פעם אחת לכל היותר (בקודקוד או בנקודת חיתוך).

בבעיה מסוג Turán על גרפים טופולוגים, אנו מבקשים לדעת את מספר הקשתות המירבי שיתכן בגרף עם n קודקודים שאינו מכיל תבנית גיאומטרית נתונה. למשל, קל

להוכיח באמצעות נוסחת אוילר כי גרף טופולוגי ללא זוג קשתות נחתכות (כלומר גרף מישורי) מכיל לכל היותר $3n - 6$ קשתות. גרפים מעין מישוריים הם הכללה טבעית של גרפים מישוריים: גרף טופולוגי נקרא k-מעין מישורי אם אינו מכיל קבוצה של k קשתות נחתכות-הדדית (כלומר, כל אחת נחתכת עם כל האחרות).

לפי השערה מפורסמת למדי מתחילת שנות ה-90, לכל קבוע k קיים קבוע אחר C_k כך שגרף k-מעין מישורי עם n קודקודים מכיל לכל היותר $C_k \cdot n$ קשתות. היות וגרפים 2-מעין מישוריים הינם גרפים מישוריים ההשערה נכונה באופן טריוויאלי בעבור המקרה $k = 2$. באמצע שנות ה-90 הוכחה ההשערה לראשונה בעבור המקרה $k = 3$, כאשר הגרף הטופולוגי פשוט [6]. מאוחר יותר פושטה הוכחה זו והוכללה גם בעבור גרפים שאינם פשוטים [84]. החסם העליון על מספר הקשתות בגרף 3-מעין מישורי שהוכח בעבודה זו היה $65n$. החסמים העליונים הטובים ביותר בעבור $k > 3$ היו $O(n \log n)$ בעבור גרפים גיאומטריים [26]; $O(n \log^{2k-6} n)$ בעבור גרפים פשוטים [6]; ו-$O(n \log^{4k-12} n)$ בעבור גרפים לא פשוטים בהכרח [84]. בעבודה זו אנו מוכיחים לראשונה כי לגרפים 4-מעין מישוריים מספר קשתות הליניארי במספר הקודקודים. אנו גם מראים את ההוכחה הפשוטה ביותר והחסם העליון הטוב ביותר הידועים בעבור מספר הצלעות המירבי בגרף 3-מעין-מישורי. יתר על כן, אנו מראים חסם עליון הדוק של $6.5\,n - O(1)$ על מספר הקשתות בגרף 3-מעין-מישורי פשוט.